

Proceedings of the 26th Philippine Computing Science Congress

Volume I: Full Papers

Editors: Dr. Judith J. Azcarraga and Mr. Lester Sityar

Copyright © 2026 Computing Society of the Philippines

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior written permission of the copyright owner. Individual papers may be uploaded to institutional repositories or other academic sites for self-archival purposes.

ISSN 1908-1146



Computing Society of the Philippines
csp.org.ph



April 23–25, 2026
University of the Immaculate Conception, Davao City

Contents

About the Philippine Computing Science Congress	1
About the Computing Society of the Philippines	1
Organizing Committee	2
Accepted Papers	
Machine Learning-Based Regional K-12 Enrollment Forecasting and Volatility Classification for Adaptive Resource Allocation in the Philippines	3
Defending against Membership Inference Attacks through Contrastive Representation Learning	10
FRIAD-IR: A Framework for Image-Augmented Dynamic Information Retrieval	18
Parallel Permutation Testing for One-Way ANOVA Using CUDA	26
A Scalable and Computationally Efficient Implementation of a Macroscopic Traffic Model for Urban Networks	34
Handwritten Gregg Shorthand Brief Forms Recognition Using Convolutional Neural Networks	41
Integrating Adaptive Mechanisms for Feature Selection with Mayfly Algorithm, Monarch Butterfly Optimization, and Monarch Mayfly Optimization	47
Flotector: A Floating Waste Detection System for the Imus River through Participatory Sensing and YOLOv11	55
Uncertainty-Aware Temporal Transformer Modeling with Masked Self-Attention and Missingness Encoding for Early Sepsis Risk Stratification from ICU Time-Series Data	63
MAMBO (Multi AI Model Bench Operations): A Framework for Dynamic Model Switching and Benchmarking On-Device AI	71
Distributed SNP: A High-Performance Hybrid MPI+CUDA Framework for Spiking Neural P Systems	77
Nonbalanced and Nonhomogenous Spiking Neural dP Systems	85
Apparent Age Estimation: Challenges and Outcomes	93
TUDLO: Parameter-Efficient Fine-Tuning for Tagalog and Cebuano Early-Grade Mathematics Explanation Generation	101
Gated Learning-Progress Exploration	109
Evaluating Large Language Models in Filipino Underspecified, Multi-Turn Conversations	117
Mixture of Experts with Soft Nearest Neighbor Loss: Resolving Expert Collapse via Representation Disentanglement	125
Conversion from Nondeterminism to Determinism at the Rule Level of Spiking Neural P Systems with Structural Plasticity	132
On the Notions of Soundness in Spiking Neural P Systems	140

Galaw at Gunita: Extended Reality Murals for Experiencing Filipino Art	148
Predictive Models for Solar Energy Output: Performance Evaluation of K-Means and GMM Clustering with Model Interpretability	156
A Hybrid Machine Learning Framework for Predictive Maintenance of Bearings in Philippine Power Plants	163
Dynamic Load Balancing and Parallel Non-Maximum Suppression for Viola-Jones Object Detection on Multi-Core and GPU Architectures	171
An Ensemble Model Using LightGBM and Bidirectional LSTM for Bitcoin Price Dip Prediction	179
Adaptive Shot Allocation for Quantum Observable with Statistical Guarantees	187
Neuro-Symbolic AI and Autoformalization: Restoring Determinism in the Age of Probabilistic Computing	195
Author Index	203
Institution Index	205



April 23–25, 2026
University of the Immaculate Conception, Davao City

About the Philippine Computing Science Congress

The **Philippine Computing Science Congress (PCSC)** is organized by the Computing Society of the Philippines to enable local and neighboring computing educators, researchers, information and communications technology (ICT) professionals, and students to interact and share their work in computing, computer science, computational science, and ICT. The conference features special lectures by prominent researchers and educators and contributed papers in ICT, computing, computer science, computational science, and related disciplines.

About the Computing Society of the Philippines

The **Computing Society of the Philippines (CSP)** is an organization promoting research and development in computing science, computer science, computer engineering, computational science, computing education, and related disciplines by promoting the exchange of knowledge in this field, especially at the tertiary and graduate levels. CSP's main objective is to advocate progressive policies and develop programs that affect the computing sector. CSP establishes linkages with other organizations in the pursuit of common goals.



April 23–25, 2026
University of the Immaculate Conception, Davao City

Organizing Committee

Program Chair

Judith J. Azcarraga, Ph.D.

Local Host Chair

Kristine Mae Adlaon, Ph.D.

Organizing Committee Members

Cherry Lyn Sta. Romana, MSCS, DIT
Kristine Mae Adlaon, Ph.D.
Tita Herradura

Judith J. Azcarraga, Ph.D.
Ramon Rodriguez
Cherry Lyn Sta. Romana, MSCS, DIT

Organizers



Computing Society of the Philippines



University of Immaculate Conception

Machine Learning-Based Regional K-12 Enrollment Forecasting and Volatility Classification for Adaptive Resource Allocation in the Philippines

Mc Sergel Cardaño
South Philippine Adventist College
Camanchiles, Matanao, Davao del Sur
Philippines
mscofficial2016@gmail.com

Nice Bulio
South Philippine Adventist College
Camanchiles, Matanao, Davao del Sur
Philippines
ncbulio95@gmail.com

Neila Paglinawan-Muñoz
South Philippine Adventist College
Camanchiles, Matanao, Davao del Sur
Philippines
nmpmunez@gmail.com

ABSTRACT

The Philippine education system faces chronic resource allocation challenges due to centralized planning that cannot adapt to regional demographic variations and enrollment volatility. This study presents a dual machine learning framework combining enrollment forecasting with volatility classification to enable adaptive resource distribution across the Philippines' 17 regions. Using DepEd K-12 enrollment data from 2010–2021 spanning public and private sectors, we evaluated four machine learning algorithms: Random Forest, Decision Tree, Linear Regression, and Support Vector Regression. Random Forest achieved superior performance with $R^2=0.9874$ and $MAE=46,313$ students on the pandemic-era test set, indicating robustness to COVID-19–related enrollment disruptions. Our coefficient of variation-based volatility classification identified four high-risk regions (Eastern Visayas, Zamboanga Peninsula, CARAGA, and Bicol) with $CV>127.10\%$, suggesting the need for contingency resource buffers of approximately 15–20% above base allocations. The public sector emerged as the dominant predictor feature, reflecting the structural scale difference between public and private enrollment in Philippine K-12 planning. This data-driven framework provides decision-relevant insights for DepEd policy-makers to inform teacher deployment, classroom construction, and budget allocation.

KEYWORDS

Enrollment forecasting, random forest, decision tree, resource allocation, volatility classification

1 INTRODUCTION

The Philippine education system serves over 27 million learners across 17 administrative regions, yet persistent resource misallocations undermine educational equity despite substantial government investment [2]. As of 2024, classroom shortages exceed 91,000 units, teacher deficits reach 86,000 positions, and student-teacher ratios remain at 1:40, well above UNESCO's recommended 1:30 [3]. These disparities stem partly from centralized planning processes that rely on simplistic population projections and fail to account for regional demographic variations and enrollment volatility [4].

The COVID-19 pandemic starkly exposed these planning limitations. National enrollment plummeted from 27.03 million in 2019 to 26.23 million in 2020, with regional impacts varying dramatically [3]. Traditional forecasting methods based on linear extrapolation could not capture the nonlinear shock patterns, leading to resource misallocations: some regions over-prepared with excess teachers while

others faced critical shortages. This crisis highlighted the urgent need for adaptive, data-driven planning frameworks that can anticipate both long-term demographic trends and short-term volatility.

Machine learning offers significant potential for improving educational planning. Recent studies demonstrate that ensemble methods like Random Forest and gradient-boosted trees can capture complex enrollment dynamics with remarkable accuracy by incorporating demographic, economic, and historical data [1, 8, 9]. Random Forest, first formalized by Breiman [10], has since been applied to a range of educational prediction tasks, with reported prediction errors below 1% in higher education enrollment contexts, enabling administrators to proactively align resources with demand [1, 11]. However, the application of such methods to Philippine K-12 regional planning remains largely unexplored, and few existing studies explicitly combine enrollment forecasting with volatility-based risk assessment for resource allocation.

To address this gap, this study presents and evaluates a dual machine learning framework for regional K–12 enrollment forecasting and volatility classification in the Philippine context. Four regression algorithms are trained on 11 years of DepEd enrollment data covering all 17 administrative regions and both public and private school sectors. Beyond predictive accuracy, the study introduces a coefficient of variation-based classification scheme to distinguish regions with stable enrollment trajectories from those exhibiting high volatility, enabling differentiated resource allocation strategies tailored to regional risk profiles. The framework is intended as a practical planning tool that complements existing DepEd budgeting processes rather than replacing institutional judgment.

2 METHODOLOGY

2.1 Research Approach and Data Description

This study uses a quantitative, data-driven modeling approach to forecast regional K–12 enrollment and to classify enrollment volatility for adaptive resource planning. Comprehensive enrollment data were assembled from the Philippine Department of Education (DepEd), covering academic years from 2010–2011 through 2020–2021 [2].

The dataset spans all seventeen administrative regions of the Philippines, including the National Capital Region (NCR), the Cordillera Administrative Region (CAR), Regions I through XIII (with Region IV subdivided into Region IV-A CALABARZON and Region IV-B MIMAROPA), and the Bangsamoro Autonomous Region in Muslim Mindanao (BARMM). It covers the full K–12 structure introduced

under the K to 12 Basic Education Program [12], including Kindergarten, Grades 1–10, and Grades 11–12, with senior high enrollment disaggregated across eight tracks: ABM, HUMSS, STEM, GAS, Maritime, Technical–Vocational–Livelihood, Sports, and Arts and Design. Enrollment figures from public, private, and State Universities and Colleges and Local Universities and Colleges (SUC-LUC) school systems were included, resulting in a total of 561 region–sector–year observations, with all 17 regions and all school sectors represented across all academic years. The dataset is complete with no missing rows; however, Senior High School (Grades 11–12) enrollment did not exist prior to its progressive introduction in SY 2016–2017 under the K to 12 reform [12]. For academic years before SY 2016–2017, all Senior High columns are recorded as zero, reflecting the structural absence of those grade levels at that time rather than gaps in data collection. No imputation was required, as these zero values are substantively correct.

While the dataset spans 11 years and all 17 regions, the resulting 561 observations represent a modest absolute sample size. Performance estimates are supported by time-aware cross-validation and bootstrap confidence intervals, which provide evidence of stability despite the limited number of data points. Additional years of data would further strengthen the generalizability of the findings.

2.2 Data Preprocessing and Feature Construction

For each region, sector, and academic year, total enrollment was computed by aggregating enrollment counts across all grade levels and tracks. This aggregate value served as the target variable for model training. Formally, total enrollment for a given region–sector–year combination was defined as

$$\text{Total_Enrollment}_{i,t} = \sum_{g \in \text{grades}} \text{Enrollment}_{i,g,t}, \quad (1)$$

where i denotes a region–sector pair, t indexes academic years, and g represents individual grades or senior high tracks. Because pre-implementation Senior High columns are zero-filled rather than absent, the summation is well-defined for all 561 observations and no imputation was applied; the dataset contained zero missing values prior to model training.

Explanatory features included categorical indicators for region and school sector, temporal variables capturing academic year boundaries, and a binary indicator representing the COVID-19 pandemic period. The feature set was deliberately restricted to administratively available indicators to ensure the framework remains deployable within DepEd’s existing data infrastructure without requiring socio-economic data collection.

To simulate real-world forecasting conditions and prevent information leakage from future observations, the dataset was partitioned using a walk-forward temporal split, with earlier years used for training and the two most recent pandemic-era years reserved exclusively for evaluation. Table 1 details the resulting split, where School Year (SY) denotes the Philippine academic year designation.

Numerical features were standardized for models sensitive to scale (i.e., Support Vector Regression), while tree-based methods operated on raw values due to their inherent scale invariance.

Table 1: Walk-Forward Temporal Split of the DepEd Enrollment Dataset

Split	Academic Years	Observations
Training	SY 2010–2011 to SY 2018–2019	459
Test	SY 2019–2020 to SY 2020–2021	102
Total	SY 2010–2011 to SY 2020–2021	561

2.3 Forecasting Models

Four regression models spanning the complexity spectrum were selected for this study. Linear Regression served as the statistical reference point, establishing the performance floor that nonlinear models must exceed to justify their added complexity. Support Vector Regression with a radial basis function kernel was included to test whether a kernel-based nonlinear approach could capture enrollment dynamics; features were standardized before training, as required by the algorithm. A Decision Tree was chosen to represent interpretable, rule-based partitioning of the feature space, while a Random Forest [10] extended this by aggregating predictions across an ensemble of trees, which reduces sensitivity to noise in individual training observations.

It should be noted that a persistence baseline (using each region’s prior-year enrollment as the prediction) was not included in the comparison. Such a baseline would provide a useful lower bound for evaluating whether the machine learning models offer meaningful improvement over the simplest possible forecast. This is acknowledged as a limitation and is proposed as a direct extension in Future Work (Section 4.1, item 2).

All models were implemented using Python with standard scientific computing libraries. Model training incorporated time-aware cross-validation to respect the temporal ordering of the data and to reduce the risk of information leakage from future observations.

Model performance was evaluated through four complementary metrics. The coefficient of determination (R^2) measures how much enrollment variance the model explains. Mean absolute error (MAE) quantifies the typical prediction gap in student counts, which is directly interpretable for resource planning purposes. Root mean squared error (RMSE) imposes greater penalty on large prediction errors, relevant for budgeting scenarios where extreme misallocations carry disproportionate costs. Mean absolute percentage error (MAPE) expresses errors as a proportion of actual enrollment, enabling comparison across regions of different sizes.

2.4 Volatility Classification Framework

To characterize enrollment instability, regional enrollment totals were first aggregated across all school sectors for each academic year. This produced a single annual enrollment time series per region, enabling consistent comparison of temporal variability.

Enrollment volatility was quantified using the coefficient of variation (CV) [14–16], defined as

$$CV_r = \frac{\sigma_r}{\mu_r} \times 100\%, \quad (2)$$

where μ_r represents the mean enrollment of region r over the observation window and σ_r denotes the corresponding sample standard deviation.

Regions were classified by comparing their CV values against a threshold derived from the empirical distribution of CV values across all 17 regions. Specifically, the 75th percentile of that empirical distribution was used as the classification boundary, yielding a cutoff of 127.10%. This percentile-based approach was chosen over a fixed arbitrary threshold because it adapts to the distributional properties of the data at hand. A fixed threshold, such as 10% or 20%, would impose an external standard that may not reflect the relative spread observed in this specific dataset. By contrast, the 75th percentile identifies the top quarter of regions by relative enrollment variability, regardless of the absolute scale, making the classification inherently data-driven and internally referenced. Regions exceeding this threshold were labeled high-volatility; the remaining regions were classified as stable.

2.5 Statistical Significance Assessment

To evaluate whether differences in predictive performance between tree-based models were statistically meaningful, a bootstrap resampling procedure was applied to the held-out test data. Repeated samples were drawn with replacement, and prediction errors were computed for each model within each resample. The distribution of error differences was then examined to determine whether observed performance gaps were distinguishable from sampling variability at conventional confidence levels.

3 RESULTS

3.1 Enrollment Distribution

Regional K-12 enrollment across the Philippines exhibits a markedly right-skewed distribution, as shown in Figure 1, with a small subset of regions accounting for a disproportionate share of total learners. Highly urbanized and economically dominant regions contribute extreme upper-tail values, while the majority of regions cluster at substantially lower enrollment levels. This pronounced heterogeneity highlights the limitations of linear modeling assumptions and motivates the use of nonlinear approaches, particularly tree-based ensemble methods.




Figure 1: Distribution of regional K-12 enrollment in the training set, illustrating right-skewness and the presence of extreme outliers across regions.

Table 2: Model Performance on Pandemic-Era Test Data (SY 2019–2020 and SY 2020–2021, $n = 102$)

Model	R^2	MAE	RMSE	MAPE [†]
Random Forest	0.9874	46,313	78,753	39.28%
Decision Tree	0.9867	46,445	80,905	28.38%
Linear Regression	0.8362	196,321	284,212	2038.43%
SVR	−0.3254	486,184	808,461	2401.08%

[†] MAPE for LR and SVR is inflated by near-zero

private-sector enrollment records in the test set;

R^2 and MAE are the primary comparison metrics.

3.2 Model Performance Comparison

Predictive performance across all evaluated models is summarized in Table 2, which reports results on the pandemic-era test set (SY 2019–2020 and SY 2020–2021, $n = 102$). The comparison reveals a clear separation between tree-based models and regression-based baselines.

The Random Forest model achieved the highest explanatory power among all evaluated approaches ($R^2 = 0.9874$). When evaluated relative to the average regional enrollment of approximately 520,536 students, the MAE of 46,313 students corresponds to an error rate of approximately 8.9%, indicating strong predictive accuracy even under pandemic-induced disruption. Bootstrap resampling of test observations produced a narrow confidence interval for the MAE, confirming stable performance across the full range of regions and enrollment levels.

The Decision Tree produced results closely comparable to Random Forest, with bootstrap comparison of error distributions confirming no statistically meaningful difference in average predictive accuracy between the two tree-based methods. The Decision Tree recorded a lower MAPE (28.38% vs. 39.28%), though Random Forest achieved slightly better R^2 (0.9874 vs. 0.9867) and lower MAE and RMSE. The ensemble mechanism provides a practical planning advantage: by averaging across many trees, Random Forest is less susceptible to localized anomalies in regional enrollment data, making it more suitable for deployment across the full range of Philippine regions.

In contrast, Linear Regression performed substantially worse across all primary evaluation metrics ($R^2 = 0.8362$, MAE = 196,321), reflecting its limited capacity to represent the complex nonlinear interactions among regional characteristics, sectoral composition, and temporal factors that tree-based methods capture naturally. Support Vector Regression exhibited negative explanatory power ($R^2 = -0.3254$), indicating predictions inferior to a simple mean baseline, which confirms it is unsuitable for this forecasting task.

3.3 Training and Test Consistency

Model generalization was further examined by comparing training and test performance for the Random Forest model. As illustrated in Figure 2, predicted enrollment values closely align with observed outcomes across the full range of regional enrollment sizes. The scatter plot appears linear because a well-fitted model produces predictions that cluster tightly around the perfect-prediction diagonal ($\hat{y} = y$), which is itself a straight line; this visual pattern is a




Figure 2: Predicted versus actual enrollment for the Random Forest model, test set SY 2019–2020 and SY 2020–2021.




Figure 3: Random Forest feature importance values for the most influential enrollment predictors.

consequence of high accuracy, not evidence of a linear relationship between the features and enrollment.

The minimal gap between training ($R^2 = 0.9937$) and test ($R^2 = 0.9874$) explanatory power indicates limited overfitting and strong temporal generalization. By comparison, Linear Regression exhibits weaker performance in both settings ($R^2_{\text{train}} = 0.8456$, $R^2_{\text{test}} = 0.8362$), confirming that the observed accuracy gap reflects structural model limitations rather than data scarcity.

3.4 Feature Importance Analysis

Feature importance analysis provides insight into the predictors captured by the Random Forest model. As shown in Figure 3, the

Table 3: Feature Importance Summary from the Random Forest Model

Feature Category	Specific Feature	Importance (%)
Sector	Public Sector	68.0
Regional Indicators	CALABARZON (Region IV-A)	10.7
	Central Luzon (Region III)	5.7
	NCR	5.1
	Western Visayas (Region VI)	2.2
	CAR	2.2
	Central Visayas (Region VII)	2.1
	Bicol Region (Region V)	1.7
	All other regions	0.5
Combined Regional Total		30.2
Temporal Trends	Academic year variables	0.3
Other Features	Private sector, SUC-LUC, others	1.5

public sector indicator is the single most important predictor, accounting for approximately 68.0% of total importance.

The dominant importance of the sector feature is expected and reflects a well-understood structural property of Philippine K-12 enrollment: public schools enroll the vast majority of learners, so the model necessarily relies on sector membership to anchor its predictions to the correct order of magnitude. This confirms that the model has correctly learned the known enrollment composition of the system.

Table 3 summarizes the contribution of key regional and temporal features. Regional indicators collectively account for 30.2% of total model importance, with economically and demographically significant regions contributing the largest individual effects. Temporal variables, including academic year boundaries, contribute minimally ($\sim 0.3\%$), suggesting that the model primarily learns cross-sectional enrollment patterns from region and sector identifiers rather than year-over-year trends. Private sector enrollment, SUC-LUC institutions, and other features collectively account for 1.5% of total importance.

Two constraints of the feature design are relevant to interpretation. First, all features were drawn exclusively from DepEd administrative records (region code, school sector, and academic year), deliberately excluding economic or demographic covariates such as regional GDP, internal migration, or fertility rates [6]. This choice keeps the framework deployable within DepEd’s existing data infrastructure, but it means the model identifies enrollment patterns rather than explaining what drives them. Second, encoding the 17 administrative regions as categorical features allows the model to learn region-specific enrollment levels, but this encoding is not portable: applying the model to a different country or a reorganized set of regions would require full retraining on locally representative data.

3.5 Regional Volatility Classification

Enrollment volatility classifications derived from aggregated regional enrollment are presented in Table 4. A subset of regions exhibits substantially higher relative variability, exceeding the 75th-percentile CV threshold of 127.10% used to distinguish high-volatility from stable regions.

Table 4: Regional Enrollment Volatility Based on Aggregate Enrollment (SY 2010–2011 to SY 2020–2021)

Region	Mean	Std Dev	CV (%)	Classification
<i>High-Volatility Regions (CV > 127.10%, 75th-percentile threshold)</i>				
Region VIII – Eastern Visayas	400,238	522,282	130.49	High Volatility
Region IX – Zamboanga Peninsula	319,686	411,035	128.57	High Volatility
CARAGA	236,374	301,531	127.57	High Volatility
Region V – Bicol Region	557,673	709,459	127.22	High Volatility
<i>Stable Regions (CV ≤ 127.10%)</i>				
Region IV-B – MIMAROPA	274,787	349,261	127.10	Stable
BARMM	296,441	373,035	125.84	Stable
Region VI – Western Visayas	635,622	777,369	122.30	Stable
Region XII – Soccsksargen	383,004	463,189	120.94	Stable
Region II – Cagayan Valley	273,955	330,261	120.55	Stable
Region XI – Davao Region	406,648	483,880	118.99	Stable
Region X – Northern Mindanao	394,849	466,661	118.19	Stable
Region I – Ilocos Region	404,872	477,036	117.82	Stable
Region VII – Central Visayas	634,556	741,437	116.84	Stable
Region III – Central Luzon	886,476	965,705	108.94	Stable
CAR – Cordillera Administrative Region	139,065	147,019	105.72	Stable
Region IV-A – CALABARZON	1,127,826	1,165,370	103.33	Stable
NCR – National Capital Region	902,568	858,243	95.09	Stable

The four high-volatility regions share contextual characteristics that are broadly consistent with their elevated enrollment variability. Eastern Visayas falls within one of the Philippines’ most typhoon-exposed geographic corridors [4], and the observed enrollment fluctuations in that region align temporally with periods of significant weather disruption. CARAGA and Zamboanga Peninsula carry documented records of socio-economic fragility [4], which may translate into less stable schooling participation over time. Among the stable regions, Western Visayas hosts several major urban centers whose diversified economies and established infrastructure appear to buffer enrollment from the kinds of external shocks that affect more vulnerable regions.

It should be noted that the regional-level aggregation used in this analysis masks intra-regional disparities between urban and rural districts within each region. A region classified as stable at the aggregate level may still contain individual divisions or schools with highly volatile enrollment. School-level or division-level analysis would provide more precise planning guidance but requires finer-grained data not yet available in DepEd’s published administrative datasets [2].

3.6 Model Accuracy Across Region Sizes

Prediction accuracy varies systematically with regional enrollment scale, as summarized in Table 5. Small regions (4 regions) achieve an average MAE of 28,341 students and an R^2 of 0.9823, reflecting lower absolute errors but slightly reduced goodness-of-fit consistent with greater sensitivity to localized shocks and limited historical observations. Medium regions (8 regions) record an average MAE of 42,189 and R^2 of 0.9867. Large regions (5 regions) exhibit the highest absolute errors at 73,492 students on average, yet achieve the strongest explanatory power at $R^2 = 0.9912$, reflecting that proportional accuracy improves with enrollment scale even as absolute deviations grow.

High-volatility regions do not exhibit systematically worse prediction accuracy, indicating that the Random Forest model effectively captures irregular enrollment patterns without introducing bias.

Table 5: Random Forest Prediction Accuracy by Region Size

Region Size Category	Number of Regions	Average MAE	Average R^2
Small	4	28,341	0.9823
Medium	8	42,189	0.9867
Large	5	73,492	0.9912

3.7 Implications for Resource Allocation

Applying the combined forecasting and volatility framework to medium-term planning scenarios illustrates how differentiated allocation strategies can be operationalized in practice. Rather than relying on uniform planning assumptions across all regions, the framework enables planners to distinguish between areas where enrollment demand follows relatively predictable trajectories and those where fluctuations are more pronounced. Regions characterized by high enrollment volatility warrant larger contingency buffers to accommodate rapid staffing adjustments, temporary classroom deployment, and short-term capacity reallocation in response to unexpected shocks. Table 6 illustrates this contrast using Eastern Visayas (high-volatility) and Central Luzon (stable) as representative examples.

Table 6: Illustrative Teacher Deployment Under Volatility-Aware Planning

Region	Class	Forecast	Allocation	Buffer	Total
Eastern Visayas	High	1,180,000	35,758	+15%	41,122
Central Luzon	Stable	3,300,000	100,000	+5%	105,000

The 15–20% contingency buffer recommendation is grounded in the forecast error characteristics of the Random Forest model. With a test-set MAE of 46,313 students against an average regional enrollment of 520,536, the model’s inherent error floor corresponds to approximately 8.9% of average enrollment. A contingency buffer must at minimum exceed this floor to provide meaningful planning headroom beyond what the model already accounts for. Table 7 presents a sensitivity analysis across five candidate buffer sizes. Buffers below 10% fall at or within the model’s own MAE range

Table 7: Sensitivity of Contingency Buffer Adequacy to Buffer Size (MAE = 46,313; mean enrollment of high-volatility regions = 378,493)

Buffer (%)	Extra Capacity (students)	MAE Headroom (students)	Volatile Coverage	Risk Level
5	18,925	-27,388	No	High (under-provisions)
10	37,849	-8,464	No	Moderate (thin margin)
15	56,774	+10,461	Most cases	Low (recommended floor)
20	75,699	+29,386	Extreme cases	Low (recommended ceiling)
25	94,623	+48,310	Yes (excess)	Wasteful (over-provisions)

Extra capacity = buffer% × mean enrollment of high-volatility regions. MAE headroom = extra capacity − MAE. Negative headroom indicates the buffer is consumed by forecast error before providing any planning margin.

and therefore offer no practical safety margin. A 15% buffer yields approximately 10,461 students of headroom above the MAE, covering the typical prediction error for most high-volatility regions. A 20% buffer extends this headroom to approximately 29,386 students, accommodating compounded errors in the most volatile region (Eastern Visayas, CV = 130.49%), which lies furthest above the 127.10% classification threshold. Buffers above 20% provide diminishing returns relative to their fiscal cost and are more likely to result in systematic over-provisioning in stable regions. The 15–20% range thus represents the interval where the buffer first meaningfully clears the forecast error floor without tipping into wasteful over-allocation. A formally optimized buffer, one incorporating utility functions, per-region fiscal constraints, and the asymmetric social costs of under-provisioning, remains an important direction for future work (Section 4.1, item 3).

In contrast, regions exhibiting more stable enrollment patterns can be managed using leaner reserve margins, closer to the 5% range, without substantially increasing operational risk, given that their enrollment trajectories fall within the model’s reliable prediction range. This differentiation allows limited resources to be allocated more efficiently, concentrating flexibility and adaptive capacity where uncertainty is highest.

Complementary infrastructure strategies further distinguish between region types, emphasizing flexibility in volatile contexts and long-term investment in stable growth regions.

4 CONCLUSION

The results of this study indicate that combining enrollment forecasting with volatility classification produces planning-relevant insights that neither approach yields on its own. The Random Forest model’s stable performance across the pandemic-era test set suggests that pattern-based forecasting retains utility even when enrollment behavior is disrupted by events outside the training window. The framework is not designed to substitute for institutional planning judgment; rather, it supplies a quantitative layer of regional differentiation that static, uniform projection methods are not structured to provide.

A key finding is that enrollment volatility is unevenly distributed across the Philippines’ 17 regions, and that this unevenness has direct consequences for resource allocation strategy. The four regions whose CV exceeded the 75th-percentile threshold of 127.10% require planning approaches that account for year-to-year fluctuation, including larger contingency reserves and more flexible staffing mechanisms. The remaining 13 stable regions can operate on tighter margins without substantially increasing the risk of

resource shortfall. Treating all regions identically, as centralized uniform allocation tends to do, misallocates resources in both directions: it over-buffers stable regions and under-prepares volatile ones.

The feature importance results further confirm that public school enrollment drives the scale of regional totals in the Philippines. This is consistent with the known enrollment composition of the Philippine K-12 system and underlines why DepEd planning efforts focused on public school capacity are likely to have the largest impact on system-wide outcomes.

Beyond its technical contributions, the study carries broader policy significance. The proposed framework aligns with international calls for adaptive, evidence-informed, and risk-sensitive education planning [5], offering a concrete mechanism through which such principles can be operationalized. Machine learning models inform the scale and direction of resource needs, while human decision-makers retain responsibility for interpreting results in light of local context, policy constraints, and equity considerations [7].

4.1 Future Work

- (1) **Granular forecasting.** Apply similar methods to division- and school-level planning once data access improves. Hierarchical models could cascade forecasts: national → regional → division → school.
- (2) **Richer feature engineering and baseline comparison.** Incorporate external socio-economic predictors (regional GDP, migration flows, fertility rates) and introduce a prior-year persistence baseline for explicit comparison with the machine learning models.
- (3) **Multi-objective optimization and cost-of-error analysis.** Extend beyond enrollment forecasting to incorporate learning outcomes, equity metrics, and fiscal constraints. A formal cost-of-error framework would quantify the trade-off between over-provisioning and under-provisioning, providing an empirical basis for the contingency buffer recommendation.
- (4) **Real-time adaptive systems.** Develop automated pipelines that retrain models quarterly, enabling continuous refinement. Integrate with early warning systems that trigger contingency plans when observed enrollment deviates from predictions.
- (5) **Cross-country validation.** Apply this framework to other Southeast Asian education systems (Indonesia, Vietnam, Thailand) to test generalizability and identify region-specific vs. universal enrollment drivers.

REFERENCES

- [1] L. Shao, C. Stripling, and P. Rayon, “Machine Learning Methods for Course Enrollment Prediction,” *Strategic Enrollment Management Quarterly*, vol. 10, no. 1, pp. 29–38, 2022.
- [2] Department of Education Philippines, “Learner Information System Data, SY 2023–2024,” 2024. [Online]. Available: <https://www.deped.gov.ph>
- [3] Trinity University of Asia, “Philippine Education Today: Challenges and Opportunities,” June 2024. [Online]. Available: <https://www.tua.edu.ph/blogs/philippine-education-statistics>
- [4] World Bank, “Philippines Basic Education Public Expenditure Review,” Report No. AUS0000920, 2021.
- [5] UNESCO International Institute for Educational Planning, “The Future of Educational Planning: The Changing Context,” Paris, 2023. [Online]. Available: <https://unesdoc.unesco.org/ark:/48223/pf0000396617>

- [6] Philippine Institute for Development Studies, "Falling fertility rate may ease school crowding, but it's not enough," PIDS News, June 2024. [Online]. Available: <https://www.pids.gov.ph>
- [7] A. Acharya and A. Vista, "Lessons in using data to improve education: The Philippines," Brookings Center for Universal Education, Jan. 2019. [Online]. Available: <https://www.brookings.edu/articles/lessons-in-using-data-to-improve-education-the-philippines/>
- [8] U. P. Inyang and E. A. Johnson, "Performance Comparison of XGBoost and Random Forest for the Prediction of Students Academic Performance," *European Journal of Computer Science and Information Technology*, vol. 13, no. 2, pp. 1–21, 2025.
- [9] M. Kumar et al., "Utilizing Random Forest and XGBoost Data Mining Algorithms for Anticipating Students' Academic Performance," *International Journal of Modern Education and Computer Science*, vol. 16, no. 2, pp. 35–50, 2024.
- [10] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [11] A. King, "Improving the Enrollment Process through Machine Learning," *EDUCAUSE Review*, Jan. 2019. [Online]. Available: <https://er.educause.edu/blogs/sponsored/2019/1/improving-the-enrollment-process-through-machine-learning>
- [12] Official Gazette of the Republic of the Philippines, "The K to 12 Basic Education Program," 2016. [Online]. Available: <https://www.officialgazette.gov.ph/k-12/>
- [13] Department of Education Philippines, "Programs and Projects Status of Implementation," 2019–2020 Report.
- [14] Corporate Finance Institute, "Coefficient of Variation," Aug. 2024. [Online]. Available: <https://corporatefinanceinstitute.com/resources/data-science/coefficient-of-variation/>
- [15] J. Frost, "Coefficient of Variation in Statistics," *Statistics By Jim*, Oct. 2020. [Online]. Available: <https://statisticsbyjim.com/basics/coefficient-variation/>
- [16] SixSigma.us, "Coefficient of Variation: Mastering Relative Variability in Statistics," Apr. 2025. [Online]. Available: <https://www.6sigma.us/six-sigma-in-focus/coefficient-of-variation/>

Defending against Membership Inference Attacks through Contrastive Representation Learning

Abien Fred Agarap
Globe Fintech Innovations Inc.
Taguig, Metro Manila, Philippines
abienfred.agarap@gcash.com

Miguel Anton Aguinaldo
Globe Fintech Innovations Inc.
Taguig, Metro Manila, Philippines
miguel-a.aguinaldo@gcash.com

Rafael Madrigal
Globe Fintech Innovations Inc.
Taguig, Metro Manila, Philippines
rafael.madrigal@gcash.com

Inigo Miguel Benavides
Globe Fintech Innovations Inc.
Taguig, Metro Manila, Philippines
inigo.benavides@gcash.com

Sara Ann Venturina
Globe Fintech Innovations Inc.
Taguig, Metro Manila, Philippines
saraann.venturina@gcash.com

ABSTRACT

The use of large-scale machine learning models in sensitive sectors such as finance and healthcare risk privacy breaches through Membership Inference Attacks (MIAs). These attacks exploit the tendency of a model to overfit, allowing an adversary to check if a specific data point was part of the training set. We propose a defense strategy that combines private synthetic data generation and contrastive representation learning. First, we generate synthetic decoy datasets using the Adaptive and Iterative Mechanism, and analyze the relationship between synthetic data quality and MIA risk. Our core defense involves training models with a joint loss function that combines a primary classification task with a contrastive learning component using the soft nearest neighbor loss and Kullback-Leibler Divergence. By *entangling* these representations, we render real and synthetic data points indistinguishable to an adversary. Results show that our approach significantly reduces the efficacy of a highly effective MIA such as Likelihood Ratio Attack (LiRA). Our approach reduces the TPR at 1% FPR of LiRA to 0.00% on the UCI Bank Marketing, Mushrooms, and Statlog German Credit datasets at $p < 0.05$, while performing on par with our baseline models in their primary tasks of binary classification.

CCS CONCEPTS

• **Security and privacy** → **Privacy protections**; • **Computing methodologies** → **Neural networks**;

KEYWORDS

contrastive learning, deep learning, membership inference attacks, privacy attacks

1 INTRODUCTION

As we scale machine learning models to adapt to the growing needs of our organizations and our users, the concerns regarding their privacy and security inevitably follow. With their ability to generalize across diverse data, these large-scale models are now widely adopted across sensitive sectors such as finance [3, 24], healthcare [12, 27], and autonomous systems [20], where working with carefully managed data is essential. However, with such strong generalization capability comes the risk of memorizing or overfitting on data, even an entirely randomly labeled ones [40]. The




Figure 1: Overview of our defense mechanism combining private synthetic data generation and a contrastive learning framework. Using some shadow models, an adversary checks if the resulting model loss originates from member or non-member distributions.

propensity of models to overfit is directly connected with the problem of membership inference. In cases where privacy concerns are of paramount importance such as in healthcare and finance, it is essential that models do not reveal any information regarding the inclusion of particular data points in their training datasets, thereby preventing any unintended exposure of private information.

A membership inference attack (MIA) checks whether a data point was used to train a machine learning model by exploiting the fact that machine learning models often yield higher probability estimates or lower loss values on previously seen data as opposed to previously unseen ones [35]. An adversary leverages these subtle differences to identify patterns that distinguish training data members from non-members. Thus, reducing overfitting to counteract MIAs has become a key approach in privacy research. Studies have explored differential privacy [9, 23], regularization [17, 18, 34, 36, 37], and adversarial training [30]. However, these methods may come with trade-offs like compromising model utility as differential privacy injects noise to uphold data privacy. We propose a synergy of differential privacy through synthetic data generation and contrastive learning.

1.1 Our contributions

We explore three defenses against MIAs. First, we use AIM to generate differentially private synthetic data to serve as decoys, and we analyze the trade-off between utility and privacy. Second, we apply the k-Winners-Take-All (kWTA) activation function to induce network sparsity. Finally, we use contrastive learning to entangle the learned representations of real and synthetic data as illustrated in Figure 1, making them indistinguishable to an adversary.

2 BACKGROUND

2.1 Membership Inference Attacks

MIA is a privacy attack in which an adversary aims to determine if specific data was used to train a target model. MIAs can be conducted in two ways: (1) in black-box settings where an adversary has only query access to the model (often via an API) and no internal knowledge of its design, and (2) in white-box settings where the full details of the model architecture and parameters are available to an adversary.

We use an updated definition of a security game for membership inference [8] that characterizes an adversary’s capabilities such as whether the adversary is allowed to train shadow models using query instances. So, this enables a clear distinction between non-adaptive (offline) and adaptive (online) settings for attack strategies, which yield different attack performance. We formally define our threat model as follows,

Definition 2.1. Membership Inference Security Game [8]. The following game between a challenger and an adversary assumes they both have access to a data distribution $\mathbb{D}$:

- (1) The challenger samples a training dataset $D \sim \mathbb{D}$, trains a target model $f_\theta \leftarrow \mathcal{T}(D)$ on the dataset D , and grants the adversary a query access to the model f_θ .
- (2) The challenger randomly selects a subset $D_a \subseteq D$ and samples a set D_b from $\mathbb{D}$. These two sets are then combined to create a query set: $D_{\text{query}} = D_a \cup D_b$, which the challenger sends to the adversary.
- (3) The adversary then responds with a set $D_g \subseteq D_{\text{query}}$ which represents that the adversary guesses that instances in D_g were used when training f_θ and instances in $D_{\text{query}} \setminus D_g$ were not used to train f_θ .

We consider the target model to be a neural network classifier in a black-box setting, i.e. the adversary is granted access to the target model f_θ but is not given access to its parameters. Concretely, the adversary can access the softmax output probabilities for any input sample given to the target model. We also assume that the adversary can construct a dataset by sampling from the data distribution $\mathbb{D}$, and use subsets of this dataset to train its shadow models. We employ the likelihood ratio attack as the MIA method [5].

2.1.1 Likelihood Ratio Attack (LiRA). The likelihood ratio attack (LiRA) [5] trains a number of “shadow models” to estimate two distributions: $Q_{\text{in}}(x, y) = \{f \leftarrow \mathcal{T}(D \cup \{(x, y)\}) \mid D \leftarrow \mathbb{D}\}$ denoting the distribution of models trained on datasets containing the target data (x, y) , and then $Q_{\text{out}}(x, y) = \{f \leftarrow \mathcal{T}(D \setminus \{(x, y)\}) \mid D \leftarrow \mathbb{D}\}$ denoting the distribution of models trained on datasets not containing the same target data (x, y) .

Since determining whether a target model f is sampled from either Q_{in} or Q_{out} is intractable, the attack is simplified by defining $\tilde{Q}_{\text{in}}$ and $\tilde{Q}_{\text{out}}$ as the distributions of shadow model losses on the target data (x, y) for models that were trained and not trained on this data point respectively. With this simplification, given the target model loss $\ell(f(x), y)$ on the target example (x, y) , an adversary tries to predict if $\ell(f(x), y)$ was sampled either from $\tilde{Q}_{\text{in}}$ or from $\tilde{Q}_{\text{out}}$, which is formally defined as the following likelihood ratio test: $p(\ell(f(x), y) \mid \tilde{Q}_{\text{in, out}}(x, y))$. The adversary then returns whichever is more likely.

In this work, we trained 16-32 shadow models per target model that was trained on the real and on the mixture of real and synthetically generated data points, and one attacker model per target model. We settled with this design choice due to computational constraints as each shadow model is trained on shadow datasets that are similarly structured to the target model training datasets. Additionally, as [5] proposed, we use their metric of true-positive rate (TPR) at a low (1%) false-positive rate (FPR) to measure the efficacy of an MIA on target models.

2.2 Defenses against MIAs

The two most prominent defenses against MIAs are differential privacy and regularization, which attempt to reduce the overfitting of machine learning models thus reducing the efficacy of MIAs. We focus on synthetic data generation (SDG) with differential privacy using Adaptive Iterative Mechanism (AIM) [26] and regularization [35, 39].

2.2.1 Private Synthetic Data Generation. We generate synthetic versions of five benchmark tabular datasets, and we use both (1) the real and (2) the mixture of real and synthetically generated datasets to train our target models, shadow models, and attacker models. To help us understand the implication of synthetic data quality w.r.t. MIA risks, we present an implication matrix in Table 1.

		Implication (Data Quality, MIA Risk)	
Low Fidelity	High Fidelity	Garbage Data, Very Low	Intelligent Filter, Very Low
	Low Fidelity	Blurry Data, Low	Ideal Replica, Moderate to High
		Low Utility	High Utility

Table 1: Implication matrix for data quality and MIA risk based on fidelity and utility of synthetically generated datasets

We discuss the entries in this data quality-MIA risk implication matrix as follows,

- (1) **Garbage Data** (Very Low risk of MIA). The synthetic data is statistically divergent and information-deprived.
- (2) **Blurry Data** (Low risk of MIA). While statistically representative, the synthetic data lacks fine-grained details, thus trading utility for lower MIA risk.
- (3) **Intelligent Filter** (Very Low risk of MIA). The generator isolated predictive signals by removing irrelevant noise and the individual “fingerprints” that MIAs exploit.
- (4) **Ideal Replica** (Moderate to High risk of MIA). This “gold standard” synthetic data is statistically similar and highly useful but its realism creates a privacy risk which may leak information if unprotected.

Dataset	# Samples	Features	Fidelity (JSD, TUC) ↓	Utility ↑	Implication (Data Quality, MIA risk)
Bank Marketing	11,162	17	(0.0032, 0.0142)	0.90	(Ideal Replica, Moderate to High)
Iranian Churn	3,150	14	(0.0137, 0.0875)	0.46	(Garbage Data, Very Low)
Mushrooms	8,124	23	(0.0011, 0.0747)	0.96	(Intelligent Filter, Very Low)
Online Shoppers	12,330	18	(0.0044, 0.0224)	0.83	(Ideal Replica, Moderate to High)
Statlog German Credit	1,000	21	(0.0060, 0.0155)	0.71	(Intelligent Filter, Very Low)

Table 2: Summary of benchmark datasets including results on data fidelity and utility evaluations between real and synthetic data across the five datasets. Low JS Divergence (JSD) and Theil’s U Coefficients (TUC) entail high fidelity while higher utility entail usability in classification tasks.

While these are descriptive interpretations for synthetic data quality implication on MIA risks, we use three quantitative metrics to evaluate the fidelity and utility of our synthetically generated datasets. Table 2 summarizes the datasets we use in this study along with the results of our SDG in terms of fidelity using Jensen-Shannon Divergence (JSD) and Theil’s U Coefficient (TUC), and utility using the average real-data test performance of top-3 classifiers trained on synthetic data, and the implication of synthetic data quality w.r.t. MIA risks.

We first preprocessed the datasets according to the feature types. Categorical features were discretized while continuous features were transformed into 32 equal-width bins [26] before discretization. Both the top-3 highly correlated features with their corresponding targets y , and the personal attributes were removed from the dataset.

We use the the OpenDP SmartNoise SDK¹ to implement the AIM algorithm for this work. All the datasets we use in this study are for binary classification tasks: (1) Bank Marketing [28] for whether a client will subscribe to a term deposit or not, (2) Statlog German Credit Data [16] for whether a customer has a good or bad credit risk, (3) Online Shoppers [32] for whether an online session has ended with a purchase or not, (4) Mushrooms [22] for whether a mushroom species is edible or poisonous, and (5) Iranian Churn [19] for whether a customer is likely to churn or not.

Table 2 assesses the fidelity (statistical similarity) and utility (machine learning performance) of our synthetic datasets. Based on these results, we classified the datasets into three types reflecting their quality and implied MIA risk: Ideal Replicas (Bank Marketing and Online Shoppers) with high fidelity and utility, suggesting higher MIA risk; Intelligent Filters (Mushrooms and Statlog German Credit) with moderate fidelity but high utility, indicating lower MIA risk; and Garbage Data (Iranian Churn) with low scores on both. The potential MIA vulnerability of the high-quality replicas justifies our use of additional regularization and contrastive learning defenses.

2.2.2 L_1 and L_2 Regularization. Regularization reduces overfitting of machine learning models to training data, thus rendering training data members and non-members less distinguishable [35]. In this work, we employ both L_1 and L_2 regularization that penalizes large model parameters θ by adding $\lambda \sum_{i=0}^m |\theta_i|$ or $\lambda \sum_{i=0}^m \theta_i^2$ to the loss function of a model. We then used a fixed value of 1×10^{-3} for the λ parameter for both L_1 - and L_2 - norm. This is to avoid severe performance degradation.




Figure 2: Lower TPR at FPR values indicate better defense: Scatter plots illustrating the TPR at 1% FPR as a function of k . On each dataset, we trained 64 different neural networks with their own k values. Then we performed LiRA MIA with 16 shadow models for each target model so we can explore the effect of k -values in mitigating MIAs.

2.2.3 k -Winners-Take-All activation. Similar to L_1 and L_2 regularization, the k -Winners-Take-All (kWTA) activation function introduces sparsity within neural networks. It achieves this by retaining only the top- k winning neurons within hidden layers [25, 31]. The kWTA activation function admits $k \geq 1$ winning neurons in a hidden layer j of a neural network, $\phi_k(z)_j = z_j \cdot \mathbb{1}(z_j \in \{\max z\}_k)$, where z is an activation output, and k is the percentage of winning neurons to retain. We set $k = 0.75$ in our comparative experiment with other target models to maintain a reasonable classification performance while aiming for membership inference defense. However, we do note that k is a tunable hyperparameter. In Figure 2, we explore $k_{i=1}^{i=64} \in [0.0, 1.0]^{\mathbb{R}}$ different values per (mixture of real and synthetic) dataset, and their effect in mitigating MIAs as measured by the TPR at 1% FPR of attacker models.

For the Mushrooms and Statlog German Credit dataset – both Intelligent Filter datasets, we see that for virtually all k -values, the target models were able to maintain low TPR at 1% FPR values. Meanwhile, for the Bank Marketing, Iranian Churn, and Online Shoppers dataset – a mixture of Ideal Replica and Garbage Data, we see no discernible pattern between varying k -values and TPR at 1% FPR values. At $k = 0.75$ as the boundary, our Mann-Whitney U tests yield no statistical significance in terms of TPR at FPR as a function of k .

¹<https://github.com/opendp/smartnoise-sdk>

3 METHODOLOGY

Neural networks learn good data representations which may involve entangling data points for better performance [1, 4, 7, 10, 14, 15, 38]. We define *entanglement* as the process of pulling structurally similar data points closer together. Entangling the learned representations of real and synthetic data will engender difficulty in membership inference. We accomplish this with a joint loss function that learns the primary task of a neural network along with a contrastive learning approach to minimize the distances between real and synthetic data points.




Figure 3: t-SNE visualization of learned representations of a neural network trained on the mixture of real and synthetic Mushrooms datasets. The KLD loss values of the **real** and **synthetic** samples, and of the **class 0** and **class 1** samples are measured during training. The KLD Loss for the membership decreases as training progresses, indicating minimization of distances between real and synthetic data points. Meanwhile, the KLD Loss for the classes increases as training progresses, indicating learning of the primary task of binary classification. While the distances between real and synthetic data points minimize, they also maximize w.r.t. the dataset classes.

3.1 Soft Nearest Neighbor Loss

The soft nearest neighbor loss [33] measures the *entanglement* or how close the structurally- or class- similar data points to each other are, relative to data points from different structures or classes. Formally defined for a batch b of samples as,

$$\mathcal{L}_{\text{sn}}(f, x, y, T) = -\frac{1}{b} \sum_{i \in 1 \dots b} \log \left(\frac{\sum_{\substack{j \in 1 \dots b \\ j \neq i, y_i = y_j}} \exp(-d(x_i, x_j)/T)}{\sum_{\substack{k \in 1 \dots b \\ k \neq i}} \exp(-d(x_i, x_k)/T)} \right) \quad (1)$$

where x may be the raw input features or the learned feature representations in the hidden layers of a neural network f , y is the class or label, T is a temperature parameter, and d is any distance metric. The temperature T dictates how distance influences representation learning, with farther distances requiring higher temperatures. We found that annealing the temperature from an initial value of $T = 1.0$ was more effective than using a fixed value [2]. We use L_1 distance due to the nature of the tabular datasets in this work, and with AIM optimizing for L_1 error metric [26]. We also find that using lowest soft nearest neighbor loss among the hidden layers of our neural networks lead to more stable computations during training.

In Figure 3, we show the effect of the soft nearest neighbor loss on entangling real and synthetic data points in juxtaposition with learning the primary task of binary classification. While the distances between the real and synthetic data points shrink, the gap between the dataset classes broadens.

3.2 KL Divergence Loss

In addition to soft nearest neighbor loss, we also apply Kullback-Leibler Divergence (KLD) as a loss function. The goal is to further entangle real and synthetic data points, thereby strengthening model defense against MIAs. It is formally defined as follows: $D_{\text{KL}}(P||Q) = \sum P(x) \log P(x) - \log Q(x)$.

In this context, $P(x)$ denotes the distribution of real data, and $Q(x)$ the distribution of synthetic data. By minimizing this divergence, we can ensure that the distributions of real and synthetic data points become similar upon training. We aim to render synthetic data indistinguishable from real data when probed by an adversary. We posit that this will hinder the ability of an adversary to infer training data membership.

Since KL Divergence is asymmetric, we minimize both $D_{\text{KL}}(P||Q)$ and $D_{\text{KL}}(Q||P)$: $\mathcal{L}_{\text{KLD}}(P, Q) = \frac{1}{2} [D_{\text{KL}}(P||Q) + D_{\text{KL}}(Q||P)]$. By reducing this bidirectional KLD loss, we force $Q(x)$ to avoid producing data points that are noticeably outside $P(x)$, in addition to covering $P(x)$. We see this bidirectional KLD loss as necessary to blur the boundaries between real and synthetic data.

3.3 Entangling Real and Synthetic Data

We define our joint loss function to entangle real and synthetic data to consist of the binary cross entropy loss $\mathcal{L}_{\text{BCE}}$ as for the primary task of binary classification across the five benchmark datasets, the soft nearest neighbor loss $\mathcal{L}_{\text{SNN}}$ as the contrastive learning component applied across all hidden layers of a neural network, and the symmetric KLD loss $\mathcal{L}_{\text{KLD}}$ to ensure the distributions between real and synthetic data points align, thus blurring their statistical “footprints”. The joint loss function is as follows,

$$\mathcal{L}(f, x, y, m) = \mathcal{L}_{\text{BCE}}(y, f^k(x)) + \alpha \cdot \min_{\mathcal{L}_{\text{SNN}}} \mathcal{L}_{\text{SNN}}(f^i(x), m) + \lambda \mathcal{L}_{\text{KLD}}(P, Q) \quad (2)$$

where α dictates the contrastive learning computation, e.g. a negative value implies disentanglement (maximization) of distances between data points of a given class (e.g. real or synthetic, or target class labels k), λ is a tunable hyperparameter that controls the amount of influence of the KLD loss to the joint loss function. We use $\alpha = \{-100.0, 100.0\}$ in our experiments but we also provide an exploration of varying α values within this range in Figure 4, and a fixed value of $\lambda = 0.5$.

By optimizing this joint loss function, our target models learn representations that are both useful for the primary task and are also challenging for adversaries attempting to infer membership.

3.4 Mathematical Analysis

We provide a theoretical justification for the efficacy of the soft nearest neighbor loss in mitigating MIAs. The core premise of an MIA, particularly LiRA, is the adversary’s ability to distinguish between the loss distributions of member samples, $\tilde{Q}_{in}$, and non-member samples, $\tilde{Q}_{out}$. The success of the attack is fundamentally

bounded by the statistical divergence between these two distributions; if $D(\hat{Q}_{in}||\hat{Q}_{out}) \approx 0$, the adversary cannot perform better than random guessing.

Our defense introduces a synthetic dataset D_{synth} into the training process to serve as a decoy. By minimizing the joint loss function $\mathcal{L} = \mathcal{L}_{BCE} + \alpha\mathcal{L}_{SNN} + \lambda\mathcal{L}_{KLD}$ with a positive α , we explicitly minimize the distance $d(x_i, x_j)$ between real samples $x_i \in D_{real}$ and synthetic samples $x_j \in D_{synth}$ within the latent feature space of the model f_θ .

Eq. 1 minimizes the intra-class variance by maximizing the likelihood that a sample’s nearest neighbors belong to the same target class y . When a training batch b consists of a mixture $D_{real} \cup D_{synth}$, minimizing $\mathcal{L}_{SNN}$ forces the latent representations to overlap regardless of their origin (real or synthetic). As the temperature T anneals, the representations of real data points $f(D_{real})$ are pulled into the local neighborhoods of synthetic data points $f(D_{synth})$. Consequently, the feature distributions converge, $P(f(x)|x \in D_{real}) \rightarrow P(f(x)|x \in D_{synth})$.

Since the synthetic data acts as a proxy for the generalized data distribution (non-members), this entanglement obscures the unique “fingerprint” of the training data. The adversary observing the loss $\ell(f(x), y)$ faces a signal-to-noise ratio where the signal (membership trace) is drowned out by the noise of the synthetic decoy variance. This mathematical blurring of boundaries explains the empirical result where the TPR at 1% FPR drops to 0.00% on high-risk datasets.

4 EXPERIMENTS

To demonstrate the effectiveness of our approach, we conducted experiments on five benchmark datasets having two sets each: (1) real datasets (denoted as “Real”) and (2) mixture of real dataset and synthetically generated data using AIM (denoted as “Mixed”). We mainly report the classification performance of our target models, and the TPR at low FPR (1%) of LiRA MIA on all the target models where each target model was trained 30 different times (each time using a different random seed) per dataset.

To better compare our baseline and experimental models, we ran Kruskal-Wallis H tests on their classification performance and on the TPR at 1% FPR results of LiRA MIA for each benchmark dataset. We report our statistical tests for $p < 0.05$ significance.

4.1 Target models

We analyze the performance of our target classifier models under two distinct training conditions: using Real and using Mixed datasets. Across our experiments for these datasets, we used 8 different target models. All these models follow the same architecture of a fully connected network with $d - 600 - 300 - 150 - 1$ units, where d refers to the dimensionality of the input features. All the hidden layers use a ReLU [29], except for the kWTA neural network, and all hidden layers were initialized with He initialization [13]. Finally, we trained all our models using Adam [21] with a learning rate of 1×10^{-3} for 100 epochs on all our datasets.

We defined our baseline approaches as (1) a fully connected network as described above, (2) a neural network with L_1 regularization, and (3) an neural network with L_2 regularization. For our experimental models, they were (4) a kWTA neural network, (5) a

Disentanglement model that uses the following objective function: $\mathcal{L}_{BCE} - \alpha\mathcal{L}_{SNN}$, (6) a Disentanglement model w/ KLD, (7) an Entanglement model, and (8) an Entanglement model w/ KLD that uses our joint loss function, $\mathcal{L}_{BCE} + \alpha\mathcal{L}_{SNN} + \lambda\mathcal{L}_{KLD}$.

Table 3 compares our baseline and experimental models in terms of predictive classification performance. In general, we observe a decrease in average accuracy rate for most models when trained on mixed datasets when compared to training only on the real datasets. This corroborates the utility scores of our synthetic datasets as per Table 2. While SDG using AIM can augment our training datasets, preliminary analysis of average accuracy rates indicates that for most models, training exclusively on real data yielded higher average accuracy. Our Kruskal-Wallis H Test (denoted as KWt) results confirm that there are significant differences in performance among the target models trained and evaluated under both real and mixed data. However, training exclusively on the real datasets may have inherently higher risk for training data membership leakage.

4.2 Attacking the target models

We evaluate the effectiveness of our proposed approaches to defend against MIAs by introducing sparsity through kWTA activation function, and by entangling or disentangling real and mixed data points during training. We employed LiRA MIA across our experiments, where we trained 32 shadow models per target model. The full results are in Table 4.

The primary observation in Table 4 is the consistent and significant reduction in MIA efficacy, as measured by TPR at 1% FPR, across all five benchmark datasets when our proposed defense mechanisms are employed. In the mixed dataset training scenario, which incorporates our AIM-generated synthetic data alongside real data, we see a drastic decrease in attack TPR. This contrasts our baseline models where LiRA frequently achieves high TPRs. For instance, on the mixed Bank Marketing dataset, the baseline model that uses L_2 regularization yield an attack TPR of 1.03. Meanwhile, our experimental model that uses entanglement w/ KLD reduces this to a TPR of 0.00. This trend is uniformly observed, notably in our entanglement and entanglement w/ KLD models.

Our entanglement defense proves effective regardless of the synthetic data quality as classified in Table 2. For high-risk “Ideal Replica” datasets Bank Marketing and Online Shoppers, it successfully decreased the attack TPR to as low as 0.00 and 0.85 respectively. For lower-risk “Intelligent Filter” and “Garbage Data” datasets, it also ensures near-perfect protection. This shows that our approach is a robust and widely applicable defense.

Table 5 presents the KWt results to show if training on mixed datasets significantly affects the vulnerability of a target model to MIAs when compared to using real data alone. While most individual models showed no significant change ($p \geq 0.05$), there were notable exceptions. More importantly, the overall results for most datasets show a statistically significant difference, indicating that the use of mixed data is not a neutral act. It measurably alters the privacy of a model, which underscores the need for careful implementation and additional defenses.

Model	Bank Marketing		Mushrooms		Iranian Churn	
	ACC (Mean (SD) in %)		ACC (Mean (SD) in %)		ACC (Mean (SD) in %)	
	Real	Mixed	Real	Mixed	Real	Mixed
Baseline	58.42 (2.47)	58.23 (2.70)	97.79 (1.18)	90.36 (1.69)	87.47 (2.23)	82.14 (2.24)
w/ L1	58.79 (2.83)	58.16 (2.83)	97.49 (1.13)	90.34 (2.01)	88.32 (2.01)	80.62 (3.38)
w/ L2	57.53 (4.49)	57.80 (4.69)	97.17 (1.25)	90.22 (2.28)	84.71 (1.72)	84.36 (1.78)
kWTA NN	58.68 (3.42)	58.78 (3.90)	98.01 (1.04)	90.53 (1.77)	87.04 (2.21)	82.68 (2.62)
Disentangled	59.01 (2.93)	59.23 (3.22)	97.68 (1.09)	90.31 (1.96)	87.16 (2.26)	82.27 (2.66)
w/ KLD	56.56 (4.02)	56.86 (3.05)	69.57 (13.41)	72.42 (10.23)	85.17 (1.60)	84.18 (1.59)
Entangled	57.89 (3.96)	57.02 (3.31)	94.36 (1.33)	88.36 (1.74)	86.99 (1.67)	83.78 (1.84)
w/ KLD	58.39 (2.93)	58.11 (3.62)	94.97 (1.70)	88.46 (2.02)	86.80 (2.11)	83.85 (2.00)
KWt (H, p)	~ (9.32, 0.23)	~ (13.48, 0.06)	* (160.54, 0.00)	* (104.90, 0.00)	* (60.16, 0.00)	* (45.12, 0.00)

Model	Online Shoppers		Statlog German Credit	
	ACC (Mean (SD) in %)		ACC (Mean (SD) in %)	
	Real	Mixed	Real	Mixed
Baseline	78.41 (3.11)	79.19 (2.26)	63.56 (2.99)	60.94 (3.85)
w/ L1	77.10 (2.91)	77.80 (2.90)	62.15 (3.79)	60.34 (3.98)
w/ L2	84.44 (1.66)	84.28 (1.63)	68.69 (2.45)	68.59 (1.99)
kWTA NN	79.91 (2.44)	79.48 (2.19)	63.68 (3.60)	60.49 (3.43)
Disentangled	79.65 (3.14)	79.31 (2.34)	63.03 (3.06)	60.51 (3.23)
w/ KLD	81.64 (7.55)	83.79 (2.20)	69.14 (1.82)	69.97 (2.11)
Entangled	83.05 (2.07)	82.73 (1.64)	53.01 (5.04)	52.60 (7.16)
w/ KLD	83.26 (2.27)	82.73 (1.78)	52.65 (4.42)	53.16 (6.89)
KWt (H, p)	* (117.75, 0.00)	* (139.06, 0.00)	* (172.72, 0.00)	* (150.16, 0.00)

Table 3: Classification performance of the target models on the five datasets (bold values represent the best results), when trained solely on the real data and when trained on mixture of the real and synthetically generated data. We report the average ACC across 30 runs of each model, along with the standard deviation of the ACC scores (in %). Using our approaches, the target models perform on par with the baseline target models. * denotes statistical significance at $p < 0.05$ and ~ denotes otherwise.

4.3 Entanglement against MIAs

In Figure 3, we illustrate that while we entangle real and synthetic data points, they become disentangled w.r.t. class labels, i.e. the gap between classes 0 and 1 has broadened but the real and synthetic data points for each class remain entangled. This led us to explore even the disentanglement capability of soft nearest neighbor loss [2, 11], which is usually more beneficial to discriminative tasks. In Figure 4, we sample $\alpha_i^{n=64} \in [-100.0, 100.0]^{\mathbb{R}}$ to train neural networks to disentangle real and synthetic data points ($-\alpha$ values), and to entangle real and synthetic data points ($+\alpha$ values). We supplement Figure 4 with quantitative results in Table 6.

Figure 4 shows increasing the entanglement between real and synthetic data (controlled by a positive α factor) leads to lower TPRs which indicates a stronger defense against MIAs. This supports that entanglement is key to mitigating these attacks.

Our results reveal dataset-dependent effect of entanglement on MIA defense. While entanglement is generally effective, Table 6 shows that for the Mushrooms and German Credit datasets, disentanglement paradoxically yields a lower attack success rate. We believe this is due to a dual objective in the model: while it tries to entangle real and synthetic data, it simultaneously works to separate the data by its actual class labels. On these datasets, the strong class separation appears to incidentally obscure the distinction between real and synthetic data within those clusters. Thus, creating a more robust albeit unintentional MIA defense. This highlights that the optimal privacy strategy depends on the interplay between the method and the dataset characteristics.

5 CONCLUSION

We proposed a defense against MIAs through SDG using AIM, kWTA activation function to induce sparsity, and contrastive learning to entangle real and synthetic data representations. Results show




Figure 4: Lower TPR at FPR values indicate better defense: Scatter plots illustrating the TPR at 1% FPR as a function of α factor for the soft nearest neighbor loss, i.e. negative values (blue) disentangle the real and synthetic data points while positive values (red) entangle the real and synthetic data points. On each dataset, we trained 64 different neural networks with varying α values. Then we performed LiRA MIA with 16 shadow models each, to explore the effect of disentanglement and entanglement in mitigating MIAs.

this method significantly reduces attack success, even on high-risk “Ideal Replica” datasets. This work repositions contrastive learning as a defensive tool, opposing prior research as an attack method [6]. We also uncovered a nuanced interplay where strong class separation can incidentally provide a better defense than direct data entanglement, as seen in a counter-intuitive result on one dataset. In conclusion, our method offers a robust privacy-preserving mechanism without significantly compromising utility.

We suggest future works on balancing the mix between real data and synthetic data, on exploring other types of data (e.g. image, text, signal), and extension of our joint loss function to traditional machine learning algorithms.

REFERENCES

- [1] Abien Fred Agarap. 2022. Self-organizing cooperative neural network experts. (2022).
- [2] Abien Fred Agarap and Arnulfo P Azcarraga. 2020. Improving k-means clustering performance with disentangled internal representations. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [3] Shamima Ahmed, Muneer M Alshater, Anis El Ammari, and Helmi Hammami. 2022. Artificial intelligence and machine learning in finance: A bibliometric review. *Research in International Business and Finance* 61 (2022), 101646.
- [4] Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence* 35, 8 (2013), 1798–1828.
- [5] Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramèr. 2022. Membership inference attacks from first principles. In *2022 IEEE symposium on security and privacy (SP)*. IEEE, 1897–1914.
- [6] Depeng Chen, Xiao Liu, Jie Cui, and Hong Zhong. 2023. Poster: membership inference attacks via contrastive learning. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 3555–3557.
- [7] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. 2020. Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems* 33 (2020), 22243–22255.
- [8] Yuntao Du, Jiacheng Li, Yuetian Chen, Kaiyuan Zhang, Zhizhen Yuan, Hanshen Xiao, Bruno Ribeiro, and Ninghui Li. 2025. Cascading and Proxy Membership Inference Attacks. *arXiv preprint arXiv:2507.21412* (2025).
- [9] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*. Springer, 265–284.

- [10] Dumitru Erhan, Aaron Courville, Yoshua Bengio, and Pascal Vincent. 2010. Why does unsupervised pre-training help deep learning?. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 201–208.
- [11] Nicholas Frosst, Nicolas Papernot, and Geoffrey Hinton. 2019. Analyzing and improving representations with the soft nearest neighbor loss. In *International conference on machine learning*. PMLR, 2012–2020.
- [12] Hafsa Habebh and Suril Gohel. 2021. Machine learning in healthcare. *Current genomics* 22, 4 (2021), 291–300.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*. 1026–1034.
- [14] Geoffrey E Hinton and Ruslan R Salakhutdinov. 2006. Reducing the dimensionality of data with neural networks. *science* 313, 5786 (2006), 504–507.
- [15] Geoffrey E Hinton and Richard Zemel. 1993. Autoencoders, minimum description length and Helmholtz free energy. *Advances in neural information processing systems* 6 (1993).
- [16] Hans Hofmann. 1994. Statlog (German Credit Data). UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5NC77>.
- [17] Hongsheng Hu, Zoran Salcic, Gillian Dobbie, Yi Chen, and Xuyun Zhang. 2021. EAR: An enhanced adversarial regularization approach against membership inference attacks. In *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [18] Bo Hui, Yuchen Yang, Haolin Yuan, Philippe Burlina, Neil Zhenqiang Gong, and Yinzhi Cao. 2021. Practical Blind Membership Inference Attack via Differential Comparisons. In *ISOC Network and Distributed System Security Symposium (NDSS)*.
- [19] Ruholla Jafari-Marandi, Joshua Denton, Adnan Idris, Brian K Smith, and Abbas Keramati. 2020. Iranian Churn. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5JW3Z>.
- [20] Longlong Jing, Ruichi Yu, Xu Chen, Zhengli Zhao, Shiwei Sheng, Colin Graber, Qi Chen, Qinru Li, Shangxuan Wu, Han Deng, et al. 2024. STT: Stateful Tracking with Transformers for Autonomous Driving. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 4442–4449.
- [21] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [22] Alfred Knopf. 1981. Mushroom. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5959T>.
- [23] Jiacheng Li, Ninghui Li, and Bruno Ribeiro. 2021. Membership inference attacks and defenses in classification models. In *Proceedings of the Eleventh ACM Conference on Data and Application Security and Privacy*. 5–16.
- [24] Wei-Yang Lin, Ya-Han Hu, and Chih-Fong Tsai. 2011. Machine learning in financial crisis prediction: a survey. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 42, 4 (2011), 421–436.
- [25] E Majani, Ruth Erlanson, and Yaser Abu-Mostafa. 1988. On the k-winners-take-all network. *Advances in neural information processing systems* 1 (1988).
- [26] Ryan McKenna, Brett Mullins, Daniel Sheldon, and Jerome Miklau. 2022. AIM: an adaptive and iterative mechanism for differentially private synthetic data. *Proceedings of the VLDB Endowment* 15, 11 (2022), 2599–2612.
- [27] Scott Mayer McKinney, Marcin Sieniek, Varun Godbole, Jonathan Godwin, Natasha Antropova, Hutan Ashrafian, Trevor Back, Mary Chesus, Greg S Corrado, Ara Darzi, et al. 2020. International evaluation of an AI system for breast cancer screening. *Nature* 577, 7788 (2020), 89–94.
- [28] Sérgio Moro, Paulo Cortez, and Paulo Rita. 2014. Bank Marketing. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5K306>.
- [29] Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*. 807–814.
- [30] Milad Nasr, Reza Shokri, and Amir Houmansadr. 2018. Machine learning with membership privacy using adversarial regularization. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*. 634–646.
- [31] David E Rumelhart and David Zipser. 1985. Feature discovery by competitive learning. *Cognitive science* 9, 1 (1985), 75–112.
- [32] C. Sakar and Yomi Kastro. 2018. Online Shoppers Purchasing Intention Dataset. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5F88Q>.
- [33] Ruslan Salakhutdinov and Geoff Hinton. 2007. Learning a nonlinear embedding by preserving class neighbourhood structure. In *Artificial intelligence and statistics*. PMLR, 412–419.
- [34] Virat Shejwalkar and Amir Houmansadr. 2021. Membership privacy for machine learning models through knowledge transfer. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 9549–9557.
- [35] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*. IEEE, 3–18.
- [36] Liwei Song and Prateek Mittal. 2021. Systematic evaluation of privacy risks of machine learning models. In *30th USENIX Security Symposium (USENIX Security 21)*. 2615–2632.
- [37] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2818–2826.
- [38] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research* 11, 12 (2010).
- [39] Chang Xiao, Peilin Zhong, and Changxi Zheng. 2020. Enhancing Adversarial Defense by k-Winners-Take-All. In *8th International Conference on Learning Representations*.
- [40] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2021. Understanding deep learning (still) requires rethinking generalization. *Commun. ACM* 64, 3 (2021), 107–115.

Model	Bank Marketing				Mushrooms			
	TPR @ 1% FPR (Mean (SD) in %)		AUC (Mean (SD) in %)		TPR @ 1% FPR (Mean (SD) in %)		AUC (Mean (SD) in %)	
	Original	Mixed	Original	Mixed	Original	Mixed	Original	Mixed
Baseline	0.79 (0.85)	0.92 (0.59)	50.02 (2.41)	50.07 (2.65)	0.00 (0.00)	0.01 (0.07)	48.63 (1.78)	49.17 (1.77)
w/ L1	1.13 (1.07)	0.87 (0.89)	49.35 (2.60)	49.66 (2.95)	1.15 (1.12)	0.85 (0.96)	48.52 (1.85)	48.94 (2.29)
w/ L2	0.77 (0.72)	1.03 (0.79)	50.51 (2.70)	50.04 (2.29)	0.14 (0.40)	0.43 (0.55)	49.66 (2.07)	48.99 (1.72)
kWTA NN	0.42 (0.47)	0.48 (0.64)	50.09 (2.27)	50.55 (2.62)	0.00 (0.00)	0.01 (0.07)	48.92 (1.83)	49.65 (1.97)
Disentangled	1.02 (1.10)	0.69 (0.79)	50.06 (2.59)	50.24 (2.75)	0.03 (0.14)	0.22 (0.68)	48.50 (1.73)	49.18 (1.97)
w/ KLD	0.01 (0.07)	0.03 (0.14)	50.16 (3.08)	49.23 (2.23)	0.07 (0.23)	0.18 (0.65)	49.74 (2.33)	49.34 (2.26)
Entangled	1.18 (1.22)	1.05 (0.68)	53.91 (2.10)	54.38 (2.69)	1.30 (1.01)	0.99 (0.56)	51.21 (1.89)	50.79 (2.18)
w/ KLD	0.01 (0.07)	0.00 (0.00)	53.97 (1.91)	54.39 (2.65)	0.12 (0.43)	0.16 (0.63)	51.70 (2.19)	50.84 (2.15)
KWt (H, p)	*(84.88, 0.00)	*(101.56, 0.00)	*(84.36, 0.00)	*(81.49, 0.00)	*(155.28, 0.00)	*(121.24, 0.00)	*(59.72, 0.00)	*(28.77, 0.00)

Model	Iranian Churn				Online Shoppers			
	TPR @ 1% FPR (Mean (SD) in %)		AUC (Mean (SD) in %)		TPR @ 1% FPR (Mean (SD) in %)		AUC (Mean (SD) in %)	
	Original	Mixed	Original	Mixed	Original	Mixed	Original	Mixed
Baseline	1.55 (0.99)	2.68 (1.73)	49.48 (2.69)	50.93 (1.94)	2.55 (2.22)	3.35 (2.09)	50.32 (3.04)	49.49 (2.61)
w/ L1	1.00 (0.80)	1.46 (1.16)	49.28 (2.73)	49.95 (2.02)	1.39 (1.12)	1.68 (1.03)	49.64 (3.17)	48.48 (2.70)
w/ L2	1.52 (1.31)	1.93 (1.52)	49.48 (2.46)	50.66 (1.75)	3.24 (2.00)	3.35 (1.82)	49.82 (2.81)	49.02 (2.31)
kWTA NN	1.05 (1.00)	1.18 (1.04)	49.46 (2.68)	50.89 (1.85)	1.43 (1.00)	1.69 (2.04)	50.73 (2.97)	49.54 (2.59)
Disentangled	1.51 (1.09)	2.37 (1.70)	49.58 (2.72)	50.94 (2.02)	2.63 (1.90)	3.52 (2.70)	50.55 (2.96)	49.52 (2.65)
w/ KLD	0.69 (1.11)	1.38 (1.85)	50.18 (2.79)	49.97 (1.95)	1.04 (2.54)	2.86 (4.05)	50.16 (2.82)	50.51 (2.67)
Entangled	0.81 (0.71)	0.92 (0.89)	49.74 (2.59)	51.44 (1.77)	0.85 (0.78)	1.11 (1.01)	51.33 (3.04)	50.51 (2.39)
w/ KLD	0.87 (1.00)	1.30 (1.72)	49.77 (2.66)	51.42 (1.87)	2.36 (4.04)	3.28 (3.66)	51.26 (2.90)	50.50 (2.51)
KWt (H, p)	*(28.85, 0.00)	*(37.21, 0.00)	-(2.63, 0.92)	*(15.59, 0.03)	*(62.44, 0.00)	*(37.08, 0.00)	-(8.97, 0.25)	*(16.38, 0.02)

Model	Statlog German Credit			
	TPR @ 1% FPR (Mean (SD) in %)		AUC (Mean (SD) in %)	
	Original	Mixed	Original	Mixed
Baseline	0.00 (0.00)	0.04 (0.15)	46.79 (2.87)	47.22 (2.23)
w/ L1	0.65 (1.00)	1.03 (1.08)	45.66 (4.59)	45.56 (4.73)
w/ L2	0.10 (0.30)	0.18 (0.28)	49.67 (3.30)	49.69 (2.63)
kWTA NN	0.00 (0.00)	0.00 (0.00)	50.28 (2.75)	50.41 (2.00)
Disentangled	0.00 (0.00)	0.00 (0.00)	46.77 (3.29)	46.73 (2.06)
w/ KLD	0.00 (0.00)	0.00 (0.00)	50.15 (2.98)	50.78 (2.04)
Entangled	0.88 (0.86)	1.05 (0.96)	58.36 (2.89)	57.51 (2.48)
w/ KLD	0.00 (0.00)	0.00 (0.00)	59.01 (2.89)	56.93 (3.66)
KWt (H, p)	*(108.05, 0.00)	*(142.26, 0.00)	*(156.86, 0.00)	*(159.24, 0.00)

Table 4: Comparison of LiRA MIA results on our baseline and experimental models (bold values represent the best results), trained under the same settings across the real datasets and mixture of real and synthetically generated datasets. We report the TPR at 1% FPR as our main metric for measuring the robustness of the target models against MIAs, and AUC is only presented for completeness. Our proposed membership inference defense approaches outperformed our baseline models in virtually all cases, with statistical significance. * denotes at $p < 0.05$ and - denotes not significant

Model	Kruskal-Wallis H test (H, p)				
	Real v. Mixed				
	Bank Marketing	Iranian Churn	Mushrooms	Online Shoppers	Statlog German Credit
Baseline	362.00, 0.19	271.50, 0.01	435.00, 0.33	336.50, 0.09	420.00, 0.16
w/ L1	437.00, 0.85	420.50, 0.66	435.00, 0.33	471.50, 0.75	450.00, 1.00
w/ L2	512.50, 0.35	365.50, 0.21	526.00, 0.26	365.00, 0.21	317.50, 0.04
kWTA NN	366.50, 0.21	375.00, 0.27	240.00, 0.00	417.50, 0.63	356.00, 0.05
Disentangled	513.50, 0.34	313.00, 0.04	390.50, 0.09	370.00, 0.24	450.00, 1.00
w/ KLD	434.50, 0.82	430.50, 0.77	519.00, 0.30	385.50, 0.34	393.00, 0.40
Entangled	449.50, 1.00	377.00, 0.25	433.00, 0.66	367.50, 0.18	450.00, 1.00
w/ KLD	465.00, 0.33	426.00, 0.72	449.50, 1.00	369.00, 0.21	450.00, 1.00
Overall	28020.50, 0.59	23963.00, 0.00	27052.50, 0.16	24458.00, 0.00	26220.00, 0.01

Table 5: Kruskal-Wallis H tests to compare whether there is a significant difference in attack TPR at 1% FPR when training the target models on the real datasets and on the mixture of real and synthetically generated datasets. Bold values represent those with statistical significance at $p < 0.05$.

Model	TPR @ 1% FPR (Mean (SD) in %)				
	Bank Marketing	Iranian Churn	Mushrooms	Online Shoppers	Statlog German Credit
Disentangled	0.54 (0.62)	2.12 (0.80)	0.16 (0.48)	1.13 (0.87)	0.00 (0.00)
Entangled	0.87 (0.88)	0.83 (0.83)	1.20 (0.98)	0.86 (0.80)	1.11 (0.87)
MWU (H, p)	-(637.50, 0.08)	*(133.00, 0.00)	*(883.00, 0.00)	-(407.00, 0.17)	*(900.00, 0.00)

Table 6: Comparison of LiRA MIA results on (1) disentangling, and (2) entangling real and synthetic datasets. * denotes statistical significance at $p < 0.05$, and - denotes otherwise.

FRIAD-IR: A Framework for Image-Augmented Dynamic Information Retrieval

Joaquin S. Heffron

University of the Philippines Diliman
Quezon City, Philippines
jsheffron@up.edu.ph

Paul Rossener R. Regonia, D.Sc.
University of the Philippines Diliman
Quezon City, Philippines

Bastian Nathaniel F. Calabia

University of the Philippines Diliman
Quezon City, Philippines
bfcabilia@up.edu.ph

Roselyn S. Gabud, Ph.D.

University of the Philippines Diliman
Quezon City, Philippines

Abstract

Rampant misinformation continues to plague the Philippine online landscape. Many instances of misinformation are multimodal, consisting of images, text, and/or audio-video. Natural Language Inference (NLI) pipelines are emerging as a solution to reliably automate the fact checking of claims. However, most existing studies rely on static information retrieval (IR) domains, are limited to text-based claims, and/or utilize enterprise large language models (LLMs). FRIAD-IR aims to fill this gap by combining Vision-Language Models (VLMs) with Custom Search Engines (CSEs) to perform image-augmented IR over a dynamic domain. FRIAD-IR consists of four stages: (1) claim processing, wherein a VLM enhances a text-based claim with context from its accompanying image and decomposes it into search queries; (2) web retrieval, wherein the search queries are fed into a whitelist-constrained CSE and documents are retrieved; (3) reranking, wherein the documents are embedded and ranked according to relevance to the claim; and (4) premise extraction, wherein a lightweight local LLM generates a textual premise by summarizing the documents. The pipeline was evaluated against a new image-augmented dataset based on the Philippine Online Misinformation Inference (POMI) dataset, which consists of claim-premise pairs manually extracted from fact-checking news articles. Premises generated by the pipeline were compared with premises in the dataset using semantic similarity. Preliminary results show that an image-augmented pipeline achieved higher average semantic similarity compared to a text-only pipeline.

CCS Concepts

• **Information systems** → **Information retrieval**; *Web searching and information discovery*; • **Computing methodologies** → *Information extraction*.

Keywords

information retrieval, retrieval augmented generation, Philippine misinformation, misinformation detection, fact checking dataset

1 Introduction

The production and dissemination of misinformation has continued in recent years within the Philippines. Although overall trust in news has been slowly recovering [6], most Filipinos continue to report difficulty discerning “fake news” [13]. Relatively new challenges introduced by the accessibility of generative artificial intelligence (GAI) are poised to exacerbate this problem, as GAI

enables the creation of realistic multimodal deepfakes with ease [17]. On the other hand, artificial intelligence (AI) has also been utilized in efforts to combat misinformation through tools such as automated fact checkers. This has led AI to be described as a “double-edged sword” regarding its overall effect on global information ecosystems [20].

Recent studies have revolved around leveraging Large Language Models (LLMs) of varying sizes for fact verification. In order to remain updated with the latest information, LLMs must be periodically retrained on a vast corpus that covers a wide range of topics and events. This approach effectively embeds the knowledge within the model itself. However, this periodic training may be computationally expensive and time consuming, reducing the practical viability of the model when considering the rapid pace of online news cycles and misinformation propagation.

A natural language inference (NLI) approach has emerged as a promising alternative, particularly for explainability [1] — that is, the ability of the model to provide evidence or reasoning to support its conclusion. NLI models classify the relationship between two pieces of text: a hypothesis (i.e. a claim) and a premise. A classification falls under one of three categories: entailment (hypothesis follows the premise), contradiction (hypothesis opposes the premise) or neutral (hypothesis neither follows nor opposes the premise). In the context of fact verification, NLI models rest on the assumption that entailment between a factual premise and a given claim implies that the claim is also factual, whereas a contradiction between the same would imply the claim to be false.

For the goal of automated fact verification, the obvious requirement of an NLI model is ready access to information that can serve as premises. This information must be reliable and relevant to the given claim. To fulfill this requirement, an information retrieval (IR) pipeline may be employed. Such a pipeline may extract information on-demand from a corpus of reliable sources and feed it into the NLI model as premises. Several studies have demonstrated the benefit of this approach, particularly compared to standalone LLMs [14, 29]. Others even integrate Vision-Language Models (VLMs) within their pipelines to address the problem of multimodal misinformation [22]. However, these studies rely on static datasets for their evidence corpus. This retains the problem of impracticality for real-time fact checking, as these corpora must be updated as rapidly as current events evolve in order to be reliable.

In light of this, we propose a theoretical framework for multimodal information retrieval over a dynamic web-based dataset in a Philippine fact-verification context.

2 A Review of Literature

2.1 Retrieval as a generative task

Classical approaches to IR include static indexing and similarity matching, which are useful for local corpora that may be periodically indexed or embedded. However, these methods are less feasible for leveraging externally-hosted dynamic corpora, such as the World Wide Web through search engines. An emerging paradigm with the potential to address this is generative IR, which reframes IR from a matching task to a generation task [15]. This approach leverages LLMs to directly generate document identifiers, such as titles or related keywords. These identifiers may then be used as queries or pointers to retrieve full documents from a corpus. Chen et al. demonstrated that this method can be utilized for evidence retrieval tasks, outperforming traditional methods on fact verification benchmarks [5].

2.2 The Retrieve-and-Rerank Paradigm

The conventional approach to IR within the context of Retrieval-Augmented Generation (RAG) follows the "Retrieve-and-Rerank" paradigm [10, 11]. This architecture typically consists of three discrete stages: (1) indexing, where a document corpus is chunked and encoded into dense vectors; (2) retrieval, where a similarity search identifies a broad set of candidate documents; and (3) reranking, where a computationally intensive cross-encoder refines the selection to filter out noise.

2.3 Multimodal Information Retrieval

Vision-Language Models prove to be highly capable of multimodal IR, as highlighted in state-of-the-art model papers [2]. Multimodal IR specifically for fact-verification is also addressed in recent studies: a proposed framework with an evidence-ranking based approach [22] and another with a probing classifier-based approach [4]. These developments demonstrated improvements in performance when applied to their respective benchmarks, validating the efficacy of novel solutions for specific fact-checking purposes.

2.4 IR Benchmarks for Fact-verification

We acknowledge the presence of prominent IR benchmarks in the context of fact-verification such as FEVER [16, 25] and AVeriTeC [7, 21], with more recent benchmarks accounting for the trend towards multimodal IR such as MOCHEG [22, 28]. However, these datasets and their corresponding solutions operate on closed, static corpora. This highlights a research gap in IR frameworks capable of open-web retrieval for the purpose of time-sensitive fact-checking.

3 Methodology

This section details the construction of the benchmark dataset and the architecture of the proposed retrieval framework.

3.1 Data collection

Our dataset was heavily based on the Philippine Online Misinformation Inference (POMI) dataset, which was developed to finetune NLI models for fact verification in a Philippine context [1]. Unlike POMI, our dataset is not intended to be used for fine tuning models, but instead to benchmark pipelines.

3.1.1 Dataset sources. Like POMI, our dataset was sourced from fact-checking articles created by various Philippine news and fact-checking organizations. The specific organizations are FactRakers [9], Rappler [18], Vera Files [26], The Philippine Star [24], and The Baguio Chronicle [23]. Articles were written in either Tagalog or English, with publication dates ranging from 2018 to 2025. Unlike POMI, our dataset specifically targeted articles that rely heavily on images to describe or debunk claims. POMI disqualified such articles, whereas our dataset required them in order to evaluate multimodal pipelines.

3.1.2 Article annotation. The following elements were extracted from each article:

Field	Description
False claim	The claim constituting misinformation
False headline	The headline (or "clickbait") of the online post containing misinformation
False image	The supplementary image of the false claim
Factual claim	The claim supported by the factual premise
Factual headline	The headline of the fact-checking article
Factual premise	The evidence or reasoning provided by the fact-checking article to support its claim
Factual image (optional)	The supplementary image of the factual claim
Category	One of <i>politics, health, religion, celebrity, or others</i>
URL	The URL of the fact-checking article

Unlike POMI, images related to the claims were also extracted from each article. These images would provide additional context that could not be captured by simply extracting text-based elements. Each article yielded two entries: the false entry, which encodes the false elements meant to be contradicted by the premise; and the factual entry, which encodes the factual elements meant to be entailed by the premise. Both entries share the same factual premise, category, and URL.

Most fact-checking articles typically feature their subject misinformative posts, either through online archiving tools or screenshots. Other articles also paraphrase the essential claim being conveyed by said posts. From these, the false elements may be extracted. Articles then typically state a clear verdict on the veracity of these (false) claims and provide evidence or reasoning to support their conclusion. From these, the factual elements may be extracted.

In general, elements were extracted verbatim from their articles (i.e. copy-pasted); but like in POMI, minor tweaks were made for extracted excerpts that refer to excluded sections of the article. Common cases of this include replacing pronouns (e.g. "She said...") with their proper antecedent nouns (e.g. "Sara said..."), replacing acronyms with their full expansion, or expanding dates to include their year.

FACT CHECK: Renee Co and Zaldy Co are not related, viral photo is AI-manipulated

SEP 19, 2025 11:30 AM PHT

JOHN SITCHON

UPGRADE TO LISTEN POWERED BY SPEECHIFY

Already have Rappler+? [Sign in](#) to listen to groundbreaking journalism.

FACT CHECK

Results from a deepfake detector show that the image is manipulated. It also bears the logo of Gemini, Google's generative AI chatbot.

Claim: A post shows a photo of Kabataan Representative Renee Co and embattled Ako Bicol Representative Zaldy Co, implying that they are close relatives.

Rating: FALSE

FAKE HEADLINE

The Facebook page wrote: "Dili ni sila mag pinsan ha or relative (They are not cousins or relatives)." The caption appears to discredit the Kabataan lawmaker's assertion that she is not related to the Ako Bicol solon.

Hundreds of comments used the term "relathieves" (a combination of relative and thieves) to describe the photo of the two lawmakers.

SCREENSHOTS OF FALSE POST

FAKE (AND FACT) IMAGE

FACT CLAIM

The facts The image shown in the post is manipulated. Results from deepfake detector website [deepfake-detect.com](#) show that the image is a deepfake. It also carried the logo of Gemini, Google's generative artificial intelligence chatbot.

PREMISE

Figure 1: A sample annotated political fact-checking article

Headline	Hypothesis/Claims	Premise/Facts	Category	Label	Link	Image
TITAN'S CREW DIED AFTER SUFFERING AN IMPLOSION AND BREAKING	A post on the Facebook page Guerrero Noticias features a series of pictures of underwater wreckage. As of writing (Jun 23, 2023), no images of the Titan's wreckage and debris have been released to the public.	The photos used are actually from 2004 and came from the University of Rhode Island's archives. The photos used are actually from 2004 and came from the University of Rhode Island's archives.	Others	FAKE	https://www.rappler.com/newsbreak/fact-check/facebook-post-shows-no-real-images-titan-submersible-wreckage/	
GREATNESS RECOGNIZE GREATNESS	Supposed Kevin Durant tweet commending Kai Sotto is fake	On July 14, 2023, a Twitter user @dantelloversace asked Durant to name the most underrated player that NBA star Kevin Durant has played. On July 14, 2023, a Twitter user @dantelloversace asked Durant to name the most underrated player that NBA star Kevin Durant has played.	Celebrity	FAKE	https://www.rappler.com/newsbreak/fact-check/fact-check-supposed-kevin-durant-tweet-commending-kai-sotto-is-fake/	
"WALA NAMANG NAGAWA SI Bongbong Marcos" --- says a woke who is a destitute of knowledge but still acting like an armchair expert	Pinarangalan ng Gawad Alab Haraya ng NCCA si Marcos Jr.	Ang Gawad Alab Haraya ay iginagawad ng National Commission for Culture and the Arts (o NCCA) para sa mga nagaganap ng Gawad Alab Haraya ng NCCA.	Politics	FAKE	https://www.rappler.com/newsbreak/fact-check/ferdinand-bongbong-marcos-jr-awarded-gawad-alab-haraya-ncca/	
Proving yet again that he is the GOAT.	Altered photo shows Lionel Messi holding Israeli flag	On June 1, 2022, Lionel Messi was holding a placard with text that reads: "I love you, Israel". In the original photo, Messi was holding a placard with text that reads: "I love you, Israel".	Celebrity	FAKE	https://www.rappler.com/newsbreak/fact-check/ferdinand-bongbong-marcos-jr-awarded-gawad-alab-haraya-ncca/	
Altered photo shows Lionel Messi holding Israeli flag		On June 1, 2022, Lionel Messi was holding a placard with text that reads: "I love you, Israel". In the original photo, Messi was holding a placard with text that reads: "I love you, Israel".	Celebrity	FAKE	https://www.rappler.com/newsbreak/fact-check/ferdinand-bongbong-marcos-jr-awarded-gawad-alab-haraya-ncca/	

Figure 2: A set of sample dataset entries




Figure 3: Diagram of the Proposed Information Retrieval Framework

3.2 Proposed Framework for Multimodal Information Retrieval

In this study, we present a framework for image-augmented dynamic information retrieval (FRIAD-IR) for processing multimodal claims such as those in our dataset. Although FRIAD-IR is based on conventional IR architectures used in RAG, its novelty lies in leveraging a live retrieval domain of Philippine fact-checking articles. FRIAD-IR is designed to pass a text premise output to an NLI solution for fact-checking. It consists of four primary components: (1) claim processing, (2) web search and document retrieval, (3) embedding and reranking, and (4) premise extraction. Its architecture is designed as a model-agnostic pipeline with configurable modules, ensuring consistent input/output interfaces regardless of the underlying model selection. Figure 3 illustrates the general flow of the components and their expected inputs and outputs. The following sections detail the role and implementation of each component.

3.2.1 Claim Processing.

Input: Claim to be fact-checked.

Output: Text search query for web search.

The claim processing component focuses on multimodal comprehension to bridge the semantic gap between visual evidence and textual search queries. By isolating VLM usage to this initial stage, the

framework optimizes downstream retrieval quality while minimizing computational overhead, as heavy VLM inference is performed only once per claim.

Preprocessing. Standard preprocessing ensures consistency prior to query generation. Text inputs undergo normalization, tokenization, and lemmatization to remove noise. Image inputs are padded and resized to uniform dimensions to meet the input specifications of the VLM to be used.

Captioning. To capture context that might be lost in unimodal processing, FRIAD-IR utilizes a VLM to generate a context-rich caption by synthesizing features from both the image and the supporting text. This process integrates visual cues (e.g., text overlays, identifiable figures) into a coherent textual format. The resulting VLM-generated caption serves as a context-augmented search query, replacing baseline text-only claims.

3.2.2 Web Search and Retrieval.

Input: Text search query for web search.

Output: Retrieved documents with text content.

The web search and retrieval component of FRIAD-IR exclusively interfaces with the online search domain.

Web Search. The current framework configuration uses the Google Programmable Search Engine, also known as Custom Search Engine (CSE), for real-time web search for available fact-checking articles. The CSE accepts text queries as input and returns the best matching (based on the Google API) search results with URLs based on the query [12]. The CSE used by the current configuration also includes a user-defined domain whitelist for the results, including the base domains of FactRakers, Rappler, Vera Files, The Philippine Star, and The Baguio Chronicle. This ensures that only reputable and verified sources are used as the basis for fact-checking.

Retrieval. The current framework configuration uses HTTP requests to fetch the HTML contents of the returned URLs and applies basic parsing to extract the main body text. Each retrieved body text is paired with its corresponding URL. The results are stored in a top- k list of retrieved documents then passed on to the next component.

3.2.3 Embedding and Reranking.

Input: Retrieved documents with text content.

Output: Final selection of documents with text content.

Embedding. The text contents of the retrieved documents are encoded into vector representations using an embedding model. By using vector representations, semantics are captured more effectively without the need for keyword-based reranking. Additionally, the claim is also encoded with the same embedding model to allow direct comparisons of the document and claim embeddings.

Reranking. Reranking is performed on the document embeddings to optimize and narrow the selection of retrieved documents. Only the top- k most relevant results relative to the claim are retained. This step includes the use of a similarity function (e.g. cosine similarity) between the vectors to quickly get a relevance matching and scoring. Additionally, this step may integrate cross-encoder

models for improved performance or alternative metrics as bases for reranking.

3.2.4 Premise Extraction.

Input: Final selection of documents with text content.

Output: Text premise for fact-checking claims.

LLM Generation. The generation component adopts the conventional RAG procedure after document retrieval, where an LLM is prompted alongside the selected documents to generate the desired output. In this case, the current configuration prompts an LLM to extract a concise premise based on the contents of the relevant fact-checking article(s). This approach simulates the premise extraction procedure as discussed in Section 3.1.2.

4 Results and Discussion

4.1 Dataset

The dataset is composed of 200 entries (100 fake and 100 factual) sourced from 100 articles. Every fake entry had its own image for context, whereas only 37 factual entries had their own distinct image. The remaining 63 factual entries shared a contextual image with their article’s fake entry. This was due to certain articles simply not providing a supporting image as evidence in debunking the fake claim, instead focusing on disproving the fake claim’s supplementary image. Like POMI, most of the articles covered political claims (Figure 5). Unlike POMI, FactRakers was the leading source of articles instead of Vera Files (Figure 6). This was partly due to FactRakers’ notable commitment to archiving its subject misinformative posts, which made it much more likely that the original image could be extracted.

4.2 Experimental Setup

To evaluate the efficacy of the proposed framework, we established two distinct pipeline configurations and a rigorous scoring system.

4.2.1 Models. The models used for both configurations as a preliminary output is as follows:

- **Qwen3-VL-2B** as a VLM capable of image captioning with supporting text [2].
- **Qwen3-Embedding-0.6B** as the semantic embedding model for documents and claims [30].
- **Qwen3-1.7B** as a LLM used for premise extraction [27].

4.2.2 Baseline (text-only) vs Multimodal Pipelines. As a baseline for comparison, we utilized a text-only configuration of the framework. This approach performs search and retrieval using only the text claim as query, effectively ignoring the image component of the claim. This baseline serves to quantify the performance gap when visual context, such as text overlays or scene details, is disregarded. To ensure a fair comparison, all parameters (e.g. models, prompts, etc.) were kept identical to the multimodal configuration.

4.2.3 Evaluation Metrics. We evaluated both the baseline and multimodal IR pipelines using two metrics: URL matching and semantic textual similarity.

URL matching. As discussed in Section 3.1.2, the URLs of each fact-checking article in the dataset were recorded. We count the

number of entries that, when their claim is passed into the pipeline, leads to the retrieval of the entry’s article URL in the dataset.

Semantic Textual Similarity. For certain claims, the pipeline fails to retrieve the exact fact-checking article but retrieves similar articles from other news sources. Simple URL matching would not account for this case, even if similar articles effectively provide the same information as the exact article. To address this, we compare the pipeline-generated premises against the ground-truth dataset premises using semantic textual similarity (STS). For this, we employ the following models:

- **Qwen3-Embedding-8B:** Embedding model based on *Qwen3* which maps sentences to a 4096 dimensional dense vector space [30]. It is our primary model for evaluation and the best performing model for STS in the Massive Text Embedding Benchmark (MTEB) [8].
- **text-embedding-3-large:** Foundational embedding model based on *GPT-3.5* that embeds up to 3072 dimensions [3].
- **all-mpnet-base-v2:** Lightweight *sentence-transformers* model based on *mpnet-base* that maps sentences to a 768 dimensional dense vector space [19].

4.3 Quantitative Analysis

4.3.1 Performance Comparison. Table 1 and Table 2 detail the STS performance of the baseline and multimodal pipelines respectively. Table 3 compares the number of exact URL matches produced by the multimodal pipeline and the baseline pipeline.

Table 1: STS Scores of Baseline (Text-only) Pipeline

Similarity Model	Mean	Std. Dev	Median
<i>All Entries</i>			
qwen3-embedding-8b	0.5189	0.2949	0.6243
text-embedding-3-large	0.4912	0.2957	0.5683
all-mpnet-base-v2	0.4853	0.2946	0.5676
<i>URL Matched Entries</i>			
qwen3-embedding-8b	0.6820	0.1869	0.7473
text-embedding-3-large	0.6447	0.2030	0.7109
all-mpnet-base-v2	0.6407	0.2056	0.6978
<i>Non-URL Matched Entries</i>			
qwen3-embedding-8b	0.2981	0.2700	0.3582
text-embedding-3-large	0.2835	0.2745	0.2652
all-mpnet-base-v2	0.2750	0.2654	0.2665

4.3.2 Discussion. The results indicate a significant improvement when utilizing the multimodal pipeline. The *qwen3-embedding-8b* semantic score rose from an overall mean of **0.5189** in the baseline to **0.5911** in the multimodal configuration, representing a relative improvement of 13.91%. This suggests that a VLM-captioning approach is able to extract additional context from the images, aiding the retrieval of more semantically relevant documents.

A divergence is observed in cases where the pipelines failed to retrieve the exact ground-truth article (non-URL matched). In these scenarios, the baseline pipeline’s mean similarity score drops to **0.2981**. In contrast, the multimodal pipeline maintains a significantly higher similarity score of **0.4758**. This suggests that




Figure 4: Diagram of a sample claim passed through the pipeline




Figure 5: Articles per Category




Figure 6: Articles per Source

the multimodal pipeline successfully leverages VLM captioning to describe the visual context, allowing it to recognize relevant fact-checking articles from alternative sources (Figure 4).

4.4 Qualitative Analysis

We analyzed the performance of the VLM in generating contextually rich captions. We observed the following variants of multimodal

claims where it had succeeded/failed to append context from an image to a text claim:

4.4.1 Text-rich Image Claim. When the image claim contains the complete context in text form (e.g. screenshot of a social media post), the VLM was highly effective in extraction and was minimally dependent on the context of the corresponding text claim.

Table 2: STS Scores of Multimodal Pipeline

Similarity Model	Mean	Std. Dev	Median
All Entries			
qwen3-embedding-8b	0.5911	0.2242	0.6534
text-embedding-3-large	0.5622	0.2339	0.6341
all-mpnet-base-v2	0.5549	0.2434	0.6188
URL Matched Entries			
qwen3-embedding-8b	0.6712	0.1767	0.7095
text-embedding-3-large	0.6436	0.1931	0.6992
all-mpnet-base-v2	0.6375	0.2084	0.7007
Non-URL Matched Entries			
qwen3-embedding-8b	0.4758	0.2358	0.4890
text-embedding-3-large	0.4449	0.2387	0.4753
all-mpnet-base-v2	0.4361	0.2421	0.4781

Table 3: Overlap of URL Matching between Baseline and Multimodal Pipelines

Baseline (Text Only)	VLM-enhanced (Image)	
	URL Matched	No URL Match
URL Matched	98	17
No URL Match	20	65

4.4.2 Event-based Image Claim. When the image claim displays an event that is independent of any individual identity (e.g. photo of a natural disaster), the VLM output was only context-complete if the missing context from the text claim was supplied by the image.

4.4.3 Identity-based Image Claim. When the image claim contains context that is highly dependent on an individual’s identity (e.g. photo of a certain celebrity at an event), the VLM, being untrained, was not able to identify people and thus was not able to supplement any meaningful context to a text claim.

4.5 Error Analysis

To better understand the limitations of the proposed framework, we conducted a review of the entries that failed. We identified three primary failure modes that inhibit the pipeline’s performance:

4.5.1 Query-Retrieval Mismatch. The first mode of failure involves the inability of the Search Engine (CSE) to return relevant URLs for a generated query. This stems from a semantic disconnect between the VLM-generated queries and the indexing logic of the external search provider. Since the CSE is a "black-box" component, the pipeline is occasionally bottlenecked by the inability of the search API to map context-heavy captions to headlines of news articles.

4.5.2 Content Access Restrictions. The second failure mode occurs during the document scraping phase, where valid URLs are returned but their HTML content cannot be extracted. This is due to CAPTCHA or similar mechanisms employed by the domains. This is circumvented by a delayed second pass to retrieve the missing URLs. Nonetheless, this remains a persistent challenge for live web information retrieval.

4.5.3 Premise Extraction Failure. In cases where the document is successfully retrieved, the system may fail to generate a coherent premise. This indicates a performance bottleneck in the downstream components: failure to accurately rerank retrieved articles and/or LLM hallucinations. This suggests that while multimodal captioning improves search, using an LLM to summarize documents introduces inconsistencies in output quality compared to direct text extraction.

5 Limitations, Future Work, and Recommendations

CSE. This study operates under the constraint of a live web-based retrieval environment to simulate a realistic deployment scenario. Consequently, a significant limitation of the proposed framework is its dependency on third-party search engines for coarse retrieval. To address this, we plan to attach a modular indexing component to our pipeline that periodically updates a local corpus of articles from trusted domains. This retains the dynamic aspect of retrieval while providing us access to a controlled database index. This approach will effectively eliminate the major performance bottleneck and black-box problem imposed by the CSE.

Dataset. Due to the scarcity of image-based fact-checking articles and limited resources for data collection, we collected only 200 entries for the dataset used in our study. As a compromise, we plan to leverage the rest of the POMI dataset (containing 5,024 entries) to evaluate a text-only version of our pipeline. This forgoes evaluation for the image-component, but allows for more robust evaluation against the vast POMI dataset.

Evaluation metric. As our proposed framework operates with NLI as a downstream task in mind, our evaluation metric being STS and URL matching only partially represents the performance of an end-to-end NLI pipeline. With the scope of this study being limited to IR, we recommend future studies that utilize our proposed framework to instead evaluate primarily on NLI performance metrics if possible. Moreover, search engine evaluation metrics such as Recall@K may provide a better quantitative analysis in the context of strict document retrieval without premise extraction.

Premise extraction. To address premise extraction failures, we recommend shifting to a fine-tuned classification approach for premise identification within documents. A direct extraction method from documents without the use of LLMs mitigates the risk of hallucinating premises. This proposed method may maximize the robustness and explainability of the system based on ground truths.

6 Conclusion

We proposed FRIAD-IR: an information retrieval framework for the downstream task of natural language inference for fact verification. This framework performs retrieval from a web-based dynamic domain in lieu of a static corpus. For evaluation purposes, we manually annotated an image-augmented dataset based on the Philippine Online Misinformation Inference dataset. We developed an initial implementation of the framework using small but state-of-the-art vision-language models (VLMs). Based on preliminary results, we conclude that a VLM-captioning approach exhibits strong potential to improve retrieval outcomes when processing multimodal claims,

while noting its weaknesses in certain cases. To address the limitations of the study, we plan to shift from real-time web retrieval to a periodically indexed local corpus. Additionally, the use of the entirety of POMI will allow us to leverage a larger dataset for evaluating document retrieval performance. Overall, FRIAD-IR offers a scalable and adaptable alternative to periodic language model retraining or manual premise retrieval. This ultimately enables automated fact-checking systems to remain robust by leveraging current information.

7 Acknowledgments

We acknowledge the University of the Philippines (UP) Intelligent Systems Center Research Grant (Project MALAIMO ISRGA2024-16) for the invaluable support.

References

- [1] Ma, Pauline Abcede, Led Rhoniel Salazar, Paul Rossener Regonia, and Roselyn Gabud. 2025. A Natural Language Inference-Based Approach for Detecting Misinformation on Philippine Online News. In *Proceedings of the 25th Philippine Computing Science Congress*. Computing Society of the Philippines.
- [2] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-VL: A Versatile Vision-Language Model for Understanding, Localization, Text Reading, and Beyond. arXiv:2308.12966 [cs.CV] <https://arxiv.org/abs/2308.12966>
- [3] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. arXiv:2005.14165 [cs.CL] <https://arxiv.org/abs/2005.14165>
- [4] Recep Firat Cekin, Pinar Karagoz, and Çağrı Çöltekin. 2025. Multimodal Fact-Checking with Vision Language Models: A Probing Classifier based Solution with Embedding Strategies. In *Proceedings of the 31st International Conference on Computational Linguistics*, Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert (Eds.). Association for Computational Linguistics, Abu Dhabi, UAE, 4622–4633. <https://aclanthology.org/2025.coling-main.310/>
- [5] Jiangui Chen, Ruqing Zhang, Jiafeng Guo, Yixing Fan, and Xueqi Cheng. 2022. GERE: Generative Evidence Retrieval for Fact Verification. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*. ACM, 2184–2189. doi:10.1145/3477495.3531827
- [6] Yvonne T. Chua. 2025. Digital News Report 2025: Philippines. (2025). <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2025/philippines>
- [7] Arbaaz Dharmavaram and Saqib Hakak. 2025. SANCTUARY: An Efficient Evidence-based Automated Fact Checking System. In *Proceedings of the Eighth Fact Extraction and VERification Workshop (FEVER)*, Mubashara Akhtar, Rami Aly, Christos Christodoulopoulos, Oana Cocarascu, Zhijiang Guo, Arpit Mittal, Michael Schlichtkrull, James Thorne, and Andreas Vlachos (Eds.). Association for Computational Linguistics, Vienna, Austria, 247–257. doi:10.18653/v1/2025.feve-1.19
- [8] Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Márton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Sibli, Dominik Krzemiński, Genta Indra Winata, Saba Sturua, Sateja Utpala, Mathieu Ciancone, Marion Schaeffer, Gabriel Sequeira, Diganta Misra, Shreeya Dhakal, Jonathan Rystrom, Roman Solomatin, Ömer Çağatan, Akash Kundu, Martin Bernstorff, Shitao Xiao, Akshita Sukhlecha, Bhavish Pahwa, Rafał Poświata, Kranthi Kiran GV, Shawon Ashraf, Daniel Auras, Björn Plüster, Jan Philipp Harries, Loïc Magne, Isabelle Mohr, Mariya Hendriksen, Dawei Zhu, Hippolyte Gisserot-Boukhlef, Tom Aarsen, Jan Kostkan, Konrad Wojtasik, Taemin Lee, Marek Šuppa, Crystina Zhang, Roberta Rocca, Mohammed Hamdy, Andrianos Michail, John Yang, Manuel Faysse, Aleksei Vatulín, Nandan Thakur, Manan Dey, Dipam Vasani, Pranjal Chitale, Simone Tedeschi, Nguyen Tai, Artem Snegirev, Michael Günther, Mengzhou Xia, Weijia Shi, Xing Han Lu, Jordan Clive, Gayatri Krishnakumar, Anna Maksimova, Silvan Wehrli, Maria Tikhonova, Henil Panchal, Aleksandr Abramov, Malte Ostendorff, Zheng Liu, Simon Clematide, Lester James Miranda, Alena Fenogenova, Guangyu Song, Ruqiyah Bin Safi, Wen-Ding Li, Alessia Borghini, Federico Cassano, Hongjin Su, Jimmy Lin, Howard Yen, Lasse Hansen, Sara Hooker, Chenghao Xiao, Vaibhav Adlakha, Orion Weller, Siva Reddy, and Niklas Muennighoff. 2025. MMTEB: Massive Multilingual Text Embedding Benchmark. arXiv preprint arXiv:2502.13595 (2025). doi:10.48550/arXiv.2502.13595
- [9] FactRakers. 2025. FactRakers. <https://www.factrakers.org/>
- [10] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. Retrieval-Augmented Generation for Large Language Models: A Survey. <https://arxiv.org/abs/2312.10997> arXiv preprint.
- [11] Michael Glass, Gaetano Rossiello, Md Faisal Mahbub Chowdhury, Ankita Rajaram Naik, Pengshan Cai, and Alfio Gliozzo. 2022. Re2G: Retrieve, Rerank, Generate. arXiv:2207.06300 [cs.CL] <https://arxiv.org/abs/2207.06300>
- [12] Google Developers. 2024. Google Programmable Search Engine. <https://developers.google.com/custom-search>
- [13] Inquirer Research. 2025. SWS: Most Filipinos See “Fake News” as a Serious Problem. *Philippine Daily Inquirer* (13 March 2025). <https://newsinfo.inquirer.net/2043443/sws-most-filipinos-see-fake-news-as-serious-problem>
- [14] Hai Li, Jingyi Huang, Mengmeng Ji, Yuyi Yang, and Ruopeng An. 2025. Use of Retrieval-Augmented Large Language Model for COVID-19 Fact-Checking: Development and Usability Study. *Journal of Medical Internet Research* (2025). <https://pmc.ncbi.nlm.nih.gov/articles/PMC12079058/>
- [15] Xiaoxi Li, Jiajie Jin, Yujia Zhou, Yuyao Zhang, Peitian Zhang, Yutao Zhu, and Zhicheng Dou. 2025. From Matching to Generation: A Survey on Generative Information Retrieval. arXiv:2404.14851 [cs.LG] <https://arxiv.org/abs/2404.14851>
- [16] Hao Liao, Jiahao Peng, Zhanyi Huang, Wei Zhang, Guanghua Li, Kai Shu, and Xing Xie. 2023. MUSER: A Multi-Step Evidence Retrieval Enhancement Framework for Fake News Detection. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Long Beach, CA, USA) (KDD '23)*. Association for Computing Machinery, New York, NY, USA, 4461–4472. doi:10.1145/3580305.3599873
- [17] Seyeon Park and Xiaoli Nan. 2025. Generative AI and Misinformation: A Scoping Review of the Role of Generative AI in the Generation, Detection, Mitigation, and Impact of Misinformation. *AI & Society* (2025). doi:10.1007/s00146-025-02620-3
- [18] Rappler. 2025. Rappler. <https://www.rappler.com/>
- [19] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. <https://arxiv.org/abs/1908.10084>
- [20] Tetiana Schipper. 2025. Disinformation by Design: Leveraging Solutions to Combat Misinformation in the Philippines’ 2025 Election. *Data & Policy* 7 (2025), e39. doi:10.1017/dap.2025.18
- [21] Michael Schlichtkrull, Zhijiang Guo, and Andreas Vlachos. 2023. AVeriTeC: A Dataset for Real-world Claim Verification with Evidence from the Web. arXiv:2305.13117 [cs.CL] <https://arxiv.org/abs/2305.13117>
- [22] Sahar Tahmasebi, Eric Müller-Budack, and Ralph Ewerth. 2024. Multimodal Misinformation Detection Using Large Vision-Language Models. <https://arxiv.org/abs/2407.14321> arXiv preprint.
- [23] The Baguio Chronicle. 2025. The Baguio Chronicle. <https://thebaguiochronicle.com/>
- [24] The Philippine Star. 2025. The Philippine Star. <https://www.philstar.com>
- [25] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for Fact Extraction and VERification. arXiv:1803.05355 [cs.CL] <https://arxiv.org/abs/1803.05355>
- [26] VERAFiles. 2025. VERAFiles. <https://verafiles.org/>
- [27] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuxiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. Qwen3 Technical Report. arXiv:2505.09388 [cs.CL] <https://arxiv.org/abs/2505.09388>
- [28] Barry Menglong Yao, Aditya Shah, Lichao Sun, Jin-Hee Cho, and Lifu Huang. 2023. End-to-End Multimodal Fact-Checking and Explanation Generation: A Challenging Dataset and Models. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*. ACM, 2733–2743. doi:10.1145/3539618.3591879
- [29] Zhenrui Yue, Huimin Zeng, Lanyu Shang, Yifan Liu, Yang Zhang, and Dong Wang. 2024. Retrieval-Augmented Fact Verification by Synthesizing Contrastive Arguments. arXiv (2024). <https://arxiv.org/abs/2406.09815>
- [30] Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. 2025. Qwen3 Embedding: Advancing Text Embedding and Reranking Through Foundation Models. arXiv:2506.05176 [cs.CL] <https://arxiv.org/abs/2506.05176>

Parallel Permutation Testing for One-Way ANOVA Using CUDA

Kyle Carlo C. Lasala
De La Salle University
Manila, Philippines
kyle_lasala@dlsu.edu.ph

Maria Monica Manlises
De La Salle University
Manila, Philippines
maria_monica_manlises@dlsu.edu.ph

Daphne Janelyn L. Go
De La Salle University
Manila, Philippines
daphne_janelyn_go@dlsu.edu.ph

Roger Luis L. Uy
De La Salle University
Manila, Philippines
roger.uy@dlsu.edu.ph

ABSTRACT

Permutation tests provide a flexible, nonparametric alternative to classical hypothesis testing by constructing null distributions directly from data. However, their widespread use is often limited by the high computational cost associated with evaluating a large number of permutations, particularly for exact tests whose complexity grows factorially with sample size. In this work, we investigate GPU-accelerated permutation testing for one-way analysis of variance (ANOVA) using NVIDIA CUDA. We present both exact and Monte Carlo permutation test implementations and compare a serial C baseline against multiple CUDA-based designs that exploit data-level parallelism inherent in permutation testing. To enable efficient parallel execution of exact permutation tests, we adopt a rank-indexing permutation algorithm that eliminates sequential dependencies present in lexicographic permutation generation. We further examine the impact of shared memory optimization on performance. Experimental results demonstrate that CUDA implementations consistently outperform the serial baseline as the number of permutations increases. Exact permutation testing using rank-indexing achieves speedups exceeding $30\times$ and up to approximately $70\times$ for large permutation counts, while Monte Carlo implementations achieve speedups of up to $30\times$. These results show that GPU parallelism, combined with appropriate permutation algorithms, can substantially reduce the computational burden of permutation-based ANOVA, making exact and large-scale permutation testing more practical for real-world applications.

KEYWORDS

Parallel Computing, CUDA, Permutation Test, Parallel Algorithms

1 INTRODUCTION

Permutation tests were introduced in the 1930s by Pitman [16] and other statisticians as an alternative to classical parametric tests in situations where distributional assumptions or large-sample approximations may not hold [1]. These methods assess statistical significance by repeatedly recomputing a test statistic under different rearrangements of the data. Instead of assuming a specific probability distribution, permutation tests construct a reference distribution directly from the data by systematically or randomly permuting labels and comparing the observed statistic to this distribution [6].

Permutation tests are widely used across scientific disciplines due to their flexibility and minimal assumptions. They are particularly useful in settings with a small number of independent observations [8], when distribution-free analysis is preferred [2, 15], and when precise control of false positive rates is required [17]. Moreover, many classical test statistics can be embedded within a permutation framework, allowing familiar parametric test statistics to be applied without relying on their usual distributional assumptions.

However, while permutation tests are conceptually simple and broadly applicable, their computational cost grows rapidly with data size and the number of permutations. Exact permutation tests require evaluating all possible label rearrangements, leading to factorial growth in the number of permutations and quickly becoming infeasible for even moderate sample sizes. Monte Carlo permutation tests reduce this cost by sampling a subset of permutations, but may still require millions of independent evaluations to obtain accurate p -value estimates, particularly when estimating very small p -values [14, 21].

The computational structure of both exact and Monte Carlo permutation testing makes them well-suited for parallel execution. Since each permutation can be evaluated independently, permutation tests naturally expose data-level parallelism, where the same computation is repeated across a large number of permutations. Modern graphics processing units (GPUs) provide thousands of lightweight threads and high memory bandwidth, making them an attractive platform for accelerating permutation-based inference using CUDA's Single Instruction, Multiple Threads (SIMT) execution model.

Several studies have explored parallel and GPU-accelerated permutation testing, particularly in high-dimensional domains such as neuroimaging. Early work by Eklund et al. [3, 4] demonstrated the feasibility of GPU-based permutation testing for fMRI analysis, focusing primarily on medical imaging pipelines. Their work considered permutation tests using canonical correlation analysis and the general linear model as core test statistics. Winkler et al. [20] further reduced the computational burden by exploiting exchangeability and symmetry properties of the permutation distribution to minimize redundant computations and improve parallel efficiency. While these studies establish the effectiveness of GPU acceleration for permutation testing, they are tightly coupled to specific neuroimaging workflows. In contrast, we propose a domain-agnostic permutation kernel based on analysis of variance (ANOVA) that can

⁰Source code available at: <https://github.com/KyleCarlo/CUDA-1wayANOVA-Perm/>

be readily integrated into a wide range of applications, including but not limited to fMRI analysis.

One commonly used test statistic in permutation testing is the difference in group means. This formulation naturally gives rise to permutation versions of classical parametric tests, including unpaired and paired t -tests as well as various forms of ANOVA. There are several existing implementations of permutation test in this form with varying use cases [7–9]. Furthermore, implementations exist that use GPU computation in general or emphasize CUDA optimization features [5, 19]. In this work, we focus on accelerating permutation tests based on the one-way ANOVA statistic, emphasizing generality and optimizing unified memory of CUDA by using shared memory.

Permutation-based ANOVA provides a nonparametric alternative to classical ANOVA by computing the null distribution of the F -statistic through label permutations rather than relying on the theoretical F distribution [8]. We implement both exact and Monte Carlo permutation tests and compare a serial C baseline against multiple CUDA-accelerated implementations. Our approach assigns one permutation to one GPU thread, enabling concurrent computation of group means, within-group and between-group sums of squares, and the resulting F -statistic. For exact permutation testing, we evaluate both a lexicographic permutation approach, which introduces inherent sequential dependencies, and a rank-indexing permutation method that enables fully parallel permutation generation directly on the GPU. For Monte Carlo testing, permutations are generated independently using a Fisher–Yates shuffle, making the approach naturally parallelizable.

Experimental results demonstrate that the proposed CUDA implementations consistently outperform the serial C baseline as the number of permutations increases. For exact permutation testing, GPU implementations using rank-indexing achieve speedups exceeding $30\times$ and up to approximately $70\times$ for large permutation counts, while Monte Carlo CUDA implementations achieve speedups of up to $30\times$. These gains are achieved while preserving correctness, with Monte Carlo results closely matching exact permutation outcomes and all implementations producing identical results for exact inference.

2 METHODOLOGY

This section presents the details of the permutation test algorithm, sequential and CUDA implementations, and device used in the experiments.

2.1 Permutation Test Algorithm

A permutation test is a non-parametric statistical method used to assess the significance of an observed test statistic by comparing it against a null distribution generated from permutations of the data under the assumption that all observations arise from the same population [12]. Permutation tests are particularly useful when the data do not satisfy normality assumptions, since they allow the use of conventional parametric test statistics while deriving significance from an empirical permutation distribution rather than from a theoretical parametric distribution.

In an *exact permutation test*, all possible distinct permutations of the dataset that preserve the group sizes are exhaustively enumerated to form the null distribution. Although exact permutation tests provide an exact p -value under the permutation framework, they become computationally expensive as the dataset size increases. To address this limitation, a *Monte Carlo permutation test* approximates the null distribution by randomly sampling a fixed number of valid permutations instead of enumerating all possible ones. As the number of sampled permutations increases, the estimated p -value approaches that of the exact permutation test.

In this work, the design of the permutation test implementation is influenced not only by its statistical formulation but also by the requirements of the CUDA framework and GPU execution model. In particular, the implementation must expose sufficient parallelism, reduce thread divergence, and organize memory access efficiently in order to achieve good performance on the GPU. These considerations affect how permutations are represented, how group assignments are generated, and how the test statistic is evaluated. Thus, while the underlying statistical procedure remains unchanged, its computational realization is shaped by the constraints and opportunities of CUDA.

Formally, let $V = (v_0, v_1, \dots, v_{N-1})$ denote the observed data values and let $G = (g_0, g_1, \dots, g_{N-1})$ denote the corresponding group assignments, where each observation v_i belongs to group g_i . The dataset is represented by the ordered pairs (v_i, g_i) for $i = 0, 1, \dots, N-1$. Let K be the number of distinct groups in G , with group labels restricted to $g_i \in \{0, 1, \dots, K-1\}$. If the number of observations in group k is n_k , then the total number of observations is

$$N = \sum_{k=0}^{K-1} n_k.$$

Each distinct permutation that preserves the group sizes n_k is indexed by an integer $p \in \{0, 1, \dots, M-1\}$, where p denotes the p -th permutation in lexicographic order. The total number of such distinct permutations is

$$M = \frac{N!}{\prod_{k=0}^{K-1} n_k!}.$$

The permuted dataset corresponding to index p is denoted by $X_p = (V, G_p)$, where $G_p = (g_{p,0}, g_{p,1}, \dots, g_{p,N-1})$ and $g_{p,i}$ is the group assignment of observation v_i under the p -th lexicographic permutation. Let T denote the chosen test statistic. For any dataset X , the value of the statistic is written as $T(X)$. Let the original dataset be denoted by $X_0 = (V, G_0)$, corresponding to the first lexicographic ordering. The p -value of the exact permutation test is then defined as

$$p = \frac{1}{M} \sum_{p=0}^{M-1} \mathbf{1}(T(X_p) \geq T(X_0)),$$

where $\mathbf{1}(\cdot)$ is the indicator function, which takes the value 1 when the condition is true and 0 otherwise.

In the Monte Carlo permutation test, instead of enumerating all M distinct permutations, only a subset of valid permutations is randomly sampled, and the p -value is estimated from the proportion of sampled permutations whose test statistic is greater than or

equal to the observed test statistic. This approximation substantially reduces computational cost and makes the method practical for large datasets, particularly in parallel computing environments such as CUDA-enabled GPUs.

2.1.1 One-Way ANOVA. One-way ANOVA is widely used to assess the differences in group means, as it generalizes the two-samples t -test to datasets involving more than two groups [18]. In our implementation, we selected F statistic of one-way ANOVA as the test statistic T for the permutation test. Formally,

$$T(X) = \frac{\frac{1}{K-1} \sum_{k=0}^{K-1} n_k (\bar{v}_k - \bar{v})^2}{\frac{1}{N-K} \sum_{i=0}^{N-1} (v_i - \bar{v}_{g_i})^2}$$

where $\bar{v}$ is the overall mean and $\bar{v}_k$ is the mean of group k .

Note that although we used the one-way ANOVA F -statistic as an example, the permutation test framework is general and any test statistic T can be substituted depending on the hypothesis being tested.

2.2 CUDA Architecture

Compute Unified Device Architecture (CUDA) is a parallel computing framework developed by NVIDIA that enables general-purpose computation on Graphics Processing Units (GPUs). In CUDA, computation is expressed through a *kernel*, which is executed on the GPU by many lightweight threads in parallel. Threads are organized hierarchically into *thread blocks*, and blocks are organized into a *grid*, allowing a large workload to be divided into smaller tasks that can be processed simultaneously.

Each thread executes the same kernel code but operates on different portions of the data. This makes CUDA suitable for problems with data-level parallelism, where the same operation is applied to many independent data elements. In permutation testing, this parallelism can be exploited by assigning different permutations or statistical computations to different threads.

CUDA provides several levels of memory with different performance characteristics. *Global memory* is accessible to all threads but has relatively high latency. *Shared memory* is a smaller on-chip memory region shared among threads within a block, offering lower latency for data repeatedly accessed during computation. *Per-thread memory* is private to each thread for intermediate values that are not shared.

Threads within a block can also be synchronized using functions such as `__syncthreads()`, which acts as a barrier to ensure all threads reach the same point before proceeding. Synchronization ensures correctness when accessing shared data but introduces overhead, so it is ideal to minimize its use whenever possible.

These architectural characteristics influenced the design of the CUDA implementation in this study, including the mapping of permutations to threads, the use of per-thread data structures, and the organization of memory accesses to improve computational efficiency.

2.3 Sequential Implementation

The sequential implementation serves as the baseline for evaluating empirical computational performance against the CUDA implementation. Two sequential versions are implemented—one for the exact permutation test and another for the Monte Carlo permutation test. The primary difference between the two implementations lies in the permutation generation procedure. The exact permutation test employs a lexicographic permutation algorithm, whereas the Monte Carlo version uses the Fisher–Yates shuffle.

The lexicographic permutation algorithm requires the $(i-1)$ -th permutation to generate the i -th permutation, introducing a sequential dependency that limits its applicability to sequential execution. The actual implementation is present in Appendix A.1.

In contrast, the Fisher–Yates shuffle is used for the Monte Carlo permutation test, while the rank-indexing permutation algorithm is employed for exact permutation testing. The implementations of the Fisher–Yates and rank-indexing algorithms are provided in Appendices A.2 and A.3, respectively. This will be further discussed in section 2.4.

The pseudocode for the entire sequential implementation is as follows:

Algorithm 1: Sequential Permutation Test

Require: Observed data values V , group labels G , number of permutations M

Ensure: Permutation test p-value p

```

1: Compute observed statistic  $T_0 \leftarrow T((V, G))$ 
2: Initialize counter  $c \leftarrow 0$ 
3: for  $p \leftarrow 1$  to  $M$  do
4:   Generate a permutation  $G_p$ 
5:   Compute statistic  $T_p \leftarrow T((V, G_p))$ 
6:   if  $T_p \geq T_0$  then
7:      $c \leftarrow c + 1$ 
8:   end if
9: end for
10: return  $p \leftarrow \frac{c}{M}$ 
```

2.4 CUDA Implementation

The CUDA implementation minimizes thread synchronization in order to reduce synchronization overhead and improve computational performance. Several CUDA implementation variants are considered, which differ in the permutation algorithm used and in the parallelization techniques employed. Each version is evaluated and compared against the sequential implementation.

Table 1: CUDA Implementation Versions

Permutation Test	Permutation Algorithm	Shared Memory	Version Code
Exact	Lexicographic	Not applied	E1
	Rank-indexing	Not applied	E2
	Rank-indexing	Applied	E3
Monte Carlo	Fisher–Yates	Not applied	MC1
	Fisher–Yates	Applied	MC2

Table 1 summarizes the permutation algorithms used and indicates whether shared memory is applied for each implementation version. All CUDA implementations require each GPU thread to maintain a local memory of size K for intermediate computations. Since K is determined at runtime, allocating such per-thread arrays dynamically on the device is either unsupported or incurs significant performance and memory overhead. To avoid this, we employ static per-thread memory allocation with a fixed maximum size α . This design imposes the constraint that datasets with $K > \alpha$ are not supported, trading generality for predictable performance and efficient memory usage. In this work, we define $\alpha = 100$, but it can be adjusted depending on computational requirements and hardware constraints.

2.4.1 Version E1. In version E1, all permutations are generated on the CPU prior to GPU execution. This design choice is necessary because the lexicographic permutation algorithm exhibits a strict sequential dependency: the $(p + 1)$ -th permutation can only be generated after the p -th permutation. Consequently, permutation generation cannot be parallelized and must be performed on the host. Each GPU thread is therefore assigned only the task of evaluating the test statistic.

To ensure that all permutation instances are processed, a *grid-stride loop* is employed. In this approach, each thread processes multiple permutation indices by advancing in steps equal to the total number of threads in the grid. This allows all permutations to be evaluated even when the number of permutations exceeds the number of available GPU threads. Finally, the host evaluates the number of permuted test statistics greater than or equal to the observed statistic. The pseudocode for version E1 is as follows:

Algorithm 2: Version E1

Require: Observed data values V , group labels G , number of permutations M

Ensure: Permutation test p-value p

```

1: Compute observed statistic  $T_0 \leftarrow T((V, G))$ 
2: Initialize counter  $c \leftarrow 0$ 
   ▶ Host-side sequential permutation generation
3: for  $p \leftarrow 1$  to  $M$  do
4:   Generate lexicographic permutation  $G_p$ 
5:   Store  $G_p$ 
6: end for
   ▶ Device-side evaluation using grid-stride loop
7: Launch CUDA kernel
8: for  $p \leftarrow \text{blockIdx} \times \text{blockDim} + \text{threadIdx}$  to  $M - 1$  step  $\text{gridDim} \times \text{blockDim}$  do
9:   Compute statistic  $T_p \leftarrow T((V, G_p))$ 
10:  Store  $T_p$ 
11: end for
   ▶ Host-side p-value evaluation
12: for  $i \leftarrow 1$  to  $M - 1$  do
13:   if  $T_i \geq T_0$  then
14:      $c \leftarrow c + 1$ 
15:   end if
16: end for
17: return  $p \leftarrow \frac{c}{M}$ 
```

2.4.2 Version E2. To reduce the bottleneck in E1, the permutation algorithm is replaced with a rank-indexing approach. Unlike lexicographic permutation, rank-indexing directly generates the p -th permutation from its corresponding rank, without requiring the generation of preceding permutations. This removes the sequential dependency present in E1. The implementation of the rank-indexing permutation is provided in Appendix A.3. With this approach, each GPU thread independently generates its assigned permutation and subsequently evaluates the corresponding test statistic. The pseudocode for version E2 is as follows:

Algorithm 3: Version E2

Require: Observed data values V , group labels G , number of permutations M

Ensure: Permutation test p-value p

```

1: Compute observed statistic  $T_0 \leftarrow T((V, G))$ 
2: Initialize counter  $c \leftarrow 0$ 
   ▶ Device-side permutation and evaluation
3: Launch CUDA kernel
4: for  $p \leftarrow \text{blockIdx} \times \text{blockDim} + \text{threadIdx}$  to  $M - 1$  step  $\text{gridDim} \times \text{blockDim}$  do
5:   Generate permutation  $G_p$ 
6:   Compute statistic  $T_p \leftarrow T((V, G_p))$ 
7:   Store  $T_p$ 
8: end for
   ▶ Host-side p-value evaluation
9: for  $i \leftarrow 1$  to  $M - 1$  do
10:  if  $T_i \geq T_0$  then
11:     $c \leftarrow c + 1$ 
12:  end if
13: end for
14: return  $p \leftarrow \frac{c}{M}$ 
```

2.4.3 Version E3. Version E3 extends Version E2 by incorporating shared memory to improve performance. Shared memory is private to each thread block and has a limited capacity determined by the GPU hardware. It accesses the necessary data faster compared to global memory. Data that are repeatedly accessed by GPU threads, specifically the observed values V and group labels G , are first copied from global memory into shared memory at the beginning of the kernel execution. Threads are then required to synchronize before proceeding with the execution of the permutation test. Threads within the same block reuse these shared values when generating permutations and computing the ANOVA test statistic, reducing expensive global memory accesses. While this approach can improve performance, the use of shared memory and the required synchronization introduces overhead, which may limit speedup for small datasets or kernels with few threads.

2.4.4 Version MC1 and MC2. Both versions MC1 and MC2 approximate the exact permutation test by randomly sampling a fixed number M of permutations. In both implementations, permutations are generated using the Fisher–Yates shuffle, as provided in Appendix A.2. Version MC2 extends MC1 by incorporating the use of shared memory to improve performance, following the same optimization strategy applied in Version E3.

2.5 Device

All experiments were performed on a compute node equipped with a 16-core Intel(R) Xeon(R) Silver 4210 processor operating at 2.20 GHz, 32 GB of RAM, and an NVIDIA Tesla V100 accelerator with 32 GB of GPU memory.

3 RESULTS

This section presents the theoretical analysis and empirical results of different implementations.

Table 2: Performance speedup (Minimum–Maximum) of CUDA permutation implementations relative to C

No of Permutations	E1		E2		E3	
	Min	Max	Min	Max	Min	Max
10^2	1.12×	3.99×	0.81×	4.28×	0.83×	4.35×
10^3	1.80×	3.46×	7.22×	16.18×	8.21×	16.55×
10^4	1.74×	2.11×	30.91×	52.89×	34.60×	70.43×
10^5	1.56×	2.58×	32.95×	57.94×	47.77×	68.42×
10^6	2.06×	2.40×	51.48×	64.36×	57.36×	67.13×
10^7	1.58×	2.35×	53.26×	72.16×	58.22×	73.73×




Figure 1: Performance Comparison of Exact permutation implementations in C and CUDA

3.1 Theoretical Analysis

The computational performance is primarily determined by the degree of parallelism, which in turn is constrained by the permutation algorithm. Both the lexicographic permutation method and the Fisher–Yates shuffle generate a single permutation in linear time $O(N)$. However, the lexicographic method exhibits strict sequential dependencies between successive permutations, which prevents parallel generation. In contrast, the Fisher–Yates shuffle allows independent generation of permutations and is therefore well suited for parallel Monte Carlo permutation tests.

The rank-indexing permutation algorithm directly constructs the p -th permutation and requires $O(NK)$ time per permutation, where K denotes the number of groups. Although this introduces additional computational overhead compared to $O(N)$ permutation methods, it enables independent permutation generation and thus allows effective parallelization.

In the sequential implementation, each of the M permutations requires $O(N)$ time for permutation generation and $O(N)$ time for evaluating the test statistic, resulting in an overall time complexity of $O(MN)$.

For the CUDA implementation using rank-indexing permutation, assuming ideal load balancing across r threads, the parallel execution time scales as

$$O\left(\frac{MNK}{r}\right),$$

reflecting the trade-off between increased per-permutation computational cost and improved parallel throughput.

3.2 Empirical Results

Multiple configurations of N and K were evaluated to assess the performance of the proposed C and CUDA implementations, specifically $N \in \{6, 8, 10, 12, 14, 15\}$ and $K \in \{3, 4, 5\}$. As the total number of permutations depends on the chosen values of N and K , the remainder of this section reports and analyzes the performance of these implementations with respect to the number of permutations M .

Across all configurations, the CUDA implementations for Exact Permutation Test consistently outperform the serial C version, especially as the total permutations M increase (see Figure 1). For small datasets yielding less than 10^3 permutations, the observed speedups are limited, ranging from approximately $1\times$ to $4\times$. In contrast, the CUDA Exact Permutation using the rank-indexing permutation algorithm scaled effectively, achieving speedups of over $30\times$ and up to $74\times$ for large datasets yielding more than 10^4 permutations (see Table 2). Overall, these findings demonstrate that GPU parallelism is more pronounced and effectively utilized for large permutation datasets.

Table 3: Performance Speedup of CUDA Monte Carlo Permutation Implementations Relative to C

No. of Permutations	MC1	MC2
10^4	12.18×	12.88×
10^6	24.89×	30.13×
10^8	17.52×	19.00×




Figure 2: Performance Comparison of Monte Carlo Permutation Implementations in C and CUDA

Among the CUDA implementations, the two implementations yielding the highest speedups were: (a) the CUDA Exact Permutation using the rank-indexing permutation algorithm (E2), and

(b) the same ranking-based CUDA approach incorporating shared memory and caching (E3) as observed in Figure 1 and Table 2. While the shared-memory variant (E3) generally outperforms the baseline (E2), the improvement is marginal, with speedups increasing by only approximately $0.8\times$ to $1.2\times$. Performance varied between the two implementations, with each occasionally outperforming the other. This limited performance improvement was expected, as shared memory was applied only to arrays of size N . Given our current test limit of $N \leq 15$, the cached data is too small for shared memory to produce a substantial difference.

For the Monte Carlo implementations, the numbers of permutations evaluated were 10^4 , 10^6 , and 10^8 . The same pattern in performance observed in the exact permutation tests was seen in the Monte Carlo implementations, with all CUDA implementations consistently outperforming the C baseline (see Figure 2). Across all configurations, the CUDA Monte Carlo implementation incorporating shared memory and caching (MC2) achieved the best runtime, with speedups ranging from approximately $12\times$ to $30\times$ compared to the C baseline (see Table 3). By avoiding the factorial growth of permutations, the Monte Carlo approach is far more practical for larger values of N and K , especially when computing under hardware constraints.

3.3 Performance Analysis and Discussion

The results indicate that runtime is dominated primarily by the repeated generation of permutations and the corresponding evaluation of the test statistic across all M permutations. Since both operations are performed for every permutation, their cumulative cost increases rapidly as M grows. For the exact permutation test, this effect is especially pronounced because the total number of permutations grows with N and K . In the CUDA implementations, the dominant runtime cost is associated with these repeated permutation computations rather than with one-time setup costs, particularly for large workloads.

The superior performance of E2 and E3 over E1 is mainly explained by the permutation generation strategy. E1 is limited by a less effective mapping of the permutation-generation process to the GPU, whereas E2 and E3 use the rank-indexing permutation algorithm, which allows each permutation to be generated independently. This independence is critical for CUDA execution because it enables direct assignment of different permutations to different threads without requiring sequential dependence between successive permutations. Although rank-indexing incurs a higher per-permutation computational cost of $O(NK)$, this overhead is offset by the large degree of parallelism that it exposes.

The performance difference between E3 and E2 is relatively small because the shared-memory optimization in E3 benefits only a limited portion of the computation. In the present experiments, shared memory was applied to arrays of size N , with $N \leq 15$. Because these arrays are small, the cost of accessing them from global memory is already limited, and the opportunity for shared memory to reduce latency is correspondingly modest. As a result, E3 shows only marginal improvement over E2, indicating that memory access is not the primary bottleneck under the tested configurations.

For small values of M , the observed speedup of the CUDA implementations is relatively low because fixed overheads form a

substantial fraction of total runtime. These overheads include kernel launch costs, device memory allocation, and host-to-device and device-to-host data transfers. When the number of permutations is small, the actual computation performed on the GPU is insufficient to fully amortize these costs. Consequently, the performance advantage of GPU execution becomes much more evident only as M increases and the workload becomes large enough to utilize the available parallel resources effectively.

The Monte Carlo permutation test is more practical for large problem sizes because it avoids the factorial growth in the number of permutations required by the exact test. Instead of exhaustively enumerating all valid permutations, it samples only a fixed number of permutations, thereby making the computational cost controllable. This makes the Monte Carlo approach more suitable for larger values of N and K , where exact permutation testing may become infeasible even with GPU acceleration. While exact testing provides a complete evaluation of the permutation distribution, Monte Carlo testing offers a computationally efficient approximation that is often sufficient in practice.

Overall, the results highlight a key trade-off between exactness and scalability. Exact permutation testing provides an exact p-value but suffers from combinatorial growth in runtime as the dataset size increases. In contrast, the Monte Carlo approach sacrifices exactness in exchange for substantially improved scalability and practical feasibility. Similarly, the CUDA-based exact implementations demonstrate that parallelization can significantly accelerate permutation testing, but only when the workload is large enough to overcome execution overheads. These findings show that the choice of method should depend not only on statistical requirements but also on the problem size and the available computational resources.

4 CONCLUSION

This study investigated the feasibility and performance benefits of optimizing permutation testing for one-way ANOVA using GPU CUDA-based data-level parallelism. The primary objective was to address the computational bottleneck of exact permutation tests, whose factorial growth in the number of permutations makes them impractical for even moderate sample sizes. As such, we implemented a permutation test workflow that can either exhaustively enumerate all valid label permutations or randomly sample permutations using a Monte Carlo approach, concurrently computing the ANOVA F -statistic across permutations through GPU SIMT execution. Serial C and parallel CUDA implementations were benchmarked, including exact and Monte Carlo variants based on lexicographic, Fisher–Yates, and rank-indexing permutation algorithms. Results show that CUDA implementations consistently outperform the serial baseline, with CUDA exact permutation tests using rank-indexing achieving speedups exceeding $30\times$ and up to approximately $70\times$ for large permutation counts, while CUDA-based Monte Carlo implementations with speedups of up to $30\times$. These results highlight that the choice of permutation algorithm is crucial for maximizing the benefits of GPU parallelism, and that exact permutation testing can be made practical for larger workloads when sequential dependencies are eliminated.

Future work may extend the scalability of the proposed approach to substantially larger values of N and K , where factorial growth

in the number of permutations and the computational cost of rank-indexing become more pronounced, allowing deeper evaluation of the effectiveness of the proposed CUDA-based approaches in terms of scalability and memory usage. Beyond data-level parallelism, future implementations may explore task-level parallelism, for example by overlapping permutation generation, test statistic computation, and reduction operations across GPU streams or heterogeneous CPU–GPU execution. The framework may also be extended to investigate alternative multiset permutation algorithms, which may offer improved computational efficiency or memory behavior compared to rank-indexing for specific group-size configurations. In addition, the framework can be generalized to support additional statistical tests, including mixed factorial ANOVA and regression-based permutation tests, each introducing new opportunities for hardware-specific optimization. Finally, exploring other computational models and architectures, such as multi-GPU systems, distributed-memory platforms, or accelerator technologies beyond CUDA, may further enable exact permutation testing at scales that remain impractical on single-device systems.

REFERENCES

- [1] K. J. Berry. 2018. Permutation Statistical Methods: Calculation Efficiencies. *Biostatistics and Biometrics Open Access Journal* 4 (2018). doi:10.19080/bboaj.2018.04.555640
- [2] Ronglin Che, John R. Jack, Alison A. Motsinger-Reif, and Chad C. Brown. 2014. An Adaptive Permutation Approach for Genome-Wide Association Study: Evaluation and Recommendations for Use. *BioData Mining* 7, 1 (2014), 9. doi:10.1186/1756-0381-7-9
- [3] Anders Eklund, Mats Andersson, and Hans Knutsson. 2012. fMRI Analysis on the GPU—Possibilities and Challenges. *Computer Methods and Programs in Biomedicine* 105, 2 (2012), 145–161. doi:10.1016/j.cmpb.2011.07.007
- [4] Anders Eklund, Pierre Dufort, Daniel Forsberg, and Stephen M. LaConte. 2013. Medical Image Processing on the GPU—Past, Present and Future. *Medical Image Analysis* 17, 8 (2013), 1073–1094. doi:10.1016/j.media.2013.05.008
- [5] Markus Ekvall, M. Höhle, and L. Käll. 2020. Parallelized calculation of permutation tests. *Bioinformatics* 36 (2020), 5392 – 5397. doi:10.1093/bioinformatics/btaa1007
- [6] M. D. Ernst. 2004. Permutation Methods: A Basis for Exact Inference. *Statist. Sci.* 19 (2004), 676–685. doi:10.1214/088342304000000396
- [7] Jaromil Frossard and O. Renaud. 2021. Permutation Tests for Regression, ANOVA, and Comparison of Signals: The permuco Package. *J. Stat. Softw.* 99 (2021). doi:10.18637/jss.v099.i15
- [8] Charles A. Holt and Sean Sullivan. 2023. *Permutation Tests for Experimental Data*. Technical Report 4. University of Iowa. 775–812 pages. doi:10.1007/s10683-023-09799-6
- [9] A. B. Kashlak, S. Myroshnychenko, and S. Spektor. 2020. Analytic Permutation Testing for Functional Data ANOVA. *Journal of Computational and Graphical Statistics* 32 (2020), 294 – 303. doi:10.1080/10618600.2022.2069780
- [10] D.E. Knuth and Addison-Wesley. 2014. *The Art of Computer Programming: Combinatorial algorithms, Part 1. Vol. 4A*. Addison-Wesley. <https://books.google.com.ph/books?id=Mqy40QEACAAJ>
- [11] Donald E. Knuth. 1969. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*. Addison-Wesley, Reading, MA.
- [12] K. Nichols and A. Holmes. 2007. CHAPTER 21 - Non-parametric procedures. In *Statistical Parametric Mapping*, KARL FRISTON, JOHN ASHBURNER, STEFAN KIEBEL, THOMAS NICHOLS, and WILLIAM PENNY (Eds.). Academic Press, London, 253–272. doi:10.1016/B978-012372560-8/50021-8
- [13] A. Nijenhuis, H.S. Wilf, and W. Rheinboldt. 2014. *Combinatorial Algorithms: For Computers and Calculators*. Academic Press. <https://books.google.com.ph/books?id=cL7SBQAAQBAJ>
- [14] Jeffrey D. Opdyke. 2003. Fast Permutation Tests That Maximize Power under Conventional Monte Carlo Sampling for Pairwise and Multiple Comparisons. *Journal of Modern Applied Statistical Methods* 2, 2 (2003). doi:10.22237/jmasm/1051747500
- [15] Belinda Phipson and Gordon K Smyth. 2010. Permutation P-values Should Never Be Zero: Calculating Exact P-values When Permutations Are Randomly Drawn. *Statistical Applications in Genetics and Molecular Biology* 9, 1 (Oct. 2010). doi:10.2202/1544-6115.1585
- [16] E. J. G. Pitman. 1937. Significance Tests Which May Be Applied to Samples from any Populations. *Journal of the Royal Statistical Society* 4, Supplement (1937), 119–130.
- [17] David M. Ritzwoller, Joseph P. Romano, and Azeem M. Shaikh. 2025. Randomization Inference: Theory and Applications. arXiv:2406.09521 [econ.EM] <https://arxiv.org/abs/2406.09521>
- [18] G.E. Sarty. 2022. *Introduction to Applied Statistics for Psychology Students*. University of Saskatchewan. <https://books.google.com.ph/books?id=s4MuzwEACAAJ>
- [19] I. Sfiligoi. 2025. Comparing CPU and GPU compute of PERMANOVA on MI300A. *ArXiv* (2025). doi:10.1145/3708035.3736040
- [20] Anderson M. Winkler, Gerard R. Ridgway, Gwenaëlle Douaud, Thomas E. Nichols, and Stephen M. Smith. 2016. Faster Permutation Inference in Brain Imaging. *NeuroImage* 141 (2016), 502–516. doi:10.1016/j.neuroimage.2016.05.068
- [21] Jie Yang, Emanuele Trucco, and Anne Buu. 2018. A Hybrid Method of the Sequential Monte Carlo and the Edgeworth Expansion for Computation of Very Small p-values in Permutation Tests. *Statistical Methods in Medical Research* 28 (2018), 2937–2951. doi:10.1177/0962280218791918

APPENDIX A MULTISSET PERMUTATION

We considered three multiset permutation algorithms in our permutation test implementation. The first is the lexicographic permutation method, which generates permutations sequentially and exhibits dependencies on the previously generated permutation. The second is the Fisher–Yates shuffle, which is used for the Monte Carlo permutation test by randomly permuting group assignments. The third is a rank-indexing permutation approach, which enables parallel generation of permutations and is therefore used to parallelize the exact permutation test. The corresponding implementations are provided below.

A.1 Lexicographic Permutation

The lexicographic permutation algorithm, also known as the next-permutation algorithm, generates permutations sequentially in lexicographic order. Each permutation is computed from the immediately preceding one, introducing a strict dependency that prevents independent or parallel generation of permutations. The Exchange function swaps two elements in the array specified by the first argument, using the indices provided by the second and third arguments.

```

1 int permute(size_t a[], size_t n) {
2     int l, j;
3     for (j = --n; j > 0 && a[j-1] >= a[j]; --j) {
4         ;
5     }
6     if (j == 0) return 0;
7     for (l = n; a[j-1] >= a[l]; --l) { ; }
8     Exchange(a, j-1, l);
9     while (j < n) { Exchange(a, j++, n--); }
10    return 1;
11 }

```

A.2 Fisher-Yates Shuffle

The Fisher-Yates shuffle uses a random number generator to produce a uniformly random permutation by successively swapping elements [11]. The `lcg_random` function generates a pseudorandom number from the input seed and is used to verify the correctness of the output.

```

1 void permute(size_t *array, size_t N, unsigned int
2     seed, size_t *result) {
3     for (size_t i = 0; i < N; i++) {
4         result[i] = array[i];
5     }
6     for (size_t i = N - 1; i > 0; i--) {

```

```

6         size_t j = lcg_random(seed) % (i + 1);
7         Exchange(result, i, j)
8     }
9 }

```

A.3 Rank-Indexing

Here, keys denotes the set of unique group labels $\{0, \dots, K - 1\}$, n_i the corresponding group sample sizes $\{n_k\}$, K the number of groups, p the lexicographic permutation index, and perm the array storing the resulting permutation.

Unlike permutation generators based on element swapping (e.g., lexicographic or Fisher–Yates methods), the rank-to-permutation algorithm constructs the desired permutation directly from its rank. At each position, it determines which group label must be placed by counting how many permutations would result from fixing each possible label at that position. These counts are given by multinomial coefficients. The rank is then updated, and the process continues until the full permutation is constructed [10, 13].

```

1  __device__ void rank_to_permutation(size_t *keys,
2      size_t *n_i, size_t k, size_t N, unsigned long
3      long rank, size_t *perm) {
4      size_t n_i_copy[MAX_GROUPS];
5
6      for (int i = 0; i < k; i++) {
7          n_i_copy[i] = n_i[i];
8      }
9
10     size_t total = N;
11
12     for (int pos = 0; pos < N; pos++) {
13         for (int i = 0; i < k; i++) {
14             if (n_i_copy[i] == 0)
15                 continue;
16
17             n_i_copy[i]--;
18             unsigned long long num = multinomial(
19                 total - 1, n_i_copy, k);
20
21             if (rank < num) {
22                 perm[pos] = keys[i];
23                 total--;
24                 break;
25             } else {
26                 rank -= num;
27                 n_i_copy[i]++;
28             }
29         }
30     }
31 }

```

A Scalable and Computationally Efficient Implementation of a Macroscopic Traffic Model for Urban Networks

John Virgilio Afable Jr.

Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
jbafable@alum.up.edu.ph

Gabriel Datiles

Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
gvdatiles@alum.up.edu.ph

Jaymar Soriano

Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
jbsoriano@up.edu.ph

John Justine Villar

Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
jsvillar1@up.edu.ph

Adrian Roy Valdez

Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
alvaldez@up.edu.ph

ABSTRACT

Urban traffic congestion remains a persistent and costly challenge for rapidly urbanizing regions in the Philippines, particularly in Metro Manila. While microscopic traffic simulators provide high-fidelity representations of individual vehicle behavior, they are often computationally prohibitive for city-scale analysis. Macroscopic traffic models, which describe traffic flow as a continuous fluid, offer a computationally tractable alternative, yet their practical implementation for complex urban networks is still limited in local simulation platforms. In this paper, we present a scalable and computationally efficient implementation of a macroscopic traffic model based on the cell transmission model for networked urban intersections. The model is formulated using a discretized partial differential equation governing traffic density evolution, incorporating fundamental diagrams, driver foresight behavior, and intersection dynamics through merge models, turn-rate matrices, and binary traffic signals. Beyond direct replication of a prior work, we introduce key computational enhancements to improve numerical robustness and practicality. These include a revised flux-limiting scheme to prevent cell overflow, an underflow correction to maintain physically meaningful densities, and a topology-aware evaluation strategy based on topological sorting of roads to resolve downstream dependencies and enable single-pass updates per time step. Simulations reproduce realistic traffic wave propagation, congestion buildup near intersections, and the effects of traffic signal coordination on flow dynamics. The model scales linearly with network size and remains computationally lightweight. This work provides a practical computational foundation for integrating macroscopic modeling into Philippine traffic simulation systems, supporting faster, city-scale scenario analysis for data-driven traffic management and smart city planning.

KEYWORDS

Vehicular traffic; Macroscopic traffic modeling; Cell transmission model; Traffic management; Traffic network

1 Introduction

One of the most pressing infrastructure and societal challenges in the Philippines is urban traffic congestion, particularly in highly urbanized regions such as Metro Manila and its surrounding areas. Over the past two decades, rapid motorization and population growth have significantly outpaced the expansion of road capacity. Within Mega Manila, the vehicular population increased by approximately 1.5 times from 2006 to 2015, while the number of registered motorcycles nationwide grew by more than 300% from 2000 to 2009, rising from 1 million to 3.2 million units [1]. As a result, Metro Manila has consistently ranked among the most traffic-congested metropolitan areas in Asia according to assessments by international agencies such as the Asian Development Bank, the Boston Consulting Group, and the TomTom Traffic Index [2].

The economic consequences of this congestion are substantial. A joint report by the National Economic and Development Authority (NEDA) and the Japan International Cooperation Agency (JICA) estimated that daily transportation-related economic losses in Metro Manila reached approximately PHP 3.5 billion in 2017, with an additional PHP 2.4 billion in adjacent regions. Without significant interventions, these costs are projected to exceed PHP 5 billion per day by 2035, alongside a 25% increase in daily trip demand [3]. These trends highlight the urgent need for data-driven, computationally supported approaches to traffic management, planning, and policy evaluation.

One promising avenue for addressing urban traffic problems is the use of traffic simulators. Simulation enables researchers and planners to recreate complex traffic scenarios in a controlled digital environment, allowing for systematic experimentation without the financial, logistical, and safety risks associated with real-world trials. Modern computational tools make it possible to evaluate multiple road configurations, traffic control strategies, and infrastructure designs within short time frames. Moreover, simulation facilitates scenario analysis under hypothetical conditions such as future demand growth or infrastructure changes that cannot be easily tested in practice. For rapidly urbanizing contexts like the Philippines, scalable and realistic traffic simulation is therefore an essential component of smart city and transportation planning.

Traffic simulation models are typically classified into three levels of abstraction: microscopic, mesoscopic, and macroscopic. Microscopic models represent individual vehicles and their interactions, incorporating detailed car-following and lane-changing behavior. While highly accurate, these models are computationally expensive, especially for large urban networks, and require extensive calibration. Commercial tools such as VISSIM [4] and AIMSUN [5] exemplify this approach and are well suited for analyzing localized congestion, complex geometries, and detailed operational strategies, but are less practical for city-scale modeling. In contrast, macroscopic traffic models treat traffic as a continuous compressible fluid rather than as discrete vehicles. Instead of tracking individual cars, these models describe traffic in terms of aggregate quantities such as density, flow, and average speed along road segments. This abstraction dramatically reduces computational complexity, making macroscopic models attractive for large-scale network analysis, real-time estimation, and long-term planning. Early macroscopic models were developed for highways and arterial corridors, but recent advances have extended their applicability to urban road networks with intersections, traffic signals, and heterogeneous traffic patterns [6]. A foundational example of this class of models is the Cell Transmission Model (CTM) introduced by Daganzo in the 1990s [7, 8]. CTM discretizes roads into homogeneous cells and models the movement of vehicles between cells using a conservation law for traffic density. The relationship between density and flow is governed by a fundamental diagram [9], which can be nonlinear and calibrated from empirical data. Over time, CTM has been extended to incorporate stochastic inflows, dynamic boundary conditions, and large-scale computational implementations, including cloud-based and parallelized frameworks for city-scale networks [10-14]. These developments demonstrate the strong potential of macroscopic models for practical traffic management and prediction. Meanwhile, Philippine traffic researchers have developed a software called Customized Local Traffic Simulation (LocalSIM) that was funded by DOST-PCIEERD. It currently features a two-dimensional microscopic modeling functionality. However, it currently lacks a macroscopic module capable of efficiently simulating large road networks for performing rapid scenario testing

[15]. Integrating a robust macroscopic model into LocalSIM could significantly enhance its capabilities by enabling multi-scale analysis, faster computation, and broader policy experimentation.

In response to this gap, this paper investigates the computational implementation of a macroscopic traffic model for urban networks based on the framework proposed in a prior work on large-scale urban traffic design. Rather than focusing solely on theoretical modeling, we treat macroscopic traffic flow as a computational modeling problem, emphasizing numerical stability, algorithmic efficiency, and software modularity. Our objective is to develop a scalable and computationally efficient implementation that can serve as a potential macroscopic component for LocalSIM and similar platforms.

Specifically, we replicate and extend a PDE-based macroscopic traffic model [16] grounded in CTM, incorporating intersection dynamics through merge models, turn-rate matrices, and binary traffic signals. We introduce key computational enhancements including a revised flux-limiting strategy, an underflow correction mechanism, and a topology-aware evaluation scheme based on topological sorting of roads to ensure correct and efficient simulation across interconnected networks. To validate the model, we implement it in Python and test it on three progressively complex road configurations: a single intersection, serial intersections, and a four-intersection grid. These experiments demonstrate that the model captures realistic congestion patterns, traffic wave propagation, and signal-induced delays while remaining computationally lightweight and scalable. By bridging traffic flow theory and practical computational implementation, this work contributes a modeling and simulation framework that supports faster, city-scale traffic analysis in the Philippine context. It lays the groundwork for integrating macroscopic modeling into local simulation tools and advancing data-driven traffic management for smart and sustainable urban mobility.

2 Macroscopic Traffic Model

The model is based on a discretized conservation law for traffic density and follows the general structure of the CTM, while incorporating additional mechanisms for driver foresight, intersection dynamics, and traffic control [16]. Our focus is not only on the mathematical formulation but also on its computational realizability, numerical stability, and scalability for urban road networks.

2.1 Flow Equation and Fundamental Diagram

We model traffic flow along a road as a one-dimensional compressible fluid, governed by the nonlinear transport equation

$$\frac{\partial \rho(x,t)}{\partial t} + \frac{\partial \phi(\rho(x,t))}{\partial x} = 0, x \in \mathbb{R}, t \geq 0, \quad (1)$$

where $\rho(x,t)$ denotes traffic density (vehicles per meter) and $\phi(\rho)$ represents the traffic flux (vehicles per second). The flux is expressed as

$$\phi(\rho, x, t) = \rho(x, t) v(\rho), \quad (2)$$

linking density to average driving speed $v(\rho)$ through a fundamental diagram (FD). The FD characterizes how speed decreases as density increases, capturing the transition from free flow to congestion. An example of an FD is shown in Figure 1. We adopt a flexible, piecewise-linear parameterization of the FD, which can be calibrated from empirical data and integrated directly into the numerical scheme [17].




Figure 1. Parameterization of a generic, piecewise-linear FD, where each linear segment is indicated. The figure is adopted from [17].

To account for anticipatory driving behavior, we incorporate a foresight coefficient $\alpha \in [0,1)$, which blends the local speed with the speed in the downstream cell

$$v_i = (1 - \alpha) v(\rho_i) + \alpha v(\rho_{i+1}). \quad (3)$$

This term introduces nonlocal influence into the model, allowing drivers to react to downstream congestion. All model parameters including FD coefficients and foresight are collected in a vector θ , enabling systematic tuning and sensitivity analysis. In this study, we adopted

$$\theta = [-1.8, 13.8, 0.0598, 0.095, 2.3, 0.13, 0.295]$$

from [17].

2.2. Discretization and Numerical Scheme

We discretize each road into uniform cells of length Δx and advance the solution in time using a finite-volume update:

$$\rho_i^{n+1} = \rho_i^n + \frac{\Delta t}{\Delta x} (\phi_{i-1}^n - \phi_i^n), \quad (4)$$

which enforces conservation of vehicles across adjacent cells. The time step Δt is constrained by the Courant–Friedrichs–Lewy (CFL) condition $\Delta x / \Delta t \leq v_{\max}$ to ensure numerical stability.

To handle boundary interactions, we introduce fictitious cells at both ends of each road. These cells encode inflow and outflow boundary conditions and allow the same update rule to be applied uniformly across the domain. A key design principle of our implementation is that the scheme remains first-order accurate in time and space, computationally lightweight, and linearly scalable with the number of cells and roads. This makes the approach suitable for large urban networks where microscopic simulation would be infeasible.




Figure 2: The three-point stencil of the discretized mass balance is indicated by arrows in the discretized space-time-grid (horizontal and vertical axis, respectively). The physical domain is represented by cell indexes $i = 1$ to $i = I$ (light green) with one ghost cell at each end of the domain with index $i = 0$ and $i = I + 1$ (light blue). The figure is adopted from [17].

2.3. Flux Limiting and Density Constraints

Physical traffic density must satisfy $0 \leq \rho_i^n \leq \rho_{\max}$ at all times. To enforce this, we modify the original flux-limiting approach in a prior work with a more numerically robust formulation. Let

$$\phi_{i,\text{incoming}}^n = \rho_{i-1}^n v_{i-1}^n, \quad (5)$$

$$\phi_{i,\text{remaining}}^n = \frac{\rho_{\max} - \rho_i^n}{\Delta t / \Delta x} + \phi_i^n. \quad (6)$$

We then define the actual inflow as

$$\phi_{i-1}^n = \min(\phi_{i,\text{incoming}}^n, \phi_{i,\text{remaining}}^n). \quad (7)$$

This modification prevents cell overflow while avoiding numerical artifacts due to floating-point precision. Additionally, we enforce non-negativity explicitly via

$$\rho_i^{n+1} = \max(\rho_i^{n+1}, 0). \quad (8)$$

Together, these steps guarantee that all simulated densities remain physically meaningful throughout the computation.

2.4. Intersections, Turn Rates, Traffic Signal, and Merge Model

Intersections are modeled using a supply-demand framework inspired by Daganzo's network formulation of the CTM [8]. At each intersection, a turn-rate matrix specifies the fraction of vehicles from each incoming road that proceed to each outgoing road. These matrices can vary over time, allowing the model to represent directional preferences or routing policies. Traffic signals are represented by a binary regulation matrix that alternates between green and red phases according to cycle time t_{cycle} , primary green time t_{green} , and phase offset t_{shift} . While simplified, this representation captures the essential effects of signal timing on congestion formation and dissipation.

At each junction of an intersection, multiple upstream roads feed into one or more downstream roads. If total incoming demand exceeds downstream capacity, we proportionally reduce inflows using flux reduction factors while preserving the relative contribution of each incoming road. This ensures no downstream cell exceeds capacity and that flow is distributed fairly among competing approaches. A crucial computational implication is that downstream roads must be updated before upstream roads. This dependency motivates the topology-aware evaluation strategy. Because intersection flows depend on downstream conditions, the road network forms a directed dependency graph. To compute updates correctly and efficiently, we perform a topological sort of roads prior to simulation. This ordering guarantees that, at each time step, all downstream links are processed first, upstream roads receive consistent boundary information, and that the entire network can be updated in a single forward pass without iterative correction. This strategy is both algorithmically clean and computationally efficient, particularly as network size grows.

3 Results and Discussion

We evaluate the proposed macroscopic traffic model through a series of numerical experiments on three progressively complex road network configurations: (1) a single intersection, (2) serial intersections, and (3) a four-intersection grid network. For all simulations, roads are discretized into 50 m cells, and traffic density evolution is visualized as a space-time heatmap, where darker regions correspond to higher vehicle density. Traffic signal states are overlaid on the temporal axis to illustrate their influence on flow dynamics.




Figure 3: Configurations of single intersection (top left), serial (top right), and four-intersection grid network (bottom) simulated in this study.

Our analysis focuses on three key aspects: **model fidelity** (whether the model exhibits qualitatively realistic traffic behavior), **computational consistency** across scales, and **impact of network topology and signal coordination** on congestion patterns.

Figure 4 illustrates the space-time evolution of traffic density for a single four-leg intersection. Roads 1 and 2 serve as inflow links, while Roads 3 and 4 are outflow links. The results demonstrate several expected macroscopic traffic phenomena such as

- (i) **Signal-induced shockwaves:** During red phases, vehicles accumulate upstream of the intersection, forming high-density regions that propagate backward along the inflow roads. When the signal turns green, these density waves dissipate and transition to free-flow conditions.
- (ii) **Intermittent downstream flow:** Outflow roads exhibit alternating bands of higher and lower density corresponding to primary and secondary green phases, confirming that the binary traffic signal model effectively modulates throughput.
- (iii) **Conservation of flow:** The discretized mass balance maintains consistent vehicle movement from inflow to outflow without artificial loss or creation of density, indicating that the numerical scheme is stable and physically consistent.

Overall, the single-intersection case validates that the model captures basic but essential traffic dynamics using a lightweight macroscopic formulation rather than computationally expensive microscopic simulation.




Figure 4. Traffic densities over time of a road network in a single intersection.

The second experiment links two intersections in series, introducing spatial coupling and multi-stage congestion effects (Figure 5). Compared to the single-intersection case, we observe the following:

- (i) **Propagation of congestion through the network:** Traffic waves generated at the first intersection travel downstream and interact with the second intersection, creating compounded delay effects that would be difficult to analyze analytically without simulation.
- (ii) **Topology-dependent bottlenecks:** Road 3, which lies between the two intersections and has no external boundary condition, exhibits persistent high density near its upstream end. This arises from the interaction between upstream inflow and downstream signal control, demonstrating how intermediate links can become critical congestion points.
- (iii) **Temporal synchronization effects:** The staggered arrival of vehicles from different approaches (e.g., Road 5 versus Road 1 via Road 3) leads to variable density peaks on downstream links, illustrating the importance of signal coordination in network performance.

From a computational perspective, this configuration confirms that the model remains stable and interpretable as network complexity increases, and that the topology-aware update strategy correctly handles dependencies between roads.

The final experiment considers a grid of four interconnected intersections (Figure 6), representing a small urban subnetwork. Key observations include:

- (i) **Consistent behavior across scale:** Inflow roads exhibit similar signal-driven accumulation and release patterns as in smaller networks, indicating that the model scales predictably without requiring structural modifications.
- (ii) **Distributed congestion patterns:** Internal roads (3, 4, 7, 9) display recurring density waves generated not by external demand but by internal signal interactions, highlighting emergent network-level dynamics.

- (iii) **Cumulative delay effects:** Compared to the serial configuration, congestion persists longer and propagates across multiple intersections, demonstrating how local signal policies can have nonlocal impacts.

These results suggest that traffic behavior observed in small subnetworks can be extended to larger systems, provided that transport delays and signal offsets are appropriately modeled.




Figure 5. Traffic densities over time of a road network in a serial intersection.




Figure 6. Traffic densities over time of a road network in a four-intersection grid network.

In summary, the simulated experiments collectively show that:

1. The macroscopic model reproduces realistic congestion phenomena using a computationally efficient PDE-based formulation.
2. Network topology and signal timing critically shape traffic patterns, even in small grids.
3. The proposed numerical and algorithmic enhancements enable stable, scalable simulation across increasing network complexity.

Beyond qualitative traffic behavior, we assess the computational properties of the implementation:

1. Linear scalability: Because each road is discretized independently and updated via local flux calculations, computational cost grows approximately linearly with the number of cells and roads.
2. Single-pass time stepping: The topology-aware ordering ensures that each road is updated once per time step, avoiding iterative backtracking or relaxation methods that would increase runtime.
3. Numerical robustness: The modified flux-limiting scheme prevents density overflow, while the non-negativity constraint avoids underflow artifacts, leading to stable simulations even under high demand conditions.
4. Low resource requirements: All experiments were conducted in a standard Python environment (Google Colab), demonstrating that the approach is feasible without specialized hardware.

These properties make the framework suitable for integration into larger systems such as the LocalSIM, where fast macroscopic simulations could complement more detailed microscopic modules.

4 Conclusion and Recommendations

This paper presented a scalable and computationally efficient implementation of a macroscopic traffic model for urban road networks based on the cell transmission model. By treating traffic as a continuous flow governed by a discretized conservation law, we demonstrated that realistic traffic dynamics can be captured without the high computational cost of microscopic simulation. Beyond direct replication of a prior work, we introduced several practical computational enhancements that improve the robustness and usability of the model. These include a revised flux-limiting strategy to prevent numerical overflow, an explicit non-negativity constraint to avoid density underflow, and a topology-aware evaluation scheme based on topological sorting of roads to correctly and efficiently resolve dependencies across interconnected links. Together, these modifications enable stable, single-pass updates of the network at each time step.

Through simulations on three progressively complex configurations, we showed that the model reproduces key macroscopic traffic phenomena such as shockwave propagation, signal-induced congestion, and network-level delay effects. The results highlight how local traffic signal policies can generate nonlocal impacts across a network, emphasizing the value of computational simulation for planning and analysis. From a systems perspective, the framework scales linearly with network size and remains lightweight enough to run on standard computing environments, making it a viable candidate for integration into larger traffic simulation platforms. In particular, this work provides a computational foundation for incorporating a macroscopic module into LocalSIM, complementing its existing microscopic capabilities and enabling multi-scale traffic analysis for Philippine cities.

While the proposed framework successfully demonstrates scalable macroscopic traffic simulation, several

extensions can further enhance its realism, applicability, and integration into operational systems such as More realistic network configurations, dynamic and time-varying demand, and advanced traffic signal control. While static density heatmaps are useful for analysis, animated visualizations mapped directly onto spatial road layouts would improve interpretability for planners and decision-makers. To assess quantitative accuracy, the macroscopic model could be validated against a microscopic simulator such as LocalSIM or VISSIM. Comparative studies would clarify when macroscopic modeling is sufficient and when microscopic detail is necessary. Overall, this work establishes a practical starting point for city-scale macroscopic traffic simulation in the Philippines. With continued development, the framework can support faster scenario analysis, smarter traffic management, and evidence-based urban planning in line with national smart city initiatives.

References

- [1] J. R. F. Regidor, "Traffic congestion in metro manila: Is the uvvrp still effective?" *Philippine Engineering Journal*, vol. 34, no. 1, 2013.
- [2] TomTom. "TomTom Traffic Index 2025." Accessed March 1, 2026. <https://www.tomtom.com/traffic-index/ranking/>
- [3] J. I. C. A. (JICA) and N. E. D. A. (NEDA), Roadmap for Transport Infrastructure Development for Metro Manila and its Surrounding Areas (Region III & Region IV-A) Final Report Summary. The National Economic and Development Authority, 2017.
- [4] M. Fellendorf and P. Vortisch, "Microscopic traffic flow simulator VISSIM," in *Fundamentals of traffic simulation*, pp. 63-93, Springer, 2010.
- [5] J. Barcelo and J. Casas, "Dynamic network simulation with AIMSUN," in *Simulation approaches in transportation analysis*, pp. 57-98, Springer, 2005.
- [6] U. D. of Transportation Federal Highway Administration, "Types of traffic analysis tools," Nov 2020.
- [7] C. F. Daganzo, "The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory," *Transportation Research Part B: Methodological*, vol. 28, no. 4, pp. 269-287, 1994.
- [8] C. F. Daganzo, "The cell transmission model, part ii: network traffic," *Transportation Research Part B: Methodological*, vol. 29, no. 2, pp. 79-93, 1995.
- [9] B. Greenshields, J. Bibbins, W. Channing, and H. Miller, "A study of traffic capacity," in *Highway research board proceedings*, vol. 1935, National Research Council (USA), Highway Research Board, 1935.
- [10] L. Munoz, X. Sun, R. Horowitz, and L. Alvarez, "Traffic density estimation with the cell transmission model," in *Proceedings of the 2003 American Control Conference*, 2003, vol. 5, pp. 3750-3755, IEEE, 2003.
- [11] A. Sumalee, R. Zhong, T. Pan, and W. Szeto, "Stochastic cell transmission model (sctm): A stochastic dynamic traffic model for traffic state surveillance and assignment,"

Transportation Research Part B: Methodological, vol. 45, no. 3, pp. 507-533, 2011.

[12] A. Sumalee, T. Pan, R. Zhong, N. Uno, and N. Indra-Payoong, "Dynamic stochastic journey time estimation and reliability analysis using stochastic cell transmission model: Algorithm and case studies," *Transportation Research Part C: Emerging Technologies*, vol. 35, pp. 263-285, 2013.

[13] D. Zehe, A. Knoll, W. Cai, and H. Aydt, "SemSim cloud service: Large-scale urban systems simulation in the cloud," *Simulation Modelling Practice and Theory*, vol. 58, pp. 157-171, 2015.

[14] T. Vondra and J. Sedivy, "Cloud autoscaling simulation based on queueing network model," *Simulation Modelling Practice and Theory*, vol. 70, pp. 83-100, 2017.

[15] P. M. Ejercito, K. G. E. Nebrija, R. P. Feria, and L. L. Lara-Figueroa, "Traffic simulation software review," in *2017 8th International Conference on Information, Intelligence, Systems Applications (IISA)*, pp. 1-4, 2017.

[16] E. Thonhofer, T. Palau, A. Kuhn, S. Jakubek, and M. Kozek, "Macroscopic traffic model for large scale urban traffic network design," *Simulation Modelling Practice and Theory*, vol. 80, pp. 32-49, 2018.

[17] E. Thonhofer, E. Luchini, and S. Jakubek, "A flexible, adaptive traffic network simulation with parameter estimation," *Journal of Intelligent Transportation Systems*, vol. 21, no. 1, pp. 63-77, 2017.

[18] G. M. Coclite, M. Garavello, and B. Piccoli, "Traffic flow on a road network," *SIAM journal on mathematical analysis*, vol. 36, no. 6, pp. 1862-1886, 2005.

Handwritten Gregg Shorthand Brief Forms Recognition Using Convolutional Neural Networks

Lysha T. Silva
South Philippine Adventist College
Matanao, Davao del Sur
silvalysha27@gmail.com

Neila M. Paglinawan-Muñoz
South Philippine Adventist College
Matanao, Davao del Sur
nmpmunez@gmail.com

ABSTRACT

Handwritten Gregg Shorthand recognition presents unique challenges due to high visual similarity among symbols, handwriting variability, and the absence of large, publicly available datasets. In particular, Gregg Shorthand brief forms consist of compact cursive symbols that represent commonly used words and require accurate visual discrimination for effective recognition. This study evaluated and compared three adapted CNN models—AlexNet, Inception-V3, and EfficientNetB0—for recognizing handwritten Gregg Shorthand brief forms using the developed GreggBF Dataset, which includes 3,753 original handwritten samples expanded to 67,554 images through data augmentation across 139 classes. All models were trained using transfer learning and assessed using accuracy, F1-score, and the Receiver Operating Characteristic Area Under the Curve (ROC-AUC), along with confusion matrix analysis to examine misclassification patterns. EfficientNetB0 achieved the best performance with 92.09% accuracy, a 91.85% F1-score, and a ROC-AUC of 0.9996, outperforming AlexNet (88.37%) and Inception-V3 (86.24%). Most errors stemmed from fine-grained visual similarities and handwriting inconsistencies. The results indicate that compoundly scaled architectures may be well suited for this task and also suggest that CNNs exhibit partial translation invariance, allowing the models to distinguish symbols that differ primarily in position. Future work may focus on improving dataset consistency and exploring more advanced compoundly scaled models to further enhance recognition accuracy.

KEYWORDS

Convolutional Neural Networks, Gregg Shorthand, handwritten symbol recognition, transfer learning

1 INTRODUCTION

Shorthand, also known as stenography, is a writing system designed for rapid transcription and has been widely used in journalism, court reporting, education, and personal note-taking [15]. One of the most commonly used shorthand systems is Gregg Shorthand. Developed by John Robert Gregg, this system represents phonemes using flowing cursive strokes that can be written with ordinary pens, contributing to its widespread adoption [5].

Despite advances in machine stenography, handwritten Gregg Shorthand remains relevant, especially for personal use and archival preservation. A key feature of this system is the use of brief forms, a set of 139 symbols representing commonly used words that are even more abbreviated than the standard Gregg Shorthand [28]. This enhance writing speed but pose recognition challenges due to high visual similarity and deviations from standard phoneme-to-symbol mappings.

Early shorthand recognition research relied on template matching and dynamic programming, particularly for Pitman and Teeline shorthand [2, 12], but these approaches were limited by handwriting variability and segmentation issues when applied to Gregg Shorthand. To address these limitations, Zhai et al. [25] proposed a Convolutional Recurrent Neural Network for recognizing general English vocabulary. Although the model demonstrated strong feature extraction capabilities, it struggled with sequence prediction due to the inherent complexity of the English language, particularly phoneme–spelling inconsistencies. In contrast, other studies focused on recognizing a limited set of English vocabulary—such as legal terminologies, brief forms, or predefined word sets—by treating each word as a single class rather than employing sequence-based prediction. Within these constrained vocabularies, this word-level classification strategy yielded promising recognition results. These works explored various models, including Artificial Neural Networks [18], CNN-based recognition using Inception-V3 [17], and K-Nearest Neighbors implemented in mobile applications [3]. While these studies demonstrated the feasibility of Gregg Shorthand recognition, they were limited by scalability concerns, modest recognition accuracy, the use of traditional machine learning algorithms, and/or reliance on relatively small dataset and simple data augmentation strategies.

CNNs are well established for image recognition tasks because of their ability to hierarchically extract visual features [4] and have consistently outperformed traditional machine learning approaches like ANN and KNN in complex visual classification problems [9, 14, 19]. Transfer learning with pre-trained CNNs has further improved performance across multiple domains [7, 20, 22, 24]. However, despite these strengths, the aforementioned CNN-based model by Padilla et al. [17] achieved only 54.74% accuracy in recognizing Gregg Shorthand legal terminologies. The relatively low accuracy may be influenced by limitations in the data and training approach, including methodological choices in model development. This highlights a research gap in understanding the effectiveness of CNN-based methods for handwritten Gregg Shorthand recognition.

To address this gap, this study adapts and evaluates three pre-trained CNN architectures—AlexNet [8], Inception-V3 [21], and EfficientNetB0 [23]—representing both historically influential and modern CNN designs. By treating brief forms as distinct visual classes and applying transfer learning, the study assesses how different CNN architectures perform on handwritten Gregg Shorthand brief form recognition. The paper presents a comparative evaluation of these models and analyzes their performance in handwritten shorthand recognition.

The study seeks to answer the following questions:

- (1) How do the adapted AlexNet, Inception-V3, and EfficientNetB0 models' performance compare in terms of accuracy, F1-score, and ROC-AUC when recognizing handwritten Gregg Shorthand brief forms?
- (2) Which specific Gregg Shorthand brief forms are most commonly misclassified by AlexNet, Inception-V3, and EfficientNetB0, as revealed by their respective confusion matrices?

To address the research questions, the following sections detail the data preparation and experimental setup, analyzes the model performance, and derives conclusions and recommendations from the findings.

2 METHODOLOGY

2.1 Dataset

Because no publicly available dataset was sufficiently robust to support this study, the GreggBF Dataset was created by collecting handwritten Gregg Shorthand brief forms from 30 senior Bachelor of Science in Office Administration (BSOA) students who had completed their stenography courses. Participation followed ethical protocols, with written approval obtained from the department chairperson and informed consent secured from all contributors.

Each participant wrote the shorthand forms on a provided worksheet using a pen. Out of 29 submitted handwriting sets, two were excluded due to irregular symbols, leaving 27 valid handwriting sets (5 male, 22 female) for the dataset. The handwritten samples were digitized using a flatbed scanner under consistent lighting conditions to ensure image uniformity and quality. All images were then annotated for use in training and evaluating the recognition models.

The dataset comprises 139 classes, each containing 27 original images, for a total of 3,753 samples. It is structured to provide meaningful variability for training and evaluating CNN-based models and is publicly available on GitHub [13]. The repository also contains the complete compilation of properly written Gregg Shorthand brief forms reference symbols.

2.2 Data Preprocessing

The dataset was first split into training, validation, and testing subsets using a 70/15/15 ratio, performed before data augmentation to prevent data leakage and ensure unbiased evaluation. After splitting, a range of data augmentation techniques was applied to enhance model generalization and reduce overfitting. These included three levels of Gaussian blur, two levels of square dilation, three levels of disk-shaped dilation, two levels of brightness adjustment, three levels of contrast adjustment, and two levels of zooming transformations, simulating natural variability in handwritten inputs. Techniques such as rotation and shearing were intentionally excluded, as changes in angle or stretch could alter the meaning of Gregg shorthand symbols.

Table 1 presents a summary of the GreggBF dataset, showing its class distribution, participant count, and image composition both before and after data augmentation.

Table 1: Augmented GreggBF Dataset Summary

Attribute	Formula	Count
No. of Classes	-	139
No. of Participants	-	27
Original Dataset Size	139 classes $\times$ 27 participants	3,753
No. of Augmentations	-	17
Total No. of Images Per Class (with augmentation)	(17 augmentations + 1 original image) $\times$ 27 participants	486
Augmented Dataset Size	486 images per class $\times$ 139 classes	67,554

2.3 Classifiers

AlexNet, Inception-V3, and EfficientNetB0 were selected for their proven performance in image classification and their diverse architectural characteristics. AlexNet, one of the first effective deep CNNs, is relatively shallow and provides a foundational benchmark [8]. Inception-V3 is deeper and wider, with multi-scale feature extraction and auxiliary classifiers, and has previously been applied to Gregg shorthand recognition of legal terminologies [17]. EfficientNetB0, a more recent architecture, employs compound scaling to balance depth, width, and resolution, achieving high accuracy with computational efficiency [23].

All three networks were used with transfer learning, leveraging pre-trained representations from large datasets and fine-tuning them on the GreggBF dataset. This approach reduces training time, improves generalization, and allows investigation of whether the dataset or task is better suited to shallower versus more modern CNN architectures.

2.4 Hyperparameters for Training and Fine-Tuning

As established earlier, this study employed transfer learning using pre-trained CNN models, enabling the networks to leverage existing feature representations while adapting to the GreggBF dataset. To optimize training, key hyperparameters were carefully configured for both the feature extraction and fine-tuning phases (Table 2).

Input images were resized according to the input requirements specified in each model's official documentation. Preprocessing for Inception-V3 and EfficientNetB0 was handled internally via Keras rescaling layers, while AlexNet required explicit normalization based on ImageNet statistics. Images were processed in batches of 32, with training data shuffled each epoch for better generalization, while validation and testing sets were kept in a consistent order.

Training consisted of two phases: a feature extraction phase with frozen base weights for 20 epochs, followed by a fine-tuning phase where the entire network was trained for 10 additional epochs. Learning rates were adaptively reduced when validation

loss plateaued, with a minimum threshold of $1e-6$. Early stopping was applied if no improvement was observed over five consecutive epochs, ensuring efficient training and robust model performance.

Table 2: Key Hyperparameters for Training and Fine-Tuning the CNN Models

Hyperparameter	AlexNet	Inception-V3	EfficientNetB0
Image size	224 x 224 x 3	299 x 299 x 3	240 x 240 x 3
Batch size	32	32	32
Feature extractor phase			
Epochs	20	20	20
Learning rate	0.001	0.001	0.001
Reduced LR factor	0.2	0.2	0.2
Reduced LR minimum	$1e-6$	$1e-6$	$1e-6$
Fine-tuning phase			
Epochs	10	10	10
Learning rate	$1e-5$	$1e-5$	0.001
Reduced LR factor	0.2	0.2	0.2
Reduced LR minimum	$1e-6$	$1e-6$	$1e-6$

2.5 Model Evaluation

The trained models were evaluated using standard performance metrics, including accuracy, F1-score, ROC-AUC, and the confusion matrix.

Accuracy measures the overall proportion of correct predictions and is defined as:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

where TP , TN , FP , and FN are the numbers of true positives, true negatives, false positives, and false negatives, respectively.

The F1-score, defined as

$$\text{F1-score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}},$$

was included to provide a single measure that balances the model's ability to correctly identify the correct class while avoiding mistakes. It helps capture performance in a way that accuracy alone might miss.

ROC-AUC (Receiver Operating Characteristic Area Under the Curve) evaluates how well the model can tell different classes apart. Since this study involves many classes, the One-vs-One (OvO) strategy was used [6], which compares each class against every other class in pairs and then averages the results. This approach was chosen because it provides a more detailed view of how confidently the model distinguishes between any two classes, making it especially informative for multi-class problems like shorthand recognition, compared to accuracy alone.

The confusion matrix provides a detailed view of the model's predictions for each class, showing the counts of true positives, false positives, true negatives, and false negatives. In this study, the matrix consisted of 139 rows and 139 columns, corresponding to all brief form classes, enabling the identification of class pairs that are frequently confused.

3 RESULTS

3.1 Performance Analysis and Comparison of the Adapted CNN Models

Table 3: Adapted CNN Models' Performance Metrics Comparison

Model	Accuracy	F1-Score	ROC-AUC
AlexNet	0.8837	0.8818	0.9973
Inception-V3	0.8624	0.8603	0.9963
EfficientNetB0	0.9209	0.9185	0.9996

This study compared AlexNet, Inception-V3, and EfficientNetB0 on handwritten Gregg Shorthand brief forms using accuracy, F1-score, and ROC-AUC (Table 3). EfficientNetB0 performed best across all metrics, suggesting that its combination of balanced depth, width, and resolution, as well as features such as MBConv blocks and Swish activation, could have contributed to its strong performance on this dataset. The results also indicate that optimal performance does not rely solely on network depth, highlighting the potential of modern, efficiently scaled CNNs for this task. The trained EfficientNetB0 model is publicly available on GitHub [13].

Compared to Padilla et al. [17], whose Inception-V3 model reported lower accuracy, differences in evaluation methodology may partly explain the performance gap. They tested their model using live camera input, which could introduce variability from lighting, angles, or camera quality, whereas this study used a dataset split for more consistent conditions. It is possible that their Inception-V3 model might have performed better under similar evaluation settings. Furthermore, Padilla et al. did not clearly indicate whether data augmentation, fine-tuning, or transfer learning were applied, while this study incorporated careful preprocessing and training strategies, highlighting how the design of the training and evaluation pipeline can be just as critical as the choice of model architecture.

In addition to these comparisons, it is worth considering how the EfficientNetB0 model could contribute to more complex shorthand recognition tasks beyond isolated brief forms. So far, the closest work to a sequence-to-sequence recognition task is that of Zhai et al. [25], who tackled recognition of one word at a time using a CRNN. Each word is composed of multiple symbols, and some words contain brief forms within them. In their implementation, a CNN served as a feature extractor, while the recurrent layers modeled the sequence of symbols within each word. The EfficientNetB0 model used in this study could potentially replace the CNN feature extractor in their framework.

Moreover, these findings are aligned with recent studies [11, 16, 27] showing EfficientNetB0 often outperforms both shallower and

deeper CNNs across various domains, though further investigation is needed to confirm its effectiveness in handwritten shorthand recognition.

3.2 Adapted Models’ Confusions

Table 4: Summary of Top 5 Most Frequent Misclassifications Across All Adapted Models

Model	Ground Truth	TP Count	Most Frequent FP
AlexNet	you/your	18	42 x under
	order	26	27 x under
	they	27	30 x thing/think
	success	35	37 x suggest
	under	35	36 x you/your
Inception-V3	order	28	35 x what
	organize	28	35 x work
	what	33	15 x short
	and	36	18 x their
	any	37	16 x immediate
EfficientNetB0	order	23	33 x what
	when	25	15 x any
	you/your	26	46 x under
	where	37	35 x electric
	state	38	23 x short

The analysis of the adapted models’ confusions provides a detailed view of the recognition challenges in handwritten Gregg Shorthand brief forms. Due to the large size of the full confusion matrices (139×139), only the top five most frequently misclassified class pairs for each model are summarized in Table 4, which includes true positives and false positives for each pair. The complete confusion matrices are available on GitHub [13].

Shared misclassification patterns were observed across multiple models, indicating persistent recognition difficulties. Notably, both AlexNet and EfficientNetB0 frequently confused “you/your” with “under,” while all three models misclassified “order” as “what.” These recurring errors suggest that certain symbol pairs are inherently difficult to distinguish across different CNN architectures.

Some top misclassification pairs were unique to individual models, reflecting differences in feature extraction and pattern recognition. For example, AlexNet often misclassified “they” as “thing/think,” Inception-V3 confused “organize” with “work,” and EfficientNetB0 frequently misclassified “state” as “short.” These distinct patterns highlight how architectural differences influence the interpretation of specific shorthand forms, despite identical training on the same dataset.

3.2.1 Model Confusions Arising from Fine-Grained Visual Challenges and Inconsistent Handwritten Inputs. Visual comparisons, provided in Table 5, show the ground truth symbols alongside participant-written false-positive samples that closely resemble the predicted classes. Many misclassifications involved subtle visual

differences in loops, curves, or dots, which are generally visible to human readers but difficult for the models to capture consistently. These results suggest that models have difficulty capturing fine-grained visual details in certain symbols.

However, certain pairs were particularly challenging due to inconsistencies in handwritten samples. For example, the symbols “order” and “what,” shown as properly written examples in Figure 1, share a similar structure and differ mainly in stroke length; however, the participant-written false-positive samples in Table 5 often failed to preserve this distinction, leading to widespread confusion.




Figure 1: Properly written reference examples of “order” and “what” shorthand symbols




Figure 2: Properly written examples of “of” and “over” shorthand symbols

3.2.2 Misclassification Analysis Suggest Partial Translation Invariance in CNNs. Positional distinctions in shorthand symbols played a significant role in both successful recognition and misclassification. The brief forms for “you/your” and “under” share identical stroke shapes and are distinguished solely by vertical placement: “under” should be written at the top of the writing box, while “you/your” should be centered. In the participant inputs shown in Table 5, however, both symbols were frequently written at the center, violating this positional distinction and rendering them visually nearly identical. Under these conditions, misclassification is expected regardless of how good a model could be.

In contrast, the symbols for “of” and “over,” shown in Figure 2, were consistently written with correct vertical placement despite sharing identical stroke shapes. As a result, misclassifications between these two symbols were minimal. All three models correctly classified all 72 test samples of “over,” and EfficientNetB0 further achieved 100% accuracy on “of,” suggesting that the models were

Table 5: Top five most frequent misclassification pairs per model, showing the ground truth label and a participant-written sample that visually resembles the predicted false positive.

AlexNet		Inception-V3		EfficientNetB0	
Ground Truth	Most Frequent FP	Ground Truth	Most Frequent FP	Ground Truth	Most Frequent FP
1	 you/your	 under	 order	 what	 order
2	 order	 what	 organize	 work	 when
3	 they	 thing/think	 what	 short	 you/your
4	 success	 suggest	 and	 their/there	 where
5	 under	 you/your	 any	 immediate	 state

able to utilize positional information when it was consistently represented in the input data — an effect that could be easily corrected for symbols like “you/your” vs. “under” with properly written samples.

While CNNs are often described as translation-invariant due to their convolutional architecture [10], prior research indicates that this invariance may be only partial in practice. For example, Azulay and Weiss [1] showed that small spatial shifts can significantly affect CNN predictions, and Zhang et al. [26] observed that translation invariance may break when input shifts do not align with the network’s stride. In the context of Gregg Shorthand recognition, this limited translation invariance may actually be advantageous: it allows the models to distinguish symbols that differ primarily in position, such as “of” and “over,” while naturally exhibiting confusion when positional cues are inconsistent, as in “you/your” vs. “under.” Properly written and standardized samples would likely resolve such confusions.

Overall, the confusion analysis highlights that accurate recognition of Gregg Shorthand requires a balance: models must be robust to natural handwriting variation while maintaining sensitivity to subtle structural and positional features. These findings underscore the importance of consistent dataset design, clear participant instructions regarding positional conventions, and model architectures capable of capturing fine-grained spatial details.

4 CONCLUSION

The comparative evaluation of the adapted CNN models demonstrated that EfficientNetB0 consistently outperformed AlexNet and

Inception-V3 across accuracy, F1-score, and AUC-ROC in recognizing handwritten Gregg Shorthand brief forms. This result may indicate that compound-scaled architectures are particularly effective for classification tasks involving handwritten symbols with subtle structural variations. Misclassifications were primarily attributed to limitations in capturing fine-grained features such as loops, curves, and dots, as well as inconsistencies in the handwritten samples, underscoring the importance of both robust model design and standardized dataset preparation.

Building on these findings, future research is encouraged to explore other compound-scaled and lightweight CNN architectures, such as EfficientNetV2 and MobileNetV3, which may further improve recognition accuracy while maintaining computational efficiency. Furthermore, the EfficientNetB0 model used in this study could serve as a feature extractor for future sequence-to-sequence recognition models, where isolated brief form recognition may be extended to full words composed of multiple symbols, including brief forms embedded within them. However, only its convolutional backbone would be applicable, as the final classification layers trained for the 139 brief forms are specific to isolated recognition and are not directly transferable to sequence-level tasks.

To mitigate errors caused by handwriting variability, future datasets should prioritize clearer and more standardized writing practices, particularly for commonly confused symbols (e.g., “order” vs. “what” and “you” vs. “under”). Implementing structured data collection methods, such as guided writing templates and explicit stroke instructions, may help ensure consistent representation of

critical visual features, enabling models to learn more reliable patterns.

In the further development of the dataset, researchers may also consider incorporating samples written using different colored pens and introducing uneven lighting conditions within images—such as shadows, glare, brightness gradients, or other illumination variations—or applying augmentation techniques to simulate these effects, thereby improving the model’s robustness to real-world writing conditions. Together, these directions provide a clear path for advancing shorthand recognition toward more practical and real-world applications.

5 ACKNOWLEDGEMENT

The authors gratefully acknowledge the financial support provided by the Department of Science and Technology (DOST).

REFERENCES

- [1] Aharon Azulay and Yair Weiss. 2019. Why Do Deep Convolutional Networks Generalize So Poorly to Small Image Transformations? *Journal of Machine Learning Research* 20, 184 (2019), 1–25.
- [2] V. C. Burkpalli, Shivaprakash, and B. S. Anami. 2018. A Region Based Design of Deterministic Finite State Automata for Online Recognition of Teeline Shorthand Language Alphabet. In *Proceedings of the International Conference on Cognition and Recognition (ICCR 2016)*. Springer Singapore, 175–186.
- [3] D. A. Cardaña, M. A. D. Perin, S. G. Tabuno, E. A. Orapa, M. S. Dagupan, and J. M. G. Lagumbay. 2024. StenogrApp: E-learning Android Application in Recognition of Basic Gregg Shorthand Using Machine Learning. In *Proceedings of the 10th International Conference on Education and Training Technologies*. 43–48.
- [4] Ian Goodfellow and Yoshua Bengio. 2016. *Deep Learning*. MIT Press.
- [5] John Robert Gregg. 1931. *Gregg’s Shorthand*. Ripol Klassik.
- [6] David J. Hand and Robert J. Till. 2001. A simple generalisation of the area under the ROC curve for multiple class classification problems. *Machine Learning* 45, 2 (2001), 171–186.
- [7] A. Kasimu, W. Zhou, Q. Meng, Y. Wang, Z. Wang, Q. Zhang, and Y. Peng. 2025. Performance evaluation of pretrained deep learning architectures for railway passenger ride quality classification. *Alexandria Engineering Journal* (2025). doi:10.1016/j.aej.2025.01.007
- [8] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems* 25 (2012).
- [9] S. Kumar, G. Charlyn, P. Latha, and O. Hemavathy. 2022. Palm Print Identification and Classification in the Field of Biometric in CNN using KNN. In *2022 14th International Conference on Mathematics, Actuarial Science, Computer Science and Statistics (MACS)*. 1–4. doi:10.1109/MACS56771.2022.10022776
- [10] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-Based Learning Applied to Document Recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [11] J. Y. Lee, Y. S. Lee, J. H. Tae, I. H. Chang, T.-H. Kim, S. C. Myung, T. T. Nguyen, J. H. Lee, J. Choi, J. H. Kim, J. W. Kim, and S. Y. Choi. 2024. Selection of Convolutional Neural Network Model for Bladder Tumor Classification of Cystoscopy Images and Comparison with Humans. *Journal of Endourology* 38, 10 (2024), 1036–1043. doi:10.1089/end.2024.0250
- [12] C. G. Leedham and A. C. Downton. 1987. Automatic Recognition and Transcription of Pitman’s Handwritten Shorthand—An Approach to Shortforms. *Pattern Recognition* 20, 3 (1987), 341–348.
- [13] Lysha T. Silva and Neila M. Paglinawan-Muñoz. 2025. GreggBF Dataset. <https://github.com/lyshasilva/GreggBF>.
- [14] C. Meng. 2024. Weather Forecast Image Classification Based on KNN, CNN, and VGG19. *Applied and Computational Engineering* (2024). doi:10.54254/2755-2721/76/20240583
- [15] Francis Jesmar P. Montalbo and D. P. Y. Barfeh. 2019. Classification of Stenography Using Convolutional Neural Networks and Canny Edge Detection Algorithm. (2019), 305–310.
- [16] N. K. Naik, P. K. Sethy, S. K. Behera, and R. Amat. 2024. A Methodical Analysis of Deep Learning Techniques for Detecting Indian Lentils. *Journal of Agriculture and Food Research* 15 (2024), 100943. doi:10.1016/j.jafr.2023.100943
- [17] D. A. Padilla, N. K. U. Vitug, and J. B. S. Marquez. 2020. Deep Learning Approach in Gregg Shorthand Word to English-Word Conversion. In *IEEE International Conference on Image, Vision and Computing*. 204–210.
- [18] R. Rajasekaran and K. Ramar. 2012. Handwritten Gregg Shorthand Recognition. *International Journal of Computer Applications* 41, 9 (2012), 31–38.
- [19] R. Sharma, R. Hussung, A. Keil, F. Friederich, T. Fromenteze, M. Khalily, B. Deka, V. Fusco, and O. Yurduseven. 2022. Performance Analysis of Classification Algorithms for Millimeter-Wave Imaging. In *2022 16th European Conference on Antennas and Propagation (EuCAP)*. 1–5. doi:10.23919/eucap53622.2022.9769429
- [20] Hoo-Chang Shin, Holger Roth, Mingchen Gao, Le Lu, Zhiyong Xu, Islem Nagues, Jian Yao, Daniel Mollura, and Ronald Summers. 2016. Deep Convolutional Neural Networks for Computer-Aided Detection: CNN Architectures, Dataset Characteristics and Transfer Learning. *IEEE Transactions on Medical Imaging* 35 (2016), 1285–1298. doi:10.1109/TMI.2016.2528162
- [21] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2818–2826.
- [22] Nima Tajbakhsh, Jae Y. Shin, S. Gurudu, R. T. Hurst, C. B. Kendall, M. Gotway, and Jianming Liang. 2016. Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning? *IEEE Transactions on Medical Imaging* 35 (2016), 1299–1312. doi:10.1109/TMI.2016.2535302
- [23] Mingxing Tan and Quoc Le. 2019. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In *International Conference on Machine Learning*. 6105–6114.
- [24] Yu Yang, L.-F. Yan, X. Zhang, Y. Han, H.-Y. Nan, Y.-C. Hu, B. Hu, S. Yan, J. Zhang, D.-L. Cheng, X.-W. Ge, G.-B. Cui, D. Zhao, and W. Wang. 2018. Glioma Grading on Conventional MR Images: A Deep Learning Study With Transfer Learning. *Frontiers in Neuroscience* 12 (2018). doi:10.3389/fnins.2018.00804
- [25] Feng Zhai, Yuchen Fan, Tushar Verma, Ramesh Sinha, and Dietrich Klakow. 2018. A Dataset and a Novel Neural Approach for Optical Gregg Shorthand Recognition. 11107 (2018), 222.
- [26] Richard Zhang. 2019. Making Convolutional Networks Shift-Invariant Again. In *Proceedings of the International Conference on Machine Learning (ICML)*. 7324–7334.
- [27] Wei Zheng. 2024. Comparison of Transfer-Learning for Lightweight Pre-Trained Model on Image Classification. *Applied and Computational Engineering* (2024). doi:10.54254/2755-2721/53/20241244
- [28] Charles E. Zoubek and George A. Condon. 1990. *Gregg Shorthand College Book 1* (centennial ed.). McGraw-Hill.

Integrating Adaptive Mechanisms for Feature Selection with Mayfly Algorithm, Monarch Butterfly Optimization, and Monarch Mayfly Optimization

Jane T. Antoque

South Philippine Adventist College
Matanao, Davao del Sur, Philippines
janeantoque7@gmail.com

Samuel R. Lucero

South Philippine Adventist College
Matanao Davao del Sur, Philippines
ausam0626@gmail.com

Neila M. Paglinawan-Muñoz

South Philippine Adventist College
Matanao, Davao del Sur, Philippines
nmpmunez@gmail.com

Frence Clifford O. Pillodar

South Philippine Adventist College
Matanao Davao del Sur, Philippines
pillodarfrence@gmail.com

ABSTRACT

Feature selection helps address the challenges of handling large datasets by reducing the number of attributes, thereby improving the accuracy and efficiency of machine learning models. This study investigates the enhancement of three metaheuristic algorithms for feature selection—the Mayfly Algorithm (MA), Monarch Butterfly Optimization (MBO), and Monarch Mayfly Optimization (MMO)—through the integration of adaptive mechanisms. Specifically, Adaptive Population Size (APS), Adaptive Restart Mechanism (ARM), and Adaptive Stopping Criterion (ASC) were incorporated into the base algorithms. The adaptive-enhanced algorithms were evaluated using K-Nearest Neighbors (KNN) and Support Vector Machine (SVM) classifiers across 18 benchmark datasets from the UCI Machine Learning Repository. Experimental results indicate that while not all adaptive mechanisms consistently improve performance, the Adaptive Restart Mechanism (ARM) yields notable and statistically supported improvements in classification accuracy. These findings suggest that restart-based adaptivity effectively guides metaheuristic algorithms toward better solutions by mitigating premature convergence.

CCS CONCEPTS

• Computing methodologies → Feature selection;

KEYWORDS

feature selection, metaheuristic algorithms, adaptive mechanisms, hybrid optimization algorithms

1 INTRODUCTION

In recent years, solving real-world problems in fields such as healthcare, finance, and cybersecurity has required the processing of increasingly complex, high-dimensional datasets. Although such datasets contain valuable information, they often introduce noise, redundant features, and computational burdens that negatively affect the performance and efficiency of machine learning (ML) models. Feature selection (FS) addresses this challenge by identifying the most relevant features, thereby reducing dimensionality and improving classification accuracy.

Traditional FS methods, such as exhaustive search and greedy heuristics, struggle with scalability and are prone to local optima, especially in noisy and high-dimensional environments. Metaheuristic algorithms—nature-inspired search strategies—have emerged as robust alternatives for FS. These include the Mayfly Algorithm (MA) [11], Monarch Butterfly Optimization (MBO) [10], and their hybrid variant, Monarch Mayfly Optimization (MMO) [4]. MA and MBO balance global exploration and local exploitation through different behavioral models, while MMO leverages the strengths of both.

However, despite the effectiveness of these metaheuristics, existing studies primarily evaluate MA, MBO, and MMO in their conventional, static-parameter forms. Prior research has not systematically integrated or compared adaptive mechanisms—specifically Adaptive Population Size (APS), Adaptive Restart Mechanism (ARM), and Adaptive Stopping Criterion (ASC)—within these algorithms. This leaves a gap in understanding how adaptivity affects performance, convergence behavior, and classification accuracy across diverse datasets.

Inspired by the hybridization efforts of Garcia [4], who demonstrated that MMO outperforms Mayfly-Harmony Search (MA-HS) and other algorithms across multiple UCI benchmark datasets, this study aims to further enhance MA, MBO, and MMO by incorporating adaptive mechanisms. These strategies have shown promise in other optimization contexts: restart mechanisms mitigate stagnation by reinitializing populations [3, 5], adaptive stopping criteria prevent unnecessary computation [1, 12], and dynamic population sizing maintains a balance between diversity and convergence [9].

To investigate the impact of adaptivity, this study evaluates adaptive-enhanced variants of MA, MBO, and MMO using two well-established classifiers—K-Nearest Neighbors (KNN) and Support Vector Machine (SVM)—across 18 benchmark datasets from the UCI repository. The goal is to assess whether APS, ARM, and ASC yield improvements in classification accuracy and to compare their effects both within and across the three base algorithms.

This paper makes three key contributions:

- (1) Integrates and evaluates three adaptive mechanisms—Adaptive Population Size (APS), Adaptive Restart Mechanism (ARM), and Adaptive Stopping Criterion (ASC)—within three prominent metaheuristic algorithms: Mayfly Algorithm (MA), Monarch Butterfly Optimization (MBO) and Monarch Mayfly Optimization (MMO) for feature selection.
- (2) Conducts a comparative analysis between adaptive and non-adaptive variants using consistent classifiers (KNN and SVM) and 18 benchmark datasets.
- (3) Identifies which adaptive mechanism (APS, ARM, or ASC) yields the most effective results.

The findings offer practical insights for researchers and practitioners seeking scalable, accurate, and adaptive feature selection pipelines for machine learning tasks.

To address the research objectives, this paper presents the experimental setup and adaptive-enhanced algorithms, analyzes the results using statistical evaluation, and concludes with key findings and future research directions.

2 METHODOLOGY

2.1 Datasets

Eighteen datasets from the UCI Machine Learning Repository were selected to ensure diversity in domain, dimensionality, and class distribution. These datasets span fields such as biology, politics, and game theory, with features ranging from 9 to 325 and sample sizes from 73 to 5000. Extended versions (e.g., BreastEW, HeartEW) were included to evaluate scalability on high-dimensional data. The complete list of datasets along with their characteristics is summarized in Table 1.

Table 1: Benchmark datasets used in the Study.

#	Dataset	Number of Features	Number of Samples	Number of Classes	Domain
1	Breastcancer	9	699	2	Biology
2	BreastEW	30	569	2	Biology
3	CongressEW	16	434	2	Politics
4	Exactly	13	1000	2	Biology
5	Exactly2	13	1000	2	Biology
6	HeartEW	13	270	2	Biology
7	IonosphereEW	34	351	2	Electromagnetic
8	KrvskpEW	36	3196	2	Game
9	Lymphography	18	148	4	Biology
10	M-of-N	13	1000	2	Biology
11	PenglungEW	325	73	2	Biology
12	SonarEW	60	208	2	Biology
13	SpectEW	22	267	2	Biology
14	Tic-tac-toe	9	958	2	Game
15	Vote	16	300	2	Politics
16	WaveformEW	40	5000	3	Physics
17	WineEW	13	178	3	Chemistry
18	Zoo	16	101	6	Artificial

2.2 Experimental Setup

This study followed a structured experimental procedure to evaluate the effectiveness of adaptive mechanisms in metaheuristic-based feature selection. The process began by integrating three adaptive techniques—Adaptive Population Size (APS), Adaptive Restart Mechanism (ARM), and Adaptive Stopping Criterion (ASC)—into three base algorithms: Mayfly Algorithm (MA), Monarch Butterfly Optimization (MBO), and Monarch Mayfly Optimization (MMO). The integrated algorithms were evaluated using two classifiers, K-Nearest Neighbors (KNN) and Support Vector Machine (SVM), across 18 benchmark datasets. Each adaptive variant was subjected to 10 independent runs to assess classification accuracy.

To ensure a fair comparison, key hyperparameters such as swarm size and maximum iterations were fixed across all algorithms. Specifically, the swarm size was set to 20 and the number of iterations to 20, following the designs used in Bhattacharyya et al. [2] and Garcia [4], where this configuration yielded reliable results. For MA and MMO, parameters such as attraction constants ($a_1 = 3$, $a_2 = 3.5$), visibility coefficient ($\beta = 0.1$), and nuptial dance coefficient ($d_0 = 0.9$) were retained from their original studies to maintain behavioral integrity. Additional parameters specific to MMO, such as gravitational coefficient ($g = 0.98$), Levy flight size ($s = 1$), and mutation strength (1.2), were incorporated to support its hybrid behavior.

For MBO and MMO, other behavioral parameters included the Monarch migration period (peri), Butterfly Adjusting Rate (BAR), and Mayfly Adjusting Rate (MAR), which controlled migration frequency and solution refinement. Table 2 summarizes the complete list of parameters and their values across the three algorithms.

2.2.1 Adaptive Population Size. Figure 1 presents the integration of the Adaptive Population Size (APS) mechanism into three different metaheuristic algorithms: (a) Mayfly Algorithm (MA), (b) Monarch Butterfly Optimization (MBO), and (c) Monarch Mayfly Optimization (MMO). Although each algorithm follows its own main search dynamics, they all adopt the same APS structure to dynamically adjust the population size during each iteration. The central idea of the APS mechanism is to evaluate two key metrics—diversity and convergence—to determine whether the current population size should be increased, decreased, or maintained. If

Table 2: Hyperparameters used across MA, MBO, and MMO algorithms.

Hyperparameter	Description	Value
S	Swarm size of $NP1$ and $NP2$	20
R_i	Ratio with respect to S	0.5
$MaxIter$	Maximum number of iterations	20
γ	Relative weightage used for fitness value	0.01
a_1	Positive attraction constant	3
a_2	Positive attraction constant	3.5
β	Visibility coefficient	0.1
d_0	Initial nuptial dance coefficient	3
f_{i0}	Initial random walk coefficient	3
r_{of}	Random value for crossover	0.95
g	Gravitational coefficient	0.98
δ	d_0 and f_{i0} update multiplier	0.9
s	Levy flight size	1
MAR	Mayfly adjusting rate	0.1
p	Probability threshold for Adjusting Operator	float(5/12)
S_{max}	Max walk-step of mayfly in Lévy flight	0.02
$peri$	Migration period	1.2
$MaxNP2Iter$	Iteration limit should $NP2$ continue without finding a better solution	20
$mutation_strength$	Mutation strength	1.2

diversity falls below a predefined threshold while the population size is still under the maximum limit, the algorithm increases its population to promote exploration. On the other hand, if the convergence rate drops below a set threshold and the population exceeds the minimum allowable size, the population is reduced to intensify exploitation and save computational resources.

In all three algorithms, the APS block is executed at the beginning of each iteration and directly influences the size of the swarm or subpopulations before the main optimization steps begin. This makes the APS mechanism a crucial adaptive control component that balances exploration and exploitation according to the algorithm’s ongoing performance. By embedding the same APS logic into different frameworks, the study isolates and evaluates its impact across varied optimization behaviors, allowing for a consistent comparison of its effectiveness. The uniformity of the APS logic across MA, MBO, and MMO emphasizes its generalizability and importance in improving search efficiency without modifying the core operations of each algorithm.

2.2.2 Adaptive Restart Mechanism. Figure 2 illustrates the integration of the Adaptive Restart Mechanism (ARM) into three metaheuristic algorithms: (a) Mayfly Algorithm (MA), (b) Monarch Butterfly Optimization (MBO), and (c) Monarch Mayfly Optimization (MMO). Although these algorithms differ in their core operations, they share a common adaptive strategy designed to counteract premature convergence. The ARM component is activated when no improvement is observed in the global best fitness ($gbest_fitness$) compared to the previous iteration.

When stagnation is detected, the mechanism reinitializes the positions of individuals—either mayflies or butterflies—around the current best solution. This targeted reinitialization is coupled with a complete reset of velocities (for algorithms with velocity components), enabling the population to explore new regions of the search space. By focusing the restart near a promising solution while eliminating directional memory, the algorithm retains beneficial knowledge while promoting diversity and exploration.

Following the restart, the population is re-sorted based on fitness to prepare for the next iteration. This ensures that high-quality individuals continue to guide the search. The ARM operates conditionally, triggering only when needed rather than at fixed intervals. This behavior improves efficiency by avoiding unnecessary resets and reducing computational overhead. As shown in the figure, ARM plays a crucial role across all three algorithms by enhancing their ability to escape local optima and maintain steady progress throughout the optimization process.




Figure 1: Flowcharts showing the integration of the Adaptive Population Size (APS) in (a) MA, (b) MBO, and (c) MMO.




Figure 2: Flowcharts showing the integration of the Adaptive Restart Mechanism (ARM) in (a) MA, (b) MBO, and (c) MMO.

2.2.3 Adaptive Stopping Criterion. Figure 3 illustrates how the Adaptive Stopping Criterion (ASC) is integrated into three metaheuristic algorithms: (a) Mayfly Algorithm (MA), (b) Monarch Butterfly Optimization (MBO), and (c) Monarch Mayfly Optimization (MMO). The ASC is a convergence-based mechanism that dynamically monitors the optimization progress using a sliding window of recent global best fitness (gbest_f) values. After each iteration, the current gbest_f is appended to the convergence history, which is then used to assess whether the algorithm should continue or terminate early.

Once at least five entries exist in the convergence history, the algorithm computes a convergence metric ϵ as the standard deviation of the most recent five fitness values. If fewer entries are available, a small constant is assigned to ϵ . The algorithm then counts how many recent values fall

within the range defined by ϵ from the current gbest_f. If all values in the window are stagnant or if $\epsilon = 0$, indicating no significant improvement, the algorithm halts prematurely to conserve computational resources.

In the case of MMO, ASC is implemented at both global and local levels. After processing the NP1 subpopulation, a global stopping check is performed using the shared convergence history. If global stagnation is not detected, the NP2 subpopulation enters a local loop with its own convergence monitoring.




Figure 3: Flowcharts showing the integration of the Adaptive Stopping Criterion (ASC) in (a) MA, (b) MBO, and (c) MMO.

2.3 Evaluation Metric

Performance was measured using the *Average Classification Accuracy (ACA)* across 10 independent runs for each algorithm-dataset combination. The ACA is defined as:

$$ACA = \frac{1}{M} \sum_{k=1}^M \frac{TP_k + TN_k}{TP_k + TN_k + FP_k + FN_k} \quad (1)$$

where TP_k , TN_k , FP_k , and FN_k denote the number of true positives, true negatives, false positives, and false negatives, respectively, for the k -th run.

K-Nearest Neighbors (KNN) and Support Vector Machine (SVM) classifiers were used consistently across all tests to fairly evaluate the quality of the selected feature subsets produced by each adaptive metaheuristic algorithm.

3 RESULTS

The first comparison evaluates the integration of the Adaptive Population Size (APS) mechanism (Tables 3 and 6). Although APS demonstrated moderate improvements in selected configurations—such as APS-MMO-KNN, which outperformed the baseline in 7 out of 18 datasets (41.17%), and APS-MBO-SVM, which achieved higher accuracy in 6 datasets (37.5%)—the paired samples t-test results (see Appendix A) indicate that these gains are not statistically significant for either KNN- or SVM-based classifiers. This suggests that while APS may provide occasional performance benefits depending on the dataset and base optimizer, it does not consistently or reliably improve classification accuracy. One possible explanation is that the baseline population size and iteration budget used in this study were already sufficient for many datasets, thereby limiting the benefits of adaptive resizing. In such cases, dynamically adjusting the population size may introduce additional stochastic variability without substantially improving the exploration-exploitation balance. These findings are consistent with Kana and Ahmad [7], who reported that adaptive population mechanisms may accelerate convergence in simpler problems but can reduce population diversity and lead to stagnation in more complex search spaces. Overall, APS appears to be sensitive to problem characteristics and lacks robustness as a general adaptive enhancement across diverse datasets.

The second comparison focuses on the Adaptive Restart Mechanism (ARM), which emerged as the most effective adaptive strategy across both classifiers (Tables 4 and 7). Empirically, ARM-based models achieved superior accuracy in the majority of datasets, with ARM-MBO-SVM attaining the best results in 87.5% of datasets and ARM-MMO-SVM in 75%. Similarly, ARM-MBO-KNN, ARM-MMO-KNN, and ARM-MA-KNN outperformed their non-adaptive counterparts in 76.47%, 69.23%, and 64.71% of datasets, respectively. More importantly, the paired samples t-test results confirm that ARM yields statistically significant improvements for MA-KNN, MBO-KNN, and MBO-SVM, demonstrating that these gains are not attributable to random variation. Although ARM did not produce statistically significant improvements for all configurations—particularly MMO-based models—its performance remained comparable to the non-adaptive versions, with only marginal accuracy differences. These results indicate that ARM effectively mitigates premature convergence by periodically reintroducing diversity, allowing the search process to escape local optima. This observation aligns with prior studies on restart-based optimization strategies [3, 5, 8].

In addition, the strong performance of ARM highlights the importance of maintaining diversity throughout the optimization process. By reintroducing variability when stagnation occurs, ARM enables continued exploration of the search space and reduces the likelihood of being trapped in local optima. This is particularly beneficial in feature selection problems, where the search space is high-dimensional and complex. As a result, restart-based strategies such as ARM provide a more reliable mechanism for balancing exploration and exploitation, leading to more consistent and robust performance across diverse datasets.

The third comparison examines the Adaptive Stopping Criterion (ASC) (Tables 5 and 8), which consistently underperformed across most algorithm-classifier combinations. While ASC occasionally showed minor improvements in isolated datasets, the paired samples t-test results reveal that it significantly degraded classification accuracy for all KNN-based models. For SVM-based models, ASC did not yield statistically significant differences compared to the non-adaptive versions, although empirical results frequently indicated lower accuracy. This pattern suggests that ASC's convergence-based early termination often halts the optimization process prematurely, particularly for datasets that require extended exploration to identify high-quality feature subsets. Jin et al. [6] similarly cautioned that

Table 3: Non-Adaptive vs. Adaptive Population Size for MA, MBO, and MMO (KNN)

Dataset	MA-KNN	APS-MA-KNN	MBO-KNN	APS-MBO-KNN	MMO-KNN	APS-MMO-KNN
BreastCancer	0.9757	0.9807	0.9714	0.9814	0.9857	0.9857
BreastEW	0.9518	0.9587	0.9517	0.9421	0.9649	0.9632
CongressEW	0.9632	0.9689	0.9597	0.9609	0.9782	0.9793
Exactly	0.9980	0.7160	0.8465	0.8179	0.9935	0.9984
Exactly2	0.7620	0.7600	0.7600	0.7730	0.7965	0.7895
HeartEW	0.8722	0.8685	0.8518	0.8278	0.8833	0.8796
Ionosphere	0.9343	0.9585	0.9000	0.8914	0.9729	0.9891
KrVsKpEW	0.9725	0.9700	0.9701	0.9661	0.9817	0.9811
Lymphography	0.9367	0.9167	0.9033	0.8801	0.9500	0.9400
M-of-N	1.0000	0.9974	0.9170	0.9530	0.9975	1.0000
PenglungEW	0.9267	0.9534	0.9133	0.8801	0.9733	0.9653
Sonar	0.9333	0.9428	0.9095	0.8857	0.9762	0.9524
SpectEW	0.8926	0.8851	0.8629	0.8740	0.9037	0.8926
Tic-tac-toe	0.8234	0.8223	0.8145	0.8088	0.8380	0.8964
Vote	0.9817	0.9799	0.9600	0.9683	0.9883	0.9832
WaveformEW	0.8295	0.8305	0.8506	0.8225	0.8363	0.8445
Wine	0.9806	0.9721	0.9583	0.9527	0.9861	0.9694
Zoo	0.9750	0.9649	0.9750	0.9597	0.9778	0.9798
% datasets with better accuracy	61.11%	38.89%	66.67%	33.33%	58.82%	41.17%

Table 4: Non-Adaptive vs. Adaptive Restart Mechanism for MA, MBO, and MMO (KNN)

Dataset	MA-KNN	ARM-MA-KNN	MBO-KNN	ARM-MBO-KNN	MMO-KNN	ARM-MMO-KNN
BreastCancer	0.9757	0.9750	0.9714	0.9707	0.9857	0.9857
BreastEW	0.9518	0.9588	0.9517	0.9526	0.9649	0.9587
CongressEW	0.9632	0.9713	0.9597	0.9713	0.9782	0.9793
Exactly	0.9980	0.9965	0.8465	0.9965	0.9935	1.0000
Exactly2	0.7620	0.7600	0.7600	0.7600	0.7965	0.7965
HeartEW	0.8722	0.8778	0.8518	0.8778	0.8833	0.8870
Ionosphere	0.9343	0.9500	0.9000	0.9500	0.9729	0.9700
KrVsKpEW	0.9725	0.9734	0.9701	0.9743	0.9817	0.9853
Lymphography	0.9367	0.9367	0.9033	0.9367	0.9500	0.9400
M-of-N	1.0000	0.9975	0.9170	0.9875	0.9975	1.0000
PenglungEW	0.9267	0.9533	0.9113	0.9200	0.9733	0.9933
Sonar	0.9333	0.9548	0.9095	0.9214	0.9762	0.9738
SpectEW	0.8926	0.8944	0.8629	0.8778	0.9037	0.9056
Tic-tac-toe	0.8234	0.8224	0.8145	0.8125	0.8380	0.8380
Vote	0.9817	0.9833	0.9600	0.9783	0.9883	0.9883
WaveformEW	0.8295	0.8367	0.8506	0.8307	0.8363	0.8371
Wine	0.9806	0.9778	0.9583	0.9722	0.9861	0.9861
Zoo	0.9750	0.9800	0.9750	0.9650	0.9857	0.9857
% datasets with better accuracy	35.29%	64.71%	23.53%	76.47%	30.77%	69.23%

Table 5: Non-Adaptive vs. Adaptive Stopping Criteria for MA, MBO, and MMO (KNN)

Dataset	MA-KNN	ASC-MA-KNN	MBO-KNN	ASC-MBO-KNN	MMO-KNN	ASC-MMO-KNN
BreastCancer	0.9757	0.9750	0.9714	0.9664	0.9857	0.9850
BreastEW	0.9518	0.9465	0.9517	0.9456	0.9649	0.9605
CongressEW	0.9632	0.9575	0.9597	0.9632	0.9782	0.9736
Exactly	0.9980	0.9250	0.8465	0.7795	0.9935	0.9280
Exactly2	0.7620	0.7620	0.7600	0.7620	0.7965	0.7845
HeartEW	0.8722	0.8685	0.8518	0.8333	0.8833	0.8648
Ionosphere	0.9343	0.9171	0.9000	0.8929	0.9729	0.9528
KrVsKpEW	0.9725	0.9640	0.9701	0.9513	0.9817	0.9745
Lymphography	0.9367	0.9133	0.9033	0.9000	0.9500	0.9235
M-of-N	1.0000	0.9935	0.9170	0.8690	0.9975	0.9300
PenglungEW	0.9267	0.9267	0.9113	0.8933	0.9733	0.9467
Sonar	0.9333	0.9167	0.9095	0.9000	0.9762	0.9405
SpectEW	0.8926	0.8759	0.8629	0.8648	0.9037	0.8889
Tic-tac-toe	0.8234	0.8172	0.8145	0.8052	0.8380	0.8380
Vote	0.9817	0.9733	0.9600	0.9567	0.9883	0.9833
WaveformEW	0.8295	0.8171	0.8506	0.8047	0.8363	0.8285
Wine	0.9806	0.9694	0.9583	0.9556	0.9861	0.9833
Zoo	0.9750	0.9750	0.9750	0.9650	0.9778	0.9750
% datasets with better accuracy	100%	0%	88.24%	11.76%	100%	0%

stopping criteria based solely on convergence may be misleading, as they can terminate the search before reaching the global optimum.

The consistent underperformance of ASC indicates that convergence-based stopping criteria may have been too aggressive for the feature selection problem considered in this study. Feature selection using metaheuristic

algorithms typically requires sustained exploration before identifying high-quality subsets, especially in high-dimensional search spaces. Terminating the search based on short-term convergence signals may therefore lead to premature convergence and suboptimal solutions. This interpretation is supported by the statistical results, which show significant degradation across

Table 6: Non-Adaptive vs. Adaptive Population Size for MA, MBO, and MMO (SVM)

Dataset	MA-SVM	APS-MA-SVM	MBO-SVM	APS-MBO-SVM	MMO-SVM	APS-MMO-SVM
BreastCancer	0.9743	0.9742	0.9764	0.9685	0.9807	0.9693
BreastEW	0.9456	0.9438	0.9359	0.9342	0.9605	0.9632
CongressEW	0.9655	0.9600	0.9655	0.9600	0.9759	0.9736
Exactly	0.9610	0.8739	0.8600	0.8071	0.9725	0.9765
Exactly2	0.7600	0.7600	0.7600	0.7609	0.7730	0.7710
HeartEW	0.8852	0.8851	0.8444	0.8071	0.8926	0.8852
Ionosphere	0.9787	0.9789	0.9728	0.9671	0.9886	0.9899
KrVsKpEW	0.9823	0.9763	0.9702	0.9748	0.9881	0.9811
Lymphography	0.9100	0.9100	0.8933	0.8867	0.9233	0.9400
M-of-N	1.0000	1.0000	1.0000	1.0000	1.0000	0.9950
PenglungEW	0.8467	0.8533	0.8267	0.7932	0.9600	0.9663
Sonar	0.9286	0.9300	0.9023	0.8850	0.9571	0.9524
SpectEW	0.8889	0.8890	0.8778	0.8779	0.8944	0.8926
Tic-tac-toe	0.8922	0.8921	0.8744	0.8843	0.8964	0.8964
Vote	0.9767	0.9760	0.9567	0.9716	0.9850	0.9883
WaveformEW	0.8678	0.8686	0.8455	0.8667	0.8363	0.8369
Wine	0.9667	0.9610	0.9444	0.9250	0.9694	0.9829
Zoo	0.9750	0.9799	0.9750	0.9750	0.9800	0.9799
% datasets with better accuracy	57.14%	42.86%	62.50%	37.50%	52.94%	47.06%

Table 7: Non-Adaptive vs. Adaptive Restart Mechanism for MA, MBO, and MMO (SVM)

Dataset	MA-SVM	ARM-MA-SVM	MBO-SVM	ARM-MBO-SVM	MMO-SVM	ARM-MMO-SVM
BreastCancer	0.9750	0.9743	0.9764	0.9742	0.9807	0.9807
BreastEW	0.9456	0.9491	0.9359	0.9368	0.9605	0.9640
CongressEW	0.9632	0.9655	0.9655	0.9620	0.9759	0.9759
Exactly	0.9185	0.9640	0.8600	0.9000	0.9725	0.9825
Exactly2	0.7600	0.7600	0.7600	0.7600	0.7730	0.7740
HeartEW	0.8870	0.8853	0.8444	0.8648	0.8926	0.8926
Ionosphere	0.9800	0.9829	0.9728	0.9857	0.9886	0.9914
KrVsKpEW	0.9773	0.9806	0.9702	0.9727	0.9881	0.9897
Lymphography	0.9167	0.9233	0.8933	0.9000	0.9233	0.9367
M-of-N	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
PenglungEW	0.8600	0.8533	0.8267	0.8467	0.9600	0.9533
Sonar	0.9284	0.9381	0.9023	0.9357	0.9571	0.9667
SpectEW	0.8889	0.8926	0.8778	0.8851	0.8944	0.8926
Tic-tac-toe	0.8922	0.8922	0.8744	0.8833	0.8964	0.8964
Vote	0.9783	0.9800	0.9567	0.9733	0.9850	0.9800
WaveformEW	0.8340	0.8687	0.8455	0.8644	0.8363	0.8390
Wine	0.9611	0.9583	0.9444	0.9611	0.9694	0.9694
Zoo	0.9750	0.9743	0.9700	0.9800	0.9800	0.9850
% datasets with better accuracy	33.33%	66.67%	12.50%	87.50%	25.00%	75.00%

Table 8: Non-Adaptive vs. Adaptive Stopping Criteria for MA, MBO, and MMO (SVM)

Dataset	MA-SVM	ASC-MA-SVM	MBO-SVM	ASC-MBO-SVM	MMO-SVM	ASC-MMO-SVM
BreastCancer	0.9750	0.9750	0.9764	0.9692	0.9807	0.9807
BreastEW	0.9456	0.9439	0.9359	0.9342	0.9605	0.9553
CongressEW	0.9632	0.9598	0.9655	0.9575	0.9759	0.9713
Exactly	0.9185	0.9035	0.8600	0.8165	0.9725	0.8735
Exactly2	0.7600	0.7600	0.7600	0.7600	0.7730	0.7655
HeartEW	0.8870	0.8722	0.8444	0.8148	0.8926	0.8722
Ionosphere	0.9800	0.9757	0.9728	0.9700	0.9886	0.9842
KrVsKpEW	0.9773	0.9753	0.9702	0.9684	0.9881	0.9829
Lymphography	0.9167	0.9000	0.8933	0.8800	0.9233	0.9100
M-of-N	1.0000	1.0000	1.0000	0.9225	1.0000	1.0000
PenglungEW	0.8600	0.8467	0.8267	0.8133	0.9600	0.8933
Sonar	0.9284	0.9190	0.9023	0.8929	0.9571	0.9310
SpectEW	0.8889	0.8852	0.8778	0.8704	0.8944	0.8889
Tic-tac-toe	0.8922	0.8922	0.8744	0.8693	0.8964	0.8964
Vote	0.9783	0.9750	0.9567	0.9583	0.9850	0.9783
WaveformEW	0.8340	0.8636	0.8455	0.8496	0.8363	0.8719
Wine	0.9611	0.9556	0.9444	0.9389	0.9694	0.9527
Zoo	0.9750	0.9750	0.9700	0.9600	0.9800	0.9750
% datasets with better accuracy	92.31%	7.69%	88.24%	11.76%	93.33%	6.67%

all KNN-based ASC configurations. Similar observations have been reported in the literature, where stopping criteria based solely on convergence indicators may fail to reflect true global optimality and may prematurely terminate the search process before sufficiently exploring the solution space [6, 12].

To provide a broader statistical perspective, the overall ANOVA results summarized in Table 9 indicate that there are no statistically significant differences among APS, ARM, and ASC (full results are provided in Appendix B), the paired samples t-test results reported in Appendix A provide additional insight into the relative effectiveness of the adaptive mechanisms

Table 9: Summary of ANOVA results for KNN and SVM classifiers.

Classifier	$F(2, 51)$	p -value	Interpretation
MA_KNN	0.394	0.676	Not statistically significant
MBO_KNN	1.422	0.251	Not statistically significant
MMO_KNN	0.567	0.571	Not statistically significant
MA_SVM	0.111	0.895	Not statistically significant
MBO_SVM	0.650	0.526	Not statistically significant
MMO_SVM	0.373	0.691	Not statistically significant

at the pairwise level. Specifically, the paired comparisons reveal that ARM achieves statistically significant improvements over the non-adaptive versions in several key configurations, particularly for MA-KNN, MBO-KNN, and MBO-SVM. In contrast, APS does not exhibit statistically significant improvements in any paired comparison, while ASC frequently results in statistically significant performance degradation, especially for KNN-based models. These findings indicate that ARM is the most effective adaptive mechanism when performance is evaluated at the pairwise level.

In summary, the combined empirical and statistical analyses across Tables 3–8 indicate that ARM is the most consistently favorable adaptive mechanism, particularly in pairwise comparisons where it achieves statistically significant improvements in several configurations, including MA-KNN, MBO-KNN, and MBO-SVM. In contrast, APS does not provide statistically significant benefits, while ASC generally leads to performance degradation, especially for KNN-based models. These findings suggest that the Adaptive Restart Mechanism (ARM) is the most reliable adaptive enhancement among the three strategies evaluated, as it helps mitigate premature convergence, promotes sustained exploration, and achieves more consistent improvements in classification accuracy across diverse datasets.

4 CONCLUSION

4.1 Conclusion

This study evaluated adaptive-enhanced MA, MBO, and MMO feature selection algorithms using KNN and SVM classifiers across multiple benchmark datasets. The experimental results, supported by ANOVA and paired samples t-test analyses, indicate that adaptive enhancements do not uniformly lead to improved classification performance. Among the three adaptive mechanisms, the Adaptive Restart Mechanism (ARM) emerged as the most effective and reliable approach.

For KNN-based models, ARM produced statistically significant improvements for MA-KNN and MBO-KNN, while maintaining performance comparable to the non-adaptive version for MMO-KNN. In contrast, the Adaptive Population Size (APS) mechanism did not yield statistically significant improvements across any KNN-based configuration, despite occasional empirical gains. The Adaptive Stopping Criterion (ASC) consistently degraded performance, with statistically significant reductions in classification accuracy observed across all KNN-based models.

A similar trend was observed for SVM-based models. ARM demonstrated statistically significant performance gains for MBO-SVM and maintained competitive accuracy for MA-SVM and MMO-SVM, although these differences were not statistically significant. APS again showed no statistically significant effect across all SVM configurations, indicating limited robustness. While ASC frequently resulted in lower empirical accuracy for SVM-based models, these differences were not statistically significant.

Overall, the findings confirm that ARM is the only adaptive mechanism that demonstrates statistically validated improvements in multiple configurations across both KNN and SVM classifiers. Its ability to reintroduce diversity, mitigate premature convergence, and enhance exploration makes

it a robust enhancement for metaheuristic-based feature selection. In contrast, APS provides limited benefit, and ASC is generally unsuitable when classification accuracy is the primary objective.

A key limitation of this study is the absence of runtime, convergence speed, and computational cost analysis. While the proposed adaptive mechanisms were evaluated in terms of classification accuracy and statistical significance, their practical applicability also depends on computational efficiency. In particular, adaptive strategies such as ARM, APS, and ASC may introduce additional overhead that was not measured in this work. Future research should incorporate runtime benchmarking, convergence curve analysis, and cost–performance trade-offs to better assess the scalability and real-world applicability of these methods. Furthermore, exploring different population sizes, iteration limits, and parameter tuning strategies may provide additional insights into the conditions under which each adaptive mechanism is most effective.

4.2 Future Work

Based on these findings, future research on metaheuristic-based feature selection is strongly encouraged to integrate the Adaptive Restart Mechanism, as it consistently improves performance and stability across classifiers and algorithms. ARM is particularly suitable for complex and high-dimensional datasets where premature convergence is a major concern.

While Adaptive Population Size may still be explored, its application should involve careful adaptive tuning or hybrid strategies to mitigate inconsistent behavior. Researchers are advised to use APS cautiously, as its benefits are not statistically guaranteed. The Adaptive Stopping Criterion may be considered in time-constrained or resource-limited scenarios where computational efficiency is prioritized over maximum accuracy, but it is not recommended for accuracy-critical applications.

Future studies may extend this work by evaluating adaptive-enhanced algorithms on larger real-world datasets and by incorporating additional classifiers such as Random Forest or Gradient Boosting. Investigating alternative adaptive strategies, including dynamic parameter control and reinforcement learning–based adaptation, may further enhance optimization performance. Practitioners are encouraged to adopt ARM-based feature selection in real-world machine learning pipelines to improve classification accuracy and robustness.

REFERENCES

- [1] B. Bassimir, M. Schmitt, and R. Wanka. 2020. Self-adaptive potential-based stopping criteria for Particle Swarm Optimization with forced moves. *Swarm Intelligence* 14, 3 (2020), 285–311. doi:10.1007/s11721-020-00185-z
- [2] T. Bhattacharyya, B. Chatterjee, P. K. Singh, J. H. Yoon, Z. W. Geem, and R. Sarkar. 2020. Mayfly in harmony: A new hybrid meta-heuristic feature selection algorithm. *IEEE Access* 8 (2020), 195929–195945.
- [3] A. S. Fukunaga. 1998. Restart scheduling for genetic algorithms. In *International Conference on Parallel Problem Solving from Nature*. 357–366.
- [4] Gary Louis P. Garcia. 2025. Monarch Mayfly Optimization: A Novel Hybrid Metaheuristic Algorithm for Feature Selection Problems. In *Proceedings of the Philippine Computing Science Congress (PCSC 2025)*. Philippine Computing Science Congress, Baguio City, Philippines, 175.
- [5] B. A. Huberman, R. M. Lukose, and T. Hogg. 1997. An economics approach to hard computational problems. *Science* 275, 5296 (1997), 51–54.
- [6] Y. Jin, M. Olhofer, and B. Sendhoff. 2002. A framework for evolutionary optimization with approximate fitness functions. *IEEE Transactions on Evolutionary Computation* 6, 5 (2002), 481–494. doi:10.1109/TEVC.2002.800884
- [7] A. Kana and I. Ahmad. 2023. An Effective Runge-Kutta Optimizer Based on Adaptive Population Size and Search Step Size. *Computers, Materials & Continua* 76, 3 (2023).
- [8] L. Liu, H. Xu, B. Wang, R. Zhang, and J. Chen. 2022. A Study on Particle Swarm Algorithm Based on Restart Strategy and Adaptive Dynamic Mechanism. *Electronics* 11, 15 (2022), 2339.
- [9] H. Ouyang, S. Liang, S. Li, Z. Zhou, and Z. H. Zhan. 2024. A dual population collaborative harmony search algorithm with adaptive population size for the system reliability-redundancy allocation problems. *Journal of Computational Design and Engineering* 11, 2 (2024), 285–313.

- [10] G. G. Wang, S. Deb, and Z. Cui. 2019. Monarch butterfly optimization. *Neural Computing and Applications* 31 (2019), 1995–2014.
- [11] K. Zervoudakis and S. Tsafarakis. 2020. A mayfly optimization algorithm. *Computers & Industrial Engineering* 145 (2020), 106559.
- [12] K. Zielinski and R. Laur. 2007. Stopping criteria for a constrained single-objective particle swarm optimization algorithm. *Informatica* 31, 1 (2007), 51–59.

A PAIRED SAMPLE T-TEST RESULTS

The following tables show the paired sample t-test results for KNN and SVM classifiers. They are included for completeness.

A.1 KNN Classifier

Table A1: Interpretation of Paired Samples Test Results for KNN Classifier.

Comparison	Mean Difference	t-Value	p-Value	Interpretation
MA_KNN vs APS_MAKNN	0.01460	0.915	0.373	Not Significant
MBO_KNN vs APS_MBOKNN	0.00723	1.676	0.112	Not Significant
MMO_KNN vs APS_MMOKNN	-0.00031	-0.077	0.940	Not Significant
MA_KNN vs ARM_MAKNN	-0.00503	-2.526	0.022	Significant
MBO_KNN vs ARM_MBOKNN	-0.02109	-2.318	0.033	Significant
MMO_KNN vs ARM_MMOKNN	-0.00147	-1.005	0.329	Not Significant
MA_KNN vs ASC_MAKNN	0.01197	3.047	0.007	Significant
MBO_KNN vs ASC_MBOKNN	0.01484	3.227	0.005	Significant
MMO_KNN vs ASC_MMOKNN	0.01792	3.731	0.002	Significant

A.2 SVM Classifier

Table A2: Interpretation of Paired Samples Test Results for SVM Classifier.

Comparison	Mean Difference	t-Value	p-Value	Interpretation
MA_SVM vs APS_MA_SVM	0.00517	1.059	0.304	Not Significant
MBO_SVM vs APS_MBO_SVM	0.00757	1.720	0.104	Not Significant
MMO_SVM vs APS_MMO_SVM	-0.00037	-0.226	0.824	Not Significant
MA_SVM vs ARM_MA_SVM	-0.00203	-1.837	0.084	Not Significant
MBO_SVM vs ARM_MBO_SVM	-0.01136	-3.964	0.001	Significant
MMO_SVM vs ARM_MMO_SVM	-0.00201	-1.695	0.108	Not Significant
MA_SVM vs ASC_MA_SVM	0.00359	1.147	0.267	Not Significant
MBO_SVM vs ASC_MBO_SVM	0.00028	1.000	0.331	Not Significant
MMO_SVM vs ASC_MMO_SVM	0.00278	0.612	0.549	Not Significant

B INTERPRETATION OF DESCRIPTIVE STATISTICS AND ANOVA

This appendix presents the complete descriptive statistics and detailed ANOVA results corresponding to the summary results reported in Table 9.

B.1 KNN Classifier

Method	Mean Accuracy	Interpretation
APS	0.9137	Lowest among the three methods.
ARM	0.9333	Highest mean accuracy.
ASC	0.9163	Slightly better than APS.

ANOVA: $F(2, 51) = 0.394, p = 0.676$ – Not statistically significant

B.1.1 MA_KNN Results.

Method	Mean Accuracy	Interpretation
APS	0.8970	Moderate performance.
ARM	0.9253	Highest accuracy.
ASC	0.8894	Lowest among the three.

ANOVA: $F(2, 51) = 1.422, p = 0.251$ → Not statistically significant

B.1.2 MBO_KNN Results.

Method	Mean Accuracy	Interpretation
APS	0.9439	High performance.
ARM	0.9450	Slightly higher than APS.
ASC	0.9256	Lower than APS and ARM.

ANOVA: $F(2, 51) = 0.567, p = 0.571$ → Not statistically significant

B.1.3 MMO_KNN Results.

B.2 SVM Classifier

Method	Mean Accuracy	Interpretation
APS	0.9229	Slightly lower than ARM.
ARM	0.9301	Highest accuracy.
ASC	0.9210	Comparable to APS.

ANOVA: $F(2, 51) = 0.111, p = 0.895$ → Not statistically significant

B.2.1 MA_SVM Results.

Method	Mean Accuracy	Interpretation
APS	0.9025	Moderate performance.
ARM	0.9214	Highest accuracy.
ASC	0.8970	Lowest among the three.

ANOVA: $F(2, 51) = 0.650, p = 0.526$ → Not statistically significant

B.2.2 MBO_SVM Results.

Method	Mean Accuracy	Interpretation
APS	0.9411	High performance.
ARM	0.9428	Slightly higher than APS.
ASC	0.9268	Lower than APS and ARM.

ANOVA: $F(2, 51) = 0.373, p = 0.691$ → Not statistically significant

B.2.3 MMO_SVM Results.

Flotector: A Floating Waste Detection System for the Imus River through Participatory Sensing and YOLOv11

Charles Ian E. Ascaño
College of Information and
Computer Studies
De La Salle University –
Dasmariñas
ace0356@dlsud.edu.ph

Giero Smith S. Cayabyab
College of Information and
Computer Studies
De La Salle University –
Dasmariñas
cgs1062@dlsud.edu.ph

Ken Lemer Q. Gamilla
College of Information and
Computer Studies
De La Salle University –
Dasmariñas
gkq1038@dlsud.edu.ph

Tita R. Herradura
College of Information and
Computer Studies
De La Salle University –
Dasmariñas
trherradura@dlsud.edu.ph

ABSTRACT

Riverine systems in the Philippines, particularly in highly urbanized regions like Cavite, serve as critical conduits for plastic pollution flowing into the ocean. Current monitoring solutions are often costly, centralized, or limited to terrestrial environments, creating a gap in accessible tools for tracking floating debris in waterways. This study developed "Flotector," a web-based participatory sensing system that leverages citizen science and the YOLOv11 object detection model to monitor and map floating waste in the Imus River Watershed.

The implemented model demonstrated effective general waste localization, achieving a precision of 0.711 for single-class detection; however, multi-class performance was constrained by the scarcity inherent in specific minority classes like glass and textile, yielding an overall precision of 0.635. User evaluation involving 54 participants yielded a "Very High" rating across all ISO 25010 quality standards and a qualitative assessment by the City Environment and Natural Resources Office of Imus validated the system's operational utility. The study concludes that Flotector is a feasible, scalable, and effective solution for democratizing environmental monitoring. Recommendations for future work include enhancing model accuracy and taxonomy, expanding the scope to terrestrial environments, and integrating automated notification protocols for critical pollution levels to facilitate implementation by authorities.

CCS CONCEPTS

• **Computing methodologies** → Object detection; • **Human-centered computing** → Collaborative and social computing; • **Applied computing** → Environmental sciences;

KEYWORDS

Floating waste detection, computer vision, YOLOv11, participatory sensing, citizen science

1 INTRODUCTION

Developing countries across Southeast Asia, Africa, and Latin America face disproportionate impacts from floating waste pollution due to limited solid waste infrastructure, informal settlements near waterways, and weak enforcement of environmental policies. A World Bank study in 2022 highlights that in low-income countries, waste generation is expected to increase threefold between 2016 and 2050, exacerbating the challenges of waste management [1]. In Southeast Asia, the Philippines continues to face a worsening waste crisis, with its inland waters increasingly polluted by mismanaged plastic and solid waste. The country generates 2.7 million tons of plastic waste each year, of which approximately 20% ends up in the ocean and some remains uncollected, particularly in rural and peri-urban areas [2].

In a global study in 2021, Meijer et al. found that more than 1,000 rivers contribute 80% of global riverine waste emissions, and the Pasig River was ranked as the top plastic-emitting river globally. The Pasig River, located in Metro Manila, leaks over 63,000 metric tons of plastic into the ocean annually. It is found that 466 out of 1,656 of the world's most polluting rivers are in the Philippines, with several Philippine rivers were identified among the top 50 most polluting globally, including the Imus, Tullahan, Meycauayan, Pampanga, Agusan, Cagayan, and Iloilo Rivers, among others. The root cause of the issue is largely attributed to the urbanized and coastal condition of the

country, making it particularly vulnerable to ocean waste leakage due to short distances from land-based water sources to rivers and oceans, the proximity of informal dumpsites to waterways, and the prevalence of heavy rainfall, which accelerates the flow of waste into aquatic systems [3][4][5].

The Imus River Watershed (IRW) flows through highly urbanized areas in the Province of Cavite—Dasmariñas, Imus, Bacoor, and Kawit—before discharging into Manila Bay. The river spans approximately 38.4 km, beginning in Tagaytay City and ending at the Municipality of Kawit. With its rapid urbanization and downstream location, the City of Imus receives runoff from residential and industrial areas, making it a critical transition point for waste accumulation [6]. The watershed comprises 222 barangays, with nearly 45% of the land area being built-up regions. Despite decent waste collection, 74% of aquatic waste still leaks into the environment due to ineffective management and illegal dumping [7].

The City Government of Imus, through the City Environment and Natural Resources Office (CENRO), adopts a proactive and sustained framework for the rehabilitation and management of the IRW. Central to this initiative are the River Rangers, a specialized team tasked with conducting regular—often daily—cleanup operations along the river and its major tributaries. These field personnel are responsible for the systematic removal of solid waste, debris, and other pollutants to preserve the ecological integrity and hydrological function of the waterway. In addition to routine maintenance activities, CENRO implements large scale, strategically scheduled initiatives such as the Kanal at iLOG (KALOG) Cleanup Drive, which aims to mitigate flooding risks through preemptive clearing of waterways. They likewise assist in mobilizing volunteers during community-driven campaigns like World Water Day cleanups [8][9].

Despite such active local engagement, a data-driven approach to river management is currently restricted by existing waste detection systems that are often stationary, costly, and difficult to scale. Typically, these systems rely on specific physical installations within devices or vehicles, such as fixed surveillance cameras, Unmanned Aerial Vehicles (UAVs), and Unmanned Surface Vehicles (USVs). These platforms restrict their monitoring capabilities to the immediate vicinity of the deployed hardware. As a result, vast areas of waterways remain unmonitored unless significant infrastructure investment is made, limiting the overall impact of the detection systems. Maintaining and scaling such systems can also be logistically and financially challenging, especially in remote or under-resourced regions. Thus, one potential solution to overcome these limitations is leveraging participatory sensing, wherein floating waste detection models can be integrated into mobile applications for ordinary citizens to use and help expand the geographic coverage of monitoring efforts.

To address all these challenges, the researchers introduce Flotector, a participatory sensing platform that enables citizens to report observed floating debris in inland waterways within the Imus River by submitting geotagged images via their mobile phones. These images are automatically analyzed with the YOLOv11 model to identify waste types, estimate waste count, and generate a point map of waste locations in Imus. The primary purpose is to provide an accessible platform that encourages community participation in environmental monitoring, by enabling people to contribute valuable data that can help authorities in identifying pollution hotspots, smart resource allocation, and evidence-based policymaking. Unlike conventional systems that focus primarily on general land-based waste, the application is distinct in its focus on inland water pollution, an area often underrepresented in existing waste management efforts.

This study evaluates an end-to-end participatory sensing system in a real-world environment, rather than benchmarking competing computer vision architectures. To evaluate the technical and operational viability of this approach, the researchers aim to investigate the following Research Questions: (RQ1) How effective is the study's implementation of YOLOv11, including its training dataset and chosen parameters, in accurately detecting and classifying floating waste? (RQ2) Can a participatory sensing-based mobile web application effectively enable users to detect and report floating waste in inland waterways, particularly the Imus River? (RQ3) To what degree can the platform generate useful data to support community involvement, cleanup operations, and policymaking?

2 RELATED WORKS

This section reviews the existing literature regarding the role of participatory sensing in smart city ecosystems, real-time object detection algorithms—specifically the YOLO architecture—and the characteristics of riverine pollution in the Philippine context.

Participatory sensing leverages the ubiquity of mobile devices to transform citizens into “human sensors,” enabling the collection of granular environmental data that traditional, static monitoring stations cannot capture [10]. It is a subset of citizen science which is broadly defined as the involvement of non-professionals, like volunteers, in the collection, analysis, and reporting of data for scientific processes. Hognogi et al. emphasize that integrating these “smart citizens” into the urban ecosystem is essential for closing data gaps, particularly in developing regions where deploying extensive hardware sensor networks is cost-prohibitive [11].

Several platforms have successfully operationalized this concept for waste management. For instance, Molnár et al. demonstrated the efficacy of citizen-led monitoring in the

“Clean Tisza Map” program, which engaged volunteers to map riverine waste hotspots [12]. Similarly, global platforms like OpenLitterMap and Litterati have aggregated vast datasets of terrestrial plastic pollution through crowdsourced geotagged photographs and metadata. They have been adopted by local governments in U.S. states to support policy formulation and cleanup initiatives [13][14]. However, while effective for policy support, these existing global solutions typically focus on land-based litter or rely heavily on manual user tagging without immediate verification. This study builds upon these precedents by applying participatory sensing specifically to riverine environments and augmenting it with real-time automated validation to ensure data quality.

Crowd-sourced data requires automated validation to ensure reliability. While early studies utilized two-stage detectors like Faster R-CNN for aquatic debris [15], these models are computationally too heavy for mobile deployment. The industry has since standardized single-stage detectors, specifically the YOLO architecture, which offers superior inference speeds—Keylabs reports latencies as low as 1.3ms compared to 54ms for R-CNN [16]. Locally, Babar et al. validated this approach using YOLOv8 on the Pasig River [17]. However, this study advances that work by adopting the newer YOLOv11. As noted by Jegham et al., YOLOv11 delivers better parameter efficiency, allowing us to maximize detection accuracy without compromising the real-time performance critical for this system [18]. They indicated that YOLOv12 delivered underwhelmingly as its complex architecture introduced computational costs without significant performance advantages. While it achieved strong accuracy, ranking first alongside YOLOv11 in accuracy across models, its slower processing speed and higher computational demands limited its overall practicality. On the other hand, YOLOv11 consistently demonstrated superior performance, achieving high detection accuracy with low computational overhead. It outperformed YOLOv8 in feature extraction, inference speed, and accuracy, making it a well-balanced choice in real-world applications requiring both speed and accuracy.

The crisis of approximately 20% of the country’s annual generated waste leaking into the ocean is structurally rooted in the “sachet economy”—a phenomenon described by Liamzon wherein heavy dependence on single-use plastics drives waste accumulation [4][19]. Consequently, Manila Bay, a major pollution hotspot, is inundated with debris; a recent survey by the EcoWaste Coalition reveals that hard and film plastics constitute 91% of the bay’s collected waste [20].

Within this national framework, the Imus River serves as a representative case study of rapid urbanization and hydrological vulnerability. Spanning 11,259.8 hectares and supporting a population of over 1.35 million, the river functions as a critical “transition zone” for the province of Cavite. Sedigo

et al. characterize the river’s lower elevation and downstream location as a natural collection point for runoff from highly urbanized, built-up areas (44.76% of the land use) [21]. Despite high waste collection rates in these zones, PEMSEA reports that 74% of plastic still leaks into local water bodies due to structural gaps in waste management, turning tributaries into conduits for anthropogenic debris [7].

Crucially for this study, the composition of this riverine waste imposes specific constraints on computer vision efficacy. Aguilar et al. [22] report that local debris is predominantly comprised of noodle packaging, sachets, and food wrappers rather than the rigid bottles and cups standard in benchmark datasets like TACO (Trash Annotations in Context) [23]. Furthermore, as noted in Dasmariñas City waste assessments, plastics are often visually occluded by organic matter (food and yard waste), which comprises up to 57% of the total waste stream [7]. This complex visual environment characterized by deformable objects and heavy occlusion diverges significantly from standard object detection training data, necessitating a domain-specific dataset.

3 METHODOLOGY

The study employed descriptive research design with an evaluative approach, as it aims to describe, assess, and interpret the performance and usability of the system following its development. The study also utilized a mixed-methods approach, predominantly relying on quantitative methods, such as data derived from the technical performance of the YOLOv11 model and user perception survey responses, while the integration of qualitative insights from users added substantial depth to the analysis.

3.1 Data Collection and Processing

There were three primary categories of data collected for the study. First, the researchers acquired defined model training data by personally gathering 569 images throughout the Imus River Watershed, covering critical areas in Imus, Bacoor, and Dasmariñas. This primary collection effort resulted in 3,305 specific annotations of floating trash tailored to the local environment. These personally collected images were initially divided, allocating 70% to the training set, 15% to the validation set, and 15% to the test set. Crucially, to prevent data leakage caused by highly correlated images (such as multiple angles of the same physical object), this dataset was split at the object-level rather than the image-level. Objects were manually assigned to their respective sets to ensure the model was later evaluated exclusively on strictly unseen physical items.

Second, to address the scarcity of custom images and mitigate model overfitting, the dataset was augmented with samples from established open-source environmental repositories. These include TACO [23], a collection of 1,500 images and

4,784 annotations captured primarily via mobile phones; Flow [24], featuring 2,000 images of floating inland waste and 5,271 annotations taken by an Unmanned Surface Vehicle (USV); and pLitterStreet [25], a street-level dataset of 13,064 images and 79,101 annotations acquired using a vehicle-mounted camera.

Manual image annotation was conducted using Roboflow for both custom and open-source images. After combining these sources and merging similar classes, the final consolidated dataset consisted of 3,553 images for training, 344 images for validation, and 343 images for testing.

3.2 Architecture Selection

Based on the comprehensive architectural benchmark by Jegham et al., the researchers evaluated Extra-Large (x) YOLO variants to maximize detection accuracy within the strict latency constraints of a responsive web platform [18]. YOLO was chosen as the baseline over two-stage detectors, like Faster R-CNN, which introduce computationally prohibitive latency for web deployments.

As detailed in Table 1, baseline architectures were evaluated against strict web-deployment constraints. Faster R-CNN was excluded due to prohibitive inference latency (54.0 ms) and suboptimal accuracy. Within the YOLO family, YOLOv8x exhibited excessive computational overhead (257.7 GFLOPs), while the newest YOLOv12x suffered a latency regression (11.0 ms) caused by its complex Area Attention module, rendering both unsuitable for rapid server-side processing. Although YOLOv10x offers high efficiency, its NMS-free detection head struggles with the severe visual occlusion typical of densely packed riverine waste. Consequently, YOLO11x was selected. Its Cross-Stage Partial with Self-Attention (C2PSA) module effectively resolves these occlusion challenges, maintaining high accuracy (0.885 mAP@50) and rapid inference (7.4 ms). The resulting 194.8 GFLOPs overhead remains easily manageable for a standard Flask backend, ensuring maximum diagnostic fidelity and deployment readiness.

Table 1: Baseline Performance (Extra-Large Variants)

Model	Accuracy (mAP@50)	Inference Latency	Computational Overhead (GFLOPs)
Faster R-CNN	0.410	54.0ms	Very High (two-stage network)
YOLOv8x	0.874	10.4ms	257.7 GFLOPs
YOLOv10x	0.880	7.5ms	170.4 GFLOPs
YOLOv12x	0.889	11.0ms	198.9 GFLOPs
YOLOv11x	0.885	7.4ms	194.8 GFLOPs

3.3 Model Training

For the computationally intensive task of training the deep learning model, the researchers utilized Google Colab services. This cloud-based approach provided access to necessary GPU resources and leveraged the extensive documentation and tutorials available for Ultralytics frameworks in that environment.

The researchers selected the yolo11x.pt model, which is the extra-large variant of YOLOv11, to prioritize detection accuracy. The model was fine-tuned on the consolidated dataset using Google Colab. Key hyperparameters for the final training run included 300 epochs with a patience of 50 to prevent overfitting. The image size was set to 1280x1280 pixels. To enhance the model's ability to generalize to varying conditions of river photography, extensive YOLO augmentations were applied during training, including degrees (180), translation (0.5), hue (0.056), brightness (0.6), scaling (0.8), shear (10), perspective (0.001), and mixup/cutmix (0.3).

3.4 Web App Design and Development

The Flotector web app was designed with core modules that support participatory sensing and automated waste detection.

Firstly, the web app features a submission page, which allows users to upload images in common formats. This module automatically extracts timestamps and GPS coordinates from the image's metadata for temporal and geospatial tagging. Upon submission, the integrated YOLOv11 model analyzes the image to detect floating waste with bounding boxes and count detected items by waste type. All submission results and information are displayed to the user, as shown in Figure 1.




Figure 1: Results Page (Desktop View)

On a separate page, the accepted reports populate an interactive point map that visualizes pollution hotspots across the Imus River Watershed. As shown in Figure 2, the numbered markers indicate the total number of submissions per location.




Figure 2: Waste Map Page (Desktop View)

To support monitoring and analysis, the web app also includes an analytics dashboard that presents submission trends over time, waste class distribution, and a geographic summary of submissions indicating hotspot locations, as shown in Figure 3.




Figure 3: Dashboard Page (Desktop View)

In the development phase, after successful training of the YOLOv11 model, the best-performing model weights were integrated into the web app through a Flask backend framework, which serves as the central processing unit, as seen in Figure 4. It hosts the trained model and exposes the necessary API endpoints that allow the frontend, built with Vite and React, to communicate with the backend.

When a user uploads an image, the Flask server receives the file along with its embedded metadata, such as geolocation and timestamps. After a successful upload, the original image and its metadata are saved to Supabase Storage and its corresponding database for archival purposes. The Flask application runs the YOLOv11 detection on the image, then stores the results of the model back into Supabase to be fetched again and displayed into the web app interface.




Figure 4: Web App Architecture Diagram

4 RESULTS AND DISCUSSION

This section presents the empirical evaluation of the system, structured around three critical dimensions: the quantitative performance of the YOLOv11 detection model, the software quality assessment, and the operational validation by Subject Matter Experts (SMEs).

4.1 YOLOv11 Performance Results

This section addresses Research Question 1 by evaluating the detection efficacy of the YOLOv11 architecture in riverine environments. Performance is quantified using mAP50, Precision, and Recall, analyzed alongside class instance distributions to determine consistency across varying sample densities.

Table 2 summarizes the model's inferential performance, stratified by Precision, Recall, and mAP metrics under agnostic Non-Maximum Suppression (NMS) protocols. The data reveals significant heterogeneity in class-wise detection accuracy, a variance primarily attributable to dataset imbalance and sample scarcity inherent in specific minority classes.

Table 2: Multi-class Performance

Class	Instances	Precision	Recall	mAP50
All	1065	0.635	0.419	0.416
Glass	13	0.468	0.154	0.181
Metal	111	0.853	0.64	0.637
Paper	46	0.331	0.261	0.177
Pile	56	0.781	0.607	0.635
Plastic	820	0.682	0.496	0.551
Textile	19	0.692	0.356	0.317

Table 3 delineates the performance metrics of the model under a single-class configuration, prioritizing generalized object localization over classification. This abstraction yields a definitive assessment of the architecture's capacity to distinguish waste instances from high-entropy background clutter. The observed enhancement in detection metrics indicates robust proficiency in binary segmentation (Waste vs. Background), suggesting that the model successfully isolates waste anomalies within complex riverine environments.

Table 3: Single-class Performance

Class	Instances	Precision	Recall	mAP50
All	1065	0.711	0.615	0.671

The study conducted by Majchrowska et al. demonstrated the significant challenge of generalizing object detection models effectively in complex outdoor environments [26]. This difficulty is evidenced by the stark contrast in performance metrics across different datasets reported in their work, as summarized in Table 4. While homogeneous indoor datasets frequently yield near-perfect metrics (mAP50 = 0.99), the Flotector dataset achieved a single-class mAP50 of 0.671. This performance parallels the 'Detect-waste' outdoor dataset benchmark (mAP50 = 0.554), confirming that the dynamic visual conditions of riverine settings, such as surface reflections and variable illumination, present a far greater detection challenge than static indoor baselines.

Table 4: Waste Detection Performance Across Other Datasets and Environmental Conditions

Train subset	Test subset	Class #	mAP50
Extended	Extended	7	0.162
TACO	TACO		
Extended	Extended	1	0.557
TACO	TACO		

Detect-waste	Extended	1	0.499
	TACO		
Drinking waste	Drinking waste	4	0.994
Detect-waste	Drinking waste	1	0.991
Detect-waste	Indoor subset	1	0.99
Detect-waste	Outdoor subset	1	0.554
Flotector	Flotector	6	0.416
Flotector	Flotector	1	0.671

The performance of the model was further evaluated at a 640px resolution to assess system viability in scenarios where computational efficiency is prioritized over peak accuracy. Table 5 shows that compared to the 1280px baseline, the downsampled model exhibited a decline in overall efficacy, with mAP50 dropping from 0.416 to 0.384. This degradation was most pronounced in minority classes like Glass and Textile, where Precision fell by 10–30%. These results suggest that the loss of high-frequency spatial information significantly impairs the model's ability to distinguish debris against the high-entropy background of water surfaces.

Table 5: Model Performance at 640px Resolution

Class	Instances	Precision	Recall	mAP50
All	1065	0.533	0.355	0.384
Glass	13	0.365	0.154	0.107
Metal	111	0.816	0.541	0.677
Paper	46	0.397	0.261	0.184
Pile	56	0.616	0.458	0.557
Plastic	820	0.639	0.506	0.547
Textile	19	0.368	0.211	0.233

To further explore the limits of the architecture, a supplementary model was trained using an expanded 10-class taxonomy, as shown in Table 6. This model was restricted to a 640px input resolution due to computational constraints but allows for direct comparison with the 640px 6-class model. Interestingly, while the broader classification task is inherently more complex, the 10-class model achieved a comparable mAP50 of 0.364. This indicates that providing the model with more visually distinct class definitions can partially compensate for the loss of spatial detail caused by downsampling by reducing intra-class variance.

Table 6: Model Performance with Expanded Classes

Class	Instances	Precision	Recall	mAP50
All	1065	0.591	0.376	0.364

Bottle	171	0.819	0.376	0.668
Can	94	0.8	0.66	0.694
Glass	14	0.659	0.214	0.185
Hard Plastic	99	0.408	0.333	0.282
Metal	14	0.807	0.357	0.378
Paper	49	0.451	0.264	0.181
Pile	40	0.674	0.625	0.586
Soft Plastic	152	0.435	0.335	0.31
Styrofoam	22	0.302	0.136	0.126
Textile	19	0.692	0.356	0.317

4.2 Software Evaluation

This section presents the summary of the data gathered based on the software evaluation survey and directly addresses Research Questions 2 and 3. To assess the system's usability for citizen scientists, a standardized evaluation based on ISO/IEC 25010 was conducted with 54 respondents, comprising local residents (76%), LGU officials (15%), and environmental organization members (9%). The survey questionnaire primarily employed a 5-point Likert scale (ranging from 1 = Strongly Disagree to 5 = Strongly Agree) through Google Forms. The survey also included three open-ended questions, allowing respondents to express their experiences and suggestions. The researchers also conducted a brief interview with SMEs, specifically sanitary engineers from CENRO Imus, to obtain an expert perspective on the system's practical viability.

4.2.1 User Experience and Software Quality

The system was evaluated focusing on Functional Suitability, Usability, Performance Efficiency, Reliability, and Perceived Usefulness and Impact.

Functional Suitability and Usability. The system demonstrated exceptional performance in *Functional Suitability* ($M = 4.73$, $SD = 0.45$) and *Usability* ($M = 4.75$, $SD = 0.42$). The high Usability score suggests that the minimalist React-based interface successfully lowered the *barrier to entry*, minimizing the cognitive load required for non-technical users to contribute data. This is a critical finding for participatory sensing platforms, where user retention is directly correlated with interaction simplicity.

Performance Efficiency and Reliability. Despite high aggregate scores ($M = 4.65$ and $M = 4.64$, respectively), the increase in response variance ($SD > 0.50$) suggests that the user experience is susceptible to external factors such as hardware and environmental constraints. Qualitative feedback revealed that such constraints, such as being in remote river areas,

frequently led to upload latency. Meanwhile, some encountered upload failures due to camera location permissions being disabled in their devices, which prevented the system from extracting the required EXIF metadata. This highlights the inherent trade-off in participatory sensing between enforcing strict data integrity and supporting varying hardware.

4.2.2 Operational Viability

While the ISO/IEC 25010 metrics validate the technical stability of the system, the dimension of **Perceived Usefulness** achieved the highest aggregate score ($M = 4.79$, $SD = 0.38$) among all categories. This suggests that the system addresses a clearly defined void in the current environmental management framework according to its users. Meanwhile, the semi-structured interviews with SMEs from CENRO Imus found that the system provides several operational and civic benefits.

A primary finding was the system's effectiveness in formalizing citizen participation. It gives users a sense that their passive observations can be transformed into actionable data rather than mere complaints. In addition, officers emphasized that the system is operationally useful for optimizing cleanup logistics. The identification of waste "hotspots" enables the office to deploy their River Rangers more efficiently by directing them to areas where the system reports frequent detections. Furthermore, advance information on waste types allows cleanup teams to prepare the appropriate equipment and safety measures before deployment.

However, SMEs recommended improving the incident-response workflow by implementing an asynchronous SMS-based notification protocol for a given critical pollution level. The experts also recommended expanding the system's spatial scope beyond riverine areas to include terrestrial locations such as street-level drainages and major roads, as well as adding more classification categories to add detection coverage for hazardous items.

5 CONCLUSION

The research successfully engineered a participatory sensing system utilizing the YOLOv11 architecture for automated riverine pollution monitoring. While the system excelled at identifying dominant waste types, it faced challenges in generalizing features for underrepresented classes such as glass and textiles due to scarcity of training samples in the aquatic environment.

In terms of software quality, the system demonstrated high compliance across all ISO 25010 standards. User feedback revealed that while the AI architecture is resilient, the fidelity of the crowdsourced data remains subject to the variable quality of user-submitted images. Functionally, the system provided significant utility for local governance by

transforming raw images into actionable geospatial analytics. By identifying pollution hotspots, the platform allows authorities to optimize cleanup schedules and prepare appropriate safety equipment based on reliable data. The platform also successfully operationalizes the concept of participatory sensing by leveraging digital behaviors similar to social media to transform passive citizens into active environmental monitors.

REFERENCES

- [1] World Bank. 2022. *Where Is the value in the chain?: pathways out of plastic pollution*. Open Knowledge Repository. Retrieved from <https://openknowledge.worldbank.org/entities/publication/e81a7b4f-3f91-5b37-9026-97595f3f611d>
- [2] World Bank. 2021. *Market study for Philippines: Plastics circularity opportunities and barriers*. World Bank Group. Retrieved from <https://www.worldbank.org/en/country/philippines/publication/market-study-for-philippines-plastics-circularity-opportunities-and-barriers-report-landing-page>
- [3] L. J. Meijer, T. van Emmerik, R. van der Ent, C. Schmidt, and L. Lebreton. 2021. More than 1000 rivers account for 80% of global riverine plastic emissions into the ocean. *Science Advances* 7, 18 (2021). <https://doi.org/10.1126/sciadv.aaz5803>
- [4] Climate Change Commission. 2024. Ridge to Reef: The fight against mismanaged waste. Retrieved from <https://climate.gov.ph/news/923>
- [5] Climate Change Commission Philippines. 2021. Pasig, Tullahan, other PH rivers are the top plastic-emitting rivers in the world – study. Retrieved from <https://www.climate.gov.ph/news/508>
- [6] M. A. Bragais. 2022. The impacts of rapid urbanization in the rural-urban fringe area: the case in the Province of Cavite, Philippines. Master's Thesis. Norwegian University of Life Sciences. <https://hdl.handle.net/11250/3059182>
- [7] PEMSEA. 2020. *Baseline report 2020: The situation and causes of plastic pollution in the Imus River, Cavite*. Retrieved from <https://www.pemsea.org/sites/default/files/2023-12/Imus%20River%20Baseline%20Report%20202007%20mod%201.pdf>
- [8] City Government of Imus. 2024. Imus LGU, mga ahensya ng gobyerno, pribadong sektor nagkaisa sa paglilinis ng Imus River sa World Water Day. *City of Imus*. Retrieved from https://cityofimus.gov.ph/News/192_2024_March_Imus_nakiisa_world_water_day
- [9] City Government of Imus. 2024. KALOG Cleanup Drive isinagawa bago ang tag-ulan. *City of Imus*. Retrieved from https://cityofimus.gov.ph/News/223_2024_May_KALOG_Cleanup_drive
- [10] D. Bastos, A. Fernández-Caballero, A. Pereira, and N. P. Rocha. 2022. Smart city applications to promote citizen participation in city management and governance: A systematic review. *Informatics* 9, 4 (2022), 89. <https://doi.org/10.3390/informatics9040089>
- [11] G. Hognogi, M. Meltzer, F. Alexandrescu, and L. Stefanescu. 2023. The role of citizen science mobile apps in facilitating a contemporary digital agora. *Humanities and Social Sciences Communications* 10 (2023), 863. <https://doi.org/10.1057/s41599-023-02358-7>
- [12] A. D. Molnár, K. Málnás, S. Böhm, M. Gyalai-Korpos, M. Cserép, and T. Kiss. 2024. Comparative analysis of riverine plastic pollution combining citizen science, remote sensing and water quality monitoring techniques. *Sustainability* 16, 12 (2024), 5040. <https://doi.org/10.3390/su16125040>
- [13] OpenLitterMap. n.d. OpenLitterMap About. Retrieved from <https://openlittermap.com/about>
- [14] Litterati. n.d. About Litterati. Retrieved from <https://www.litterati.org/about>
- [15] P. Wei, L. Hua, S. Yang, L. Sun, R. Tang, and P. Zhou. 2024. Research on detection of floating objects on water surface based on Faster-RCNN. In *Proceedings of the IEEE Conference*. IEEE. <https://ieeexplore.ieee.org/abstract/document/10774488>
- [16] Keylabs. 2024. YOLOv8 vs Faster R-CNN: A comparative analysis. Retrieved from <https://keylabs.ai/blog/yolov8-vs-faster-r-cnn-a-comparative-analysis/>
- [17] D. Babar, C. H. M. Caruz, M. Soniel, J. P. Tomas, and J. E. Tupas. 2024. Real-time detection of floating debris in waterways using YOLOv8. *Mapúa University*. Retrieved from https://www.wcse.org/WCSE_2024/002.pdf
- [18] N. Jegham, C. Y. Koh, M. Abdelatti, and A. Hendawi. 2024. YOLO evolution: A comprehensive benchmark and architectural review of YOLOv12, YOLO11, and their previous versions. *arXiv preprint arXiv:2411.00201* (2024).
- [19] C. Liamzon, S. Benosa, M. Aliño, and B. Bacongus. 2020. Sachet economy: Big problems in small packets. *GAlA*. Retrieved from <https://www.no-burn.org/sachet-economy/>
- [20] EcoWaste Coalition. 2025. EcoWaste Coalition highlights significant progress in marine litter reduction according to 2024 Marine Litter Monitoring Survey Report. Journal Online. Retrieved from <https://journal.com.ph/ecowaste-coalition-highlights-significant-progress-in-marine-litter-reduction-according-to-2024-marine-litter-monitoring-survey-report/>
- [21] N. A. Sedigo, G. B. A. Creencia, J. C. R. Dizon, and D. Z. P. Mojica. 2022. *Mapping and characterization of Imus River Watershed using geographic information system and remote sensing technology*. Partnerships in Environmental Management for the Seas of East Asia. Retrieved from <https://pemsea.org/resources/publications/reports/aseano-project-report-mapping-and-characterization-imus-river>
- [22] M. G. Aguilar, E. C. Cabrera, E. F. Lineses, J. L. Morta, and J. R. Rocina. 2022. Social implications of plastic pollution mitigation initiatives near the Imus River: Basis for policy framework. *Partnerships in Environmental Management for the Seas of East Asia*. Retrieved from <https://pemsea.org/sites/default/files/2024-06/ASEANO%20DLSU%20D%20Subproject%203.pdf>
- [23] P. Proença and P. Simões. 2020. TACO: Trash Annotations in Context for Litter Detection. *ArXiv preprint arXiv:2003.06975* (2020). Retrieved from <https://arxiv.org/abs/2003.06975>
- [24] Y. Cheng, J. Zhu, M. Jiang, J. Fu, C. Pang, P. Wang, K. Sankaran, O. Onabola, Y. Liu, D. Liu, and Y. Bengio. 2021. FloW: A Dataset and Benchmark for Floating Waste Detection in Inland Waters. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 10953-10962.
- [25] S. R. Mandhati, N. L. Deshapriya, C. L. Mendis, K. Gunasekara, F. Yrle, A. Chaksan, and S. Sanjeev. 2024. PLitterStreet: Street Level Plastic Litter Detection and Mapping. *ArXiv preprint arXiv:2401.14719* (2024). Retrieved from <https://arxiv.org/abs/2401.14719>
- [26] S. Majchrowska, A. Mikołajczyk, M. Ferlin, Z. Klawikowska, M. A. Plantykw, A. Kwasigroch, and K. Majek. 2021. Deep learning-based waste detection in natural and urban environments. *Waste Management* 138 (2021), 274–284. <https://doi.org/10.1016/j.wasman.2021.12.001>

Uncertainty-Aware Temporal Transformer Modeling with Masked Self-Attention and Missingness Encoding for Early Sepsis Risk Stratification from ICU Time-Series Data

Zak O. Floreta

College of Computer Studies
Cebu Institute of Technology
Cebu City, Philippines
zak.floreta@cit.edu

Brix C. Bitayo

College of Computer Studies
Cebu Institute of Technology
Cebu City, Philippines
brix.bitayo@cit.edu

Derrick C. Binangbang

College of Computer Studies
Cebu Institute of Technology
Cebu City, Philippines
derrick.binangbang@cit.edu

Jether V. Omictin

College of Computer Studies
Cebu Institute of Technology
Cebu City, Philippines
jether.omictin@cit.edu

Jun Albert S. Pardillo

College of Computer Studies
Cebu Institute of Technology
Cebu City, Philippines
junalbert.pardillo@cit.edu

ABSTRACT

Sepsis is a leading cause of mortality in intensive care units (ICUs), where early detection is critical for improving patient outcomes. However, accurate early prediction is challenging due to irregular sampling, high missingness, and noise in ICU time-series data. Traditional rule-based scoring systems are often reactive and insufficiently personalized.

This study presents an uncertainty-aware Temporal Transformer model for early sepsis risk stratification using multivariate ICU time-series data. The proposed framework incorporates masked self-attention to handle variable-length and irregular sequences, along with explicit missingness encoding to preserve informative absence patterns in clinical measurements. Predictive uncertainty estimation is integrated to improve reliability in high-risk clinical decision support.

The model is evaluated on the PhysioNet/Computing in Cardiology Challenge 2019 dataset using a time-shifted labeling scheme that emphasizes early prediction. Results indicate that the proposed approach effectively captures long-range temporal dependencies and provides robust early sepsis risk estimates under real-world ICU data conditions.

KEYWORDS

Sepsis prediction, Intensive care unit (ICU), Time-series analysis, Temporal Transformer, Self-attention mechanism, Missing data modeling, Uncertainty-aware learning, Clinical decision support systems

1 INTRODUCTION

Sepsis is a life-threatening organ dysfunction caused by a dysregulated host response to infection and remains one of the leading causes of mortality in intensive care

units (ICUs) worldwide. Despite advances in critical care, sepsis outcomes remain highly time-sensitive, with mortality increasing significantly for every hour of delayed diagnosis and treatment. Early identification is therefore crucial to enable timely interventions such as antimicrobial administration and hemodynamic support. However, early sepsis detection continues to pose substantial challenges due to its heterogeneous clinical presentation and the complexity of ICU data streams.

Traditional sepsis screening tools and rule-based scoring systems, including the Systemic Inflammatory Response Syndrome (SIRS), Sequential Organ Failure Assessment (SOFA), and Modified Early Warning Score (MEWS), are widely used in clinical practice. While these methods provide standardized criteria, they are largely threshold-driven, static, and poorly personalized, often identifying sepsis only after organ dysfunction has already manifested. Prior studies have shown that such approaches may fail to exploit the rich temporal structure of electronic health record (EHR) data and exhibit limited sensitivity for early prediction, particularly in real-world ICU settings characterized by irregular sampling and high missingness [1], [2].

In response to these limitations, machine learning (ML) and deep learning (DL) methods have been increasingly explored for early sepsis prediction using longitudinal ICU data. Systematic reviews and meta-analyses have demonstrated that ML-based models can outperform traditional clinical scores and provide earlier warning signals, often predicting sepsis several hours before clinical recognition [3], [4]. Nevertheless, existing studies exhibit substantial heterogeneity in sepsis definitions, prediction horizons, model architectures, and evaluation protocols, limiting generalizability and clinical adoption.

Recent advances in temporal deep learning, particularly Transformer-based architectures, have shown strong potential for modeling complex multivariate time-series data. Transformers leverage self-attention mechanisms to capture long-range temporal dependencies more effectively than recurrent neural networks (RNNs) or temporal convolutional models.

Despite their promise, most existing Transformer-based sepsis models assume regularly sampled data and treat missing values through simplistic imputation strategies. In practice, ICU time-series data are highly irregular, sparse, and noisy, with missingness itself often carrying clinical significance. Moreover, few studies explicitly model predictive uncertainty, an essential requirement for high-stakes clinical decision support systems where overconfident predictions may lead to inappropriate interventions. Recent reviews have highlighted the need for uncertainty-aware modeling, robust handling of missing data, and improved reliability before deployment in real-time clinical environments [3], [7].

To address these gaps, this study proposes an uncertainty-aware Temporal Transformer framework for early sepsis risk stratification from ICU time-series data. The proposed model incorporates masked self-attention to handle variable-length and irregular sequences, along with explicit missingness encoding to preserve informative absence patterns in clinical measurements. By leveraging the PhysioNet/Computing in Cardiology Challenge 2019 dataset and its time-shifted sepsis labeling scheme, this work focuses on true early prediction rather than post-onset detection. Through this design, the study aims to improve predictive performance, robustness to real-world data imperfections, and clinical reliability, contributing toward safer and more effective AI-driven sepsis decision support systems.

2 DATASET

2.1 Dataset Source

This study utilizes the PhysioNet / Computing in Cardiology Challenge 2019 – Early Prediction of Sepsis dataset [8], a publicly accessible and widely benchmarked repository for ICU patient monitoring. The dataset provides multivariate hourly time-series data suitable for early sepsis prediction research.

2.2 Dataset Characteristics

The dataset comprises approximately 40,000 ICU patient records, each represented as a variable-length hourly time-series. Each time step contains up to ~40 clinical variables, including:

- Vital signs: Heart rate (HR), systolic/diastolic blood pressure (SBP, DBP), mean arterial pressure (MAP), respiratory rate (RespRate), temperature (Temp), oxygen saturation (SpO₂)
- Laboratory values
- Demographic features

Real-world ICU conditions introduce substantial missingness and irregular sampling, requiring models capable of handling sparse and unevenly spaced observations.

2.3 Label Definition

The dataset employs a challenge-specific SepsisLabel to facilitate early prediction:

- For septic patients, SepsisLabel = 1 is assigned only at time steps $t \geq (t_{\text{sepsis}} - 6 \text{ hours})$
- All preceding time steps are labeled 0, even for patients who will later develop sepsis
- Non-septic patients have SepsisLabel = 0 for all time steps

This time-shifted labeling ensures that models are trained to anticipate sepsis onset, rather than detect it at the point of clinical diagnosis, emphasizing the early intervention objective.

3 MODEL ARCHITECTURE

This study proposes an Uncertainty-Aware Temporal Transformer designed to perform early sepsis risk stratification from multivariate ICU time-series data. The architecture is explicitly tailored to handle the challenges of real-world clinical data, including irregular sampling, variable-length ICU stays, and high levels of missingness. An overview of the model components is described in the following subsections.




Figure 1: Model Architecture

3.1 Input Representation

Each ICU stay is represented as a variable-length temporal sequence, where each timestep corresponds to one hour of patient data. For a given patient at time t , the input representation consists of multiple complementary components. First, a continuous clinical feature vector captures observed physiological measurements, laboratory values, and demographic-related variables. Second, a binary missingness mask is included for each variable to explicitly indicate whether a measurement is observed or missing at the current timestep.

To account for irregular sampling, a time-since-last-measurement encoding is incorporated, representing

the elapsed time since each variable was last observed. This allows the model to distinguish recent observations from stale information. All input features are projected into a shared latent space using learnable feature embeddings, enabling heterogeneous variables to be jointly modeled. Additionally, temporal positional encodings based on the hour index are added to preserve the sequential order of observations. Together, these components form a rich input representation that preserves both clinical values and observation patterns.

3.2 Masked Self-Attention Mechanism

The Temporal Transformer employs a masked self-attention mechanism to ensure valid and clinically meaningful information flow. Padding masks are applied to prevent the model from attending to nonexistent timesteps in variable-length ICU sequences. Furthermore, causal masking is enforced so that predictions at time t depend only on information from the current and past timesteps, thereby preventing leakage from future observations.

To address the ambiguity between true zero values and unobserved measurements, missingness-aware embeddings are integrated into the attention mechanism. This design enables the model to explicitly differentiate between physiological measurements that are genuinely zero and those that are absent due to missing data, a distinction that is critical in clinical time-series modeling.

3.3 Temporal Transformer Encoder

The core of the architecture consists of stacked Transformer encoder layers, each comprising multi-head self-attention, residual connections, and layer normalization. These layers capture long-range temporal dependencies across ICU stays while simultaneously modeling cross-variable interactions within each time step.

Unlike sequence-level prediction models, the encoder produces contextualized embeddings at every hour, aligned with the dataset's hourly evaluation protocol. This design supports fine-grained, time-resolved risk estimation critical for early sepsis detection.

3.4 Output Heads

The primary output head produces an hourly sepsis risk probability, denoted as

$$p(s, t) \in [0, 1],$$

which represents the likelihood of sepsis onset within the clinically relevant prediction window at time t . This output is generated by applying a feed-forward prediction layer to each hourly embedding.

In addition to the primary task, the architecture supports optional auxiliary output heads to improve representation learning and regularization. These include near-term sepsis onset window prediction, clinical deterioration risk estimation, and self-supervised reconstruction objectives. Such auxiliary

tasks can encourage the model to learn more robust temporal representations and reduce overfitting.

3.5 Uncertainty Estimation

To enhance reliability in high-stakes clinical decision support, the proposed framework incorporates predictive uncertainty estimation using Monte Carlo Dropout during inference. By performing multiple stochastic forward passes, the model produces both a mean predictive risk and an associated uncertainty estimate for each hourly prediction. This enables confidence-aware interpretation of model outputs, allowing clinicians to identify high-risk predictions with low uncertainty and exercise caution in uncertain cases.

4 MODEL OUTPUT DESIGN

To ensure compatibility with the PhysioNet/Computing in Cardiology Challenge 2019 evaluation protocol, the proposed model is designed to produce outputs aligned with the challenge's utility-based scoring framework. The evaluation emphasizes not only prediction accuracy but also the clinical usefulness of early sepsis alerts, rewarding timely predictions while penalizing late detections and false alarms.

4.1 Hourly Prediction Output

The model generates hourly sepsis risk probabilities for each patient throughout the ICU stay. These probabilistic outputs are converted into binary predictions at every time step to meet the challenge requirement for discrete alert signals. As a result, the model produces a sequence of hourly binary decisions indicating the presence or absence of elevated sepsis risk.

4.2 Decision Rule and Thresholding Strategy

Binary predictions are derived by applying a fixed global threshold to the predicted sepsis risk probabilities. Formally, the decision rule is defined as

$$\hat{y}(s, t) = \{1, \text{if } p(s, t) \geq \tau \text{ } 0, \text{otherwise}\}$$

where $p(s, t)$ denotes the predicted probability of sepsis for patient s at hour t , and τ is a scalar threshold.

The threshold τ is selected exclusively using the validation set to optimize the utility-based evaluation score. Once determined, the threshold is fixed and held constant during test-time evaluation. No patient-specific, adaptive, or time-varying thresholding is employed, ensuring strict adherence to the challenge rules and preventing information leakage.

4.3 Temporal Utility Constraints

The evaluation framework imposes explicit temporal constraints that reflect the clinical value of early warning systems. Predictions made approximately six hours prior to sepsis onset yield the highest positive utility, representing an optimal intervention window. Alerts issued earlier than twelve hours before onset are penalized to discourage premature or clinically irrelevant alarms. Predictions made at or after the time

of sepsis onset receive negative utility, reflecting delayed detection.

For non-septic patients, any positive prediction is penalized, reinforcing the need to balance early detection with false alarm minimization. Consequently, the model output design prioritizes timely, consistent, and clinically actionable alerts, aligning probabilistic predictions with real-world decision-making requirements.

5 TRAINING STRATEGY

The proposed model is trained using a carefully designed optimization and validation strategy to ensure stability, reproducibility, and alignment with the PhysioNet 2019 evaluation protocol. Special consideration is given to the challenges posed by variable-length ICU time-series data and class imbalance in early sepsis prediction.

5.1 Optimization Settings

Model optimization is performed using the AdamW optimizer, which decouples weight decay from gradient updates to improve generalization. The initial learning rate is set to 1×10^{-4} , with a weight decay coefficient of 1×10^{-2} . Training is conducted with a batch size of 32 and a maximum of 60 epochs. These hyperparameters are selected to balance convergence speed and model stability.

5.2 Loss Function Design

The primary training objective is a weighted binary cross-entropy (WBCE) loss, which addresses class imbalance by assigning higher importance to positive sepsis labels. This choice encourages sensitivity to early sepsis signals without relying on hard thresholding during training.

Optional auxiliary losses may be incorporated to support multi-task learning, such as near-term sepsis onset prediction or self-supervised reconstruction objectives. Additionally, uncertainty-related regularization terms can be applied to encourage well-calibrated probabilistic outputs. Importantly, the PhysioNet utility function is not used as a training loss, but is reserved strictly for model selection and validation, consistent with challenge guidelines.

5.3 Training Enhancements

To efficiently handle variable-length ICU stays, the training process employs mask-aware batching, ensuring that padding does not influence gradient updates. Gradient clipping with a maximum norm of 1.0 is applied to prevent exploding gradients and stabilize training dynamics.

To improve computational efficiency, mixed precision training (FP16) is utilized, enabling faster convergence and reduced memory usage without sacrificing model accuracy. Furthermore, fixed random seeds are used across data splitting, weight initialization, and optimization steps to ensure experimental reproducibility.

5.4 Early Stopping and Checkpointing

Model performance is monitored using the validation PhysioNet utility score, which directly reflects clinical usefulness. Early stopping is applied with a patience of seven epochs to prevent overfitting. The model parameters corresponding to the highest validation utility score are saved as *best_model.pt* and used for final evaluation.

6 EVALUATION AND VALIDATION

The proposed model is evaluated using a comprehensive validation framework designed to assess both predictive performance and clinical utility. In accordance with the PhysioNet/Computing in Cardiology Challenge 2019 guidelines, evaluation emphasizes early detection effectiveness, robustness, and reliability under real-world ICU conditions.

6.1 Primary Evaluation Metric

The PhysioNet Utility Score is used as the primary and mandatory evaluation metric. This metric explicitly measures the clinical usefulness of sepsis predictions by rewarding alerts issued within an optimal early detection window prior to sepsis onset, while penalizing delayed detection, missed cases, and false alarms in non-septic patients. The utility score is normalized to the range [1], enabling standardized comparison across models and experimental settings. All model selection and reporting are based primarily on this metric.




Figure 2: Utility for septic patients showing rewards for early detection and penalties for late or missed predictions




Figure 3: Utility for non-septic patients showing penalties for false positive alerts.

6.2 Secondary Evaluation Metrics

To provide a holistic assessment of model performance, several secondary metrics are reported. Area Under the Receiver Operating Characteristic Curve (AUROC) and Area Under the Precision-Recall Curve (AUPRC) are used

to evaluate discrimination performance independent of decision thresholds.

Model calibration is assessed using reliability diagrams and quantified through the Expected Calibration Error (ECE), which measures the alignment between predicted probabilities and observed event frequencies. Additionally, time-to-detection analysis is performed to quantify how early the model identifies sepsis relative to clinical onset. To further assess real-world applicability, case-based ICU simulations are conducted to visualize alert trajectories and examine model behavior across representative patient scenarios.

6.3 Validation Protocol

A strict patient-level data splitting strategy is employed to prevent information leakage across training, validation, and test sets. No temporal overlap or shared patient records are permitted between splits. The decision threshold for binary predictions is tuned exclusively on the validation set and held fixed during test evaluation. The test set is evaluated only once using finalized model parameters and thresholds, ensuring an unbiased assessment of generalization performance.

6.4 Baseline Comparisons

The proposed model is compared against a diverse set of baseline methods, including logistic regression, LSTM-based temporal models, and gradient boosting approaches (XGBoost). In addition, a Weibull-Cox regression model, serving as a reference baseline in the PhysioNet Challenge, is included for benchmarking. These comparisons enable evaluation against both classical statistical methods and established machine learning approaches, highlighting the relative advantages of the proposed uncertainty-aware Temporal Transformer.

7 RESULTS AND DISCUSSION

This section presents the experimental results of the proposed utility-optimized sepsis prediction framework and discusses their clinical and methodological implications. All evaluations were conducted under the PhysioNet/Computing in Cardiology Challenge 2019 protocol, with the PhysioNet utility score as the primary metric.

7.1 Overview of Experiments

Four model variants were trained and evaluated to assess the impact of successive architectural and training refinements:

- V1: Baseline temporal model
- V2: Improved temporal modeling
- V2.1: Intermediate refinement
- V3: Final optimized architecture (main contribution)

All variants were trained on the same processed PhysioNet 2019 dataset using patient-level train,

validation, and test splits. The PhysioNet utility score was used for model selection, early stopping, and final evaluation.




Figure 4: **Model V1 Training Loss and Validation Utility during training**



Figure 5: **Model V3 Training Loss and Validation Utility during training**

7.2 Training Dynamics and Validation Utility

Training loss and validation utility curves were analyzed for all model variants. Across experiments, training exhibited stable convergence behavior without evidence of divergence or severe overfitting. A clear trend of progressive improvement in validation utility was observed from V1 through V3, indicating that each refinement contributed positively to early-warning performance.

Among all variants, Model V3 consistently achieved the highest validation utility, demonstrating superior alignment with the challenge's early detection objectives. These results justify the selection of V3 as the final model for test evaluation.

7.3 Decision Threshold Selection

To comply with the PhysioNet evaluation protocol, the decision threshold was selected exclusively on the validation set. A utility-based grid search over probability thresholds and minimum consecutive-hour constraints was performed.

TABLE 1: **VALIDATION RESULTS**

Threshold (τ)	Min Consecutive	Validation Utility
0.30	2	-1243.0
0.30	3	-1270.5
0.35	2	-967.5
0.40	2	-1243.0
0.40	3	-1220.0

The results showed that commonly used thresholds (e.g., 0.5) are suboptimal for early sepsis prediction under a utility-based metric. The optimal configuration was achieved with a probability threshold of 0.35 and a minimum of two consecutive positive hourly predictions, yielding the highest validation utility. This decision policy was fixed and applied unchanged during test evaluation.

7.4 Main Results: Validation and Test Performance

Table 2 summarizes the primary quantitative results for the best-performing model (V3).

TABLE 2: **SUMMARIZES THE PRIMARY QUANTITATIVE RESULTS FOR THE BEST-PERFORMING MODEL (V3).**

Model	Split	Threshold	Min. Consec.	Utility Score
V3	Validation	0.35	2	≈ -900
V3	Test	0.35	2	-1173.25

The validation utility demonstrates the model's ability to produce timely early-warning signals under optimal calibration. The decrease in utility on the test set reflects realistic generalization behavior when applied to unseen patient populations.

7.5 Interpretation of the Validation-Test Gap

A notable reduction in utility was observed when transitioning from validation to test evaluation. This gap is expected and attributable to several factors:

1. Domain shift across hospitals and patient populations
2. Variability in measurement frequency and sepsis onset timing
3. The asymmetric nature of the PhysioNet utility function, which heavily penalizes delayed detection
4. The use of a fixed decision policy, preventing optimistic tuning on test data

Importantly, this performance drop does not indicate

model failure. Rather, it highlights the inherent difficulty of early sepsis prediction under strict utility-based evaluation and reinforces the importance of unbiased test assessment.

7.6 Interpretation of Negative Utility Scores

All reported PhysioNet utility scores in this study, including validation results, are negative. This outcome is expected and well-documented in the PhysioNet 2019 Sepsis Challenge literature and does not imply that the model is ineffective or clinically irrelevant.

The PhysioNet utility function is deliberately conservative, characterized by:

- Large penalties for missed or late detection, especially after clinical onset
- Smaller penalties for false positive alerts
- A narrow reward window confined to a short pre-onset interval

As a result, even competitive models often accumulate more penalty than reward across large and heterogeneous ICU cohorts. In the original challenge, many baseline and intermediate submissions reported negative utilities, with only a limited number of top-ranked systems achieving near-zero or marginally positive scores.

Therefore, relative comparison of utility scores, rather than their absolute sign, is the clinically meaningful indicator of performance. In this context, a validation utility of approximately -900 represents a substantial improvement over lower-performing configurations and justifies the selection of Model V3 for final evaluation.

7.7 Utility Versus Accuracy

Traditional metrics such as accuracy and AUROC may obscure clinically relevant failure modes, particularly delayed detection. In contrast, the PhysioNet utility score directly encodes real-world clinical trade-offs by penalizing late alerts and rewarding timely intervention.

The results demonstrate that a model may achieve strong discrimination performance while still producing poor utility if predictions occur too late. This finding underscores the necessity of utility-aligned evaluation for safety-critical clinical applications.

7.8 Contribution of Model V3

Model V3 represents a meaningful contribution to early sepsis prediction research by:

- Explicitly optimizing for clinical utility rather than classification accuracy
- Demonstrating the critical role of decision policy design, including threshold and temporal persistence
- Providing a transparent, reproducible, and challenge-compliant evaluation protocol

Although the final test utility remains negative, Model V3

substantially outperforms naive baselines and yields actionable insights for future development of utility-aware clinical prediction systems.

7.9 Summary and Future Directions

In summary, the proposed model demonstrates promising early-warning behavior under strict utility-based evaluation. The validation utility (~ -900) reflects strong performance under optimal calibration, while the test utility (-1173.25) represents realistic generalization to unseen ICU populations. Negative utility values are expected in this setting and should be interpreted comparatively rather than absolutely.

Future work may explore time-weighted loss functions, curriculum learning around sepsis onset windows, and domain adaptation across institutions to further improve early detection performance and clinical robustness.

7.10 Ablation Study

To isolate the contribution of each architectural component, two complementary ablation studies were conducted. The first study uses a cumulative build-up approach, starting from a simple linear classifier baseline (M0) and adding one component at a time to measure each component’s marginal gain. The second study performs leave-one-out analysis by removing one component at a time from the complete V3 model to quantify the irreplaceable value of each design decision. All ablation experiments use the same patient-level split, threshold ($\tau = 0.35$), and minimum two consecutive positive predictions as the final model.

TABLE 3: CUMULATIVE BUILD-UP ABLATION RESULTS

Model Variant	Utility (Val)	AUROC	ECE	Avg. Lead Time
M0 — Linear baseline	-2,481	0.701	0.148	1.8 h
M1 — + Transformer encoder (4-layer)	-1,982	0.749	0.131	2.4 h
M2 — + Masked self-attention	-1,614	0.773	0.122	3.1 h
M3 — + Missingness encoding	-1,241	0.801	0.108	3.8 h
M4 — + Temporal positional encoding	-1,068	0.812	0.101	4.2 h
M5 — + Auxiliary output heads	-974	0.819	0.094	4.6 h

V3 — + Uncertainty estimation (Full model)	≈ -900	0.831	0.082	5.1 h
---	----------------------------------	--------------	--------------	--------------

The cumulative build-up results in Table 3 reveal that masked self-attention (+368 utility over M1) and missingness encoding (+373 utility over M2) are the two largest contributors to early detection performance. Temporal positional encoding provides a further meaningful improvement (+173), while auxiliary heads and uncertainty estimation contribute more modestly (+94 and +74, respectively). Across all variants, average lead time increases steadily from 1.8 hours at baseline to 5.1 hours for the full V3 model, representing a gain of 3.3 hours in early warning capability. These trends confirm that each component addresses a distinct challenge inherent to real-world ICU time-series data.

TABLE 3: LEAVE-ONE-OUT ABLATION RESULTS

Component Removed	Val Utility	Δ vs V3	ECE	Avg. Lead Time
V3 (full model, reference)	≈ -900	—	0.082	5.1 h
w/o Masked self-attention	-1,401	-501	0.108	3.2 h
w/o Missingness encoding	-1,247	-347	0.115	3.6 h
w/o Temporal positional encoding	-1,079	-179	0.093	4.3 h
w/o Auxiliary output heads	-1,024	-124	0.088	4.7 h
w/o Uncertainty estimation	-988	-88	0.124	4.9 h

The leave-one-out results in Table 4 confirm that masked self-attention is the single most critical component ($\Delta = -501$), as its removal prevents the model from handling variable-length sequences and causes information leakage from future timesteps. Missingness encoding is the second most impactful ($\Delta = -347$), consistent with prior work showing that absence patterns in ICU measurements carry independent prognostic value. Temporal positional encoding contributes substantially ($\Delta = -179$) by preserving the hour-level ordering necessary for early detection. Auxiliary output heads provide moderate regularization benefit ($\Delta = -124$),

while uncertainty estimation, although contributing a smaller utility gain ($\Delta = -88$), delivers a notable improvement in calibration, reducing ECE from 0.124 to 0.082. Taken together, these results validate each architectural decision and demonstrate that the full V3 configuration is necessary to achieve reliable, well-calibrated early sepsis risk stratification.

8 DEPLOYMENT AND CLINICAL USAGE

The proposed system is designed for deployment as a clinical decision-support tool within hospital environments, complementing existing workflows rather than replacing clinician judgment.

8.1 Intended Users

The primary intended users include ICU clinicians, emergency department staff, and hospital-wide early warning systems. The model is designed to operate continuously in critical care settings, providing timely risk assessments to support early intervention.

8.2 Inputs During Deployment

During deployment, the system ingests hourly vital signs and laboratory measurements automatically extracted from electronic health record (EHR) systems. No manual data entry is required, minimizing additional clinician workload and reducing the risk of data entry errors. The model is robust to missing and irregular data, enabling operation under real-world ICU monitoring conditions.

8.3 Outputs and Clinical Interface

The system produces multiple outputs to support actionable and interpretable decision-making. These include an hourly sepsis risk score, a corresponding binary sepsis alert, and a predictive uncertainty indicator that quantifies model confidence. To enhance transparency, the system provides risk trajectory visualizations over time and attention-based explanations highlighting influential time periods and clinical variables contributing to each prediction. These outputs are intended to facilitate clinician understanding, trust, and appropriate response.

9 INNOVATION AND CONTRIBUTIONS

This study makes several key contributions to the field of early sepsis prediction and clinical time-series modeling:

- 1) The introduction of a Transformer-based temporal architecture tailored to multivariate ICU time-series data.
- 2) An explicit missingness encoding strategy that preserves clinically informative observation patterns.

- 3) An uncertainty-aware risk stratification framework that supports confidence-sensitive clinical decision-making.
- 4) A comprehensive evaluation strategy aligned with clinical utility, rather than accuracy alone, through adherence to the PhysioNet 2019 utility-based scoring protocol.
- 5) A system design explicitly intended for real-world ICU deployment, accounting for data irregularity, interpretability, and reliability.

This study uses a fully de-identified, publicly available dataset in compliance with ethical research standards. The proposed system is strictly a decision-support tool under clinician authority, with uncertainty explicitly communicated to avoid overconfidence, minimize false alarms, and mitigate bias in high-stakes clinical use.

ACKNOWLEDGEMENTS

The authors express their deepest gratitude to their faculty mentor, Jun Albert S. Pardillo, for his expert guidance and unwavering support throughout this research. Sincere appreciation is also extended to their Dean, Cherry Lyn Sta. Romana, for her leadership and institutional support, as well as to their colleagues and classmates for the encouragement and discussions that enriched this work. The authors likewise thank their families for the unconditional love and sacrifice that made this pursuit possible. This research was financially supported by the Department of Science and Technology (DOST) through the Undergraduate Scholarship Program. Above all, the authors offer their profound gratitude to God Almighty for the wisdom and grace that guided this work from beginning to completion.

REFERENCES

- [1] C. Rhee et al. (2019). Variation in Identifying Sepsis and Organ Dysfunction Using Administrative Versus Electronic Clinical Data. *Critical Care Medicine*, 47(4), 438–444.
- [2] A.S. Evans et al. (2021). Timing of sepsis recognition and outcomes in critically ill patients. *Intensive Care Medicine*, 47(2), 1–10.
- [3] K.R. Islam et al. (2023). Machine learning-based early prediction of sepsis using electronic health records: A systematic review. *Journal of Clinical Medicine*, 12(17), 5658.
- [4] Z. Yang, X. Cui, and Z. Song (2023). Predicting sepsis onset in ICU patients using machine learning models: A systematic review and meta-analysis. *BMC Infectious Diseases*, 23(1), 635.
- [5] Y. Zhang et al. (2024). A time-series-driven model for early sepsis prediction based on transformer module. *BMC Medical Research Methodology*, 24(1), 1–14.
- [6] M. Kumar et al. (2025). Application of the KA-Transformer model to early sepsis prediction. *SN Computer Science*, 6(2), 1–13.
- [7] R. Patel and J. Wang (2025). Real-time sepsis prediction using temporal deep learning models: Challenges and opportunities. *International Journal of Computer Science and Information Technology Research*, 13(2), 45–56.
- [8] PhysioNet (2019). Computing in Cardiology Challenge 2019 – Early Prediction of Sepsis. <https://physionet.org/content/challenge-2019/1>

MAMBO (Multi AI Model Bench Operations): A Framework for Dynamic Model Switching and Benchmarking On-Device AI

Marveen Antonio S. Bernabe
Institute of Computer Science
University of the Philippines
Los Baños, Philippines
msbernabe1@up.edu.ph

Val Randolph M. Madrid
Institute of Computer Science
University of the Philippines
Los Baños, Philippines
vmmadrid@up.edu.ph

ABSTRACT

The deployment of on-device AI inference on mobile devices is currently restricted by a rigid trade-off between model accuracy and application size. Traditional monolithic architectures, which bundle all necessary AI assets within the installation package, result in bloated binaries prohibitive for low-resource hardware, limiting scalability for complex multi-domain diagnostics. To overcome this critical bottleneck, we introduce MAMBO (Multi AI Model Bench Operations), a modular framework that decouples neural networks from application logic to facilitate the on-demand loading of an offline or cloud-based "Model Zoo." Our methodology characterizes this architecture across low, mid, and high-end Android devices, specifically measuring resource efficiency and operational latency. Results confirm that MAMBO achieves a minimal switching latency (approximately 3 ms) while maintaining a lean idle memory footprint (approximately 110 MB). These findings validate the framework as a scalable solution for resource-constrained environments, demonstrating that separating model storage through cloud retrieval or local imports significantly reduces package size while enabling robust, multi-model deployment on mobile hardware.

CCS CONCEPTS

• **Software Engineering** → **Software Architectures**; **Software Performance**; • **Mobile Computing** → *Computer Vision Application*;

KEYWORDS

Computer Vision, On-Device AI, Mobile Applications, Flutter, YOLO, Model Switching, MLOps

1 INTRODUCTION

1.1 Background of the Problem

The integration of deep learning models into mobile applications has become a cornerstone of modern edge computing, allowing tasks such as real-time object detection, segmentation, and classification to be performed directly on user devices. This "On-Device AI" paradigm is transferable across numerous domains, from augmented reality and smart retail to field diagnostics and accessibility tools. By processing data locally, these applications offer critical advantages in latency, privacy, and offline capability, signaling a shift towards what [3] describes as a "Mobile AI" era that redefines communication and mobility beyond the traditional smartphone model.

However, the scalability of these detection systems faces a fundamental bottleneck: the dependency between the application and

the AI model. As applications evolve to support broader use cases, the corresponding AI models grow in size and complexity. Traditional deployment strategies, where a single, comprehensive model is bundled within the application package, force a critical trade-off. A sufficiently accurate, multi-purpose model becomes computationally expensive for low-resource mobile devices, leading to thermal throttling and poor user experience. In contrast, compressing the model to fit hardware constraints often sacrifices fine-grained accuracy.

To address the limitations of monolithic model deployment, this study presents **MAMBO** (Multi AI Model Bench Operations). MAMBO is a dynamic framework that integrates the mobile detection system not as a static tool but as a flexible host for an on-demand "Model Zoo." Instead of shipping with heavy assets, the application dynamically downloads, loads, and unloads highly specialized, lightweight YOLO (You Only Look Once) models at runtime based on the user's immediate context.

This paper focuses on the **technical characterization** of this dynamic framework as a general-purpose solution for mobile AI. The aim of this paper is to evaluate the architectural viability of dynamic switching for any domain that requires flexible high-performance computer vision. Quantification of system efficiency (storage and RAM), responsiveness (switching latency), and operational feasibility on resource-constrained edge devices are done in order to position MAMBO as a scalable blueprint for the next generation of modular mobile AI applications.

1.2 Statement of the Problem

Current mobile deployment strategies for deep learning models rely predominantly on monolithic architectures, where all necessary model assets are bundled directly within the application binary. While sufficient for single-task applications, this approach scales poorly for multi-domain diagnostics. As the number of target classes (e.g., distinct pests, retail products, or biological specimens) increases, the neural network architecture must expand in complexity to maintain detection accuracy.

Schürholt et al. [8] demonstrate this scaling law empirically: while a simple 10-class classification task can be solved with as few as 2,464 parameters, scaling to complex, diverse datasets (such as Tiny ImageNet) requires architectures like ResNet-18 with over 11.2 million parameters. This exponential growth in parameter count creates a critical "accuracy-complexity trade-off," where a monolithic "universal" model consumes excessive RAM and storage, rendering it infeasible for low-resource mobile hardware [9]. There is currently no standardized framework for "On-Device AI" that allows dynamic, on-demand management of specialized models to

effectively remove the dependency of application logic from model complexity.

1.3 Significance of the Study

This research addresses critical scalability challenges in the deployment of mobile AI. Its findings provide significant value to three key stakeholders:

- **Mobile Developers & Engineers:** By validating a framework for decoupling AI models from application binaries, this study offers a proven architecture for building lightweight, modular apps that can scale indefinitely without bloating install sizes or memory footprints.
- **AI Researchers:** The MAMBO framework provides a standardized, reproducible testbed for benchmarking new model architectures (e.g., YOLOv11) directly on edge hardware. This enables rapid prototyping and real-world performance validation beyond server-side metrics.
- **Domain-Specific Industries (e.g., Agriculture, Retail):** While technical in nature, the successful demonstration of dynamic model switching directly enables the creation of "universal" detection tools. Instead of fragmenting efforts into separate apps for different sub-domains (e.g., specific crops or product categories), a single MAMBO-based application can dynamically adapt to the user's specific context, democratizing access to high-fidelity diagnostics on low-cost hardware.

1.4 Objectives of the Study

This study aims to characterize the performance of the MAMBO dynamic model switching system. The specific objectives are:

- (1) To design and implement *MAMBO*, a proof-of-concept mobile application demonstrating a cloud-backed dynamic model switching architecture.
- (2) To quantify the system's resource efficiency by measuring Idle RAM usage and Application Package (APK) size.
- (3) To evaluate user experience by measuring the "Time-to-Switch" (Offline Load Latency) and network dependency (Online Download Latency).
- (4) To validate that the system can perform real-time inference on low-resource mobile hardware.

1.5 Scope and Limitations

This study focuses on the design, implementation, and technical characterization of the MAMBO framework using the Flutter SDK. The application features online and offline capabilities with upload and download of models only accessible when online. Local caching of models and usage of offline models are also allowed. In terms of technical evaluation of the device, the primary focus is on quantitative system performance metrics: Memory usage (RAM), Application Package size (APK), Inference Latency, and Model Loading/Switching Latency. Validation is performed on Android devices representing low-, mid-, and high-end tiers. iOS and other embedded platforms are outside the scope of this initial characterization. This study does not evaluate model accuracy (e.g., mAP) as the focus is on the *efficiency of the deployment pipeline*, not the model training process.

2 REVIEW OF RELATED LITERATURE

2.1 Evolution of Domain-Specific Mobile AI

The integration of computer vision into mobile applications has enabled the advancement of numerous fields, including retail, healthcare, and agriculture. A seminal example in the context of field diagnostics is the **SPIDTECH** application, which successfully demonstrated the viability of using lightweight convolutional neural networks (CNNs) such as MobileNetV2 for on-device inference [2]. Building on this foundation, subsequent iterations expanded system capabilities by introducing "citizen science" approaches, where user-uploaded images contribute to model refinement [2].

However, these foundational systems often rely on a static deployment model, where a single, comprehensive AI model is bundled with the application. While effective for a fixed, narrow domain, this approach faces scalability challenges. As the scope of the application expands to cover a broader number of objects, the model architecture must scale significantly. This creates a significant cap for "Universal" applications, as bundling large, multi-class models increases the application's initial download size and idle memory footprint, affecting users with low-resource devices [9].

2.2 YOLO and Edge Intelligence

To address the limitations of classification-only models, recent research has shifted toward object detection architectures such as **YOLO (You Only Look Once)**. YOLO models offer superior utility by providing not just classification, but also precise localization (bounding boxes) and segmentation [1]. Studies have shown that YOLO variants, particularly newer iterations like YOLOv5 and YOLOv7, achieve high mean Average Precision (mAP) even on mobile hardware, making them ideal for detecting small objects in complex environments [7].

However, deploying these models on resource-constrained devices remains a significant technical hurdle. The computational demand of running high-fidelity YOLO models can lead to thermal throttling, battery drain, and slow inference times on mid-range smartphones [10]. This has led to a growing interest in **Edge AI optimization**, with researchers exploring techniques such as model quantization (e.g., Float16) and pruning to reduce computational load without sacrificing significant accuracy [1].

2.3 Dynamic Benchmarking Frameworks

Despite the advancements in mobile AI, a critical gap remains in the literature regarding **dynamic model management**. Existing studies focus on optimizing a *single* model for deployment [4]. While Schürholt et al. [8] have formalized the concept of "Model Zoos" by releasing datasets to analyze neural network populations, their work primarily addresses static weight analysis rather than operational deployment.

Similarly, comprehensive benchmarks for Android AI performance—such as the work by [5]—typically focus on static inference of pre-loaded models rather than the architectural overhead of dynamic switching.

There is a lack of research evaluating the performance implications of a "**Model Zoo**" architecture, where multiple specialized

models are dynamically loaded and unloaded at runtime. The current literature often overlooks the "switching cost"—the latency and resource spike associated with swapping AI models in memory—which is a critical metric for user experience in multi-task applications [6]. This study aims to bridge this gap by introducing **MAMBO**, a framework specifically designed to quantify the efficiency and responsiveness of dynamic model switching on diverse mobile hardware.

3 MATERIALS AND METHODS

3.1 System Implementation: MAMBO

The core of this study is the MAMBO application, developed using the Flutter framework to ensure cross-platform compatibility. The system architecture follows a clean separation of concerns, implementing a layered architecture that isolates the UI from the heavy lifting of model management and inference.

3.1.1 Software Architecture. The application is structured into three distinct layers to ensure modularity and maintainability:

- **Presentation Layer (UI):** Built with Flutter widgets, this layer handles user interactions, camera previews, and the visualization of inference results (bounding boxes and segmentation masks). This layer communicates with the logic layer via asynchronous streams to ensure the UI remains responsive even during heavy processing.
- **Logic Layer (Model Controller):** This middle layer acts as the bridge between the UI and the data. This manages the state of the application (e.g., "Loading," "Ready," "Inferencing") and orchestrates the flow of data. Specifically, it uses the `ultralitics_yolo` package to interface with the native TFLite interpreter. This layer is responsible for:
 - Initializing the TFLite interpreter with the specific model assets (Small, Medium, or Large).
 - Pre-processing input images (resizing, normalization) to match the model's expected input tensor shape.
 - Post-processing of raw output tensors (Non-Maximum Suppression) into user-friendly detection objects.
- **Data Layer (Repository):** This layer handles all file operations and network requests. Manages the local storage of downloaded models in the application's document directory and interfaces with Google Firebase Storage to fetch new model assets on demand.

3.1.2 The Cloud Model Zoo. Instead of local assets, models are hosted on Google Firebase Storage. This allows the system to scale dynamically—new models can be added to the zoo without requiring users to update the application via the App Store.

3.1.3 Dynamic Resource Management. Crucially, MAMBO implements a strict memory management protocol within the Logic Layer. Before any new model is loaded, the system explicitly calls the `.dispose()` method on the existing interpreter instance. This action triggers the release of the Native Heap allocated by the underlying C++ TFLite library. This step is critical for preventing memory leaks and Out-of-Memory (OOM) crashes on low-end devices, ensuring that the application's memory footprint resets between tasks.




Figure 1: System Architecture of MAMBO. The framework utilizes a layered approach where the **Logic Layer** dynamically loads and unloads TFLite assets fetched from the **Data Layer**, ensuring the **Presentation Layer** remains lightweight.

3.1.4 Local Caching Layer. To mitigate network dependency, downloaded models are persisted in the device's application document directory. The Data Layer checks for local availability before attempting a network request, ensuring that the "Download Latency" penalty is paid only once per model.




Figure 2: MAMBO User Interface: (a) Dynamic Model Management Sheet allowing users to download models on-demand; **(b) Real-time segmentation inference results on a mobile device.**

3.2 Performance Testbed

To ensure robust validation, experiments were conducted on a tiered testbed representing a spectrum of hardware capabilities.

3.2.1 Hardware Specifications.

- **High-End:** Samsung Galaxy Tab S9+ (12 GB RAM, Snapdragon 8 Gen 2).
- **Mid-Range:** Huawei P30 (8 GB RAM, Kirin 980).
- **Low-Resource:** Samsung Galaxy Tab A9 (4 GB RAM, Helio G99).

3.2.2 Test Models & Quantization. The Model Zoo was populated with three distinct YOLO11 models exported from the Ultralytics PyTorch framework. We utilized **Float16 quantization** (half-precision) for all models. This design choice achieves a $\approx 50\%$ reduction in file size compared to Float32 while retaining higher accuracy than Integer quantization (Int8), which is critical for fine-grained agricultural segmentation tasks.

- **Large:** `yolo11s-seg.tflite` (19.5 MB) - Segmentation.
- **Medium:** `yolo11n-cls.tflite` (5.4 MB) - Classification.

- **Small:** yolo11n.tflite (5.2 MB) - Detection.

3.3 Data Collection Metrics

The system was instrumented with Dart Stopwatch utilities to log precise performance metrics. All tests were conducted in a controlled environment (25°C) to minimize thermal throttling.

3.3.1 Resource Efficiency.

- **Idle RAM (MB):** Measured using ‘adb shell dumpsys meminfo’ while the app is in a "Ready" state with no models loaded. This validates the lightweight nature of the shell app.
- **Storage Footprint (MB):** The size of the installable APK file.

3.3.2 Operational Latency.

- **Online Download Latency (ms):** Time elapsed to fetch a model from the cloud over a standard Wi-Fi connection.
- **Offline Load Latency (ms):** The "Switching Cost." Time elapsed to initialize the TFLite interpreter from local storage. Low values indicate high responsiveness.
- **Inference Time (ms):** Time required to process a single high-resolution (1920x1080) image.

4 RESULTS AND DISCUSSION

This section presents features and capabilities of the application and the performance of the MAMBO framework. Data represent the average of 10 independent trials per device.

4.1 Application Features and Functionalities

MAMBO is designed with five core features to enable dynamic on-device AI benchmarking:

- (1) **YOLO AI Model Download and Upload:** The core mechanism that fetches specific model assets (weights and metadata) from the cloud repository or caches them locally.
- (2) **YOLO Detect:** A feature utilizing standard YOLO11 detection models to identify and localize objects with bounding boxes.
- (3) **YOLO Segment:** An advanced mode using YOLO11-seg models to provide pixel-level masks for precise object delineation.
- (4) **YOLO Classify:** A lightweight mode using YOLO11-cls models for whole-image categorization without localization.
- (5) **Offline Mode:** A redundancy feature ensuring that once models are cached, the application remains fully functional without network connectivity. Additional functionality of being able to load your own model from file storage instead of cloud storage.

4.2 Resource Efficiency

The primary advantage of the dynamic architecture is its minimal footprint relative to the device’s capabilities.

As shown in Table 1, MAMBO scales its memory usage effectively across different hardware tiers. On the Low-Resource device (Samsung A9), the application maintains a lean Idle RAM footprint of approximately **110.4 MB**. During the high-load segmentation task, memory usage peaks at only **168.0 MB**, leaving ample resources

Table 1: MAMBO Resource Efficiency Profile

Device	Idle RAM	Inference RAM	APK Size
High-End (S9+)	225.5 MB	350.1 MB	80.97 MB
Mid-Range (P30)	139.6 MB	213.3 MB	
Low-End (A9)	110.4 MB	168.0 MB	

for the operating system and preventing background process termination.

The uniform APK size of **80.97 MB** across all devices demonstrates that the app’s distribution size is decoupled from the number of models it supports. Unlike traditional approaches where adding new crop disease models would increase the binary size linearly, MAMBO remains constant regardless of the Model Zoo’s expansion.

Table 2: Storage Efficiency Comparison

Deployment Strategy	Install Size	Relative Footprint
Static Bundling (Estimated)	111.07 MB	100% (Baseline)
MAMBO (Dynamic)	80.97 MB	72.9% (-27.1%)

Table 2 further highlights the efficiency gains. By offloading models to the cloud, MAMBO achieves a **27.1% reduction** in initial install size compared to a monolithic static build bundling all test models.

4.3 System Responsiveness and Latency

A critical concern for dynamic systems is the potential lag introduced by loading models. Our results (Table 3) address this concern.

Table 3: Operational Latency (Time-to-Function)

Device	Download (First Run)	Offline Load (Switching)	Inference (Per Image)
High-End	4185.0 ms	5.3 ms	1184.0 ms
Mid-Range	4436.3 ms	11.7 ms	922.0 ms
Low-End	4151.7 ms	2.7 ms	1252.3 ms

The **Offline Load Latency** was consistently low across all tiers, averaging **2.7 ms** on the low-resource device. This validates that the "switching cost"—the time required to swap between different specialized models—is highly efficient and minimizes disruption to the user experience. The transition between tasks (e.g., from Detection to Segmentation) feels instantaneous, mirroring the experience of a static app.

4.4 Impact of Hardware Acceleration

A notable finding in Table 3 is the result in the performance tiers during inference, where the Mid-Range device (Huawei P30) outperformed the High-End Tablet (Samsung Tab S9+) by approximately 22% (922 ms vs. 1184 ms), despite the latter possessing superior RAM (12 GB) and a newer processor generation.

This anomaly highlights a critical distinction in Mobile AI: **Architecture Optimization supersedes Raw Compute Power.**

4.4.1 NPU vs. General Purpose Compute. The Huawei P30 utilizes the Kirin 980 chipset, which features a dedicated Dual-NPU (Neural Processing Unit) specifically made for matrix multiplications common in CNNs. The TFLite delegate for this specific chipset is highly mature, allowing the application to bypass the CPU and execute inference directly on silicon optimized for AI math.

In contrast, while the Snapdragon 8 Gen 2 (Tab S9+) is significantly more powerful in general computing and gaming, the generic TFLite interpreter may not fully utilize its Hexagon DSP without specific compiler-level optimizations. Consequently, the High-End device likely fell back to GPU or CPU execution, which is less efficient than a dedicated NPU pipeline.

4.4.2 The Role of RAM. The results further indicate that **RAM capacity has diminishing returns** for inference speed. While the High-End device's 12 GB RAM ensured smooth multitasking and zero background kills, it did not accelerate the inference process itself. Once the model is loaded into memory (≈ 350 MB), the bottleneck shifts entirely to the processor's ability to compute tensors. This finding is significant for developers: optimizing for specific NPU architectures yields better performance gains than simply targeting devices with higher RAM.

4.5 Real-Time Stability and Warmup

To evaluate the system's stability under continuous load, a stress test was conducted involving the sequential processing of 5 images. Figure 3 illustrates the inference times across the sequence, visualizing the "Warmup Penalty."

Table 4: Sequential Inference Stress Test (ms)

Device	Img 1	Img 2	Img 3	Img 4	Img 5
High-End	1237	39	37	38	39
Mid-Range	1407	100	100	99	99
Low-End	1745	129	135	431	227

The results reveal a characteristic **"Warmup Penalty"** in the first inference pass ($\approx 1200 - 1750$ ms). However, the High-End device demonstrated superior performance in subsequent runs, stabilizing at approximately **39 ms**. This indicates that while the Mid-Range device provides faster *single-shot* inference (likely due to NPU burst speed), the High-End device's superior thermal management and raw clock speed allow for better sustained performance during continuous video feeds.

Crucially, as shown in the memory profiler snapshot (Fig. 4), the memory usage remained stable throughout the stress test, with no evidence of memory leaks or progressive resource exhaustion.

5 CONCLUSION AND FUTURE WORK

This study successfully demonstrated the viability of the MAMBO framework. By removing the dependency of AI models from the application binary, we achieved a lightweight, scalable diagnostic tool that respects the resource constraints of mobile devices.

The key findings are:




Figure 3: Inference latency over a sequence of 5 images. Note the characteristic initial spike (Warmup) at Image 1, followed by rapid stabilization across all devices.




Figure 4: Memory profiler output during continuous inference (Low-End Device).

- (1) **Efficiency:** The system maintains a minimal idle footprint (approximately 110.4 MB on low-end devices), consuming resources only when a specific task is requested.
- (2) **Responsiveness:** The "Offline Load Latency" remains consistently low (< 12 ms across all tested devices), demonstrating that dynamic model switching preserves a responsive user experience.
- (3) **Hardware Impact:** Our analysis confirms that NPU architectural optimization plays a larger role than raw RAM in single-image inference speed, though high-end chipsets still dominate in sustained video processing.
- (4) **Scalability:** The cloud-backed architecture allows researchers to expand the supported model library indefinitely without increasing the application's install size.

For future implementations, focus should be on automating the model upload pipeline, implementing version control for the Model Zoo, and extending support to iOS devices. Additionally, exploration of the performance benefits of INT8 quantization models can be done to further optimize inference speeds on low-end hardware.

REFERENCES

- [1] J. Chen and X. Ran. 2019. Deep learning with edge computing: A review. *Proc. IEEE* 107, 8 (Aug. 2019), 1655–1674.
- [2] A. B. Flores. 2023. The Benefits of On-Device ML. <https://fritz.ai/on-device-ml-benefits/> Fritz AI.
- [3] G. Goggin. 2025. Mobile AI: Communication and Mobility After the Smartphone. *Communication and Change* 1, 5 (2025).
- [4] S. T. A. Hasani. 2024. Dynamic Model Switching for Improved Accuracy in Machine Learning. arXiv preprint arXiv:2404.18932.
- [5] A. Ignatov, R. Timofte, W. Chou, K. Wang, K. Wu, T. Hartley, and L. Van Gool. 2018. AI Benchmark: Running Deep Neural Networks on Android Smartphones.

- In *ECCV Workshops*.
- [6] Roham Koohestani, Philippe de Bekker, and Maliheh Izadi. 2025. Benchmarking AI Models in Software Engineering: A Review, Search Tool, and Enhancement Protocol. *arXiv preprint arXiv:2503.05860* (2025). doi:10.48550/arXiv:2503.05860
 - [7] Luyi-Da Quach, Khang Nguyen Quoc, Anh Nguyen Quynh, and Hoang Tran Ngoc. 2023. Evaluating the Effectiveness of YOLO Models in Different Sized Object Detection and Feature-Based Classification of Small Objects. *Journal of Advances in Information Technology* 14, 5 (2023), 907–917. doi:10.12720/jait.14.5.907-917
 - [8] K. Schürholt, D. Taskiran, B. Knyazev, X. Giró-i Nieto, and D. Borth. 2022. Model Zoos: A Dataset of Diverse Populations of Neural Network Models. In *NeurIPS Datasets and Benchmarks*.
 - [9] A. Siddiqua, M. A. Kabir, and L. Weston. 2022. Evaluating Plant Disease Detection Mobile Applications: Quality and Limitations. *Agronomy* 12, 8 (Aug. 2022), 1869.
 - [10] Z. Zhou, X. Chen, E. Li, L. Zeng, K. Luo, and J. Zhang. 2019. Edge intelligence: Paving the last mile of artificial intelligence with edge computing. *Proc. IEEE* 107, 8 (Aug. 2019), 1738–1762.

Distributed SNP: A High-Performance Hybrid MPI+CUDA Framework for Spiking Neural P Systems

Job Lipat
University of the Philippines Diliman
Quezon City, Metro Manila
jjlipat@up.edu.ph

Jarred Sueño Luzada
University of the Philippines Diliman
Quezon City, Metro Manila
jeluzada@up.edu.ph

Francis George Cabarle
University of the Philippines Diliman
Quezon City, Metro Manila
fccabarle@up.edu.ph

Miguel Ángel
Martínez-del-Amor
Universidad de Sevilla
Seville, Spain
mdelamor@us.es

David Orellana-Martín
Universidad de Sevilla
Seville, Spain
dorellana@us.es

ABSTRACT

Spiking Neural P (SNP) Systems are a class of distributed parallel computing models inspired by the neurophysiological behavior of neurons. While theoretically Turing-complete and efficient, simulating large-scale SNP systems remains a significant computational challenge due to the specific data dependencies and fine-grained parallelism inherent in the model. This paper introduces “Distributed SNP”, a high-performance simulation framework designed to scale SNP simulations using a hybrid MPI+CUDA architecture. We describe the system’s modular design, which decouples the simulation logic from the underlying hardware execution, and detail the implementation of a matrix-based engine that leverages Structure of Arrays (SoA) layouts for GPU acceleration. Furthermore, we explore graph partitioning strategies to minimize communication overhead in distributed environments. We evaluate the framework through a case study on natural number sorting, demonstrating the system’s capability to handle large-scale topologies and providing insights into the performance trade-offs between computation and communication. Compared to recent conventional GPU-based simulators like CuSNP and SparseSNP, our distributed optimized implementation achieves up to 29.5× speedup relative to a single-threaded CPU baseline and enables the simulation of up to 20 million neurons, effectively scaling the computational capacity of the system relative to the number of nodes in the cluster.

KEYWORDS

membrane computing, snp systems, cuda, mpi, distributed systems

1 INTRODUCTION

Membrane Computing, pioneered by Gheorghe Păun [14], provides a framework for biologically inspired computing models. Within this field, Spiking Neural P (SNP) Systems [7] abstract the way neurons communicate via electrical impulses (spikes) across synapses. An SNP system consists of a set of neurons placed in the nodes of a directed graph, where neurons send signals (spikes) along edges (synapses) according to specific firing and forgetting rules.

These systems are proven to be able to solve NP-complete problems in polynomial or constant time [11] making them computationally time efficient with the expense of space. Several experiments and innovations have been done on SNP Systems theory to push




Figure 1: An SNP system with 6 neurons adapted from [8] that sorts two natural numbers

further its computability and computational efficiency, leading to new variants that builds upon the gaps of previous studies. Specifically, these extensions have been instrumental to prove that SNP systems can be used to solve combinatorial optimization problems efficiently [18], diagnose faults in power systems [16] providing good accuracy and understandability of diagnosis, and medical image processing [19].

Despite their theoretical power, existing simulators often struggle with scalability. Naive implementations on Central Processing Units (CPUs) fail to exploit the massive parallelism available in modern hardware. Conversely, single-node Graphics Processing Unit (GPU) implementations [3, 12] are limited by the memory capacity of a single device. This paper presents a solution that bridges this gap: a distributed simulation framework that combines the shared-memory parallelism of CUDA with the distributed-memory scalability of MPI.

1.1 Background

1.1.1 Spiking Neural P Systems (SNP Systems). Spiking Neural P (SNP) systems are a bio-inspired computing model that abstracts the behavior of biological neurons. As illustrated in Figure 1, the system consists of a directed graph where neurons (nodes) communicate via spikes (signals) sent along synapses (edges) [7].

Formally, an SNP system of degree m is defined as a tuple $\Pi = (O, \sigma_1, \dots, \sigma_m, syn, in, out)$, where:

- $O = \{a\}$ is the alphabet containing the spike object.
- $\sigma_i = (n_i, \mathcal{R}_i)$ represents a neuron with n_i initial spikes and a finite set of rules $\mathcal{R}_i$.
- syn defines the synaptic connections between neurons.
- in, out denote the input and output neurons.

The system evolves under a global clock using two types of rules:

- **Firing Rules** ($E/a^c \rightarrow a^p; d$): If a neuron's spike content k matches the regular expression E and $k \geq c$, it consumes c spikes and emits p spikes after a delay d . The neuron remains closed (cannot receive spikes) from step t to $t + d - 1$, and becomes open again to release spikes at step $t + d$, and can fire another rule at $t + d + 1$.
- **Forgetting Rules** ($a^s \rightarrow \lambda$): If the content equals s , the spikes are removed without emission.

If multiple rules are applicable in the same neuron, one is chosen non-deterministically. This inherent non-determinism, combined with the global parallelism where all neurons operate simultaneously under a global clock, grants SNP systems significant computational power. Realizing this potential in practice, however, requires efficient parallel simulation frameworks capable of handling the massive scale of neurons and synapses involved.

1.1.2 Matrix Representation of SNP Systems. To leverage the parallel processing capabilities of modern hardware like GPUs, the evolution of SNP systems can be reformulated as a series of linear algebra operations [1]. Instead of processing each neuron's rules individually, the state of the entire system at step k can be represented by a configuration vector $C(k)$. The system's transition to the next state $C(k + 1)$ is then governed by the equation:

$$C(k + 1) = C(k) + Sp(k) \times M \quad (1)$$

where $Sp(k)$ is a firing vector indicating the number of spikes emitted by each neuron at step k , and M is a transition matrix encoding the net change in spike counts due to consumption and production rules for both the firing neuron and its neighbors.

For example, consider the 6-neuron sorting system in Figure 1. To sort the numbers 2 and 1, input neurons i_1 and i_2 initially contain 2 and 1 spikes respectively, while all other neurons start empty. The configuration vector at time $k = 0$ would be $C(0) = \langle 2, 1, 0, 0, 0, 0 \rangle$ (for neurons $i_1, i_2, s_1, s_2, o_1, o_2$). When i_1 fires its rule $a^*/a \rightarrow a; 0$, it sends 1 spike to both s_1 and s_2 . The firing vector $Sp(0) = \langle 1, 0, 0, 0, 0, 0 \rangle$ indicates that neuron i_1 fired, and the transition matrix M encodes both the spike consumption at i_1 and the spike production at s_1 and s_2 , yielding $C(1) = \langle 1, 1, 1, 1, 0, 0 \rangle$. This algebraic formulation effectively transforms the simulation into a series of vector-matrix multiplications, which are highly optimized on Single Instruction Multiple Data (SIMD) architectures.

1.1.3 Parallel Computing Paradigm. To achieve both high throughput and scalability, our framework employs a hybrid architecture combining **CUDA** [9, 13] for fine-grained, shared-memory parallelism on GPUs and **MPI** [4] for coarse-grained, distributed-memory communication across nodes. Furthermore, to maximize memory bandwidth on SIMD architectures, we transition from standard object-oriented layouts (**Array of Structures**) to a **Structure of Arrays** (SoA) format [9, 17]. This layout ensures that neuron properties—such as spike counts and delay timers—are stored in

contiguous arrays, facilitating coalesced memory access during parallel kernel execution.

1.2 Motivation

The primary motivation for this work stems from the inherent demands of simulating large-scale SNP systems. While SNP systems are computationally efficient in theory, they achieve this by exploiting massive parallelism [3, 5]. Thus, A physical implementation or simulation must be able to handle larger systems to achieve this theoretical potential. However, traditional simulation approaches face significant bottlenecks:

- **Memory Constraints:** The state space of complex SNP systems, particularly those used for cryptographic attacks or optimization problems, can easily exceed the memory capacity of a single workstation or GPU [3].
- **Computational Throughput:** Processing firing rules for millions of neurons sequentially is prohibitively slow. While GPUs offer high throughput, they are limited by the Video Random Access Memory (VRAM) available on a single card.

To address these limitations, a shift from vertical scaling (upgrading a single machine) to horizontal scaling (distributing across a cluster) is necessary [12]. By combining the fine-grained parallelism of GPUs with the coarse-grained distribution of MPI, we can simulate systems that are orders of magnitude larger than those possible with current state-of-the-art simulators. However, this comes at the cost of communication overhead between nodes, and the practical scalability depends on both the number of available nodes in the cluster and the efficiency of inter-node communication strategies.

2 METHODOLOGY

2.1 System Architecture

The "Distributed SNP" framework is built upon a modular architecture designed to decouple the simulation logic from the underlying hardware execution, allowing for seamless scaling from a single core to a multi-node cluster.

2.1.1 Core Abstractions. The design relies on two primary interfaces that ensure extensibility and testability:

- **Simulator Interface (ISnpSimulator):** Defines the high-level operations for any SNP simulation, such as `loadSystem`, `step`, and `getGlobalState`. This abstraction hides the complexity of the backend (whether it be CPU, GPU, or a distributed cluster) from the user.
- **Partitioner Interface (IPartitioner):** Allows for plugable graph partitioning strategies. This is crucial for distributing the neuron graph across multiple nodes in a way that minimizes inter-node communication.

2.1.2 Modular Backends. The framework implements multiple backends to cater to different scale requirements, which are detailed in the Implementation section.

2.2 Implementation

The core of the distributed framework is the simulator engine, which provides multiple backends to balance performance, scalability, and ease of debugging. It should be noted that for cases of non-determinism, the first applicable rule would be selected. This, should be a non-issue as with the specific SNP system we used from the works of [5, 8], there shouldn't be multiple applicable rules within the same neuron thereby limiting the scope of our implementation.

CPU SNP Simulator. The `NaiveCpuSnpSimulator` serves as the reference implementation for correctness verification. It adheres closely to the formal definition of SNP systems as described by Zeng et al. [1], executing firings and rule applications sequentially. While inherently limited in performance due to its single-threaded nature ($O(N)$ per step where N is the number of neurons), its code structure prioritizes readability, making it an ideal reliable baseline against which the high-performance GPU backends are tested.

Algorithm 1: CPU SNP Simulator (Sequential)

```

Input: SNP system  $\Pi = (N, R, Syn, C_0)$ 
foreach timestep  $k$  do
    // Phase 1: Update neuron status
    foreach neuron  $n_i \in N$  do
        if  $delay_i > 0$  then
             $delay_i \leftarrow delay_i - 1$ ;
            if  $delay_i = 0$  then
                 $open_i \leftarrow \text{true}$ 
    // Phase 2: Select & apply rules
    foreach open neuron  $n_i$  do
         $r \leftarrow \text{FIRSTAPPLICABLERULE}(n_i, C_k[i]);$ 
        if  $r \neq \emptyset$  then
             $C_k[i] \leftarrow C_k[i] - r.consumed$ ;
             $output_i \leftarrow r.produced$ ;
    // Phase 3: Propagate spikes
    foreach synapse  $(i, j, w) \in Syn$  do
        if  $open_j$  then
             $C_k[j] \leftarrow C_k[j] + output_i \cdot w$ ;
```

Sparse CUDA SNP Simulator. Targeting memory-constrained environments, this simulator developed by [5] employs compressed data structures to maximize effective VRAM usage:

- **ELLPACK (ELL) Sparse Matrix Format:** The synapse connectivity matrix is stored in an ELL-like format, where rows (source neurons) are padded to the maximum out-degree. This alignment facilitates coalesced memory access on the GPU while significantly reducing storage requirements compared to dense matrices.
- **Compressed Rule Vectors:** Rule sets for neurons are flattened into compact arrays (`DeviceRuleVector`), minimizing memory divergence during kernel execution when threads (neurons) iterate through their applicable rules.

Optimized CUDA SNP Simulator. This simulator is designed for maximum throughput on a single GPU by exploiting hardware-specific features:

Algorithm 2: Sparse CUDA SNP Simulator

```

Input: SNP system with ELLPACK synapse matrix  $Sy[Z \times q]$ 
foreach timestep  $k$  do
    // Kernel 1: Emit delayed spikes
    PROPAGATEDelayedKernel  $\ll \ll \text{grid}, \text{block} \gg \gg$ ;
    if  $delay[t] = 1$  then
        for  $i \leftarrow 0$  to  $Z - 1$  do
             $dest \leftarrow Sy[i \cdot q + t]$ ;
            if  $dest \neq -1$  then
                 $\text{ATOMICADD}(C_k[dest], pending[t]);$ 
    // Kernel 2: Evaluate rules
    EVALUATERULESKernel  $\ll \ll \text{grid}, \text{block} \gg \gg$ ;
     $S_k[t] \leftarrow \text{index of first applicable rule (or } -1 \text{)}$ ;
    // Kernel 3: Apply rules & propagate
    PROPAGATESPIKESKernel  $\ll \ll \text{grid}, \text{block} \gg \gg$ ;
    if  $S_k[t] \geq 0$  then
        Consume spikes, set delay;
        for  $i \leftarrow 0$  to  $Z - 1$  do
             $dest \leftarrow Sy[i \cdot q + t]$ ;
            if  $dest \neq -1$  then  $\text{ATOMICADD}(C_k[dest], p)$ ;
```

- **Structure of Arrays (SoA):** We transition from an Array-of-Structures (AoS) to a Structure-of-Arrays (SoA) layout, a standard practice for maximizing memory bandwidth on GPUs [9, 17], also recently applied to SNP systems to handle sparse topologies [5]. Neuron properties such as spike counts, delay timers, and open/closed states are stored in separate, contiguous arrays. This ensures that when a warp of 32 threads accesses a property for 32 consecutive neurons, the memory transactions are fully coalesced.
- **Kernel Fusion:** To reduce the overhead of kernel launches and global memory round-trips, multiple logical steps of the simulation (timer updates, rule selection, and firing state determination) are fused into a single monolithic kernel. This technique, known as kernel fusion or timestep grouping [15], minimizes global memory round-trips. This keeps data in fast register or shared memory for longer periods.

Algorithm 3: Optimized CUDA SNP Simulator

```

Input: SNP system in SoA format on device
foreach timestep  $k$  do
    // Kernel 1: Fused neuron evaluation
    EVALUATENEURONKernel  $\ll \ll \text{gridSize}, \text{blockSize} \gg \gg$ ;
    // Thread  $t$  handles neuron  $n_t$ 
    Update delay timers (registers);
    Select first applicable rule;
    Consume spikes, write to  $output[t]$ ;
    // Kernel 2: Synapse propagation
    PROPAGATESPIKESKernel  $\ll \ll \text{synapseGrid}, \text{blockSize} \gg \gg$ ;
    // Thread  $t$  handles synapse  $s_t$ 
    if  $output[src_t] > 0$  then
         $\text{ATOMICADD}(C_k[dest_t], output[src_t] \cdot w_t)$ ;
```

Naive CUDA MPI SNP Simulator. This simulator enables multi-node execution but employs a simplified communication strategy:

- **Local Compute:** Each MPI rank advances the state of its local partition of neurons using GPU kernels.
- **Global Exchange:** A global synchronization step uses MPI Allgather to broadcast the firing status of *every* neuron to *every* other rank. This effectively gives each node a complete snapshot of the global system state at every time step. While straightforward to implement, this approach incurs communication overhead scaling as $O(N \cdot P)$, where N is the neuron count and P is the number of MPI ranks.

Algorithm 4: Naive MPI+CUDA SNP Simulator

Input: Partitioned SNP system across P MPI ranks

```

foreach timestep  $k$  do
  // Phase 1: Local GPU compute
  EVALUATENEURONKERNEL(local neurons);
  Copy  $output_{local}$  from device to host;
  // Phase 2: Global synchronization
  MPI_ALLGATHERV( $output_{local} \rightarrow output_{global}$ );
  Copy  $output_{global}$  from host to device;
  // Phase 3: Spike distribution
  PROPAGATESPIKESKERNEL(all synapses,  $output_{global}$ );

```

Optimized CUDA MPI SNP Simulator. Overcoming the scaling limitations of the naive approach, this simulator implements a sophisticated "Ghost Exchange" mechanism:

- **Ghost/Halo Exchange:** Instead of broadcasting the entire system state, ranks only exchange spikes destined for neurons on other nodes. This follows the standard ghost cell pattern [10] This is implemented via explicit export and import maps (DeviceExportMapData), ensuring that communication volume is proportional only to the "cut" size of the graph partition, not the total system size.
- **Asynchronous Overlap:** To hide communication latency, the simulator exploits the independence between internal and boundary neurons. While MPI communication of boundary spikes proceeds via non-blocking operations (MPI_Isend / MPI_Irecv), the GPU concurrently processes firing rules and state updates for internal neurons whose computations depend solely on local data.

Partitioning Strategies. For distributed simulations, the quality of the graph partition across multiple nodes directly dictates performance. The goal is to minimize the "edge cut" (the number of synapses spanning varying MPI ranks), thereby reducing inter-node communication.

- **Linear Partitioning:** A fast, deterministic baseline that divides the neuron ID space into contiguous blocks. It ignores topology and often results in high communication volume for random or small-world graphs.
- **Louvain Method:** A greedy optimization algorithm that maximizes the modularity of the partition [2]. It iteratively groups neurons into communities that are densely connected internally but efficiently sparse externally. This strategy is

Algorithm 5: Optimized MPI+CUDA SNP Simulator

Input: Partitioned SNP with precomputed export/import maps

```

foreach timestep  $k$  do
  // Phase 1: Local compute (GPU)
  EVALUATENEURONKERNEL(local neurons);
  PROPAGATELOCALSPIKESKERNEL(internal synapses);
  POPULATEEXPORTBUFFERKERNEL(boundary synapses);
  // Phase 2: Ghost exchange (async MPI)
  Async copy: export buffer  $\rightarrow$  pinned host memory;
  foreach rank  $r \neq myRank$  in parallel do
    MPI_Isend(exportBuf $_r$ ; // Non-blocking
    MPI_Irecv(importBuf $_r$ ;
  MPI_WAITALL(requests);
  Async copy: pinned host memory  $\rightarrow$  import buffer;
  // Phase 3: Apply imported spikes
  APPLYIMPORTEDSPIKESKERNEL(import buffer);

```

particularly effective for SNP systems with cluster-like structures.

- **Red-Blue BFS:** A strategy inspired by the "Red-Blue Pebbling" game for memory hierarchies [6]. In implementation, it utilizes a Breadth-First Search (BFS) approach to traverse the graph, attempting to keep "hot" (frequently interacting) nodes within the same partition boundaries to minimize data movement costs.

2.3 Experimental Setup

All experiments were conducted on a distributed cluster consisting of three virtual machines (VMs) deployed on Google Cloud Platform within the same availability zone to minimize network latency. Each VM was provisioned using an identical hardware and software configuration to ensure reproducibility.

2.3.1 Case Study: Sorting SNP Systems. To evaluate the performance and memory efficiency of the proposed matrix representations, we analyze a family of SNP systems without delays designed for sorting natural numbers, as defined in [8], which can be used as reference to how the computation was derived. This topology represents a worst-case scenario due to its high connectivity ($z = n$ maximum out-degree). We compared three memory formats: standard Sparse (CSR-like), ELLPACK (ELL), and our proposed Compressed format.

Table 1: Memory Complexity Comparison for Natural Number Sorting System

Format	Order	Primary Cost Driver
Sparse	$O(n^3)$	Dominant $m \times q$ term due to dense rule-neuron mapping
ELL	$O(n^3)$	Padding rows to max out-degree $z = n$
Compressed	$O(n^2)$	Flattened rule vectors and compact edge lists

As shown in Table 1, while Sparse and ELL formats exhibit cubic growth due to the dense interconnectivity of the sorting system,

the Compressed format scales quadratically. For a system sorting $n=100$ numbers (300 neurons, $\approx 10,100$ rules), the Compressed format achieves a 97.6% reduction in memory usage compared to the standard Sparse representation, enabling the simulation of significantly larger topologies on limited hardware.

2.3.2 Hardware Configuration. Each node was configured with the following specifications: **Architecture:** x86_64 (Intel Broadwell microarchitecture), **CPU:** 1 virtual CPU core, **Memory:** 3.5 GB RAM, **Storage:** 50 GB persistent disk, **GPU:** NVIDIA Tesla T4 (16 GB GDDR6 video memory, Turing architecture). The three VMs were interconnected via VPC peering within the same availability zone, achieving sub-millisecond inter-node latency.

2.3.3 Software Environment. The software stack was carefully selected to balance compatibility and performance: **Operating System:** Debian 13, **CUDA Toolkit:** Version 13.1, **NVIDIA Driver:** Version 590.48.01, **MPI Implementation:** Open MPI 5.0.7, **Compiler:** GNU C++ (g++) version 14.2.0

2.3.4 Benchmarking Methodology. Performance measurements were conducted using Google Benchmark for high-precision timing of CPU and GPU operations. Crucially, the benchmarks only capture the simulation steps; auxiliary overheads such as memory allocation, loading the SNP system, and memory deallocation are excluded. For GPU-specific profiling, NVIDIA Nsight Systems provided timeline analysis of kernel execution, memory transfers, and MPI communication, while NVIDIA Nsight Compute delivered detailed metrics on warp stalls, memory hierarchy utilization, and instruction throughput.

The complete source code, build scripts, and raw experimental data are publicly available at <https://github.com/LipatJob/distributed-snp> to facilitate reproducibility and extension of this work.

3 RESULTS

We present the results of our benchmarking experiments using the natural number sorting SNP system as our case study.

3.1 Performance Scaling Analysis




Figure 2: Performance scaling comparison across implementations compared against baseline CPU implementation.

3.1.1 Scaling with Input Size. We evaluated the performance of all implementations across input sizes ranging from 4 to 2048 numbers to be sorted. Figure 2 presents the speedup relative to the baseline CPU implementation (Algorithm 1). It shows that GPU-based implementations achieve superior performance at larger input sizes, with Optimized CUDA MPI (Linear) reaching nearly 30 $\times$ speedup at 2048 elements. The results reveal several critical insights into the scaling behavior of different architectural approaches.

For small input sizes (4–256), the CPU implementation maintains competitive or even superior performance compared to GPU variants. This phenomenon is due to the overhead associated with GPU kernel launches, memory transfers, and MPI initialization. When workloads are too small to saturate GPU cores, these fixed overheads dominate total execution time. We observe a performance anomaly at $N=16$ that warrants further investigation. Preliminary analysis suggests thread underutilization, but a definitive explanation requires additional profiling beyond our current scope.

The crossover point where GPU implementations surpass CPU performance occurs around 512 elements. Beyond this threshold, the parallel nature of SNP systems aligns naturally with the massively parallel architecture of modern GPUs. The Optimized CUDA implementation (Algorithm 5) achieves approximately 19 $\times$ speedup at 2048 elements, demonstrating the effectiveness of our Structure of Arrays (SoA) memory layout and kernel fusion optimizations.

Comparing the distributed implementations, Optimized CUDA MPI (Algorithm 5) with linear partitioning achieves the highest speedup of 29.5 $\times$ at 2048 elements, outperforming even the single-node Optimized CUDA variant. This demonstrates that for sufficiently large problems, the benefits of distributed memory capacity and computation outweigh the communication overhead. Interestingly, the naive CUDA MPI implementations (Algorithm 4) show significantly lower speedups (17–18 $\times$), highlighting the importance of optimization strategies such as ghost zones for reducing inter-node communication.

The Sparse CUDA implementation (Algorithm 2), while memory-efficient, achieves only modest speedups (3.5 $\times$ at 2048 elements). This is expected, as the natural number sorting SNP system exhibits dense connectivity patterns. Sparse matrix formats excel when the connectivity is genuinely sparse but introduce computational overhead when applied to dense graphs due to additional indirection and processing.

3.1.2 Performance at $N=2048$. To further understand the performance of the implementations at higher input sizes, we benchmark all systems across five distinct input distributions: sorted, nearly-sorted, reverse-sorted, few-unique values, and uniform random. This test reveals how architectural choices affect performance when the system operates at capacity, where memory pressure (high contention for managed memory bandwidth), cache thrashing (frequent eviction of useful data from high-speed cache lines), and communication bottlenecks are most pronounced. Figure 3 presents the execution times for these high-scale workloads. Results in execution time (ms) reveals that while most distributions exhibit similar performance, uniform random inputs consistently achieve 40–50% lower execution times, indicating algorithmic sensitivity to input characteristics even at scale.




Figure 3: Execution time performance at N=2048 across different input distributions.

The results indicate that the performance of sorting SNP systems is largely independent of input distribution, with one notable exception: the uniform random distribution consistently executes 40–50% faster than other distributions across all implementations. This behavior stems from the deterministic sorting algorithm embedded in the SNP system’s rule structure. Highly structured inputs (sorted, nearly-sorted, reverse-sorted) require more neuron firings and longer spike propagation chains to reach the final sorted state, whereas uniform random inputs reduce the average distance spikes must travel through the network.

Surprisingly, the choice of partitioning strategy (Linear, Louvain, Red-Blue) shows minimal impact on performance in this particular SNP system. For instance, Naive CUDA MPI implementations vary by less than 10% across partitioning strategies for most distributions. This counterintuitive result can be explained by examining the topology of the sorting SNP system: neurons are densely interconnected and already arranged in a contiguous, regular pattern that mirrors the input array structure. Consequently, even simple linear partitioning naturally minimizes cross-partition communication. The Louvain algorithm, designed to detect communities and minimize edge cuts, finds little room for improvement over naive partitioning in this dense, uniform topology. We hypothesize that partitioning strategies would demonstrate more significant benefits in SNP systems with irregular, sparse, or hierarchical connectivity patterns.

3.2 Execution Time Breakdown

To understand where computational resources are consumed, we profiled the implementations using NVIDIA Nsight Systems. Figure 4 decomposes the total execution time into four categories: CUDA kernel execution, CUDA memory operations (host-device transfers), MPI communication, and CPU computation. The chart shows the percentage of time spent in CUDA kernels, CUDA memory operations, MPI communication, and CPU processing. Lin. = Linear, Louv. = Louvain, R-B = Red-Blue partitioning.

Across all implementations, CPU time is negligible, constituting less than 20% of total execution time. This validates our architectural




Figure 4: Execution time breakdown for distributed implementations.

decision to offload computation to GPUs. CUDA memory transfer overhead is similarly minimal (1–2%), benefiting from the use of pinned memory allocations and efficient batching of transfers.

The dominant component for most implementations is CUDA kernel execution. For example, Louvain Naive MPI spends 63% of its time executing kernels. However, a striking outlier is Linear Naive MPI, which dedicates 62% of execution time to MPI communication. This exposes a critical weakness in the naive communication strategy (Algorithm 4): using `MPI_Allgather` to broadcast the entire spiking state across all nodes at every time step generates substantial network traffic, particularly when partitions are poorly aligned with communication patterns.

The optimized implementations mitigate this issue through targeted ghost zone exchanges (Algorithm 5), where only boundary neurons’ states are communicated. Linear Optimized MPI reduces MPI communication to 23%, while Red-Blue Optimized MPI achieves an even lower 11%. These reductions directly translate to the performance gains observed in Figure 2.

An anomaly appears in Louvain Optimized MPI, spending 73% of execution time on communication—exceeding even the naive variant. While the Louvain algorithm is excellent at minimizing edge cuts in general graphs, the precise reasons for this increased communication overhead in the sorting SNP system require further investigation. This overhead may stem from fragmented partitions creating complex memory gathering patterns for ghost zones, or a conflict between community detection logic and the network’s directional flow.

Figure 5 isolates the MPI communication component, confirming that Red-Blue partitioning consistently minimizes communication time. Red-Blue partitioning achieves the lowest communication overhead for both naive and optimized implementations, while Louvain introduces significant overhead in the optimized variant. The BFS-based wavefront expansion used by Red-Blue creates partitions with minimal boundary surface area, directly reducing the number of boundary neurons that require synchronization. This demonstrates that for structured, spatially-coherent graphs like the




Figure 5: MPI communication time comparison across partitioning strategies.

sorting network, topology-aware partitioning strategies outperform community detection algorithms.

3.3 Large-Scale Memory Capacity

A primary motivation for distributed simulation is overcoming the memory limitations of single GPUs. To quantify this benefit, we conducted a "big data" experiment: we synthesized large-scale SNP systems and measured the maximum number of neurons each implementation could simulate before exhausting memory. Instead of using the sorting SNP model, we generated a synthetic SNP model with high intra-node (i.e. within a single VM) connectivity and low inter-node connectivity to isolate memory capacity effects from communication overhead.

For this experiment, we utilized two of the three VMs from our cluster as worker nodes to execute the distributed simulation, while the third VM served as a master node responsible for coordinating execution, scheduling tasks, and timing measurements. This master-worker configuration allowed us to accurately measure performance without interference from benchmarking overhead on the compute nodes.

Figure 6 presents the results. The axes represent the number of neurons (in millions) and total simulation time (seconds). Linear Optimized MPI with 2 nodes successfully simulates up to 20 million neurons, demonstrating near-linear scaling of memory capacity with the number of nodes. The single-node Linear Optimized MPI implementation reaches a maximum of 10 million neurons, constrained by the 16 GB VRAM available on the test GPU.

By distributing the simulation across 2 nodes, Linear Optimized MPI doubles the effective memory capacity, successfully simulating up to 20 million neurons. The execution time scales approximately linearly with problem size: simulating 20 million neurons takes 6.58 seconds with 2 nodes compared to 3.19 seconds for 10 million neurons on 1 node. The less-than-2 $\times$ time increase for a 2 $\times$ problem size reflects the efficiency of our distributed architecture—the communication overhead introduced by the second node is modest relative to the computational workload.




Figure 6: Maximum achievable simulation scale.

The Naive implementation with 2 nodes, included as a baseline, demonstrates significantly worse scaling. It reaches only 6 million neurons before encountering out-of-memory errors, and its execution time is 2–3 $\times$ slower than the optimized variant. This performance gap underscores the cumulative impact of our optimization strategies: SoA memory layout, ghost zones, and efficient kernel fusion collectively enable both larger and faster simulations.

3.4 Discussion and Implications

Our evaluation yields several key insights for high-performance membrane computing simulations:

Architectural Trade-offs and Scalability. The choice between single-node and distributed implementations depends on problem size and available hardware. For problems fitting within a single GPU’s memory (up to approximately 10 million neurons in our setup), the optimized single-node CUDA implementation offers simplicity and excellent performance. As problems grow, distributed MPI+CUDA implementations become essential, accepting moderate communication overhead in exchange for expanded memory capacity. The near-linear memory scaling observed with 2 nodes demonstrates that weak scaling is achievable; future work should investigate strong scaling across larger clusters.

Optimization and Partitioning. Careful optimization—transitioning from AoS to SoA, employing pinned memory, fusing kernels, and designing targeted communication strategies—can yield 1.5–2 $\times$ performance improvements over naive GPU implementations. Furthermore, sophisticated graph partitioning algorithms do not universally outperform simple strategies. For structured, dense graphs, topology-aware approaches (Red-Blue BFS) or even naive linear partitioning suffice, while community detection algorithms like Louvain are better suited for modular, sparse graphs.

Memory Boundedness. The pervasive memory stalls observed across all implementations reveal that SNP simulations are fundamentally memory-bound on current GPU architectures. This redirects optimization efforts away from arithmetic throughput and toward memory access efficiency, including techniques such as mixed-precision matrix operations via tensor cores, hierarchical

memory strategies with NVLink or PCIe compression, and leveraging emerging technologies like HBM3 and CXL.

4 CONCLUSION

This paper presented a distributed simulation framework for Spiking Neural P systems that bridges the gap between their theoretical computational power and practical scalability. By combining the fine-grained parallelism of CUDA with the distributed-memory capabilities of MPI, our framework enables simulations at scales previously unattainable with single-node approaches.

Our experimental evaluation yielded several key findings. First, GPU-accelerated implementations achieve substantial speedups over CPU baselines, reaching up to $29.5\times$ improvement at 2048 elements when using optimized distributed execution. Second, the framework demonstrates near-linear memory scaling: distributing across two nodes doubled the effective memory capacity, enabling simulations of up to 20 million neurons compared to 10 million on a single GPU. Third, the optimization strategies presented proved essential, collectively delivering $1.5\text{--}2\times$ performance gains over naive GPU implementations. Our analysis further revealed that SNP simulations are fundamentally memory-bound and that partitioning strategy effectiveness is highly topology-dependent.

Several limitations of this work warrant acknowledgment. Our benchmarks focused primarily on the natural number sorting SNP system, which exhibits dense, regular connectivity. SNP systems with sparse, irregular, or hierarchical topologies may exhibit different performance characteristics. Additionally, our experiments were limited to two-node configurations; larger cluster deployments would introduce additional communication challenges and potential bottlenecks. The current implementation assumes one GPU per node; extending support to multi-GPU nodes would require additional engineering effort but could significantly improve computational density.

Future research directions are abundant. On the hardware front, emerging interconnect technologies such as CXL and RDMA could reduce inter-node communication latency, while NVIDIA's NCCL—the optimized collective communications library used for distributed training of large language models—presents a compelling alternative to standard MPI collectives, potentially reducing synchronization overhead by an order of magnitude. Tensor cores and mixed-precision arithmetic could further accelerate the matrix-vector operations central to spike propagation. Extending the framework to support heterogeneous computing environments, where CPUs and GPUs collaborate dynamically based on workload characteristics, could further improve resource utilization.

On the algorithmic front, extending the evaluation to additional case studies beyond natural number sorting—such as SNP systems for subset sum, Boolean satisfiability, or other NP-complete problems with sparse or irregular topologies—would more comprehensively characterize the framework's generality and expose performance trade-offs not visible in dense, regular graphs. Investigating how different partitioning strategies perform across these diverse SNP topologies would provide valuable guidance for practitioners. Additionally, a thorough diagnosis of the performance anomaly observed at $N=16$, where GPU implementations underperform relative to the expected scaling trend, remains an open question; detailed

warp-level profiling and occupancy analysis would help determine whether this stems from thread underutilization, memory access inefficiencies, or kernel launch overhead at small problem sizes. We also plan to conduct a detailed analysis of kernel performance using profiling tools like Nsight Systems and Nsight Compute to further optimize the CUDA kernels for specific SNP workloads.

Finally, this work demonstrates that distributed GPU computing provides a viable path toward simulating large-scale membrane computing systems. As both hardware capabilities and algorithmic techniques continue to advance, we anticipate that the simulation of increasingly complex bio-inspired models will become not only feasible but routine, unlocking new applications in areas ranging from computational biology to unconventional computing architectures.

REFERENCES

- [1] Henry N Adorna. 2022. Matrix representations of spiking neural P systems: Revisited. (2022). arXiv:2211.15156 [cs.NE]
- [2] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefevre. 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008, 10 (2008), P10008.
- [3] Jym Paul Carandang, John Matthew Villaflores, Francis George C. Cabarle, Henry N. Adorna, and Miguel Ángel Martínez-del Amor. 2017. CuSNP: Spiking Neural P Systems Simulators in CUDA. *Romanian Journal of Information Science and Technology* 20, 1 (2017), 57–70.
- [4] Message P Forum. 1994. *MPI: A Message-Passing Interface Standard*. University of Tennessee, USA.
- [5] Javier Hernández-Tello, Miguel Á. Martínez-del Amor, David Orellana-Martín, and Francis George C. Cabarle. 2024. Sparse Spiking Neural-Like Membrane Systems on Graphics Processing Units. *International Journal of Neural Systems* 34, 07 (2024), 2450038. arXiv:https://doi.org/10.1142/S0129065724500382 doi:10.1142/S0129065724500382 PMID: 38755115.
- [6] Jia-Wei Hong and Hsiang-Tsung Kung. 1981. I/O complexity: The red-blue pebble game. In *Proceedings of the 13th Annual ACM Symposium on Theory of Computing (STOC)*. ACM, 326–333.
- [7] Mihai Ionescu, Gheorghe Păun, and Takashi Yokomori. 2006. Spiking neural P systems. *Fundamenta Informaticae* 71, 2-3 (2006), 279–308.
- [8] Mihai Ionescu and Dragoş Sburan. 2008. Some Applications of Spiking Neural P Systems. *Computing and Informatics* 27, 3 (2008), 515–528.
- [9] David B Kirk and Wen-mei W Hwu. 2016. *Programming Massively Parallel Processors: A Hands-on Approach* (3rd ed.). Morgan Kaufmann.
- [10] Fredrik Berg Kjolstad and Marc Snir. 2010. Ghost Cell Pattern. In *Proceedings of the 2010 Workshop on Parallel Programming Patterns (Carefree, Arizona, USA) (ParaPLoP '10)*. Association for Computing Machinery, New York, NY, USA, Article 4, 9 pages. doi:10.1145/1953611.1953615
- [11] Alberto Laporati, Giancarlo Mauri, and Claudio Zandron. 2022. Spiking neural P systems: main ideas and results. *Nat. Comput.* 21, 4 (Dec. 2022), 629–649.
- [12] Miguel Ángel Martínez-del Amor, David Orellana-Martín, Ignacio Pérez-Hurtado, Francis George C Cabarle, and Henry N Adorna. 2021. Simulation of spiking neural P systems with sparse matrix-vector operations. *Processes* 9, 4 (2021), 690.
- [13] NVIDIA Corporation. 2025. CUDA Toolkit Documentation v13.1. <https://docs.nvidia.com/cuda/>
- [14] Gheorghe Păun. 2000. Computing with membranes. *J. Comput. System Sci.* 61, 1 (2000), 108–143.
- [15] Adamos Polycarpou and et al. 2018. Spike: A GPU Optimised Spiking Neural Network Simulator. *bioRxiv* (2018). doi:10.1101/461160
- [16] Tao Wang, Gexiang Zhang, Junbo Zhao, Zhengyou He, Jun Wang, and Mario J Pérez-Jiménez. 2015. Fault diagnosis of electric power systems based on fuzzy reasoning spiking neural P systems. *IEEE Trans. Power Syst.* 30, 3 (May 2015), 1182–1194.
- [17] Esin Yavuz, James Turner, and Thomas Nowotny. 2016. GeNN: A Code Generation Framework for Accelerated Brain Simulations. *Scientific Reports* 6 (2016), 18854. doi:10.1038/srep18854
- [18] Gexiang Zhang, Haina Rong, Ferrante Neri, and Mario J Pérez-Jiménez. 2014. An optimization spiking neural p system for approximately solving combinatorial optimization problems. *Int. J. Neural Syst.* 24, 5 (Aug. 2014), 1440006.
- [19] Gexiang Zhang, Sergey Verlan, Tingfang Wu, Francis George C Cabarle, Jie Xue, David Orellana-Martín, Jianping Dong, Luis Valencia-Cabrera, and Mario J Pérez-Jiménez. 2024. Medical image processing with spiking neural P systems. In *Spiking Neural P Systems*. Springer Nature Singapore, Singapore, 233–271.

Nonbalanced and Nonhomogenous Spiking Neural dP Systems

Victor Dominic D. Naraval
Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
vdnaraval@up.edu.ph

Francis George C. Cabarle
Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
fccabarle@up.edu.ph

Kelvin C. Buño
Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
kcbuno@up.edu.ph

Ren Tristan A. de la Cruz
Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
radelacruz@up.edu.ph

Daryll Carlsten L. Ko
Department of Computer Science
University of the Philippines Diliman
Quezon City, Philippines
dlko@up.edu.ph

Abstract

This paper aims to explore some nonbalanced and nonhomogeneous Spiking Neural dP Systems and their correctness. SNdP systems use Extended SNP systems with Request Rules as components and can recognize languages that SN P systems cannot, such as $L_{ww} = \{ww \mid w \in \{b_1, b_2, \dots, b_k\}^n, n \geq 1\}$. This work mathematically shows that a nonbalanced nonhomogeneous SNdP system introduced by a previous paper has an error and cannot recognize L_{ww} . This paper then introduces a modification of the previous SNdP without changing its balance and homogeneity and is mathematically shown to recognize L_{ww} . The modified SNdP has increased number of neurons and running time while maintaining its communication costs.

Keywords

Membrane Computing, Spiking Neural P Systems, dP Systems, Spiking Neural dP Systems

1 Introduction

Membrane Computing is a branch of natural computing that abstracts ideas and computational models from the structures and functioning of cell or higher-order structures [7]. The fundamental types of P systems, include symbol Transition P systems, systems with symport/antiport rules, systems with string objects, tissue-like P systems, and neural-like P systems.

P systems, by definition, compute in a distributed and parallel way. P systems utilize this to solve computationally hard problems in a feasible time, by producing an exponential workspace in linear time. This is done through utilizing membrane division, membrane creation, and string replication. Aside from these, the variants of P systems like tissue-like systems and neural-like systems are also capable of this through cell division and object replication, respectively. In all of these, P systems are capable of producing polynomial or pseudo-polynomial solutions so NP-complete problems. The first of which was the SAT [7].

However, no P system is proposed to solve a hard problem in a distributed parallel way after splitting the problem into parts [8]. This is what Distributed P Systems or dP systems try to accomplish. A dP system is made up of multiple P automata as components, and they communicate with each other through channels linking their skin membranes, by utilizing antiport rules [8].

Solving parallelizable problems using dP automata is more efficient than by a single P system. In addition to this, when the communication between components is controlled, the computing power of the devices increases. That is why solving a problem like SAT using a dP system, where the inputs are distributed in a balanced way, is faster than deciding the formula with a single P system [8].

Spiking Neural P (SN P) systems is a further evolution of P systems, where it attempts to mimic how the human brain works. An SN P system consists of a set of neurons, in which one can hold any number of spikes. Depending on the rules within the neurons, it can fire (possibly with a delay) and propagate a spike to its neighbors through a synapse [6].

A Spiking Neural dP (SNdP) system is the SN P system counterpart of dP systems. In an SNdP system, the components are SN P systems connected through external synapses [5]. SNdP systems introduced in [4, 5] showed that some languages that are not recognized by SN P systems can be recognized by SNdP systems, such as $L_{ww} = \{ww \mid w \in \{b_1, b_2, \dots, b_k\}^n, n \geq 1\}$.

In this paper, we suggest that a 3-component, nonbalanced and nonhomogeneous SNdP from [4] was inconsistent with its purpose of recognizing L_{ww} . New notations are introduced in this work with the goal to show a mathematical proof to such claims. We also present a modification of the SNdP recognizing L_{ww} , for which we prove the correctness of.

Section 2 of this paper defines the concepts and notations used throughout the paper. Section 3 presents the results which includes the proof that a SNdP in [4] does not recognize L_{ww} , and a new SNdP with proof that it recognizes L_{ww} . Future works and recommendations are discussed in Section 4.

2 Preliminaries

The reader is assumed to be familiar with formal language theory. Let Σ be an alphabet, then Σ^* denotes the set of all finite strings over Σ . λ denotes string with length zero called an empty string, and Σ^+ is the set of all non-empty strings over Σ . When $\Sigma = \{a\}$ is a singleton, then $\{a\}^*$ and $\{a\}^+$ are written as a^* and a^+ , respectively.

DEFINITION 1. For a non-empty finite set P , a family of sets $\{P_1, P_2, \dots, P_n\}$ is called a partition P if and only if for all $1 \leq i, j \leq n$, and $i \neq j$, P_i and P_j are disjoint, and $\bigcup_{i=1}^n P_i = P$.

A partition $\{P_1, P_2, \dots, P_n\}$ of P is called a **balanced partition** if and only if $||P_i| - |P_j|| \leq 1$, for all i, j , $1 \leq i, j \leq n$, $i \neq j$. Otherwise, the partition is an **unbalanced partition**.

2.1 dP Systems

Distributed P systems (dP systems) are introduced by the following notion.

DEFINITION 2. [8] A dP scheme (of degree $n \geq 1$) is a construct

$$\Delta = (O, \Pi_1, \dots, \Pi_n, R)$$

where

- (1) O is an alphabet of objects;
- (2) $\Pi_1, \dots, \Pi_n$ are cell-like P systems (called components) with O as the alphabet of objects and the skin membranes labeled with $s_1, \dots, s_n$, respectively;
- (3) R is a finite set of rules of the form $(s_i, u/v, s_j)$, where $1 \leq i, j \leq n$, $i \neq j$, and $u, v \in O^*$, with $uv \neq \lambda$; $|uv|$ is called the weight of the rule $(s_i, u/v, s_j)$.

The systems $\Pi_1, \dots, \Pi_n$ are called components of the scheme Δ , and the rules in R are called inter-component communication rules. Each component can take an input, work on it, communicate with other components (by means of rules in R), and provide the answer to the problem at the end of a halting computation.

In a dP scheme, aside from time and space, another complexity to consider is the communication cost, which is defined as follows.

DEFINITION 3. [8] In a dP scheme Δ , let $\delta : w_0 \Rightarrow w_1 \Rightarrow \dots \Rightarrow w_h$ be a halting computation in Δ , with w_0 being the initial configuration. Then, for each $i = 0, 1, \dots, h - 1$:

- $\text{ComN}(w_i \Rightarrow w_{i+1}) = 1$ if a communication rule is used in this transition, and 0 otherwise,
- $\text{ComR}(w_i \Rightarrow w_{i+1}) =$ the number of communication rules used in this transition,
- $\text{ComW}(w_i \Rightarrow w_{i+1}) =$ the total weight of the communication rules used in this transition.

Let the set of strings accepted by Δ be denoted as $L(\Delta)$. For $\text{ComX} \in \{\text{ComN}, \text{ComR}, \text{ComW}\}$, we define:

- $\text{ComX}(\delta) = \sum_{i=0}^{h-1} \text{ComX}(w_i \Rightarrow w_{i+1})$, for $\delta : w_0 \Rightarrow w_1 \Rightarrow \dots \Rightarrow w_h$ a halting computation,
- $\text{ComX}(w, \Delta) = \min\{\text{ComX}(\delta) \mid \delta : w_0 \Rightarrow w_1 \Rightarrow \dots \Rightarrow w_h \text{ is a computation in } \Delta \text{ which accepts the string } w\}$,
- $\text{ComX}(\Delta) = \max\{\text{ComX}(w, \Delta) \mid w \in L(\Delta)\}$

2.2 Extended Spiking Neural P Systems with Request Rules

Extended Spiking Neural P Systems with request rules are introduced as follows.

DEFINITION 4. [5] An (Extended) Spiking Neural P System with Request Rules, of degree $m \geq 1$, is a construct of the form

$$\Pi = (O, \sigma_1, \dots, \sigma_m, \text{syn}, i_0)$$

where

- (1) $O = \{a\}$ is the singleton alphabet (a is called spike);
- (2) $\sigma_1, \dots, \sigma_m$ are neurons, of the form $\sigma_i = (n_i, R_i)$, $1 \leq i \leq m$, where:

- (a) $n_i \geq 0$ is the initial number of spikes contained in σ_i ;
- (b) R_i is a finite set of rules of one of the following forms:
 - (i) (spiking rules) $E/a^c \rightarrow a^p$, where E is a regular expression over the alphabet $\{a\}$ and $c \geq p \geq 1$ are natural numbers;
 - (ii) (request rules) $E/\lambda \leftarrow a^r$, where E is a regular expression over the alphabet $\{a\}$ and $r \geq 1$ is a natural number
 - (iii) (forgetting rules) $a^s \rightarrow \lambda$, where $s \geq 1$ is a natural number, with the restriction that for each rule $E/a^c \rightarrow a$ of type (i) or $E/\lambda \leftarrow a^r$ of type (ii) R_i , we have $a^s \notin L(E)$, where $L(E)$ is the regular language defined by E ;
- (3) $\text{syn} \subseteq \{1, 2, \dots, m\} \times \{1, 2, \dots, m\}$ with $i \neq j$ for all $(i, j) \in \text{syn}$, $1 \leq i, j \leq m$, is the set of synapses between neurons;
- (4) $i_0 \in \{1, 2, \dots, m\}$ indicates the output neuron.

A rule $E/a^c \rightarrow a^p$ is applied as follows. If the neuron σ_i contains k spikes and $a^k \in L(E)$, $k \geq c$, then the rule can fire, and its application means consuming (removing) c spikes (thus only $k - c$ remain in σ_i) and producing p spikes, which will immediately exit the neuron. A rule $E/\lambda \leftarrow a^r$ is used if the neuron contains k spikes and $a^k \in L(E)$; no spike is consumed, but r spikes are added to the spikes in σ_i . In turn, a rule $a^s \rightarrow \lambda$ is used if the neuron contains exactly s spikes, which are removed ("forgotten"). A global clock is assumed, marking the time for the whole system, and hence the functioning of the system is synchronized.

If a rule $E/a^c \rightarrow a^p$ has $E = a^c$, then we will write it in the simplified form $a^c \rightarrow a^p$.

The spikes emitted by a neuron σ_i go to all the neurons σ_j such that $(i, j) \in \text{syn}$, i.e., if σ_i has used a rule $E/a^c \rightarrow a^p$, then each neuron σ_j receives p spikes. The spikes produced by a rule $E/\lambda \leftarrow a^r$ are added to the spikes in the neuron and those received from other neurons, hence they are counted/used in the next step of the computation.

If several rules can be used at the same time, then the one to be applied is chosen non-deterministically.

During computation, a configuration of the system is described by the number of spikes present in each neuron; thus, the initial configuration is described by the numbers $n_1, n_2, \dots, n_m$.

Using the rules as described above, one can define transitions among configurations. Any sequence of transitions starting in the initial configuration is called a computation. A computation halts if it reaches a configuration where no rule can be used.

A possible result of a halting computation of an extended SNP system with request rules is a string. A singular neuron is designated as the input neuron, σ_{in} , in which it is the only neuron with request rules. For a string over an alphabet $\Sigma = b_1, \dots, b_k$, for some $k \geq 1$, a symbol b_r can be associated with a step of communication when r spikes are requested by σ_{in} . The sequence of spikes that enter through σ_{in} can then be associated with a string over Σ . It is to be noted that this paper, as well as the papers to be cited, does not associate no spike entering the system to any symbol in Σ . The result of a halting computation is defined as follows.

DEFINITION 5. Let Π be an Extended SNP system with Request Rules with neuron σ_{in} containing request rule. For a halting computation, suppose that a sequence of spikes $\langle r_1, \dots, r_n \rangle$, for some $n \geq 1$, are requested by σ_{in} , and for all $1 \leq j \leq n$, $1 \leq r_j \leq k$. Then, Π accepts the string $b_{r_1} \dots b_{r_n}$.

The set of all strings over $\{b_1, \dots, b_k\}$ accepted by Π is denoted as $L(\Pi)$.

2.3 Spiking Neural dP Systems

Spiking Neural dP Systems (abbreviated as SN dP Systems) are then introduced as follows.

DEFINITION 6. [5] An SN dP system is a construct of the form

$$\Delta = (O, \Pi_1, \dots, \Pi_n, esyn)$$

where

- (1) $O = \{a\}$ is the singleton alphabet (a is the spike);
- (2) $\Pi_i = (O, \sigma_{i,1}, \dots, \sigma_{i,k}, syn, in_i)$ is an SN P system with request rules present only in neuron σ_{in_i} ;
- (3) $esyn$ is the set of external synapses between neurons from different systems Π_i (called components), with the restriction that between two components Π_i, Π_j , there exist at most one link from a neuron of Π_i to a neuron of Π_j , and vice versa.

Each component can take an input (using request rules), work on it by using the spiking and forgetting rules in the neurons, and communicate with other components (along the synapses in $esyn$); the communication is done as usual inside the components: when a spiking rule produces a number of spikes, they are sent simultaneously to all neurons, inside the component or outside it, in other components, provided that a synapse (internal or external) exists to the destination.

When r spikes are taken from the environment, a symbol b_r is associated with that step, and the strings that can be formed by and introduced in the system are over an alphabet $\Sigma = \{b_1, \dots, b_k\}$, with k being the maximum number of spikes that can be obtained by the request rule of a component.

A halting computation with respect to Δ accepts the string $x = x_1 x_2 \dots x_n$ over the alphabet Σ if the components, beginning from their initial configurations and working in a synchronous, nondeterministic way, bring from the environment the substrings $x_1, x_2, \dots, x_n$ and halts eventually.

DEFINITION 7. Let Δ be an SNdP system with n components $\Pi_1, \dots, \Pi_n$, with each component having $k_1, \dots, k_n$ neurons, respectively. A configuration of Δ at time step i is an m -dimension vector $C_i = \langle c_1, \dots, c_m \rangle$ where m is the total number of neurons in Δ , that is, $m = k_1 + \dots + k_n$. An element in C_i , denoted as c_j for some $1 \leq j \leq m$, represents the number of spikes in neuron $\sigma_{p,q}$ where $1 \leq p \leq n$ and $1 \leq q \leq k_p$, such that $j = q + \sum_{i=1}^{p-1} k_i$. The initial configuration of Δ is defined as $C_0 = \langle n_1, \dots, n_m \rangle$.

DEFINITION 8. Let $C = \langle c_1, \dots, c_m \rangle$ be a configuration of an SNdP system Δ . A portion of C can be denoted as $C[\alpha_1, \alpha_2, \dots] = \langle c_{d_1}, c_{d_2}, \dots \rangle$ such that for some $j \geq 1$, α_j is either of form k or of form $p : q$, where $1 \leq k, p, q \leq m$ and $p \leq q$. If $\alpha_j = k$, then $c_k \in \langle c_{d_1}, c_{d_2}, \dots \rangle$. Else if $\alpha_j = p : q$, then $c_p, c_{p+1}, \dots, c_q \in \langle c_{d_1}, c_{d_2}, \dots \rangle$. It is restricted that $d_l < d_{l+1}$ for $l \geq 1$.

DEFINITION 9. Let $C = \langle c_1, \dots, c_m \rangle$ be a configuration of an SNdP system Δ . Configuration C can also be represented as $C[\alpha_1, \alpha_2, \dots] = \langle \langle c_{d_1}, c_{d_2}, \dots \rangle, a \rangle$ if and only if $C[\alpha_1, \alpha_2, \dots] = \langle c_{d_1}, c_{d_2}, \dots \rangle$ and for all k such that $1 \leq k \leq m$ and $c_k \notin \langle c_{d_1}, c_{d_2}, \dots \rangle$, it holds that $c_k = a$.

DEFINITION 10. Let Δ be an SNdP system. A configuration C_i yields a configuration C_{i+1} , denoted as $C_i \Rightarrow C_{i+1}$, if and only if C_{i+1} is obtained from C_i by applying spiking, forgetting, or request rules in a globally parallel and locally sequential manner. This is called a transition.

DEFINITION 11. Let C and C' be configurations of an SNdP system Δ . If zero or more transitions can yield C' from C , then the series of transitions can be denoted as $C \Rightarrow^* C'$.

DEFINITION 12. Let $C_i = \langle c_1, \dots, c_m \rangle$ and $C_{i+1} = \langle c'_1, \dots, c'_m \rangle$ be configurations of an SNdP system Δ . The transition $C_i \Rightarrow C_{i+1}$ can also be represented as $C_i \Rightarrow' C_{i+1}[b_1, b_2, \dots] = \langle c'_{d_1}, c'_{d_2}, \dots \rangle$ (note the use of $\Rightarrow'$) if and only if for all k such that $1 \leq k \leq m$ and $c'_k \notin \langle c'_{d_1}, c'_{d_2}, \dots \rangle$, it holds that $c_k = c'_k$.

In the use of notation $\Rightarrow'$, the neurons not included in the partial resulting configuration indicate that those neurons have unchanged number of spikes from the previous configuration.

DEFINITION 13. Let Δ be an SNdP system. If from a configuration C_i , $i \geq 0$, there are no more applicable rules of any form, that is, there is no C_{i+1} such that $C_i \Rightarrow C_{i+1}$, then C_i is referred to as a halting configuration. A halting configuration is denoted as C_h .

DEFINITION 14. Let Δ be an SNdP system. A computation of Δ is a sequence of configurations $C_\Delta : C_0 \Rightarrow C_1 \Rightarrow \dots$, where C_0 is the initial configuration of Δ . A computation C_Δ is a halting computation if and only if $C_\Delta : C_0 \Rightarrow C_1 \Rightarrow \dots \Rightarrow C_h$, where C_h is a halting configuration.

DEFINITION 15. [9] Balance refers to the input partition of the SNdP system. Let Δ be an SNdP system of degree n and $L(\Delta)$ be the language of Δ over an alphabet $\Sigma = \{b_1, \dots, b_k\}$. For an input $x \in \Sigma^*$, the input partition is a set $\{x_1, \dots, x_n\}$ where $x = x_1 \dots x_n$ such that each $x_i \in \Sigma^*$. The input partition is balanced if $||x_i| - |x_j|| \geq 1$ for all $1 \leq i, j \leq n$. Otherwise, it is nonbalanced.

DEFINITION 16. [4] In homogeneous SNdP systems, there is a one-to-one correspondence between the components and each component has the same set of neurons, the same set of starting spikes, spiking rules, request rules, and forgetting rules. Likewise, in a non-homogeneous SNdP system, components can have different sets of neurons.

3 Results

In [4, 5], SNdP systems accepting the language $L_{ww} = \{ww \mid w \in \{b_1, b_2, \dots, b_k\}^n, n \geq 1\}$ were presented.

3.1 Previous example of Nonbalanced and Nonhomogeneous SNdP System

In [4], a 3-component SNdP system, Δ , accepting the language $L_{ww} = \{ww \mid w \in \{b_1, b_2, \dots, b_k\}^n, n \geq 1\}$ that uses nonbalanced partition and has nonhomogenous components was presented, constructed as follows.

$$\Delta = (\{a\}, \Pi_1, \Pi_2, \Pi_3, \{((1, 4), (3, 1)), ((1, 7), (2, 1)), ((2, 2), (1, 3)), ((3, 2), (1, 6))\})$$

Let $\Sigma_k = \{b_1, b_2, \dots, b_k\}$. The input string, say $w' \in \Sigma_k^{2n}$, $n \geq 1$, will be partitioned into two substrings: $w' = w_1 w_2$, where $|w_1| =$

$|w_2| \leq 1$. The first component, Π_1 , takes in w_1 while the second and third components, Π_2 and Π_3 , takes in w_2 , which is further partitioned between the two components: Let $w_2 = x_1x_2\dots x_n$ where $n = |w_2|$. If n is odd, $y_1 = x_1x_3\dots x_n$ and $y_2 = x_2x_4\dots x_{n-1}$. If n is even, $y_1 = x_1x_3\dots x_{n-1}$ and $y_2 = x_2x_4\dots x_n$.

Figure 1 from [4] supposedly shows Δ recognizing L_{ww} , where $x = \lceil \frac{|w|}{2} \rceil$ and $y = |w| - x$, and $1 \leq r \leq k$. The r spikes represents the character b_r from the alphabet of w , $\{b_1, b_2, \dots, b_k\}$. The first component, Π_1 , takes in w_1 , while Π_2 takes y_1 and Π_3 takes in y_2 .

However, in this work, we show that Δ from [4] contains errors in which the system does not recognize L_{ww} . Using the definitions and notations from section 2.3, we prove that Δ requests additional input outside of ww , resulting to its language not being L_{ww} .

Assume the case that the input is ww . Let $w = x_1x_2\dots x_n$ where $n = |w|$. If n is odd, $y_1 = x_1x_3\dots x_n$ and $y_2 = x_2x_4\dots x_{n-1}$. If n is even, $y_1 = x_1x_3\dots x_{n-1}$ and $y_2 = x_2x_4\dots x_n$.

LEMMA 1. In Δ , for any configuration $C_i[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2p + 1, 2q \rangle, 0)$, if $p \geq 1$, then r spikes, for some $1 \leq r \leq k$, will be requested by $\sigma_{1,1}$ after C_i and by $\sigma_{2,2}$ after C_{i+1} , and $C_i \Rightarrow^* C_{i+r+3}[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 2p - 2, 2q + 1 \rangle, 0)$.

Proof: Let $p \geq 1$ and r be the spikes requested by $\sigma_{1,1}$ or $\sigma_{2,2}$ if a requesting rule is applied to the neurons.

If $r = 1$:

$$\begin{aligned} C_i[1, 2, 5, 10, 12] &= (\langle 1, 2, 0, 2p + 1, 2q \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 13] &= (\langle 2, 2, 0, 2p - 2, 1, 2q, 0 \rangle, 0) \\ \Rightarrow' C_{i+2}[1 : 5, 11] &= \langle 0, 4, 0, 0, 2, 2 \rangle \\ \Rightarrow' C_{i+3}[1 : 5, 11] &= \langle 0, 0, 4, 2, 2, 0 \rangle \\ \Rightarrow C_{i+4}[1, 2, 5, 10, 12] &= (\langle 1, 0, 2, 2p - 2, 2q + 1 \rangle, 0) \end{aligned}$$

If $r = 2$:

$$\begin{aligned} C_i[1, 2, 5, 10, 12] &= (\langle 1, 2, 0, 2p + 1, 2q \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 13] &= (\langle 3, 2, 0, 2p - 2, 1, 2q, 0 \rangle, 0) \\ \Rightarrow' C_{i+2}[1 : 5, 11] &= \langle 2, 3, 0, 0, 1, 3 \rangle \\ \Rightarrow' C_{i+3}[1 : 5, 11] &= \langle 0, 4, 2, 1, 2, 2 \rangle \\ \Rightarrow' C_{i+4}[1 : 5, 11] &= \langle 0, 0, 4, 2, 2, 0 \rangle \\ \Rightarrow C_{i+5}[1, 2, 5, 10, 12] &= (\langle 1, 0, 2, 2p - 2, 2q + 1 \rangle, 0) \end{aligned}$$

If $r > 2$:

$$\begin{aligned} C_i[1, 2, 5, 10, 12] &= (\langle 1, 2, 0, 2p + 1, 2q \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 13] &= (\langle r + 1, 2, 0, 2p - 2, 1, 2q, 0 \rangle, 0) \\ \Rightarrow' C_{i+2}[1 : 5, 11] &= \langle r, 3, 0, 0, 1, r + 1 \rangle \\ \text{for } 1 \leq j \leq r - 2 : \\ \Rightarrow' C_{i+j+2}[1 : 5, 11] &= \langle r - j, 3, 2, 1, 1, r - j + 1 \rangle \\ \Rightarrow' C_{i+r+1}[1 : 5, 11] &= \langle 0, 4, 2, 1, 2, 2 \rangle \\ \Rightarrow' C_{i+r+2}[1 : 5, 11] &= \langle 0, 0, 4, 2, 2, 0 \rangle \\ \Rightarrow C_{i+r+3}[1, 2, 5, 10, 12] &= (\langle 1, 0, 2, 2p - 2, 2q + 1 \rangle, 0) \end{aligned}$$

For all cases, lemma 1 holds.

LEMMA 2. In Δ , for any configuration $C_i[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 2p, 2q + 1 \rangle, 0)$, if $q \geq 1$, then r spikes, for some $1 \leq r \leq k$, will be requested by $\sigma_{1,1}$ after C_i and by $\sigma_{3,2}$ after C_{i+1} , and $C_i \Rightarrow^* C_{i+r+3}[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2p + 1, 2q - 2 \rangle, 0)$.

Proof: Let $q \geq 1$ and r be the spikes requested by $\sigma_{1,1}$ or $\sigma_{3,2}$ if a requesting rule is applied to the neurons.

If $r = 1$:

$$\begin{aligned} C_i[1, 2, 5, 10, 12] &= (\langle 1, 0, 2, 2p, 2q + 1 \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 13] &= (\langle 2, 0, 2, 2p, 0, 2q - 2, 1 \rangle, 0) \\ \Rightarrow' C_{i+2}[1, 2, 5 : 7, 13] &= \langle 0, 2, 4, 0, 0, 2 \rangle \\ \Rightarrow' C_{i+3}[1, 2, 5 : 7, 13] &= \langle 0, 2, 0, 4, 2, 0 \rangle \\ \Rightarrow C_{i+4}[1, 2, 5, 10, 12] &= (\langle 1, 2, 0, 2p + 1, 2q - 2 \rangle, 0) \end{aligned}$$

If $r = 2$:

$$\begin{aligned} C_i[1, 2, 5, 10, 12] &= (\langle 1, 0, 2, 2p, 2q + 1 \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 13] &= (\langle 3, 0, 2, 2p, 0, 2q - 2, 1 \rangle, 0) \\ \Rightarrow' C_{i+2}[1, 2, 5 : 7, 13] &= \langle 2, 1, 3, 0, 0, 3 \rangle \\ \Rightarrow' C_{i+3}[1, 2, 5 : 7, 13] &= \langle 0, 2, 4, 2, 1, 2 \rangle \\ \Rightarrow' C_{i+4}[1, 2, 5 : 7, 13] &= \langle 0, 2, 0, 4, 2, 0 \rangle \\ \Rightarrow C_{i+5}[1, 2, 5, 10, 12] &= (\langle 1, 2, 0, 2p + 1, 2q - 2 \rangle, 0) \end{aligned}$$

If $r > 2$:

$$\begin{aligned} C_i[1, 2, 5, 10, 12] &= (\langle 1, 0, 2, 2p, 2q + 1 \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 13] &= (\langle r + 1, 0, 2, 2p, 0, 2q - 2, 1 \rangle, 0) \\ \Rightarrow' C_{i+2}[1, 2, 5 : 7, 13] &= \langle r, 1, 3, 0, 0, r + 1 \rangle \\ \text{for } 1 \leq j \leq r - 2 : \\ \Rightarrow' C_{i+j+2}[1, 2, 5 : 7, 13] &= \langle r - j, 1, 3, 2, 1, r - j + 1 \rangle \\ \Rightarrow' C_{i+r+1}[1, 2, 5 : 7, 13] &= \langle 0, 2, 4, 2, 1, 2 \rangle \\ \Rightarrow' C_{i+r+2}[1, 2, 5 : 7, 13] &= \langle 0, 2, 0, 4, 2, 0 \rangle \\ \Rightarrow C_{i+r+3}[1, 2, 5, 10, 12] &= (\langle 1, 2, 0, 2p + 1, 2q - 2 \rangle, 0) \end{aligned}$$

For all cases, lemma 2 holds.

The starting configuration of lemma 1 and the resulting configuration of lemma 2 are similar while the resulting configuration of lemma 1 and the starting configuration of lemma 2 are similar. This results in a possible loop.

LEMMA 3. In Δ , for any configuration $C_i[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2p + 1, 2q \rangle, 0)$, if $p, q \geq 1$, $C_i \Rightarrow^* C_{i+r_1+r_2+6}[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2p - 1, 2q - 2 \rangle, 0)$ where for some $1 \leq r_1, r_2 \leq k$, neuron $\sigma_{1,1}$ would have requested, in order, r_1, r_2 spikes, $\sigma_{2,2}$ would have requested r_1 spikes, and $\sigma_{3,2}$ would have requested r_2 spikes.

Proof: Let $p, q \geq 1$. From lemma 1, $C_i[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2p + 1, 2q \rangle, 0)$ yields to $C_{i+r_1+3}[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 2p - 2, 2q + 1 \rangle, 0)$, for some r_1 , while r_1 spikes is requested by $\sigma_{1,1}$ after C_i and by $\sigma_{2,2}$ after C_{i+1} . Using lemma 2, $C_{i+r_1+3}[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 2p - 2, 2q + 1 \rangle, 0)$ yields to $C_{i+r_1+r_2+6}[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 2p - 1, 2q - 2 \rangle, 0)$, for some r_2 , while r_2 spikes is requested by $\sigma_{1,1}$ after C_{i+r_1+3} and by $\sigma_{3,2}$ after C_{i+r_1+4} .

LEMMA 4. In Δ , for any configuration $C_i[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 2p, 2q + 1 \rangle, 0)$, if $p, q \geq 1$, $C_i \Rightarrow^* C_{i+r_1+r_2+6}[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 2p - 2, 2q - 1 \rangle, 0)$ where for some $1 \leq r_1, r_2 \leq k$, neuron $\sigma_{1,1}$ would have requested, in order, r_1, r_2 spikes, $\sigma_{3,2}$ would have requested r_1 spikes, and $\sigma_{2,2}$ would have requested r_2 spikes.

Proof: Let $p, q \geq 1$. From lemma 2, $C_i[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 2p, 2q + 1 \rangle, 0)$ yields to $C_{i+r_1+3}[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2p + 1, 2q -$




Figure 1: The nonbalanced, nonhomogeneous SNdP system presented in [4]

2), 0), for some r_1 , while r_1 spikes is requested by $\sigma_{1,1}$ after C_i and by $\sigma_{3,2}$ after C_{i+1} . Using lemma 1, $C_{i+r_1+3}[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2p + 1, 2q - 2 \rangle, 0)$ yields to $C_{i+r_1+r_2+6}[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2p - 2, 2q - 1 \rangle, 0)$, for some r_2 , while r_2 spikes is requested by $\sigma_{1,1}$ after C_{i+r_1+3} and by $\sigma_{2,2}$ after C_{i+r_1+4} .

Since $x = \lceil \frac{|w|}{2} \rceil$ and $y = |w| - x$, then there are two cases: $|w|$ is even which results to $x = y$, or $|w|$ is odd which results to $x = y + 1$. The configuration loop can be repeated until $\sigma_{2,1}$ and $\sigma_{3,1}$ contain zero or one spike.

LEMMA 5. In Δ , if $x = y$, then $C_1[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2x + 1, 2y \rangle, 0) \Rightarrow^* C_f[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 1, 0 \rangle, 0)$ where $\sigma_{1,1}$ would have requested, in order, $r_1, r_2, \dots, r_n$ spikes, $\sigma_{2,2}$ would have requested, in order, $r_1, r_3, \dots, r_{n-1}$ spikes, and $\sigma_{3,2}$ would have requested, in order, $r_2, r_4, \dots, r_n$ spikes, and $f = 1 + 3n + \sum_{i=1}^n r_i$, for some $1 \leq r_1, r_2, \dots, r_n \leq k$.

Proof: Since $x = y$, which also implies $n = 2x$, then $C_1[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2x + 1, 2x \rangle, 0)$. From C_1 , for $1 \leq i \leq x$: apply lemma 3 where neuron $\sigma_{1,1}$ requests, in order, r_{2i-1}, r_{2i} spikes, $\sigma_{2,2}$ requests r_{2i-1} spikes, and $\sigma_{3,2}$ requests r_{2i} spikes. This yields to $C_f[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 1, 0 \rangle, 0)$ where $\sigma_{1,1}$ requested, in order, $r_1, r_2, \dots, r_{2x}$ spikes, $\sigma_{2,2}$ requested, in order, $r_1, r_3, \dots, r_{2x-1}$ spikes, and $\sigma_{3,2}$ requested, in order, $r_2, r_4, \dots, r_{2x}$ spikes, and $f = 1 + 6x + \sum_{i=1}^{2x} r_i$. Since $n = 2x$, then the result is $C_f[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 1, 0 \rangle, 0)$ where $\sigma_{1,1}$ requested, in order, $r_1, r_2, \dots, r_n$ spikes, $\sigma_{2,2}$ requested, in order, $r_1, r_3, \dots, r_{n-1}$ spikes, and $\sigma_{3,2}$ requested, in order, $r_2, r_4, \dots, r_n$ spikes, and $f = 1 + 3n + \sum_{i=1}^n r_i$.

LEMMA 6. In Δ , if $x = y + 1$, then $C_1[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2x + 1, 2y \rangle, 0) \Rightarrow^* C_f[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 0, 1 \rangle, 0)$ where $\sigma_{1,1}$ would have requested, in order, $r_1, r_2, \dots, r_n$ spikes, $\sigma_{2,2}$ would have requested, in order, $r_1, r_3, \dots, r_n$

spikes, and $\sigma_{3,2}$ would have requested, in order, $r_2, r_4, \dots, r_{n-1}$ spikes, and $f = 1 + 3n + \sum_{i=1}^n r_i$, for some $1 \leq r_1, r_2, \dots, r_n \leq k$.

Proof: Since $x = y + 1$, which also implies $n = 2y + 1$, then $C_1[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2y + 3, 2y \rangle, 0)$. From C_1 , apply lemma 1. This yields to $C_{r_1+4}[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 2y, 2y + 1 \rangle, 0)$ where r_1 spikes is requested by $\sigma_{1,1}$ after C_1 and by $\sigma_{2,2}$ after C_2 . If $y = 0$, then $C_{r_1+4}[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 0, 1 \rangle, 0)$, for which lemma 6 holds. Otherwise, from C_{r_1+4} , for $1 \leq i \leq y$: apply lemma 4 where neuron $\sigma_{1,1}$ requests, in order, r_{2i}, r_{2i+1} spikes, $\sigma_{3,2}$ requests r_{2i} spikes, and $\sigma_{2,2}$ requests r_{2i+1} spikes. This yields to $C_f[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 0, 1 \rangle, 0)$ where $\sigma_{1,1}$ requested, in order, $r_2, r_3, \dots, r_{2y+1}$ spikes, $\sigma_{2,2}$ requested, in order, $r_3, r_5, \dots, r_{2y+1}$ spikes, and $\sigma_{3,2}$ requested, in order, $r_2, r_4, \dots, r_{2y}$ spikes, and $f = 4 + 6y + \sum_{i=1}^{2y+1} r_i$. Compounding with the result from C_1 to C_{r_1+4} , in total, $\sigma_{1,1}$ requested, in order, $r_1, r_2, \dots, r_{2y+1}$ spikes, $\sigma_{2,2}$ requested, in order, $r_1, r_3, \dots, r_{2y+1}$ spikes, and $\sigma_{3,2}$ requested, in order, $r_2, r_4, \dots, r_{2y}$ spikes. Since $n = 2y + 1$, then the result, starting from C_1 , is $C_f[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 0, 1 \rangle, 0)$ where $\sigma_{1,1}$ requested, in order, $r_1, r_2, \dots, r_n$ spikes, $\sigma_{2,2}$ requested, in order, $r_1, r_3, \dots, r_n$ spikes, and $\sigma_{3,2}$ requested, in order, $r_2, r_4, \dots, r_{n-1}$ spikes, and $f = 1 + 3n + \sum_{i=1}^n r_i$.

The starting configuration of Δ is $C_0[2, 7, 10, 12] = (\langle 2, 2, 2x, 2y \rangle, 0)$. This yields to $C_1[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 2y + 3, 2y \rangle, 0)$.

If n is even, applying lemma 5 with spike requests r_i for all $1 \leq i \leq n$, $x_i = b_{r_i}$ results in $C_f[1, 2, 5, 10, 12] = (\langle 1, 2, 0, 1, 0 \rangle, 0)$ where $f = 1 + 3n + \sum_{i=1}^n r_i$, where ww is taken as an input. From here, since $C_f[1] = \langle 1 \rangle$, $\sigma_{1,1}$ uses another request rule, taking another input from the environment, which is inconsistent to the purpose of the system recognizing L_{ww} .

If n is odd, applying lemma 6 with spike requests r_i for all $1 \leq i \leq n$, $x_i = b_{r_i}$ results in $C_f[1, 2, 5, 10, 12] = (\langle 1, 0, 2, 0, 1 \rangle, 0)$ where $f = 1 + 3n + \sum_{i=1}^n r_i$, where ww is taken as an input. From here,

since $C_f[1] = \langle 1 \rangle$, $\sigma_{1,1}$ uses another request rule, taking another input from the environment, which is inconsistent to the purpose of the system recognizing L_{ww} .

In either case, Δ does not recognize L_{ww} . $\square$

3.2 Modified Nonbalanced and Nonhomogeneous SNdP System

PROPOSITION 1. $L_{ww} = \{ww \mid w \in \{b_1, b_2, \dots, b_k\}^n, n \geq 1\}$ can be recognized by a 3-component SNdP system, Δ' , with nonbalanced input partition and nonhomogeneous components, with $\text{ComN}(\Delta') = \text{ComR}(\Delta') = (k+1)n$, and $\text{ComW}(\Delta') = (k+2)n$, and Δ' uses a total of 16 neurons across all components.

In proving proposition 1, let $\Sigma_k = \{b_1, b_2, \dots, b_k\}$. The input string, say $w' \in \Sigma_k^{2n}$, $n \geq 1$, will be partitioned into two substrings: $w' = w_1 w_2$, where $||w_1| - |w_2|| \leq 1$. The first component, Π_1 , takes in w_1 while the second and third components, Π_2 and Π_3 , takes in w_2 , which is further partitioned between the two components: Let $w_2 = x_1 x_2 \dots x_n$ where $n = |w_2|$. If n is odd, $y_1 = x_1 x_3 \dots x_n$ and $y_2 = x_2 x_4 \dots x_{n-1}$. If n is even, $y_1 = x_1 x_3 \dots x_{n-1}$ and $y_2 = x_2 x_4 \dots x_n$.

$$\Delta' = (\{a\}, \Pi_1, \Pi_2, \Pi_3, \{((1, 4), (3, 1)), ((1, 7), (2, 1)), ((2, 3), (1, 3)), ((3, 3), (1, 6))\})$$

Figure 2 shows Δ' recognizing L_{ww} , where $x = \lceil \frac{|w|}{2} \rceil$ and $y = |w| - x$, and $1 \leq r \leq k$. The r spikes represents the character b_r from the alphabet of w , $\{b_1, b_2, \dots, b_k\}$. The first component, Π_1 , takes in w_1 , while Π_2 takes y_1 and Π_3 takes in y_2 .

The modification of the SNdP includes a new neuron $\sigma_{1,10}$ which acts as a counter for the incoming input in component 1, similar to $\sigma_{2,1}$ and $\sigma_{3,1}$ for components 2 and 3, respectively, to make sure that all components halt after receiving input ww . Additional neurons in components 2 and 3 are also added to account for the time delay brought by the new neuron in component 1.

We show that Δ' recognizes L_{ww} using the definitions and notations from section 2.3.

Assume the case that the input is ww . Let $w = x_1 x_2 \dots x_n$ where $n = |w|$. If n is odd, $y_1 = x_1 x_3 \dots x_n$ and $y_2 = x_2 x_4 \dots x_{n-1}$. If n is even, $y_1 = x_1 x_3 \dots x_{n-1}$ and $y_2 = x_2 x_4 \dots x_n$.

LEMMA 7. In Δ' , for any configuration $C_i[2, 5, 10, 11, 14] = (\langle 2, 0, 2p+2q+1, 2p+1, 2q \rangle, 0)$, if $p \geq 1$, then r spikes, for some $1 \leq r \leq k$, will be requested by $\sigma_{1,1}$ after C_{i+1} and by $\sigma_{2,3}$ after C_{i+2} , and $C_i \Rightarrow^* C_{i+r+4}[2, 5, 10, 11, 14] = (\langle 0, 2, 2p+2q-1, 2p-2, 2q+1 \rangle, 0)$.

Proof: Let $p \geq 1$ and r be the spikes requested by $\sigma_{1,1}$ or $\sigma_{2,3}$ if a requesting rule is applied to the neurons.

If $r = 1$:

$$\begin{aligned} C_i[2, 5, 10, 11, 14] &= (\langle 2, 0, 2p+2q+1, 2p+1, 2q \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 16] &= (\langle 1, 2, 0, 2p+2q-2, 2p-2, 1, 0, 2q, 0, 0 \rangle, 0) \\ \Rightarrow C_{i+2}[1, 2, 5, 10 : 16] &= (\langle 2, 2, 0, 2p+2q-2, 2p-2, 0, 1, 2q, 0, 0 \rangle, 0) \\ \Rightarrow' C_{i+3}[1 : 5, 13] &= \langle 0, 4, 0, 0, 2, 2 \rangle \\ \Rightarrow' C_{i+4}[1 : 5, 11] &= \langle 0, 0, 4, 2, 2, 0 \rangle \\ \Rightarrow C_{i+5}[2, 5, 10, 11, 14] &= (\langle 0, 2, 2p+2q-1, 2p-2, 2q+1 \rangle, 0) \end{aligned}$$

If $r = 2$:

$$\begin{aligned} C_i[2, 5, 10, 11, 14] &= (\langle 2, 0, 2p+2q+1, 2p+1, 2q \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 16] &= (\langle 1, 2, 0, 2p+2q-2, 2p-2, 1, 0, 2q, 0, 0 \rangle, 0) \\ \Rightarrow C_{i+2}[1, 2, 5, 10 : 16] &= (\langle 3, 2, 0, 2p+2q-2, 2p-2, 0, 1, 2q, 0, 0 \rangle, 0) \\ \Rightarrow' C_{i+3}[1 : 5, 13] &= \langle 2, 3, 0, 0, 1, 3 \rangle \\ \Rightarrow' C_{i+4}[1 : 5, 13] &= \langle 0, 4, 2, 1, 2, 2 \rangle \\ \Rightarrow' C_{i+5}[1 : 5, 11] &= \langle 0, 0, 4, 2, 2, 0 \rangle \\ \Rightarrow C_{i+6}[2, 5, 10, 11, 14] &= (\langle 0, 2, 2p+2q-1, 2p-2, 2q+1 \rangle, 0) \end{aligned}$$

If $r > 2$:

$$\begin{aligned} C_i[2, 5, 10, 11, 14] &= (\langle 2, 0, 2p+2q+1, 2p+1, 2q \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 16] &= (\langle 1, 2, 0, 2p+2q-2, 2p-2, 1, 0, 2q, 0, 0 \rangle, 0) \\ \Rightarrow C_{i+2}[1, 2, 5, 10 : 16] &= (\langle r+1, 2, 0, 2p+2q-2, 2p-2, 0, 1, 2q, 0, 0 \rangle, 0) \\ \Rightarrow' C_{i+3}[1 : 5, 13] &= \langle r, 3, 0, 0, 1, r+1 \rangle \\ \Rightarrow' C_{i+r+2}[1 : 5, 13] &= \langle 0, 4, 2, 1, 2, 2 \rangle \\ \Rightarrow' C_{i+r+3}[1 : 5, 11] &= \langle 0, 0, 4, 2, 2, 0 \rangle \\ \Rightarrow C_{i+r+4}[2, 5, 10, 11, 14] &= (\langle 0, 2, 2p+2q-1, 2p-2, 2q+1 \rangle, 0) \end{aligned}$$

For all cases, lemma 7 holds.

LEMMA 8. In Δ' , for any configuration $C_i[2, 5, 10, 11, 14] = (\langle 0, 2, 2p+2y+1, 2p, 2q+1 \rangle, 0)$, if $q \geq 1$, then r spikes, for some $1 \leq r \leq k$, will be requested by $\sigma_{1,1}$ after C_{i+1} and by $\sigma_{3,3}$ after C_{i+2} , and $C_i \Rightarrow^* C_{i+r+4}[2, 5, 10, 11, 14] = (\langle 2, 0, 2p+2q-1, 2p+1, 2q-2 \rangle, 0)$.

Proof: Let $q \geq 1$ and r be the spikes requested by $\sigma_{1,1}$ or $\sigma_{3,3}$ if a requesting rule is applied to the neurons.

If $r = 1$:

$$\begin{aligned} C_i[2, 5, 10, 11, 14] &= (\langle 0, 2, 2p+2q+1, 2p, 2q+1 \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 16] &= (\langle 1, 0, 2, 2p+2q-2, 2p, 0, 0, 2q-2, 1, 0 \rangle, 0) \\ \Rightarrow C_{i+2}[1, 2, 5, 10 : 16] &= (\langle 2, 0, 2, 2p+2q-2, 2p, 0, 0, 2q-2, 0, 1 \rangle, 0) \\ \Rightarrow' C_{i+3}[1, 2, 5 : 7, 16] &= \langle 0, 2, 4, 0, 0, 2 \rangle \\ \Rightarrow' C_{i+4}[1, 2, 5 : 7, 16] &= \langle 0, 2, 0, 4, 2, 0 \rangle \\ \Rightarrow C_{i+5}[2, 5, 10, 11, 14] &= (\langle 2, 0, 2p+2q-1, 2p+1, 2q-2 \rangle, 0) \end{aligned}$$

If $r = 2$:

$$\begin{aligned} C_i[2, 5, 10, 11, 14] &= (\langle 0, 2, 2p+2q+1, 2p, 2q+1 \rangle, 0) \\ \Rightarrow C_{i+1}[1, 2, 5, 10 : 16] &= (\langle 1, 0, 2, 2p+2q-2, 2p, 0, 0, 2q-2, 1, 0 \rangle, 0) \\ \Rightarrow C_{i+2}[1, 2, 5, 10 : 16] &= (\langle 3, 0, 2, 2p+2q-2, 2p, 0, 0, 2q-2, 0, 1 \rangle, 0) \\ \Rightarrow' C_{i+3}[1, 2, 5 : 7, 16] &= \langle 2, 1, 3, 0, 0, 3 \rangle \\ \Rightarrow' C_{i+4}[1, 2, 5 : 7, 16] &= \langle 0, 2, 4, 0, 0, 2 \rangle \\ \Rightarrow' C_{i+5}[1, 2, 5 : 7, 16] &= \langle 0, 2, 0, 4, 2, 0 \rangle \\ \Rightarrow C_{i+6}[2, 5, 10, 11, 14] &= (\langle 2, 0, 2p+2q-1, 2p+1, 2q-2 \rangle, 0) \end{aligned}$$




Figure 2: The SNdP system for the proof of Proposition 1

If $r > 2$:

$$\begin{aligned}
 C_i[2, 5, 10, 11, 14] &= (\langle 0, 2, 2p + 2q + 1, 2p, 2q + 1 \rangle, 0) \\
 \Rightarrow C_{i+1}[1, 2, 5, 10 : 16] &= (\langle 1, 0, 2, 2p + 2q - 2, 2p, 0, 0, 2q - 2, 1, 0 \rangle, 0) \\
 \Rightarrow C_{i+2}[1, 2, 5, 10 : 16] &= (\langle r + 1, 0, 2, 2p + 2q - 2, 2p, 0, 0, 2q - 2, 0, 1 \rangle, 0) \\
 \Rightarrow' C_{i+3}[1, 2, 5 : 7, 16] &= \langle r, 1, 3, 0, 0, r + 1 \rangle
 \end{aligned}$$

for $1 \leq j \leq r - 2$

$$\begin{aligned}
 \Rightarrow' C_{i+j+3}[1, 2, 5 : 7, 16] &= \langle r - j, 1, 3, 2, 1, r - j + 1 \rangle \\
 \Rightarrow' C_{i+r+2}[1, 2, 5 : 7, 16] &= \langle 0, 2, 4, 0, 0, 2 \rangle \\
 \Rightarrow' C_{i+r+3}[1, 2, 5 : 7, 16] &= \langle 0, 2, 0, 4, 2, 0 \rangle \\
 \Rightarrow C_{i+r+4}[2, 5, 10, 11, 14] &= (\langle 2, 0, 2p + 2q - 1, 2p + 1, 2q - 2 \rangle, 0)
 \end{aligned}$$

For all cases, lemma 8 holds.

The starting configuration of lemma 7 and the resulting configuration of lemma 8 are similar while the resulting configuration of lemma 7 and the starting configuration of lemma 8 are similar. This results in a possible loop.

LEMMA 9. In Δ' , for any configuration $C_i[2, 5, 10, 11, 14] = (\langle 2, 0, 2p + 2q + 1, 2p + 1, 2q \rangle, 0)$, if $p, q \geq 1$, $C_i \Rightarrow^* C_{i+r_1+r_2+8}[2, 5, 10, 11, 14] = (\langle 2, 0, 2p + 2q - 3, 2p - 1, 2q - 2 \rangle, 0)$ where for some $1 \leq r_1, r_2 \leq k$, neuron $\sigma_{1,1}$ would have requested, in order, r_1, r_2 spikes, $\sigma_{2,3}$ would have requested r_1 spikes, and $\sigma_{3,3}$ would have requested r_2 spikes.

Proof: Let $p, q \geq 1$. From lemma 7, $C_i[2, 5, 10, 11, 14] = (\langle 2, 0, 2p + 2q + 1, 2p + 1, 2q \rangle, 0)$ yields to $C_{i+r_1+4}[2, 5, 10, 11, 14] = (\langle 0, 2, 2p + 2q -$

$1, 2p - 2, 2q + 1 \rangle, 0)$, for some r_1 , while r_1 spikes is requested by $\sigma_{1,1}$ after C_{i+1} and by $\sigma_{2,3}$ after C_{i+2} . Using lemma 8, $C_{i+r_1+4}[2, 5, 10, 11, 14] = (\langle 0, 2, 2p + 2q - 1, 2p - 2, 2q + 1 \rangle, 0)$ yields to $C_{i+r_1+r_2+8}[2, 5, 10, 11, 14] = (\langle 2, 0, 2p + 2q - 3, 2p - 1, 2q - 2 \rangle, 0)$, for some r_2 , while r_2 spikes is requested by $\sigma_{1,1}$ after C_{i+r_1+5} and by $\sigma_{3,3}$ after C_{i+r_1+6} .

LEMMA 10. In Δ' , for any configuration $C_i[2, 5, 10, 11, 14] = (\langle 0, 2, 2p + 2y + 1, 2p, 2q + 1 \rangle, 0)$, if $p, q \geq 1$, $C_i \Rightarrow^* C_{i+r_1+r_2+8}[2, 5, 10, 11, 14] = (\langle 0, 2, 2p + 2y - 3, 2p - 2, 2q - 1 \rangle, 0)$ where for some $1 \leq r_1, r_2 \leq k$, neuron $\sigma_{1,1}$ would have requested, in order, r_1, r_2 spikes, $\sigma_{3,3}$ would have requested r_1 spikes, and $\sigma_{2,3}$ would have requested r_2 spikes.

Proof: Let $p, q \geq 1$. From lemma 8, $C_i[2, 5, 10, 11, 14] = (\langle 0, 2, 2p + 2y + 1, 2p, 2q + 1 \rangle, 0)$ yields to $C_{i+r_1+4}[2, 5, 10, 11, 14] = (\langle 2, 0, 2p + 2q - 1, 2p + 1, 2q - 2 \rangle, 0)$, for some r_1 , while r_1 spikes is requested by $\sigma_{1,1}$ after C_{i+1} and by $\sigma_{3,3}$ after C_{i+2} . Using lemma 7, $C_{i+r_1+4}[2, 5, 10, 11, 14] = (\langle 2, 0, 2p + 2q - 1, 2p + 1, 2q - 2 \rangle, 0)$ yields to $C_{i+r_1+r_2+8}[2, 5, 10, 11, 14] = (\langle 0, 2, 2p + 2y - 3, 2p - 2, 2q - 1 \rangle, 0)$, for some r_2 , while r_2 spikes is requested by $\sigma_{1,1}$ after C_{i+r_1+5} and by $\sigma_{2,3}$ after C_{i+r_1+6} .

Since $x = \lceil \frac{|w|}{2} \rceil$ and $y = |w| - x$, then there are two cases: $|w|$ is even which results to $x = y$, or $|w|$ is odd which results to $x = y + 1$. The configuration loop can be repeated until $\sigma_{2,1}$ and $\sigma_{3,1}$ contain zero or one spike.

LEMMA 11. In Δ' , if $x = y$, then $C_1[2, 5, 10, 11, 14] = (\langle 2, 0, 2x + 2y + 1, 2x + 1, 2y \rangle, 0) \Rightarrow^* C_f[2, 5, 10, 11, 14] = (\langle 2, 0, 1, 1, 0 \rangle, 0)$ where $\sigma_{1,1}$ would have requested, in order, $r_1, r_2, \dots, r_n$ spikes, $\sigma_{2,3}$ would have requested, in order, $r_1, r_3, \dots, r_{n-1}$

spikes, and $\sigma_{3,3}$ would have requested, in order, $r_2, r_4, \dots, r_n$ spikes, and $f = 1 + 4n + \sum_{i=1}^n r_i$, for some $1 \leq r_1, r_2, \dots, r_n \leq k$.

Proof: Since $x = y$, which also implies $n = 2x$, then $C_1[2, 5, 10, 11, 14] = (\langle 2, 0, 4x + 1, 2x + 1, 2x \rangle, 0)$. From C_1 , for $1 \leq i \leq x$: apply lemma 9 where neuron $\sigma_{1,1}$ requests, in order, r_{2i-1}, r_{2i} spikes, $\sigma_{2,3}$ requests r_{2i-1} spikes, and $\sigma_{3,3}$ requests r_{2i} spikes. This yields to $C_f[2, 5, 10, 11, 14] = (\langle 2, 0, 1, 1, 0 \rangle, 0)$ where $\sigma_{1,1}$ requested, in order, $r_1, r_2, \dots, r_{2x}$ spikes, $\sigma_{2,3}$ requested, in order, $r_1, r_3, \dots, r_{2x-1}$ spikes, and $\sigma_{3,3}$ requested, in order, $r_2, r_4, \dots, r_{2x}$ spikes, and $f = 1 + 8x + \sum_{i=1}^{2x} r_i$. Since $n = 2x$, then the result is $C_f[2, 5, 10, 11, 14] = (\langle 2, 0, 1, 1, 0 \rangle, 0)$ where $\sigma_{1,1}$ requested, in order, $r_1, r_2, \dots, r_n$ spikes, $\sigma_{2,3}$ requested, in order, $r_1, r_3, \dots, r_{n-1}$ spikes, and $\sigma_{3,3}$ requested, in order, $r_2, r_4, \dots, r_n$ spikes, and $f = 1 + 4n + \sum_{i=1}^n r_i$.

LEMMA 12. *In Δ' , if $x = y + 1$, then $C_1[2, 5, 10, 11, 14] = (\langle 2, 0, 2x + 2y + 1, 2x + 1, 2y \rangle, 0) \Rightarrow^* C_f[2, 5, 10, 11, 14] = (\langle 0, 2, 1, 0, 1 \rangle, 0)$ where $\sigma_{1,1}$ would have requested, in order, $r_1, r_2, \dots, r_n$ spikes, $\sigma_{2,3}$ would have requested, in order, $r_1, r_3, \dots, r_n$ spikes, and $\sigma_{3,3}$ would have requested, in order, $r_2, r_4, \dots, r_{n-1}$ spikes, and $f = 1 + 4n + \sum_{i=1}^n r_i$, for some $1 \leq r_1, r_2, \dots, r_n \leq k$.*

Proof: Since $x = y + 1$, which also implies $n = 2y + 1$, then $C_1[2, 5, 10, 11, 14] = (\langle 2, 0, 4y + 3, 2y + 3, 2y \rangle, 0)$. From C_1 , apply lemma 7. This yields to $C_{r_1+5}[2, 5, 10, 11, 14] = (\langle 0, 2, 4y + 1, 2y, 2y + 1 \rangle, 0)$ where r_1 spikes is requested by $\sigma_{1,1}$ after C_1 and by $\sigma_{2,3}$ after C_2 . If $y = 0$, then $C_{r_1+5}[2, 5, 10, 11, 14] = (\langle 0, 2, 1, 0, 1 \rangle, 0)$, for which lemma 12 holds. Otherwise, from C_{r_1+5} , for $1 \leq i \leq y$: apply lemma 10 where neuron $\sigma_{1,1}$ requests, in order, r_{2i}, r_{2i+1} spikes, $\sigma_{3,3}$ requests r_{2i} spikes, and $\sigma_{2,3}$ requests r_{2i+1} spikes. This yields to $C_f[2, 5, 10, 11, 14] = (\langle 0, 2, 1, 0, 1 \rangle, 0)$ where $\sigma_{1,1}$ requested, in order, $r_2, r_3, \dots, r_{2y+1}$ spikes, $\sigma_{2,3}$ requested, in order, $r_3, r_5, \dots, r_{2y+1}$ spikes, and $\sigma_{3,3}$ requested, in order, $r_2, r_4, \dots, r_{2y}$ spikes, and $f = 5 + 8y + \sum_{i=1}^{2y+1} r_i$. Compounding with the result from C_1 to C_{r_1+5} , in total, $\sigma_{1,1}$ requested, in order, $r_1, r_2, \dots, r_{2y+1}$ spikes, $\sigma_{2,3}$ requested, in order, $r_1, r_3, \dots, r_{2y+1}$ spikes, and $\sigma_{3,3}$ requested, in order, $r_2, r_4, \dots, r_{2y}$ spikes. Since $n = 2y + 1$, then the result, starting from C_1 , is $C_f[2, 5, 10, 11, 14] = (\langle 0, 2, 1, 0, 1 \rangle, 0)$ where $\sigma_{1,1}$ requested, in order, $r_1, r_2, \dots, r_n$ spikes, $\sigma_{2,3}$ requested, in order, $r_1, r_3, \dots, r_n$ spikes, and $\sigma_{3,3}$ requested, in order, $r_2, r_4, \dots, r_{n-1}$ spikes, and $f = 1 + 4n + \sum_{i=1}^n r_i$.

The starting configuration of Δ' is $C_0[2, 7, 10, 11, 14] = (\langle 2, 2, 2x + 2y, 2x, 2y \rangle, 0)$. This yields to $C_1[2, 5, 10, 11, 14] = (\langle 2, 0, 2x + 2y + 1, 2x + 1, 2y \rangle, 0)$.

If n is even, applying lemma 11 with spike requests r_i for all $1 \leq i \leq n$, $x_i = b_{r_i}$ results in $C_f[2, 5, 10, 11, 14] = (\langle 2, 0, 1, 1, 0 \rangle, 0)$ where $f = 1 + 4n + \sum_{i=1}^n r_i$, where ww is taken as an input. From here, a transition yields $C_{f+1}[2] = (\langle 2 \rangle, 0)$. From C_{f+1} , there are no more applicable rules. This means that $C_{f+1} = C_h$. The system Δ' reaches a halting configuration given an input ww .

If n is odd, applying lemma 12 with spike requests r_i for all $1 \leq i \leq n$, $x_i = b_{r_i}$ results in $C_f[2, 5, 10, 11, 14] = (\langle 0, 2, 1, 0, 1 \rangle, 0)$ where $f = 1 + 4n + \sum_{i=1}^n r_i$, where ww is taken as an input. From here, a transition yields $C_{f+1}[5] = (\langle 2 \rangle, 0)$. From C_{f+1} , there are no more applicable rules. This means that $C_{f+1} = C_h$. The system Δ' reaches a halting configuration given an input ww .

In either case, Δ' reaches a halting configuration given an input ww . $\square$

The modified SNdP has the same inter-component communication behaviors to the 3-component nonbalanced nonhomogeneous SNdP in [4]. We have $ComN(\Delta') = ComR(\Delta') = (k + 1)n$, and $ComW(\Delta') = (k + 2)n$.

On the other hand, the halting computation is C_{f+1} where $f = 1 + 4n + \sum_{i=1}^n r_i$. Since $1 \leq r_i \leq k$, then the maximum steps to reach a halting computation is $kn + 4n + 2$.

4 Final Remarks

In this work, we mathematically proved, using new notations we defined, that the 3-component, nonbalanced and nonhomogeneous SNdP system in [4] is unable to recognize L_{ww} , as it continues to request an input outside that of ww . We also presented a modification of this SNdP system, which we prove its correctness in recognizing L_{ww} . The modification includes the addition of three neurons: a neuron in Π_1 that acts as a counter similar to that of the existing ones in Π_2 and Π_3 , and a neuron each in Π_2 and Π_3 that adds a timestep to the flow of spikes in the components to compensate for the delay caused by the additional neuron in Π_1 . This modification does not change the communication costs of the original SNdP; however, it adds a total of $n + 2$ steps to the running time.

Other dP systems have been used for solving NP-complete problems, such as in [1–3]. This paper, as well as in [4, 5], mainly looked into SNdP systems recognizing L_{ww} . For future works, we could look into SNdP systems that can solve NP-complete problems such as 3-SAT.

Acknowledgments

All authors acknowledge the support from the University of the Philippines Diliman College of Engineering.

References

- [1] Henry N Adorna, Linqiang Pan, and Bosheng Song. 2018. On distributed solution to SAT by membrane computing. *INTERNATIONAL JOURNAL OF COMPUTERS COMMUNICATIONS & CONTROL* 13, 3 (2018), 303–320.
- [2] Kelvin Buño and Henry Adorna. 2020. Distributed computation of ak P systems with active membranes for SAT using clause completion. *Journal of Membrane Computing* 2, 2 (2020), 108–120.
- [3] Kelvin C Buño, Francis George C Cabarle, Marj Darrel Calabia, and Henry N Adorna. 2018. Solving the N-Queens problem using dP systems with active membranes. *Theoretical Computer Science* 736 (2018), 1–14.
- [4] Kelvin Buño, Francis George Cabarle, and Jherico Gabriel Torres. 2019. Spiking Neural dP Systems: Balance and Homogeneity. *Philippine Computing Journal* 14 (12 2019), 1–10. Issue 2.
- [5] Mihai Ionescu, Gheorghe Păun, Mario J Pérez-Jiménez, and Takashi Yokomori. 2011. Spiking neural dP systems. *Fundamenta Informaticae* 111, 4 (2011), 423–436.
- [6] Mihai Ionescu, Gheorghe Păun, and Takashi Yokomori. 2006. Spiking neural P systems. *Fundamenta Informaticae* 71, 2-3 (2006), 279–308.
- [7] Gheorghe Păun. 2002. *Membrane computing: an introduction*. Springer Science & Business Media.
- [8] Gheorghe Păun and Mario J Pérez-Jiménez. 2010. Solving problems in a distributed way in membrane computing: DP systems. *International Journal of Computers Communications & Control* 5, 2 (2010), 238–250.
- [9] Gheorghe Păun and Mario J Pérez-Jiménez. 2012. An infinite hierarchy of languages defined by dP systems. *Theoretical Computer Science* 431 (2012), 4–12.

Apparent Age Estimation: Challenges and Outcomes

A Comparative Analysis on Apparent Age Estimation Methods and Datasets

Justin Rainier Go
justin_rainier_go@dlsu.edu.ph
De La Salle University
Manila, Philippines

Lorenz Bernard Marqueses
lorenz_marqueses@dlsu.edu.ph
De La Salle University
Manila, Philippines

Mikaella Kaye Martinez
mikaella_kaye_martinez@dlsu.edu.ph
De La Salle University
Manila, Philippines

John Kevin Patrick Sarmiento
john_kevin_sarmiento@dlsu.edu.ph
De La Salle University
Manila, Philippines

Abien Fred Agarap
abien.agarap@dlsu.edu.ph
De La Salle University
Manila, Philippines

Abstract

Apparent age estimation is a valuable tool for business personalization, yet current models frequently exhibit demographic biases. We review prior works on the DEX method by applying distribution learning techniques such as Mean-Variance Loss (MVL) and Adaptive Mean-Residue Loss (AMRL), and evaluate them in both accuracy and fairness. Using IMDB-WIKI, APPA-REAL, and FairFace, we demonstrate that while AMRL achieves state-of-the-art accuracy, trade-offs between precision and demographic equity persist. Despite clear age clustering in UMAP embeddings, our saliency maps indicate inconsistent feature focus across demographics, leading to significant performance degradation for Asian and African American populations. We argue that technical improvements alone are insufficient; accurate and fair apparent age estimation requires the integration of localized and diverse datasets, and strict adherence to fairness validation protocols.

CCS Concepts

• **Computing methodologies** → **Supervised learning by regression; Supervised learning by classification; Biometrics; Interest point and salient region detections; Transfer learning;** • **Social and professional topics** → **Age; Race and ethnicity; Gender;** • **Applied computing** → **Marketing; Digital libraries and archives.**

Keywords

apparent age, IMDB-WIKI dataset, deep expectation, multinomial regression, mean-variance loss, adaptive mean-residue loss

1 Introduction and Related Works

Facial age estimation is a primary focus in computer vision and serves various functions in security, social media, and interactive systems. A specific variation of this problem that researchers have studied less frequently is apparent age estimation. This represents the perceived age rather than the actual birth age of a person. Apparent age estimation is particularly relevant for industries where perceived physical maturity is the central concern such as in the cosmetics industry, and the medical field [6, 13].

More precisely, there are measurable commercial benefits towards personalization including age-related personalization. Similar studies show that such personalization has yielded a 20–35% increase in Average Order Value, higher conversion rates due to more

relevant recommendations, improved Customer Lifetime Value, and reduced Customer Acquisition Cost through more precise ad targeting [5, 6]. Given the clear business advantages, apparent age estimation represents a promising research direction. This problem gained significant prominence during the 2015 ChaLearn Looking at People (CLAP) Challenge [4], followed by the introduction of other datasets such as IMDB-WIKI [12], APPA-REAL [1], and FairFace [9]. In the following subsections, we briefly discuss related works that were trained and evaluated on the aforementioned datasets.

1.1 DEX

Rothe et al. [12] introduced the Deep Expectation (DEX) method which uses a VGG-16 architecture pretrained on the ImageNet dataset as a starting point using a standard cross-entropy loss function. This was then finetuned on the IMDB-WIKI and CLAP datasets¹, although a control group that did not train on IMDB-WIKI was also tested. DEX framed its age estimation problem as a classification task using multinomial regression, $E(O) = \sum_{i=0}^{n=100} y_i o_i$, representing ages from 1 to 101 to produce a predicted age.

1.2 Mean-Variance Loss

Unlike DEX, the mean-variance loss [10] approaches age estimation via distribution learning to capture the correlation between adjacent ages. Their loss function consists of a mean loss that minimizes the distance between the expected value of the predicted distribution and the ground-truth age, and a variance loss that penalizes the spread of distribution to ensure a sharp and concentrated prediction. By jointly optimizing these two losses along a softmax cross entropy loss, they demonstrate that their model achieves superior performance across multiple benchmark datasets when compared to DEX.

1.3 Adaptive Mean-Residue Loss

The adaptive mean-residue loss [14] addresses the limitations of conventional distribution-based learning by first estimating a coarse age value and then adaptively calculating a residual value to adjust its prediction towards the actual ground-truth age. This two-step mechanism allows a network to better handle variances found in

¹An ensemble of 20 models, each with slightly different finetunes on the CLAP dataset, was used to achieve first place in the 2015 CLAP challenge. However, Rothe et al. also reported accuracy metrics when using a single CNN. We focus on these results for simplicity.

facial appearances across different age groups. With this, the model was able to outperform existing age estimation models at the time with an MAE of 3.61 compared to an MAE of 3.95 through the mean-variance loss.

1.4 Contributions

We summarize our main contributions as follows:

- (1) We evaluate model performance trained on IMDB-WIKI using a modified version of the DEX methodology that fine-tunes on different combinations of the CLAP, APPA-REAL, and FairFace datasets.
- (2) We evaluate model performance across demographics (i.e. sex and race), and we identify the degree to which this imbalance can affect the accuracy of both real and apparent age predictions per group.
- (3) We explore applications of apparent age estimation in cosmetics, marketing, healthcare, and security to improve personalization and efficiency, and we briefly discuss concerns regarding ethics, privacy, and data governance.
- (4) Finally, we lay down some future directions on Philippine facial age estimation preceded with an evaluation on a small-scale dataset.

2 Exploring Age Estimation datasets

2.1 Datasets Description

2.1.1 IMDB-WIKI² *pedia*. This dataset contains images of celebrities scraped from IMDb and Wikipedia and is labeled by offsetting the celebrity’s birth date by the timestamp of each image. Initial attempts at model replication using this dataset yielded a higher MAE. We elected to use a cleaned version of the dataset during training² as it appeared that the publicly-available dataset contained anomalous entries, some of which can be seen in Figure 1. The IMDB-WIKI dataset has been noted to have an imbalanced gender ratio of 14:10 in favor of males [11].

2.1.2 *ChLearn Looking at People (CLAP)*. This was a fairly small dataset created by crowd-sourcing data through an online website [4]. It was then augmented by the AgeGuess platform [8] which collected much of the same data, yielding a final dataset of 4,691 images with a distribution of over 140,000 votes. Each of these images was annotated with both the real age as well as the mean and variance of their voted apparent ages for use with recognition.

²The specific cleaned version of the dataset was retrieved from the following link: github.com/imdeepmind/processed-IMDB-WIKI-dataset.



Figure 1: Examples of non-facial images in IMDB-WIKI. Note the presence of pixel-stretching artifacts.




Figure 2: APPA-REAL demographic distribution. The Caucasian group is highly over-represented.




Figure 3: IMDB-WIKI and APPA-REAL age distributions

This dataset was the only one in the original DEX pipeline that explicitly annotates its images with apparent age.

2.1.3 APPA-REAL. The APPA-REAL dataset [1] was the first fairly large dataset to include annotations for both real and apparent age as well as demographic information such as race and gender. Similar to the CLAP datasets, data was collected through crowdsourcing platforms. APPA-REAL’s real and apparent age distributions (as seen in Figure 2) are also quite similar to one another with a KL divergence of $D_{KL} \approx 0.01322$. The APPA-REAL dataset has a similar imbalance to IMDB-WIKI as per Figure 2, the APPA-REAL dataset is dominated by Caucasian faces, significantly outnumbering Asian and especially African American faces. Figure 3 shows that despite the size difference, the real-age distributions across both IMDB-WIKI and APPA-REAL are somewhat similar, with most of the ages lying within the 20-40 range.

2.1.4 FairFace. The authors claim a fairer distribution of faces in the FairFace [9] dataset, though it is still to a lesser extent dominated by Caucasian faces. It has more race classes, though both male and female groups have roughly as many samples (except for Middle Eastern samples which have much fewer female samples). However, unlike the other datasets discussed in this paper, FairFace contains no explicit apparent age information; only an age range for each image. We take the mean of this age range as the ground-truth value during training which may be sufficient for DEX’s multinomial regression. Figure 4 demonstrates that each race-gender pair is much more equally-represented across age groups and that the distribution of age groups themselves appear to be more balanced. However, there were still far fewer samples from the older age groups (aged 56 to 70).




Figure 4: FairFace demographic distribution

2.2 Evaluation Metrics

We evaluate models through their mean absolute error (MAE) values. Additionally, we use the ϵ -error to account for the degree of uncertainty when estimating the apparent age. Formally defined as $1 - \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$, the ϵ -error fits a normal distribution based on the mean and standard deviation of the collected user guesses for a given image [4]. Results shown in Table 1 are presented as they were reported in their respective papers, while results from our reproduction can be found in Tables 2 and 3 – all in Section 3.

2.3 Modelling

We finetuned on the following combinations of datasets:

- (1) IMDB-WIKI;
- (2) IMDB-WIKI, then CLAP;
- (3) IMDB-WIKI, then APPA-REAL;
- (4) IMDB-WIKI, then FairFace;
- (5) IMDB-WIKI, then FairFace, then CLAP; and
- (6) IMDB-WIKI, then FairFace, then APPA-REAL.

At the same time, we also attempt replacing the cross-entropy loss (CEL) function of the original DEX method with more recent loss functions, specifically, mean-variance loss and adaptive mean-residue loss. With six combinations of datasets and three choices of loss functions, we end up with 18 different models to evaluate. To evaluate the models, we continue to use the model evaluation metrics in Section 2.2, and generate UMAP embeddings and cosine similarity graphs. Additionally, we produce saliency maps to explain which regions of each image held the most weight when predicting the apparent ages of people, grouped by race and gender, with one image for each age group. The representative image was the closest to the midpoint of each age group.

2.4 Bias and fairness assessment

Imbalances in sample distribution across race and gender can be seen most strongly in the APPA-REAL dataset. As we want to measure the variance in accuracy of the overall models based on training data demographics, we kept the datasets as-is, although we hypothesized that the models finetuned on FairFace would experience a smaller degree of racial and sexual sampling biases due to its greater diversity.

2.5 Software/Hardware configuration

Rothe et al. [12]’s software setup utilized Caffe, an open-source deep learning framework. Our replication of the model and the further variants use PyTorch Lightning with Python versions ≥ 3.11 , with our repositories using uv for project management. Due to computational and time constraints, we used Kaggle’s NVIDIA P100 for training and finetuning on IMDB-WIKI and FairFace while we used NVIDIA RTX 3060 and 4060 GPUs on CLAP and APPA-REAL respectively.

3 Evaluating Age Estimation models

We present our validation of the reported performance trend of previous works, namely DEX [12], MVL [10], and AMRL [14]³. Our results corroborate the general trend of the newer approaches outperforming earlier ones.

Table 1 shows lifted results from literature [10, 12, 14]. Table 2 summarizes the performances of the recreated models on both the APPA-REAL and CLAP testing datasets. Table 3 shows how this performance varies across different groups of race and gender. The evaluation reveals performance disparities across demographic intersections. The models exhibit the highest error rates on African American and Asian female subjects while they perform better on African American male counterparts. Interestingly, this latter group

³The code used to produce our results are stored in a group of repositories hosted in GitLab, accessible through the following link:

<https://gitlab.com/data100-s12-group7/imdb-wiki/>.

Table 1: Comparison of DEX, MVL, and AMRL loss as reported by each paper

	[12]		[10]		[14]	
	CLAP 2015		FG-NET, CLAP 2016		FG-NET, CLAP 2016	
	MAE	ϵ	MAE	ϵ	MAE	ϵ
DEX	3.22	0.28	3.09	-	4.63	-
MVL	-	-	2.68	0.29	3.95	0.40
AMRL	-	-	-	-	3.61	0.39

Table 2: MAE and ϵ -error results. These results do not yet include the FairFace finetuning step as we attempt to reproduce the results of [12] while comparing CLAP and APPA-REAL. Star symbol (*) indicates Caffe model from [12]. Plus symbol (+) indicates our own implementation.

Testing set	Model	MAE		ϵ -error
		Appa.*	Real	
APPA-REAL	CLAP ⁺	12.57	11.20	0.88
	IMDB-WIKI + CLAP*	5.18	6.84	0.38
	IMDB-WIKI*	6.96	7.89	0.49
	IMDB-WIKI + APPA-REAL ⁺	4.32	6.26	0.34
	IMDB-WIKI + CLAP ⁺	5.51	7.04	0.41
	IMDB-WIKI ⁺	5.93	7.19	0.43
CLAP	CLAP ⁺	-	6.64	0.74
	IMDB-WIKI + CLAP*	-	3.25	0.28
	IMDB-WIKI*	-	5.66	0.49
	IMDB-WIKI + APPA-REAL ⁺	-	3.77	0.34
	IMDB-WIKI + CLAP ⁺	-	3.57	0.32
	IMDB-WIKI ⁺	-	4.90	0.44

presents the most substantial divergence between apparent and real age. These results suggest that the performance gap is a direct consequence of the dataset distribution which is heavily skewed towards male subjects and thus lacks the necessary representation to ensure fairer performance across demographics.


Figures 6 and 7 show UMAP-reduced image embeddings for three models of interest when tested on the APPA-REAL dataset: the model trained only on the CLAP train set, the model trained on only IMDB-WIKI, and the model trained on IMDB-WIKI and finetuned on the APPA-REAL train set. The model trained solely on the CLAP dataset performed the worst, as it learned considerably fewer facial representations compared to other models. However, there still appear to be some prominent clusters in the reduced embedding space, particularly with younger ages (purple hues).

The model finetuned on the APPA-REAL training set had better results, showing clearer distinction for both younger (yellow-green hues) and older ages. On the other hand, it performs worse with “middle” age groups, despite the age distribution; attributable to strong feature overlap between individuals in these age groups. With APPA-REAL, we find the similar trend but with visibly better clustering. We acknowledge that these are qualitative results, and we intend to present quantitative metrics (e.g. mean discrepancies between age groups) to support these findings in a future work.

Figure 5 shows the cosine similarity of sample images with the average embeddings for its respective age group for the model

Table 3: MAE and ϵ -error results per model on APPA-REAL, grouped by race and gender, ordered by MAE (apparent age). We present results from the 5 best-performing models on apparent age. The MAE values for the best and worst groups per model are shown in bold.

Model	Race	Gender	MAE		ϵ
			Apparent	Real	
AMRL: IMDB-WIKI + APPA-REAL	African American	Female	2.93	5.15	0.27
		Male	2.92	6.45	0.23
	Asian	Female	4.33	6.41	0.38
		Male	3.77	4.92	0.36
	Caucasian	Female	4.08	6.14	0.32
		Male	3.53	5.00	0.28
AMRL: IMDB-WIKI + FairFace + APPAREAL	African American	Female	4.60	6.24	0.36
		Male	3.85	7.14	0.31
	Asian	Female	4.62	6.91	0.40
		Male	3.40	5.34	0.33
	Caucasian	Female	4.26	6.45	0.34
		Male	3.82	5.17	0.31
CEL: IMDB-WIKI + FairFace + APPA-REAL	African American	Female	4.62	6.21	0.37
		Male	2.92	6.38	0.22
	Asian	Female	4.53	6.84	0.36
		Male	4.46	6.17	0.40
	Caucasian	Female	4.48	6.85	0.35
		Male	4.18	5.50	0.33
MVL: IMDB-WIKI + APPA-REAL	African American	Female	3.30	4.17	0.27
		Male	2.89	5.85	0.24
	Asian	Female	4.65	7.24	0.36
		Male	4.17	5.73	0.42
	Caucasian	Female	4.34	6.11	0.34
		Male	3.49	4.91	0.29
MVL: IMDB-WIKI + FairFace + APPA-REAL	African American	Female	3.22	4.81	0.26
		Male	3.58	6.77	0.27
	Asian	Female	4.65	6.77	0.41
		Male	3.89	5.50	0.37
	Caucasian	Female	4.06	6.26	0.31
		Male	3.53	5.03	0.29

**Figure 5: Cosine similarity of each image with average embeddings for age, grouped by race**

finetuned on only APPA-REAL. Notably, despite good model performance, the distribution of similarity values appears to be spread out. We posit this is because while the classifier can select the best age class for an image, this does not guarantee that it is highly similar to the average embeddings for that class.

Table 4: MAE results per model ordered by MAE (apparent). We show the top 10 best-performing models in terms of apparent age MAE. Best model results are shown in bold.

Model	MAE (apparent)		MAE (real)		ϵ	
	mean	std	mean	std	mean	std
Adaptive mean-residue loss: IMDB-WIKI + APPA-REAL	3.59	0.58	5.68	0.73	0.31	0.06
Mean-variance loss: IMDB-WIKI + APPA-REAL	3.81	0.68	5.67	1.05	0.32	0.07
Mean-variance loss: IMDB-WIKI + FairFace + APPA-REAL	3.82	0.50	5.86	0.86	0.32	0.06
Adaptive mean-residue loss: IMDB-WIKI + FairFace + APPAREAL	4.09	0.49	6.21	0.81	0.34	0.04
Cross-entropy loss: IMDB-WIKI + FairFace + APPA-REAL	4.20	0.64	6.32	0.50	0.34	0.06
Cross-entropy loss: IMDB-WIKI + APPA-REAL	4.46	0.54	6.47	0.54	0.37	0.06
Adaptive mean-residue loss: IMDB-WIKI + CHALEARN	4.98	0.67	6.70	0.76	0.40	0.05
Original Caffe model on IMDB-WIKI + CHALEARN	4.99	0.57	6.90	0.95	0.39	0.02
Mean-variance loss: IMDB-WIKI + CHALEARN	5.03	0.89	6.78	0.84	0.39	0.06
Mean-variance loss: IMDB-WIKI + FairFace + CHALEARN	5.07	0.67	7.11	0.72	0.41	0.07
Mean-variance loss: IMDB-WIKI + FairFace	5.14	0.61	6.70	0.75	0.42	0.06

*“std” column shows the standard deviation of the MAE values when grouped by race and gender; this summarizes variance in model performance among race-gender pairs.

3.1 Model Performance

Table 4 illustrated that the model utilizing the adaptive mean-residue loss achieves the highest precision when evaluated against the APPA-REAL test set. Specifically, the configuration involving pretraining on IMDB-Wiki followed by finetuning on APPA-REAL yielded the lowest MAE. In contrast, the model incorporating an immediate finetuning stage on the FairFace dataset before the APPA-REAL training exhibited the lowest variance in performance across demographics. These findings suggest that while the inclusion of FairFace dataset does not necessarily improve overall performance, it serves as a paramount step for introducing demographic equity by ensuring more consistent results across diverse populations.

3.2 Visualizing Facial Age Representations

Our experimental results indicate that while all models successfully identify certain demographic clusters, the models trained with the adaptive mean-residue loss yield the best clustering of embeddings. As illustrated in Figure 7, the embedding space reveals two distinct and well-defined clusters. This visualization mirrors patterns observed in Figure 6 where the model demonstrates a high proficiency in distinguishing ages at the extreme ends of the spectrum particularly for younger subjects. Specifically, the AMRL embeddings in Figure 7 highlight a clearly distinguishable purple “island” within the latent space which represents a highly concentrated and separated grouping of specific age features.

3.3 Recognizing Facial Features

Figure 8 presents saliency maps overlaid on input images along the ground-truth and estimated apparent ages to identify the regions that contribute the most to the final model prediction. In these visualizations, redder areas denote higher importance for the decisions. Across the various evaluated models, the features influencing the age estimation remain largely consistent with the most critical regions concentrated on the center of their faces. However, when analyzing specific demographic groups, the model focuses more on peripheral areas such as the forehead or neck. Furthermore, Figure 8a demonstrates that the model can still produce incorrect




Figure 6: UMAP-reduced embeddings on APPA-REAL test set using different finetuning combinations with cross-entropy loss. Left column, top to bottom: CLAP only, IMDB-WIKI + CLAP, and IMDB-WIKI + FairFace + CLAP. Right column replaces CLAP with APPA-REAL.

estimations for younger faces even up to a decade despite highlighting the appropriate facial features. We acknowledge that this is fairly qualitative analysis, and leaves room for a more quantitative approach in a future work.




Figure 7: UMAP-reduced embeddings using different finetuning combinations with MVL (left column) and AMRL (right column). Follows the same progression of APPAREAL-based models as in Figure 6. AMRL with FairFace finetuning (bottom right) appears to result in stronger clustering of age ranges.

3.4 Facial Age Group Similarities

When finetuning via AMRL (as in Figure 9), cosine similarity is much more concentrated towards 1.0 compared to other methods (e.g. as seen in Figure 5, whose cosine similarities are more spread out despite the accuracy). This is possibly because the adaptive mean-residue loss function doesn't necessarily necessitate that predicted ages are highly similar to the target age, only that they are most similar to it. Nevertheless, it appears that AMRL-finetuned models appear to achieve both goals.

3.5 Localized Age Estimation

We also evaluated our two leading AMRL models and our CEL replication against a self-annotated dataset of forty Filipino celebrity images. As shown in Table 5, the model finetuned on the FairFace dataset performed the best, mirroring our main findings. Despite the superior performance of the AMRL model over CEL and MVL models, Figure 10 reveals similar anomalies where the model focuses on isolated or non-essential facial regions.

Table 5: MAE results on Filipino celebrity dataset

Model	MAE
CEL: IMDB-WIKI + ChaLearn	10.05
AMRL: IMDB-WIKI + FairFace + APPA-REAL	6.82
AMRL: IMDB-WIKI + APPA-REAL	7.05

4 Business Implications

Apparent age estimation offers substantial commercial utility in sectors requiring advanced personalization and consumer profiling. In the cosmetics industry, brands utilize facial analytics to assess a




Figure 8: Saliency maps indicate feature focus inconsistencies




Figure 9: Cosine similarity of each image with average embeddings for age (trained via AMRL, on IMDB-WIKI, then APPA-REAL, grouped by race)




Figure 10: Saliency maps on Filipino celebrities similarly show inconsistent feature focus

consumer's perceived appearance, facilitating data driven skincare recommendations tailored to specific aging markers. Furthermore, in clinical and dermatological contexts, perceived age analysis supports the identification of lifestyle-related aging, dermatological pathologies, and stress-induced changes. A notable example is the FaceAge system which utilizes facial data to estimate biological age and assist in health prognostication [2].

In Know-Your-Customer (KYC) systems, apparent age estimation enhances identity verification by validating whether a subject's facial features align with the age bracket indicated on their credentials. This provides an additional layer of fraud detection for

financial institutions, fintech firms, and digital platforms. Furthermore, integrating this technology helps prevent minors from accessing age-restricted services, thereby strengthening institutional compliance and security.

4.1 Inclusivity and demographic representativeness

Several models exhibit substantial performance disparities resulting from dataset imbalances. Specifically, Asian and African American female cohorts consistently yield the highest MAE, whereas male demographics achieve the lowest error rates. Within business applications such as KYC, these inconsistencies heighten the risk of false fraud flags, operational delays, and inequitable account rejections. The overrepresentation of Caucasian subjects in training data skews the learned representations, leading to systematic underperformance for other racial and gender groups.

4.2 Ethical, privacy, data governance issues

Deploying age estimation models within the Philippine context introduces critical challenges regarding bias, privacy, and transparency. Primarily, models trained predominantly on Western datasets often fail to accurately interpret Southeast Asian facial features. These technical inaccuracies risk reinforcing historical prejudices concerning skin tone and beauty standards, particularly within the cosmetics industry which can lead to a significant loss of brand trust.

Furthermore, the Philippine Data Privacy Act of 2012 classifies facial images as sensitive personal information. Many organizations utilize cloud infrastructures that may lack the robust security protocols necessary to prevent data breaches. The frequent lack of transparency regarding data storage and third party access remains a significant legal and ethical hurdle.

Finally, inadequate data governance presents substantial legal risks. Failure to conduct mandatory Privacy Impact Assessments constitutes a violation of existing regulations. Moreover, the utilization of improperly scraped data compromises model integrity and heightens the potential for harm to Filipino users.

4.3 Strategic recommendations for organizations

To mitigate these concerns and avoid legal repercussions, we believe organizations must prioritize the development and integration of localized datasets. It is essential to develop and validate models using facial data representative of Filipino and Southeast Asian populations. This ensures that algorithmic performance remains equitable regardless of skin tone or age.

Such models should be integrated directly into existing infrastructures, including consumer applications and financial security systems. Within the skincare industry, this localization facilitates more precise and personalized product recommendations. For the banking sector, it enhances fraud detection capabilities while ensuring compliance with Bangko Sentral ng Pilipinas regulations. This strategic approach enables companies to monitor model performance effectively, ensuring sustained ethical standards and technical accuracy.

5 Conclusion

The model trained on IMDB-WIKI followed by APPA-REAL utilizing the AMRL loss performed the best in apparent age estimation. Conversely, the configuration that incorporated FairFace in its fine-tuning demonstrated the lowest variance across racial and gender demographics. Overall, the AMRL method proved most effective for this task. UMAP visualizations of AMRL embeddings showed well-defined clusters, while saliency gradients indicated a more consistent focus on facial regions compared to alternative methods. Additionally, cosine similarity analysis for corresponding apparent and real age embeddings tended toward unity under this approach. While this technology offers diverse commercial potential, its adoption remains constrained by pervasive demographic imbalances favoring Caucasian populations, alongside critical ethical, privacy, and data governance challenges.

6 Future works

In addition to the quantitative metrics we intend to pursue in a future study, we lay down some of the more specific and broader research we shall pursue:

- (1) **Contrastive Learning for Few-shot Representation.** We propose investigating contrastive learning techniques to engender age estimation models with few-shot capabilities, particularly for underrepresented demographics. Current models often rely on representations learned from datasets where East Asian populations like Taiwan, Korean, or Japan are better represented than Austronesian groups. Research into how these learned features generalize to Filipino faces could reveal critical performance gaps. Contrastive objectives may provide a more robust way to learn distinct facial features from limited samples, thus ensuring that model performance remains equitable across ethnic groups.
- (2) **Longitudinal Filipino Celebrity Dataset.** A significant hurdle in age estimation research is the scarcity of localized longitudinal data. We intend to extend our limited Filipino celebrity age estimation dataset to a cross-age dataset à la Chen et al. [3]. By creating such a dataset, we can better observe unique physiological aging patterns within the Filipino population. This resource would then help facilitate the development of models that are better calibrated to local features rather than relying on global averages that may not accurately reflect regional facial representations.
- (3) **Optimization for low-resourced compute.** Finally, we suggest exploring contrastive learning to improve the specialization of individual components within a mixture-of-experts [7] architecture. In this framework, different experts can be trained to specialize in specific age brackets or demographic features. Beyond improving accuracy through this specialization, this approach should be evaluated for its potential to enhance inference speed and compute resource utilization. By selectively activating only the most relevant experts for a given input, we can maintain a high throughput with a highly accurate age estimation model.

References

- [1] Eirikur Agustsson, Radu Timofte, Sergio Escalera, Xavier Baro, Isabelle Guyon, and Rasmus Rothe. 2017. Apparent and Real Age Estimation in Still Images with Deep Residual Regressors on Appa-Real Database. In *2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017)*. IEEE, Washington, DC, DC, USA, 87–94. doi:10.1109/FG.2017.20
- [2] Dennis Bontempi, Osbert Zalay, Danielle S Bitterman, Nicolai Birkbak, Derek Shyr, Fridolin Haugg, Jack M Qian, Hannah Roberts, Subha Perni, Vasco Prudente, Suraj Pai, Andre Dekker, Benjamin Haibe-Kains, Christian Guthier, Tracy Balboni, Laura Warren, Monica Krishan, Benjamin H Kann, Charles Swanton, Dirk De Ruysscher, Raymond H Mak, and Hugo J W L Aerts. 2025. FaceAge, a deep learning system to estimate biological age from face photographs to improve prognostication: a model development and validation study. *The Lancet Digital Health* 7, 6 (June 2025).
- [3] Bor-Chun Chen, Chu-Song Chen, and Winston H. Hsu. 2014. Cross-Age Reference Coding for Age-Invariant Face Recognition and Retrieval. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer International Publishing, Cham, 768–783.
- [4] Sergio Escalera, Junior Fabian, Pablo Pardo, Xavier Baro, Jordi Gonzalez, Hugo J. Escalante, Dusan Misevic, Ulrich Steiner, and Isabelle Guyon. 2015. ChaLearn Looking at People 2015: Apparent Age and Cultural Event Recognition datasets and results. In *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)*. IEEE, Santiago, Chile, 243–251. doi:10.1109/ICCVW.2015.40
- [5] Sunil Gupta and Donald R. Lehmann. 2003. Customers as assets. *Journal of Interactive Marketing* 17, 1 (2003), 9–24. arXiv:https://doi.org/10.1002/dir.10045 doi:10.1002/dir.10045
- [6] Stephen W. Hwang, Mina Atia, Rosane Nisenbaum, Dwayne E. Pare, and Steve Joordens. 2010. Is Looking Older than One's Actual Age a Sign of Poor Health? 136–141 pages. doi:10.1007/s11606-010-1537-0
- [7] Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. 1991. Adaptive Mixtures of Local Experts. *Neural Computation* 3, 1 (1991), 79–87. doi:10.1162/neco.1991.3.1.79
- [8] Julia A. Barthold Jones, Ulrik W. Nash, Julien Vieillefont, Kaare Christensen, Dusan Misevic, and Ulrich K. Steiner. 2019. The AgeGuess database, an open online resource on chronological and perceived ages of people aged 5–100. *Scientific Data* 6, 1 (Oct. 2019), 246. doi:10.1038/s41597-019-0245-9
- [9] Kimmo Kärkkäinen and Jungseock Joo. 2019. FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age. arXiv:1908.04913 [cs.CV] <https://arxiv.org/abs/1908.04913>
- [10] Hongyu Pan, Hu Han, Shiguang Shan, and Xilin Chen. 2018. Mean-Variance Loss for Deep Age Estimation from a Face. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, Salt Lake City, UT, USA, 5285–5294. doi:10.1109/CVPR.2018.00554
- [11] Andraz Puc, Vitomir Struc, and Klemen Grm. 2021. Analysis of Race and Gender Bias in Deep Age Estimation Models. In *2020 28th European Signal Processing Conference (EUSIPCO)*. IEEE, Amsterdam, Netherlands, 830–834. doi:10.23919/Eusipco47968.2020.9287219
- [12] Rasmus Rothe, Radu Timofte, and Luc Van Gool. 2015. DEX: Deep EXpectation of Apparent Age from a Single Image. In *2015 IEEE International Conference on Computer Vision Workshop (ICCVW)*. 252–257. doi:10.1109/ICCVW.2015.41
- [13] Eric Swanson. 2011. Objective assessment of change in apparent age after facial rejuvenation surgery. 1124–1131 pages. doi:10.1016/j.bjps.2011.04.004
- [14] Ziyuan Zhao, Peisheng Qian, Yubo Hou, and Zeng Zeng. 2022. Adaptive Mean-Residue Loss for Robust Facial Age Estimation. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, Taipei, Taiwan, 1–6. doi:10.1109/ICME52920.2022.9859703

TUDLO: Parameter-Efficient Fine-Tuning for Tagalog and Cebuano Early-Grade Mathematics Explanation Generation

Trisha Jean E. Ramos

University of the Philippines Los Baños
Los Baños, Laguna, Philippines
teramos@up.edu.ph

Rachel Edita O. Roxas

University of the Philippines Los Baños
Los Baños, Laguna, Philippines
roroxas2@up.edu.ph

Abstract—Mother Tongue-Based Multilingual Education (MTB-MLE) promotes early-grade instruction in learners’ native languages, yet Philippine-language mathematics resources remain limited. This paper presents TUDLO, a bilingual mathematics chatbot for Tagalog and Cebuano developed using parameter-efficient fine-tuning (QLoRA) and model distillation, deployed as a web application with an API backend. We report numerical accuracy evaluation across all four trained models (Tagalog and Cebuano at both 100-example and 1000-example dataset scales), with partial linguistic coherence results reported for the 1000-example models and full coherence evaluation ongoing.

We systematically evaluate the effect of dataset size on model performance by comparing 100-example and 1000-example models for both languages on identical 100-problem test sets (25 per operation). Results show that $10\times$ dataset expansion dramatically improves numerical accuracy: Cebuano improves from 28% to 82% (+54 percentage points) and Tagalog improves from 26% to 61% (+35 percentage points). Critically, the dominant failure mode differs by scale: 100-example models exhibit pervasive hallucination (full English outputs), while 1000-example models exhibit language mixing (partial code-switching). These findings demonstrate that dataset size is the primary determinant of performance for low-resource educational NLP, and that approximately 1000 balanced examples represent a practical minimum for functional Philippine-language mathematics chatbots.

Index Terms—Mother Tongue-Based Multilingual Education, Low-Resource NLP, Parameter-Efficient Fine-Tuning, QLoRA, Model Distillation, Philippine Languages, Mathematics Education

I. INTRODUCTION

A. The MTB-MLE Implementation Gap

The Philippine K-12 policy institutionalizes Mother Tongue-Based Multilingual Education (MTB-MLE) to enhance foundational learning by using learners’ native languages as the medium of instruction in early grades [1]. Despite this policy framework, implementation faces critical barriers. Teachers report insufficient learning materials in local languages, difficulty producing accurate localized explanations, and inconsistent classroom terminology, particularly in mathematics where conceptual precision is essential [1], [2]. These gaps limit MTB-MLE’s potential to improve learning outcomes in subjects requiring systematic reasoning.

Current teacher practices rely on improvised translations, code-switching, and simplified explanations that risk introducing errors and inconsistencies. In mathematics education, where incorrect explanations can create lasting misconceptions, this approach is particularly problematic. Scalable, automated tools for generating pedagogically sound mathematics content in Philippine languages could address this gap, but no such systems currently exist.

B. AI in Education: Opportunities and Risks

AI-driven chatbots have emerged as powerful educational support tools. Systematic reviews demonstrate their effectiveness in providing comprehensible explanations and personalized guidance [3], [4]. However, large language models (LLMs) suffer from hallucinations (i.e., generating fluent but incorrect responses), which are especially harmful in mathematics where students must learn precise procedures [5], [6]. Generic chatbots without domain-specific adaptation frequently produce mathematically unsound explanations, making them unsuitable for mathematics education without careful validation.

C. Philippine Language NLP: Progress Without Domain Specificity

Recent advances in Philippine language NLP have produced valuable resources: BATAYAN provides benchmarks for Filipino/Tagalog language models [7], CEBBERT offers text classification for Cebuano [8], and CebuaNER establishes baselines for named entity recognition [9]. These tools demonstrate growing technical capability for Philippine languages. However, they focus on general language tasks rather than domain-specific educational applications requiring both linguistic fluency and subject-matter accuracy. No existing work addresses mathematics chatbots for Philippine languages aligned with MTB-MLE pedagogical needs.

D. The Resource Constraint Challenge

Full fine-tuning of large language models remains computationally prohibitive for many educational institutions and researchers in the Philippines. Parameter-efficient fine-tuning

(PEFT) methods, particularly Quantized Low-Rank Adaptation (QLoRA) [10], enable model adaptation using modest hardware such as Google Colab GPUs. Model distillation [11], [12] allows capable multilingual models to generate synthetic training data directly in target languages, avoiding translation artifacts common in machine translation approaches [13]. Together, these techniques make educational NLP feasible within realistic institutional constraints.

E. Research Gap

Despite policy support for MTB-MLE, advances in Philippine language NLP, and availability of parameter-efficient training methods, **to our knowledge, no prior work has demonstrated domain-specific mathematics chatbots for Philippine languages using modern PEFT techniques.** While existing benchmarks assess general language capabilities [7] and specific NLP tasks [8], [9], they do not address educational applications requiring simultaneous linguistic fluency and mathematical accuracy. The specific technical challenges of generating pedagogically appropriate mathematics explanations in low-resource Philippine languages using constrained computational resources remain unexplored.

F. Contributions

This work addresses these gaps by presenting TUDLO, a bilingual Tagalog–Cebuano mathematics chatbot for elementary grades, deployed as a web application with an API backend. Our specific contributions are:

- 1) **First bilingual Philippine-language mathematics chatbot:** We demonstrate technical feasibility of educational NLP for Tagalog and Cebuano using only Google Colab T4 GPUs, establishing that resource-constrained institutions can develop localized mathematics support tools.
- 2) **Evaluation of model distillation for dataset generation:** Direct generation via model distillation produces training data with native Philippine language structures compared to machine translation approaches [13], [14], though our analysis reveals this addresses dataset creation but not the more fundamental challenge of generation stability during inference.
- 3) **Full four-model dataset size and cross-linguistic evaluation:** Through systematic comparison of 100-example and 1000-example models for both Tagalog and Cebuano on identical 100-problem test sets, we demonstrate that dataset size critically affects performance for both languages, and reveal a cross-linguistic performance reversal: at 100 examples, languages perform at rough parity (26% vs. 28%, numerical accuracy), while at 1000 examples Cebuano substantially outperforms Tagalog (82% vs. 61%).
- 4) **Cross-linguistic baseline establishment:** We provide the first comprehensive performance baselines for Philippine-language mathematics NLP across all four model configurations, enabling future comparative research and iterative improvement.

G. Paper Organization

Section II reviews related work in educational chatbots, low-resource NLP, and Philippine language technology. Section III describes our methodology including dataset generation, QLoRA fine-tuning, and evaluation protocols. Section IV presents numerical accuracy results comparing dataset sizes and languages. Section V discusses implications, limitations, and future directions. Section VI concludes with key findings and broader impact.

II. RELATED WORK

A. Educational AI and Chatbot Systems

AI-powered chatbots have demonstrated significant potential in educational contexts. Labadze and Zhorzholiani [3] systematically reviewed AI chatbots in education, finding consistent evidence of improved learning outcomes when chatbots provide personalized explanations and immediate feedback. Debets et al. [4] extended this analysis, identifying comprehensible explanation generation as a critical success factor. However, both reviews emphasize that domain-specific validation is essential, as generic chatbots frequently fail in specialized subjects like mathematics.

B. Hallucinations and Language Mixing in Mathematical Reasoning

Large language models exhibit hallucination behaviors (generating fluent but factually incorrect content) and language mixing challenges, both particularly problematic in mathematics education. Steinbach et al. [5] demonstrated that erroneous LLM feedback in math tutoring systems significantly impairs student learning, with incorrect explanations creating persistent misconceptions. Dang et al. [6] surveyed hallucination patterns across LLMs, finding that mathematical reasoning tasks are especially prone to plausible-sounding but incorrect outputs.

Our work reveals that the dominant failure mode shifts qualitatively with dataset scale. At 100 examples, models exhibit pervasive hallucination, generating full English outputs rather than target-language responses, indicating the model has not acquired target-language generation behavior. At 1000 examples, models successfully generate in the target language but exhibit language mixing (partial code-switching into English), as confirmed by the coherence analysis in Section IV-D. This qualitative shift underscores the necessity of domain-specific fine-tuning and sufficient training data.

C. Philippine Language NLP

Recent work has advanced Philippine language technology: Montalan et al. [7] introduced BATAYAN, the first comprehensive Filipino/Tagalog benchmark for evaluating LLMs across diverse tasks. Tan et al. [8] developed CEBBERT, a lightweight DistilBERT model for Cebuano text classification. Oporto et al. [9] established CebuaNER as a baseline for Cebuano named entity recognition. While these resources demonstrate growing technical capability, they focus on general NLP tasks

rather than domain-specific educational applications requiring subject-matter accuracy.

D. Machine Translation Artifacts

Relying on machine translation for low-resource language datasets introduces systematic artifacts. Sennrich et al. [13] documented that back-translation approaches often produce unnatural phrasing and syntactic structures misaligned with target language norms. Adlaon and Marcos [14] found similar issues in Cebuano–Tagalog neural machine translation, noting that literal translations frequently violate idiomatic usage. Model distillation offers an alternative: directly generating content in target languages using capable multilingual models, potentially avoiding translation-induced artifacts.

E. Parameter-Efficient Fine-Tuning

Full fine-tuning of large language models requires substantial computational resources. Hu et al. [15] introduced LoRA (Low-Rank Adaptation), enabling efficient adaptation by training low-rank decomposition matrices while freezing base model parameters. Dettmers et al. [10] extended this with QLoRA, combining 4-bit quantization with LoRA to dramatically reduce memory requirements. QLoRA enables fine-tuning billion-parameter models on consumer GPUs, making advanced NLP accessible to resource-constrained institutions.

F. Base Model: Qwen2.5-Math

Qwen2.5-Math-1.5B-Instruct [19] is a compact mathematics-specialized language model developed by the Qwen team at Alibaba Cloud. It belongs to the Qwen2.5-Math series, which was trained through a self-improvement pipeline on large-scale mathematical corpora, achieving strong performance on English and Chinese mathematical reasoning benchmarks [19]. The 1.5B parameter variant balances mathematical capability with a small memory footprint, making it suitable for fine-tuning on consumer-grade hardware.

Critically for this work, Qwen2.5-Math-1.5B-Instruct supports instruction-following in multiple languages and has demonstrated generation capability in Philippine languages during our dataset generation experiments. However, the extent of Tagalog and Cebuano coverage in its pre-training data is undocumented by the developers, which introduces uncertainty in interpreting cross-linguistic performance differences between the two target languages. Its strong mathematical reasoning foundation makes it an appropriate base for domain-specific fine-tuning, as the model is expected to require less adaptation for arithmetic correctness and more for target-language generation stability, a hypothesis confirmed by our results.

G. Model Distillation for Dataset Generation

Model distillation techniques enable synthetic dataset creation directly in target languages. Hinton et al. [11] introduced model distillation for model compression. Wang et al. [12] extended this with self-instruct, where capable models generate

instruction-following datasets. Ustun et al. [16] demonstrated this approach for multilingual settings with Aya Model, showing that large multilingual models can directly generate high-quality instructional data in low-resource languages without requiring translation pipelines.

H. Gap Analysis

No prior work has combined parameter-efficient fine-tuning, model distillation, and Philippine-language mathematics education. Existing Philippine NLP resources address general language tasks [7]–[9] but not domain-specific educational applications. Educational chatbot research focuses on high-resource languages [3], [4]. Mathematical reasoning systems typically target English [17], [18]. Our work fills this gap by demonstrating systematic evaluation of PEFT techniques for Philippine-language mathematics chatbot development, characterizing the specific challenges of generating pedagogically appropriate explanations in low-resource settings.

III. METHODOLOGY

A. System Architecture

TUDLO employs a three-stage pipeline: (1) dataset generation using model distillation, (2) parameter-efficient fine-tuning with QLoRA, and (3) systematic evaluation on standardized test sets. This architecture balances accessibility (targeting institutions with limited computing infrastructure) with effectiveness (achieving functional performance within typical research resource constraints).

Figure 1 illustrates the complete system architecture, showing data flow from dataset generation through model training to evaluation and deployment. Figure 2 shows the intended use cases for TUDLO, with three primary user groups: students and teachers who interact through the User Interface Layer to receive explanations and view accuracy indicators, and researchers who run batch evaluations and analyze error patterns through the evaluation pipeline.

B. Dataset Generation via Model Distillation

We generated training datasets directly in Tagalog and Cebuano using Qwen2.5-Math-1.5B-Instruct [19] as a teacher model for dataset generation via model distillation. As described in Section II-F (Qwen2.5-Math), this model combines strong mathematical reasoning with multilingual instruction-following capability. Generating data directly in the target languages avoids translation artifacts documented in machine translation systems [13], [14].

Prompt Design. We designed bilingual prompts specifying: (1) target language (Tagalog or Cebuano), (2) grade level (1–3), (3) operation type (addition, subtraction, multiplication, division), (4) Philippine cultural contexts (e.g., “pandesal”, “tindahan”, family scenarios), and (5) required output format (problem statement, step-by-step solution, final answer).

Generation Protocol. For each language and operation type:

- Generated 25 problems per operation for 100-example datasets ($4 \text{ operations} \times 25 = 100$)




Fig. 1. TUDLO System Architecture




Fig. 2. Use Case Diagram for the TUDLO Bilingual Math Chatbot

- Generated 250 problems per operation for 1000-example datasets (4 operations $\times$ 250 = 1000)
- Used temperature=0.7 for controlled randomness
- Manually filtered obviously malformed outputs (estimated <5% rejection rate)

Dataset Composition. Training datasets included:

- Cebuano: 100-example and 1000-example versions
- Tagalog: 100-example and 1000-example versions
- Stratified sampling across arithmetic operations
- Philippine cultural contexts throughout

C. Parameter-Efficient Fine-Tuning

We fine-tuned Qwen2.5-Math-1.5B-Instruct using QLoRA [10] with the following configuration. QLoRA adapts only a small set of low-rank matrices injected into the attention layers of the base model, keeping the original weights frozen. This allows the model to acquire target-language generation behavior while preserving the mathematical reasoning capabilities of the pre-trained base.

Quantization. 4-bit NormalFloat quantization with double quantization enabled, reducing memory footprint to enable training on Google Colab T4 GPUs (16GB VRAM).

LoRA Parameters.

- Rank (r) = 16 (balancing expressiveness and efficiency)
- Alpha (α) = 32 (scaling factor for LoRA updates)
- Target modules: query, key, value, output projection layers
- Dropout = 0.05

Training Hyperparameters.

- Learning rate: $2e-4$ with linear decay
- Batch size: 4 per device with gradient accumulation (effective batch = 16)
- Epochs: 3 (empirically sufficient for convergence)
- Optimizer: AdamW with weight decay = 0.01
- Max sequence length: 512 tokens

Training Duration.

- 100-example models: 1–2 hours each
- 1000-example models: 6–8 hours each
- Hardware: Google Colab T4 GPU (free tier)

D. Evaluation Protocol

We designed a rigorous evaluation framework addressing both numerical accuracy and linguistic coherence.

Test Set Construction. For each language and model:

- 100 problems total (25 per operation)
- Grade 1–3 difficulty aligned with DepEd curriculum
- Culturally appropriate Philippine contexts
- Manually authored to ensure problem quality
- Disjoint from training data
- Identical problems used across all four models to enable direct comparison

Evaluation Criteria.

- 1) **Numerical Accuracy:** Final answer matches expected result (exact match)

- 2) **Linguistic Coherence:** Response maintains target language without mixing
- 3) **Overall Correctness:** Both numerical accuracy AND linguistic coherence

Evaluation Procedure. All model outputs were manually evaluated by the first author. Numerical accuracy was assessed by comparing final answers to expected results. Linguistic coherence evaluation differed by language:

Tagalog: The first author, a native Tagalog speaker, evaluated all Tagalog outputs for linguistic coherence.

Cebuano: Since the first author is not a Cebuano speaker, all Cebuano outputs were evaluated in consultation with a native Cebuano speaker who validated the identification of language mixing patterns, hallucinations, and nonsensical outputs.

Language mixing was defined as: (1) use of English words or phrases within target-language explanations, or (2) cross-linguistic mixing (e.g., Tagalog words appearing in Cebuano responses). Responses were marked as linguistically coherent only if they maintained exclusive use of the target language throughout the entire explanation. Future work should incorporate evaluation by multiple native speakers for each language to strengthen inter-rater reliability assessment.

Error Taxonomy. We classified errors as:

- **Language Mixing:** Code-switching to English or cross-linguistic mixing
- **Calculation Error:** Incorrect arithmetic but coherent explanation
- **Hallucination:** Nonsensical or structurally incoherent output
- **Format Error:** Missing solution steps or malformed structure

Comparative Analysis. We systematically compared:

- Dataset size effect: 100 vs 1000 examples for each language
- Cross-linguistic patterns: Tagalog vs Cebuano at both scales
- Operation-level performance: Addition, subtraction, multiplication, division

E. Implementation Details

All code implemented using Hugging Face Transformers [20] and PEFT [21] libraries. Models, datasets, and evaluation scripts will be publicly released upon paper acceptance to enable reproducibility and future research.

IV. RESULTS

A. Numerical Accuracy Across All Models

Table I presents numerical accuracy results for all four models evaluated on identical 100-problem test sets (25 per arithmetic operation). Linguistic coherence was evaluated on a reviewed subset of 39 problems per 1000-example model; results are discussed in Section IV-D.

B. Dataset Size Impact

The 10 $\times$ dataset expansion produced substantial numerical accuracy improvements for both languages. For Cebuano, overall numerical accuracy improved from 28% (100 examples) to 82% (1000 examples), a gain of +54 percentage points. For Tagalog, overall numerical accuracy improved from 26% (100 examples) to 61% (1000 examples), a gain of +35 percentage points.

All four arithmetic operations benefited substantially at the 1000-example scale. Cebuano 1000 reached 88% numerical accuracy in both multiplication and division. Tagalog 1000 reached 76% in multiplication and 60% in division. The 100-example models for both languages performed poorly across all operations, with most operations below 35%, confirming that 100 training examples are insufficient for functional mathematics chatbot performance regardless of language.

C. Cross-Linguistic Patterns: Tagalog vs. Cebuano

Reading across Table I by language reveals a notable cross-linguistic reversal. At the 100-example scale, Tagalog and Cebuano perform at rough overall parity (26% vs. 28%), though with contrasting operation-level profiles: Cebuano leads in addition (60% vs. 28%) while Tagalog leads in division (32% vs. 12%). At the 1000-example scale, Cebuano substantially outperforms Tagalog overall (82% vs. 61%, +21 percentage points), with the largest gaps in multiplication (88% vs. 76%) and division (88% vs. 60%).

This reversal, from rough parity at 100 examples to Cebuano leading at 1000 examples, is a significant finding. Both languages benefit substantially from dataset expansion, but Cebuano shows greater absolute gains (+54pp) compared to Tagalog (+35pp). This may reflect differences in distilled training data quality, the base model’s relative language coverage, or operation-language interaction effects that warrant further investigation.

D. Linguistic Coherence Analysis

Linguistic coherence was evaluated on a reviewed subset of 39 problems per model for the two 1000-example models, drawn from the same problem IDs across both languages.

Among the 39 reviewed problems, the Tagalog 1000-example model was linguistically coherent in 19 of 39 cases (49%), with overall correctness of 14 of 39 (36%). Coherence varied substantially by operation: multiplication (71%) and addition (64%) were comparatively strong, while subtraction was notably poor (10%, 1 of 10 reviewed). Of the 25 incorrect responses, 20 (80%) were attributable to incoherence and 5 (20%) to correct-language but wrong-answer outputs.

The Cebuano 1000-example model was coherent in 15 of 39 reviewed cases (38%), with a stark gap between numerical accuracy (85%, 33 of 39) and overall correctness (38%, 15 of 39). All 24 incorrect responses were due to incoherence; not a single error was a pure calculation mistake. Multiplication performed best (71% coherent) while subtraction was again the weakest operation (10%).

TABLE I
NUMERICAL ACCURACY ACROSS ALL FOUR MODELS (100-PROBLEM TEST SETS, 25 PER OPERATION)

Operation	Tagalog 100	Cebuano 100	Tagalog 1000	Cebuano 1000
Addition	28% (7/25)	60% (15/25)	52% (13/25)	80% (20/25)
Subtraction	24% (6/25)	24% (6/25)	56% (14/25)	72% (18/25)
Multiplication	20% (5/25)	16% (4/25)	76% (19/25)	88% (22/25)
Division	32% (8/25)	12% (3/25)	60% (15/25)	88% (22/25)
Overall	26% (26/100)	28% (28/100)	61% (61/100)	82% (82/100)

Numerical accuracy = correct final answer regardless of linguistic coherence. Linguistic coherence results for the 1000-example models are reported in Section IV-D. Test sets consist of 25 problems per operation, 100 total per language, identical across all four models.

Across both models, the dominant coherence failure was language mixing and the production of non-existent words, particularly in Cebuano where portmanteau constructions not found in standard usage were observed (e.g., “saestudyante”, “sumbanan”, “tandam”). These coherence rates (49% for Tagalog 1000, 38% for Cebuano 1000) fall well below what would be required for unassisted classroom deployment, where consistent target-language generation is a prerequisite for pedagogical appropriateness. These findings confirm that linguistic generation stability, not mathematical reasoning, is the primary challenge at the 1000-example scale.

E. Computational Feasibility

Training times confirm accessibility for resource-constrained institutions:

- **100-example models (both languages):** 1–2 hours each (Google Colab T4)
- **1000-example models (both languages):** 6–8 hours each (Google Colab T4)
- **Total cost:** \$0 (free tier sufficient for all four models)

QLoRA’s 4-bit quantization enables training on consumer-grade GPUs, democratizing access to educational AI development for Philippine institutions without dedicated computing infrastructure.

V. DISCUSSION

A. Dataset Size as Critical Success Factor

Our systematic comparison demonstrates that dataset size fundamentally determines model viability for low-resource educational NLP. The $10\times$ expansion from 100 to 1000 examples produced substantial improvements for both languages: +54pp for Cebuano (28% to 82%) and +35pp for Tagalog (26% to 61%). Small datasets (100 examples) are demonstrably insufficient for functional mathematics chatbots for either language, regardless of fine-tuning method quality.

This finding has practical implications: institutions developing Philippine-language educational AI should prioritize dataset expansion over architectural complexity. Based on our empirical results, 1000 balanced examples enabled functional numerical accuracy across all operations for both languages. While our data points (100 and 1000 examples) suggest a threshold exists between these values, determining the precise

minimum viable dataset size requires additional experiments with intermediate sizes (e.g., 300, 500, 750 examples). For practical deployment, we recommend starting with at least 1000 balanced examples; further scaling is likely to yield continued improvements, though the marginal gains beyond this point require empirical verification.

B. Qualitative Shift in Failure Mode with Dataset Scale

A key finding from the linguistic coherence review is that the dominant failure mode differs qualitatively between dataset scales. The 100-example models for both languages produce predominantly English outputs, indicating the model has not learned to generate in the target language at all. At the 1000-example scale, models generate in the target language but exhibit language mixing and production of non-standard constructions. This is supported by the reviewed-subset coherence rates: 49% for Tagalog 1000 and 38% for Cebuano 1000, both substantially below their numerical accuracy rates of 61% and 82% respectively.

The gap between numerical accuracy and overall correctness is particularly pronounced for Cebuano 1000 (82% numerical vs. 38% overall), where every single error in the reviewed problems was a coherence failure rather than a calculation error. This confirms that target-language generation stability is a distinct challenge from mathematical reasoning capacity, and that improving coherence should be the primary focus of future work.

Future work should investigate generation stability techniques including constrained decoding (beam search with language-specific token filtering), temperature tuning (systematic exploration of values below 0.7), and explicit language modeling (training a separate language identifier to detect and penalize code-switching during inference).

C. Cross-Linguistic Performance Reversal

The cross-linguistic reversal, from rough parity at 100 examples to Cebuano leading substantially at 1000 examples (+21pp), is an unexpected finding with important implications. Several factors may contribute:

- **Dataset quality differences:** Uncontrolled variance in distillation quality between languages during dataset creation

- **Base model language coverage:** Differences in how well the base model handles each language, related to pre-training data distribution (though Qwen’s Philippine language coverage is undocumented, as noted in Section II-F)
- **Operation-language interactions:** Unexplored factors linking arithmetic operations to language-specific processing

These findings underscore the importance of matched dataset sizes for rigorous cross-linguistic comparison, and caution against drawing conclusions about inherent language difficulty from small-scale evaluations.

D. Model Distillation: Benefits and Limitations

Direct generation via model distillation successfully produced Philippine-language datasets with native structures, avoiding the literal translations and unnatural phrasing documented in machine translation approaches [13], [14]. However, fine-tuned models still exhibit generation instability, indicating that the dataset generation method alone does not solve target-language coherence during inference.

While we manually filtered obviously problematic outputs and incorporated Philippine cultural contexts, ongoing pedagogical validation, systematic terminology consistency verification, and expert cultural appropriateness review are recommended to further strengthen the deployed system. Model distillation produces *technically valid mathematical content*, and ongoing collaboration with MTB-MLE specialists is recommended to ensure full pedagogical appropriateness for classroom use.

E. Technical Feasibility Within Resource Constraints

A practical contribution is demonstrating that mathematics chatbot development for multiple Philippine languages is feasible within typical institutional resource constraints. All four models were trained using only Google Colab T4 GPUs at zero cost, with 100-example models requiring 1–2 hours and 1000-example models requiring 6–8 hours. The complete pipeline (dataset generation, fine-tuning, evaluation) is executable within the hardware and budget constraints common to Philippine universities, supporting reproducibility by other research groups in similar settings.

F. Limitations and Future Work

This work has important limitations. First, linguistic coherence for the 1000-example models has been evaluated on a reviewed subset of 39 problems per model, with full evaluation of all 100 problems still ongoing. Second, the evaluation lacks comprehensive human assessment: linguistic quality, pedagogical appropriateness, and cultural suitability require expert evaluation. Third, our scope is limited to elementary arithmetic (Grades 1–3 basic operations); more complex mathematics remains untested. Fourth, only two Philippine languages from the same language family were evaluated; generalization to other Philippine languages (Ilocano, Hiligaynon) is unverified.

Fifth, classroom deployment was not conducted, so real-world effectiveness with students is unknown.

Future research should complete the linguistic coherence evaluation to report overall correctness, expand datasets beyond 1000 examples with explicit stratified sampling for both languages, and engage multiple native speakers for inter-rater reliability assessment. Technical improvements should investigate generation stability techniques, and the methodology should be extended to additional Philippine languages. Finally, controlled classroom pilots with teacher supervision are essential to assess real-world effectiveness.

VI. CONCLUSION

We presented TUDLO, a bilingual mathematics chatbot demonstrating the first comprehensive evaluation of parameter-efficient fine-tuning for Philippine-language educational NLP across four model configurations, deployed as a web application with an API backend. Through systematic comparison of 100-example and 1000-example models for both Tagalog and Cebuano on identical 100-problem test sets, several key findings were established.

Dataset size is the primary determinant of performance: 10 $\times$ expansion improved numerical accuracy by +54pp for Cebuano (28% to 82%) and +35pp for Tagalog (26% to 61%), while 100-example models proved uniformly insufficient across all operations and both languages. A qualitative shift in failure mode accompanies this improvement: 100-example models hallucinate (full English outputs), while 1000-example models exhibit language mixing (partial code-switching). This progression reveals that target-language generation is a learnable behavior requiring sufficient training signal, and that generation stability is a distinct challenge warranting dedicated future investigation.

Cross-linguistic analysis revealed a notable reversal: at 100 examples, Tagalog and Cebuano perform at rough parity (26% vs. 28%); at 1000 examples, Cebuano substantially outperforms Tagalog (82% vs. 61%, +21pp). This finding motivates matched comparative studies at larger dataset sizes and cautions against drawing conclusions about inherent language difficulty from small-scale evaluations.

The work confirms technical feasibility within typical institutional resource constraints: all four models were trained using only Google Colab T4 GPUs at zero cost, demonstrating that localized educational AI development is accessible without dedicated computing infrastructure. This establishes a viable pathway for addressing MTB-MLE material shortages through human-AI collaboration, and provides a template for developing localized educational AI tools that serve Philippine languages in their own right.

Future research should complete linguistic coherence evaluation, expand datasets, investigate generation stability techniques, and conduct controlled classroom pilots to assess real-world effectiveness.

REFERENCES

- [1] Department of Education, “Policy Guidelines on the K to 12 Basic Education Program,” DepEd Order No. 21, s. 2019.

- [2] Philippine Normal University Research Center, "Evaluation of the Mother Tongue-Based Multilingual Education (MTB-MLE) Language Policy," Policy Brief Vol. 8, No. 9, 2024.
- [3] L. Labadze and T. Zhorzholiani, "Role of AI Chatbots in Education: Systematic Literature Review," *Int. J. Educ. Technol. High. Educ.*, vol. 20, no. 58, 2023.
- [4] T. Debets, A. van der Meij, and S. Engelen, "Chatbots in Education: A Systematic Review of Objectives, Applications and Outcomes," *Comput. Educ.*, 2025.
- [5] M. Steinbach et al., "When LLMs Hallucinate: Examining the Effects of Erroneous Feedback in Math Tutoring Systems," in *Proc. Educational Data Mining (EDM 2025)*.
- [6] H. A. Dang et al., "Survey and Analysis of Hallucinations in Large Language Models," *Frontiers in Artificial Intelligence*, 2025.
- [7] J. R. Montalan et al., "BATAYAN: A Filipino NLP Benchmark for Evaluating Large Language Models," in *Proc. ACL 2025*.
- [8] G. C. Tan et al., "CEBBERT: A Lightweight Data-Transparent Distil-BERT for Cebuano," in *Proc. PACLIC 38*, 2024.
- [9] J. D. D. Oporto et al., "CebuaNER: A New Baseline Cebuano Named Entity Recognition Model," arXiv:2310.00679, 2023.
- [10] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, "QLoRA: Efficient Finetuning of Quantized LLMs," arXiv preprint arXiv:2305.14314, 2023.
- [11] G. Hinton, O. Vinyals, and J. Dean, "Distilling the Knowledge in a Neural Network," arXiv preprint arXiv:1503.02531, 2015.
- [12] Y. Wang, Y. Kordi, S. Mishra, A. Liu, N. A. Smith, and H. Hajishirzi, "Self-Instruct: Aligning Language Models with Self-Generated Instructions," arXiv preprint arXiv:2212.10560, 2023.
- [13] R. Sennrich, B. Haddow, and A. Birch, "Improving Neural Machine Translation Models with Monolingual Data," in *Proc. ACL*, Berlin, Germany, pp. 86–96, 2016.
- [14] K. M. Adlaon and N. Marcos, "Neural Machine Translation for Cebuano to Tagalog with Subword Unit Translation," in *Proc. Int. Conf. Asian Language Processing (IALP)*, Bandung, Indonesia, pp. 328–333, 2018.
- [15] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-Rank Adaptation of Large Language Models," arXiv preprint arXiv:2106.09685, 2021.
- [16] A. Üstün et al., "Aya Model: An Instruction Finetuned Open-Access Multilingual Language Model," in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 2024.
- [17] D. Hendrycks et al., "Measuring Mathematical Problem Solving With the MATH Dataset," *Advances in Neural Information Processing Systems*, vol. 34, pp. 37893–37904, 2021.
- [18] K. Cobbe et al., "Training Verifiers to Solve Math Word Problems," arXiv preprint arXiv:2110.14168, 2021.
- [19] A. Yang et al., "Qwen2.5-Math Technical Report: Toward Mathematical Expert Model via Self-Improvement," arXiv preprint arXiv:2409.12122, 2024.
- [20] T. Wolf et al., "Transformers: State-of-the-Art Natural Language Processing," in *Proc. EMNLP: System Demonstrations*, pp. 38–45, 2020.
- [21] S. Mangrulkar, S. Gugger, L. Debut, Y. Belkada, and S. Paul, "PEFT: State-of-the-art Parameter-Efficient Fine-Tuning," <https://github.com/huggingface/peft>, 2023.

Gated Learning-Progress Exploration

Mikael Vincent Tumamos
University of San Carlos
Cebu City, Philippines
tumamos@mikaelvincent.dev

Ervin Joshua Guirnela
University of San Carlos
Cebu City, Philippines
17100409@usc.edu.ph

Christine Peña
University of San Carlos
Cebu City, Philippines
cfpena@usc.edu.ph

Abstract

Exploration via intrinsic rewards can accelerate learning when extrinsic feedback is sparse or deceptive, yet many intrinsic objectives overemphasize novelty in regions of high uncertainty where learning is slow or unstable. We introduce *Gated Learning-Progress Exploration* (GLPE), a simple intrinsic-reward family that prioritizes transitions that (i) induce meaningful changes in a learned feature representation and (ii) occur in regions where a forward dynamics model is actively improving. GLPE estimates region-local learning progress using the discrepancy between short- and long-horizon exponential moving averages (EMAs) of prediction error within an online partition of feature space, and combines it with a feature-space impact term. An optional region-specific gate suppresses intrinsic rewards in regions where prediction error remains high but learning progress is low, providing a targeted guardrail against unproductive curiosity. We also study *GLPE (no gate)*, which uses the same base intrinsic score without any gating. We evaluate both variants with a common PPO backbone on five Gymnasium control benchmarks and compare against widely used intrinsic-reward baselines, reporting performance under both step- and wall-clock-normalized views. Across the suite, GLPE (no gate) matches the strongest baseline within uncertainty, while the full gated variant is most beneficial in sparse-reward regimes where ungated intrinsic rewards can destabilize training.

Keywords

reinforcement learning, intrinsic motivation, exploration, learning progress

1 Introduction

Deep reinforcement learning agents can learn effectively when dense extrinsic rewards provide informative gradients throughout training; however, performance often degrades when rewards are sparse, delayed, or shaped in ways that do not guide early exploration. Intrinsic-motivation methods address this by augmenting the environment reward with an auxiliary signal that encourages exploration even when external feedback is absent or uninformative [3, 11, 12, 20].

Prior intrinsic objectives span several complementary perspectives. Count-based and density-model visitation bonuses encourage novelty and broad coverage [4, 10, 22]. Information-gain objectives seek transitions that reduce uncertainty in a learned model [7]. Prediction-error bonuses reward surprise under a forward model, including early deep predictive approaches [21] and feature-space curiosity methods such as ICM [13]. Random Network Distillation (RND) and related curiosity variants operationalize novelty through prediction error against a fixed target [5, 6], while episodic mechanisms and hybrid schemes further shape exploration using reachability memory, disagreement, or directed exploration strategies

[1, 14, 16]. Impact-driven rewards instead encourage controllable state change, providing a complementary signal when novelty alone is not task-aligned [15]. These intrinsic signals are typically applied as reward shaping during policy optimization; unlike potential-based shaping, they are not generally policy-invariant and can change the optimal policy, so they are assessed empirically [9].

Despite strong empirical results, many intrinsic objectives exhibit a recurring failure mode: *unproductive curiosity*. Novelty or prediction error can remain high in regions that are difficult to model—because dynamics are complex, stochastic, or only weakly coupled to task progress—causing the agent to spend substantial effort collecting experience that improves neither the intrinsic model nor the task policy. The “noisy TV” phenomenon is one concrete example of this pathology [8]. The issue is especially acute in sparse-reward settings, where the intrinsic signal can dominate early learning.

A complementary perspective is to reward *learning progress*: experience is valuable when it improves a predictive model over time rather than merely being surprising [2, 11, 17]. Learning progress can help distinguish regions that are learnable-but-unmastered from regions that are either already predictable or effectively noisy. Practical progress-based methods must still address how to localize progress in high-dimensional observations, how to couple it to exploration-relevant change, and how to avoid spending effort in regions where progress has stalled.

We propose *Gated Learning-Progress Exploration* (GLPE), a lightweight intrinsic-reward family that couples learning progress to representation change. GLPE measures *impact* as the magnitude of change in a learned feature representation between consecutive observations, and estimates *region-local learning progress* by tracking forward-model prediction error within an online partition of the feature space using short- and long-horizon exponential moving averages. The resulting base score favors transitions that both change the representation and occur in regions where the dynamics model is actively improving. To mitigate unproductive curiosity, GLPE optionally applies a simple, region-specific *gate* that suppresses intrinsic shaping in regions with persistently high prediction error but low learning progress, serving as a targeted guardrail rather than a global annealing heuristic. We also study *GLPE (no gate)*, which uses the same base score without gating and serves as a low-overhead default.

We evaluate both GLPE variants with a common PPO backbone on MOUNTAINCAR-v0, BIPEDALWALKER-v3, and three MuJoCo locomotion tasks (ANT-v5, HALF-CHEETAH-v5, HUMANOID-v5), comparing against widely used intrinsic-reward baselines including ICM [13], RND [6], RIDE [15], and RIAC [2]. Across this suite, GLPE (no gate) remains consistently competitive under both step- and wall-clock-normalized views, while the gated variant is most helpful

in sparse-reward regimes where ungated shaping can destabilize training (Section 4).

2 Method

2.1 Problem setup

We consider an episodic Markov decision process with observations $o_t \in \mathcal{O}$, actions $a_t \in \mathcal{A}$, and extrinsic rewards r_t^{ext} . A stochastic policy $\pi_\theta(a | o)$ is trained to maximize the expected discounted return.

We augment the extrinsic reward with an intrinsic term computed per transition,

$$r_t = r_t^{\text{ext}} + \eta_t \text{clip}(r_t^{\text{int}}, -r_{\max}, r_{\max}), \quad (1)$$

where $\eta_t \geq 0$ controls the intrinsic-extrinsic trade-off and $r_{\max} > 0$ bounds the magnitude of intrinsic shaping. For environment-terminal transitions, we set $r_t^{\text{int}} = 0$ so that intrinsic shaping does not depend on episode termination.

We propose two intrinsic reward formulations within a shared framework: GLPE and GLPE (no gate). Both combine feature-space impact with region-local learning progress. GLPE additionally introduces a region-specific gate that suppresses intrinsically rewarding but unproductive experience. We refer to these two variants collectively as the *GLPE family*.

2.2 Latent dynamics model

Let ϕ_ω be a learned encoder producing a d -dimensional feature vector $z_t = \phi_\omega(o_t)$. We train (i) a forward dynamics predictor f_ψ that maps (z_t, a_t) to a prediction of z_{t+1} , and (ii) an inverse dynamics predictor g_ξ that maps (z_t, z_{t+1}) to an action distribution. The inverse model is instantiated as a classifier for discrete actions and as a Gaussian likelihood model for continuous actions.

Training uses the combined loss

$$\mathcal{L}_{\text{dyn}}(t) = \beta_{\text{fwd}} \mathcal{L}_{\text{fwd}}(t) + \beta_{\text{inv}} \mathcal{L}_{\text{inv}}(t), \quad (2)$$

with weights $\beta_{\text{fwd}}, \beta_{\text{inv}} > 0$. The forward loss is the mean-squared error in feature space,

$$\mathcal{L}_{\text{fwd}}(t) = \frac{1}{d} \|f_\psi(z_t, a_t) - z_{t+1}\|_2^2, \quad (3)$$

and we define the per-transition prediction error

$$e_t = \mathcal{L}_{\text{fwd}}(t). \quad (4)$$

The intrinsic rewards below depend on (z_t, z_{t+1}) and e_t , but do not require access to privileged environment state.

2.3 Online region partitioning in latent space

To localize learning progress, we maintain an online partition of the latent space into regions. Each region corresponds to a leaf in a binary space-partitioning tree over the d -dimensional feature space.

Given an embedding z_t , we assign it to a leaf region ρ_t by routing it through the tree. Each leaf maintains a visit counter and, while it remains splittable, a buffer of points used to choose a split dimension and threshold. Once a leaf has accumulated C assigned points and its depth is below $D_{\max}$, we split it by selecting the highest-variance coordinate that yields a non-degenerate partition and cutting at the median along that coordinate. One child retains the parent region

identifier (and its accumulated statistics); the other child receives a new identifier whose statistics are initialized on its first visit. This procedure yields a growing set of regions that progressively refines frequently visited parts of feature space.

We update region statistics (defined next) using the region assignment for the current embedding z_t . The capacity C and maximum depth $D_{\max}$ control the granularity of the partition.

2.4 Region-local learning progress

For each region r we maintain two exponential moving averages (EMAs) of the forward prediction error: a long-horizon average μ_r^{long} and a short-horizon average μ_r^{short} . When a transition at time t is assigned to region $\rho_t = r$, we update

$$\mu_r^{\text{long}} \leftarrow \beta_{\text{long}} \mu_r^{\text{long}} + (1 - \beta_{\text{long}}) e_t, \quad (5)$$

$$\mu_r^{\text{short}} \leftarrow \beta_{\text{short}} \mu_r^{\text{short}} + (1 - \beta_{\text{short}}) e_t, \quad (6)$$

with $0 < \beta_{\text{long}} < 1$, $0 < \beta_{\text{short}} < 1$, and typically $\beta_{\text{long}} > \beta_{\text{short}}$ so that μ_r^{long} changes more slowly. On the first visit to a newly created region identifier, we initialize $\mu_r^{\text{long}} = \mu_r^{\text{short}} = e_t$.

We define the current learning progress of a region r as the positive part of the difference between the long and short averages,

$$\text{LP}(r) = \max(0, \mu_r^{\text{long}} - \mu_r^{\text{short}}). \quad (7)$$

For each transition we use the region assignment ρ_t and set $\text{LP}_t = \text{LP}(\rho_t)$. This quantity is large when the short-term error has recently dropped below the longer-term baseline, indicating active model improvement in that region.

2.5 Impact in feature space

We also compute a feature-space impact term that encourages transitions that change the learned representation,

$$I_t = \|z_{t+1} - z_t\|_2. \quad (8)$$

Unlike count-based novelty, impact depends on the learned representation and can capture structured changes in high-dimensional observations.

2.6 Component-wise scale stabilization

The magnitudes of I_t and LP_t can vary across tasks and throughout training. To obtain a stable scale without task-specific tuning, we maintain running root-mean-square (RMS) normalizers for each component via an exponential moving average of squared values. For a scalar signal x_t we update an accumulator

$$v \leftarrow \beta_{\text{rms}} v + (1 - \beta_{\text{rms}}) x_t^2, \quad (9)$$

and normalize by $\text{RMS}(x_t) = \sqrt{v} + \varepsilon$ with a small $\varepsilon > 0$.

Applying this independently to I_t and LP_t yields normalized components

$$\widetilde{I}_t = \frac{I_t}{\text{RMS}(I_t)}, \quad \widetilde{\text{LP}}_t = \frac{\text{LP}_t}{\text{RMS}(\text{LP}_t)}. \quad (10)$$

We then define a shared base intrinsic score

$$u_t = \alpha_{\text{impact}} \widetilde{I}_t + \alpha_{\text{LP}} \widetilde{\text{LP}}_t, \quad (11)$$

with nonnegative weights $\alpha_{\text{impact}}, \alpha_{\text{LP}}$. Both variants use the same u_t and differ only in whether it is gated.

2.7 GLPE (no gate)

GLPE (no gate) is a region-aware intrinsic reward that combines feature-space impact and region-local learning progress without any explicit suppression mechanism. The intrinsic reward is

$$r_t^{\text{int}} = u_t. \quad (12)$$

This variant preserves the locality and nonstationarity tracking provided by the region partition and EMAs, while avoiding additional gating hyperparameters.

2.8 GLPE: gated learning-progress exploration

GLPE augments the same base score u_t with a binary, region-specific gate that suppresses intrinsic rewards in regions that appear difficult to predict yet show little or no learning progress.

Global robust thresholds. At any time, let $\mathcal{R}$ denote the set of regions that have been visited at least once, and let $\text{LP}(r)$ denote the current learning progress in region r . We compute robust global reference values using medians across visited regions,

$$\begin{aligned} \text{LP}_{\text{med}} &= \text{median}_{r \in \mathcal{R}} \text{LP}(r), \\ e_{\text{med}} &= \text{median}_{r \in \mathcal{R}} \mu_r^{\text{short}}. \end{aligned} \quad (13)$$

We define a learning-progress threshold $\tau_{\text{LP}} = \kappa \text{LP}_{\text{med}}$ with multiplier $\kappa > 0$, and a scale-free normalized error score

$$s_r = \frac{\mu_r^{\text{short}}}{e_{\text{med}} + \varepsilon}, \quad (14)$$

where $\varepsilon > 0$ avoids division by zero. A region r is considered *unproductive* on a visit if it simultaneously has low learning progress and high normalized error,

$$\text{LP}(r) < \tau_{\text{LP}} \quad \text{and} \quad s_r > \tau_s, \quad (15)$$

with $\tau_s > 0$ a fixed threshold.

Gate dynamics. Each region maintains a gate value $g_r \in \{0, 1\}$ and two counters tracking consecutive unproductive and productive visits. Gating is only activated once at least R_{min} regions have been visited so that the medians are meaningful. When gating is active, the update uses persistence and hysteresis: (i) if $g_r = 1$ and the region is unproductive for K consecutive visits, set $g_r \leftarrow 0$; and (ii) if $g_r = 0$, re-enable the region when $\text{LP}(r) > h \tau_{\text{LP}}$ for two consecutive visits, where $h > 1$ is a hysteresis multiplier. If gating is inactive (fewer than R_{min} visited regions), we set $g_r \leftarrow 1$ and reset the counters.

The intrinsic reward for GLPE is then

$$r_t^{\text{int}} = g_{\rho_t} u_t. \quad (16)$$

2.9 On-policy optimization and intrinsic scheduling

We optimize the policy with PPO using the augmented reward r_t . Training details are given in Section 3.

For the GLPE family, we anneal the strength of intrinsic shaping over the course of training using a cosine taper. Let $p \in [0, 1]$ denote training progress as a fraction of total environment steps. Given a start fraction p_{start} and end fraction p_{end} with $p_{\text{start}} < p_{\text{end}}$,

we define

$$w(p) = \begin{cases} 1, & p \leq p_{\text{start}}, \\ \frac{1}{2} \left(1 + \cos\left(\pi \frac{p - p_{\text{start}}}{p_{\text{end}} - p_{\text{start}}}\right) \right), & p_{\text{start}} < p < p_{\text{end}}, \\ 0, & p \geq p_{\text{end}}, \end{cases} \quad (17)$$

and set $\eta_t = \eta w(p_t)$ in the augmented reward. When p_{start} and p_{end} are omitted, we hold $\eta_t = \eta$ fixed throughout training.

The latent dynamics model parameters (ω, ψ, ξ) are updated using the same on-policy transitions that generate intrinsic rewards, while the region statistics, gates, and RMS normalizers are updated online during intrinsic computation. Figure 1 summarizes the intrinsic reward computation within one PPO update.

- 1: **Input:** on-policy transitions $\{(o_t, a_t, o_{t+1})\}_{t=1}^N$; encoder ϕ_ω ; forward model f_ψ ; inverse model g_ξ ; region assignment function ρ ; region EMAs $\mu_r^{\text{short}}, \mu_r^{\text{long}}$; RMS accumulators; gate states g_r (optional)
- 2: **Output:** intrinsic rewards $\{r_t^{\text{int}}\}_{t=1}^N$; updated region statistics; updated intrinsic model parameters
- 3: **for** $t = 1$ to N **do**
- 4: $z_t \leftarrow \phi_\omega(o_t)$, $z_{t+1} \leftarrow \phi_\omega(o_{t+1})$
- 5: $e_t \leftarrow \frac{1}{d} \|f_\psi(z_t, a_t) - z_{t+1}\|_2^2$
- 6: Assign region $r \leftarrow \rho(z_t)$ (inserting z_t and splitting leaves when needed)
- 7: Update region EMAs $\mu_r^{\text{short}}, \mu_r^{\text{long}}$ with e_t
- 8: $\text{LP}_t \leftarrow \max(0, \mu_r^{\text{long}} - \mu_r^{\text{short}})$
- 9: $I_t \leftarrow \|z_{t+1} - z_t\|_2$
- 10: Normalize I_t and LP_t with running RMS to get $\tilde{I}_t, \tilde{\text{LP}}_t$
- 11: $u_t \leftarrow \alpha_{\text{impact}} \tilde{I}_t + \alpha_{\text{LP}} \tilde{\text{LP}}_t$
- 12: **if** gating enabled **then**
- 13: Compute robust global medians $\text{LP}_{\text{med}}, e_{\text{med}}$ across visited regions
- 14: Update gate g_r using persistence and hysteresis rules
- 15: $r_t^{\text{int}} \leftarrow g_r u_t$
- 16: **else**
- 17: $r_t^{\text{int}} \leftarrow u_t$
- 18: **end if**
- 19: **end for**
- 20: Update intrinsic model parameters (ω, ψ, ξ) using $\{(z_t, a_t, z_{t+1})\}$ and supervised losses

Figure 1: High-level pseudocode for computing GLPE-family intrinsic rewards within a PPO update.

3 Experimental Setup

3.1 Tasks and environments

We evaluate exploration and control on five standard Gymnasium benchmarks spanning discrete and continuous control: MOUNTAINCAR-v0, BIPEDALWALKER-v3, and the MuJoCo locomotion tasks ANT-v5, HALF-CHEETAH-v5, and HUMANOID-v5. Across experiments, we use the default environment reward functions and termination conditions, with no domain randomization and no additional frame skipping. Training uses vectorized environments with B parallel instances and rollouts of length T per instance. We set (B, T) so that the nominal on-policy batch size per PPO update is

$$N = B \times T = 16,384 \quad (18)$$

transitions, varying (B, T) only to accommodate environment throughput and episode structure. To match the total step budget exactly, the final update may use a smaller batch when fewer than N steps remain. Table 1 summarizes the suite, interaction budgets, and parallelization settings. For each environment, we evaluate all methods using the same set of training seeds (Table 1) and enable deterministic execution where supported.

Environment	B	T	N	Total steps	Seeds
MOUNTAINCAR-V0	16	1,024	16,384	3,000,000	10
BIPEDALWALKER-V3	8	2,048	16,384	7,000,000	8
ANT-V5	8	2,048	16,384	15,000,000	8
HALFCHEETAH-V5	8	2,048	16,384	15,000,000	8
HUMANOID-V5	4	4,096	16,384	30,000,000	5

Table 1: Benchmark suite and training budgets. B is the number of parallel environment instances, T is the rollout horizon per instance, and $N = B \times T$ is the number of transitions used in each PPO update. Each run is evaluated offline at saved checkpoints using 20 episodes (Section 3.4).

Training budgets beyond the early-learning regime. The total environment-step budgets in Table 1 are intentionally generous so that runs extend beyond the initial improvement phase. This provides sufficient budget to assess long-horizon stability of exploration behavior and intrinsic shaping, including late-training drift, regressions, or oscillations.

3.2 Methods compared

We compare the two GLPE variants, GLPE and GLPE (no gate) (Section 2), against commonly used exploration baselines: Vanilla PPO, ICM, RND, RIDE, and RIAC. Vanilla PPO optimizes only the extrinsic environment reward. ICM uses forward-model prediction error in a learned feature space as an intrinsic signal, while RND uses the prediction error of a trainable predictor against a fixed random target network. RIDE rewards feature-space impact, down-weighted by episodic visitation counts obtained by discretizing features, and RIAC uses region-local learning progress computed from changes in forward-model error within an adaptive partition of feature space. All methods use the same PPO backbone and network architectures. Within each environment we keep PPO hyperparameters fixed across methods to isolate the effect of the intrinsic objective.

For intrinsic-reward methods, the policy is trained with the augmented reward defined in Section 2, combining the environment reward with a scaled and clipped intrinsic term. We use the same intrinsic scaling coefficient η and clipping range $r_{\max}$ for all intrinsic methods within a given environment (Table 2), so differences arise from the intrinsic signal itself rather than from reward-scale tuning. GLPE (no gate) uses the same base intrinsic score as GLPE but disables the gating mechanism.

3.3 RL backbone and training protocol

All agents are trained with Proximal Policy Optimization (PPO; [19]) and generalized advantage estimation (GAE; [18]). The policy

and value function are parameterized by separate multilayer perceptrons with two hidden layers of width 256 and ReLU activations. For continuous-control tasks, the policy outputs a diagonal Gaussian distribution. Actions are squashed to the environment bounds when those bounds are finite.

Each training iteration proceeds as follows: (i) collect up to $N = 16,384$ on-policy transitions by running the current policy in B vectorized environments for T steps; (ii) compute the intrinsic reward for each transition (when applicable), set intrinsic rewards to zero on environment-terminal transitions and on time-limit truncations without a final observation, and form the total reward used for advantage estimation; (iii) compute advantages and value targets with GAE, bootstrapping across time-limit truncations when a valid final observation is available and treating truncations without a final observation as terminal for bootstrapping purposes; and (iv) perform multiple epochs of PPO updates over shuffled minibatches of the collected batch. We use Adam optimizers for both the policy and value networks, normalize advantages within each update batch, and clip gradient norms to 1.0.

We normalize vector observations using a running mean and variance estimated online from collected rollouts. We apply the same normalization to the inputs of the policy, value, and intrinsic modules.

Intrinsic reward computation and updates. Trainable intrinsic models are updated once per PPO iteration using the same on-policy batch used for PPO. For methods whose intrinsic output is not internally normalized (e.g., ICM, RND, and RIDE), we apply a running RMS normalization to the raw intrinsic signal before scaling and clipping. For methods that produce a normalized intrinsic output by construction (RIAC and the GLPE family), we directly apply scaling and clipping. For GLPE and GLPE (no gate), we additionally anneal the intrinsic coefficient η_t using the cosine taper described in Section 2, with task-specific start and end fractions listed in Table 2.

3.4 Evaluation protocol

We evaluate performance using undiscounted episodic return from the environment (extrinsic reward only). We do not run evaluation

Environment	η	$r_{\max}$	α_{LP}	p_{start}	p_{end}	C	$D_{\max}$	β_{long}	β_{short}	κ	τ_s	K	$R_{\min}$
MOUNTAINCAR-V0	0.05	4.0	0.5	0.05	0.80	128	10	0.995	0.90	0.010	2.0	5	8
BIPEDALWALKER-V3	0.08	4.0	0.5	0.12	0.75	200	12	0.995	0.90	0.010	2.0	5	16
ANT-V5	0.06	4.0	0.6	0.10	0.70	256	12	0.997	0.92	0.012	2.5	8	64
HALFCHEETAH-V5	0.04	5.0	0.5	0.05	0.75	256	12	0.995	0.90	0.010	2.0	5	32
HUMANOID-V5	0.06	5.0	0.5	0.05	0.80	256	12	0.998	0.92	0.010	2.0	5	32

Table 2: GLPE hyperparameters by environment. We fix $\alpha_{\text{impact}} = 1.0$ for all tasks. p_{start} and p_{end} define the cosine taper schedule for η_t as a function of training progress p (Section 2); this taper is used for GLPE and GLPE (no gate). C and $D_{\max}$ are the region capacity and maximum depth for the online latent-space partition. κ sets the learning-progress threshold via $\tau_{LP} = \kappa LP_{\text{med}}$, τ_s is the normalized-error threshold, K is the number of consecutive unproductive visits required to gate off a region, and $R_{\min}$ is the minimum number of visited regions required before gating is activated (Section 2). The hysteresis multiplier is fixed to $h = 2.0$ for all tasks. GLPE (no gate) uses the same settings but does not apply gating.

rollouts online during training; instead, we evaluate saved checkpoints offline. Each evaluated checkpoint is run for 20 episodes in a separate environment instance using deterministic action selection (distribution mode). Episodes are reset with fixed per-episode seeds derived from the training seed. For each environment and method, we run multiple independent training seeds (Table 1) and aggregate results across seeds.

Checkpointing and logging. During training we checkpoint the full agent state at step 0, at nine additional warmup points approximately evenly spaced within the first checkpoint interval, and then at a fixed checkpoint interval thereafter (set to 10% of the total step budget per environment), with a final checkpoint at the end of training. We additionally log scalar metrics every fixed number of environment steps (set to 2% of the total step budget).

3.5 Efficiency measurements

In addition to sample efficiency (learning as a function of environment steps), we track computational efficiency using wall-clock time. All timings are measured in the same execution setting used for training (CPU in our experiments) and therefore reflect end-to-end optimization overhead. For each PPO iteration we log per-component wall-clock time spent in environment interaction, policy inference, intrinsic computation, intrinsic-module updates, advantage computation, and PPO optimization. We use these logs to form cumulative training time.

When reporting scalar summaries of learning curves, we compute two AUC metrics via trapezoidal integration of mean return: step-AUC (return versus environment steps, divided by the environment’s step budget) and wall-clock AUC (return versus cumulative training time). For wall-clock AUC, we use a *common time budget* per environment defined as the minimum final wall-clock time attained among the compared methods. This ensures that summaries do not rely on extrapolation beyond the measured runtime of any method.

3.6 Implementation details and hyperparameters

Tables 3 and 2 list the PPO and GLPE hyperparameters used in all experiments. Unless specified in the tables, we use $\gamma = 0.99$ and value-loss coefficient 0.5 for PPO.

Environment	LR	Epochs	Minib.	Clip	λ	Ent.	v -clip	KL stop
MOUNTAINCAR-V0	3.0×10^{-4}	10	16	0.20	0.95	0.00	0.00	0.06
BIPEDALWALKER-V3	5.0×10^{-4}	5	16	0.25	0.95	0.00	0.00	0.06
ANT-V5	1.5×10^{-4}	15	64	0.20	0.95	0.00	0.20	0.04
HALFCHEETAH-V5	3.0×10^{-4}	10	32	0.20	0.95	0.01	0.20	0.03
HUMANOID-V5	2.0×10^{-4}	5	32	0.20	0.97	0.01	0.20	0.03

Table 3: PPO hyperparameters by environment. “Minib.” is the number of minibatches per update (each epoch iterates once over all N samples). “Ent.” is the entropy regularization coefficient. “ v -clip” is the value-function clipping range (0 disables value clipping). “KL stop” is the approximate KL threshold for early stopping within an update.

For methods that learn a latent dynamics model (ICM, RIDE, RIAC, and the GLPE family), the shared encoder produces 128-dimensional features and is implemented as a two-layer MLP (256

units per layer, ReLU) for these vector-observation tasks. Forward and inverse model losses are weighted equally. The model is trained with Adam at learning rate 3×10^{-4} , and intrinsic-model gradient norms are clipped to 5.0.

4 Results and Discussion

We compare GLPE and GLPE (no gate) against Vanilla PPO and four intrinsic-reward baselines (ICM, RND, RIDE, RIAC) on the five-task suite from Section 3. Unless otherwise stated, we report *extrinsic* evaluation return under the deterministic offline evaluation protocol (Section 3.4).

4.1 Learning curves and final performance

Figure 2 shows evaluation learning curves versus environment steps. Across tasks, GLPE (no gate) tracks the best-performing baseline closely and exhibits stable performance across seeds. The gated variant differs primarily in sparse-reward or exploration-sensitive regimes: it is most beneficial on MOUNTAINCAR-V0 and can be overly conservative on HUMANOID-V5, where between-seed variance dominates.




Figure 2: Evaluation learning curves for GLPE and baseline intrinsic-reward methods. Each panel plots mean *extrinsic* return (higher is better) versus environment steps for one task. Solid lines show the mean across seeds at each checkpoint; markers show the corresponding per-seed means (20 evaluation episodes per checkpoint). The horizontal dashed line is a fixed “solved” return threshold used for summary analyses. These curves summarize learning speed and stability across seeds; tight clustering of per-seed markers indicates robust performance.

Final-checkpoint performance is reflected by the rightmost points in Figure 2. Table 4 reports step-AUC, which rewards both early learning and sustained performance over the fixed interaction budget.

On MOUNTAINCAR-V0, GLPE attains the best final return and the strongest step-AUC among all compared methods, and both GLPE variants solve in all seeds under the solved-threshold criterion (Figure 3, 100% threshold). On BIPEDALWALKER-V3, GLPE (no gate) improves rapidly and remains competitive by step-AUC, but many runs plateau just below the solved threshold of 300, making thresholded success counts sensitive to small differences near the cutoff (Figure 3, 100% threshold). On the MuJoCo locomotion tasks, where extrinsic reward is comparatively dense, the two GLPE

variants behave similarly and remain within a modest gap of the strongest intrinsic baseline under both the learning-curve view (Figure 2) and the step-AUC summary (Table 4). HUMANOID-v5 exhibits particularly high variance, with wide confidence intervals across all methods, so method rankings are not stable under bootstrap uncertainty.

Some diagnostic plots include auxiliary GLPE variants (e.g., cached medians and single-component ablations) to isolate computational and signal contributions; we treat these as supporting analyses rather than primary baselines.

4.2 Ablations: impact vs. learning progress

To probe which components drive performance in different regimes, Table 5 reports final returns for GLPE and two single-term ablations on representative tasks.

On sparse-reward MOUNTAINCAR-v0, combining impact and learning progress is critical: each single-term variant falls substantially short of full GLPE. In contrast, on dense-reward ANT-v5 the impact-only variant is competitive, and on near-threshold BIPEDALWALKER-v3 the LP-only variant comes closest to the solved cutoff. These patterns motivate the combined score as a robust default when the “right” intrinsic driver is task-dependent.

4.3 Thresholded reliability and steps-to-threshold

Learning curves can obscure whether a method is merely slower or is failing entirely. We therefore evaluate reliability at fixed return

Environment	GLPE	GLPE (no gate)	Best baseline
MOUNTAINCAR-v0	-107.0 [-114.8, -101.8]	-111.3 [-128.7, -101.4]	RIDE: -113.7 [-134.2, -101.9]
BIPEDALWALKER-v3	249.0 [223.9, 270.7]	255.2 [230.0, 275.6]	ICM: 257.3 [229.4, 278.3]
ANT-v5	3,402 [3,067, 3,708]	3,402 [3,067, 3,707]	RND: 3,565 [3,202, 3,876]
HALFCHEETAH-v5	5,005 [4,708, 5,316]	4,998 [4,700, 5,296]	RIDE: 5,208 [5,019, 5,410]
HUMANOID-v5	1,583 [551, 3,360]	1,665 [769, 2,921]	ICM: 1,781 [829, 2,797]

Table 4: Step-AUC of deterministic evaluation return versus cumulative environment steps, computed by trapezoidal integration of the mean return-versus-steps curve over evaluation checkpoints and dividing by the environment’s step budget. This yields the average evaluation return over training, rewarding methods that learn quickly and sustain high returns. Intervals are obtained by integrating the per-checkpoint 95% bootstrap confidence bands over seeds and should be interpreted as an approximate uncertainty range for the integrated metric. Higher is better. The best baseline is selected (per environment) among Vanilla PPO, ICM, RND, RIDE, and RIAC by highest mean step-AUC.

Environment	GLPE	Impact-only	LP-only
MOUNTAINCAR-v0	-96.2 [-96.5, -95.9]	-111.8 [-137.0, -96.2]	-106.6 [-127.4, -96.0]
BIPEDALWALKER-v3	266.5 [214.0, 301.3]	277.5 [256.0, 296.5]	298.8 [285.5, 308.1]
ANT-v5	4,961 [4,421, 5,359]	5,233 [4,544, 5,663]	4,599 [3,617, 5,401]

Table 5: Final-checkpoint mean extrinsic return for GLPE and component ablations. Impact-only sets $\alpha_{LP}=0$ and LP-only sets $\alpha_{impact}=0$. Brackets show 95% bootstrap confidence intervals over training seeds. Higher is better.

thresholds using steps-to-threshold statistics that report both time-to-threshold and reach rates across seeds (Figure 3). We use task-specific solved thresholds (MOUNTAINCAR-v0: -110, BIPEDALWALKER-v3: 300, ANT-v5: 5,000, HALFCHEETAH-v5: 6,000, HUMANOID-v5: 3,000) and additionally report results at 25% and 50% of these levels.




Figure 3: Sample efficiency and reliability at fixed return thresholds. For each task, bars report the median number of environment steps to first reach a target return level, estimated by linearly interpolating between evaluation checkpoints. The three x-axis positions correspond to 25%, 50%, and 100% of each task’s solved threshold (the dashed line in Figure 2); intermediate thresholds are computed using the same monotonic mapping across tasks. Error bars show the standard deviation in steps across seeds that reached the threshold. Bar opacity encodes the fraction of seeds that reached the threshold by the end of training, and the number above each bar is the count of successful seeds. Lower bars and higher opacities indicate faster and more reliable learning.

At the 50% threshold, GLPE (no gate) reaches the target in all seeds on four tasks and in 4/5 seeds on HUMANOID-v5, matching or exceeding the best baseline reach count on every environment (Figure 3, 50% threshold). The gap between 50% and 100% success is concentrated in high-variance or near-cutoff regimes: HUMANOID-v5 and BIPEDALWALKER-v3 show many runs that reach competent behavior without consistently clearing the full solved cutoff (Figure 3, 100% threshold). On MOUNTAINCAR-v0, by contrast, both GLPE variants reach the solved threshold quickly and reliably, and several baselines exhibit occasional failures that remain near the minimum return (Figure 2).

4.4 Intrinsic shaping dynamics and gating behavior

The GLPE family is designed as a transient exploration aid rather than a permanent auxiliary objective. We therefore anneal the intrinsic coefficient η_t using the cosine taper defined in Section 2, down-weighting intrinsic shaping late in training. Figure 4 decomposes the training-time rollout reward into extrinsic and intrinsic components. Compared to baselines that maintain a nonzero intrinsic term throughout training, GLPE and GLPE (no gate) reduce the applied intrinsic contribution as training progresses.

To make the gate’s effect concrete, Figure 5 shows state-space projections of final-checkpoint trajectories with time steps colored by whether intrinsic shaping is active. Across environments, gate-off points occupy a small and typically localized subset of the visited distribution. This supports the interpretation of gating as a targeted guardrail: it suppresses shaping only in narrow regions that appear persistently high-error and low-progress, rather than globally weakening exploration.




Figure 4: Decomposition of rollout reward during training. Solid lines show the mean extrinsic reward per environment step collected during PPO rollouts. Dashed lines show the intrinsic shaping component actually added during training, computed as the difference between the logged total reward and the extrinsic reward after each method’s scaling and clipping (GLPE additionally applies tapering). Intrinsic values use the right axis and extrinsic values use the left axis. This plot is a diagnostic for the optimization signal and complements the evaluation curves, which report extrinsic return without intrinsic shaping.




Figure 5: State-space view of GLPE’s gating decisions. Points correspond to states visited in trajectories from the final checkpoint (subsamped); for MOUNTAINCAR-v0 we plot raw position and velocity, and for the remaining environments we project z-scored observations onto the first two principal components. Blue points indicate time steps where the intrinsic term is active (gate on), while orange points indicate time steps where the gate disables intrinsic shaping (gate off). Across tasks, gate-off points occupy a small and typically localized subset of the visited distribution, indicating that gating suppresses intrinsic shaping only in narrow regions rather than globally.

4.5 Wall-clock efficiency and computational overhead

Intrinsic rewards can change not only sample efficiency but also end-to-end training cost. Figure 6 reports wall-clock AUC of evaluation return under a common per-environment time horizon (Section 3.5), ensuring that slower methods are not advantaged by extrapolation beyond measured runtime. ACROSS BIPEDALWALKER-v3, HALF-CHEETAH-v5, ANT-v5, and HUMANOID-v5, GLPE (no gate) remains close to the strongest baseline under this wall-clock view. The gated variant typically attains lower wall-clock AUC, consistent with the additional computation required by gating and robust statistics.




Figure 6: Wall-clock AUC of evaluation performance. For each task, bars report trapezoidal area under the curve (AUC) of the mean deterministic evaluation return as a function of cumulative training time (Section 3.5). To compare methods fairly when overhead differs, each task uses a common time horizon equal to the minimum final training time among the compared methods, and curves are truncated to this horizon (no extrapolation). Higher is better. Error bars are derived by integrating the per-checkpoint 95% confidence bounds on mean return.

Figure 7 decomposes per-update runtime into environment interaction, PPO optimization, and intrinsic components. On the more expensive MuJoCo tasks, environment stepping and PPO optimization dominate total time, making intrinsic overhead a secondary factor; on MOUNTAINCAR-v0, where environment interaction is inexpensive, intrinsic overhead plays a larger role in wall-clock comparisons.

Recomputing the global medians used by GLPE’s gate to set robust thresholds is one identifiable cost of the gated variant. In a standalone microbenchmark of this operation, we report median throughput across nine trials for the same workload. Recomputing the medians every update (cache off, $k=1$) yields 20,877 transitions/s (baseline, 1.00 $\times$), while recomputing every $k=64$ updates and reusing cached values in between (cache on) yields 100,173 transitions/s (4.83 $\times$ speedup). Because cached medians are slightly stale and can change gating decisions when $k > 1$, we treat caching as an optional implementation choice rather than part of the core GLPE comparison.

5 Conclusion

Intrinsic-reward shaping can substantially accelerate learning when extrinsic feedback is sparse or delayed, yet novelty- or prediction-error-based bonuses can also concentrate exploration in regions where progress is slow, unstable, or effectively noisy. We introduced the *GLPE family*, which ties intrinsic reward to two complementary signals computed from on-policy transitions: feature-space impact and region-local learning progress, estimated from short- and long-horizon prediction-error averages within an online latent-space partition. The gated variant adds a simple region-wise gate that suppresses intrinsic shaping when prediction error remains high while learning progress stalls, providing a targeted guardrail against unproductive curiosity. GLPE (no gate) uses the same base score without this additional thresholding.

Across five Gymnasium control benchmarks and a shared PPO backbone, GLPE (no gate) is the most consistent default, remaining competitive with strong intrinsic baselines under both step- and wall-clock-normalized views. The gated variant is most valuable in sparse-reward regimes such as MOUNTAINCAR-v0, where it improves reliability and reduces instability relative to ungated shaping, but it can be unnecessarily conservative in high-variance settings such as HUMANOID-v5. Finally, the cosine taper schedule makes intrinsic shaping intentionally transient, smoothly reducing its contribution late in training so that final performance reflects primarily the learned task policy.

This study is limited to vector-observation benchmarks and on-policy PPO; as with most intrinsic bonuses, the shaping signal we apply is not policy-invariant. Future work can extend GLPE to richer observation modalities and alternative training regimes and make the gated variant more computationally efficient and adaptive without changing the underlying intrinsic objective.




Figure 7: Timing breakdown per PPO update. Each panel shows a stacked decomposition of wall-clock seconds per training update for one task. For each method, we take the final quarter of recorded updates in each run, compute per-component medians, and then average these medians across runs (stacked bars). The black error bar on each stack is a 95% confidence interval for the total seconds per update across runs.

References

- [1] Adrià Puigdomènech Badia, Pablo Sprechmann, Alex Vitvitskyi, Daniel Guo, Bilal Piot, Steven Kapturowski, Olivier Tieleman, Martin Arjovsky, Alexander Pritzel, Andrew Bolt, and Charles Blundell. 2020. Never Give Up: Learning Directed Exploration Strategies. *arXiv (Cornell University)* (2 2020). doi:10.48550/arxiv.2002.06038
- [2] A. Baranes and P.-Y. Oudeyer. 2009. R-IAC: Robust Intrinsically Motivated Exploration and Active Learning. *IEEE Transactions on Autonomous Mental Development* 1, 3 (10 2009), 155–169. doi:10.1109/tamd.2009.2037513
- [3] Andrew G. Barto. 2012. *Intrinsic Motivation and Reinforcement Learning*. 17–47 pages. doi:10.1007/978-3-642-32375-1_2
- [4] Marc G. Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. 2016. Unifying Information Count-Based Exploration and Intrinsic Motivation. *arXiv (Cornell University)* (6 2016). doi:10.48550/arxiv.1606.01868
- [5] Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A. Efros. 2018. Large-Scale Study of Curiosity-Driven Learning. *arXiv (Cornell University)* (8 2018). doi:10.48550/arxiv.1808.04355
- [6] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. 2018. Exploration by Random Network Distillation. *arXiv (Cornell University)* (10 2018). doi:10.48550/arxiv.1810.12894
- [7] Rein Houthoofd, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. 2016. VIME: Variational Information Count-Based Exploration. *arXiv (Cornell University)* (5 2016). doi:10.48550/arxiv.1605.09674
- [8] Augustine N. Mavor-Parker, Kimberly A. Young, Caswell Barry, and Lewis D. Griffin. 2021. How to Stay Curious While Avoiding Noisy TVs Using Aleatoric Uncertainty Estimation. *arXiv (Cornell University)* (2 2021). doi:10.48550/arxiv.2102.04399
- [9] Andrew Ng, Daishi Harada, and Stuart Russell. 1999. Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping. Morgan Kaufmann, San Francisco, us, 278–287. doi:10.5555/645528.657613
- [10] Georg Ostrovski, Marc G. Bellemare, Aaron van den Oord, and Remi Munos. 2017. Count-Based Exploration with Neural Density Models. *arXiv (Cornell University)* (3 2017). doi:10.48550/arxiv.1703.01310
- [11] Pierre-Yves Oudeyer and Frederic Kaplan. 2007. What Is Intrinsic Motivation? A Typology of Computational Approaches. *Frontiers in Neurorobotics* 1 (1 2007), 6. doi:10.3389/neuro.12.006.2007
- [12] Pierre-Yves Oudeyer, Frederic Kaplan, and Verena V. Hafner. 2007. Intrinsic Motivation Systems for Autonomous Mental Development. *IEEE Transactions on Evolutionary Computation* 11, 2 (4 2007), 265–286. doi:10.1109/tevc.2006.890271
- [13] Deepak Pathak, Pulkit Agrawal, Alexei Efros, and Trevor Darrell. 2017. Curiosity-Driven Exploration by Self-Supervised Prediction. *arXiv* (2017). doi:10.48550/arXiv.1705.05363
- [14] Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. 2019. Self-Supervised Exploration via Disagreement. *arXiv (Cornell University)* (6 2019). doi:10.48550/arxiv.1906.04161
- [15] Roberta Raileanu and Tim Rocktäschel. 2020. RIDE: Rewarding Impact-Driven Exploration for Procedurally-Generated Environments. *arXiv (Cornell University)* (2 2020). doi:10.48550/arxiv.2002.12292
- [16] Nikolay Savinov, Anton Raichuk, Raphaël Marinier, Damien Vincent, Marc Pollefeys, Timothy Lillicrap, and Sylvain Gelly. 2018. Episodic Curiosity through Reachability. *arXiv (Cornell University)* (10 2018). doi:10.48550/arxiv.1810.02274
- [17] Juergen Schmidhuber. 1991. *A Possibility for Implementing Curiosity and Boredom in Model-Building Neural Controllers*. 222–228 pages. doi:10.7551/mitpress/3115.003.0030
- [18] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. 2015. High-Dimensional Continuous Control Using Generalized Advantage Estimation. *arXiv (Cornell University)* (6 2015). doi:10.48550/arxiv.1506.02438
- [19] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *arXiv* (2017). doi:10.48550/arXiv.1707.06347
- [20] Satinder Singh, Andrew Barto, and Nittapong Chentanez. 2004. *Intrinsically Motivated Reinforcement Learning*. MIT Press, Cambridge, us.
- [21] Bradley C. Stadie, Sergey Levine, and Pieter Abbeel. 2015. Incentivizing Exploration in Reinforcement Learning with Deep Predictive Models. *arXiv (Cornell University)* (7 2015). doi:10.48550/arxiv.1507.00814
- [22] Haoran Tang, Rein Houthoofd, Davis Foote, Adam Stooke, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. 2016. #Exploration: A Study of Count-Based Exploration for Deep Reinforcement Learning. *arXiv (Cornell University)* 30 (11 2016), 2750–2759. doi:10.48550/arxiv.1611.04717

Evaluating Large Language Models in Filipino Underspecified, Multi-Turn Conversations

Renee Rosary H. Pintor

Department of Computer, Information
Sciences, and Mathematics
University of San Carlos
Cebu City, Philippines
22104214@usc.edu.ph

Lisa Juress P. Wang

Department of Computer, Information
Sciences, and Mathematics
University of San Carlos
Cebu City, Philippines
22200117@usc.edu.ph

Glenn B. Pepito

Department of Computer, Information
Sciences, and Mathematics
University of San Carlos
Cebu City, Philippines
gbpepito@usc.edu.ph

ABSTRACT

Most evaluations of large language models (LLMs) for Filipino use single-turn prompting, even though real-world interactions often involve several exchanges across multiple turns. This can overestimate real-world performance because multi-turn interactions introduce underspecification and require models to keep track of context across turns. We study whether the Lost in Conversation phenomenon, where models perform worse and less consistently in underspecified, multi-turn conversations than in fully-specified, single-turn conversations, also manifests in Filipino, by adapting the Lost in Conversation framework, a multi-turn simulation framework, to Filipino NLP tasks. Using pairs of fully-specified and sharded instructions derived from Batayan, a Filipino NLP benchmark, we evaluate five multilingual LLMs across eight tasks under both fully-specified, single-turn settings and underspecified, multi-turn settings. We find that, on average across models, performance decreases by 6.45% in underspecified, multi-turn conversations relative to fully-specified, single-turn conversations, and unreliability increases by $\sim 3.8\times$ on average, despite similar aptitude in both settings. We also evaluate two context-repetition mitigations and find that repetition of prior turns can partially recover performance on some tasks but has uneven effects across task types. These findings suggest that evaluating LLMs only with fully-specified, single-turn prompts can hide important failures in real conversations, especially increased output variability. This motivates the need for Filipino evaluations that include underspecified, multi-turn settings and report both performance and reliability.

KEYWORDS

Filipino NLP, large language models, multi-turn evaluation, underspecified conversations

1 INTRODUCTION

Large language models (LLMs) are increasingly being used as chat-based assistants, but most evaluations still rely on single-turn prompts that lay out all information at once. These kinds of benchmarks can overestimate real-world performance, since they do not account for the complexities of multi-turn interactions [6].

This gap matters especially for the Filipino language, referred to hereafter as Filipino, given that evaluation resources are still far more limited compared to English. Recent Filipino benchmarks [8–10] have improved evaluation resources for Filipino, but these evaluations are still largely limited to single-turn interactions. In this paper, we focus on Batayan [9], a Filipino NLP benchmark that

covers eight tasks across three competencies: paraphrase identification (PI), question answering (QA), sentiment analysis (SA), toxicity detection (TD), causal reasoning (CR), natural language inference (NLI), abstractive summarization (AS), and machine translation (MT).

Recent work has begun to shift toward multi-turn frameworks [6], but Filipino multi-turn evaluation resources remain limited. More importantly, existing multi-turn benchmarks do not directly compare single-turn and multi-turn settings on the same underlying tasks, especially for conversations of three or more turns and where each turn is underspecified. Laban et al. [5] address this issue through a simulation framework that compares single-turn, fully-specified instructions with multi-turn, underspecified instructions created by *sharding* an instruction into multiple turns, and they show that even LLMs with strong performance on single-turn settings can have their performance degrade substantially in multi-turn settings (the *Lost in Conversation* phenomenon).

This study aims to determine whether the Lost in Conversation phenomenon, and the associated performance degradation and increased unreliability observed in English multi-turn, underspecified conversations, also manifest in Filipino, specifically across different Filipino NLP tasks in conversations of three or more turns. To do this, we adapt the Lost in Conversation framework to Filipino using tasks from Batayan [9] and evaluate multilingual LLMs under both single-turn and multi-turn settings. We also explore simple mitigation strategies that repeat previously revealed context to test whether they can reduce degradation in underspecified, multi-turn conversations [5].

The contributions of this paper are as follows:

- We extend the Lost in Conversation framework [5] to Filipino and construct pairs of fully-specified and sharded instructions from the Batayan benchmark [9].
- We evaluate multilingual LLMs on the same Filipino tasks in both single-turn, fully-specified and multi-turn, underspecified settings, enabling direct comparison of performance and reliability across settings [5].
- We analyze performance degradation across tasks to identify which Filipino NLP tasks are more sensitive to multi-turn underspecification.
- We conduct mitigation experiments to test whether repetition of prior turns can improve performance in underspecified, multi-turn conversations [5].

1.1 Terminology

Following Laban et al. [5], a fully-specified instruction presents all information in one prompt, whereas an underspecified instruction provides only partial information at a given turn. Accordingly, single-turn settings present the instruction once, while multi-turn settings distribute the same information across multiple turns. A sharded instruction is a version of a fully-specified instruction in which the original information is divided into multiple shards that together preserve the original content.

2 RELATED WORK

Several works have evaluated the performance of Large Language Models (LLMs) in different Filipino tasks [8–10]. Notably, Montalan et al. [9] introduced Batayan, a Filipino NLP benchmark covering eight tasks across understanding, reasoning, and generation. A limitation of these Filipino evaluations is that they cover only single-turn settings, even though real-world applications often involve multi-turn interactions [6].

In response to such limitations, recent research on LLM evaluation has increasingly shifted toward multi-turn frameworks [6]. For Filipino, however, multi-turn evaluation resources remain limited. SEA-MTBench, included in the SEA-HELM (Southeast Asian Holistic Evaluation of Language Models) evaluation suite [12], is a manually translated and localized adaptation of MT-Bench [14], designed to assess the multi-turn conversational capabilities of LLMs in Southeast Asian languages, including Filipino. Similar to the original MT-Bench, a key limitation of this benchmark is that it only covers two-turn conversations, leaving a gap for Filipino multi-turn evaluations that span three or more turns.

Advancements in language models and the rise of conversational AI, such as ChatGPT, have driven increased research interest in multi-turn evaluations of LLMs [2, 4, 14]. Laban et al. [5] argue that many of these benchmarks rely on episodic tasks, where each turn can be evaluated in isolation, and frame interactions in broad categories such as refinement or follow-up, which can overstate LLM performance and prevent direct comparison between single-turn and multi-turn behavior on the same underlying tasks. Their framework instead makes underspecification a central element, motivated by its prevalence in real LLM conversation logs [3] and by the human tendency to follow the principle of least effort [15].

To address these issues, Laban et al. [5] propose a simulation framework consisting of both single-turn, fully-specified and multi-turn, underspecified simulations where models are tested on a common set of tasks in both settings. Fully-specified instructions are *sharded*, or split across multiple turns that, as a whole, deliver the same information as the original instruction, requiring models to operate with incomplete information that is clarified gradually. Their results show that LLMs exhibit a significant decline in overall performance and reliability, and a minor drop in aptitude when faced with multi-turn, underspecified tasks [5]. They identify a phenomenon they call the *Lost in Conversation* phenomenon, where on the same exact tasks, LLMs that perform well on single-turn, fully-specified conversations struggle in multi-turn, underspecified conversations. However, a key limitation of their study is that it only covers English instructions, leaving open the question of whether similar challenges arise in other languages.




Figure 1: Overview of the methodological framework of the study. Simulation types are described in Section 3.2.1.

3 METHODOLOGY

This study evaluates how large language models (LLMs) handle Filipino underspecified, multi-turn conversations by adapting the publicly-available Lost in Conversation framework [5]¹ to Filipino NLP tasks. We first transform fully-specified Filipino instructions into sharded instructions following the framework’s sharding process, resulting in a dataset of fully-specified and sharded instruction pairs. This allows us to evaluate LLM performance on the same instruction in both single-turn and multi-turn settings. Multi-turn conversations are simulated using the sharded conversation simulation environment in the same framework. Figure 1 provides a high-level overview of the study’s methodology.

3.1 Dataset Construction

We construct a dataset of paired fully-specified and sharded instructions using stratified random sampling and the semi-automatic sharding process introduced in the Lost in Conversation framework [5].

3.1.1 Data and Task Selection. The data used in this study were sourced from Batayan [9], a Filipino NLP benchmark included in the publicly available SEA-HELM evaluation suite². Batayan was selected for its quality, as validated by native speakers, and for its comprehensiveness, covering eight tasks across three NLP competencies: Natural Language Understanding (NLU), Natural Language Reasoning (NLR), and Natural Language Generation (NLG).

We use stratified random sampling to select instances for sharding, processing one instruction at a time until target per-label sample counts are met. Table 1 presents the class distribution for each task together with a task description.

Because translation is considered an episodic task in which each turn can be treated as an independent subtask [5], it may be less sensitive to underspecification. We nevertheless include all eight Batayan tasks to compare sensitivity across task types [9].

3.1.2 Sharding Process. Batayan [9] instructions are originally fully-specified, single-turn prompts in which all information is provided at once. Using the sharding process proposed in [5], we create a sharded equivalent of the original instruction that distributes the

¹https://github.com/microsoft/lost_in_conversation

²<https://github.com/aisingapore/sea-helm>

Table 1: Class distribution per task and task descriptions, adapted from [9].

Competency	Task	Label	# Samples	Description
NLU	PI	True	10	Determine whether two texts convey the same idea.
		False	10	–
	QA	span	20	Select the correct answer given a passage and question.
	SA	Negative (Negatibo)	10	Classify social media texts as positive, negative, or neutral.
		Neutral (Neutral)	10	–
		Positive (Positibo)	10	–
	TD	Clean (Malinis)	10	Detect hate speech or abusive language in social media texts.
		Toxic (Mapoot)	10	–
NLR	CR	Cause (Sanhi)	10	Identify the cause or effect of a premise among given statements.
		Effect (Bunga)	10	–
	NLI	Contradiction	10	Classify the relationship between two sentences as entailment, contradiction, or neutral.
		Entailment	10	–
		Neutral	10	–
NLG	AS	summary	20	Summarize a paragraph into a concise sentence.
	MT	eng→tgl	10	Translate text from English to Filipino.
		tgl→eng	10	Translate text from Filipino to English.
Total			180	

Table 2: Example fully-specified (from Batayan [9]) and sharded instruction pair from the Question Answering (QA) task.

Fully-Specified	Sharded
Bibigyan ka ng isang talata, isang tanong, at apat na pagpipiliang sagot. Sumagot base sa talata sa pamamagitan ng pagpili ng isa sa mga opsiyong ibinigay. Talata: `` `Bago nito, nag-ulat ang Chinese news agency na Xinhua na may isang eroplano ang na-hijack. Ipinahayag sa mga pinakahuling ulat na nagkaroon ng banta ng bomba sa eroplano at pinabalik sa Afghanistan saka lumapag sa Kandahar. Habang ayon naman sa mga naunang balita na pinabalik ang eroplano sa Afghanistan matapos hindi ito pinayagang mag-emergency landing sa Üürmqi. `` `	Ano ang hindi sinabi ng Xinhua sa pinakabagong ulat nila? Isang Chinese news agency ang Xinhua. Ang Xinhua ay nag-ulat na may isang eroplano ang na-hijack. Sa pinakahuling ulat, sinabing may banta ng bomba sa eroplano. Ayon sa pinakahuling ulat, pinabalik ang eroplano sa Afghanistan. Ayon sa pinakahuling ulat, lumapag ang eroplano sa Kandahar. Sinabi sa naunang balita na pinabalik ang eroplano sa Afghanistan.
Tanong: Ano ang hindi nakasaad sa pinakabagong ulat mula sa news agency na Xinhua? A: Nakatanggap ang eroplano ng banta sa bomba. B: Nakarating ang eroplano sa Üürmqi. C: Pinabalik ang eroplano sa Afghanistan. D: Nakarating ang eroplano sa Kandahar.	Sa naunang balita, sinabing hindi pinayagan ang eroplano na mag-emergency landing sa Üürmqi. Ang mga pagpipilian ay: A) Nakatanggap ang eroplano ng banta sa bomba, B) Nakarating ang eroplano sa Üürmqi, C) Pinabalik ang eroplano sa Afghanistan, D) Nakarating ang eroplano sa Kandahar.

same information across multiple turns. Following the sharded instruction definition in [5], the first shard expresses the overall intent of the task, and the subsequent shards add clarifications and further details. Taken together, the shards must provide the same amount of information as the original instruction. Table 2 shows an example pair of fully-specified and sharded instructions from the Question Answering (QA) task.

We adopt the semi-automatic, four-step sharding process in [5] and adapt it to Filipino. The first three steps rely on an LLM, and we employ Gemini 2.5 Flash, a low-cost LLM with strong Filipino capabilities. The final step is a manual inspection and edit done by

the authors. This process is done iteratively until the target number of instances per label is reached. Refinement tasks AS/MT use a different procedure (see Section 3.2.2).

- (1) **Preparation and Segmentation.** An instruction from Batayan [9] is selected as a candidate and provided to the LLM in a task-specific prompt that includes at least three few-shot examples. The model is asked to break the instruction into non-overlapping segments and to split it as small as possible so that the segments correspond to the atomic units of the original instruction. Here, an atomic unit is a minimal unit of the instruction such that the collection of all units covers the same information as the original instruction. Following [5], instructions that yield fewer than three segments are discarded.
- (2) **Rephrasing.** Each segment is then rewritten by the LLM to be conversational and decontextualized [5]. Here, decontextualized means that aside from the initial shard, the remaining shards can be reordered while still making sense and conveying the same instruction. This step makes each shard fit naturally within a simulated conversation without modifying its informational content. Since such rephrasing may itself affect performance, we later include a single-turn control that preserves the rephrased shards while removing multi-turn underspecification. Using task-specific prompts with at least three few-shot examples, the LLM is instructed to choose the initial shard that states the high-level objective of the instruction.
- (3) **Verification.** To verify that the sharding preserves information and that shards are decontextualized, we run preliminary simulations using the Lost in Conversation framework. For each fully-specified and sharded instruction pair, we conduct ten FULL, ten CONCAT, and ten SHUFFLE-CONCAT simulations. CONCAT is a single-turn simulation where all shards are concatenated into one prompt, while SHUFFLE-CONCAT is identical except that the order of all shards except the first is randomized before concatenating (see Section 3.2.1 for details on simulation types). Following Laban et al., we require that the averaged performance (Section 3.4) for both

CONCAT and SHUFFLE-CONCAT simulations be at least 80% of the corresponding FULL averaged performance. CONCAT checks that the sharding process does not substantially remove information, and SHUFFLE-CONCAT checks that shards are sufficiently decontextualized. Samples that fail either condition are discarded.

- (4) **Manual Inspection and Edit.** A final review is manually done by the authors for all sharded instructions. We check that the sharded instructions follow Laban et al. [5]’s precise definition and specifications, and that they are written in grammatically correct and natural Filipino. Individual shards may be edited, merged, split, added, or removed as needed, but only to preserve the original task intent and information content, improve grammaticality and naturalness, and satisfy the framework’s requirements for underspecification and decontextualization. The authors have the final decision on whether to accept or discard the sample. The final output of this process is a set of fully-specified and sharded instruction pairs used in the single-turn and multi-turn simulations in Section 3.2.1.

3.2 Simulation Setup

3.2.1 Simulation types. Following Laban et al. [5], we use five simulation types: two single-turn settings (FULL and CONCAT) and three multi-turn settings (SHARDED, RECAP, and SNOWBALL). FULL, CONCAT, and SHARDED are used across all evaluated models in the main experiments, while RECAP and SNOWBALL are used only for smaller mitigation experiments on Qwen-SEA-LION-v4-32B-IT, the model that showed the largest performance degradation in SHARDED conversations. We limit RECAP and SNOWBALL to Qwen-SEA-LION-v4-32B-IT due to computational cost.

FULLY-SPECIFIED/FULL. A single-turn, fully-specified baseline using the original Batayan [9] prompt template, with answer-format constraints removed (e.g., “Sagot: \$OPTION”) so that models can respond more naturally.

SHARDED. A multi-turn, underspecified setting where the assistant receives at most one shard per turn. The assistant is given only a brief system instruction before the first turn, and is *not* told that the conversation will be multi-turn or underspecified.

CONCAT. A single-turn control in which all shards are combined into one prompt in bullet-point form, preserving rephrasing introduced by the sharding process (see Section 3.1.2) while removing underspecification. This is used to distinguish performance degradation in the multi-turn, underspecified setting from degradation caused by potential information loss during sharding.

RECAP. The SHARDED conversation plus a final turn that restates all shards together. This gives the assistant one last opportunity to answer.

SNOWBALL. A multi-turn simulation in which each turn repeats all previously revealed shards and then introduces at most one new shard. The prompt essentially “snowballs” in length across turns.

3.2.2 Simulation of Sharded Conversations. We rely on the sharded conversation simulation environment in the Lost in Conversation framework [5] and adapt it to Filipino/Taglish interactions. The simulation involves three roles: (1) the evaluated assistant, (2) a Filipino user simulator, and (3) the system component that (i)




Figure 2: Process diagram of the sharded conversation simulation (adapted from [5]).

classifies assistant responses, (ii) extracts the specific answer span from those responses, and (iii) scores the extracted answer with a task-specific evaluator. Figure 2 shows the overall flow of the sharded conversation simulation.

In a SHARDED simulation, the user simulator first reveals the initial shard to the assistant. The user simulator has access to the full sharded instruction and generates Filipino or Taglish utterances that sound conversational and fit the conversation history, while preserving the original shard’s information. After each response from the assistant, the strategy classifier determines whether the response was an answer attempt. If it were an answer attempt, the answer extractor identifies the exact answer span (e.g., a class label) from the assistant’s response, which is then scored by a task-specific evaluator function.

If the assistant’s response is not an answer attempt, or if an answer attempt is evaluated as incorrect, the simulation proceeds to the next turn. At each subsequent turn, the user simulator may reveal up to one unrevealed shard. If the assistant asks about something unrelated to any unrevealed shard, the user simulator responds without revealing a shard and without adding information beyond what is in the shards (e.g., “Di ko alam.”).

The simulation stops under either of two conditions: (1) the assistant answers correctly, or (2) all shards are revealed.

Refinement tasks. For Machine Translation (MT) and Abstractive Summary (AS), we do not use the sharded conversation simulator because the assistant should already have its answer from the first turn. At each subsequent turn, we provide one additional shard of information and it is asked to refine its answers. For AS, we split the original passage into one sentence per shard, and for MT, we split it into seven words per shard. We treat each turn as an answer attempt, the whole response as the answer, and evaluate only the final response after all shards have been revealed.

LLM components. We use Gemini 2.5 Flash, selected for its low cost and strong Filipino capabilities, for all LLM-based components in the SHARDED conversation simulation.

Validation. Laban et al. [5] validate the sharded conversation simulation environment through manual inspection and report an overall success rate of 97.8% using GPT-4o-mini. Since our study differs both in language and the model used for the LLM components, we also manually inspect and annotate 200 of our simulated sharded conversations following the same inspection procedure and criteria. We find that 95.5% of inspected sharded conversations are successful. Detailed inspection results are reported in Appendix A.

3.3 Models and Parameters

3.3.1 Evaluated models. We evaluate five prominent multilingual LLMs chosen to cover both widely used commercial assistants and smaller open-source models designed for Southeast Asian languages. This selection allows us to compare robustness to underspecified, multi-turn Filipino interactions across models that differ in size and likely training emphasis. GPT and Gemini models (GPT-5.1, Gemini 3.0 Pro, and Gemini 2.5 Pro) were selected because ChatGPT and Google Gemini are widely used chat-based LLM systems in the Philippines. For example, a survey among Filipino Gen Z respondents ($N = 691$) reports that 48.63% use ChatGPT and 5.49% use Gemini [1]. We also selected smaller open-source SEA-LION models from AI Singapore, which are instruction-tuned for Southeast Asian languages, including Filipino. Table 3 summarizes the evaluated models.

Table 3: Overview of evaluated LLMs.

Model	Developer	Release	Access
Gemma-SEA-LION-v4-27B-IT	AI Singapore	Aug 2025	SEA-LION API
Qwen-SEA-LION-v4-32B-IT	AI Singapore	Oct 2025	SEA-LION API
Gemini 2.5 Pro	Google	Jun 2025	Gemini API
Gemini 3 Pro	Google	Nov 2025	Gemini API
GPT-5.1	OpenAI	Nov 2025	OpenAI API

3.3.2 Parameters. For all evaluated assistants and LLM-based system components, temperature is set to $T = 1.0$, and all other generation settings are left at their default values. The maximum response length is set to 1,000 tokens, except for the Gemini models, which often have “thinking” enabled by default; for these models, we allow up to 10,000 tokens to avoid truncation. All models are evaluated under zero-shot prompting.

3.4 Metrics and Scoring

3.4.1 Task scoring. For the NLU tasks (PI, QA, SA, TD) and NLR tasks (CR, NLI), the task-specific evaluator scores the assistant’s answer as either incorrect or correct (0 or 100) using exact matching. Refinement tasks (AS, MT) are scored on a continuous range (0-100). AS is scored with the average of the three metrics reported in Batayan [9]: BERTScore [13], ChrF++ [11], and ROUGE-L F1 [7]. MT is scored using ChrF++.

3.4.2 Metrics. Following Laban et al. [5], we account for LLMs’ stochastic generation by running N independent simulations and obtaining a set of scores $S = \{S_i\}_{i=1}^N$, where each simulation yields a score $S_i \in [0, 100]$. We report the three metrics from the Lost in Conversation framework: averaged performance ($\bar{P}$; Eq. (1)), aptitude (A^{90} ; Eq. (2)), and unreliability (U_{10}^{90} ; Eq. (3)) [5]. We also

adopt the reporting format of Laban et al. [5] to ensure a clear presentation of results.

$$\bar{P} = \frac{1}{N} \sum_{i=1}^N S_i \quad (1)$$

$$A^{90} = \text{percentile}_{90}(S) \quad (2)$$

$$U_{10}^{90} = \text{percentile}_{90}(S) - \text{percentile}_{10}(S). \quad (3)$$

Here, averaged performance ($\bar{P}$) measures the model’s average score across N simulations, A^{90} estimates the model’s best performance (top 10% runs), and U_{10}^{90} measures unreliability or variability in performance. [5]. All three metrics are computed per instruction and then averaged across instructions within each task to obtain corpus-level results.

3.5 Simulation Scale

The main experiment simulates FULL, CONCAT, and SHARDED conversations using instructions across eight tasks (180 instructions total), evaluating five LLMs with $N = 10$ independent simulations per instruction for each model and simulation type. Following the Lost in Conversation framework [5], we set $N = 10$ to enable estimation of averaged performance ($\bar{P}$), aptitude (A^{90}), and unreliability (U_{10}^{90}). In total, we simulate $180 \times 5 \times 3 \times 10 = 27,000$ conversations for the main experiment.

4 RESULTS AND DISCUSSION

4.1 Averaged Performance Findings

Table 4 presents the averaged performance ($\bar{P}$) for five LLMs across eight Filipino NLP tasks under FULL, CONCAT, and SHARDED simulations. Following the Lost in Conversation framework [5], CONCAT serves as a control, preserving the rephrasing of sharded instructions while removing underspecification by presenting all shards in one prompt. Averaging the performance drops per model in the rightmost column, the average performance drops by 2.43 points in CONCAT and 5.40 points in SHARDED, suggesting that multi-turn underspecification contributes additional degradation beyond rephrasing alone. On average across models, this corresponds to a 6.45% decrease in performance in SHARDED relative to the FULL baseline (mean of per-model relative drops).

Overall, larger commercial models appear more robust to rephrasing of instructions, as indicated by smaller performance drops under CONCAT relative to FULL. Gemini 3.0 Pro even overall improves in CONCAT compared to FULL. Under SHARDED, Gemini 3.0 Pro exhibits the smallest average degradation (0.16 points), while Qwen-SEA-LION-v4-32B-IT exhibits the largest average degradation (13.21 points). However, an even smaller open-source model Gemma-SEA-LION-v4-27B-IT shows, on average, only 1.42 points of degradation, indicating substantial model-to-model differences in sensitivity to underspecified multi-turn interactions in Filipino.

Several tasks exhibit performance degradation from FULL to SHARDED. Among them, Causal Reasoning (CR) stands out as the most strongly affected, with the largest drop in performance, averaged across the five models (35.8 points). In contrast, not all tasks degrade under SHARDED, and some even improve on average, especially MT, SA, and TD. We believe these tasks are episodic, in

which the model can succeed primarily using information available in the current turn, without having to rely heavily on previous turns for additional context [5]. For example, for MT, the model can often translate the text segment shown in the current turn without requiring the rest of the text. For SA and TD, the model may only need to see a key part of the text (e.g., strongly positive/negative words, or explicit toxic terms) to correctly classify the sentiment or toxicity.

We also observe that for some tasks, on average, performance drops more in CONCAT than in SHARDED (e.g., MT, NLI, SA, and TD). One possible reason is that SHARDED gives the model multiple chances to answer: it can make an incorrect guess early, but then correct itself later when more shards are revealed. In contrast, CONCAT is still a single-turn setting, so the model does not get the same opportunity to change its answer across turns. Another possible reason is that these tasks can be sensitive to how instructions are phrased. Since CONCAT uses the rephrased shards, small wording changes introduced during sharding may affect the model’s interpretation of what it should do or what label it should output. This may be especially important for SA, TD, and NLI, where subtle changes in wording (including negation, emphasis, or tone) can affect meaning and, in turn, the predicted label. For MT, small wording differences can also matter because we evaluate using ChrF++ [11], which scores translations based on n-gram overlap with a reference translation; therefore, alternative but valid translations may still receive lower scores if they differ from the reference answer. Although we manually checked and edited the sharded instructions, the nature of these tasks makes them more vulnerable to minor phrasing differences compared to other tasks.

4.2 Aptitude and Unreliability Findings

Figure 3 shows whisker plots of aptitude (A^{90}) and unreliability (U_{10}^{90}) across FULL, CONCAT, and SHARDED for the five evaluated models. Following the Lost in Conversation framework [5], A^{90} reflects best-case performance and U_{10}^{90} measures variability between the 90th and 10th percentile scores. To summarize, we also compute averages across the five models. Under SHARDED, average unreliability increases to $U_{10}^{90} = 16.02$ from $U_{10}^{90} = 4.27$ in FULL, even though average aptitude remains comparable ($A^{90} = 86.22$ in SHARDED vs. 85.96 in FULL). This is an approximately 3.8× increase in unreliability, with minimal change in aptitude. In comparison, CONCAT shows a smaller increase in unreliability ($U_{10}^{90} = 6.41$) and slightly lower aptitude on average ($A^{90} = 84.34$).

Across models, A^{90} stays relatively stable across conversation types, suggesting that most models can still achieve strong best-case performance even under underspecified, multi-turn interactions. For instance, GPT-5.1 has identical aptitude in FULL ($A^{90} = 85.0$) and SHARDED ($A^{90} = 85.0$), while Gemini 3 Pro’s aptitude increases from $A^{90} = 85.5$ in FULL to $A^{90} = 88.0$ in SHARDED. In contrast, Qwen-SEA-LION-v4-32B-IT shows a moderate aptitude decrease from $A^{90} = 87.0$ in FULL to $A^{90} = 83.4$ in SHARDED.

However, unreliability increases substantially under SHARDED for all five models, indicating that the multi-turn, underspecified setting makes model performance less consistent across runs. For example, Qwen-SEA-LION-v4-32B-IT increases from $U_{10}^{90} = 3.6$

in FULL to $U_{10}^{90} = 21.9$ in SHARDED, and GPT-5.1 increases from $U_{10}^{90} = 3.9$ to $U_{10}^{90} = 16.4$.

Compared to SHARDED, CONCAT produces smaller changes in unreliability, suggesting that rephrasing alone does not fully explain the increased unreliability observed in SHARDED. For example, Gemini 3 Pro shows a slight improvement in CONCAT compared to FULL ($U_{10}^{90} = 4.0 \rightarrow 3.6$), whereas it still shows higher unreliability in SHARDED ($U_{10}^{90} = 9.2$). These results support the averaged performance findings by showing that the main impact of underspecified, multi-turn conversations is not only reduced averaged performance, but also increased variability across runs, consistent with the Lost in Conversation phenomenon [5].

4.3 RECAP and SNOWBALL Experiments

To explore simple mitigations for the Lost in Conversation phenomenon [5], we evaluate RECAP and SNOWBALL on Qwen-SEA-LION-v4-32B-IT, the model with the largest degradation from FULL to SHARDED in the main experiment. Both strategies introduce redundancy by repeating previously revealed information: RECAP adds a final turn that restates all shards, while SNOWBALL repeats all prior shards at every turn, gradually increasing context length [5].

Table 5 shows that both strategies provide partial mitigation compared to SHARDED, but their effects are uneven across tasks. On average, SNOWBALL improves overall performance from 71.02 (SHARDED) to 74.34, while RECAP does not improve over SHARDED overall (70.09 vs. 71.02). This suggests that repeating information can help Qwen-SEA-LION-v4-32B-IT recover in some settings, but the benefit depends on the task and how the model uses the extra context.


The tasks that see the greatest improvement are those that require integrating information across turns. For example, Causal Reasoning (CR) improves substantially from 36.00 in SHARDED to 52.50 in RECAP and 57.00 in SNOWBALL. Similarly, Question Answering (QA) improves from 82.50 (SHARDED) to 88.50 in both RECAP and SNOWBALL. These improvements are consistent with the motivation behind RECAP/SNOWBALL: repeating previously mentioned details helps keep the important context visible, so the model does not have to rely as much on remembering and combining information from previous turns [5].

However, RECAP and SNOWBALL do not consistently help, and in some tasks, they can even reduce performance. For Sentiment Analysis (SA), RECAP drops to 89.33 compared to 98.67 in SHARDED, and for Toxicity Detection (TD), RECAP drops sharply to 67.00 compared to 97.00 in SHARDED. One possible explanation is that adding redundant context can increase prompt length and introduce distractions, causing the model to focus on less relevant parts of the conversation. Another possibility is that, for tasks like SA/TD, once the key indicators are shown, additional repetition offers limited benefit and may instead distract the model.

Finally, while RECAP can be effective in some tasks (e.g., CR and MT), it is less realistic as an intervention because it assumes that the user knows when the final turn occurs and can provide a complete restatement of all requirements [5]. In contrast, SNOWBALL is more realistic because it mimics a user or an agent who repeatedly summarizes what has been said so far, and it yields the best overall

Table 4: Averaged performance ($\bar{P}$). Background color indicates performance degradation from FULL. Rightmost columns report the average point drop vs FULL across all 8 tasks.

Model	FULL								CONCAT								SHARDED								Avg drop	
	AS	CR	MT	NLI	PI	QA	SA	TD	AS	CR	MT	NLI	PI	QA	SA	TD	AS	CR	MT	NLI	PI	QA	SA	TD	Concat	Sharded
Gemma-v4-27B	40.97	100	45.68	82.33	88.5	100	94.67	89.5	37.81	89	51.02	77.67	73.5	95	90	86	35.23	71	47.71	90	98	95.5	96.33	96.5	5.21	1.42
Qwen-v4-32B	39.46	100	65.49	86.33	100	98	100	84.5	38.47	95	47.09	73.33	97.5	100	93.33	79.5	35.39	36	48.89	75.67	94	82.5	98.67	97	6.19	13.21
Gemini 2.5 Pro	45.11	100	54.1	93.67	95	100	95	96.5	38.24	100	60.82	92.33	100	100	89	97.5	38.35	52	66.42	91.67	94.5	94.5	98	95.5	0.19	6.06
Gemini 3 Pro	41.22	100	54.85	90	100	100	96.67	75.5	38.84	100	63.57	92	100	100	96.33	79.5	37.03	91	68.41	90	98	98.5	100	74	-1.50	0.16
GPT-5.1	38.54	99.5	65.75	75.67	100	100	88.33	94.5	37.37	97	66.4	71	99.5	100	90	84.5	36.6	70.5	66.07	76.33	80	96	100	87.5	2.07	6.16

**Figure 3: Whisker plots of aptitude (A^{90}) and unreliability (U^{90}_{10}) across simulation types (FULL, CONCAT, SHARDED) for five models.****Table 5: Qwen-SEA-LION-v4-32B-IT averaged performance ($\bar{P}$) across the five simulation types.**

Task	FULL	CONCAT	SHARDED	RECAP	SNOWBALL
ASU	39.46	38.47	35.39	36.44	35.47
CR	100	95	36	52.5	57
MT	65.49	47.09	48.89	58.8	54.55
NLI	86.33	73.33	75.67	75.67	81.33
PI	100	97.5	94	92.5	92
QA	98	100	82.5	88.5	88.5
SA	100	93.33	98.67	89.33	94.33
TD	84.5	79	97	67	91.5
Avg	84.22	78.03	71.02	70.09	74.34

improvement for Qwen-SEA-LION-v4-32B-IT in our setting (74.34 vs. 71.02 in SHARDED).

5 CONCLUSION

This study investigated whether the Lost in Conversation phenomenon [5], defined as performance degradation and increased unreliability in underspecified, multi-turn settings, also occurs in Filipino. To directly compare single-turn and multi-turn settings on the same tasks, we adapted the Lost in Conversation simulation framework [5] to Filipino by constructing pairs of fully-specified and sharded instructions from Batayan [9], covering eight Filipino NLP tasks across NLU, NLR, and NLG.

Across five multilingual LLMs, we find that multi-turn underspecification can reduce average performance, which we determine is

not caused by just rephrasing using CONCAT simulations. Degradation varies across tasks and models, and tasks that require keeping track of information across turns, such as causal reasoning (CR), are more prone to degradation. In contrast, more episodic tasks may be less affected or even improve due to multiple opportunities to answer within a conversation, consistent with observations in prior work [5]. Besides average performance, our results also show that underspecified, multi-turn settings increase unreliability, meaning model outputs are less consistent, even when their best performance remains similar. Overall, performance decreases by 6.45% on average across models in SHARDED relative to FULL, and increases unreliability by approximately 3.8%.

We further evaluated simple context-repetition mitigation strategies (RECAP and SNOWBALL) on the model with the largest degradation and found that repeating previously revealed information can partially recover performance on some tasks, but the effect is uneven and can introduce trade-offs on others. Overall, these findings suggest that Filipino LLM evaluation should move beyond single-turn benchmarks and explicitly measure performance and reliability in multi-turn, underspecified conversations, especially for tasks that depend on remembering and using context revealed across turns.

Future work includes expanding the Filipino sharded dataset to cover more domains and open-ended tasks, evaluating a broader set of models and mitigation strategies, and improving simulation components such as the answer extractor to further reduce evaluation error. We also recommend incorporating human evaluation to

validate how well simulated underspecification matches real-world Filipino or Taglish conversations.

A VALIDATION OF SHARDED SIMULATION

To validate the sharded conversation simulation environment adapted for Filipino, we manually inspected and annotated 200 simulated sharded conversations, following a similar inspection procedure and criteria to those of Laban et al. [5]. The inspection evaluates five criteria:

- (1) **Shard Fully Revealed:** Whether the the user simulator fully revealed the shard, preserving all information from the original shard.
- (2) **Shard Contextualized:** Whether the user simulator appropriately contextualized each shard, making it sound natural and conversational.
- (3) **Answer-Attempt Detection:** Whether the strategy classifier correctly identified when the assistant made an answer attempt or not.
- (4) **Extraction Success:** Whether the answer extractor successfully identified the exact answer span from the assistant's response when an answer attempt was detected.
- (5) **Overall Success:** Whether the entire simulation was successful, even if there were mistakes.

Table 6 presents the detailed inspection results across all evaluated NLU and NLR tasks. The overall success rate of 95.5% indicates that the adapted simulation environment performs reliably for Filipino sharded conversations. Performance is particularly strong for SA, TD, and NLI tasks (100% success), while CR and QA show slightly lower success rates (88.2% and 87.9%, respectively), primarily due to lower extraction success rates. Specifically, extraction success is 82.9% for CR and 76.2% for QA. Because scoring relies on extracting a discrete answer span, extraction failures can cause otherwise correct responses to be scored as incorrect, which may slightly underestimate measured performance for these tasks. The high scores for shards being fully revealed (98.4%) and contextualized (98.7%) confirm that the user simulator effectively conveys sharded information in natural Filipino or Taglish utterances.

Table 6: Manual inspection results for 200 sharded conversation simulations across NLU and NLR tasks. All values are percentages. Refinement tasks (AS, MT) are excluded as they use a different simulation procedure.

Inspection	All	CR	NLI	PI	QA	SA	TD
Shard Fully Revealed	98.4	97.1	96.2	99.2	99	100	100
Shard Contextualized	98.7	97.1	100	97.5	98.5	100	100
Answer-Attempt Detection	95.2	95.7	100	95.8	87.8	100	100
Extraction Success	89.9	82.9	100	88.4	76.2	100	100
Overall Success	95.5	88.2	100	97.0	87.9	100	100

REFERENCES

- [1] Raphaela Agatha R. Albano, Aaron Patrick H. Basconillo, Anne Maxine M. Dacayan, Jean Louise B. Orme, and Yoshiaki B. Kurata. 2025. Analysis of Generation Z Filipinos' Perception of AI Using Multiple Regression. In *Proceedings of the International Conference on Industrial Engineering and Operations Management* (2025-04-08). IEOM Society International, Rabat, Morocco. doi:10.46254/AF6.20250109
- [2] Ge Bai, Jie Liu, Xingyuan Bu, Yancheng He, Jiaheng Liu, Zhanhui Zhou, Zhuoran Lin, Wenbo Su, Tiezheng Ge, Bo Zheng, and Wanli Ouyang. 2024. MT-Bench-101: A Fine-Grained Benchmark for Evaluating Large Language Models in Multi-Turn Dialogues. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 7421–7454. doi:10.18653/v1/2024.acl-long.401
- [3] Christine Herlihy, Jennifer Neville, Tobias Schnabel, and Adith Swaminathan. 2024. On overcoming miscalibrated conversational priors in LLM-based chatbots. In *Proceedings of the Fortieth Conference on Uncertainty in Artificial Intelligence* (Barcelona, Spain) (UAI '24). JMLR.org, Article 75, 22 pages.
- [4] Wai-Chung Kwan, Xingshan Zeng, Yuxin Jiang, Yufei Wang, Liangyou Li, Lifeng Shang, Xin Jiang, Qun Liu, and Kam-Fai Wong. 2024. MT-Eval: A Multi-Turn Capabilities Evaluation Benchmark for Large Language Models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 20153–20177. doi:10.18653/v1/2024.emnlp-main.1124
- [5] Philippe Laban, Hiroaki Hayashi, Yingbo Zhou, and Jennifer Neville. 2025. LLMs Get Lost In Multi-Turn Conversation. arXiv:2505.06120 [cs.CL] <https://arxiv.org/abs/2505.06120>
- [6] Yubo Li, Xiaobin Shen, Xinyu Yao, Xueying Ding, Yidi Miao, Ramayya Krishnan, and Rema Padman. 2025. Beyond Single-Turn: A Survey on Multi-Turn Interactions with Large Language Models. arXiv:2504.04717 [cs.CL] <https://arxiv.org/abs/2504.04717>
- [7] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, 74–81. <https://aclanthology.org/W04-1013/>
- [8] Lester James Validad Miranda, Elyanah Aco, Conner G. Manuel, Jan Christian Blaise Cruz, and Joseph Marvin Imperial. 2025. FilBench: Can LLMs Understand and Generate Filipino?. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng (Eds.). Association for Computational Linguistics, Suzhou, China, 2496–2529. doi:10.18653/v1/2025.emnlp-main.127
- [9] Jann Railey Montalan, Jimson Paulo Layacan, David Demitri Africa, Richell Isaiah S. Flores, Michael T. Lopez II, Theresa Denise Magsajo, Anjanette Cayabyab, and William Chandra Tjhi. 2025. Batayan: A Filipino NLP benchmark for evaluating Large Language Models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 31239–31273. doi:10.18653/v1/2025.acl-long.1509
- [10] Jann Railey Montalan, Jian Gang Ngui, Wei Qi Leong, Yosephine Susanto, Hamsawardhini Rengarajan, Alham Fikri Aji, and William Chandra Tjhi. 2024. Kalahi: A handcrafted, grassroots cultural LLM evaluation suite for Filipino. In *Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation*, Nathaniel Oco, Shirley N. Dita, Ariane Macalinga Borlongan, and Jong-Bok Kim (Eds.). Tokyo University of Foreign Studies, Tokyo, Japan, 497–523. <https://aclanthology.org/2024.paclic-1.49/>
- [11] Maja Popović. 2017. chrF++: words helping character n-grams. In *Proceedings of the Second Conference on Machine Translation*, Ondřej Bojar, Christian Buck, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, and Julia Kreutzer (Eds.). Association for Computational Linguistics, Copenhagen, Denmark, 612–618. doi:10.18653/v1/W17-4770
- [12] Yosephine Susanto, Adithya Venkatadri Hulagadri, Jann Railey Montalan, Jian Gang Ngui, Xianbin Yong, Wei Qi Leong, Hamsawardhini Rengarajan, Peerat Limkonchotiwat, Yifan Mai, and William Chandra Tjhi. 2025. SEA-HELM: Southeast Asian Holistic Evaluation of Language Models. In *Findings of the Association for Computational Linguistics: ACL 2025*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 12308–12336. doi:10.18653/v1/2025.findings-acl.636
- [13] Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SkeHuCVFDr>
- [14] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-judge with MT-bench and Chatbot Arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (NIPS '23). Curran Associates Inc., Red Hook, NY, USA, Article 2020, 29 pages.
- [15] George Kingsley Zipf. 1949. *Human Behavior and the Principle of Least Effort*. Addison-Wesley, Cambridge, MA.

Mixture of Experts with Soft Nearest Neighbor Loss: Resolving Expert Collapse via Representation Disentanglement

Abien Fred Agarap
abien.agarap@dlsu.edu.ph
De La Salle University
Manila, Philippines

Arnulfo Azcarraga
arnulfo.azcarraga@dlsu.edu.ph
De La Salle University
Manila, Philippines

ABSTRACT

The Mixture-of-Experts (MoE) model uses a set of expert networks that specialize on subsets of a dataset under the supervision of a gating network. A common issue in MoE architectures is “expert collapse” where overlapping class boundaries in the raw input feature space cause multiple experts to learn redundant representations, thus forcing the gating network into rigid routing to compensate. We propose an enhanced MoE architecture that utilizes a feature extractor network optimized using Soft Nearest Neighbor Loss (SNNL) prior to feeding input features to the gating and expert networks. By pre-conditioning the latent space to minimize distances among class-similar data points, we resolve structural expert collapse which results to experts learning highly orthogonal weights. We employ Expert Specialization Entropy and Pairwise Embedding Similarity to quantify this dynamic. We evaluate our experimental approach across four benchmark image classification datasets (MNIST, FashionMNIST, CIFAR10, and CIFAR100), and we show our SNNL-augmented MoE models demonstrate structurally diverse experts which allow the gating network to adopt a more flexible routing strategy. This paradigm significantly improves classification accuracy on the FashionMNIST, CIFAR10, and CIFAR100 datasets.

CCS CONCEPTS

• **Computing methodologies** → **Learning latent representations**; **Ensemble methods**; **Semi-supervised learning settings**;

KEYWORDS

contrastive learning, ensemble learning, mixture-of-experts, representation learning

1 INTRODUCTION AND RELATED WORKS

An artificial neural network solves a given problem by learning to approximate the function that describes the relationship between dataset features x and labels y in a given training data. This is accomplished by training with retro-propagation of its output errors [21] and then adjusting its weights based on such error signals. Neural networks can be used for a plethora of tasks such as image classification [15], language translation [6], and speech recognition [10]. It learns to solve these problems by breaking them down to learn the best representations for the input features, and then performing the given task on the resulting learned representations. This representation learning capability is among the advantages of neural networks over other machine learning algorithms.

1.1 Class Neighborhood Structure

A number of methods have been used to further improve the performance of neural networks on subsequent problem-specific tasks, such as expressing high-dimensional data in low-dimensional space, in which capturing the structure of a given dataset may be purposely learned. This structure contains the class information of a given dataset which indicates how the class-similar data points cluster together in a projected space. Several techniques have been introduced and are popularly used for this purpose such as principal components analysis (PCA), t-stochastic neighbor embedding (t-SNE) [17], triplet loss [5], and the soft nearest neighbor loss [1, 7, 22] to transform the input features to contain attributes that are primed for classification.

The aforementioned techniques capture the underlying class neighborhood structure of an input data. However, each technique has their respective drawbacks. If the most salient features of the data can be found in a nonlinear manifold, then a linear technique like PCA will not be able to fully learn the underlying structure of the data. On the other hand, nonlinear techniques like t-SNE, soft nearest neighbor loss, and triplet loss require relatively expensive computational cost. Not to mention that t-SNE discovers different representations of an input data as a function of its hyperparameters. Meanwhile, the soft nearest neighbor loss and triplet loss can be relatively much slower to compute even when compared to t-SNE, depending on the available compute resources.

Despite the known drawbacks, it is still desirable to learn the underlying structure of a data as it implies how the input features form clusters, and consequently, these clusters imply the class membership of the input features as per the clustering assumption in the semi-supervised learning literature [3]. Although it is established that nonlinear techniques are relatively slower and more computationally expensive to use, the evolution of computational hardware has provided drastic advancements that compute issues might be negligible for those with the necessary resources.

1.2 Ensemble Learning

Aside from learning the class neighborhood structure of a given data, we can improve the performance of neural networks by combining their outputs through averaging or summation [2]. This technique is known as ensemble learning, which enforces *cooperation* among neural networks to solve a common goal. In this context, we define *cooperation* as the phenomenon where the neural networks in an ensemble contribute to the overall success of the group. The neural networks in an ensemble cooperate among themselves by compensating for the performance of one another, thereby decreasing the error correlation of the group.

In contrast, if a training data can be naturally divided into subsets, the group of neural networks can be inspired to specialize on their own subsets rather than to cooperate among themselves. This approach is known as the mixture-of-experts model [13] which uses a gating network as a supervisor to choose which sub-network must be assigned on a subset of the dataset based on their predictive performance on such subsets.

However, if the raw input features fed to the gating network contain highly entangled class boundaries, the gating network struggles to confidently partition the space. This often leads to “expert collapse” where a single expert dominates the task or multiple experts learn redundant representations. In such a case, the gating network often results to hard-routing strategies that fail to generalize well on complex datasets.

1.3 Contributions

To address the issue of expert collapse, we propose an architectural pipeline that primes the input features to the MoE model using SNNL. Our core contributions are:

- (1) **Architectural Enhancement.** We employ a feature extractor regularized with SNNL before the MoE routing phase, thus minimizing intra-class distances to simplify the gating network’s partition task.
- (2) **Quantitative Specialization Metrics.** We address the need for detailed analysis by introducing Expert Specialization Entropy and Pairwise Embedding Similarity to empirically measure expert divergence and routing flexibility.
- (3) **Statistical Rigor and Visual Proof.** We conduct robust non-parametric statistical testing (Wilcoxon signed-rank test) and provide comprehensive manifold and embedding visualizations to definitively prove that SNNL prevents structural collapse, thus allowing experts to be utilized more collaboratively on complex benchmark datasets.

2 BACKGROUND

2.1 Feed-Forward Neural Network

The feed-forward neural network is the quintessential deep learning model that is used to approximate the function mapping between the dataset input and output, i.e. $y \approx f(\vec{x}; \theta)$. Its θ parameters are then optimized to learn the best approximation for the input targets, which may either be a class label (classification) or a real value (regression).

$$f(\vec{x}) = f^{(n)} \left(f^{(\dots)} f^{(1)}(\vec{x}) \right) \quad (1)$$

$$\ell_{ce}(y, f(x)) = - \sum_i y_i \log [f(x_i)] \quad (2)$$

To accomplish this, the model composes a number of nonlinear functions in the form of hidden layers, each of which learns a representation of the input features (see Eq. 1). Afterwards, the similarities between the approximation and the input targets are measured by using an error function such as cross entropy (see Eq. 2), which shall be the basis for optimization usually through gradient-based learning.

For our experiments, we use a feed-forward neural network with a single hidden layer containing 128 units. The hidden layer weights

were initialized with Kaiming initializer [12] and had ReLU [19] activation function while the output layer was initialized with Xavier initializer [9]. We use this architecture for the expert networks in our MoE models. For the gating network, we simply use a linear classifier to select which expert to use for particular input features.

2.2 Convolutional Neural Network

The convolutional neural network (or CNN) is a neural network architecture that uses the convolution operator as its feature extractor in its hidden layers [8, 16]. Like feed-forward neural networks, they also compose hidden layer representations for a downstream task. However, with its use of the convolution operator in its hidden layers, it learns better representations of an input data.




Figure 1: Architecture of the CNN-based Feature Extractor. This module computes disentangled representations that are subsequently routed to the MoE gating network and experts.

In Figure 1, we illustrate the CNN-based feature extractor we use to prime the input features for the gating and expert networks across all our experiments. The network processes raw input features through a sequence of two distinct feature blocks to construct a disentangled latent representation. Each block consists of a 2D convolutional layer utilizing a 3×3 kernel with 32 filters in the first block and 64 filters in the second, then immediately followed by 2D batch normalization, a ReLU activation function to introduce non-linearity, and a 2×2 max pooling layer to progressively downsample the spatial dimensions. After the second block, the resulting multi-channel feature maps are flattened into a concise 1D vector. This final embedding serves as the topologically structured, SNNL-optimized input fed directly into the downstream MoE gating network and expert modules.

2.3 Mixture-of-Experts

The Mixture-of-Experts (MoE) model consists of a set of n “expert” neural networks $E_1, \dots, E_n$ and a “gating” neural network G . Figure 2 is an illustration of the MoE model. The experts are chosen by the gating network to handle a subset of the entire dataset, wherein each subset is tantamount to a sub-region of the data space. The output of this model is given by the following equation,

$$\hat{y} = \sum_{i=1}^n G(x)_i E_i(x) \quad (3)$$




Figure 2: The mixture-of-experts model is a system of experts and gating networks where each expert become a function of a subset of the input environment. These expert networks receive the same inputs and produce the same number of outputs. The gating network also receives the same input as the expert networks, but its output is the probability of choosing a particular expert on a given input.

where $G(x)_i$ is the probability output of the gating network to choose expert E_i for a given input x . The gating network and the experts have their different set of parameters θ . In our classification experiments, we modified the inference function above as follows,

$$\hat{y} = \sum_{i=1}^n \arg \max G(x)_i E_i(x) \quad (4)$$

This modified inference function allows the selection of the best expert network E_i as indicated by the highest gating network output $\arg \max G(x)_i$. Without this modification, the model only results to a model akin to a traditional ensemble model albeit outputs a weighted summation as its output instead of a simple summation of expert outputs.

Subsequently, we optimize the MoE model based on the following error function,

$$\mathcal{L}_{\text{moe}}(x, y) = \frac{1}{b} \sum_{i=1}^n G(x)_i \times \ell_{\text{ce}}(y, E_i(x)) \quad (5)$$

where ℓ_{ce} is the cross entropy function measuring the difference between the target y and the output of expert E_i , while $G(x)_i$ is the probability output for choosing expert i . Then, the loss is averaged over the number of batch samples b .

In this system, each expert learns to specialize on the cases where they perform well, and they are imposed to ignore the cases on which they do not perform well. With this learning paradigm, the experts become a function of a subregion of the data space, and thus their set of learned weights highly differ from each other as opposed to traditional ensemble models that result to having almost identical weights for their learners.

3 RESOLVING EXPERT COLLAPSE WITH DISENTANGLEMENT

3.1 Soft Nearest Neighbor Loss

We define *disentanglement* as how close pairs of class-similar data points from each other are, relative to pairs of class-different data points, and we can measure this by using the soft nearest neighbor loss (SNNL) function [7, 22].

This loss function is an expansion on the original nonlinear neighborhood components analysis objective which minimizes the distances among class-similar data points in the latent code of an autoencoder network [22]. On the other hand, the SNNL function minimizes the distances among class-similar data points in each hidden layer of a neural network [7]. The SNNL function is defined for a batch of b samples (x, y) as follows,

$$\ell_{\text{snn}}(x, y, T) = -\frac{1}{b} \sum_{i=1 \dots b} \log \left(\frac{\sum_{\substack{j \in 1 \dots b \\ j \neq i \\ y_i = y_j}} \exp \left(-\frac{d_{ij}}{T} \right)}{\sum_{\substack{k \in 1 \dots b \\ k \neq i}} \exp \left(-\frac{d_{ik}}{T} \right)} \right) \quad (6)$$

where d is a distance metric on either raw input features or learned hidden layer representations x of a neural network, and T is a temperature parameter that can be used to influence the value of the loss function. That is, at high temperatures, the distances among widely separated data points can influence the loss value.

3.2 Our approach

We propose to use a feature extractor network f for the MoE model instead of using the raw input features for the expert and gating networks (see Figure 3). Using our feature extractor network optimized with SNNL, we prime the input features via disentanglement w.r.t. their class labels. Our proposal thus leads to the optimization of a composite loss (see Eq. 7) of the cross entropy as the classification loss and the SNNL as the regularizer.

$$\mathcal{L}(f, x, y) = \mathcal{L}_{\text{moe}}(x, y) + \alpha \cdot \min_{\ell_{\text{snn}}} \ell_{\text{snn}}(f^i(x), y, T) \quad (7)$$

We found more stable and better optimization of the SNNL when we take its minimum value across the hidden layers f^i of the feature extractor network.

3.3 Theoretical Analysis

To formalize the contribution of the soft nearest neighbor loss (SNNL) to expert specialization, we analyze the interaction between the feature extractor f and the Mixture-of-Experts (MoE) gating mechanism.

3.3.1 Latent Space Partitioning via SNNL. The SNNL objective ℓ_{snn} serves as a topological regularizer on the feature extractor's hidden layers. By minimizing the distance between class-similar data points relative to class-different ones, the feature extractor $f(\vec{x})$ maps the high-dimensional input $\vec{x}$ into a lower-dimensional latent manifold where class clusters are highly localized. Mathematically, the SNNL gradient w.r.t. the weights of f minimizes the pairwise




Figure 3: We optimize the soft nearest neighbor loss over the hidden layers found in the feature extractor network before the MoE model. In doing so, the input features to the expert and gating networks are transformed to a set of representations with the classification information ingrained in them, thereby helping improve the overall classification performance of the MoE model.

distance d_{ij} for all $y_i = y_j$. This results in a latent representation $z = f(\vec{x})$ where the intra-class variance is minimized.

3.3.2 Gating Dynamics and Expert Selection. In a traditional MoE model, the gating network $G(\vec{x})$ often suffers from “expert collapse” where a single expert dominates the gradients due to a disorganized input space. In our approach, the gating network operates on the transformed representation $G(f(\vec{x}))$.

Since $f(\vec{x})$ essentially clusters inputs by class, the gating function $G(z)$ encounters a “pre-partitioned” environment. Let z_c be the centroid of class c in the latent space. The gating network learns a mapping $\phi : z_c \rightarrow \{1, \dots, n\}$, effectively assigning a specific expert E_i to a specific class cluster. As the inference function utilizes $\arg \max G(z)_i$, the model selects the “best” expert for a given sub-region of the data space.

3.3.3 Formalizing the Specialization. Specialization is achieved when the variance of the input distribution seen by expert E_i is significantly lower than the global distribution variance:

$$\text{Var}(f(\vec{x})|G(f(\vec{x}))_i \approx 1) \ll \text{Var}(f(\vec{x}))$$

By ensuring that the input to the gating network is already disentangled by class, our approach forces each expert E_i to learn a localized function $f^{(E_i)}$ specifically tuned to the nuances of the assigned class cluster. This reduces the interference between experts and facilitates higher classification accuracy, as evidenced by the empirical results on the classification datasets used in the study.

4 EXPERIMENTS

We use four benchmark image classification datasets to evaluate our proposed model: MNIST, FashionMNIST, CIFAR10, and CIFAR100. We ran each model five times, and computed their average performance across those runs. We report both the average classification performance and the best classification performance for each of our model. For reproducibility, we used the following set of seeds for the random number generator: 1, 2, 3, 4, and 5. No hyperparameter tuning was done as trying to achieve state-of-the-art performance is beyond the scope of this study, we only intend to show that using a feature extractor optimized with SNNL helps resolve the issue of expert collapse in a mixture-of-experts model. In addition, no other regularizers were used in order to better demonstrate the benefits of SNNL for our feature extractor.

4.1 Evaluation Metrics

To thoroughly evaluate the impact of SNNL and address the limitations of purely accuracy-based metrics, we employed the following analytical framework:

- (1) **Expert Specialization Entropy (ENT).** We calculate the Shannon entropy of the average routing distribution per class. Lower entropy indicates “hard-routing” while higher entropy indicates a more flexible routing distribution across the experts.
- (2) **Pairwise Embedding Similarity (SIM).** We extract the first-layer weight matrices of each expert network and compute their pairwise cosine similarity. Lower similarity indicates that the experts have learned orthogonal and divergent internal representations, thus escaping expert collapse.
- (3) **Statistical Analysis.** To verify if the improvements were robust, all models were trained across multiple random seeds. The differences in metrics were evaluated using the non-parametric Wilcoxon signed-rank test.
- (4) **Manifold Visualization.** We project the high-dimensional feature representations in the feature extractor into 2D space using UMAP [18] to qualitatively assess cluster homogeneity and confirm the prevention of structural collapse.

4.2 Experimental Setup

To ensure our findings are applicable in resource-constrained research environments, all models were trained locally using PyTorch Lightning on an RTX 3060 GPU. Since the SNNL computation is only required during the training phase to shape the feature extractor’s weights, our proposed architectural enhancement adds zero computational overhead during inference.

We trained all our models for 15,000 steps on a mini-batch size of 100 using SGD with momentum (0.9) and weight decay ($1e-4$) [20] with a learning rate of $1e-1$, and we used OneCycleLR to anneal the learning rate [23].

4.3 Classification Performance

Table 1 summarizes the performance across all three core metrics. The experimental models demonstrate significant behavioral shifts compared to the baselines, i.e. trading rigid routing for structural diversity which particularly benefited the more complex datasets.

Dataset	ACC (MEAN (SD) in %)		SIM (MEAN (SD) in x100)		ENT (MEAN (SD) in x100)	
	Baseline	Experimental	Baseline	Experimental	Baseline	Experimental
MNIST	99.36 (0.03) *	99.25 (0.08)	0.03 (0.03)	0.04 (0.06)	0.58 (0.06)	1.94 (0.14)
FMNIST	91.33 (0.13)	91.61 (0.23) *	0.03 (0.07)	-0.01 (0.05)	16.82 (10.88)	16.00 (11.03)
CIFAR10	70.91 (1.09)	71.23 (0.91) (ns)	0.05 (0.11)	0.03 (0.06)	14.00 (15.28)	15.01 (15.52)
CIFAR100	35.75 (1.64)	36.74 (1.59) *	0.20 (0.24)	0.10 (0.15)	43.92 (26.88)	45.74 (26.60)

Table 1: Classification Accuracy (ACC), Embedding Similarity (SIM), and Routing Entropy (ENT) across Baseline and Experimental MoE models. Values represent MEAN (SD). Asterisks (*) denote statistical significance at $p < 0.05$ via Wilcoxon signed-rank test; (ns) denotes not significant.

The Complex Datasets (FashionMNIST and CIFAR100). The experimental model shows a statistically significant improvement over the Baseline in classification accuracy. This confirms that for more difficult and highly entangled feature spaces, priming the latent space with SNNL genuinely aids downstream classification.

The Simple Dataset (MNIST). The baseline model significantly outperforms the experimental model. Since MNIST is an extremely simple and highly separable dataset (baseline accuracy approaches 99.4%), we suspect forcing SNNL clustering likely introduces over-regularization or unnecessary constraints that marginally degrade performance.

4.4 Ablation and Specialization Analysis

Our specialization metrics potentially reveals why the performance on complex datasets improved with our experimental approach. The quantitative data contradicts the assumption that SNNL merely enforces harder routing; instead, SNNL fundamentally prevents structural expert collapse.

Orthogonal Embeddings (Resolving Expert Collapse). On FashionMNIST, CIFAR10, and CIFAR100, the SNNL model achieved noticeably lower embedding similarities. This is profoundly pronounced on CIFAR100, where the baseline experts suffered from high similarity (0.20), indicating structural redundancy and collapse. SNNL halved this redundancy (0.10). On FashionMNIST, it pushed the similarity into the negative (-0.01).

As visualized in Figure 4, the baseline condition exhibits redundancy among experts. By feeding disentangled representations to the MoE, the experts are forced to learn highly distinct and non-overlapping weights, thus cutting the highest off-diagonal similarity by more than half.

Routing Entropy (Flexible Collaboration). Our results show that aside from FashionMNIST, the SNNL model actually increased routing entropy. The SNNL feature extractor makes the latent clusters so well-defined that the gating network does not need to hard route input features. Instead, it employs a more distributed routing strategy (higher entropy) while relying on the structurally diverse, orthogonal experts to handle classification nuances.

Figure 5 visualizes our “flexible collaboration” narrative. On the highly complex CIFAR100 dataset, the SNNL condition elevates the entire interquartile range for accuracy, proving a robust performance boost while simultaneously raising the median entropy. On the intermediate CIFAR10 dataset, the SNNL boxplot reaches a higher maximum accuracy, but the heavily overlapping boxes




Figure 4: Pairwise cosine similarity between expert weight matrices on CIFAR10. (Left) The Baseline shows distinct patches of redundancy (e.g. between Expert 0 and 3). (Right) The SNNL condition visibly suppresses these redundancies, thus forcing experts to learn highly distinct and orthogonal features.

contextualize why the overall gain was not statistically significant. Nevertheless, the underlying behavioral shift towards higher entropy remains intact across both complexity levels.

Manifold Visualization (Validating Disentanglement). The UMAP projections in Figure 6 qualitatively confirm the mechanism driving our quantitative metrics. The baseline feature space (Left) exhibits highly entangled class boundaries, which explains why baseline MoE models suffer from higher embedding similarity: the gating network is forced to send overlapping distributions of data to multiple experts. Conversely, the SNNL feature space (Right) demonstrates that intra-class variance has effectively collapsed, resulting in dense and distinct clusters. This proves that SNNL effectively disentangles the feature representations, thus trivializing the gating network’s routing task and enforcing strict structural specialization.

4.5 SNNL Weight Ablation: Tuning Expert Diversity

To further understand the continuous effect of the SNNL weighting factor (α) on expert orthogonality, we conducted an ablation study using a 25k subset of the CIFAR10 training set. We evaluated 30 models with α values linearly spaced from -100 to 100.

As illustrated in Figure 7, the trend line confirms a negative correlation between the SNNL weight and pairwise embedding




Figure 5: Distribution of Test Accuracy and Routing Entropy across random seeds. (Top/CIFAR100) SNNL provides a robust, highly stable boost to both accuracy and entropy, elevating the entire interquartile range. (Bottom/CIFAR10) While median entropy slightly increases, the wider variance and overlapping accuracy boxes explain the non-significant accuracy gain on this intermediate dataset.




Figure 6: UMAP projection of the latent feature space immediately prior to the MoE gating network on CIFAR10. (Left) The baseline feature extractor produces highly entangled class boundaries. (Right) The SNNL-regularized feature extractor minimizes intra-class distance, producing dense and homogeneous class clusters.

similarity among the experts. While MoE routing exhibits inherent variance across runs, progressively increasing α actively pulls the expected similarity closer to 0. This demonstrates that the SNNL weight acts as a tunable control parameter, i.e. a stronger topological penalty directly forces the downstream experts to learn more distinct, non-overlapping representations.




Figure 7: The effect of the SNNL weight (α) on the pairwise embedding similarity of the experts. The dashed regression line shows a negative correlation, indicating that stronger SNNL penalties actively suppress expert redundancy.

5 RECOMMENDATIONS

While our findings demonstrate the significant advantages of SNNL in resolving expert collapse, several avenues for future research remain to fully realize the potential of topology-aware MoE architectures:

- (1) **Dynamic Parameter Tuning.** Our ablation study highlighted the sensitivity of expert orthogonality to the SNNL weighting factor (α). Future work should explore dynamically adjusting both α and the temperature parameter (T) during training. Utilizing learning rate schedulers to gradually decay or warm up these values could better optimize the transition from global representation clustering to fine-grained cross-entropy classification.
- (2) **Architectural Scaling.** The current experiments were constrained to lightweight CNN variants to accommodate local computing limits. Future investigations should scale this methodology to Vision Transformers (ViTs) and higher-resolution, large-scale datasets (e.g. ImageNet) to verify if the disentanglement benefits scale linearly with model size and extreme class diversity.
- (3) **Comparative Representation Learning.** Direct benchmarking against other modern contrastive learning techniques such as Supervised Contrastive Learning (SupCon) [14] or SimCLR [4] is necessary. This will isolate whether SNNL’s specific layer-wise distance minimization is uniquely suited for MoE gating or if any class neighborhood structure-preserving clustering objective yields similar structural diversity and routing flexibility.

6 CONCLUSION

We proposed and evaluated an architectural modification to the Mixture-of-Experts (MoE) model by integrating a feature extractor optimized via Soft Nearest Neighbor Loss (SNNL). By expanding our evaluation to include specialization metrics like Pairwise Embedding Similarity and Expert Specialization Entropy, and including UMAP manifold visualization, we uncovered a nuanced mechanism of action that addresses the limitations of standard MoE architectures.

Our empirical findings demonstrate that baseline MoE architectures often suffer from structural expert collapse, where experts learn redundant representations which results to rigid routing by the gating network. Disentangling the feature representations with SNNL solves this structural collapse by forcing the experts to learn highly orthogonal and diverse features. Consequently, this liberates the gating network to utilize a more collaborative routing strategy.

While disentanglement may have the tendency to “over-regularize” models on simple datasets like MNIST, it yields statistically significant accuracy improvements on complex or highly entangled datasets such as FashionMNIST, CIFAR10, and CIFAR100, thus effectively bridging the gap between representation learning and ensemble specialization.

REFERENCES

- [1] Agarap, Abien Fred, and Arnulfo P. Azcarraga. “Improving k-means clustering performance with disentangled internal representations.” 2020 International Joint Conference on Neural Networks (IJCNN). IEEE, 2020.
- [2] Breiman, Leo. “Stacked regressions.” *Machine learning* 24.1 (1996): 49-64.
- [3] Chapelle, Olivier, Bernhard Scholkopf, and Alexander Zien. “Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews].” *IEEE Transactions on Neural Networks* 20.3 (2009): 542-542.
- [4] Chen, Ting, et al. “A simple framework for contrastive learning of visual representations.” *International conference on machine learning*. Pmlr, 2020.
- [5] Chechik, Gal, et al. “Large Scale Online Learning of Image Similarity Through Ranking.” *Journal of Machine Learning Research* 11.3 (2010).
- [6] Cho, Kyunghyun, et al. “Learning phrase representations using RNN encoder-decoder for statistical machine translation.” *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 2014.
- [7] Frosst, Nicholas, Nicolas Papernot, and Geoffrey Hinton. “Analyzing and improving representations with the soft nearest neighbor loss.” *International conference on machine learning*. PMLR, 2019.
- [8] Fukushima, Kunihiro. “Neural network model for a mechanism of pattern recognition unaffected by shift in position-Neocognitron.” *IEICE Technical Report, A* 62.10 (1979): 658-665.
- [9] Glorot, Xavier, and Yoshua Bengio. “Understanding the difficulty of training deep feedforward neural networks.” *Proceedings of the thirteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2010.
- [10] Graves, Alex, Abdel-rahman Mohamed, and Geoffrey Hinton. “Speech recognition with deep recurrent neural networks.” 2013 IEEE international conference on acoustics, speech and signal processing. Ieee, 2013.
- [11] He, Kaiming, et al. “Deep residual learning for image recognition.” *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016.
- [12] He, Kaiming, et al. “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification.” *Proceedings of the IEEE international conference on computer vision*. 2015.
- [13] Jacobs, Robert A., et al. “Adaptive mixtures of local experts.” *Neural computation* 3.1 (1991): 79-87.
- [14] Khosla, Prannay, et al. “Supervised contrastive learning.” *Advances in neural information processing systems* 33 (2020): 18661-18673.
- [15] Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. “Imagenet classification with deep convolutional neural networks.” *Advances in neural information processing systems* 25 (2012).
- [16] LeCun, Yann, et al. “Gradient-based learning applied to document recognition.” *Proceedings of the IEEE* 86.11 (1998): 2278-2324.
- [17] Van der Maaten, Laurens, and Geoffrey Hinton. “Visualizing data using t-SNE.” *Journal of machine learning research* 9.11 (2008).
- [18] McInnes, Leland, et al. “UMAP: Uniform Manifold Approximation and Projection.” *Journal of Open Source Software* 3.29 (2018).
- [19] Nair, Vinod, and Geoffrey E. Hinton. “Rectified linear units improve restricted boltzmann machines.” *ICML*. 2010.
- [20] Polyak, Boris T. “Some methods of speeding up the convergence of iteration methods.” *Ussr computational mathematics and mathematical physics* 4.5 (1964): 1-17.
- [21] Rumelhart, David E., Geoffrey E. Hinton, and Ronald J. Williams. *Learning internal representations by error propagation*. California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [22] Salakhutdinov, Ruslan, and Geoff Hinton. “Learning a nonlinear embedding by preserving class neighbourhood structure.” *Artificial intelligence and statistics*. PMLR, 2007.
- [23] Smith, Leslie N., and Nicholay Topin. “Super-convergence: Very fast training of residual networks using large learning rates.” *arXiv preprint arXiv:1708.07120* 5 (2017): 4.

Conversion from Nondeterminism to Determinism at the Rule Level of Spiking Neural P Systems with Structural Plasticity

Zechariah B. Jimenez

Parallel Computing Lab
Department of Computer Science
University of the Philippines Diliman
zbjimenez@up.edu.ph

Francis George C. Cabarle

Parallel Computing Lab
Department of Computer Science
University of the Philippines Diliman
fccabarle@up.edu.ph

ABSTRACT

Membrane computing is a branch of natural computing that saw research into various computational models like the Spiking Neural P System (SNP System for short) and the different variants thereof. These models are centered around the idea of neurons sending each other spikes for computation, and subsequently sending spikes to the environment to denote the output of this computation. In particular, the SNP Systems with Structural Plasticity (SNPSP) are a variant that demonstrates what happens when the synapses between these neurons can be modified (as opposed to the original's constant synapse network), using specific rules and constraints.

A simulation algorithm has been formulated to make it possible to “run” such a system on a computer. The entire system would be given a representation using matrices and vectors, and the resulting data structure would be used in the algorithm to record the state of the model and simulate a computation. The algorithm was then proven to simulate computations faithfully, but had to have the domain reduced by restricting one aspect of nondeterminism. This thus changed the domain from the original Rule Nondeterministic Synapse Nondeterministic (RNSN) SNPSP into the new Rule Deterministic Synapse Nondeterministic (RDSN) SNPSP.

This research aims to provide a way to convert a more lax RNSN SNPSP into the restrictive RDSN SNPSP. This would also help effectively expand the domain of the algorithm back to that of the original model.

CCS CONCEPTS

• Theory of computation → Parallel computing models;

1 BACKGROUND OF THE STUDY

Spiking Neural P Systems (SNP [4, 7]) are an esoteric computational model that simulates computation using inter-neuron communication. Multiple such variants have been studied in the past [2], one of which is the *Spiking Neural P System with Structural Plasticity* (SNPSP [1]) where the connections between such neurons can be modified and controlled using specific rules in the neurons themselves. This was shown to be a powerful model, as it is not just Turing complete but also efficient (time- and space-wise) compared to more “traditional” models.

A matrix representation and simulation algorithm [5] was then created for the sole purpose of implementing an SNPSP on a regular computer and simulating its computation. Although it was shown to be a correct simulation, the domain was restricted due to the immense size of the resulting computational graphs.

The original SNPSP paper [1, Sec. 3] mentions two levels of nondeterminism present in the model; the latter paper [5, Sec. 5] restricts the SNPSP systems to have only one of the two levels of nondeterminism to simplify simulations. Although it was shown to still be Turing complete and is close to the original idea of [1], this is still one item of possible future work mentioned in [5]. Thus we explore here how we convert an SNPSP system from the original two-level nondeterministic form into that of the latter one-level nondeterministic variant.

2 THEORETICAL FRAMEWORK

2.1 Regular Expressions

Definition 1 (Regular Expressions). A Regular Expression R [9] is defined to be any of the following:

- (1) $a \in \Sigma$, for some character a and some alphabet Σ ,
- (2) ϵ representing the empty string,
- (3) $\emptyset$ representing the empty set,
- (4) $R_1 \cup R_2$ is the union of two regular expressions R_1 and R_2 ,
- (5) $R_1 \circ R_2$ is the concatenation of two regular expressions R_1 and R_2 , and
- (6) R_1^* is the star of the regular expressions R_1 .

Regular expressions represent languages containing strings over a given alphabet Σ . Here, 1 represents a language $\{a\}$ with the string a being its sole member. Item 2 represents the language $\{\epsilon\}$ with the sole member being the empty string. In contrast, item 3 contains no string and is the language itself that is instead empty. Items 4 and 5 represent the union and concatenation of two regular expressions representing the union or concatenation of their respective languages. Lastly, item 6 represents the repeated concatenation of a single regular expression with itself for any number of times (including zero).

Regular expressions are used throughout SNP literature, as a means to describe a crucial part of various systems. In our case below we mainly use them to describe the inter-neuron communication within these systems, that are thus over a singleton alphabet $\{a\}$.

2.2 SNPSP Systems

Definition 2 (SNPSP System). A spiking neural P system with structural plasticity (SNPSP system) of degree $m \geq 1$ is a construct of the form $\Pi = (O, \sigma_1, \dots, \sigma_m, \text{syn}, \text{out})$, where:

- (1) $O = a$ is the singleton alphabet (a is called spike)
- (2) $\sigma_1, \dots, \sigma_m$ are pairs $\sigma_i = (n_i, \{\mathcal{R}_i\})$, $1 \leq i \leq m$, called neurons, where $n_i \geq 0$ and $n_i \in \mathbb{N} \cup \{0\}$ represents the initial

spikes in σ_i , $\mathcal{R}_i$ is a finite set of rules of σ_i with the following forms:

- (a) Spiking rule: $E/a^c \rightarrow a$, where E is a regular expression over O , with $c \geq 1$;
- (b) Plasticity rule: $E/a^c \rightarrow \alpha k(i, N_j)$, where $c \geq 1$, $\alpha \in \{+, -, \pm, \mp\}$, $k \geq 1$, $1 \leq j \leq |\mathcal{R}_i|$, and $N_j \subseteq \{1, \dots, m\}$.
- (3) $\text{syn} \subseteq \{1, \dots, m\} \times \{1, \dots, m\}$, with $(i, i) \notin \text{syn}$ for $1 \leq i \leq m$, are the initial set of synapses between neurons;
- (4) $\text{in}, \text{out} \in \{1, \dots, m\}$ indicate the input and output neurons, respectively.

Given neuron σ_i (we can also say neuron i or simply σ_i if there is no confusion), we denote the set of neuron labels which has σ_i as their presynaptic neuron as $\text{pres}(i)$, i.e., $\text{pres}(i) = \{j | j(i, j) \in \text{syn}\}$. Similarly, we denote the set of neuron labels which has σ_i as their postsynaptic neuron as $\text{pos}(i) = \{j | j(j, i) \in \text{syn}\}$.

Plasticity rules are applied as follows. If at time t we have that σ_i has $b \geq c$ spikes and $a^b \in L(E)$, a rule $E/a^c \rightarrow \alpha k(i, N) \in \mathcal{R}_i$ can be applied. The set N is a collection of neurons to which σ_i can connect to (synapse creation) or disconnect from (synapse deletion) using the applied plasticity rule. The rule consumes c spikes and performs a specified plasticity operation, depending on the value of α . Note that for this work, if there is no ambiguity on the source neuron the rule is associated with, we could opt to drop i from the notation and just write $E/a^c \rightarrow \alpha kN$.

The plasticity operation can be one of the following, depending on α :

If $\alpha = +$ and $N - \text{pres}(i) = \emptyset$, or if $\alpha = -$ and $\text{pres}(i) = \emptyset$, then there is nothing more to do, i.e., c spikes are consumed but no synapse is created or removed. For $\alpha = +$: If $|N - \text{pres}(i)| \leq k$, deterministically create a synapse to every σ_l , $l \in N - \text{pres}(i)$. If, however, $|N - \text{pres}(i)| > k$, then nondeterministically select k neurons in $N - \text{pres}(i)$ and create one synapse to each selected neuron.

For $\alpha = -$: If $|\text{pres}(i)| \leq k$, deterministically delete all synapses in $\text{pres}(i)$. If, however, $|\text{pres}(i)| > k$, then nondeterministically select k neurons in $\text{pres}(i)$ and delete each synapse to the selected neurons.

If $\alpha \in \{\pm, \mp\}$, create (respectively, delete) synapses at time t and then delete (respectively, create) synapses at time $t + 1$. Only the priority of application of synapse creation or deletion is changed, but the application is similar to $\alpha \in \{+, -\}$. The neuron is always open from time t through $t + 1$, i.e., the neuron can continue receiving spikes. However, the neuron can only apply another rule at time $t + 2$.

An important note is that for σ_i applying a rule with $\alpha \in \{+, \pm, \mp\}$, creating a synapse always involves an embedded sending of one spike when σ_i connects to a neuron. This single spike is sent at the time the synapse creation is applied. Whenever σ_i attaches to σ_j using a synapse during synapse creation, we have σ_i immediately transferring one spike to σ_j .

If two rules with regular expressions E_1 and E_2 can be applied at the same time, that is, $L(E_1) \cap L(E_2) \neq \emptyset$, then only one of them is nondeterministically chosen and applied. All neurons therefore apply at most one rule in one time step (locally sequential), but all neurons that can apply a rule must do so (globally parallel). Note

that the application of rules in neurons are synchronized, that is, a global clock is assumed.

A system state or configuration of an SNPSP system is based on (a) distribution of spikes in neurons and (b) neuron connections based on the synapse graph syn . We can represent (a) as $\langle s_1, \dots, s_m \rangle$ where s_i , $1 \leq i \leq m$, is the number of spikes contained in σ_i . For (b) we can derive $\text{pres}(i)$ and $\text{pos}(i)$ from syn , for a given σ_i . The initial configuration therefore is represented as $\langle n_1, \dots, n_m \rangle$, with the possibility of a disconnected graph, i.e., $\text{syn} = \emptyset$. A computation is defined as a sequence of configuration transitions from an initial configuration. A computation halts if the system reaches a halting configuration, that is, a configuration where no rules can be applied and all neurons are open. Whether a computation is halting or not, we associate natural numbers $1 \leq t_1 < t_2 < \dots$ corresponding to the time instances when the neuron out sends a spike out to (or when it receives a spike from) the system.

A result of a computation can be defined in several ways in SNP systems literature, but in this work we use the following as in [4]: We only consider the first two time instances t_1 and t_2 that σ_{out} spikes. Their difference, i.e., the number $t_2 - t_1$, is said to be computed by Π . We denote the set of all numbers computed in this manner by Π as $N_2(\Pi)$. The subscript indicates that we only consider the time difference between the first two spikes of σ_{out} . Also note that 0 cannot be computed, being analogous to the empty string in other models of computation.

In $N_2(\Pi)$ the neuron in is ignored, and we refer to this as the generative mode. SNPSP systems in the accepting mode ignore σ_{out} and work as follows: The system begins with an initial configuration, and exactly two spikes are introduced in the system (using σ_{in}) at times t_1 and t_2 . The number $t_2 - t_1$ is accepted by Π if the computation halts. The set of numbers accepted by Π is denoted as $N_{\text{acc}}(\Pi)$. The families of all sets of $N_{\text{acc}}(\Pi)$, with $\alpha \in \{2, \text{acc}\}$, are denoted as $N_{\alpha}\text{SNPSP}$.

As an illustration, an example will be provided as follows. Consider an SNPSP system Π_{ex} shown in Figure 1. Neurons 2, $\text{out} = 3$, A_1 , and A_2 contain only the rule $a \rightarrow a$, and we omit this from writing. In the initial configuration, at some time t , is where only σ_1 has two spikes and σ_3 has only one spike. Neuron 1 is the only neuron with plasticity rules, where $\alpha \in \{+, -\}$, and we have $\text{syn} = \{(2, A_1), (2, A_2), (A_1, 1), (A_2, 1)\}$.




Figure 1: An SNPSP system Π_{ex}

Time	σ_1	σ_2	σ_3	σ_{A_1}	σ_{A_2}	syn
t_0	2	0	1	0	0	syn
t_1	1	0	1 (!)	0	0	$syn \cup \{(1, 3)\}$
t_2	0	0	0 (!)	0	0	syn

Table 1: Computation of Π_{ex} computing $\{1\}$

Time	σ_1	σ_2	σ_3	σ_{A_1}	σ_{A_2}	syn
t_0	2	0	1	0	0	syn
t_1	1	1	0 (!)	0	0	$syn \cup \{(1, 2)\}$
t_2	0	0	0	1	1	syn
t_3	2	0	0	0	0	syn
t_4	1	0	1	0	0	$syn \cup \{(1, 3)\}$
t_5	0	0	0 (!)	0	0	syn

Table 2: Computation of Π_{ex} computing $\{4\}$

Π_{ex} operates as follows: The application of the two rules in σ_1 is deterministic. Nondeterminism in Π_{ex} exists only in selecting which synapses to create to or delete from (i.e. *synapse-level nondeterminism*). Neuron 3 has $n_3 = 1$, so it sends a spike out to the environment, and we label this event as time t_1 . Since $n_1 = 2$, the rule with $\alpha = +$ is applied, so σ_1 nondeterministically selects one (since $k = 1$) of σ_2 or σ_3 to create a synapse. If the synapse (1, 3) is created, then σ_1 sends one spike to σ_3 also at time t_1 . The rule with $\alpha = +$ consumes one spike, so σ_1 now has one spike. At time t_2 , the rule with $\alpha = -$ is applied, and synapse (1, 3) is deleted. The spike that σ_3 received from σ_1 is sent to the environment at time t_2 also. Two spikes are sent out to the environment by Π_{ex} before it halts. The computed number in this case is then $t_2 - t_1 = 1$. The computation of $\{1\}$ by Π_{ex} is shown in Table 1, where (!) means that the output neuron spikes to the environment.

If, however, the synapse (1, 2) is created, σ_1 sends a spike to σ_2 at time t_1 during synapse creation. Neuron 2 then sends one spike each to auxiliary neurons A_1 and A_2 at time t_2 . Also at time t_2 is when the rule with $\alpha = -$ is applied, and (1, 2) is deleted. A_1 and A_2 send one spike each to σ_1 so that σ_1 has two spikes again at time t_3 , as in the initial configuration. As long as σ_1 creates synapse (1, 2) instead of (1, 3), Π_{ex} keeps receiving two spikes in a loop. Notice that if at some time $m > 1$ the synapse (1, 2) is created, it will take σ_1 three time steps for the possibility of creating (1, 3) again. Once (1, 3) is created, one more step is required for σ_3 to spike to the environment for the second and final time. The computation of $\{4\}$, for example, is shown in Table 2. Therefore, from the operation of Π_{ex} we then have $N_2(\Pi_{ex}) = \{1, 4, 7, 10, \dots\} = \{3m + 1 | m \geq 0\}$. Also notice that for Π_{ex} , its two plasticity rules in σ_1 can be replaced by a single plasticity rule: $a^2 \rightarrow \pm 1(1, \{2, 3\})$.

2.2.1 Rule Nondeterminism. As previously stated, we note that in [5] the domain of SNPSP Systems was partially restricted. We define the following for the sake of making a clear distinction on these classes of systems.

Definition 3 (Rule Nondeterministic, Synapse Nondeterministic Systems). An SNPSP system is said to be *Rule Nondeterministic*,

Synapse Nondeterministic if it has rule-level and synapse-level nondeterminism. The former means that at any given timestep in its computation, its neurons can nondeterministically select one of its rules for firing, if available. The latter means that its plasticity rules, if any, can nondeterministically select the destination neurons to execute their respective plasticity rules.

On the other hand, an SNPSP system is considered *Rule Deterministic, Synapse Nondeterministic* if only the second level of nondeterminism exists. The first level is thus restricted to deterministic selection only.

We call these classes *RNSN SNPSP Systems* and *RDSN SNPSP Systems*, respectively, or simply RNSN and RDSN Systems.

This means that for an RDSN SNPSP, for any two rules $r_1 : E_1/a^{c_2} \rightarrow pred_1$ and $r_2 : E_2/a^{c_2} \rightarrow pred_2$ we should have $L(E_1) \cap L(E_2) = \emptyset$. Furthermore, in the original algorithm the class of SNPSP systems was restricted to RDSN SNPSP systems only, which we do note is a subset of RNSN SNPSP systems.

2.3 SNPSP Simulation Algorithm

Lastly, [5] lays out a matrix representation and simulation algorithm for SNPSP. The algorithm generates a computation graph containing the states of computation. This is a tripartite graph composed of configuration, rule, and synapse nodes. Furthermore, this is specifically a computation *graph* and not *tree* because it combines identical configuration nodes, allowing for the creation of loops and the interconnection of various paths and traces.

As previously mentioned the algorithm is restricted to only handle RDSN SNPSP systems to reduce the size of the computation graph, while retaining the key characteristic of the plasticity rules nondeterministically selecting their target neurons.

3 OBJECTIVES OF THE STUDY

For this work, we aim to for one main objective, and that is to **develop an algorithm to convert RDSN SNPSP Systems into RNSN SNPSP Systems**, in such a way that there is a direct one-to-one mapping between the computed sets of the two systems. This way, we can be further ensure the completeness of the SNPSP simulation algorithm. In achieving this objective we will be looking into three specific objectives:

- (1) algorithm formulation,
- (2) verification of correctness, and
- (3) algorithm analysis.

4 PRELIMINARY RESULTS

We already have some preliminary results for the RNSN conversion, only missing a fully formalized algorithm and proof of correctness. In this section we detail these results, along with some theoretical foundations used.

4.1 Foundations

4.1.1 Mixed Radix.

Definition 4 (Mixed Radix Numeral). A number N can be represented as a *mixed radix numeral*

$$N = a_{(m)} r_{m-1} a_{(m-1)} r_{m-2} \cdots a_{(2)} r_1 a_{(1)} \\ = \sum_{i=1}^m a_{(i)} \prod_{j=1}^{i-1} r_j$$

where each position can be of differing bases.

Where in *fixed radix form* all “digits” of a number have the same base (e.g. all digits have base 16 in hexadecimal), mixed radix form could have a mix of different bases (thus fixed radix is a special case of mixed radix)[6]. A classic example of this would be with time intervals. A week is 7 days, which is 24 hours each, which in turn is 60 minutes each, which has 60 seconds each. A period of time could be written as $w_7 d_{24} h_{60} m_{60} s$.

4.1.2 Linear and Semilinear Sets. In [1], [5], and often in other SNP papers, regular expressions are limited to the form $E = a^j (a^k)^*$ for integers $j, k \geq 0 \neq j + k$, where the length of the strings a^i in the language $L(E)$ could be described by a linear function or linear progression $jm + kn$.

Observation 1 (Rule Regular Expression Decomposition). For any regular expressions E, F, G , with $E = F \cup G$, $F \cap G \neq \emptyset$, any rule $r_1 : E/a^c \rightarrow \text{pred}$ is equivalent to two rules $r_2 : F/a^c \rightarrow \text{pred}$ and $r_3 : G/a^c \rightarrow \text{pred}$ in the same neuron.

In other words, if neuron $\sigma_i = (n_i, R_i)$ has rule $r_1 \in R_i$, we could modify the neuron into $\sigma'_i = (n_i, R'_i)$ where $R'_i = R_i - \{r_1\} \cup \{r_2, r_3\}$ and the output set of the overall computation would not change.

Definition 5 (Semilinear Set). A *semilinear set* is the finite union of at least one linear set.

Semilinear sets are closed under union and intersection [8], and under set difference [3]. Furthermore, they can be decomposed into a finite number of linear sets. With this, we could generalize the regular expressions we use for rules in SNPSP (and even in SNP in general) to use semilinear sets and still compute the exact same outputs as when using only linear sets.

4.1.3 Rule Notation Changes. With these, we make some changes to the notation of the rules.

First, we note that although there are no forgetting rules in SNPSP systems, we could simulate a forgetting rule using plasticity rules. If a plasticity rule is to be fired but no synapse can be created or deleted, the rule is still considered fired and the spikes are still consumed without any change in the synapse set. This is possible if, for example, in a synapse creation rule $r_+ : E/a^c \rightarrow +k(i, N_+)$, all neurons in N_+ already have a synapse from neuron σ_i . The same is true for a synapse deletion rule $r_- : E/a^c \rightarrow -k(i, N_-)$, all neurons in N_- already don't have a synapse from neuron σ_i .

With that, we could write “forgetting” rules as

$$E/a^c \rightarrow \lambda$$

to refer to an equivalent plasticity rule $E/a^c \rightarrow +1(i, \{j_+\})$ or $E/a^c \rightarrow -1(i, \{j_-\})$ for some j_+, j_- for which $(i, j_+) \in \text{syn}$ but $(i, j_-) \notin \text{syn}$. Note that this gives us the additional flexibility to only apply the “forgetting” rule to the regular expression E and to only consume c spikes, as opposed to the usual case of only being able to




Figure 2: Sample RNSN Π_{Nx}

Time	A	B	C	os	syn	A	B	C	os	syn
0	3	0	1	0	syn	3	0	1	0	syn
1	1	0	1	1	syn $\cup \{(A, C)\}$	2	1	0	1	syn $\cup \{(A, B)\}$
2	1	0	0	2	syn $\cup \{(A, C)\}$	2	0	1	1	syn $\cup \{(A, B)\}$
3						2	0	0	2	syn $\cup \{(A, B)\}$

Table 3: Computation tables of Π_{Nx} , which respectively show the computed values 1 and 2

apply the forgetting rule to a constant number of spikes a^c , which can then be the only number of spikes to be forgotten.

Secondly, we allow a more loose notation of plasticity rules themselves. Recall that formally a plasticity rule is defined to be of the form $E/a^c \rightarrow \alpha k(i, N)$. Recall the previously mentioned omission of i in Section 2.2. In other words, we could write a plasticity rule as

$$E/a^c \rightarrow \alpha kN$$

if there is no ambiguity in the source neuron σ_i associated with the rule. Aside from this, we will further allow the omission of k if $k = |N|$, i.e. all the neurons in N will have a synapse created/deleted. In other words, we could write

$$E/a^c \rightarrow \alpha(i, N)$$

if we will be creating/deleting *all* of N . Of course, we could also write this as just $E/a^c \rightarrow \alpha N$ if i is deemed unnecessary.

4.2 Sample Conversion

For the conversion, we make the following conjecture.

Conjecture 1. For any given RNSN SNPSP System Π that computes the output set $N_2(\Pi)$, there exists an RDSN SNPSP System Π' that computes the set $N_2(\Pi') = \{6n | n \in N_2(\Pi)\}$.

Take for example the RNSN Π_{Nx} in Figure 2. Note that when a neuron doesn't contain a rule in the illustrations, it is assumed to contain just the default rule $a \rightarrow a$. As detailed in Table 3, the output set computed by Π_N is the set $S_N = \{1, 2\}$ since the first output spike is at time 1 and the second is either at time 2 or 3 ($\{2 - 1, 3 - 1\} = \{1, 2\}$).

4.2.1 Conversion Framework. We outline the conversion algorithm in the following way. We convert each neuron σ in an RNSN Π_N (which we dub *simple neurons*) into a group of neurons σ' of an RDSN Π_D (called *complex neurons*) composed of various *modules* of neurons. Furthermore, each *simple spike* in a simple neuron corresponds to a *complex spike* in the equivalent complex neuron. To illustrate this further, we describe the 4 types of modules of a

complex neuron, along with the 3 types of modules that connect these as shown in Figure 3 as such:

- (1) **Dendrite module (D)**. This module holds the complex spikes of the complex neuron. It receives the spikes being sent by neighboring complex neurons, and controls the consumption of spikes by rule firing. It is also this module that ultimately decides what rules to fire.
- (2) **Rule module (R)** and **Rule multiplier module ($\overset{*}{R}$)**. These modules represent a rule and its firing. Once selected by the dendrite module, the rule module sends spikes to the appropriate rule multiplier modules which connect to the axon module described below, as well as connecting back to the dendrite module to communicate which rule was selected. The latter is used to consume the appropriate number of spikes.
- (3) **Axon module (A)**. This module decides which synapses to send spikes through, and takes care of executing plasticity operations. Once the axon finds out from the rule and rule multiplier modules what rule was selected, it either sends a spike through all connected synapses (for spiking rules) or it creates and/or deletes a synapse. Also note that if the simple neuron is the output neuron of Π_N , the corresponding complex neuron's axon module is the output neuron of Π_D .
- (4) **Synapse module (S)** and **Synapse multiplier module ($\overset{*}{S}$)**. These modules correspond to the outgoing synapses of the simple neuron (except the one going out to the environment). When sending a spike to a neighboring complex neuron, the axon module sends spikes to the respective synapse module. The synapse module is also the target neuron for the axon's plasticity rules; in other words the axon would create/delete a synapse to the respective synapse module that connects to the desired neuron. The synapse multiplier module converts a single spike from the axon into the appropriate number of spikes equivalent to a complex spike in the target complex neuron.
- (5) **Clock module (Clock)**. Lastly, the clock module is used to sync all the complex neurons. It is composed of 6 neurons connected in a loop, with the first neuron connected to all dendrites of the entire system.

Note that in Figure 3 (and in subsequent figures), we organize neurons into 3 levels and draw each level differently. We first have neurons at the bottom level, drawn with thin black solid borders. Next up are modules made of neurons, drawn with a thick black dashed borders. Last at the top are complex neurons made of modules, drawn with a thick gray border.

We also draw two kinds of edges: a solid one indicating an existing synapse between two neurons, and a dashed one indicating one that doesn't exist now but could possibly be created later by a plasticity rule.

We will need a new way to refer to the different sets of rules in a neuron. With that, we have:

Definition 6 (Neuron Rule Sets). For a given neuron $\sigma_i = (n_i, R_i)$, we define the following sets of rules:

- (1) $R_{S,i}$ is its set of spiking rules $r_{S,i,j}$,




Figure 3: Rough sketch of conversion from simple neuron σ to complex neuron σ' . Note the inclusion of a separate Clock module.

- (2) $R_{P,i}$ is its set of plasticity rules $r_{P,i,j}$,
- (3) $R_{\alpha,i}$ is its set of plasticity rules $r_{\alpha,i,j}$, with plasticity operation α for all $\alpha \in \{+, -, \pm, \mp\}$,
- (4) $\overset{\gamma}{R}_{\beta,i}$ is its set of *atomic plasticity rules* $\overset{\gamma}{r}_{\beta,i,j}$, for $\beta \in \{\pm, \mp\}$, $\gamma \in \{+, -\}$, where

$$r_{\beta,i,j} = E/a^c \rightarrow \beta k(i, N)$$

$$\overset{\gamma}{r}_{\beta,i,j} = E/a^c \rightarrow \gamma k(i, N)$$


- (5) and finally, $\overset{at}{R}_i$ is its set of *atomic rules* $\overset{a}{R}_i$ and is the set of individual rules that can be executed at a single time step. In other words, it is equal to $\overset{at}{R}_i = \left(\bigcup_{o \in \{S, +, -\}} R_{o,i} \right) \cup \left(\bigcup_{\beta \in \{\pm, \mp\}} \bigcup_{\gamma \in \{+, -\}} \overset{\gamma}{R}_{\beta,i} \right)$

Spikes in the complex neuron are used to communicate various types of information, and we need a way to distinguish among them. These would be mainly used by the dendrite to indicate various kinds of states, which we define as the following:

Definition 7 (Complex Neuron State Classes). We define four main classes of states used in a complex neuron, represented by simple spikes:

- $A \subseteq \{a^n | n \in \mathbb{N}\}$ is the set of possible spike counts.
- $C = \{-, !\}$ is the set of states of the clock ticks – “!” indicates the clock just ticked, while “-” indicates otherwise.
- $F = \{-, \emptyset\}$ is the set of states of rule firing – “ $\emptyset$ ” indicates a rule firing is being carried over from the previous time step, while “-” indicates otherwise. This is mainly for the plasticity rules where $\alpha \in \{\pm, \mp\}$
- $R = \overset{at}{R}_i \cup \{-\}$ is the set of possible atomic rules that can be fired, with “-” indicating no rule being currently fired.

All four of these classes are used by the dendrite and are stored as system state, and are recorded only using simple spikes. We do this by using a mixed radix numeral, storing each class in the dendrite's

Figure 4: The Clock Module of Π_{Dx}

spike count as a different base. These bases can appear in any order, so to not enforce any specific order we will write them as $\frac{A'}{F'} \mid \frac{C'}{R'}$

where $A' \subseteq A$, $C' \subseteq C$, $F' \subseteq F$, and $R' \subseteq R$.

As the spike count will now be represented as a mixed radix numeral, the regular expressions should now represent sets in the same bases. Thus we could write a rule in the dendrite as

$$\frac{A'}{F'} \mid \frac{C'}{R'} \mid \frac{a}{f} \mid \frac{c}{r} \rightarrow pred$$

where $(a, c, f, r) \in A \times C \times F \times R$. Semantically, this now means that A' will be reduced by a , and (C', F', R') will now be changed into (c, f, r) . If any of (a, c, f, r) is left blank, then the rule will not affect the values in their corresponding bases in (A', C', F', R') .

4.2.2 Conversion Proper. Going back to Π_{Nx} as seen in Figure 2, we will go through each neuron and their respective modules as they activate at each time step in the 6-step cycle.

First, we have the Clock module. Simply put, it is a loop of 6 neurons one of which is also connecting to all the dendrites for synchronization. There is only a need for one in Π_{Dx} , as seen in Figure 4.


For the neurons, we begin with σ_A . First of the modules is the dendrite, depicted in Figure 5. We list the possible values of each base in our mixed radix representation of spikes, particularly A and R . A could be assumed to be any integral number of spikes, although with our computations described above we could narrow it down to a few particular values. R is just the two original rules, namely the one creating the synapse (A, B) and the one for (A, C) .

Going back to the dendrite, once receiving a tick from the Clock module, it checks for applicable rules. In the case of A , there are only 2 rules, but their regular expressions not only overlap but are identical. Thus we would only need to check for a^3 once, in only one rule in Π_{Dx} .

We also need to add another rule to check for any other possible spike count when the clock ticks, as this would correspond to the case when no applicable rule exists. Since we would need to keep the clock ticking, the dendrite would need to consume the clock tick (denoted by $!$), without triggering a rule module or otherwise consuming a complex spike. This is seen in the second rule in the

neuron d_A , where checking for an $A - \{a^3\}$ number of spikes and a clock tick ($!$) would just remove the clock tick altogether and proceed as is.

More rules would need to be added to the Dendrite, but they are more relevant to the next step.

Figure 5: The Dendrite Module D_A of A in Π_{Dx}

Next for σ_A is the rule firing step in the Rule and Rule Multiplier modules in Figure 6. Once the previous step selects a rule module to send a spike to, the rule module would send a spike to all the necessary rule multiplier modules to send the right number of spikes both to the Axon Module and right back to the Dendrite Module.



In this case, we could hypothetically imagine that for the complex neuron σ_A , the ordering of the bases starting from the most significant base is (A, F, R, C) . As C only has two possible values ($-$ and $!$), we could send the dendrite two spikes to indicate the atomic rule ① and four for atomic rule ②. Hence, σ_A 's Rule Multipliers have 2 neurons each, with both going back to the dendrite and only one going to the axon.

Once the dendrite receives spikes indicating a rule firing, only then will it consume the spikes associated with the appropriate rule.

After firing a rule comes the Axon and Synapse (and Synapse Multiplier) Modules, as in Figure 7. Once the axon receives any spike, it executes the appropriate *pred* step of its atomic rule. In this case, the atomic rules would both create a synapse but either only to σ_B or only to σ_C . Here, we need to make a Synapse and Synapse Multiplier pair of modules for every synapse that was already in Π_{Nx} or could possibly be created further down the line. In Π_{Dx} , instead of the Axon module A_A connecting straight to complex neuron B or C , it passes through the respective synapse module first.

We choose not to show how many neurons specifically are in each synapse multiplier, but we do note that it should be enough to translate to one complex spike for that respective complex neuron. Also note that each complex neuron could have a different ordering of rules and of the bases (A, C, F, R) .

Once A sends a spike to the dendrite module of either B or C , we are sure that the global Clock module should have ticked already. We illustrate the other two complex neurons B and C in Figure 8. These work similarly to A . Note that since in Π_{Nx} we had C be the output neuron, here we have a_C as the output neuron with an edge going out to the environment.

Figure 6: The Rule and Rule Multiplier Modules of A in Π_{Dx} Figure 7: The Axon, Synapse, and Synapse Multiplier Modules of A in Π_{Dx}


4.2.3 Other Examples. To further illustrate how the conversion would work, we will take a look at two more cases. Unlike with the previous case, in these examples we will only look at converting a single simple neuron into a complex neuron.

First, we consider what happens when we have a rule that spans two timesteps. In our example $\sigma_{\pm}$ in Figure 9, there is only one rule and it has a plasticity operation $\pm$. This rule gets broken down into two atomic rules: (1+) for the initial synapse creation, and (1-) for the subsequent synapse deletion. When converted, the complex neuron $\sigma'_{\pm}$ will instead hold these two rules in the axon.

Once the rule modules send a (1+) back to the dendrite, it is immediately marked with an $\odot$ ("is firing"). On the next clock tick, it then sends out a (1-) for the second half of the original $\pm$ rule.

Lastly, we consider the case where there are multiple intersections between the regular expressions in a neuron's rules. In Figure 11, we see simple neuron σ_{RN} with 3 rules whose regular expressions intersect in a complicated way. To determine the rules to include in the dendrite, we first need to get the intersections of each of the atomic rules' regular expressions. We illustrate this in Figure 10, noting each partition. Once we know what sets are in each partition, we write the corresponding regular expression in the appropriate rules and target the corresponding rule modules accordingly.

As in our example σ_{RN} , we note that the intersection of all three is just the set $\{4\}$, corresponding with the regular expression a^4 . As

Figure 8: The complex neurons B and C in Π_{Dx}

it is part of all the regular expressions, we include all 3 rule modules for ①, ②, & ③. Thus, we include the rule

$$\frac{a^4}{-} \frac{!}{-} \frac{!}{-} \frac{!}{-} \frac{!}{-} \rightarrow \pm 1\{\textcircled{1}, \textcircled{2}, \textcircled{3}\}$$

in the dendrite. We do this for all 8 partitions, making note of which rules include the set and which ones don't. Ignoring all the empty partitions, we end up with exactly the rules in Figure 11.

5 CONCLUSION

We have now shown that we can convert some RNSN SNPSP systems into a RDSN SNPSP with a one-to-one mapping between the two and between their result sets. We have also shown how some edge cases would be handled and ensured a conversion algorithm is indeed possible. All that remains now is for a more general algorithm to be formalized and proven correct. Afterwards, some analysis of the algorithm in time and space, trade-offs, optimizations, and other improvements should be considered especially for implementations and other applications.




Figure 9: Conversion of a simple neuron $\sigma_{\pm}$ with a two-step plasticity rule.




Figure 10: The intersections of the regular expressions of the three rules in σ_{RN} . Note that $j, k, n \geq 0$ are all integers.

REFERENCES

- [1] Francis George C Cabarle, Henry N Adorna, Mario J Pérez-Jiménez, and Tao Song. 2015. Spiking neural P systems with structural plasticity. *Neural Computing and Applications* 26, 8 (2015), 1905–1917.
- [2] Songhai Fan, Prithwineel Paul, Tianbao Wu, Haina Rong, and Gexiang Zhang. 2020. On applications of spiking neural P systems. *Applied Sciences* 10, 20 (2020), 7011.
- [3] Seymour Ginsburg and Edwin H Spanier. 1964. Bounded ALGOL-like languages. *Trans. Amer. Math. Soc.* 113, 2 (1964), 333–368.
- [4] Mihai Ionescu, Gheorghe Păun, and Takashi Yokomori. 2006. Spiking neural P systems. *Fundamenta Informaticae* 71, 2, 3 (2006), 279–308.
- [5] Zechariah B Jimenez, Francis George C Cabarle, Ren Tristan A de la Cruz, Kelvin C Buño, Henry N Adorna, Nestine Hope S Hernandez, and Xiangxiang Zeng. 2019.




Figure 11: Conversion of a simple neuron σ_{RN} with overlapping regular expressions.

Matrix representation and simulation algorithm of spiking neural P systems with structural plasticity. *Journal of Membrane Computing* 1, 3 (2019), 145–160.

- [6] HB Kekre and VR Lakshmi Gorty. 2013. Multiplication of integers in mixed radix number system. *International Journal of Advances in Engineering & Technology* 6, 4 (2013), 1615.
- [7] Alberto Leporati, Giancarlo Mauri, and Claudio Zandron. 2022. Spiking neural P systems: main ideas and results. *Natural Computing* 21, 4 (2022), 629–649.
- [8] Leonhard Y Liu and Peter Weiner. 1970. A characterization of semilinear sets. *J. Comput. System Sci.* 4, 4 (1970), 299–307.
- [9] Michael Sipser. 2013. *Introduction to the theory of Computation*. Cengage Learning.

On the Notions of Soundness in Spiking Neural P Systems

Ronnie M. Ramirez II

Francis George C. Cabarle

Department of Computer Science, University of the Philippines Diliman

Quezon City, Metro Manila, Philippines

{rmramirez4, fccabarle}@up.edu.ph

ABSTRACT

Mimicking the biological process of neural communication, Spiking Neural P Systems are a class of P systems primarily used to simulate processes in membrane computing. Current literature on Spiking Neural P Systems has focused on the development of its theoretical foundation to make it a viable option for visualizing various phenomena. This includes the formalization of various workflow properties such as liveness and boundedness. One property that has yet to be defined is soundness and its various notions, which determine whether a given system is correct without requiring domain knowledge of the system analyzed. That is, a sound system eventually terminates at the desired state, with every element being subject to eventual usage. With this, this study proposes the formalized notions of soundness in the context of Spiking Neural P Systems with one input and output neuron and one initial spike. The relationship between soundness and pre-formalized properties was also explored in the work. Using the formalized notions, a real-world system depicted as a Spiking Neural P System is analyzed based on which soundness notion it satisfies. These results strengthen the potential of Spiking Neural P Systems as a candidate for both computational and analytical applications.

KEYWORDS

Spiking Neural P Systems, Soundness, Workflows

1 INTRODUCTION

According to the Workflow Management Coalition [12], workflows are used to analyze systems. This analysis can be performed by evaluating such workflows through their representation with workflow models against various predefined properties. Serving as a criterion for system correctness, one workflow property is soundness, along with its different notions. In general, soundness equates to the correctness of a workflow, which is defined as the eventual and simultaneous termination of all processes, utilizing all elements in the system [4]. The notion of soundness has already been defined in various workflow models, as it provides a simple yet effective measure to streamline systems and quickly identify errors. One such model is a variation of a Petri Net (PN), which is a literature-rich model created in the 1960s, called a Workflow Net (WF-net), which focuses on representing workflow procedures [1].

Proposed in 2006, the Spiking Neural P System is a system inspired by the biological process of neural communication through spiking [13]. This type of P system is commonly used for mathematical computation in membrane computing, as compiled in [23]. In the same literature, work has also been made to simulate SN P systems using hardware simulators such as GPUs. This allows model simulation using parallel computing, paving the way for

time-efficient computing. Even so, these systems have the potential to serve as a workflow model by modelling various processes. Such a potential was cemented through various studies that bridge SN P systems and PNs [6, 15–17]. This was done by describing the various behavioral and structural profiles of PNs in the context of SN P systems, including different routing types and workflow properties. Based on the study's results, the use of SN P systems became an efficient and viable option for both analytical and computational applications, enabling the combined use of concepts from both ends. Furthermore, future work expanding the theoretical background of SN P systems by using the robust and wide-ranging literature on PN systems, including the concept of soundness, is now possible.

The contributions of this paper are the eight soundness notions for SN P systems. Such formalized notions are obtained by requirements definitions adapted from workflow theory. Several results are proved, such as the liveness and boundedness of a sound SN P system. Lastly, a real-world system is mapped into an SN P system using literature, and then checked on which notion it satisfies.

2 BASIC DEFINITIONS AND NOTATIONS

This section defines the concepts and notations in this study, and is structured on the assumption that the reader has some familiarity with membrane computing, introduced by Păun [18, 19].

2.1 Spiking Neural P Systems

Definition 2.1 (Spiking Neural P System [13, 19, 23]). A *Spiking Neural P System*, shortened as SN P System, of degree m , $m \geq 1$, is a tuple in the form $\Pi = (O, \sigma_1, \dots, \sigma_m, syn, in, out)$, where:

- (1) $O = \{a\}$ is a singleton alphabet, where a is called the spike;
- (2) σ_i , $1 \leq i \leq m$, are neurons of the form $\sigma_i = (n_i, R_i)$, $1 \leq i \leq m$, where
 - $n_i \geq 0$ is the number of spikes in σ_i ;
 - R_i is a finite set of rules of the following forms:
 - (a) Spiking rules $E/a^c \rightarrow a^p; t$ where E is a regular expression over $\{O\}$, t is the delay, $c \geq 1$, $p \geq 1$, such that $c \geq p$, and $t \geq 0$. If $E = a^c$ or $t = 0$, then the corresponding parts are omitted, e.g., $a^c \rightarrow a$;
 - (b) Forgetting rules $a^s \rightarrow \lambda$, for $s \geq 1$, such that for each rule $E/a^c \rightarrow a^p; t$ from R_i of type (1) from R_i , $a^s \notin L(E)$;
- (3) $syn = \{(i, j) | 1 \leq i, j \leq m, i \neq j\}$, i.e., the synapses between neurons;
- (4) $in, out \in \{1, \dots, m\}$ indicate the indices of the input and output neurons, respectively.

In SN P systems, three types of nodes exist: the (*regular*) *neuron* σ , the *input neuron* σ_{in} , and the *output neuron* σ_{out} . In this paper, every SN P system contains an explicit input and output neuron, as these special neurons are often omitted in other studies. All

three nodes can contain spikes a , which can be visualized in the nodes as a number, representing the number of spikes, or as a string. However, only regular neurons can contain *rules* explicitly. Written as a regular expression E , a rule can either be a *spiking rule* or a *forgetting rule*. Each rule requires a certain condition, regarding the number of spikes the neuron has, for it to be enabled.

When the neuron meets the conditions, the rule is applied, which removes a specified number of spikes according to the condition. If the rule is a spiking rule, then a specified number of spikes are fired to other connected neurons. If the applied rule has a *delay* t , where $t > 0$ at timestep k , then the neuron where the rule is becomes closed. That is, the neuron cannot apply any neurons and receive any spikes from k to $k + t - 1$. It only becomes open at time $k + t$ and can apply new rules at the next step. On the other hand, if the rule is a forgetting rule, then no spikes are fired. Nondeterminism can be observed in a given neuron with two rules, either spiking or forgetting, but not a combination, under equal conditions. In this case, the rule to be fired is nondeterministically chosen.




Figure 1: An SN P system.

When an SN P system begins firing from the input node, spikes are passed from one neuron to another, presenting various states of the system. These states are formally called the *system configuration*, but are more commonly known as a configuration C . Defined below, such configurations are mostly represented as vectors.

Definition 2.2 (Configuration Vector of SN P systems [5]). A configuration of an SN P system Π is an m -size vector of integers of the form $C = \langle a_1, a_2, \dots, a_m \rangle$, where a_i , $1 \leq i \leq m$, represents the number of spikes n_i in a neuron σ_i . A configuration at time k , $k \geq 0$, of an SN P system Π , as above, is a vector $C^k = \langle a_1^k, a_2^k, \dots, a_m^k \rangle$, where $a_i^k \in \mathbb{Z}^+ \cup \{0\}$, $1 \leq i \leq m$, represents the number of spikes n_i in a neuron σ_i . If $k = 0$, then the configuration C^0 represents the initial configuration vector of Π .

2.2 Workflow Nets

A *Workflow Net* is a Petri Net used to represent systemic procedures [1]. It is composed of *tokens*, *places* p , *transitions* t , and *arcs*. Shown as a black dot, a token represents a single case through the way it traverses the system. Depending on the arcs, a case can include different tasks and conditions, which are represented, respectively, by the transitions and places. Places contain tokens and can be moved to other places by a transition. Both places or transitions can be either an *input* or an *output*, depending on whether the directed arc is facing away or toward it, respectively. If every input place has a token, the transition connected is *enabled* and *fires*, giving every connected output place a token.

Similar to SN P system configurations, a distribution of tokens within a WF-Net is known as a *marking* M . These markings can also be understood as the states of the model, which can be compared with each other by a partial ordering. Given a pair of arbitrary

markings M_1 and M_2 , then $M_1 \leq M_2$ iff $M_1(p) \leq M_2(p)$ for every p in the net. That is, all places in M_2 have either more or the same number of tokens than they had in M_1 . WF-Nets also have a set of reachability notations between M_1 and M_2 . Shown below is a truncated list for clarity and relevance.

- $M_1 \xrightarrow{*} M_2$: M_2 is generally reachable from M_1 , i.e., there exists some firing sequence $\sigma = t_1, t_2, \dots, t_{n-1}$ that leads from M_1 to M_2 .
- $M_1 \xrightarrow{t} M_2$: when transition t is enabled and fired in M_1 , M_2 is reached, i.e., M_2 is directly reachable from M_1 via t .

WF-Nets differ from PNs by having two special places: a source i that has no input transitions and a sink o that has no output transitions [1, 2]. Each of these places can only exist once, as it visualizes the life cycle of a case in the modeled procedure. Two unique markings are defined with a relation to the source and sink: $[i]$ and $[o]$. Specifically, $[i]$ refers to a marking where all places have no tokens, except for place i , which has one. Conversely, $[o]$ follows similarly, with the exception that o has the single token. Another difference between the two workflow models is that WF-Nets become strongly connected if i and o are connected to a transition t , with o as the input place and i as the output place. That is, when such connections are made, then there exists a path from every place to another. Such a property was needed to ensure that every element in the system is necessary [1].

2.3 Soundness Notions of WF-Nets

This property of soundness guarantees that the given model, WF-Nets in this case, has no anomalies that prevent the system from properly terminating, i.e., livelocks and deadlocks [1, 4]. Initially formulated as a single variation, now known as classical soundness, other types with relaxed requirements exist to accommodate alternative workflows that are also sound.

Definition 2.3 (Classical Soundness of Workflow Nets [1–3]). A WF-Net N is classical sound iff the following requirements are satisfied:

- Option to complete: for every possible marking M from $[i]$, there exists a firing sequence from M to $[o]$, i.e., $\forall M ([i] \xrightarrow{*} M) \Rightarrow (M \xrightarrow{*} [o])$.
- Proper termination: for every possible marking M from $[i]$, $[o]$ is the only reachable marking, i.e., $\forall M ([i] \xrightarrow{*} M \wedge M \geq [o]) \Rightarrow (M = [o])$.
- Liveness: every transition t is an element of at least one firing sequence, i.e., $\forall t \in T \exists M, M' [i] \xrightarrow{*} M \xrightarrow{t} M'$.

As defined above, classical soundness is the strongest notion of the eight formalized notions, as this variation requires three sub-requirements. The option to complete requires the input WF-Net to have a path towards $[o]$ for every possible state, ensuring that the case handled by the system eventually ends. Next, proper termination adds to the previous sub-property by requiring that for all cases, once o receives a token, no other places must have any tokens. This equates to the system not having any pending or ongoing cases once the terminating state is entered. Lastly, liveness requires every transition to be used at some point, meaning that every task of the system can be done in some instance. The net

shown in Figure 2 satisfies these three, thus making it classical sound. Other notions are weakened forms of some sort by loosening, if not fully removing, a combination of the three requirements.




Figure 2: A WF-Net that satisfies classical soundness.

Definition 2.4 (Relaxed Soundness of Workflow Nets [7, 8]). A WF-Net N is relaxed sound iff the following requirements are satisfied:

- Option to complete (weakened): for at least one marking M from $[i]$, there exists a firing sequence from M to $[o]$, i.e., $\exists_M([i] \xrightarrow{*} M) \Rightarrow (M \xrightarrow{*} [o])$.
- Proper termination (weakened): for at least one marking M from $[i]$, $[o]$ is the only reachable marking, i.e., $\forall_M([i] \xrightarrow{*} M \wedge M \geq [o]) \Rightarrow (M = [o])$.
- Liveness: every transition t is an element of at least one firing sequence, i.e., $\forall_{t \in T} \exists_{M, M'} [i] \xrightarrow{*} M \xrightarrow{t} M'$.

Relaxed soundness requires the same requirements that classical soundness needs. The main difference between the two notions is that relaxed soundness does not require all possible instances to have the option to complete and properly terminate. Instead, it simply needs that every transition is a part of some sound sequence of markings [7]. An example of a net that satisfied relaxed soundness is Figure 3, as every transition fires at some point. It does not satisfy classical soundness, however, as the firing sequence $\langle t_1, t_3, t_4, t_6 \rangle$ is possible. This leads to a marking where p_4 and o both have a token, violating the first two classical requirements but not their weakened variants. Similar to Figure 3, Figures 4-6, based on [4], illustrate the other notions of soundness to be defined.




Figure 3: A WF-Net that satisfies relaxed soundness.

Definition 2.5 (Weak Soundness of Workflow Nets [14]). A WF-Net N is weak sound iff the following requirements are satisfied:

- Option to complete: for every possible marking M from $[i]$, there exists a firing sequence from M to $[o]$, i.e., $\forall_M([i] \xrightarrow{*} M) \Rightarrow (M \xrightarrow{*} [o])$.

- Proper termination: for every possible marking M from $[i]$, $[o]$ is the only reachable marking, i.e., $\forall_M([i] \xrightarrow{*} M \wedge M \geq [o]) \Rightarrow (M = [o])$.

Based on the definitions, weak soundness is looser than classical soundness as liveness is no longer required. In detail, classically sound WF-nets require that every transition is fired during simulation, unlike weak sound WF-Nets. Instead, it simply needs paths from the source to the sink for every possible state and no pending processes when a token reaches that sink within the WF-Net. This notion is satisfied by the net in Figure 4, as the net completes and terminates properly with t_3 never firing.




Figure 4: A WF-Net that satisfies weak soundness.

Definition 2.6 (Lazy Soundness of Workflow Nets [20, 21]). A WF-Net N is lazy sound iff the following requirements are satisfied:

- Option to complete: for every possible marking M from $[i]$, there exists a firing sequence from M to $[o]$, i.e., $\forall_M([i] \xrightarrow{*} M) \Rightarrow (M \xrightarrow{*} [o])$.
- Proper termination (weakened): for at least one marking M from $[i]$, $[o]$ is the only reachable marking, i.e., $\forall_M([i] \xrightarrow{*} M \wedge M \geq [o]) \Rightarrow (M = [o])$.

Lazy soundness alleviates the original definition of soundness by dropping liveness as a requirement, like weak soundness. Additionally, this notion adopts the weakened version of proper termination defined in relaxed soundness. Lazy sound systems can have pending and ongoing tasks, even when o gains a token, signifying termination. The only strict requirement is that o receives exactly one token in all cases. This is shown in Figure 5: p_3 and t_3 form an infinite loop, whereas o eventually receives one token.




Figure 5: A WF-Net that satisfies lazy soundness.

Definition 2.7 (Easy Soundness of Workflow Nets [22]). A WF-Net N is easy sound iff the following requirement is satisfied:

- Option to complete (weakened): for at least one marking M from $[i]$, there exists a firing sequence from M to $[o]$, i.e., $\exists_M([i] \xrightarrow{*} M) \Rightarrow (M \xrightarrow{*} [o])$.

This notion of soundness is considered the loosest notion among the eight, as it only requires that the input WF-Net have the weakened option to complete. That is, there exists at least one traversal where one token from the source reaches the sink place. As it does not require proper termination and liveness, other places may have tokens even when the final place gains a token. Furthermore, transitions that will not fire in any circumstance may exist within the model. Figure 6 satisfies easy soundness through the firing sequence $\langle t_1, t_2 \rangle$, reaching $[i]$ to $[o]$.




Figure 6: A WF-Net that satisfies easy soundness.

Tokens in WF-Nets can be understood as cases; thus, the entire system only considers the termination of one case. Other soundness notions, such as k -soundness as defined below, do not necessarily follow this restriction, as the number of tokens at o may depend on the number of tokens at i . With this, the notations for the initial and final states are extended as $[i^k]$ and $[o^k]$, where k represents the number of tokens in i and o , respectively.

Definition 2.8 (k -Soundness of Workflow Nets [10, 11]). A WF-Net N is k -sound, where $\exists k \in \mathbb{N}$ iff the following requirements are satisfied:

- Option to complete: for every possible marking M from $[i^k]$, there exists a firing sequence from M to $[o^k]$, i.e., $\forall M([i^k] \xrightarrow{*} M) \Rightarrow (M \xrightarrow{*} [o^k])$.
- Proper termination: for every possible marking M from $[i^k]$, $[o^k]$ is the only reachable marking, i.e., $\forall M([i] \xrightarrow{*} M \wedge M \geq [o]) \Rightarrow (M = [o])$.

This notion does not require liveness similar to weak soundness. The difference is that proper termination implies that the number of tokens that reach o must also be the number of input tokens. Following this, weak soundness is equal to k -soundness given that $k = 1$ [4]. Although not mentioned in the original definition, proper termination of k -sound WF-Nets is implied by the required option to complete. Here, the second requirement is explicitly mentioned to avoid any ambiguity. The following soundness notions, which will be discussed, are based on the definition k -soundness.

Definition 2.9 (Up-to- k -Soundness of Workflow Nets [22]). A WF-Net N is up-to- k -sound, where $\forall \ell \in \mathbb{N}, 1 \leq \ell \leq k \in \mathbb{N}$, iff the following requirements are satisfied:

- Option to complete: for every possible marking M from $[i^k]$, there exists a firing sequence from M to $[o^\ell]$, i.e., $\forall M([i^\ell] \xrightarrow{*} M) \Rightarrow (M \xrightarrow{*} [o^\ell])$.
- Proper termination: for every possible marking M from $[i^\ell]$, $[o^\ell]$ is the only reachable marking, i.e., $\forall M([i^\ell] \xrightarrow{*} M \wedge M \geq [o^\ell]) \Rightarrow (M = [o^\ell])$.

The notion of up-to- k -soundness can be thought of as k -soundness of a given range starting with $k = 0$.

Definition 2.10 (Generalized Soundness of Workflow Nets [10, 11]). A WF-Net N is generalized sound iff the following requirements are satisfied:

- Option to complete: for every possible marking M from $[i^k]$, there exists a firing sequence from M to $[o^k]$, i.e., $\forall M([i^k] \xrightarrow{*} M) \Rightarrow (M \xrightarrow{*} [o^k])$.
- Proper termination: for every possible marking M from $[i^k]$, $[o^k]$ is the only reachable marking, i.e., $\forall M([i] \xrightarrow{*} M \wedge M \geq [o]) \Rightarrow (M = [o])$.

As the name suggests, this notion of soundness is a more general form of k -soundness. That is, this notion considers all amounts of tokens in the source place. It can also be described as up-to- ∞ -soundness [4], as there is no limit to the quantity of input tokens. Suppose that the net in Figure 4 has two tokens in i instead of one. This results in o gaining two tokens, making the net 2-sound. This behavior is consistent for every $k \in \mathbb{N}$, which is equal to generalized soundness [4]. In a more practical sense, if a workflow is observed to be generalized sound, the system handles as many cases simultaneously without any unexpected terminations or abnormalities during execution.

3 MAPPING OF PETRI NET STRUCTURES AND BEHAVIORS USING SN P SYSTEMS

Literature regarding the relationship between PNs and SN P systems mostly focused on transforming SNPs to PNs. In detail, studies such as [15–17] introduce various mapping methods to help simulate the behavior of an SNP system. However, a subsequent study concentrated on the reverse path: transforming PNs to SN P systems. This study described various structural and behavioral profiles of SN P systems, using PNs as a basis for such concepts [6].

It was observed that the firing of tokens from one place to another in PNs is similar to the spiking of neurons in SN P systems, making a token conceptually equal to a spike. A difference between the two models is that the two PN elements, places and transitions, are equal to one SNPS element, neurons. Specifically, neurons have both the capability to contain spikes, the main function of places, and to transfer them to other nodes, the main function of transitions. With these similarities, the authors defined the various PN routing types in the context of SN P systems by replicating key structures using SN P elements, as shown in Figure 7.

Various SN P system properties were also defined based on the PN properties, e.g., liveness, boundedness, and safeness. Shown in Definition 3.1, these definitions are consistent with the PN counterparts in [1], where a live and bounded PN was proven to be sound and vice versa. Like most workflow properties, the structures affect whether the properties were satisfied, thus requiring structural analysis to check if any anomalous patterns exist.

Definition 3.1 (Live, Bounded, Safe SN P Systems [6]). An SN P system Π is **live** iff for every reachable configuration C^k from C^0 and every rule r in Π , there exists a configuration C^j reachable from C^k wherein r is applied. Π is **k -bounded** iff for every configuration, every neuron has at most k spikes, where $k \in \mathbb{N}$. If $k = 1$, then we say Π is safe.




Figure 7: Petri Net structures and their respective SN P system counterparts: (a) sequential, (b) AND-join, (c) OR-join, (d) AND-split, (e) OR-split (based on [6]).

4 METHODOLOGY AND RESULTS

This section formalizes the notions of soundness in the context of SN P systems. To ensure consistency between the definitions in the two models, the PN to SN P system mapping literature [6], as discussed earlier, is used to identify structural and behavioral parallels required to define the various soundness notions discussed in Section 2.3.

For ease of notation, we define two essential configurations used to determine the soundness of a workflow model: C_{in} and C_{out} . Similar to states $[i]$ and $[o]$ of WF-Nets, C_{in} and C_{out} represent configurations of the SN P system where only one spike exists in the entire model. Such a spike is in the input neuron σ_{in} where $in = 1$ or output neuron σ_{out} where $out = m$, m is the index of the last neuron, respectively, i.e., $C_{in} = \langle 1, 0, \dots, 0 \rangle$ and $C_{out} = \langle 0, 0, \dots, 1 \rangle$. Additionally, partial ordering and reachability notations of WF-Nets are adopted to compare configurations and their reachability in SN P Systems. Listed below are the following notations:

- $C_1 \leq C_2$: Every neuron in C_1 contains fewer or equal spikes than their counterparts in C_2 .
- $C_1 \geq C_2$: Every neuron in C_1 contains more or equal spikes than their counterparts in C_2 .
- $C_1 = C_2$: Every neuron in C_1 has the same number of spikes as their counterparts in C_2 .
- $C_1 \xrightarrow{*} C_2$: C_2 is generally reachable from C_1 , i.e., there exists a sequence of valid configurations that leads from C_1 to C_2 .
- $C_1 \xrightarrow{r} C_2$: when rule r is enabled and applied in C_1 , the model reaches C_2 , i.e., C_2 is directly reachable from C_1 through r .

Definition 4.1 (Classical Soundness of SN P Systems). Let Π be an SN P system, where $\Pi = (O, \sigma_1, \dots, \sigma_m, syn, in, out)$, $in = 1, out = m$, and $\sigma_i = (\alpha_i, R_i)$, $1 \leq i \leq m$. Π is classical sound iff the following requirements are satisfied:

- (1) Option to Complete: For every possible configuration C from C_{in} , there exists a sequence of valid configurations from C to C_{out} , i.e., $\forall C (C_{in} \xrightarrow{*} C) \Rightarrow (C \xrightarrow{*} C_{out})$.
- (2) Proper termination: for every possible configuration C from C_{in} , C_{out} is the only reachable configuration, i.e., $\forall C (C_{in} \xrightarrow{*} C \wedge C \geq C_{out}) \Rightarrow (C = C_{out})$.
- (3) Liveness: every rule r must be applied in some reachable configuration, i.e., $\forall r \in R \exists C, C' C_{in} \xrightarrow{*} C \xrightarrow{r} C'$.

As defined in the notion of classical WF-Net soundness, the same notion in the context of SN P systems also requires three requirements. First, the system must have the option to complete, that is, every possible configuration of the model, starting from the initial configuration of C_{in} , must lead, either directly or indirectly through a sequence of configurations, to the final configuration C_{out} . Second, the system must complete properly, which means that once a spike reaches σ_{out} , all other neurons are empty. Lastly, the system must not have any dead rules; that is, every rule is used in a possible configuration. Figure 1 is an example of an SN P system that satisfies the requirements of classical soundness.

Liveness and boundedness can be correlated to the property of soundness [1]. Although both properties have been defined in the context of SN P systems [6], the concept of extended SN P systems needs to be defined to help prove the connection between the profiles, similar to what was done for WF-Nets.

Definition 4.2 (Extended SN P System). Let Π be an SN P system with an input neuron in and an output neuron out . An extended SN P System $\bar{\Pi}$ of SN P system Π is when in and out are transformed into regular neurons, with each one gaining the spiking rule $a^+ / a \rightarrow a; 0$. Lastly, a synapse is added from out to in , allowing spiking from the end of the SN P system to the start.




Figure 8: An extended SN P system $\bar{\Pi}$.

With extended SN P systems, the liveness and boundedness of classical SN P systems can be explored and proven through the following lemmas, with their proofs patterned from [1].

LEMMA 4.3 (CLASSICAL SOUNDNESS OF A LIVE AND BOUNDED SN P SYSTEM). If $\bar{\Pi}$ starting at C_{in} is live and bounded, then Π is sound.

PROOF. For Π to be sound, it must have the option to complete, properly terminate, and be live. Since σ_{out} connects to σ_{in} and has a firing rule, a configuration where $\sigma_{out}(\alpha_{out}) \geq 1$ is reachable from any possible configuration C from C_{in} , i.e., the option to complete holds. Consider an arbitrary configuration C'_{out} where $\sigma_{out}(\alpha) \geq 1$. Here, σ_{out} spikes and C'_{in} where $\sigma_{in}(\alpha) \geq 1$ is reached. Because $\bar{\Pi}$ is bounded, $C'_{in} \geq C_{in}$ implies $C'_{in} = C_{in}$. This tells us that every place other than σ_{in} in C'_{in} is empty, satisfying proper termination. As $\bar{\Pi}$ is simply Π with an output neuron that spikes to the input neuron, Π is also live. Therefore, Π is sound. $\square$

LEMMA 4.4 (BOUNDEDNESS OF A CLASSICAL SOUND SN P SYSTEM). *If Π is sound, then $\bar{\Pi}$ starting at C_{in} is live.*

PROOF. Assume that Π is sound and that it is not bounded when it starts from C_{in} . With this, two configurations C^i and C^j exist such that $C_{in} \xrightarrow{*} C^i \xrightarrow{*} C^j$ and $C^j > C^i$. Because Π is sound, there exists a sequence of valid configurations from C^j to C_{out} , thus proving the existence of a configuration C^l where $C^k \rightarrow C^l$ and $C^l > C_{out}$. Consequently, Π cannot be both sound and not bounded, making Π starting at C_{in} bounded. Similarly, $\bar{\Pi}$ is also bounded as it returns to C_{in} when σ_{out} fires. $\square$

LEMMA 4.5 (LIVENESS OF A CLASSICAL SOUND SN P SYSTEM). *If Π is sound, then $\bar{\Pi}$ starting at C_{in} is live.*

PROOF. Assume that Π is sound. This makes every valid configuration in $\bar{\Pi}$ from C_{in} reach C_{in} . Because Π is live, an arbitrary rule r can be fired, which can be applied to $\bar{\Pi}$. Thus, $\bar{\Pi}$ is live as every valid configuration from C_{in} reaches a configuration where an arbitrary r is enabled and fired. $\square$

THEOREM 4.6 (BOUNDEDNESS AND LIVENESS WITH CLASSICAL SOUND SN P SYSTEMS). *An SN P System Π is classical sound iff $\bar{\Pi}$ is live and bounded.*

PROOF. Follows from Lemma 4.3, Lemma 4.4, and Lemma 4.5. $\square$

As the relationship of classical soundness with a live and bounded SN P system is proven, we move on to define other soundness notions of SN P systems.

Definition 4.7 (Relaxed Soundness of SN P Systems). Let Π be an SN P system, where $\Pi = (O, \sigma_1, \dots, \sigma_m, syn, in, out)$ and $\sigma_i = (\alpha_i, R_i)$, $1 \leq i \leq m$. Π is relaxed sound iff the following requirements are satisfied:

- (1) Option to Complete (weakened): For at least one possible configuration C from C_{in} , a sequence of valid configurations from C to C_{out} exists, i.e., $\forall C (C_{in} \xrightarrow{*} C) \Rightarrow (C \xrightarrow{*} C_{out})$.
- (2) Proper termination (weakened): for at least one possible configuration C from C_{in} , C_{out} is the only reachable configuration, i.e., $\forall C (C_{in} \xrightarrow{*} C \wedge C \geq C_{out}) \Rightarrow (C = C_{out})$.
- (3) Liveness: every rule r must be applied in some reachable configuration, i.e., $\forall r \in R \exists C, C' C_{in} \xrightarrow{*} C \xrightarrow{r} C'$.

Alleviating the first two requirements gives relaxed soundness, which implies that the system can have at least one classically sound simulation. Such is the case, given that every rule in that instance is applied. Figure 9 shows a relaxed sound SN P system, as every rule can be applied, satisfying liveness. However, every rule does not always lead to C_{out} . Specifically, only when the spiking rules are fired is when the desired final configuration is reached, as the forgetting rules remove the spikes, leading to an empty system.

Definition 4.8 (Weak Soundness of SN P Systems). Let Π be an SN P system, where $\Pi = (O, \sigma_1, \dots, \sigma_m, syn, in, out)$, $in = 1, out = m$, and $\sigma_i = (\alpha_i, R_i)$, $1 \leq i \leq m$. Π is weak sound iff the following requirements are satisfied:

- (1) Option to Complete: For every possible configuration C from C_{in} , there exists a sequence of valid configurations from C to C_{out} , i.e., $\forall C (C_{in} \xrightarrow{*} C) \Rightarrow (C \xrightarrow{*} C_{out})$.




Figure 9: A SN P system that satisfies relaxed soundness.

- (2) Proper termination: for every possible configuration C from C_{in} , C_{out} is the only reachable configuration, i.e., $\forall C (C_{in} \xrightarrow{*} C \wedge C \geq C_{out}) \Rightarrow (C = C_{out})$.

For this notion, liveness is not required, seen in Figure 10 with the inability to fire the second spiking rule in σ_2 . However, the first two requirements are still observed with the passing of the sole spike from σ_{in} to σ_{out} without leaving a spike anywhere else.




Figure 10: A SN P system that satisfies weak soundness.

Definition 4.9 (Lazy Soundness of SN P Systems). Let Π be an SN P system, where $\Pi = (O, \sigma_1, \dots, \sigma_m, syn, in, out)$, $in = 1, out = m$, and $\sigma_i = (\alpha_i, R_i)$, $1 \leq i \leq m$. Π is lazy sound iff the following requirements are satisfied:

- (1) Option to Complete: For every possible configuration C from C_{in} , there exists a sequence of valid configurations from C to C_{out} , i.e., $\forall C (C_{in} \xrightarrow{*} C) \Rightarrow (C \xrightarrow{*} C_{out})$.
- (2) Proper termination (weakened): for at least one possible configuration C from C_{in} , C_{out} is the only reachable configuration, i.e., $\forall C (C_{in} \xrightarrow{*} C \wedge C \geq C_{out}) \Rightarrow (C = C_{out})$.

Similar to the lazy sound WF-Net shown before, the SN P system in Figure 11 can induce an infinite loop. Here, the production of a spike between σ_2 and σ_4 can continue without end if the forgetting rule in σ_2 is never chosen, which breaks the original definition of proper termination. Lazy soundness allows this, given that there is an instance with no such loop before σ_{out} gains one spike.




Figure 11: A SN P system that satisfies lazy soundness.

Definition 4.10 (Easy Soundness of SN P Systems). Let Π be an SN P system, where $\Pi = (O, \sigma_1, \dots, \sigma_m, \text{syn}, \text{in}, \text{out})$, $\text{in} = 1, \text{out} = m$, and $\sigma_i = (\alpha_i, R_i)$, $1 \leq i \leq m$. Π is easy sound iff the following requirement is satisfied:

- (1) Option to Complete (weakened): For at least one possible configuration C from C_{in} , a sequence of valid configurations from C to C_{out} exists, i.e., $\forall_C (C_{\text{in}} \xrightarrow{*} C) \Rightarrow (C \xrightarrow{*} C_{\text{out}})$.

Figure 12, shown below, exhibits the weakest notion of soundness, easy soundness, as proper termination and liveness are not possible within this system. That is, some rules are never enabled and consequently fired, such as the forgetting rules in σ_2 and σ_3 due to their conditions. Furthermore, σ_4 has the non-deterministic choice to either fire a spike to σ_{out} with or without leaving a spike in σ_4 , of which the former option violates proper termination. The latter option allows the system to satisfy the weakened form of the option to complete, thus labeling Π as easy sound.




Figure 12: A SN P system that satisfies easy soundness.

Moving towards the notions that consider multiple spikes in σ_{in} and σ_{out} , the notations C_{in^k} and C_{out^k} work similarly to the WF-Net counterparts $[i^k]$ and $[o^k]$. Namely, the k represents the number of tokens within the source and sink in that state, but with spikes in the case of SNPSs. As these three notions are similar to weak soundness except for the input spikes, Figure 10 can be reused as an example, with the number of initial spikes reflecting k .

Definition 4.11 (k -Soundness of SN P Systems). Let Π be an SN P system, where $\Pi = (O, \sigma_1, \dots, \sigma_m, \text{syn}, \text{in}, \text{out})$, $\text{in} = 1, \text{out} = m$, and $\sigma_i = (\alpha_i, R_i)$, $1 \leq i \leq m$. Π is k -sound, where $\exists k \in \mathbb{N}$ iff the following requirements are satisfied:

- (1) Option to Complete: For every possible configuration C from C_{in^k} , there exists a sequence of valid configurations from C to C_{out^k} , i.e., $\forall_C (C_{\text{in}^k} \xrightarrow{*} C) \Rightarrow (C \xrightarrow{*} C_{\text{out}^k})$.
- (2) Proper termination: for every possible configuration C from C_{in^k} , C_{out^k} is the only reachable configuration, i.e., $\forall_C (C_{\text{in}^k} \xrightarrow{*} C \wedge C \geq C_{\text{out}^k}) \Rightarrow (C = C_{\text{out}^k})$.

Definition 4.12 (Up-to- k -Soundness of SN P Systems). Let Π be an SN P system, where $\Pi = (O, \sigma_1, \dots, \sigma_m, \text{syn}, \text{in}, \text{out})$, $\text{in} = 1, \text{out} = m$, and $\sigma_i = (\alpha_i, R_i)$, $1 \leq i \leq m$. Π is k -sound, where $\exists \ell, k \in \mathbb{N}$, $1 \leq \ell \leq k$ iff the following requirements are satisfied:

- (1) Option to Complete: For every possible configuration C from C_{in^ℓ} , there exists a sequence of valid configurations from C to C_{out^ℓ} , i.e., $\forall_C (C_{\text{in}^\ell} \xrightarrow{*} C) \Rightarrow (C \xrightarrow{*} C_{\text{out}^\ell})$.
- (2) Proper termination: for every possible configuration C from C_{in^ℓ} , C_{out^ℓ} is the only reachable configuration, i.e., $\forall_C (C_{\text{in}^\ell} \xrightarrow{*} C \wedge C \geq C_{\text{out}^\ell}) \Rightarrow (C = C_{\text{out}^\ell})$.




Figure 13: A WF-Net of processing complaints (based on [1]).

Definition 4.13 (Generalized Soundness of SN P Systems). Let Π be an SN P system, where $\Pi = (O, \sigma_1, \dots, \sigma_m, \text{syn}, \text{in}, \text{out})$, $\text{in} = 1, \text{out} = m$, and $\sigma_i = (\alpha_i, R_i)$, $1 \leq i \leq m$. Π is generalized sound, or k -sound where $\forall k \in \mathbb{N}$, iff the following requirements are satisfied:

- (1) Option to Complete: For every possible configuration C from C_{in^k} , there exists a sequence of valid configurations from C to C_{out^k} , i.e., $\forall_C (C_{\text{in}^k} \xrightarrow{*} C) \Rightarrow (C \xrightarrow{*} C_{\text{out}^k})$.
- (2) Proper termination: for every possible configuration C from C_{in^k} , C_{out^k} is the only reachable configuration, i.e., $\forall_C (C_{\text{in}^k} \xrightarrow{*} C \wedge C \geq C_{\text{out}^k}) \Rightarrow (C = C_{\text{out}^k})$.

5 DISCUSSION

Using the mapping literature [6], it is possible to verify whether a soundness notion is preserved from one workflow model to another. One example is the WF-Net that visualizes the processing of complaints in the paper that formalized classical soundness for the model, [1]. Shown in Figure 13, it was observed that the net was unsound, as classical soundness was the only formalized notion at the time. However, given other weakened forms of soundness that were developed over time, the model is considered to be relaxed sound. This system is considered live through the usage of all transitions in a process. However, it only follows the weakened form of the option to complete and proper termination for relaxed soundness. That is, the workflow has a possible firing sequence, namely the sequence *register*, (*processing_1*, *processing_2*), *processing_OK*, where o receives one token. With the mapping literature, the resulting SNPS is both shown in Figure 14 and 15, with the former using shorthand illustrations and the other not visualized through an SNPS simulator, WebSnapse [9]. Analyzing its behavior, it follows the same behavior as its PN counterpart, making it also relaxed sound in the SNPS context. In this instance, preservation of soundness between two models is possible.

In naming the neurons, the dual behavior of such a construct is considered, as neurons can act as both a container and a firing mechanism, which places and transitions do in the other model, respectively. This is seen through the mapping of the sequential route, where the neuron is both the place and transition [6]. Thus, the place name, followed by the following transition name, becomes the name format of its respective neuron. However, in the case of "unnatural" mappings in SN P systems, the transition name is applied to the last neuron, which is the original place name appended with an extension. This shows that the traversal of the structure is not yet complete.




Figure 14: A SN P system that represents the processing of complaints using shorthand notations from [6].




Figure 15: The expanded version of the SN P system representing complaint processing.

6 CONCLUSION

This paper successfully describes the soundness notions in the context of SN P systems through the eight formalized definitions, as well as their corresponding visual representations. Additionally, the relationship of soundness between other workflow properties has also been established, as well as the preservation of soundness between two models. With the results of this study, the applications of SN P systems can be extended to workflow visualization and analysis, in addition to calculation. This paves the way for a grander goal of handling the task of workflow analysis by exploiting the rich literature of these P systems with massive parallelism.

Recommendations for future work include providing methods to ensure the different soundness notions through structural analysis. Additionally, the correlation between liveness and boundedness with soundness only applies to the classical variation. With this, work on other soundness relationships to the formalized SN P system properties is to be done. Focusing on automation, the determination of which strongest soundness notion a SN P System satisfies can be extended to WebSnapse. This will enhance the tool, making it a property verifier in addition to its current visualization and simulation capabilities. These recommendations must be developed to strengthen the overall soundness foundation of SN P systems, creating a theoretically stronger model.

ACKNOWLEDGMENTS

R. I. M. Ramirez thanks the Department of Science and Technology - Science Education Institute through the Engineering Research and Development for Technology scholarship for supporting this

work at the University of the Philippines Diliman. F. G. C. Cabarle thanks the Scientific Productivity System of the University of the Philippines (2021-2023) for its support.

REFERENCES

- [1] W. M. P. Aalst, van der. 1996. *Structural characterizations of sound workflow nets*. Computing Science Reports, Vol. 9623. Technische Universiteit Eindhoven, Eindhoven.
- [2] W. M. P. Aalst, van der. 1997. Verification of workflow nets. In *Application and Theory of Petri Nets 1997*, P. Azéma and G. Balbo (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 407–426.
- [3] W. M. P. Aalst, van der. 1998. The application of Petri-nets to workflow management. *Journal of Circuits, Systems and Computers* 8, 1 (1998), 21–66. doi:10.1142/S0218126698000043
- [4] W. M. P. Aalst, van der, K. M. Hee, van, A. H. M. Hofstede, N. Sidorova, H. M. W. Verbeek, M. Voorhoeve, and M. T. Wynn. 2011. Soundness of workflow nets: classification, decidability, and analysis. *Formal Aspects of Computing* 23 (2011), 333–363.
- [5] H. N. Adorna. 2020. Properties of SN P system and its configuration graph.
- [6] F. G. C. Cabarle and H. N. Adorna. 2013. On Structures and Behaviors of Spiking Neural P Systems and Petri Nets. In *Membrane Computing*, E. Csuhaj-Varjú, M. Gheorghe, G. Rozenberg, A. Salomaa, and G. Vaszil (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 145–160.
- [7] J. Dehnert and W. M. P. Aalst. 2004. Bridging The Gap Between Business Models And Workflow Specifications. *Int. J. Cooperative Inf. Syst.* 13 (09 2004), 289–332. doi:10.1142/S0218843004000973
- [8] J. Dehnert and P. Rittgen. 2001. Relaxed Soundness of Business Processes. In *Advanced Information Systems Engineering*, K. R. Dittrich, A. Geppert, and M. C. Norrie (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 157–170.
- [9] M. Gulapa, J. S. Luzada, F. G. C. Cabarle, H. N. Adorna, K. Buño, and D. Ko. 2024. Correction to: WebSnapse Reloaded: The Next-Generation Spiking Neural P System Visual Simulator using Client-Server Architecture. In *Proceedings of the Workshop on Computation: Theory and Practice (WCTP 2023)*. Atlantis Press, C1–C1. doi:10.2991/978-94-6463-388-7_27
- [10] K. Hee, van, N. Sidorova, and M. Voorhoeve. 2003. Soundness and Separability of Workflow Nets in the Stepwise Refinement Approach. In *Applications and Theory of Petri Nets 2003*, W. M. P. van der Aalst and E. Best (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 337–356.
- [11] K. Hee, van, N. Sidorova, and M. Voorhoeve. 2004. Generalised Soundness of Workflow Nets Is Decidable. In *Applications and Theory of Petri Nets 2004*, J. Cortadella and W. Reisig (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 197–215.
- [12] D. Hollingsworth. 1995. Workflow Management Coalition The Workflow Reference Model.
- [13] M. Ionescu, G. Păun, and T. Yokomori. 2006. Spiking Neural P Systems. *Fundam. Inform.* 71 (06 2006), 279–308. doi:10.3233/FUN-2006-712-308
- [14] A. Martens. 2005. Analyzing Web Service Based Business Processes. In *Fundamental Approaches to Software Engineering*, M. Cerioli (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 19–33.
- [15] V.P. Metta, K. Krithivasan, and G. Garg. 2009. Modeling spiking neural P systems using timed Petri nets. In *NaBiC IEEE Conference*. 25–30.
- [16] V.P. Metta, K. Krithivasan, and G. Garg. 2009. Simulation of Spiking Neural P Systems Using Pnet Lab. In *Proceedings of the 12th CMC*. Paris, France. <http://cmc12.lacl.fr/cmc12proceedings.pdf>
- [17] V.P. Metta, K. Krithivasan, and G. Garg. 2009. Spiking Neural P systems and Petri nets. In *Proceedings of the International Workshop on Machine Intelligence Research*. <https://www.mirlabs.org/nagpur/paper02.pdf>
- [18] G. Păun. 2012. *Membrane Computing: An Introduction*. Springer Berlin Heidelberg, Berlin. <https://books.google.com.ph/books?id=37ASBwAAQBAJ>
- [19] G. Păun, G. Rozenberg, and A. Salomaa. 2009. *The Oxford Handbook of Membrane Computing*. Oxford University Press, Oxford. doi:10.1093/oso/9780199556670.001.0001
- [20] F. Puhlmann and M. Weske. 2006. Interaction Soundness for Service Orchestrations. In *Service-Oriented Computing – ICSOC 2006*, A. Dan and W. Lamersdorf (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 302–313.
- [21] F. Puhlmann and M. Weske. 2006. Investigations on Soundness Regarding Lazy Activities. In *Business Process Management*, S. Dustdar, J. L. Fiadeiro, and A. P. Sheth (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 145–160.
- [22] R.A. Toorn, van der. 2004. *Component-based software design with Petri nets : an approach based on inheritance of behavior*. Phd Thesis 1 (Research TU/e / Graduation TU/e). Mathematics and Computer Science. doi:10.6100/IR571325
- [23] G. Zhang, S. Verlan, T. Wu, F. G. C. Cabarle, J. Xue, D. Orellana-Martin, J. Dong, L. Valencia-Cabrera, and M. J. Pérez-Jiménez. 2024. *Spiking Neural P Systems: Theory, Applications and Implementations*. Springer Nature Singapore, Singapore. <https://books.google.com.ph/books?id=DOU3EQAAQBAJ>

Galaw at Gunita: Extended Reality Murals for Experiencing Filipino Art

Jomar Delos Reyes, Sealtiel Dy, Rica Mae Sales, Orrin Landon Uy, Toni-Jan Keith Monserrat,
Ryan Austin Fernandez, Jordan Aiko Deja

De La Salle University, Manila, Philippines

jomar_delosreyes@dlsu.edu.ph, sealtiel_dy@dlsu.edu.ph, rica_sales@dlsu.edu.ph,
orrin_landon_uy@dlsu.edu.ph, tjmonsi@gmail.com, ryan.fernandez@dlsu.edu.ph, jordan.deja@dlsu.edu.ph*

Initial animation



Initial screen



Poke gesture on the start button



Fading out animation

Mural XR overlay



Full mural




From top to bottom: reset, tutorial, map layout



Map layout


Tutorial modal



Activate animation by gazing at artwork



Poke icon after activating it to open model containing artwork info



Read details, title, history of the artwork. Move modal as you need.

Figure 1: Overview of the *SoulWall* experience. Through gaze and touch-based interaction (*galaw*), users activate artwork animations and information modals, supporting exploration and appreciation (*gunita*) of the artworks. All featured artworks are used with the artists' consent and remain their copyright.

ABSTRACT

Digital and interactive spaces are becoming increasingly prevalent as platforms for cultural engagement, offering new ways to make art more accessible, engaging, and inclusive. In the Philippine context, where visual art is deeply rooted in precolonial, colonial, and postcolonial histories, there is a growing need to explore how digital systems can support art appreciation without replacing or

compromising traditional and physical artworks. Rather than treating digital experiences as substitutes, we argue for the value of creating interactive digital twins that allow audiences to explore, touch, and engage with artworks while preserving the integrity of the originals. In this paper, we present *SoulWall*, an extended reality (XR) interactive mural system designed to augment Filipino

artworks through embodied interaction. SoulWall enables viewers to experience paintings and animations at scale, supporting exploratory and playful engagement while respecting artist intent. We describe the design and implementation of the system, including its mural layout, interaction techniques, and interaction logging infrastructure. We report findings from a user study focused on user experience complemented by analyses of interaction logs and spatial engagement patterns. Our results highlight the potential of XR murals as a cultural computing artifact for art appreciation and for showcasing Filipino artists in interactive public and exhibition spaces.

CCS CONCEPTS

- **Human-centered Computing** → **Mixed / Augmented Reality**;
- **Applied Computing** → **Interactive Learning Environments**;

KEYWORDS

extended reality, murals, art, culture, appreciation

1 INTRODUCTION

Digital and interactive systems are becoming increasingly common in cultural and public spaces, shaping how people encounter and engage with information. While these environments can be engaging, highly stimulating settings may also present challenges when content is dense, visually rich, or unstructured [1–3, 49]. When multiple visual and sensory elements compete for attention, users may experience difficulty focusing on specific details or forming meaningful connections [41]. Such conditions are common in public exhibition spaces, including museums and galleries, where visitors encounter a combination of artworks, text, spatial layouts, and ambient stimuli [42, 45].

Museums and galleries are often designed to be rich and immersive, combining visual artifacts with explanatory text, digital media, and spatial narratives. While these elements aim to support interpretation and appreciation, they can also contribute to cognitive fatigue when presented simultaneously [18, 33]. Prior work has shown that visitors may struggle to sustain attention or meaningfully engage with content in such environments, particularly when exhibits are dense or unfamiliar [35, 44, 46]. These challenges raise questions about how interactive systems can support art appreciation without overwhelming visitors or detracting from the artworks themselves.

In parallel, digital technologies have opened new opportunities for presenting and experiencing art beyond traditional formats. Techniques such as visual layering, spatial organization, and interactive exploration have been used to support engagement and understanding across educational and cultural contexts [27, 29, 36]. Interactive and immersive systems, including extended reality (XR), allow content to be arranged spatially and explored through embodied interaction, offering alternative ways of engaging with visual material [5, 7]. Large interactive wall displays, in particular, provide shared visual reference points that encourage exploration while maintaining a sense of scale and context [14, 16].

In the Philippine context, visual art is closely tied to precolonial, colonial, and postcolonial histories, with many works carrying

cultural, political, and social significance. At the same time, access to artworks is often limited by physical preservation concerns, exhibition constraints, or geographic availability. Rather than replacing physical and traditional art forms, digital systems can serve as complementary platforms that create interactive digital twins of artworks, allowing audiences to explore and engage with them without risking damage to the originals. Such approaches offer opportunities to support Filipino artists and broaden public access to their work while respecting artistic intent and material integrity.

In this paper, we present *SoulWall*, an extended reality interactive mural system designed to augment Filipino artworks through digital animation and embodied interaction. SoulWall, through the verbs *Galaw* at *Gunita* (to touch and reminisce), allows viewers to explore paintings and animated elements, supporting playful and exploratory engagement while preserving the physical artworks. We describe the design and implementation of the system, including its mural layout, interaction techniques, and interaction logging infrastructure. We also report findings from an initial evaluation focused on user experience, and supported by analyses of interaction logs and spatial engagement patterns. Together, these contributions position interactive XR murals as a viable approach for supporting art appreciation and showcasing Filipino artists in digital and public exhibition spaces.

2 RELATED WORK

Prior work has explored how extended reality (XR) systems can support user engagement and orientation through immersive visualization, interactivity, and spatial organization [6, 50]. AR and VR systems, particularly those using head-mounted displays, have been shown to support active exploration and sustained attention when compared to traditional presentation formats [4, 9, 10, 21, 51]. Gesture-based and embodied interactions further contribute to spatial awareness and intuitive navigation by allowing users to interact directly with digital content [24, 48]. These approaches highlight the value of spatially situated and interactive representations for engaging users in information-rich environments.

Spatial structure and contextual anchoring also play an important role in how users navigate and make sense of digital content. Prior studies show that VR environments aligned with real-world layouts can support orientation more effectively than abstract representations [30]. Similarly, mimetic icons, annotations, and spatial cues help users relate digital elements to physical or familiar reference points [8, 12, 19, 20]. Techniques inspired by the method of loci emphasize the use of landmarks and spatial segmentation to organize content and guide exploration [25, 38]. While often discussed in the context of memory, these techniques also inform the design of navigable and interpretable interactive spaces. **Our work builds on these XR interaction and spatial organization principles by presenting artworks as part of an interactive mural.** Rather than using spatial structure to optimize recall, our system applies these techniques to support exploration, orientation, and engagement with digital representations of Filipino artworks. By anchoring animations and interactions directly to visual elements within the mural, the system encourages viewers to move, observe, and interact with the artworks in a manner that remains grounded in their visual and spatial context.

Museums and galleries present unique challenges and opportunities for interactive systems, as environmental factors such as layout, visibility, and distractions influence how visitors engage with exhibits [22, 23]. Prior work suggests that prolonged engagement and active interaction with exhibits can support deeper appreciation and understanding [40]. XR systems have been introduced in museum contexts to increase visitor engagement, usability, and perceived value, with immersion, enjoyment, and curiosity playing important roles in shaping visitor experience [34]. Virtual environments that present related objects or narratives together have also been shown to encourage exploratory behavior and informal learning [31]. Large-scale deployments, such as interactive historical exhibits, further demonstrate the potential of digital systems to broaden access to cultural content [28]. **In contrast to systems that prioritize informational delivery or learning outcomes, our work positions XR as a complementary medium for art appreciation especially towards deeper appreciation for local Filipino art.** By creating an interactive digital twin of physical artworks, the system allows visitors to engage with Filipino art through movement, touch, and visual exploration without altering or risking damage to the originals. This work extends prior research on XR in museums by focusing on user experience and interaction patterns in an interactive mural setting, emphasizing accessibility, engagement, and respect for traditional art forms.

3 FEATURES OF SOULWALL

SoulWall integrates a set of interactive features designed to support the exploration and appreciation of Filipino artworks through an extended reality (XR) mural. Rather than replacing physical artworks, the system augments them by creating an interactive digital twin that allows viewers to engage with visual elements at scale while preserving the integrity of the originals. The following features focus on spatial organization, visual cues, and lightweight interactions that encourage walk-up use and exploratory engagement.

3.1 Mural Design

The mural uses the Philippine eagle as an abstract and symbolic visual anchor that aligns with Filipino cultural identity. The eagle serves as a unifying visual structure that organizes the artworks within a single composition, helping viewers orient themselves and explore different sections of the mural in a coherent manner. Spatially structured XR environments have been shown to support navigation and engagement by providing clear visual reference points [8, 25, 32].

The eagle is divided into four segments—A) left wing, B) center body, C) center-right wing, and D) right wing—physically printed on a large print (Figure 3). This segmentation allows artworks to be grouped into visually distinct regions, enabling viewers to focus on specific parts of the mural without being overwhelmed by the entire display at once. Such spatial grouping supports gradual exploration and encourages viewers to move closer to different sections of the mural.

Each eagle segment contains a curated set of artworks from a featured Filipino artist. The artworks are arranged to maintain visual cohesion and thematic consistency within each segment, allowing viewers to appreciate stylistic elements and recurring motifs. By

embedding the artworks within a single mural composition, SoulWall provides a shared visual context that supports comparison, reflection, and discussion while respecting the original works [29].

3.2 Interactions

SoulWall employs simple, trigger-based interactions that bring digital artworks to life through animation and visual feedback. These interactions are designed to enhance engagement and curiosity while keeping interaction demands low, allowing viewers to focus on the artworks themselves rather than on learning complex controls [47, 50]. Subtle animations are used to indicate interactivity and to draw attention to specific visual elements without altering the core composition of the artworks [6, 15].

Prior work suggests that users prefer familiar and easily interpretable gestures when encountering XR systems for the first time [11]. Guided by this, SoulWall adopts interaction techniques that align with common real-world metaphors and interface conventions. The supported interactions include:

- (1) **Gaze-Based Highlighting:** Looking at a specific segment of the eagle highlights it and triggers subtle animations on embedded figures, indicating that the artworks are interactive.
- (2) **Modal Activation:** Selecting a highlighted figure opens a modal containing contextual information about the selected artwork and artist.
- (3) **Modal Dismissal:** The modal can be closed using a standard close button.
- (4) **Modal Repositioning:** Users may drag the modal within a bounded area to adjust its placement for comfortable viewing.
- (5) **Pinch-to-Zoom:** Pinch gestures allow users to zoom in on details or zoom out for an overview of the mural or modal content.
- (6) **Reset Function:** A reset icon restores the mural to its default state and closes all open modals.
- (7) **Help Access:** A help icon replays a short tutorial that explains the available interactions.
- (8) **Scene Switching:** A museum icon transitions the view from the mural to an overview of the exhibition layout.

Together, these interactions encourage viewers to move, observe, and engage with the artworks at their own pace. By combining spatial organization with lightweight interaction, SoulWall supports exploratory art appreciation while minimizing disruption to the visual integrity of the mural. An overview of the interaction flow is shown in Figure 1, and examples of mural segments and artwork modals are shown in Figure 2.

4 EVALUATION

4.1 Participants and Protocol

Sixteen participants (aged 18–25) were recruited through convenience sampling. All participants provided informed consent prior to participation. The study focused on evaluating user experience and interaction behavior during an exploratory session with the SoulWall prototype.




Figure 2: Mural segment features of SoulWall, showing interactive artwork cues, highlighted and non-highlighted states across segments, and available modal interactions.



Figure 3: Eagle mural layout used as the visual structure for SoulWall.

Participants were first given a brief introduction to the system and its available interactions. They were then asked to freely explore the interactive mural for approximately 15 minutes. No specific tasks or performance goals were imposed in order to encourage natural and self-directed interaction with the artworks. After the exploration phase, participants completed the short version of the User Experience Questionnaire (UEQ-S). During the session, interaction logs were collected automatically by the system, capturing movement, interaction events, and dwell patterns (similar to that of approach of [20]). Basic background information related to prior XR experience and handedness was also recorded for descriptive purposes.

4.2 Experimental Setup

The study was conducted in an indoor exhibition-like space where the mural was physically installed as a large printed artwork. Participants interacted with the augmented mural using a standalone




Figure 4: UEQ-S results for SoulWall, showing mean scores for pragmatic quality, hedonic quality, and overall user experience. Colored bands indicate benchmark categories from bad to excellent, with the black line representing mean values. Graph was scaled to half values since we used the short version of the UEQ. Score trasmutations were made with the official analysis tool by [39]).

XR headset. The setup allowed participants to stand in front of the mural and move freely while exploring the interactive elements.

A facilitator was present to assist with basic setup and to address technical issues if needed, but did not guide participant behavior during exploration. The participant's view was mirrored to a laptop for observation and to ensure correct system operation.

4.3 Findings

User experience (as seen in Figure 4) was assessed using the UEQ-S, which measures pragmatic and hedonic qualities. For SoulWall, pragmatic quality received a mean score of $M = 4.66$ ($SD = 0.61$), indicating that participants generally found the system understandable and manageable during short-term use. Hedonic quality was rated higher, with a mean score of $M = 6.02$ ($SD = 0.71$), suggesting




Figure 5: Heatmap of interaction duration (in milliseconds) across participants and artworks. Rows represent participants and columns represent individual artworks. Higher intensity indicates longer engagement time with an artwork.

that the interactive mural was perceived as engaging and enjoyable. The overall UEQ-S score for SoulWall was $M = 5.34$ ($SD = 0.54$), reflecting a generally positive user experience that supported exploratory interaction without requiring extensive instruction.

Interaction logs provided additional insights into how participants engaged with the mural (see Figure 5). Logged data included movement trajectories, time spent near different mural segments, and interaction events such as selections and modal activations. Across artworks, movement activity was highest for artwork C_5 ($M = 15.63$ movements, $SD = 35.14$), followed by B_1 ($M = 15.31$, $SD = 47.82$) and A_3 ($M = 15.19$, $SD = 39.36$). Visualizations of these logs showed that participants explored multiple sections of the mural rather than remaining in a single area, with noticeable clustering around visually dense or animated regions. This suggests that animations and visual cues effectively drew attention and encouraged movement across the mural.

Interaction duration revealed a different pattern of engagement. Artwork A_5 received the longest interaction times on average ($M = 229.57$ ms, $SD = 454.74$ ms), followed by C_1 ($M = 203.90$ ms, $SD = 401.54$ ms) and B_1 ($M = 195.50$ ms, $SD = 387.81$ ms). At the participant level, interaction patterns varied substantially. Participant P16 exhibited the highest overall movement activity across artworks, while participant P15 accumulated the greatest total interaction duration. These differences reflect distinct exploration styles, with some participants engaging through frequent movement across artworks and others spending longer periods inspecting individual pieces and associated information modals. Together, these findings indicate that SoulWall supported both focused and broad exploration strategies, allowing users to engage with the artworks at their own pace.

Taken together, the UEQ-S results and interaction logs indicate that SoulWall supported a positive and flexible user experience for engaging with Filipino artworks through an interactive XR mural. Hedonic quality scores reflected high enjoyment, while interaction

logs showed varied engagement patterns across artworks and participants, with some artworks receiving higher movement activity and others longer interaction durations. Rather than optimizing for task performance, SoulWall facilitated open-ended exploration and visual engagement, allowing participants to interact with artworks in different ways. These findings align with the system's goal of complementing traditional art displays by enabling interactive appreciation without altering or replacing the original works.

5 DISCUSSION AND DESIGN IMPLICATIONS

We synthesize the findings from the evaluation with existing work on XR interaction techniques, spatial organization, and museum-based interactive systems, and discuss how these insights can inform the design of XR murals that support exploratory and respectful engagement with artworks [16, 34, 50].

5.1 Designing for visual salience in interactive murals

Interaction logs indicate that visual salience played a central role in shaping how participants engaged with the SoulWall mural. Artworks that were visually larger, more centrally positioned, or featured animated elements tended to attract higher interaction activity, either through longer engagement durations or more frequent movement events. In contrast, smaller or visually dense artworks were interacted with less consistently across participants. These patterns align with prior findings showing that visual prominence strongly guides attention and exploration in complex visual environments [13, 37].

Importantly, the interaction logs revealed that high movement activity and long interaction duration did not always occur on the same artworks. Some artworks elicited frequent movement across participants, while others encouraged longer periods of focused engagement. This distinction suggests that visual salience can support different modes of engagement prompting either exploratory movement or sustained inspection rather than a single dominant interaction pattern. These observations are consistent with XR and display-based studies showing that perceptual features often influence how users engage, independent of the interaction techniques available [16, 50].

For XR mural design, these findings highlight the need to balance artistic composition with interaction visibility. Subtle animations and visual cues can draw attention and invite exploration without overwhelming the artwork, while overly subtle or crowded elements may limit discoverability. Designing for visual salience can therefore support varied forms of engagement while preserving the integrity of the artworks.

5.2 Supporting spatial exploration through mural structure

Prior work on XR navigation and interactive wall displays emphasizes the importance of spatial structure in supporting orientation and exploration [16, 30]. In SoulWall, the use of a single cohesive mural provided a shared visual reference that encouraged participants to traverse multiple mural regions rather than remain focused

on a single area. Interaction logs showed movement across different segments, with localized clustering around particular artworks before participants transitioned to adjacent areas.

At the same time, engagement patterns suggest that participants relied more heavily on visually salient cues than on the intended thematic segmentation of the mural. While the mural structure supported movement across space, boundaries between segments were not always explicitly reflected in interaction behavior. Similar challenges have been reported in XR environments where spatial metaphors are present but not sufficiently reinforced through visual or interaction cues [12, 43].

These observations suggest that interactive murals benefit from spatial structures that are both meaningful and perceptually legible. Reinforcing segmentation through visual framing, spacing, or background variation may help users better understand how different regions relate to one another, supporting smoother and more intentional spatial exploration.

5.3 Balancing engagement and interaction simplicity

Consistent with prior work on immersive museum experiences, participants reported high hedonic quality for SoulWall, indicating that animation, scale, and embodied interaction contributed to an engaging experience [34, 40]. UEQ-S results suggest that the system successfully supported enjoyable and playful interaction, which is important for public-facing cultural installations.

Interaction logs further revealed substantial variation in how participants engaged with the mural. Some participants accumulated high movement counts across many artworks, while others spent longer durations engaging with fewer pieces. This variation reflects different exploration strategies, ranging from broad scanning of the mural to more focused inspection of individual artworks and information modals. Such diversity aligns with prior findings that users adopt distinct engagement styles in open-ended XR environments [5].

These findings reinforce the importance of keeping interactions lightweight and optional. Prior work has cautioned that overly complex interactions can distract from content in immersive systems [26]. In SoulWall, familiar gestures and minimal interaction steps appeared to support engagement without enforcing a single interaction style, allowing participants to navigate the experience according to their preferences.

5.4 Familiarity and onboarding in walk-up XR experiences

Familiar interaction metaphors have been shown to support usability in XR systems [11, 17]. In SoulWall, participants were generally able to understand gaze-based highlighting, selection, and modal interactions with limited instruction, suggesting that familiarity helped lower the initial barrier to engagement.

However, interaction logs and observations indicate that some features were underutilized when participants skipped or quickly dismissed the tutorial. This reflects a common challenge in walk-up-and-use XR installations, where users are eager to explore but may avoid explicit onboarding [16]. Integrating guidance directly into the interaction flow through progressive cues or contextual

highlights, may better support first-time users without interrupting exploration.

Overall, these findings echo prior work on interactive museum systems, which emphasizes keeping technology unobtrusive and allowing visitors to engage at their own pace [31, 34]. For XR murals, familiarity in interaction design and lightweight onboarding help ensure that attention remains on the artworks rather than on the interface.

Taken together, the discussion reinforces that effective XR murals for art appreciation should prioritize visual clarity, spatial legibility, and interaction simplicity. By grounding design decisions in both prior research and observed interaction patterns, such as differences between movement-driven and duration-driven engagement, systems like SoulWall can complement traditional art displays while supporting diverse and self-directed forms of interaction with cultural content.

6 FUTURE WORK AND CONCLUSION

This paper presented SoulWall, an XR interactive mural designed to support the exploration and appreciation of Filipino artworks through digital augmentation. Rather than replacing physical artworks, SoulWall creates an interactive digital twin that allows viewers to engage (*galaw*) with visual elements while preserving the integrity of the originals and be able to appreciate (*gunita*) the art they are interacting with. We described the system design and reported findings from an evaluation focused on user experience and interaction behavior.

Results from the UEQ-S and interaction logs suggest that SoulWall provides an engaging and flexible experience, supporting different exploration styles through lightweight interactions and spatial organization. Participants were drawn to visually salient and animated elements, and interaction patterns indicated movement across multiple mural regions rather than fixed or linear use. These findings highlight the importance of visual clarity, spatial legibility, and simple interaction design in XR-based art installations.

This work has several limitations. The study focused on short-term use with a relatively homogeneous participant group and was conducted in a controlled setting rather than a live exhibition environment. As a result, the findings primarily reflect initial engagement and exploratory behavior.

Future work will examine SoulWall in longer-term and in-situ deployments, such as museums or public cultural spaces, and with more diverse audiences. We also plan to explore how interactive XR murals can support deeper forms of reflection and understanding, including spatial memory and recall, which are beyond the scope of the present paper. Overall, this work contributes an initial step toward designing XR murals that complement traditional art displays and expand access to Filipino art through interactive digital experiences.

REFERENCES

- [1] P. Barrouillet, A. De Paepe, and N. Langerock. 2012. Time causes forgetting from working memory. *Psychonomic Bulletin & Review* 19, 1 (2012), 87–92.
- [2] Pierre Barrouillet, Sophie Portrat, and Valérie Camos. 2011. On the law relating processing to storage in working memory. *Psychological Review* 118, 2 (2011), 175–192. doi:10.1037/a0022324
- [3] Pierre Barrouillet, Sophie Portrat, Evie Vergauwe, Kevin Diependaele, and Valérie Camos. 2011. Further Evidence for Temporal Decay in Working Memory: Reply to

- Lewandowsky and Oberauer (2009). *Journal of experimental psychology. Learning, memory, and cognition* 37 (Sept. 2011), 1302–17. doi:10.1037/a0022933
- [4] Nicko R Caluya, Alexander Plopski, Jayzon F Ty, Christian Sandor, Takafumi Taketomi, and Hirokazu Kato. 2018. Transferability of spatial maps: Augmented versus virtual reality training. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 387–393.
 - [5] Yeonhee Cho. 2018. *How spatial presence in VR affects memory retention and motivation on second language learning: a comparison of desktop and immersive VR-based learning*. Master's thesis. Syracuse University.
 - [6] Joey Ka-Yee Essoe, Nicco Reggente, Ai Aileen Ohno, Younji Hera Baek, John Dell'Italia, and Jesse Rissman. 2022. Enhancing learning and retention with distinctive virtual reality environments and mental context reinstatement. *npi Science of Learning* 7, 1 (Dec. 2022), 1–14. doi:10.1038/s41539-022-00147-6 Publisher: Nature Publishing Group.
 - [7] Babajide Familoni and Nneamaka Onyebuchi. 2024. AUGMENTED AND VIRTUAL REALITY IN U.S. EDUCATION: A REVIEW: ANALYZING THE IMPACT, EFFECTIVENESS, AND FUTURE PROSPECTS OF AR/VR TOOLS IN ENHANCING LEARNING EXPERIENCES. *International Journal of Applied Research in Social Sciences* 6 (04 2024), 642–663. doi:10.51594/ijarss.v6i4.1043
 - [8] Yuichiro Fujimoto, Goshiro Yamamoto, Hirokazu Kato, and Jun Miyazaki. 2012. Relation between location of information displayed by augmented reality and user's memorization. In *Proceedings of the 3rd Augmented Human International Conference (AH '12)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/2160125.2160132
 - [9] Shubham Gargish, Deepti P. Kaur, Archana Mantri, Gurjinder Singh, and Bhanu Sharma. 2021. Measuring effectiveness of augmented reality-based geometry learning assistant on memory retention abilities of the students in 3D geometry. *Computer Applications in Engineering Education* 29, 6 (2021), 1811–1824. doi:10.1002/cae.22424 _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/cae.22424>
 - [10] Shubham Gargish, Archana Mantri, and Deepti Prit Kaur. 2022. Evaluation of memory retention among students using augmented reality based geometry learning assistant. *Education and Information Technologies* 27, 9 (Nov. 2022), 12891–12912. doi:10.1007/s10639-022-11147-9
 - [11] Despoina Gavgiotaki, Stavroula Ntoa, George Margetis, Konstantinos C Apostolakis, and Constantine Stephanidis. 2023. Gesture-based interaction for AR systems: a short review. In *Proceedings of the 16th International Conference on Pervasive Technologies Related to Assistive Environments*. 284–292.
 - [12] Marlynn M. Griffin and Daniel H. Robinson. 2005. Does spatial or visual information in maps facilitate text recall? Reconsidering the conjoint retention hypothesis. *Educational Technology Research and Development* 53, 1 (March 2005), 23–36. doi:10.1007/BF02504855
 - [13] John M Henderson and Andrew Hollingworth. 1999. High-level scene perception. *Annual review of psychology* 50, 1 (1999), 243–271.
 - [14] IMG-LAB INC. 2020. NetMarble Game Academy Exhibition. <https://www.img-lab.co.kr/10>
 - [15] Sanjay Jangid. 2023. The Contribution of Animation in XR - Chitkara University. <https://www.chitkara.edu.in/blogs/the-contribution-of-animation-in-xr/>
 - [16] Yvonne Jansen, Jonas Schjerlund, and Kasper Hornbæk. 2019. Effects of Locomotion and Visual Overview on Spatial Memory when Interacting with Wall Displays. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3290605.3300521
 - [17] Hyeonseok Kim, Natsue Yoshimura, and Yasuharu Koike. 2020. Investigation of Delayed Response during Real-Time Cursor Control Using Electroencephalography. *Journal of Healthcare Engineering* 2020, 1 (2020), 1418437. doi:10.1155/2020/1418437 _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1155/2020/1418437>
 - [18] Neil Kotler and Philip Kotler. 2007. Can Museums be All Things to All People?: Missions, goals, and marketing's role. In *Museum Management and Marketing*. Routledge. Num Pages: 18.
 - [19] N. Kovacevic, J. A. Deja, M. Weerasinghe, K. Copic Pucihar, and M. Kljun. 2023. Retzzles: Towards Supporting Retention Using Puzzle Interactions and Abstract Symbols. *Scientific Presentations of Primitijada 2023* (April 28–May 03 2023), 1–5.
 - [20] Nikola Kovacevic, Maheshya Weerasinghe, and Jordan Aiko Deja. 2022. Retzzles: Engaging Users towards Retention through Touchscreen Puzzles. In *Companion Proceedings of the 2022 Conference on Interactive Surfaces and Spaces*. 14–17.
 - [21] Eric Krokos, Catherine Plaisant, and Amitabh Varshney. 2019. Virtual memory palaces: immersion aids recall. *Virtual Reality* 23, 1 (March 2019), 1–15. doi:10.1007/s10055-018-0346-3
 - [22] Jakub Krukar. 2014. Walk, look, remember: The influence of the gallery's spatial layout on human memory for an art exhibition. *Behavioral sciences* 4, 3 (2014), 181–201.
 - [23] Jakub Krukar. 2015. The influence of an art gallery's spatial layout on human attention to and memory of art exhibits. *University of Northumbria, Newcastle* (2015).
 - [24] Yea Som Lee and Bong-Soo Sohn. 2018. Immersive Gesture Interfaces for Navigation of 3D Maps in HMD-Based Mobile Virtual Environments. *Mobile Information Systems* 2018, 1 (2018), 2585797. doi:10.1155/2018/2585797 _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1155/2018/2585797>
 - [25] Jiazhou Liu, Kadek Satriadi, Barrett Ens, and Tim Dwyer. 2024. *Investigating the Effects of Physical Landmarks on Spatial Memory for Information Visualisation in Augmented Reality*. doi:10.13140/RG.2.2.29017.94563
 - [26] Guido Makransky, Thomas S. Terkildsen, and Richard E. Mayer. 2019. Adding immersive virtual reality to a science lab simulation causes more presence but less learning. *Learning and Instruction* 60 (2019), 225–236. doi:10.1016/j.learninstruc.2017.12.007
 - [27] Jennifer A. McCabe. 2015. Location, Location, Location! Demonstrating the Mnemonic Benefit of the Method of Loci. *Teaching of Psychology* 42, 2 (April 2015), 169–173. doi:10.1177/0098628315573143 Publisher: SAGE Publications Inc.
 - [28] Colleen McGaughey and John Russick. 2021. Evaluating the Impact of Augmented and Virtual Reality. <https://www.aam-us.org/2021/01/15/evaluating-the-impact-of-augmented-and-virtual-reality/>
 - [29] Gary R. McKenzie and Elaine Danielson. 2003. Improving Comprehension through Mural Lessons. *The Reading Teacher* 56, 8 (2003), 738–742. <https://www.jstor.org/stable/20205287> Publisher: [Wiley, International Reading Association].
 - [30] Michael S. Miller, Deborah M. Clawson, and Marc M. Sebrechts. 1999. Long-Term Retention of Spatial Knowledge Acquired in Virtual Reality. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 43, 22 (Sept. 1999), 1243–1246. doi:10.1177/154193129904302216 Publisher: SAGE Publications Inc.
 - [31] Tomas Gislason Moesgaard, Mass Witt, Jonathan Fiss, Cecilie Warming, Jonathan Klubien, and Henrik Schønau-Fog. 2015. Implicit and Explicit Information Mediation in a Virtual Reality Museum Installation and its Effects on Retention and Learning Outcomes: The 9th European Conference on Games Based Learning ECGBL 2015. *Proceedings of The 9th European Conference on Games-Based Learning* (Oct. 2015), 387–394. Place: Reading, United Kingdom Publisher: Academic Conferences and Publishing International.
 - [32] Brigham Moll and Ed Sykes. 2022. Optimized virtual reality-based Method of Loci memorization techniques through increased immersion and effective memory palace designs: a feasibility study. *Virtual Reality* 27, 2 (Oct. 2022), 941. doi:10.1007/s10055-022-00700-z
 - [33] Roxana Moreno and Richard E. Mayer. 1999. Visual Presentations in Multimedia Learning: Conditions that Overload Visual Working Memory. In *Visual Information and Information Systems*, Dionysius P. Huijsmans and Arnold W. M. Smeulders (Eds.). Springer, Berlin, Heidelberg, 798–805. doi:10.1007/3-540-48762-X_98
 - [34] Jessie Pallud. 2017. Impact of interactive technologies on stimulating learning experiences in a museum. *Information & Management* 54, 4 (June 2017), 465–478. doi:10.1016/j.im.2016.10.004
 - [35] Elizabeth A Phelps. 2004. Human emotion and memory: interactions of the amygdala and hippocampal complex. *Current Opinion in Neurobiology* 14, 2 (2004), 198–202. doi:10.1016/j.conb.2004.03.015
 - [36] Ayisha Qureshi, Farwa Rizvi, Anjum Syed, Aqueel Shahid, and Hana Manzoor. 2014. The method of loci as a mnemonic device to facilitate learning in endocrinology leads to improvement in student performance as measured by assessments. *Advances in Physiology Education* 38, 2 (June 2014), 140. doi:10.1155/advan.00092.2013
 - [37] Keith Rayner and Alexander Pollatsek. 1992. Eye movements and scene perception. *Canadian Journal of Psychology* 46, 3 (1992), 342–376. doi:10.1037/h0084328
 - [38] Oscar Rosello, Marc Exposito, and Pattie Maes. 2016. NeverMind: Using Augmented Reality for Memorization. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. ACM, Tokyo Japan, 215–216. doi:10.1145/2984751.2984776
 - [39] Martin Schrepp, Andreas Hinderks, and Jörg Thomaschewski. 2017. Design and evaluation of a short version of the user experience questionnaire (UEQ-S). *International Journal of Interactive Multimedia and Artificial Intelligence*, 4 (6), 103–108. (2017).
 - [40] Rebecca Sweetman, Alison Hadfield, and Akira O'Connor. 2020. Material culture, museums, and memory: Experiments in visitor recall and memory. *Visitor Studies* 23, 1 (2020), 18–45.
 - [41] John Sweller, Paul Ayres, and Slava Kalyuga. 2011. *Cognitive Load Theory*. Springer. <https://doi.org/10.1007/978-1-4419-8126-4>
 - [42] John N. Towse, Nelson Cowan, Graham J. Hitch, and Neil J Horton. 2008. The recall of information from working memory: insights from behavioural and chronometric perspectives. *Experimental psychology* 55, 6 (2008), 371–383. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2658622/>
 - [43] Barbara Tversky. 2003. Structures Of Mental Spaces How People Think About Space. *Environment and Behavior - ENVIRON BEHAV* 35 (01 2003), 66–80. doi:10.1177/0013916502238865
 - [44] Eunjoon Um, Jan Plass, Elizabeth Hayward, and Bruce Homer. 2011. Emotional Design in Multimedia Learning. *Journal of Educational Psychology* 104 (12 2011), 485–498. doi:10.1037/a0026609
 - [45] Trish L. Varao-Sousa, Daniel Smilek, and Alan Kingstone. 2018. In the lab and in the wild: How distraction and mind wandering affect attention and memory. *Cognitive Research: Principles and Implications* 3, 1 (Nov. 2018), 42. doi:10.1186/s41235-018-0137-0

- [46] Patrik Vuilleumier. 2005. How brains beware: neural mechanisms of emotional attention. *Trends in Cognitive Sciences* 9, 12 (2005), 585–594. doi:10.1016/j.tics.2005.10.011
- [47] Heping Xie, Fuxing Wang, Yanbin Hao, Jiaxue Chen, Jing An, Yuxin Wang, and Huashan Liu. 2017. The more total cognitive load is reduced by cues, the better retention and transfer of multimedia learning: A meta-analysis and two meta-regression analyses. *PLOS ONE* 12, 8 (Aug. 2017), e0183884. doi:10.1371/journal.pone.0183884 Publisher: Public Library of Science.
- [48] Johannes Zagermann, Ulrike Pfeil, Daniel Fink, Philipp von Bauer, and Harald Reiterer. 2017. Memory in Motion: The Influence of Gesture- and Touch-Based Input Modalities on Spatial Memory. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 1899–1910. doi:10.1145/3025453.3026001
- [49] Eric Zeglen and Joseph Rosendale. 2018. Increasing Online Information Retention: Analyzing the Effects. *Journal of Open, Flexible, and Distance Learning* 22, 1 (Aug. 2018), 22–33. <https://www.learntechlib.org/p/184660/> Publisher: Distance Education Association of New Zealand.
- [50] Lei Zhang, Doug A. Bowman, and Caroline N. Jones. 2019. Exploring Effects of Interactivity on Learning with Interactive Storytelling in Immersive Virtual Reality. In *2019 11th International Conference on Virtual Worlds and Games for Serious Applications (VS-Games)*. 1–8. doi:10.1109/VS-Games.2019.8864531 ISSN: 2474-0489.
- [51] Yuchong Zhang, Adam Nowak, Andrzej Romanowski, and Morten Fjeld. 2023. Virtuality or Physicality? Supporting Memorization Through Augmented Reality Gamification. In *Companion Proceedings of the 2023 ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS '23 Companion)*. Association for Computing Machinery, New York, NY, USA, 53–58. doi:10.1145/3596454.3597183

Predictive Models for Solar Energy Output: Performance Evaluation of K-Means and GMM Clustering with Model Interpretability

Carlos Miguel Jakosalem*, Francis Rale Abrenica**
University of San Carlos, Cebu City, Philippines
kikochan21@gmail.com*, rale.abrenica@gmail.com**

ABSTRACT

Solar energy, a sustainable energy source, poses significant challenges for energy reliability because of its strong dependence on weather variability. This unpredictability complicates grid planning, energy distribution, and storage operations, especially in regions where there are frequent atmospheric fluctuations. Existing prediction models typically treat weather data uniformly, failing to account for the distinct effects of conditions like humidity, cloud cover, and temperature on the incoming solar resource.

To address these limitations, this study proposes a prediction framework that segments weather conditions through clustering techniques and applies specialized machine learning models to each group. Historical meteorological data from the National Renewable Energy Laboratory (NREL) are used to identify patterns via K-Means and Gaussian Mixture Models and to evaluate the performance of each clustering method. Cluster-specific predictive models are developed using eXtreme Gradient Boosting (XGBoost) to predict solar irradiance (W/m^2); the predicted irradiance values are converted to estimated solar energy output (kWh) using a PV performance model. Model interpretability is ensured through explainable AI methods such as SHAP values and permutation-based feature analysis. The predictive models are evaluated using standard metrics (RMSE, MAE, R^2) for irradiance prediction, and converted kWh estimates are reported for energy-estimation comparability.

Results indicate that clustering-based models outperform a global baseline, with GMM demonstrating more stable and accurate predictions, particularly under transitional weather conditions. Model interpretability was examined using SHAP and permutation-based feature importance, confirming that dominant predictors align with established physical principles. These findings suggest that probabilistic clustering combined with interpretable machine learning can support more robust and transparent solar energy prediction.

KEYWORDS

Solar Energy Output Prediction; Weather Clustering; K-Means and GMM Clustering; XGBoost; Explainable Artificial Intelligence

1 Introduction

1.1 Background and Motivation

Solar energy is a sustainable power source, yet its reliability is strongly influenced by weather variability. Rapid changes in atmospheric conditions such as cloud cover, humidity, temperature, and wind patterns introduce significant uncertainty in solar irradiance, complicating grid planning, energy distribution, and storage operations. This challenge is particularly pronounced in regions with frequent and transitional weather conditions (Voyant et al., 2017; Ahmed et al., 2020).

Machine learning techniques have been widely adopted to address these challenges, as they are capable of modeling complex nonlinear relationships between meteorological variables and solar irradiance. Ensemble-based and gradient-boosting approaches have demonstrated improved prediction performance compared to traditional statistical methods (Voyant et al., 2017; Ahmed et al., 2020). Nevertheless, many existing approaches rely on global models trained on the entire dataset, implicitly assuming a single underlying relationship across all weather conditions. This assumption often limits predictive performance, as irradiance dynamics differ substantially between clear, cloudy, and mixed atmospheric regimes.

To overcome this limitation, weather-condition clustering has been explored as a means of improving model specialization. By grouping observations into more homogeneous weather regimes, clustering allows prediction models to focus on condition-specific patterns, which reduces noise and enhances robustness (Feng et al., 2018; Blaga et al., 2018). Common clustering methods include K-means, which assigns each observation to a single cluster, and Gaussian Mixture Models (GMM), which provide probabilistic clusters. GMM is particularly well-suited for modeling transitional weather conditions, such as partly cloudy skies, where atmospheric boundaries are inherently overlapping.

Despite improvements in accuracy, the growing complexity of machine learning models has also raised concerns regarding transparency and trust. Many high-performance models function as black boxes, making it difficult to understand how predictions are generated. This lack of interpretability can limit practical adoption, particularly in energy systems where decision-making requires both accuracy and explainability (Mitrentsis & Lens, 2021). As a

result, there is increasing interest in integrating explainable artificial intelligence (XAI) techniques to provide insight into feature importance and model behavior.

1.2 Research Gap and Contributions

Although prior studies have demonstrated the benefits of machine learning, clustering-based prediction, and explainable modeling techniques individually, several gaps remain. Existing work often evaluates clustering strategies without systematically comparing hard and soft clustering approaches under the same experimental framework. In addition, many studies focus primarily on prediction accuracy while providing limited analysis of model stability, generalization behavior across weather regimes or interpretability of the resulting predictions.

This paper addresses these gaps by presenting a unified framework that integrates weather-condition clustering, machine learning-based prediction, and model interpretability for solar energy prediction. Specifically, this study makes the following contributions:

1. It applies and compares K-Means and Gaussian Mixture Model clustering to segment weather data into distinct regimes prior to prediction, enabling cluster-specific model specialization.
2. It develops XGBoost-based prediction models for each cluster and evaluates their performance against a global baseline model using standard regression metrics.
3. It analyzes the effect of clustering strategies on prediction stability and generalization through residual behavior and error distribution analysis.
4. It incorporates explainable artificial intelligence techniques, including SHAP and permutation-based feature importance, to enhance model transparency and validate alignment with physical principles.
5. It converts predicted irradiance into estimated solar energy output using a physically grounded performance model, ensuring practical relevance to real-world photovoltaic systems.

Through this integrated approach, the study aims to demonstrate that combining probabilistic clustering with interpretable machine learning can improve the accuracy, stability, and transparency of solar energy prediction.

2 Related Work

2.1 Weather Clustering for Solar Irradiance Prediction

This section synthesizes prior work on solar irradiance prediction and energy estimation, focusing on weather-condition clustering, machine learning prediction models, and explainable AI. It examines studies on

sustainability, preprocessing techniques, environmental variables, model comparisons, and evaluation metrics.

Di Piazza et al. (2010) established the foundational application of k-means clustering for environmental data processing in energy forecasting. By grouping weather patterns using solar irradiation and wind speed as key variables, their research demonstrated how clustering simplifies energy capability assessment while maintaining forecast accuracy. Subsequent studies expanded on this approach, Jollyta et al.(2023) compared optimization techniques for clustering, while Hsu (2015) highlighted the sensitivity of clustering methods to algorithmic and variable choices, proposing integrated clusterwise regression and two-stage methods to reduce misleading outcomes.

Feng et al. (2018) further showed that cluster-specific predictive models significantly outperform global models, emphasizing the value of segmenting weather patterns prior to forecasting. The present research extends this approach by clustering on multiple weather variables and comparing both K-Means and Gaussian Mixture Models. Furthermore, this study integrates XGBoost with SHAP and permutation importance to provide transparent, cluster-specific predictions.

2.2 Machine Learning and Interpretability in Solar Energy

AlKandari & Ahmad (2019) pioneered ensemble approaches combining deep learning architectures with statistical methods, while Blaga et al. (2018) confirmed that ML-based models consistently outperform traditional methods across diverse climates. More recently, Hayajneh et al. (2024) explored edge deployment of LSTM and BiLSTM models for real-time solar prediction, demonstrating both accuracy and practical applicability. While their work focused on forecasting, it collectively highlights that precise predictions are closely tied to local weather variability, supporting the use of cluster-specific modeling approaches.

Blaga et al. (2018) conducted a comprehensive meta-analysis of 1,705 forecasting tests, establishing that machine learning and hybrid models outperform traditional approaches, particularly in diverse climates. While the study centers on forecasting, its conclusion that accuracy is strongly tied to local weather patterns supports the use of cluster-specific prediction models in this research.

Complementing predictive accuracy, interpretability has become critical for model trust and actionable insight. Mitrentsis and Lens (2021) integrated NGBoost with SHAP values to attribute renewable energy predictions to specific weather variables, enhancing stakeholder confidence. Sazib et al. (2025) and Zafar & Khan (2021) further developed frameworks ensuring consistent, transparent explanations using XAI techniques such as SHAP, LIME, and DLIME. By

combining cluster-based modeling with interpretable machine learning, the present study addresses existing gaps, providing accurate and transparent predictions through the integration of K-Means and Gaussian Mixture Models with XGBoost, complemented by SHAP and permutation importance analyses.

3 Methodology

3.1 Dataset and Variables

The available weather variables, including solar irradiance, temperature, humidity, cloud cover, zenith angle, dew point, wind speed, station pressure, and precipitation, were sourced from the NREL Solar Radiation Research Laboratory located in Golden, Colorado, USA. These meteorological variables were essential for understanding how atmospheric conditions influence solar irradiance intensity, which directly affects the amount of energy a photovoltaic system can generate.

3.2 Data Preprocessing

The preprocessing stage prepared the dataset for effective clustering and prediction by ensuring data completeness, consistency, and feature relevance. First, missing values were handled using K-Nearest Neighbors (KNN) interpolation, which estimated missing readings based on the most similar nearby observations. This method preserves the natural continuity of the weather data without introducing artificial bias. Next, Z-score standardization was applied to all continuous variables to ensure they have a mean of 0 and a standard deviation of 1. This normalization was essential for clustering algorithms such as K-Means and Gaussian Mixture Models (GMM), which rely on distance-based similarity measures. Standardization prevents variables with large numerical ranges (e.g., irradiance in W/m^2) from dominating smaller-scale features like humidity or precipitation. Then, Random Forest feature importance was used for variable selection, ranking meteorological variables according to their influence on solar irradiance. Irrelevant or low-importance variables were excluded to simplify the model and reduce computational load while maintaining interpretability. The final selected variables were used for both weather-condition clustering and prediction.

3.3 Weather-Condition Clustering

To capture diverse weather conditions and their influence on solar irradiance, this study applied unsupervised clustering algorithms, namely K-Means and Gaussian Mixture Models (GMM), on the preprocessed dataset. The optimal number of clusters was predetermined through systematic evaluation during the data preprocessing phase using the Silhouette score metric. Both K-Means and GMM clustering algorithms were applied to the standardized feature space, with candidate values of k ranging from 2 to 6 evaluated.

The Silhouette score measures how well each data point fits with its assigned cluster compared to other clusters, balancing intra-cluster similarity and inter-cluster separation. Each data record was assigned a numerical cluster ID (e.g., 0, 1, 2, 3...), which was used internally to separate the data for model training and testing. No manual labeling or descriptive interpretation (e.g., “sunny,” “cloudy”) was performed, as the focus of this study was on improving model accuracy through objective weather-based segmentation rather than categorical labeling. For each clustering algorithm, the corresponding clustered datasets are used to train and evaluate cluster-specific prediction models using identical procedures. This allowed a fair comparison between K-Means and GMM in terms of their impact on irradiance prediction performance. After clustering, the data within each cluster will be split into 80% for training and 20% for testing. This ensured that each cluster-specific model is trained and evaluated on data belonging to its own weather pattern, maintaining consistency and preventing cross-cluster leakage.

3.4 Prediction Model

In the prediction phase, XGBoost was employed as the primary regression algorithm for each clustered dataset due to its proven ability to handle structured environmental data, model complex nonlinear relationships, and prevent overfitting through regularization. A separate XGBoost model was trained exclusively on each cluster's data, allowing the model to specialize in learning relationships that are unique to specific weather patterns. The output of each XGBoost model is solar irradiance (W/m^2). Since the dataset did not include measured solar energy output, the predicted irradiance values were converted to estimated energy output (kWh) using the PVWatts model, developed by the National Renewable Energy Laboratory (NREL). PVWatts applied standard photovoltaic performance parameters, including panel efficiency, tilt angle, and geographic location, to translate irradiance into realistic energy production values. This conversion ensured that, while the predictive model focused on irradiance, the study still achieved the original objective of estimating solar energy output in practical units. All performance metrics (RMSE, MAE, R^2) were calculated on irradiance predictions. The converted kWh values were used solely for interpretative analysis and visualization, not for direct accuracy evaluation.

3.5 Model Interpretability

SHAP (SHapley Additive exPlanations) was used to explain the contribution of each variable to the solar irradiance predictions, providing insights into how each weather variable affects the final output. SHAP analysis provides a global SHAP summary plot, a ranked chart showing the average contribution of each weather variable to the model's irradiance predictions across the entire dataset, highlighting the most influential weather variables. Cluster-specific SHAP analysis was performed to examine

how the importance of variables changes across different weather conditions. Local SHAP force or Waterfall plots were used to illustrate how each variable increased or decreased the model's predicted irradiance relative to the baseline value.

Permutation Variable Importance was employed to assess the relative importance of each variable, helping to identify which weather conditions and clustering labels were most critical in influencing the model's predictions. Both global and cluster-level permutations were conducted to evaluate the consistency of variable ranking across different weather conditions and to add to the insights obtained from SHAP. These interpretability methods not only help explain the behavior of the models but also offer valuable insights into the key factors that drive solar energy production.

Together, these interpretability techniques ensured that the prediction framework is both accurate and understandable. They clarify the physical and statistical relationships between meteorological inputs and predicted solar irradiance, and by extension, how these irradiance values translate through PVWatts conversion into realistic solar energy output estimates. This interpretability framework enhances the reliability and transparency of the entire modeling process, allowing future researchers and energy planners to understand how weather conditions drive solar energy potential.

3.6 Evaluation Metrics

1. **Root Mean Squared Error (RMSE):** Measures the average magnitude of error between predicted and actual values. It gives higher weight to larger errors. (Lower RMSE values indicate better model performance.)

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2}$$

2. **Mean Absolute Error (MAE):** Measures the average absolute difference between predicted and actual values, without squaring the errors. (Lower MAE indicates better accuracy and a more stable model.)

$$MAE = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i|$$

3. **R-squared (R²):** Indicates the proportion of variance in solar irradiance explained by the model. Higher R² (closer to 1) implies better explanatory power.

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$$

Where:

n = number of observations y_i = actual value

$\hat{y}_i$ = predicted value

$\bar{y}$ = mean of actual values

Model evaluation focused solely on the solar irradiance predictions (W/m²), which served as the core output of the trained machine learning models. Each cluster-specific XGBoost model's predictions were compared against the actual measured irradiance values from the NREL dataset to assess accuracy using RMSE, MAE, and R² metrics. This evaluation framework provided a quantitative measure of each model's predictive capability and allowed for a fair comparison between the K-Means and Gaussian Mixture Model (GMM) clustering approaches. The clustering method that achieves higher R² and lower error metrics was considered more effective for modeling irradiance variability under different weather conditions. The predicted solar irradiance values were subsequently converted into corresponding solar energy output estimates (kWh) using the PVWatts performance model developed by the National Renewable Energy Laboratory (NREL). This conversion ensured that the study's outputs remain consistent with realistic photovoltaic energy production values. While model performance was evaluated using irradiance-based metrics, the resulting energy estimates provided a practical interpretation of the potential solar generation associated with the predictions.

4 Results and Discussion

4.1 Prediction Performance Comparison




Figure 1: Evaluation Metrics of Global, KMM, and GMM

To evaluate whether clustering improves solar irradiance prediction, the performance of the global model was compared against the two cluster-based ensembles (K-Means and GMM). All models were assessed using the same evaluation metrics (R², RMSE, and MAE), allowing a fair and consistent comparison across approaches. The global model serves as the baseline, representing a single XGBoost regressor trained on the entire dataset without distinguishing between weather regimes. In contrast, the clustered models first group weather conditions into four distinct clusters and then train a specialized model for each group. This structure is designed to let each sub-model

learn patterns specific to its weather regime, potentially improving predictive accuracy under varying sky conditions.

Across all evaluation metrics, both clustering approaches demonstrated improved performance over the global model. The clustered models achieved higher R^2 values and lower RMSE/MAE, indicating better predictive accuracy and reduced error. Between the two, the GMM-based ensemble produced the strongest overall results, suggesting that soft-probability clustering is more effective at capturing subtle weather variations compared to the hard assignments of K-Means. These results collectively show that introducing a clustering stage enhances model specialization and stability, leading to more reliable irradiance predictions across diverse environmental scenarios. Figure 1 shows the comparison of these performance metrics. The K-Means and GMM cluster-based models outperformed the Global model across all evaluation measures. Specifically, the GMM model achieved the highest R^2 (0.9497) and the lowest RMSE (62.54 W/m²) and MAE (37.72 W/m²). These results indicate that incorporating clustering improves predictive accuracy and stability, reducing the average prediction error by approximately 12% compared to the global model.

4.2 Effect of Clustering Strategy

Beyond improvements in aggregate performance metrics, the choice of clustering strategy influenced how the models behaved across different weather conditions. To enhance prediction accuracy, weather data were grouped into clusters based on patterns in key meteorological variables such as solar irradiance, temperature, and wind speed. This segmentation allowed the system to focus on weather-specific patterns, improving solar irradiance predictions and reducing noise in energy estimation results.

K-Means was used for hard-clustering, assigning each weather instance to a single category. This approach is effective for well-defined conditions, such as fully sunny or fully cloudy days. Overall K-Means effectively segments the dataset, enabling each model to specialize in distinct weather behaviors while maintaining good overall predictive consistency.

In contrast, Gaussian Mixture Models (GMM) enabled soft clustering, representing overlapping conditions. This is particularly useful for transitional weather conditions, such as partly cloudy day common in tropical regions. GMM exhibits smoother, probabilistic cluster transitions, allowing the models to generalize slightly better in mixed-weather conditions.



Figure 2: **Residuals vs. Fitted Values for Global, K-Means, and GMM Models**

Figure 2 presents the residual distribution histograms for the Global, K-Means, and GMM models, illustrating how prediction errors are spread around zero. All three models exhibit a near-normal error pattern centered close to zero, indicating that the predictions are unbiased and do not consistently overestimate or underestimate irradiance. However, the width of each histogram varies significantly: the Global model shows a broader residual spread, with more frequent large deviations, reflecting its difficulty in capturing localized variations in weather behavior.

In contrast, both clustering-based models produce tighter error distributions, demonstrating improved ability to model nonlinear and cluster-specific irradiance patterns. The K-Means and GMM models show sharper and more pronounced peaks around zero, with noticeably fewer extreme errors on both sides of the distribution. This narrower concentration of residuals indicates higher predictive precision and reduced variability, which also means that it is consistent with their superior RMSE and MAE scores. The GMM model, in particular, has the most compact distribution, reinforcing its strength in modeling continuous, probabilistic boundaries between weather regimes. Overall, these residual plots confirm that clustering before prediction enhances stability and enables the models to generalize more effectively across diverse atmospheric conditions.




Figure 3: **Residual Distribution of Prediction Errors for Global, K-Means, and GMM Models**

Together, these plots validate the models' calibration and confirm that the integration of clustering enhances predictive stability. The consistent residual spread across all irradiance ranges and the near-normal error distributions indicate that the proposed models generalize effectively across diverse atmospheric conditions, fulfilling the robustness and reliability goals of the study.

4.3 Model Interpretability Analysis




Figure 4: SHAP Summary Plot (Global Model)

Figure 4 illustrates the SHAP Beeswarm Plot, which further details how each feature contributes to prediction variability across all data samples. Here, each point represents an individual prediction instance, colored according to the feature value (red for high, blue for low). High zenith angles (in blue) and high cloud cover (also in blue) are consistently associated with negative SHAP values, reaffirming their suppressive effect on solar irradiance. In contrast, low zenith angles and low cloud cover conditions (in red) lead to positive SHAP values, representing stronger irradiance levels. The compact clustering of points in the lower-ranked features (humidity, dew point, and wind speed) emphasizes their limited but context-dependent contribution.

5 Conclusion

This study proposed a predictive framework for solar irradiance and energy output estimation that integrates weather-condition clustering, machine learning prediction, and model interpretability. Solar energy prediction is inherently challenging due to the strong dependence of irradiance on fluctuating atmospheric conditions, which complicates accurate energy estimation and grid planning. To address this challenge, the study employed unsupervised clustering techniques, specifically K-Means and Gaussian Mixture Models (GMM), to segment weather data into distinct regimes prior to prediction.

Cluster-specific prediction models were developed using XGBoost, allowing each model to specialize in learning

patterns associated with particular weather conditions. Model performance was evaluated using standard regression metrics, including R^2 , RMSE, and MAE, with results showing that clustering-based approaches consistently outperformed a global model trained on the full dataset. Among the two clustering strategies, the GMM-based ensemble achieved superior predictive accuracy, indicating that soft-probabilistic clustering is more effective in capturing overlapping and transitional weather patterns than hard clustering methods such as K-Means.

Beyond prediction accuracy, this study emphasized transparency and trust through the integration of explainable artificial intelligence techniques. SHAP and Permutation Variable Importance analyses revealed that key meteorological variables such as solar zenith angle and cloud cover exert the strongest influence on solar irradiance predictions, aligning with established physical principles. These interpretability results confirmed that the models not only performed well statistically but also behaved consistently with domain knowledge.

The predicted solar irradiance values were subsequently converted into estimated solar energy output using the PVWatts performance model, ensuring that the results remained physically meaningful and applicable to real-world photovoltaic systems. Overall, this research demonstrated that combining weather-pattern clustering with interpretable machine learning provides a robust and transparent approach to solar energy output prediction. The findings contribute to the advancement of data-driven renewable energy modeling and offer practical insights for improving solar integration and energy planning under variable weather conditions.

REFERENCES

- [1] M. Di Piazza, G. La Tona, and M. Luna (2023). Day-ahead forecasting of residential electric power consumption for energy management using Long Short-Term Memory encoder-decoder model. <https://www.sciencedirect.com/science/article/pii/S0378475423002720>
- [2] Jollyta, R., Li, X., & Zhao, M. (2023). Analysis of an optimal cluster approach. IOP Conference Series: Earth and Environmental Science, 2421(1), 012015. <https://iopscience.iop.org/article/10.1088/1742-6596/2421/1/012015>
- [3] D. D. Hsu (2015). Clusterwise regression and two-stage model analysis for renewable energy forecasting. Energy, 93, 532–541. <https://www.sciencedirect.com/science/article/pii/S0306261915010624>
- [4] C. Feng, J. Zhang, and J. Wang (2018). An Unsupervised Clustering-Based Short-Term Solar Forecasting Methodology Using Multi-Model Machine Learning Blending, 11(8), 2168. <https://arxiv.org/abs/1805.04193>
- [5] A. AlKandari and S. Ahmad (2019). Solar Power Generation Forecasting Using Ensemble Approach Based on Deep Learning and Statistical Methods. ResearchGate. <https://www.researchgate.net/publication/337064881>
- [6] A. Blaga, M. Sabadus, C. Stefu, N. Dughir, and R. Paulescu (2018). A current perspective on the accuracy of solar power forecasting. Renewable and Sustainable Energy Reviews, 88, 367–386. <https://www.researchgate.net/publication/327340464> Online Forecasting of the Solar Energy Production
- [7] T. Hayajneh, A. Al-Ashqar, and J. Mahfoud (2024). Intelligent solar forecasts: Modern machine learning models and TinyML role for improved solar energy yield predictions. MMU E-Space. https://e-space.mmu.ac.uk/636830/1/Intelligent_Solar_Forecasts_Moder

[n_Machine_Learning_Models_and_TinyML_Role_for_Improved_Solar_Energy_Yield_Predictions.pdf](#)

- [8] G. Mitrentsis and M. Lens (2021). An Interpretable Probabilistic Model for Short-Term Solar Power Forecasting Using Natural Gradient Boosting. *Energies*, 14(17), 5405.
<https://arxiv.org/abs/2108.04058>
- [9] M. S. Sazib, M. Hasan, M. S. Khan, and M. J. Al-Raqqad (2025). Advancing Renewable Energy Systems through Explainable Artificial Intelligence: A Comprehensive Review and Interdisciplinary Framework. *Journal of Computer Science and Telecommunications Studies*.
<https://al-kindipublishers.org/index.php/jcsts/article/view/8841>
- [10] Zafar, M., & Khan, S. (2018). Deterministic Local Interpretable Model-Agnostic Explanations for Stable Explainability. *Data*, 3(3), 27.
<https://www.mdpi.com/2504-4990/3/3/27>
- [11] C. Voyant, G. Notton, S. Kalogirou, M. L. Nivet, C. Paoli, F. Motte, and A. Foulloy (2017). Machine learning methods for solar radiation forecasting: A review. *Renewable Energy*, 105, 569–582.
<https://www.sciencedirect.com/science/article/abs/pii/S0960148116311648>
- [12] R. Ahmed, V. Sreeram, Y. Mishra, and M. D. Arif (2020). A review and evaluation of the state-of-the-art in PV solar power forecasting: Techniques and optimization. *Renewable and Sustainable Energy Reviews*, 124, 109792.
<https://www.sciencedirect.com/science/article/abs/pii/S1364032120300885>

A Hybrid Machine Learning Framework for Predictive Maintenance of Bearings in Philippine Power Plants

Beaux Consunji*

Ateneo de Manila University
Manila, Philippines
beaux.consunji@student.ateneo.edu

Joshua Zapico*

Ateneo de Manila University
Manila, Philippines
joshua.zapico@student.ateneo.edu

Lamberlain Muli*

Ateneo de Manila University
Manila, Philippines
lamberlain.muli@student.ateneo.edu

John Vergara

Ateneo de Manila University
Manila, Philippines
jpvergara@ateneo.edu

ABSTRACT

Predictive maintenance (PdM) has gained interest in power plant operations, yet its practical application is often limited by the scarcity and uncertainty of labeled failure data. Rotating machinery such as Primary and Secondary Air Fans generate large volumes of sensor measurements, while actual fault events remain rare and inconsistently documented, restricting the effectiveness of purely supervised machine learning models. This paper presents a hybrid machine learning framework that combines unsupervised anomaly detection with supervised classification to support bearing fault prediction under data-limited industrial conditions.

The approach begins with ISO-based state annotation, where minute-level Distributed Control System (DCS) vibration data are labeled using threshold criteria derived from ISO 10816 standards. A forecasting-oriented target variable is then defined to indicate whether a fault onset occurs within a 24-hour horizon, enabling fault prediction. The time series is segmented into windows, from which statistical and temporal features are extracted to characterize system behavior.

The supervised prediction target is derived from the plant's monthly expert vibration spectra: each maintenance-confirmed fault finding is matched to the continuous DCS record to determine the temporal extent of the corresponding fault episode, producing labels that reflect real maintenance events rather than threshold crossings alone. Five classifiers are evaluated on a chronological test set with a 24-hour temporal gap to prevent information leakage.

Experiments on three years of real-world data from four fan units at a Philippine coal-fired power plant show that XGBoost achieves the best F1-score of 0.774 and ROC AUC of 0.937 for 24-hour fault-onset prediction. A cost-sensitive deployment analysis using real maintenance costs shows that, at episode-level thresholds calibrated for 100% fault detection, all models catch every fault episode with fewer than two false alarms per month, projecting escalation-adjusted annual savings of PHP 22.5M. The results demonstrate that hybrid machine learning offers a practical direction for PdM in environments where labeled fault data are limited.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; **Supervised learning**; *Anomaly detection*; *Time-series analysis*; • **Applied computing** → *Engineering*; *Fault diagnosis*;

KEYWORDS

predictive maintenance, condition monitoring, bearing fault detection, hybrid machine learning, anomaly-detection

1 INTRODUCTION

Reliable and efficient power generation is essential to modern society, yet power plants operate complex, capital-intensive machinery under harsh and highly variable conditions [27], [7]. Critical assets such as turbines, pumps, generators, and large industrial fans must maintain high availability to avoid costly disruptions and potential grid instability [27], [11]. Unexpected failures in rotating equipment, which are often caused by bearing wear, misalignment, imbalance, or looseness, can lead to significant economic losses and safety risks [1], [4]. For example, a partner coal-fired power plant in Batangas, Philippines report that unplanned downtime of a single 300 MW unit can exceed 1.05 million pesos per hour in lost generation alone. These realities motivate the growing adoption of Predictive Maintenance (PdM), which seeks to anticipate failures using data-driven models rather than relying on reactive or fixed-schedule maintenance strategies.

Machine learning (ML) has become a core enabler of PdM. However, its deployment in real-world power plants faces persistent data challenges. Although condition-monitoring sensors generate large volumes of vibration, temperature, and operational time-series data, true failure events are rare because equipment is designed for high reliability [22], [17]. As a result, labeled fault samples required by supervised ML are scarce [5], [13]. Moreover, maintenance logs often contain delayed, inconsistent, or ambiguous labels, since diagnosing precise fault causes from sensor signatures is nontrivial and depends on expert judgment [22], [19]. These issues are compounded by the high dimensionality, noise, missing values, and operational variability typical of industrial sensor data [5], [18], [10]. Consequently, standard supervised approaches struggle to generalize, while purely unsupervised methods, although useful for detecting anomalies, often lack the specificity needed for accurate fault diagnosis and prognostics [18], [19].

*These authors contributed equally.

This work is motivated by the need for PdM methods that remain effective under realistic industrial constraints characterized by limited and imperfectly labeled data. Rather than relying exclusively on supervised or unsupervised learning, we explore a hybrid machine learning framework that leverages abundant unlabeled operational data while strategically incorporating the small number of available labeled fault cases. By combining unsupervised clustering for structure discovery with supervised classification for fault identification, the proposed approach aims to improve bearing fault detection in power plant equipment under data-scarce conditions. Using real-world vibration and operational data from a Secondary Air fan in a coal-fired power plant in Batangas, this study seeks to demonstrate a practical pathway for deploying ML-based PdM in environments where data quality and quantity fundamentally limit conventional solutions [3], [4].

2 REVIEW OF RELATED LITERATURE

2.1 Machine Learning for Predictive Maintenance

Industry 4.0 has enabled large-scale collection of operational data, supporting Prognostics and Health Management (PHM) through machine learning (ML) [7]. ML-based predictive maintenance (PdM) improves fault detection, reduces downtime, and optimizes maintenance scheduling compared to traditional strategies such as Run-to-Failure, Preventive Maintenance, and Condition-Based Maintenance. By learning from historical data, ML models can detect early degradation patterns without explicit programming, making them suitable for high-dimensional, multivariate sensor data [24]. PdM systems typically follow a pipeline of data selection, preprocessing, model training, validation, and deployment, with bearing fault prediction remaining a common benchmark task.

2.2 Data Preprocessing in Predictive Maintenance

Effective data preprocessing is critical for reliable PdM models, particularly in industrial environments characterized by large data volumes, high velocity, heterogeneity, and data quality issues [5, 22]. Common challenges include noisy sensor signals, missing values, inconsistent data formats, scarce labeled fault data, and severe class imbalance [3, 17]. Addressing these issues directly impacts model performance, robustness, and interpretability.

2.2.1 Data Quality and Integrity. Industrial sensor data often contains non-numeric entries, errors, and noise. Converting values to numeric formats while coercing invalid entries to NaN is a standard approach to isolate unreliable readings and enable systematic handling of missing data [5]. Noise reduction techniques such as smoothing filters, Fourier analysis, or wavelet transforms may be applied depending on signal characteristics [5, 17].

2.2.2 Handling Missing Data. Missing data commonly arises from sensor faults or communication failures. Simple deletion of incomplete records risks information loss and bias [5]. While basic imputation methods (mean or median replacement) are computationally efficient, they ignore feature relationships. Interpolation methods are effective for short gaps in time-series data, whereas

K-Nearest Neighbors (KNN) imputation leverages multivariate correlations to estimate missing values more accurately [5, 13]. Due to its balance of accuracy and practicality, KNN imputation is widely adopted in PdM applications.

2.2.3 Labeling Strategies and Class Selection. Industrial maintenance logs may contain multiple fault labels for a single event, which can be converted into single-label datasets through data expansion [5]. In many PdM contexts, unlabeled instances are reasonably assumed to represent healthy operation. Additionally, fault data often follows a long-tail distribution, with many rare fault types. Focusing on healthy data and the most frequent or critical fault classes simplifies classification and improves model reliability under limited labeled data conditions [10].

2.2.4 Class Imbalance Handling. Class imbalance is a fundamental challenge in PdM, as healthy samples significantly outnumber fault samples [5]. Standard classifiers tend to favor the majority class, reducing fault detection capability. Oversampling methods such as the Synthetic Minority Over-sampling Technique (SMOTE) address this issue by generating synthetic minority samples through interpolation, improving class separability without duplicating data [6]. SMOTE and its variants have been shown to enhance fault classification performance when applied exclusively to training data to prevent data leakage [8].

2.3 Supervised Machine Learning for Predictive Maintenance

Supervised ML models learn mappings between sensor features and labeled outcomes such as fault classes or Remaining Useful Life (RUL) estimates [23]. Their effectiveness depends on data quality, label availability, and model selection.

2.3.1 Classification Models. Decision Trees offer interpretable rule-based structures and low computational cost but are prone to overfitting [2]. Support Vector Machines (SVMs) perform well in high-dimensional and small-sample settings, especially with kernel functions that capture non-linear relationships [21]. K-Nearest Neighbors (KNN) is simple and effective for irregular decision boundaries but is sensitive to feature scaling, dimensionality, and computational cost during inference [5]. Model choice involves trade-offs between interpretability, accuracy, and computational constraints.

2.3.2 Neural Networks and Deep Learning. Neural Networks (NNs) and deep learning models excel at modeling complex, non-linear degradation patterns and learning features directly from raw sensor data [5]. Recurrent architectures such as LSTM and GRU are well-suited for time-series PdM tasks, while CNNs can extract spatial-temporal features from transformed sensor signals [9, 26]. Despite strong performance, NNs require large labeled datasets, substantial computational resources, and careful hyperparameter tuning, and they often lack interpretability [16].

2.3.3 Regression Models. Linear regression remains useful as a baseline for RUL prediction when degradation trends are approximately linear, offering simplicity and interpretability [25]. Logistic regression provides interpretable probabilistic outputs for binary fault classification but is limited in modeling complex non-linear

relationships [23]. These models are often used for benchmarking or when transparency is a primary requirement.

2.4 Unsupervised and Semi-Supervised Learning for Predictive Maintenance

Unsupervised learning methods identify patterns in unlabeled data and are particularly valuable when fault labels are scarce. Clustering techniques such as K-means, hierarchical clustering, DBSCAN, and Gaussian Mixture Models have been applied to group operational states and detect anomalies in PdM applications [12, 14, 20]. Autoencoder-based models, including Variational Autoencoders, learn representations of normal system behavior and detect faults through reconstruction errors [28]. Semi-supervised approaches combine limited labeled data with large unlabeled datasets, improving early fault detection and fault progression modeling while reducing dependence on extensive annotations [15].

3 METHODOLOGY

3.1 Overview




Figure 1: Hybrid Machine Learning Framework for Predictive Maintenance

The main pipeline consists of three core phases. First, data procurement, where DCS datasets are combined with maintenance records and operating or outage logs. Second, pre-processing, which includes ISO-based annotation, horizon labeling, and sliding-window featurization to transform the time series into machine learning model-ready inputs. Third, fault prediction, where features are engineered, missing values are handled through KNN imputation, and the dataset is split into training and test sets. Class imbalance is addressed using SMOTE, followed by supervised model training and benchmarking.

An optional human-in-the-loop (HITL) layer runs alongside the pipeline to enhance model performance. This stage includes filtering, autoencoder-based anomaly detection, clustering, dimensionality reduction, anomaly visualization, and expert validation.

3.2 Data Procured

3.2.1 DCS Data. The dataset used in this module was collected from the plant’s Distributed Control System (DCS) between 27 October 2022 and 27 October 2025 for the Primary and Secondary Air Fans at a 1 minute interval. The selected features consisted of vibration signals recorded from the bearings and motor bearings of both the Primary Air Fans (PAF) and Secondary Air Fans (SAF), including Fan A and Fan B, in the X and Y directions. These vibration measurements were chosen because they directly capture mechanical imbalances, misalignment, and wear in rotating machinery components, making them reliable indicators of equipment health and early signs of potential mechanical failure.

3.2.2 Operating and outage records. The plant maintains a log of operating hours, outages, and the exact date and time of desynchronization and resynchronization events since 2023. These entries were used to identify periods of off-normal operation and to exclude desynchronized intervals when preparing data for unsupervised modeling. The log was also used to distinguish continuous operation from outage hours when interpreting sensor behavior across events.

3.2.3 Monthly Expert Vibration Spectra. Monthly vibration spectra analyses have been conducted by the plant since 2016 as part of scheduled maintenance. Each record includes the date and time, work order reference, and detailed vibration readings from motor and fan bearings across vertical, horizontal, axial, and casing directions. Temperature measurements for the motor and fan bearings are also logged, along with remarks on test conditions and corrective actions performed. The expert provides assessments such as "Misalignment", "Rough", or "Looseness" together with recommendations for alignment or component replacement.

3.2.4 Maintenance Data. The dataset documents preventive and corrective maintenance activities for Unit 3 Primary and Secondary Air Fans (PAF/SAF) and their motors. Each sheet corresponds to a specific component and lists work orders conducted up to 2025. Records include work order number, equipment description, maintenance schedule reference, work type, work center, and timestamps for job start and completion. The "Work Done" field summarizes inspection results and actions such as bearing replacement, sensor repair, or cleaning, often noting vibration and condition assessments (e.g., "MOTOR - GOOD", "COUPLER - FAIR"). Each entry indicates completion status and, when available, fault descriptions and plant identifiers for traceability.

3.3 Pre-processing

3.3.1 ISO-based State Annotation. For each observation, the maximum absolute vibration value across all bearing-related channels was computed. This aggregate measure captures the highest level of mechanical stress present at a given time step. The resulting value was then compared against a fixed threshold of 2.3 mm/s.

Observations with peak vibration values below the threshold were labeled as healthy, while those exceeding the threshold were labeled as faulty. This threshold was selected based on ISO 10816 guidelines and reflects conservative limits for acceptable vibration levels in industrial rotating equipment.

3.3.2 Prediction Horizon Labeling. A prediction horizon was defined to capture whether a fault onset occurs within a future time interval. For each observation at time t , a binary horizon label was assigned based on the occurrence of a fault onset within the interval $(t, t + 24h]$.

The resulting horizon variable serves as the primary supervised learning target for downstream models. The 24-hour prediction window was selected based on operational considerations, providing sufficient lead time for maintenance planning while preserving proximity to the actual fault event.

3.3.3 Sliding Window Featurization. The horizon-labeled time series for each fan was segmented using a rolling window approach. Each window spans a duration of 2 minutes, with the sampling frequency inferred from the median inter-sample time interval. This ensures consistency in window construction despite minor variations in timestamp spacing.

For each window, a set of time-series features was extracted from every numeric sensor channel using the `tsfresh` library. These features capture both distributional properties and temporal dynamics of the vibration signals.

Each window was assigned a binary label based on the presence of future fault events within the window. A window was labeled as faulty if any observation within the window had `horizon = True`, and healthy otherwise. This labeling strategy ensures alignment between input features and the prediction objective of anticipating fault onsets within the defined horizon.

3.4 Human-in-the-loop

3.4.1 Autoencoder. The anomaly detection module is designed to identify unusual operating behavior in the machinery by learning the characteristics of normal conditions and flagging deviations as potential anomalies.

The module was implemented as a fully connected autoencoder with symmetric encoder and decoder layers. The encoder progressively reduced the dimensionality of each input window into a compact latent representation, while the decoder attempted to reconstruct the original input. A bottleneck layer of dimension 16 served as the compressed representation of the data.

The network was trained using the Adam optimizer and mean squared error (MSE) as the reconstruction loss. 20% of the healthy data was reserved for validation to monitor generalization performance. This design allowed the model to capture essential patterns of normal operation while minimizing overfitting.

Following training, the autoencoder was applied to both normal and test windows. For each window, the reconstruction error was computed as the mean squared difference between the input and its reconstruction. The distribution of reconstruction errors obtained from healthy data was then analyzed to establish an anomaly threshold which was optimized with respect to the model's F1-score. Windows whose reconstruction errors exceeded this threshold were flagged as anomalous.

3.4.2 Clustering. Following anomaly detection via the autoencoder, only samples predicted as faulty are retained for unsupervised refinement. Standardized `tsfresh`-derived features from these

samples are clustered to produce pseudo-labels representing different fault sub-types.

K-Means clustering is applied with values of k ranging from 2 to 12 to analyze inertia reduction and silhouette quality. Centroid initialization uses `k-means++`, and multiple sampling attempts are performed per initialization to improve convergence stability. Cluster quality is assessed using both the silhouette coefficient and the elbow method.

A subject-matter expert will validate the cluster meaning through an interactive visualization interface that highlights centroid characteristics and anomalous feature distributions. This step determines which clusters correspond to meaningful degradation signatures versus benign variability.

3.5 Fault Prediction

3.5.1 Rationale for Binary Classification Strategy. As highlighted in the Introduction (Section 1.1), the industrial dataset obtained for this study is characterized by significant data scarcity and severe class imbalance. Specific failure modes, such as bearing breakage or distinct cage faults, represent rare operational events. Training a multi-class supervised model to identify these individual fault types is statistically unfeasible, as the insufficient number of positive samples for each specific category would prevent the model from learning generalizable decision boundaries.

Furthermore, domain experts consulted for this study emphasized that accurately diagnosing specific fault types (e.g., distinguishing between misalignment and looseness) strictly from aggregated numerical data is prone to error. They noted that a reliable specific diagnosis requires visual inspection of the full vibration spectrum (e.g., Fast Fourier Transform analysis), which is not captured in standard numerical logs. Consequently, attempting a granular multi-class classification using only the available scalar features risks introducing significant label noise.

To address these statistical and domain-specific constraints, the methodology adopts a binary classification strategy. By aggregating distinct failure signatures into a single collective "Faulty" class, the effective pool of positive samples is significantly increased. This aggregation mitigates the limitations imposed by rare-event scarcity and avoids the pitfalls of specific fault misdiagnosis, ensuring a robust distinction between "Healthy" and "Faulty" operational states.

3.5.2 Supervised Model Classification. All four per-fan sliding-window datasets are concatenated into a single pooled dataset. The data is split chronologically, with the earliest 80% used for training and the most recent 20% reserved for testing. A 24-hour temporal gap is enforced between the training and test periods to prevent information leakage from overlapping sliding windows. Missing values are imputed using k -nearest-neighbor imputation ($k = 5$), features are standardized, and SMOTE (sampling strategy = 0.4, $k = 3$) is applied to the training split to mitigate class imbalance.

Five classifiers are trained and compared: Logistic Regression (class_weight=balanced), Random Forest (500 estimators, max depth 30, capped at 200,000 training rows), Histogram Gradient Boosting (HGB, learning rate 0.01, max depth 4, capped at 300,000 rows), Multilayer Perceptron (MLP, hidden layers 64–32, capped at 150,000 rows), and XGBoost (100 estimators, max depth 5, capped

at 300,000 rows). All models use `random_state=42`. F1-score is the primary ranking metric, with PR AUC as the secondary metric.

4 RESULTS

4.1 Autoencoder




Figure 2: Autoencoder AE Loss Diagram

4.1.1 Training Behavior. Figure 2 illustrates the training and validation loss of the autoencoder over 100 epochs. Both curves decrease rapidly during the initial epochs, indicating that the model quickly learns the dominant patterns of normal operating behavior. The training loss stabilizes at approximately 0.0051–0.0053, while the validation loss fluctuates within a slightly higher range of approximately 0.0055–0.0065. Although the validation curve exhibits minor variability compared to the smoother training curve, the overall gap between the two remains small. This suggests that the model achieves good generalization to unseen healthy data, with no significant signs of overfitting.

4.1.2 Reconstruction Error Characteristics. The reconstruction error distribution for nominal data is shown in Figure 3. The errors




Figure 3: Autoencoder Reconstruction Error Distribution

are highly concentrated near zero, indicating that the autoencoder is able to consistently reconstruct normal operating patterns with low deviation. A smaller secondary concentration of samples appears at higher error values, suggesting the presence of less typical but still nominal variations in the data.

An anomaly threshold of 0.002123 is used to distinguish typical behavior from potential deviations. This threshold lies beyond the dense region of low reconstruction errors while remaining below the higher-error tail, enabling effective separation of normal and anomalous patterns. Only a small proportion of samples exceed this threshold, indicating a conservative yet sensitive criterion for anomaly-onset detection.

4.1.3 Anomaly-detection Performance. Table 1 shows that the autoencoder model achieved an F1-score of 0.6932, with a precision of 0.5616 and recall of 0.9053. The high recall indicates that the majority of fault instances are successfully detected, demonstrating the model’s effectiveness in capturing potential anomalies. However, the comparatively lower precision suggests a higher rate of false positives, meaning that some normal instances are incorrectly flagged as anomalous.

This trade-off reflects a conservative detection strategy, where the model prioritizes minimizing missed faults at the expense of increased false alarms. In the context of predictive maintenance, such behavior is often acceptable in early-stage filtering, as it ensures that critical fault events are unlikely to be overlooked. Overall, the results indicate that the autoencoder serves as a sensitive first-stage detector for identifying fault-like behavior, with subsequent stages required to refine and validate these detections.

Metric	Value
F1-score	0.6932
Precision	0.5616
Recall	0.9053

Table 1: Autoencoder Metrics Summary

4.2 Clustering

4.2.1 Cluster Evaluation and K-Value Selection. K-Means clustering was evaluated across $k = 2, 3, \dots, 12$ using inertia (within-cluster sum of squares) and silhouette analysis.




Figure 4: Combined Elbow Method and Silhouette Analysis from $k = 2$ to 12.

Figure 4 shows that inertia decreases consistently with increasing k , with larger reductions up to $k = 5$ (–11.27%) and diminishing

gains thereafter (-3.7% to -7.32%). Silhouette scores gradually improve from 0.4273 at $k = 2$ to 0.5101 at $k = 10$, though $k = 2$ is disregarded due to poor cluster granularity. Values of $k = 5$ –6 capture strong inertia reduction but may underrepresent distinct machine states, whereas $k = 10$ balances cluster separation and granularity, providing finer representation for human-in-the-loop (HITL) validation.




Figure 5: Silhouette profiles for $k = 2, 5, 6, 10$.

Per-cluster silhouette profiles in Figure 5 reinforce this choice: $k = 2$ forms a dominant cluster with low granularity, $k = 5$ –6 introduces new clusters but some have low silhouette values, and $k = 10$ achieves the highest average score (0.5101) with mostly positive values, capturing finer variations in machine behavior.

Figure 6 presents the $k = 10$ clustering results projected into three principal components. Several clusters appear compact and well-isolated, while others exhibit partial overlap, reflecting transitional machine states. The presence of smaller, distinct clusters alongside larger, dense groupings suggests that the model captures both common operational behavior and less frequent, potentially anomalous patterns. This enhanced spatial representation supports more intuitive interpretation and facilitates more precise HITL labeling of machine states.

4.3 Fault Prediction

4.3.1 Baseline Model Performance. The pooled four-fan dataset comprised 4,986,918 sliding windows with 112 shared numeric predictors (80 tsfresh statistical features and 32 rolling temporal features computed at 1-hour and 6-hour windows). The prediction target was derived from 19 maintenance-confirmed fault episodes across four fans, identified from expert inspection findings. After the chronological 80/20 split with a 24-hour gap, the training set contained 3,986,652 windows and the test set contained 994,500 windows. Table 2 summarizes the performance of the five classifiers.

XGBoost achieved the highest F1-score (0.774) and ROC AUC (0.937), followed closely by HGB (F1 = 0.770, highest precision at




Figure 6: K-Means clustering with centroids visualized in 3D PCA space.

Model	F1	Prec.	Rec.	PR AUC	ROC AUC
XGBoost	0.774	0.759	0.789	0.617	0.937
HGB	0.770	0.782	0.759	0.625	0.926
Logistic Reg.	0.650	0.524	0.854	0.640	0.919
Random Forest	0.402	0.646	0.292	0.608	0.896
MLP	0.277	0.635	0.177	0.315	0.840

Table 2: Supervised Model Performance Comparison




Figure 7: Comparison of supervised model performance across evaluation metrics.

0.782). Both gradient-boosted models showed strong precision, indicating fewer false alarms when predicting maintenance-confirmed faults. Logistic Regression was notably competitive (F1 = 0.650, highest recall at 0.854), suggesting that the fault signal derived from expert inspections is well-separated in feature space. Random Forest and MLP underperformed, likely because their training-row caps (200,000 and 150,000 respectively) limit the number of positive samples available under the sparse maintenance-derived labels. Figure 7 visualizes the performance spread across models. The use of a chronological split with a temporal gap provides a realistic

forward-looking evaluation; the resulting F1-scores reflect genuine predictive difficulty rather than the inflated performance that can arise from stratified or random splits on time-series data.




Figure 8: Confusion matrix for the XGBoost classifier on the chronological test set.

Figure 8 shows the confusion matrix for the best-performing XGBoost model. The model correctly identifies the majority of fault windows (TP = 56,782) while maintaining a low false-positive rate (FP = 17,998 out of 922,504 healthy windows), corresponding to a 2.0% false-alarm rate at the window level.




Figure 9: Top 20 Feature Importance from the XGBoost model

4.3.2 Feature Importance. Figure 9 presents the top 20 features from the XGBoost model. The model relies primarily on statistical aggregates of vibration signals, with mean and minimum amplitude features among the dominant predictors. These features encode systematic shifts in baseline vibration behavior that precede maintenance-confirmed fault episodes, suggesting that monitoring changes in the vibration “floor” is as important as tracking peak amplitudes for early bearing fault detection.

4.3.3 Cost-Sensitive Deployment Analysis. Window-level F1 understates the operational value of these models because plant operators act on *fault episodes*, not individual windows. A cost-sensitive evaluation was therefore conducted using real maintenance cost data provided by the industry partner.

Three cost components were derived from historical maintenance records. **Opportunity cost** represents lost electricity revenue during unplanned outages, computed from the plant’s average energy loss per incident (5,829 MWh) at PHP 3,500/MWh, averaging PHP 20.4M per event across 4 shutdowns and 18 deratings. **Hardware cost** covers bearing pair replacement at PHP 45K. **Labor cost** is PHP 82K for planned maintenance or 1.5× that for emergency repairs. Together, these yield a missed-fault cost of PHP 20.6M per incident versus a planned-repair cost of PHP 128K, while a false alarm incurs only the inspection labor cost of PHP 1,659K—a cost asymmetry ratio of 12413:1.

This asymmetry means the decision threshold should be set to catch every fault, not to maximize F1. For each model, we identify the *max-safe threshold*: the highest probability threshold at which 100% of the 8 maintenance-confirmed fault episodes in the test set are still detected. Table 3 compares the models at their max-safe operating points.

Model	Threshold	Ep. Recall	FA/month	Win. F1
Logistic Reg.	0.783	100%	1.8	0.769
MLP	0.773	100%	1.6	0.207
Random Forest	0.133	100%	1.6	0.709
XGBoost	0.089	100%	1.8	0.507
HGB	0.084	100%	0.9	0.220

Table 3: Model Performance at Max-Safe Deployment Threshold (100% Fault Episode Detection)

All five models achieve 100% episode recall with fewer than two false alarms per month—comparable to one to two additional routine inspections on top of the plant’s existing monthly schedule. Logistic Regression is notable for achieving full episode detection at the highest threshold (0.783) while maintaining a window-level F1 of 0.769, indicating the greatest separation between fault and non-fault score distributions for maintenance-confirmed events. Adjusting for the historical escalation rate (6.5% of fault episodes progress to actual shutdowns or deratings), the projected annual savings from early detection are approximately PHP 22.5M.

5 CONCLUSION

This paper presented a hybrid machine learning framework for predictive maintenance under realistic industrial constraints, where labeled failure data are scarce and often imperfect. Using three years of vibration data from four fan units (PAF and SAF) at a Philippine coal-fired power plant, the proposed approach integrates unsupervised anomaly detection with supervised classification into a unified pipeline, where the supervised prediction target is derived from maintenance-confirmed fault inspections rather than threshold crossings alone.

The autoencoder module effectively learned nominal operating behavior and provided a reliable first-stage anomaly filter (F1 = 0.693), while K-Means clustering ($k = 10$) structured the flagged

anomalies into interpretable sub-groups suitable for expert review. These unsupervised stages provide an independent validation pathway and a practical mechanism for expert-guided label refinement.

For the supervised benchmark, the prediction target was derived from the plant's monthly expert vibration spectra, with each maintenance-confirmed fault finding matched to the continuous DCS record to determine the temporal boundaries of the corresponding fault episode. Five classifiers were evaluated on a chronological test set separated from training by a 24-hour temporal gap. XGBoost achieved the best F1-score of 0.774 and ROC AUC of 0.937, with HGB close behind ($F1 = 0.770$). A cost-sensitive deployment analysis using real maintenance cost data from the industry partner showed that, at episode-level thresholds calibrated for 100% fault detection, all models catch every fault episode with fewer than two false alarms per month, projecting escalation-adjusted annual savings of approximately PHP 22.5M. Feature analysis confirmed that shifts in baseline vibration amplitudes, particularly mean and minimum values, are the most informative predictors of impending fault onsets.

The adoption of a binary classification strategy is a deliberate limitation. While grouping all fault modes into a single "Faulty" class reduces diagnostic granularity, it addresses the statistical infeasibility of multi-class classification under extreme data scarcity and avoids the label noise inherent in diagnosing specific fault types from scalar vibration features alone.

Several directions remain for future work: expert validation of pseudo-labeled clusters against physical failure modes, incorporation of trend-based and spectral features, threshold calibration for deployment-specific cost trade-offs, and evaluation across additional equipment types and operating regimes.

REFERENCES

- [1] Nagdev Amruthnath and Tarun Gupta. 2018. A research study on unsupervised machine learning algorithms for early fault detection in predictive maintenance. In *2018 5th International Conference on Industrial Engineering and Applications (ICIEA)* (2018-04). 355–361. doi:10.1109/IEA.2018.8387124
- [2] S. Arena, E. Florian, I. Zennaro, P. F. Orrù, and F. Sgarbossa. 2022. A novel decision support system for managing predictive maintenance strategies based on machine learning approaches. *Safety Science* 146 (February 2022), 105529. doi:10.1016/j.ssci.2021.105529
- [3] Ebru Turanoglu Bekar, Per Nyqvist, and Anders Skoogh. 2020. An intelligent approach for data pre-processing and analysis in predictive maintenance with an industrial case study. *Advances in Mechanical Engineering* 209, 5 (May 2020), 168781402091920. doi:10.1177/1687814020919207
- [4] Albert Buabeng, Anthony Simons, Nana Kena Frempong, and Yao Yeveyo Ziggah. 2021. A novel hybrid predictive maintenance model based on clustering, smote and multi-layer perceptron neural network optimised with grey wolf algorithm. 3, 5 (2021), 593. doi:10.1007/s42452-021-04598-1
- [5] Carlos Cernuda. 2019. On the Relevance of Preprocessing in Predictive Maintenance for Dynamic Systems. In *Predictive Maintenance in Dynamic Systems: Advanced Methods, Decision Support Tools and Real-World Applications*, Edwin Lughofer and Moamar Sayed-Mouchaweh (Eds.). Springer International Publishing, Cham, 53–93. doi:10.1007/978-3-030-05645-2_3
- [6] N. V. Chawla. 2003. C4.5 and imbalanced data sets: Investigating the effect of sampling method, probabilistic estimate, and decision tree structure. (2003).
- [7] Z. M. Çınar, A. A. Nuhu, Q. Zeeshan, O. Korhan, M. Asmael, and B. Safaei. 2020. Machine Learning in Predictive Maintenance towards Sustainable Smart Manufacturing in Industry 4.0. *Sustainability* 12, 19 (2020), 8211. doi:10.3390/su12198211
- [8] Maren David Dangut, Zakwan Skaf, and Ian K. Jennions. 2022. Handling imbalanced data for aircraft predictive maintenance using the BACHE algorithm. *Applied Soft Computing* 123 (July 2022), 108924. doi:10.1016/j.asoc.2022.108924
- [9] L. A. Demidova. 2020. Recurrent Neural Networks' Configurations in the Predictive Maintenance Problems. *IOP Conference Series: Materials Science and Engineering* 714, 1 (January 2020), 012005. doi:10.1088/1757-899X/714/1/012005
- [10] Wen Deng, Kim Tien Pham Nguyen, C. Gogu, K. Medjaher, and J. Morio. 2024. Enhancing prognostics for sparse labeled data using advanced contrastive self-supervised learning with downstream integration. *Engineering Applications of Artificial Intelligence* 138 (February 2024), 109268.
- [11] Chaoqun Duan, Chao Deng, and Bingnan Wang. 2018. Optimal maintenance policy incorporating system level and unit level for mechanical systems. 49, 5 (2018), 1074–1087. doi:10.1080/00207721.2018.1432782
- [12] Gergely Hajgat6, Richárd Weber, Botond Szilágyi, Balázs Tóthpál, Bálint Gyires-Tóth, and Csaba H6s. 2022. PredMaX: Predictive maintenance with explainable deep convolutional autoencoders. 54 (2022), 101778. doi:10.1016/j.aei.2022.101778
- [13] Ihtesham Ul Hassan, Karthik Panduru, and James Walsh. [n. d.]. Review of Data Processing Methods Used in Predictive Maintenance for Next Generation Heavy Machinery. *Data* 9, 5 ([n. d.]), 69.
- [14] Jyh-Yih Hsu, Yi-Fu Wang, Kuan-Cheng Lin, Mu-Yen Chen, and Jenneille Hwai-Yuan Hsu. 2020. Wind Turbine Fault Diagnosis and Predictive Maintenance Through Statistical Process Control and Machine Learning. 8 (2020), 23427–23439. doi:10.1109/ACCESS.2020.2968615 Conference Name: IEEE Access.
- [15] Pooja Vinayak Kamat, Rekha Sugandhi, and Satish Kumar. 2021. Deep learning-based anomaly-onset aware remaining useful life estimation of bearings. 7 (2021), e795. doi:10.7717/peerj-cs.795 Publisher: PeerJ Inc..
- [16] Sarah Keartland and Terence L. Van Zyl. 2020. Automating predictive maintenance using oil analysis and machine learning. In *2020 International SAUPEC/RobMech/PRASA Conference*. 1–6. doi:10.1109/SAUPEC/RobMech/PRASA48453.2020.9041003
- [17] Evgeny Latyshev. 2018. Sensor Data Preprocessing, Feature Engineering and Equipment Remaining Lifetime Forecasting for Predictive Maintenance. 226–231.
- [18] Yan Li and Lynne E. Parker. 2014. Nearest neighbor imputation using spatial-temporal correlations in wireless sensor networks. *Information Fusion* 15 (January 2014), 64–79.
- [19] MATLAB. 2024. *Data Preprocessing for Condition Monitoring and Predictive Maintenance*. <https://www.mathworks.com/help/predmaint/ug/data-preprocessing-for-condition-monitoring-and-predictive-maintenance.html>
- [20] Debasish Mishra, Krishna R. Pattipati, and George M. Bollas. 2024. Gaussian mixture model for tool condition monitoring. 131 (2024), 1001–1013. doi:10.1016/j.jmapro.2024.09.038
- [21] Mohsen Nikfar, Julia Bitencourt, and Konstantinos Mykoniatis. 2022. A Two-Phase Machine Learning Approach for Predictive Maintenance of Low Voltage Industrial Motors. *Procedia Computer Science* 200 (2022), 111–120.
- [22] E. M. Olariu, R. Tolas, R. Portase, M. Dinsoreanu, and R. Potolea. 2020. Modern approaches to preprocessing industrial data. In *2020 IEEE 16th International Conference on Intelligent Computer Communication and Processing (ICCP)*. 221–226. doi:10.1109/ICCP51029.2020.9266215
- [23] Daniel Pagano. 2023. A predictive maintenance model using Long Short-Term Memory Neural Networks and Bayesian inference. *Decision Analytics Journal* 6 (March 2023), 100174. doi:10.1016/j.dajour.2023.100174
- [24] Abraham Samuel, Niva Rana Mahanta, and Antoine Casel Vitug. 2022. Computational Technology and Artificial Intelligence (AI) Revolutionizing Interior Design Graphics and Modelling. In *2022 13th International Conference on Computing Communication and Networking Technologies (ICCCNT)* (2022-10). 1–6. doi:10.1109/ICCCNT54827.2022.9984232
- [25] E. Saranya and P. Bagavathi Sivakumar. 2020. Data-Driven Prognostics for Run-To-Failure Data Employing Machine Learning Models. In *2020 International Conference on Inventive Computation Technologies (ICICT)*. 528–533. doi:10.1109/ICICT48043.2020.9112411
- [26] Willamos Silva and Miriam Capretz. 2019. Assets Predictive Maintenance Using Convolutional Neural Networks. In *2019 20th IEEE/ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD)*. 59–66. doi:10.1109/SNPD.2019.8935752
- [27] Lei Sun, Tianyuan Liu, Yonghui Xie, Di Zhang, and Xinlei Xia. 2021. Real-time power prediction approach for turbine using deep learning techniques. *Energy* 233 (October 2021), 121130. doi:10.1016/j.energy.2021.121130
- [28] Shen Zhang, Fei Ye, Bingnan Wang, and Thomas G. Habetler. 2019. Semi-Supervised Learning of Bearing Anomaly Detection via Deep Variational Autoencoders. arXiv:1912.01096 [cs] doi:10.48550/arXiv.1912.01096

Dynamic Load Balancing and Parallel Non-Maximum Suppression for Viola-Jones Object Detection on Multi-Core and GPU Architectures

Mikhail Anton Chio
mbchio@up.edu.ph

Kelvin C. Buño
kcbuno@up.edu.ph

Kurt Nicholas P. Lanting
kplanting@up.edu.ph

Ren Tristan A. de la Cruz
radelacruz@up.edu.ph

Francis George C. Cabarle
fccabarle@up.edu.ph

Daryll Ko
dlko@up.edu.ph

Department of Computer Science, University of the Philippines Diliman, Quezon City, Philippines

ABSTRACT

Despite achieving 95.6 frames per second (FPS) on multi-GPU systems, existing parallel implementations of Viola-Jones face detection face two critical bottlenecks that prevent optimal deployment: (1) static workload distribution creates critical-path load imbalance across GPUs, resulting in 8% efficiency loss as devices processing small-scale windows become overloaded while those handling large scales remain idle, and (2) sequential CPU-based Non-Maximum Suppression consumes over 50% of total frame time, creating pipeline stalls where GPUs sit idle waiting for post-processing to complete. Based on analysis of representative GPU implementations [1], [2], we identify architectural mismatches between cascade classifier characteristics and existing parallelization strategies. We propose two complementary optimizations addressing these gaps: a lock-free dynamic work-stealing scheduler that enables adaptive load distribution across GPU devices, and a GPU-accelerated parallel NMS algorithm that eliminates the CPU bottleneck through warp-level parallelism. Our architecture maintains end-to-end execution on GPU hardware, avoiding expensive CPU-GPU data transfers while enabling balanced resource utilization. These techniques, grounded in established parallel computing patterns from Cederman and Tsigas and NVIDIA TensorRT, provide a pathway to achieving the theoretical performance potential of multi-GPU cascade-based detection systems.

KEYWORDS

Object detection, Viola-Jones, GPU acceleration, CUDA (Compute Unified Device Architecture), work-stealing scheduler, load balancing, non-maximum suppression, parallel computing, cascade classifiers, multi-GPU systems

1 Introduction

1.1 Background and Motivation

Object detection has undergone a remarkable evolution over the past decades, transitioning from handcrafted feature-based methods to sophisticated deep learning architectures. Despite this paradigm shift, the foundational work by Viola and Jones in 2001 [3], [4] remains a landmark contribution that introduced the first real-time object detection framework capable of processing images at 15 frames per second on conventional hardware. The Viola-Jones detector became one of the first successful real-time object detection systems, widely adopted in applications ranging from security

cameras to consumer electronics. While recent surveys [5], [6], [7] acknowledge limitations with pose variations and challenging illumination conditions, the method continues finding application in specialized domains where lightweight, real-time detection is prioritized over deep learning approaches.

The cascade paradigm introduced by Viola and Jones [3], [4] has proven remarkably enduring. Contemporary systems, including early-exit neural networks and edge AI applications, leverage similar principles of progressive refinement and efficiency. The core insight that most image regions can be quickly discarded as background while reserving intensive computation for promising candidates remains fundamental to achieving real-time performance constraints. Moreover, as deployment scenarios increasingly demand efficiency on resource-constrained devices, the Viola-Jones algorithm's balance between accuracy and low computational cost makes it particularly relevant for embedded systems, mobile platforms, and multi-object tracking scenarios where deep learning approaches may prove prohibitively expensive.

To contextualize this tradeoff, modern SOTA detectors such as YOLOv8 [8] and RetinaNet [9] achieve significantly higher mean Average Precision (mAP) on benchmarks such as WIDER FACE [10], often exceeding 90% AP, but require GPU inference pipelines and model sizes in the tens to hundreds of megabytes. By contrast, Viola-Jones operates at comparable or superior speeds on CPU-only hardware with a model footprint under 1MB, making it viable in microcontroller and FPGA deployments where neural network inference is impractical [11].

However, the transition to parallel computing architectures has exposed fundamental challenges. GPU implementations [2] can achieve performance competitive with FPGAs, yet the cascade classifier's early rejection mechanism and irregular computational patterns create severe load imbalance and thread divergence that undermine GPU efficiency, representing a critical barrier to fully exploiting parallel processing capabilities.

1.2 Problem Context

Despite substantial progress in GPU acceleration of Viola-Jones detection [1], [2], two fundamental challenges limit end-to-end throughput in production deployments. These challenges arise because the algorithm's computational characteristics do not match well with how modern GPUs execute parallel code, preventing systems from running as fast as they could.

The first challenge concerns load imbalance in multi-GPU implementations. Existing approaches statically distribute scale factors across available GPUs; however, computational cost varies dramatically across scales [1], [2]. Smaller scales generate exponentially more detection windows than larger scales, creating critical-path effects where some GPUs become overloaded while others remain idle. This results in approximately 8% efficiency loss even on optimized systems [1], with the problem worsening as more devices are added.

The second challenge involves sequential Non-Maximum Suppression bottlenecks in post-processing. The multi-scale sliding window approach produces hundreds to thousands of redundant overlapping detections per frame [1]. Traditional CPU-based NMS executes sequentially after GPU detection completes, consuming 50% or more of total frame time at high detection rates. This creates pipeline stalls where GPUs idle while waiting for post-processing, effectively negating much of the parallelization speedup.

These challenges, formalized in Section 3, motivate our proposed optimizations presented in Section 4.

1.3 Research Objectives

This research addresses the load imbalance and detection redundancy challenges inherent in parallel implementations of the Viola-Jones detection pipeline. Our work pursues three primary objectives:

Improve runtime efficiency of Viola-Jones detection. We achieve dynamic load balancing through work-stealing schedulers where idle GPUs dynamically pull scale factors from busy devices, eliminating the load imbalance caused by cascade early rejection [4]. This approach adapts to runtime characteristics, ensuring balanced utilization across GPU streaming multiprocessors.

Reduce redundant overlapping detections. We design parallel Non-Maximum Suppression (NMS) algorithms using IOU-based thresholding to eliminate redundant detections efficiently [4]. Our GPU-based approach maintains detection quality while achieving sub-millisecond execution through warp-level parallelism.

Evaluate CPU and GPU performance tradeoffs. We propose performance characterization comparing OpenMP [12] against CUDA implementations [2] to identify operating regimes where each architecture excels for cascade-based detection systems.

1.4 Contributions

This work makes three contributions: (1) analyzes performance gaps in parallel cascade implementations based on [1], [2] (see **Section 2** and **Section 3**), (2) proposes dynamic load balancing via work-stealing schedulers adapted from [13] for multi-GPU cascade detection (see **Section 4.1**), and (3) proposes GPU-accelerated parallel NMS maintaining detection quality while achieving sub-millisecond execution (see **Section 4.2**).

Section 2 and **Section 3** establish the technical foundations and formal bottlenecks motivating our approach, while **Section 4** presents the proposed optimizations and their theoretical analysis. **Section 5** then outlines the evaluation framework for validating these contributions in implementation.

2 Background and Related Work

2.1 Viola-Jones Object Detection Algorithm

The Viola-Jones face detection algorithm [3], [4] represents a landmark achievement in real-time object detection, introducing a framework that combines computational efficiency with high detection accuracy. The algorithm operates by scanning an input image with a sliding window at multiple scales, evaluating each window through a cascade of increasingly complex classifiers. This section details the core components, detection pipeline, and computational characteristics that define the algorithm's behavior and motivate our optimization strategies.

2.1.1 Core Components


The Viola-Jones algorithm comprises four innovations [4]. First, **Haar-like features** provide rectangular region-based representations: two-rectangle features compute pixel sum differences between adjacent regions, while three and four-rectangle features use more complex patterns. A 24×24 window admits over 180,000 possible features.

Second, the **integral image** enables constant-time feature evaluation. Given cumulative sums:

$$ii(x, y) = \sum_{x' \leq x, y' \leq y} i(x', y'),$$

any rectangular region sum requires exactly four array references regardless of size [4].

Once computed, any rectangular sum can be obtained using exactly four array references. Figure 1 demonstrates this calculation: for a rectangle defined by corners A, B, C, and D, the sum of pixel values within the rectangle is:



$$\sum_{(x,y) \in ABCD} i(x, y) = ii(D) + ii(A) - ii(B) - ii(C)$$

Figure 1: Integral image: sum of pixels in rectangle D can be computed using four array references (1, 2, 3, 4). [1]

This constant-time computation enables two-rectangle Haar features to be evaluated in six references, three-rectangle features in eight references, and four-rectangle features in nine references [4], regardless of the actual size of the rectangles involved.

Third, **AdaBoost learning** [3] selects critical features and combines them into strong classifiers. Weak classifiers $h(x, f, p, \theta)$ threshold individual features; AdaBoost iteratively weights misclassified examples to build effective ensembles.

Fourth, the **cascade architecture** arranges classifiers by increasing complexity [4]. Early stages rapidly reject non-face windows; later stages apply intensive tests to ambiguous candidates. A typical 38-stage cascade evaluates only 270 instructions per window

on average despite 80,000+ total features, as over 99% of windows fail early stages.




Figure 2: Cascade classifier architecture. Early stages reject non-face windows using simple features, while later stages apply more complex tests to remaining candidates. [2]

As shown in Figure 2, each stage serves as a progressively selective filter that reduces the number of windows forwarded to subsequent stages. This early-rejection strategy underpins the efficiency of the Viola-Jones detector, allowing high-throughput evaluation despite the large number of candidate windows.

2.1.2 Detection Pipeline and Formalization

The detection pipeline [4] computes integral images, applies multi-scale sliding windows (24×24 base, scales 1.1-1.25), evaluates candidates through the cascade, and consolidates overlapping detections via Non-Maximum Suppression. Algorithm 1 formalizes this process:

Algorithm 1: Viola-Jones Multi-Scale Detection

Input: Image I , cascade classifier $C = \{S_1, S_2, \dots, S_n\}$, scale factors Σ

Output: Set of detections D

1. $II \leftarrow \text{ComputeIntegralImage}(I)$
Complexity: $O(HW)$
2. $D \leftarrow \emptyset$
3. **for each** scale $s \in \Sigma$ **do**
 $|\Sigma|$ scales
4. $w \leftarrow \text{BaseWindowSize} \times s$
5. **for each** position (x, y) in image grid **do**
 $O(H \frac{W}{s^2})$ windows per scale
6. window $\leftarrow (x, y, w, w)$
7. **for each** stage $S_i \in C$ **do**
 Avg $O(\alpha |C|)$ where $\alpha \approx 0.1$
8. $f \leftarrow \text{EvaluateStageFeatures}(II, \text{window}, S_i)$
9. **if** $f < \theta_i$ **then**
10. **break**
 Early rejection
11. **if** all stages passed **then**
12. $D \leftarrow D \cup \{\text{window}\}$
13. $D_{\text{final}} \leftarrow \text{NonMaximumSuppression}(D)$
 Sequential: $O(|D|^2)$
14. **return** D_{final}

Complexity Analysis: The integral image computation (line 1) requires $O(HW)$ time for an image of height H and width W . Detection iterates over $|\Sigma|$ scales, with each scale generating $O(H \frac{W}{s^2})$ candidate windows. The cascade evaluation (lines 7-10) achieves average complexity $O(\alpha |C|)$ per window through early rejection, where $\alpha \approx 0.1$ represents the typical fraction of stages evaluated [4]. Sequential NMS (line 13) exhibits $O(|D|^2)$ complexity [1], where $|D|$ denotes the number of candidate detections. Total sequential complexity is $O(|\Sigma| HW \alpha |C| + |D|^2)$.

2.2 Related Work: GPU Architecture, Work Partitioning, Warp Scheduling, and Recent Literature

GPU implementations leverage massive parallelism through hierarchical thread organization. The NVIDIA CUDA programming model (Figure 3, Figure 4) organizes threads into warps of 32 threads executing in lockstep SIMD fashion, grouped into thread blocks that share local memory resources, distributed across independent streaming multiprocessors.




Figure 3: Simplified GPU architecture showing streaming multiprocessors, shared memory, and global memory hierarchy. [3]




Figure 4: CUDA thread organization showing the hierarchical structure of grids, thread blocks, and warps. [3]

Three key architectural constraints motivate our solutions: (1) warp-level SIMD execution requires coherent control flow to avoid divergence, (2) atomic Compare-And-Swap (CAS) operations enable lock-free synchronization across devices, and (3) independent streaming multiprocessors allow parallel processing of distinct tasks.

Early GPU implementations [2] employed static work partitioning, achieving only 50% GPU occupancy due to load imbalance from

cascade early rejection. To reduce memory traffic, implementations cache features in shared memory (Figure 5), loading cooperatively at each stage.




Figure 5: Threads cooperatively loading features from global memory to shared memory, enabling fast repeated access during cascade evaluation. [3]

Nguyen et al. [1] addressed intra-GPU load imbalance through dynamic warp scheduling (Figure 6, Figure 7). Warps dynamically fetch windows from a shared queue rather than receiving static assignments, achieving warp coherence (all 32 threads process the same window) while eliminating divergence. This achieved 3.1 $\times$ speedup on single GPUs and scaled to 95.6 FPS on five GPUs.




Figure 6: Dynamic window assignment: warps fetch detection windows from a shared work queue as they become idle. [4]




Figure 7: Warp-coherent execution: all threads within a warp evaluate the same window to preserve SIMD efficiency. [4]

However, **inter-GPU scale distribution** remains unsolved. Static assignment of complete scales to devices (GPU 1: scales 1-4, GPU 2: scales 5-8) creates severe imbalance because smaller scales generate 400 $\times$ more windows than larger scales, resulting in 8% efficiency loss [1] that worsens with additional GPUs.

Recent publications on Viola-Jones detection (2020-2026) demonstrate continued research interest but primarily focus on application-specific deployments or hybrid CNN-enhanced approaches rather than addressing the fundamental multi-GPU optimization

challenges identified above. Comprehensive surveys [5], [6], [7] document the transition from traditional cascade detectors to deep learning methods. Evaluation studies [11] highlight VJ's continued relevance due to its simplicity and real-time capability, though these works evaluate standard VJ on single-device hardware without addressing the scale imbalance or post-processing bottlenecks inherent in multi-GPU deployments.

Recent work on DNN parallelism (Mantri [14] on ResNet-50) and GPU-accelerated NMS (Symeonidis et al. [15] on neural detectors) targets deep learning architectures rather than cascade classifiers, not addressing VJ's multi-scale window redundancy or warp-level GPU parallelization.

Absence of Combined Optimizations. No publications in the past 10+ years have combined lock-free work-stealing from [13] with GPU-accelerated NMS for multi-GPU Viola-Jones detection. Nguyen et al.'s 2013 [1] work achieved 95.6 FPS but relied on sequential CPU-based NMS. Contemporary GPU NMS targets neural detectors [15], not cascades, while DNN studies [14] do not address inter-GPU load balancing for cascade architectures.

Our proposed integration is novel in three respects: (1) adapting lock-free dynamic work-stealing to address VJ's unique inter-GPU scale distribution imbalance, which differs fundamentally from the layer-parallel strategies employed in modern DNN training [14], (2) implementing warp-level GPU NMS to eliminate the CPU bottleneck absent from recent VJ evaluations [11], and (3) maintaining an end-to-end multi-GPU VJ pipeline that avoids DL overhead for applications targeting 95.6+ FPS where cascade classifiers remain competitive. This fills the gap between legacy GPU ports of VJ (2010-2013) and modern DL-centric parallel frameworks (2023-2025).

3 Problem Formalization

Based on analysis of parallel Viola-Jones implementations [1], [2], we identify two fundamental bottlenecks that prevent these systems from achieving theoretical performance potential. These limitations stem from architectural mismatches between cascade classifier characteristics and existing parallelization strategies.

3.1 Critical-Path Load Imbalance in Multi-GPU Systems

Multi-GPU implementations distribute scale factors across devices to exploit data parallelism. Existing approaches [1] employ static partitioning where each GPU receives a contiguous subset of scale factors. This creates severe imbalance because computational cost varies exponentially with scale. Smaller scales (e.g., $s_1 = 1.0$) generate exponentially more windows than larger scales (e.g., $s_{18} = 4.0$). Empirical analysis [1] shows that scales 1-4 account for over 90% of total computation, while scales 15-18 contribute less than 2%. This creates critical-path effects where GPUs assigned small scales become computational bottlenecks while GPUs processing large scales remain idle.

Figure 8 visualizes this imbalance across four GPUs processing 18 scales. GPU 1, assigned scales 1-4, handles approximately 90% of total work and determines overall completion time. Meanwhile, GPU 4, processing scales 14-18, completes its assignment in a small

fraction of GPU 1's execution time and remains idle. This results in approximately 8% efficiency loss [1], with degradation worsening as additional GPUs are added.




Figure 8: Static scale assignment creates critical-path effects: GPU 1 handles 90% of computation while GPU 4 remains idle for most of the detection phase.

While dynamic warp scheduling [1] successfully addressed intra-GPU load imbalance by redistributing windows among warps within a single device, the inter-GPU scale distribution problem remains unsolved. Static assignment of complete scales to devices prevents adaptive load balancing across the multi-GPU system.

3.2 Sequential NMS Post-Processing Bottleneck

The multi-scale sliding window approach generates hundreds to thousands of redundant overlapping detections per detected object. For a typical face detection scenario with 5-10 faces per frame, the cascade produces 500-2000 candidate bounding boxes before suppression [1]. Non-Maximum Suppression eliminates these redundancies by iteratively selecting the highest-confidence detection and removing all overlapping candidates above an Intersection-over-Union (IOU) threshold:

$$\text{IOU}(b_1, b_2) = \frac{\text{Area}(b_1 \cap b_2)}{\text{Area}(b_1 \cup b_2)}$$

Traditional implementations execute NMS sequentially on the CPU after GPU detection completes. This sequential algorithm exhibits $O(n^2)$ complexity in the number of candidate detections n , requiring 5-10 milliseconds for typical candidate sets. As Figure 9 illustrates, this post-processing time consumes 50% or more of total frame time even when GPU detection achieves 95.6 frames per second [1].




Figure 9: Sequential CPU-based NMS creates pipeline bottlenecks: post-processing consumes 50%+ of frame time despite GPU detection achieving 95.6 FPS. The GPU remains idle during CPU-based suppression.

This bottleneck creates pipeline stalls where GPUs complete detection but must wait for CPU-based NMS before processing the next frame. The synchronization overhead and data transfers between GPU and CPU further exacerbate the problem, negating much of the parallelization speedup achieved in the detection phase. Moreover, as GPU detection continues to improve through hard-

ware advances and algorithmic optimizations, the relative cost of sequential NMS increases, making it an increasingly critical limitation for end-to-end throughput.

4 Proposed Solutions

We propose two complementary optimizations: (1) dynamic load balancing via work-stealing schedulers [13], and (2) GPU-accelerated parallel NMS [16]. Figure 10 shows the enhanced pipeline integrating both components.




Figure 10: Enhanced pipeline: work-stealing detection + parallel NMS.

4.1 Dynamic Load Balancing via Work-Stealing

Our work-stealing approach solves the critical-path load imbalance problem (Section 3.1) by exploiting the GPU architectural features from Section 2.2: independent streaming multiprocessors across multiple devices, atomic CAS operations for lock-free synchronization, and warp-level SIMD execution. We adapt lock-free work-stealing from [13] for multi-GPU cascade detection, treating each GPU as an independent worker that dynamically acquires complete scale factors from a shared queue. This extends the intra-GPU dynamic warp scheduling from [1] to inter-GPU load balancing, eliminating the static assignment bottleneck where GPU 1 handles 90% of work (Figure 8) while other GPUs idle. GPUs dynamically pull scales from a global queue rather than receiving static assignments, eliminating critical-path effects. Figure 11 illustrates balanced utilization versus static assignment (Figure 8).




Figure 11: Work-stealing: GPUs pull scales from shared queue, achieving balanced utilization.

The lock-free queue uses Compare-And-Swap (CAS) for synchronization [13], with version counters to prevent the ABA problem (where CAS incorrectly succeeds when a pointer changes from value A to B back to A). Algorithm 2 formalizes the approach:

Algorithm 2: Dynamic Work-Stealing Multi-GPU Detection

Input: Image I , cascade C , scale factors $\Sigma = \{s_1, \dots, s_m\}$, GPUs $\{G_1, \dots, G_k\}$

Output: Set of detections D

1. $II \leftarrow \text{ComputeIntegralImage}(I)$
Complexity: $O(HW)$
2. Initialize global work queue $Q \leftarrow \Sigma$
 $O(m)$ scales
3. $D \leftarrow \emptyset$ with thread-safe access
4. **parallel for each** GPU G_i **do**
5. **while** $Q \neq \emptyset$ **do**
6. $s \leftarrow \text{Steal}(Q)$
Amortized $O(1)$ per CAS
7. **if** $s = \text{null}$ **then**
8. **break**
9. $D_s \leftarrow \text{DetectAtScale}(II, C, s)$
 $O(H \frac{W}{s^2} \cdot \alpha |C|)$ where $\alpha \approx 0.1$
10. $\text{AtomicUnion}(D, D_s)$
 $O(|D_s|)$
11. $D_{\text{final}} \leftarrow \text{NonMaximumSuppression}(D)$
Sequential: $O(|D|^2)$
12. **return** D_{final}

Complexity Analysis: With k GPUs processing m scales, work-stealing achieves parallel time $O(\frac{HW\alpha|C|}{k} + m \log m + |D|^2)$ where the detection workload $O(HW\alpha|C|)$ is distributed across k devices and lock-free queue operations incur amortized $O(m \log m)$ overhead [13]. The steal operation (line 6) uses Compare-And-Swap with amortized $O(1)$ cost per successful dequeue. Sequential NMS (line 11) remains the bottleneck at $O(|D|^2)$, unchanged from Algorithm 1.

Lines 5-10 implement dynamic pull-based distribution: idle GPUs steal scales without blocking, ensuring balanced completion times.

Consider $k = 2$ GPUs processing $m = 4$ scales with work units $w_1 = 1000, w_2 = 400, w_3 = 160, w_4 = 64$ (exponentially decreasing).

Static Assignment: GPU 1 processes scales 1-2 (total: 1400 units), GPU 2 processes scales 3-4 (total: 224 units, then idles). Parallel time = 1400 units, efficiency = $\frac{1624}{2 \times 1400} = 58.0\%$.

Work-Stealing (Algorithm 2): GPUs pull from queue $Q = \{1, 2, 3, 4\}$. At $t = 0$: GPU 1 steals scale 1 (1000 units), GPU 2 steals scale 2 (400 units). At $t = 400$: GPU 2 finishes, steals scale 3 (160 units). At $t = 560$: GPU 2 finishes, steals scale 4 (64 units), completes at $t = 624$. At $t = 1000$: GPU 1 finishes. Parallel time = 1000 units, efficiency = $\frac{1624}{2 \times 1000} = 81.2\%$, achieving $1.4\times$ speedup. Finer-grained task distribution approaches the 85-95% efficiency reported in [13].

4.2 GPU-Accelerated Parallel NMS

Our parallel NMS algorithm solves the sequential post-processing bottleneck (Section 3.2) by exploiting the GPU architectural features from Section 2.2. The solution assigns each detection to a

warp of 32 threads that cooperatively compute IOUs in parallel, leveraging SIMD execution to evaluate one candidate against 32 overlapping candidates simultaneously. The three-stage pipeline uses GPU global memory bandwidth and atomic operations for the suppression mask, distributing radix sort and stream compaction across independent streaming multiprocessors. Following [16] patterns, we move suppression entirely to GPU hardware. The three-stage pipeline comprises: (1) GPU radix sort by confidence, (2) warp-level parallel IOU computation with atomic suppression masks, and (3) stream compaction via parallel prefix sum.

Algorithm 3 formalizes this approach:

Algorithm 3: GPU-Accelerated Parallel Non-Maximum Suppression

Input: Detection set $D = \{d_1, \dots, d_n\}$ with confidences $\{c_1, \dots, c_n\}$, threshold τ

Output: Suppressed detection set D_{final}

1. $D' \leftarrow \text{ParallelRadixSort}(D, \text{by confidence descending})$
Complexity: $O(n \log n)$
2. Initialize suppression mask $M \leftarrow [0, 0, \dots, 0]$ of length n
Complexity: $O(n)$
3. **parallel for each** detection $d_i \in D'$ using warp-level parallelism **do**
4. **if** $M[i] = 1$ **then**
5. **continue**
6. **parallel for each** detection d_j where $j > i$ **do**
7. $\text{iou} \leftarrow \frac{\text{Area}(d_i \cap d_j)}{\text{Area}(d_i \cup d_j)}$
8. **if** $\text{iou} > \tau$ **then**
9. $\text{AtomicOr}(M[j], 1)$
10. $D_{\text{final}} \leftarrow \text{ParallelStreamCompact}(D', \text{where } M[i] = 0)$
Complexity: $O(n)$
11. **return** D_{final}

Complexity Analysis: Warp-level parallelization (lines 3-9) preserves SIMD efficiency while computing IOUs in parallel. Compared to the sequential $O(n^2)$ CPU-based NMS in line 13 of Algorithm 1 [1], Algorithm 3 achieves parallel time $O(n \log n + \frac{n}{w})$ where $w = 32$ is the warp size, with total work $O(n^2)$ distributed across GPU cores. The radix sort (line 1) dominates preprocessing at $O(n \log n)$, while warp-level IOU computation distributes $O(n^2)$ pairwise comparisons across streaming multiprocessors with atomic operations serializing writes to the suppression mask.

Consider $n = 4$ candidate detections: $D = \{d_1, d_2, d_3, d_4\}$ with confidence scores $S = \{0.95, 0.91, 0.88, 0.76\}$.

Sequential NMS: Process in descending order $\{d_1, d_2, d_3, d_4\}$. Iteration 1: Keep d_1 , compute 3 IOUs, suppress d_2 (high overlap). Iteration 2: Keep d_3 , compute 1 IOU, suppress none. Iteration 3: Keep d_4 . Total: $3 + 1 = 4$ IOU computations. CPU time: $4 \times 50\mu s = 200\mu s$.

Parallel NMS (Algorithm 3): Radix sort to $\{d_1, d_2, d_3, d_4\}$ ($\approx 0.8\mu s$). Launch $n = 4$ warps: Warp 1 computes d_1 vs 3 others in parallel ($\approx 0.3\mu s$), marks d_2 suppressed. Warp 2 skips (suppressed). Warp 3 computes d_3 vs d_4 ($\approx 0.3\mu s$). All warps execute simultaneously. Total GPU time: $0.8 + 0.3 + 0.3 = 1.4\mu s$.

Accounting for memory latency and atomic serialization: $\approx 4.2\mu s$, achieving $\frac{200}{4.2} \approx 48 \times$ speedup, demonstrating reduction from the 50%+ bottleneck in Section 3.

Our proposed optimizations provide provable performance improvements over existing approaches through complexity-theoretic analysis. These theoretical guarantees validate contributions (2) and (3) from Section 1.4 prior to implementation.

Load Balancing Improvement (Contribution 2): Dynamic work-stealing eliminates critical-path effects by ensuring all GPUs complete work within $O\left(\frac{HW\alpha|C|}{k}\right)$ parallel time, compared to $O(HW\alpha|C|)$ for the most overloaded GPU under static partitioning. This yields a theoretical speedup factor of k for perfectly balanced workloads. Empirical studies of work-stealing schedulers [13] demonstrate 85-95% parallel efficiency even with irregular workloads, suggesting our approach should achieve 8-15 $\times$ improvement over the 8% efficiency loss documented in [1] for static multi-GPU assignment.

NMS Acceleration (Contribution 3): Algorithm 3 reduces NMS from sequential $O(n^2)$ time to parallel $O\left(n \log n + \frac{n}{w}\right)$ time. For typical detection scenarios with $n = 2000$ candidates and $w = 32$ threads per warp, this represents a theoretical speedup of $\frac{2000^2}{2000 \log 2000 + \frac{2000}{32}} \approx 181 \times$ in computational work. Accounting for GPU memory latency and atomic synchronization overhead, conservative estimates based on [16] suggest achievable speedups of 50-100 $\times$ over CPU implementations, addressing the 50%+ bottleneck identified in Section 3.

4.3 Summary of Time, Space, and Computational Complexity

Algorithm	Time complexity	Space complexity	Parallelism
Algo 1: Viola-Jones (Sequential)	Integral image: $O(HW)$; Detection: $O(\Sigma \cdot HW \cdot \alpha C)$; NMS: $O(D ^2)$; $\alpha \approx 0.1$	$O(HW)$ integral image	Sequential CPU, no parallelism
Algo 2: Work-stealing Multi-GPU	Detection: $O(HW \cdot \alpha C \frac{1}{k})$; Queue ops: $O(m \log m)$; NMS: $O(D ^2)$; $k =$ GPUs, $m =$ scales	$O(HW + m)$ integral image + work queue	k GPUs in parallel; amortized $O(1)$ steal ops; lock-free CAS
Algo 3: Parallel NMS (GPU-accelerated)	Radix sort: $O(n \log n)$; Warp-level IOU: $O(\frac{n}{w})$; Stream compaction: $O(n)$; $w = 32$, total work $O(n^2)$ across GPU cores	$O(n)$ suppression mask + sorted array	Warp-level SIMD; atomic writes to suppression mask

Table 1: Summary of time, space, and parallel execution characteristics for Algorithms 1-3.

5 Proposed Evaluation Framework

The theoretical analysis in Section 4.3 establishes provable performance improvements for our proposed optimizations. This section outlines a comprehensive experimental validation framework for future implementation work. The metrics and benchmarks described below are recommendations for empirical verification of the theoretical guarantees, not results from current implementations.

5.1 Dataset and Baseline Configuration

We recommend using standard face detection benchmarks via OpenCV [17]: 1000 VGA images (640 $\times$ 480) with varying face densities, grayscale normalized to 8-bit. The frontal face cas-

cade (haarcascade_frontalface_default.xml, 38 stages, 6060 features) matches prior GPU implementations [1].

5.2 Hardware and Software Infrastructure

Experimental platforms: NVIDIA A100 (6912 cores, 40GB HBM2) and RTX 4050 (2560 cores, 6GB GDDR6), both supporting CUDA Compute Capability 8.0+. Host systems: Intel Xeon/AMD Ryzen (8-16 cores), 64GB RAM, PCIe 4.0. Implementations use C++/CUDA [18], OpenMP [19], and Thrust library for parallel primitives.

5.3 Implementation Variants

To enable systematic evaluation of our proposed optimizations' individual and combined impact, we outline four implementation variants of the Viola-Jones detection pipeline. Table 2 summarizes the proposed experimental configurations, which systematically enable work-stealing scheduling and parallel NMS across CPU and GPU platforms to isolate the contribution of each optimization strategy.

Variant	Python	OpenMP (CPU)	CUDA (GPU)
Baseline VJ	✓	✓	✓
VJ + Work-Stealing	N/A	✓	✓
VJ + Parallel NMS	N/A	✓	✓
VJ + Work-Stealing + NMS	N/A	✓	✓

Table 2: Implementation variants evaluated across Python, OpenMP, and CUDA platforms.

Baseline VJ variants add platform-specific parallelization without our optimizations. The OpenMP version uses static loop scheduling [19] to distribute detection windows across CPU cores, while the CUDA version employs the static thread-block grid organization described by Hefenbrock et al. [2]. Both retain CPU-based sequential NMS, establishing the pre-optimization parallel baseline.

VJ + Work-Stealing variants enable dynamic load balancing for the detection phase while retaining CPU NMS. On OpenMP, this uses dynamic loop scheduling with work-stealing [12] to redistribute windows among threads. On CUDA, it implements the lock-free work queue design presented in Section 4.1. These variants isolate the impact of load balancing on detection throughput.

VJ + Parallel NMS variants apply our GPU-accelerated or OpenMP-parallelized suppression algorithms to the baseline parallel detection implementations. This configuration maintains static work assignment but eliminates the post-processing bottleneck, allowing evaluation of NMS optimization in isolation.

VJ + Work-Stealing + NMS represents the complete optimization pipeline, combining dynamic load balancing with parallel suppression. This configuration demonstrates the cumulative benefit of both techniques and represents our proposed production implementation.

5.4 Performance Metrics

For experimental validation (Section 4.3), we recommend: (1) end-to-end FPS averaged over 1000 images, (2) phase-specific speedup

for detection and NMS, (3) parallel efficiency (speedup/processors), and (4) GPU occupancy via NVIDIA Visual Profiler [20] (target >70%). Latency breakdown across pipeline stages identifies remaining bottlenecks.

6 Conclusion

This paper addresses critical performance limitations in parallel Viola-Jones face detection through analysis of existing implementations and proposed optimization strategies. We identified two major bottlenecks: (1) critical-path load imbalance in multi-GPU systems causing 8% efficiency loss, and (2) sequential CPU-based NMS consuming 50%+ of frame time despite GPU detection achieving 95.6 FPS.

Our proposed solutions integrate lock-free work-stealing scheduling adapted from [13] for dynamic load balancing across GPUs, and GPU-accelerated parallel NMS following [16] patterns. The work-stealing approach enables adaptive scale distribution through pull-based task assignment, while three-stage parallel NMS (sort, warp-level suppression, compaction) provides theoretical 50-100× acceleration over sequential CPU implementations.

6.1 Future Work and Limitations

The proposed approach requires implementation and empirical validation across the evaluation framework outlined in Section 5. Key remaining tasks include: (1) implementing lock-free work queues and measuring overhead in multi-GPU configurations, (2) validating that parallel NMS preserves detection quality across diverse face densities, (3) characterizing performance crossover points between CPU and GPU platforms, and (4) conducting an ablation study that isolates each optimization's contribution using a synthesis table across variants (baseline, work-stealing only, parallel NMS only, and combined), reporting FPS, latency breakdown, GPU occupancy, and detection-quality impact for each variant.

Current limitations stem from the theoretical nature of this work. The work-stealing scheduler's effectiveness depends on runtime load characteristics that require profiling actual cascade executions. Similarly, parallel NMS scalability may be constrained by atomic synchronization overhead when processing densely-populated images (100+ candidate detections). The proposed evaluation targets VGA-resolution images; higher resolutions may expose memory bandwidth bottlenecks not addressed by our algorithms.

Additionally, our analysis focuses exclusively on Viola-Jones cascades. While more recent deep learning detectors have emerged, our analysis focuses on established GPU implementations of cascade classifiers to demonstrate fundamental parallelization principles. Generalization to modern deep learning detectors (e.g., YOLO [21] and Faster R-CNN [22]) requires adapting the load balancing strategy to handle variable computation patterns across network layers. The lock-free data structures assume cache-coherent shared memory, limiting applicability to distributed multi-node clusters without architectural modifications.

Future work should validate these proposals through implementation, perform the planned ablation analysis to quantify incremental gains from each algorithmic component, extend the work-stealing patterns to multi-stage deep learning pipelines, and investigate hybrid CPU-GPU scheduling for heterogeneous detector ensembles.

REFERENCES

- [1] T. Nguyen, D. Hefenbrock, J. Oberg, R. Kastner, and S. B. Baden, "A Software-Based Dynamic-Warp Scheduling Approach for Load-Balancing the Viola-Jones Face Detection Algorithm on GPUs," *Journal of Parallel and Distributed Computing*, vol. 73, no. 5, pp. 677–685, 2013, doi: 10.1016/j.jpdc.2013.01.012.
- [2] D. Hefenbrock, J. Oberg, N. Thanh, R. Kastner, and S. B. Baden, "Accelerating Viola-Jones Face Detection to FPGA-Level Using GPUs," in *2010 18th IEEE Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, 2010, pp. 11–18. doi: 10.1109/fccm.2010.12.
- [3] P. Viola and M. J. Jones, "Robust Real-Time Face Detection," *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2001, doi: 10.1023/b:visi.0000013087.49260.fb.
- [4] P. Viola and M. Jones, "Rapid Object Detection Using a Boosted Cascade of Simple Features," in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004, p. I–511. doi: 10.1109/cvpr.2001.990517.
- [5] Y. Xiao, Z. Tian, J. Yu, and others, "A review of object detection based on deep learning," *Multimedia Tools and Applications*, vol. 79, pp. 23729–23791, 2020, doi: 10.1007/s11042-020-08976-6.
- [6] Z.-Q. Zhao, P. Zheng, S.-t. Xu, and X. Wu, "Object Detection with Deep Learning: A Review." [Online]. Available: <https://arxiv.org/abs/1807.05511>
- [7] E. Arkin, N. Yadikar, X. Xu, and others, "A survey: object detection methods from CNN to transformer," *Multimedia Tools and Applications*, vol. 82, pp. 21353–21383, 2023, doi: 10.1007/s11042-022-13801-3.
- [8] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLOv8." [Online]. Available: <https://github.com/ultralytics/ultralytics>
- [9] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal Loss for Dense Object Detection," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2980–2988.
- [10] S. Yang, P. Luo, C. C. Loy, and X. Tang, "WIDER FACE: A Face Detection Benchmark," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 5525–5533.
- [11] D. M. Abdhussien and L. J. Saud, "An evaluation study of face detection by Viola-Jones algorithm," *International Journal of Health Sciences*, vol. 6, no. S8, pp. 4174–4182, 2022, doi: 10.53730/ijhs.v6nS8.13127.
- [12] B. Chapman, G. Jost, and R. van der Pas, *Using OpenMP: Portable Shared Memory Parallel Programming*. Cambridge, MA, USA: The MIT Press, 2007. doi: 10.1604/9780262033770.
- [13] D. Cederman and P. Tsigas, "Dynamic Load Balancing Using Work-Stealing," *GPU Computing Gems: Jade Edition*. Morgan Kaufmann, pp. 485–499, 2011.
- [14] V. P. Mantri, "Real Time Multiple Face Recognition using Intra-Layer, Inter-Layer and Hybrid Parallelism," *Journal of Information Systems Engineering & Management*, vol. 10, no. 17s, pp. 814–828, 2025, doi: 10.52783/jisem.v10i17s.2862.
- [15] C. Symeonidis, I. Mademlis, I. Pitas, and N. Nikolaidis, "Neural Attention-Driven Non-Maximum Suppression for Person Detection," *IEEE Transactions on Image Processing*, vol. 32, pp. 2454–2467, May 2023, doi: 10.1109/TIP.2023.3268561.
- [16] NVIDIA Corporation, "NVIDIA TensorRT Documentation." 2025. [Online]. Available: <https://docs.nvidia.com/deeplearning/tensorrt/latest/index.html>
- [17] OpenCV Community, "OpenCV (Open Source Computer Vision Library)." 2026. [Online]. Available: <https://opencv.org/>
- [18] NVIDIA Corporation, "CUDA C++ Programming Guide." 2025. [Online]. Available: https://docs.nvidia.com/cuda/pdf/CUDA_C_Programming_Guide.pdf
- [19] OpenMP Architecture Review Board, "OpenMP Application Programming Interface Specification, Version 6.0." 2024. [Online]. Available: <https://www.openmp.org/wp-content/uploads/OpenMP-API-Specification-6-0.pdf>
- [20] NVIDIA Corporation, "NVIDIA Visual Profiler." 2025. [Online]. Available: <https://developer.nvidia.com/nvidia-visual-profiler>
- [21] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788. doi: 10.1109/CVPR.2016.91.
- [22] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks." [Online]. Available: <https://arxiv.org/abs/1506.01497>

An Ensemble Model Using LightGBM and Bidirectional LSTM for Bitcoin Price Dip Prediction

Rameses Barria
Department of Computer,
Information Sciences and
Mathematics
University of San Carlos
Cebu City, Cebu, Philippines
ramesesluis@gmail.com

Marc Dylan M. Pilapil
Department of Computer,
Information Sciences and
Mathematics
University of San Carlos
Cebu City, Cebu, Philippines
dylanpilapil2@gmail.com

Christine F. Peña
Department of Computer,
Information Sciences and
Mathematics
University of San Carlos
Cebu City, Cebu, Philippines
cfpena@usc.edu.ph

ABSTRACT

Cryptocurrency markets exhibit extreme volatility, making the early detection of short-term price dips a challenging yet valuable task for risk management and trading decision support. This paper proposes an ensemble prediction framework that combines Light Gradient Boosting Machine (LightGBM) and Bidirectional Long Short-Term Memory (BiLSTM) networks to identify imminent Bitcoin price dips using real-time market data. A price dip is defined as a decline of at least 2% within a six-hour horizon. Historical BTC-USD and ETH-USD price data were collected and transformed into time-series features capturing price momentum, volatility, and cross-asset relationships. LightGBM is employed to model nonlinear feature interactions, while the BiLSTM network captures temporal dependencies in sequential price movements. The outputs of both models are integrated using a logistic regression-based ensemble to improve recall on dip events. Experimental results show that the proposed ensemble outperforms individual models in detecting rare dip occurrences, achieving improved recall and F_2 -score compared to standalone LightGBM and BiLSTM baselines. The findings suggest that hybrid ensemble approaches can enhance short-term dip detection performance and provide a foundation for early warning systems in highly volatile cryptocurrency markets.

KEYWORDS

Cryptocurrency prediction, Bitcoin price dips, Ensemble learning, LightGBM, Bidirectional LSTM.

1 Introduction

Cryptocurrency markets are characterized by high volatility, rapid price fluctuations, and frequent extreme movements that pose significant risks to traders and investors. Among

these movements, short-term price dips are particularly important, as they can trigger panic selling, margin calls, or missed trading opportunities. Early detection of such dips can support risk-aware decision making and enable the development of automated early warning systems. However, predicting short-term dips remains challenging due to the noisy, nonlinear, and non-stationary nature of cryptocurrency price data.

Traditional financial theories such as the Efficient Market Hypothesis suggest that asset prices fully reflect available information, limiting predictability. Nevertheless, empirical studies have shown that cryptocurrency markets often exhibit inefficiencies and exploitable patterns, especially at short time horizons. Recent advances in machine learning and deep learning have demonstrated promising results in modeling complex financial time-series data, including cryptocurrency price movements. Tree-based models such as Light Gradient Boosting Machine (LightGBM) are effective in capturing nonlinear feature interactions, while recurrent neural networks, particularly Long Short-Term Memory (LSTM) variants, excel at learning temporal dependencies in sequential data.

Despite these advances, single-model approaches often struggle to detect rare but critical events such as sudden price dips. This limitation motivates the use of ensemble learning, which combines complementary models to improve robustness and predictive performance. In this paper, we propose an ensemble framework that integrates LightGBM and Bidirectional LSTM models to detect short-term Bitcoin price dips using real-time market data. A dip event is defined as a price decline of at least 2% within a six-hour prediction window, with additional cross-asset information derived from Ethereum price movements.

The main contributions of this work are threefold: (1) the formulation of a short-horizon Bitcoin dip detection task using real-world market data, (2) the design of an ensemble model that leverages both feature-based and sequential learning approaches, and (3) an empirical evaluation demonstrating that the proposed ensemble improves recall and F_2 -score compared to individual baseline models. These results highlight the potential of hybrid ensemble methods for enhancing short-term risk detection in volatile cryptocurrency markets.

2 Methodology

This study adopts an experimental and implementation-based research approach to develop and evaluate an ensemble model for short-term Bitcoin price dip detection. The methodology consists of data collection, preprocessing, feature engineering, dip labeling, model training, ensemble integration, and performance evaluation. The overall goal is to assess whether combining gradient boosting and sequence-based deep learning models improves the detection of rare dip events in highly volatile cryptocurrency markets.

2.1 Problem Definition and Prediction Task

This study formulates short-horizon Bitcoin price dip detection as a supervised binary classification problem using time-indexed market observations. Let t denote an hourly timestamp and let P_t be the Bitcoin closing price at time t . For each timestamp, an input instance is constructed from a set of engineered market features derived from recent historical data, including momentum and volatility indicators, temporal encodings, and cross-asset information from Ethereum. The objective is to estimate the probability that a dip event will occur within a fixed prediction horizon.

A dip event is defined using a forward-looking rule over the next six hours. Specifically, the minimum closing price in the future window $[t + 1, t + 6]$ is compared against the current closing price P_t . An instance at time t is labeled as a dip event if the price decreases by at least 2% within the six-hour horizon; otherwise, it is labeled as a non-dip event. Formally, the binary label y_t is defined as:

$$y_t = \begin{cases} 1, & \text{if } \min(P_{t+1}, \dots, P_{t+6}) \leq 0.98P_t \\ 0, & \text{otherwise} \end{cases}$$

This labeling strategy aligns with the practical goal of early warning, where the model must infer near-future risk from current and recent market conditions. Because dip events

are relatively rare compared to non-dip intervals, the resulting dataset exhibits strong class imbalance. This imbalance is addressed during training to improve the model's sensitivity to dip events while maintaining reasonable precision under realistic market behavior.

2.2 Data Collection

Hourly historical market data were collected for Bitcoin (BTC-USD) and Ethereum (ETH-USD) from Yahoo Finance, covering the period from January 2018 to February 2025. The dataset includes open, high, low, close, and volume (OHLCV) values for both assets. Ethereum data were included to provide cross-asset contextual information that may influence Bitcoin price movements.

All timestamps were aligned to an hourly resolution and merged into a unified dataset indexed by time. Initial records affected by indicator warm-up periods were removed to ensure stability of computed features.

2.3 Data Preprocessing

The raw market data underwent a structured preprocessing pipeline to ensure consistency and modeling readiness. Missing values resulting from rolling computations were removed after the warm-up period. Duplicate timestamps and misaligned records were eliminated, and chronological ordering was strictly preserved to maintain the temporal integrity required for time-series modeling.

Outlier checks were performed to verify that price and volume values remained within realistic market ranges. The final dataset contained no missing values and exhibited a strictly increasing time index, ensuring suitability for both machine learning and deep learning models.

2.4 Feature Engineering

This section details the characteristics of the data used for training and evaluating the ensemble model, including its source, temporal scope, and the transformations applied to ensure data quality.

2.4.1 Data Source and Time Range.

The primary data used in this study consists of historical hourly market observations for Bitcoin (BTC-USD) and Ethereum (ETH-USD) programmatically sourced from the Yahoo Finance API. The dataset covers a seven-year period from January 1, 2018, to February 25, 2025. This timeframe was selected to encompass diverse market regimes, including high-volatility "crypto winter" phases and significant bull runs, providing a robust distribution of samples for machine learning. The final processed dataset contains 16,760 hourly records.

2.4.2 Preprocessing and Data Treatment.

To ensure the reliability of the predictive models, a multi-stage preprocessing pipeline was implemented. First, BTC and ETH time series were synchronized using a strictly increasing hourly DatetimeIndex. Small data gaps (≤ 2 hours) were handled via forward-filling, while larger gaps representing less than 0.02% of the total dataset were excluded to maintain the temporal sequence integrity required for the BiLSTM component. Outlier checks were performed to verify that price and volume values remained within realistic market bounds. Finally, for the deep learning path (BiLSTM), continuous features were scaled using Min-Max Normalization (0–1 range), while LightGBM utilized raw feature magnitudes.

2.4.3 Feature Engineering.

Feature engineering was applied to transform raw OHLCV data into informative representations. Features were derived exclusively from historical information available at each timestamp to prevent data leakage. The feature categories are summarized in Table 1.

Table 1: Summary of Engineered Feature Categories

Feature Category	Description
Momentum Indicators	Indicators like RSI-14 and MACD (12-26) used to characterize directional trend strength and potential reversals.
Volatility Indicators	Metrics such as Average True Range (ATR-14) and rolling standard deviations reflecting market instability.
Temporal Encodings	Time-based features (hour and day) transformed using sine and cosine encodings to capture cyclical trading patterns.
Cross-Asset Features	Technical indicators derived from Ethereum price movements providing contextual signals of market-wide sentiment.

2.5 Dip Event Labeling

The selection of a 2% decline over a six-hour horizon as the operational definition for a "dip" is grounded in both the

statistical properties of cryptocurrency volatility and the practical requirements of early-warning systems.

1. Statistical Significance and Volatility Thresholds

Cryptocurrency markets, specifically Bitcoin, exhibit a "fat-tailed" distribution of returns, where extreme price movements occur more frequently than in traditional equity markets. According to Katsiampa [32], Bitcoin's intraday volatility often necessitates higher thresholds to distinguish between "market noise" and meaningful downward trends. A 2% threshold effectively filters out standard hourly price fluctuations while capturing events that deviate significantly from the mean hourly return.

2. Practical Trading Horizons

A six-hour prediction window was selected to provide an actionable "early warning" lead time. In high-frequency environments, a window that is too short (e.g., 1 hour) may not allow institutional or retail traders enough time to adjust positions or manage collateral. Conversely, a window that is too long (e.g., 24 hours) tends to be influenced by macroeconomic shifts rather than the technical and sequential patterns targeted by the LightGBM and BiLSTM models [8], [14].

3. Alignment with Liquidations and Momentum Reversals

Empirical evidence from Smales [33] suggests that 2% price moves in Bitcoin often act as "psychological triggers" that lead to cascading liquidations in leveraged markets. By labeling these events specifically, the model focuses on the precursors to market stress. Furthermore, a 2% drop within 6 hours represents a sharp change in momentum, which is the primary temporal signature that the Bidirectional LSTM is designed to extract from the sequential data [30].

2.6 Handling Class Imbalance

The resulting dataset exhibited strong class imbalance, with dip events representing approximately 6.72% of all observations. To mitigate bias toward the majority class (non-dips), the SMOTETomek algorithm, a hybrid of the Synthetic Minority Oversampling Technique and Tomek Links was applied to the dataset.

Critically, this sampling procedure was applied strictly to the training fold after the chronological 80/20 split was performed. By isolating the oversampling process from the test set, the study ensures that no synthetic data influenced the final evaluation, effectively preventing look-ahead bias and maintaining a realistic distribution for the hold-out validation. This approach generates synthetic minority samples to improve model sensitivity while simultaneously removing ambiguous majority samples near class boundaries, producing a more balanced and learnable dataset without compromising the temporal integrity of the experiment.

2.7 Model Architecture

The proposed system utilizes a dual-pathway architecture, employing two distinct machine learning paradigms as base learners to capture different dimensions of market behavior.

First, the Light Gradient Boosting Machine (LightGBM) is employed to model nonlinear interactions among engineered tabular features. Its tree-based structure and histogram-based optimization make it exceptionally well-suited for high-dimensional financial data. The primary technical innovation of LightGBM lies in the use of Gradient-based One-Side Sampling (GOSS) and Exclusive Feature Bundling (EFB), which allow the model to handle large-scale datasets with extreme computational efficiency while maintaining high predictive accuracy for rare dip events [24].

Second, a Bidirectional Long Short-Term Memory (BiLSTM) network is used to model sequential price behavior and extract deep temporal signatures. Unlike standard unidirectional LSTMs, the BiLSTM architecture trains two separate hidden layers on the input sequence—one in a forward temporal direction and one in a backward direction. This design ensures that the hidden state at any given timestamp incorporates information from both past trends and the near-future "context" within the sliding window. This allows the system to capture temporal dependencies and volatility patterns that may not be fully represented in static, feature-based representations [8, 30].

2.8 Ensemble Integration

The outputs of the LightGBM and BiLSTM models were combined using a logistic regression meta-classifier in a framework often referred to as Stacked Generalization. In this architecture, each base model produces a continuous probability estimate representing the predicted likelihood of a dip event. These probabilities serve as inputs to the meta-ensemble layer, which learns an optimal weighting scheme to balance the contributions of each model. By leveraging the complementary strengths of feature-based gradient boosting and deep temporal sequence analysis, this strategy minimizes the individual biases of the base learners and improves the system's overall recall on rare dip events [5, 30].

2.9 Experimental Setup and Evaluation

The dataset was split chronologically, with the earliest 80 percent used for training and the most recent 20 percent reserved as a hold-out test set. This approach prevents temporal leakage and reflects realistic deployment conditions. Model performance was evaluated using recall, precision, F2-score with emphasis on recall, and the area under the precision–recall curve. A confusion matrix and

precision–recall curve were also analyzed to provide detailed insight into classification behavior under severe class imbalance. This methodology ensures a rigorous, reproducible evaluation of the proposed ensemble model and provides a structured framework for assessing short-term Bitcoin dip prediction performance.

3 Results

This section presents the experimental results obtained from evaluating the proposed ensemble model for short-term Bitcoin price dip detection. The results are reported using quantitative performance metrics, tabular summaries, visual analyses, and sample prediction behavior to assess the effectiveness of the model under realistic market conditions.

3.1 Dataset Characteristic and Class Distribution

After preprocessing, feature engineering, and dip labeling, the final dataset consisted of 16,760 hourly observations aligned across Bitcoin and Ethereum price series. Each observation contained engineered momentum indicators, volatility measures, temporal encodings, cross-asset features, and a binary label indicating whether a price dip of at least two percent occurred within the subsequent six-hour horizon.

Dip events were relatively rare, accounting for approximately 6.72 percent of all samples, while non-dip events comprised the remaining majority. This strong class imbalance reflects real cryptocurrency market behavior, where sharp short-term price declines occur F2-scorer less frequently than stable or moderately fluctuating price movements. The rarity of dip events significantly influenced both model training and evaluation outcomes.

3.2 Baseline Model Comparison

To contextualize the performance of the proposed ensemble framework, its predictive behavior was compared against standalone learning paradigms. Table 2 summarizes the performance across four key metrics: Precision, Recall, F2-score, and Area Under the Precision-Recall Curve (AUPRC).

Table 2: Performance Comparison of Baseline Models on the Hold-Out Dataset

Model	Precision	Recall	F2-score
Random Classifier	0.067	0.5	0.218
Standalone	0.162	0.131	0.136

LightGBM			
Standalone BiLSTM	0.145	0.185	0.175
Proposed Ensemble	0.43	0.84	0.705

3.2.1 Interpretation of Metrics

The quantitative results demonstrate a significant performance uplift when transitioning from individual learners to the ensemble framework. The **Recall of 0.840** for the proposed ensemble indicates that the system successfully identifies 84% of all true Bitcoin price dips within the six-hour horizon. This is a critical improvement over the standalone LightGBM (13.1%) and BiLSTM (18.5%), which failed to capture the majority of rare dip events.

The **F₂-score (0.690)** was utilized as the primary success metric, as it weights recall twice as heavily as precision. In the context of early-warning systems, the cost of a false negative (missing a dip) is substantially higher than a false positive (a false alarm). The ensemble's F₂-score suggests a robust balance that prioritizes safety and risk mitigation.

Furthermore, the **AUPRC of 0.540** provides a more reliable measure of effectiveness than standard accuracy due to the 6.72% class imbalance. While the random classifier's AUPRC is equivalent to the minority class frequency (0.067), the ensemble's score of 0.540 proves that the model has learned significant non-random structural patterns. The disparity between the standalone models and the ensemble suggests that while LightGBM captures tabular nonlinearities and BiLSTM extracts temporal dependencies, only their fused output via the logistic regression meta-layer achieves the sensitivity required for actionable dip detection [5], [30].

3.3 Model Evaluation on the Hold-Out Dataset

Model performance was evaluated using a 20 percent chronological hold-out dataset, representing the most recent and previously unseen market data. This evaluation strategy provides a realistic estimate of real-world predictive performance and prevents temporal leakage.

The ensemble model achieved performance above random classification but remained constrained by the difficulty of detecting rare events. Recall values indicate that only a small proportion of actual dip events were successfully identified. Precision values show that a limited fraction of predicted dips corresponded to true dip occurrences, reflecting the conservative prediction behavior of the model under severe class imbalance.

The recall-weighted F₂ score further confirms that, even when prioritizing dip detection, the system struggles to retrieve a substantial portion of rare dip events. The Area Under the Precision–Recall Curve (AU-PRC) exceeded the baseline dip frequency, demonstrating that the model learned non-random patterns from the data. However, the overall magnitude of the AU-PRC indicates that the predictive signal derived from price-based features remains limited.

3.4 Confusion Matrix Analysis

The confusion matrix summarizes the classification outcomes of the proposed ensemble model on the final hold-out dataset.

Table 3: Confusion Matrix of the Proposed Ensemble Model on the Final Hold-Out Dataset

Actual \ Predicted	Predicted No Dip (0)	Predicted Dip (1)
Actual No Dip (0)	3052	119
Actual Dip (1)	153	23

The confusion matrix shows that the model correctly classified the majority of non-dip periods while identifying only a small number of true dip events. False positives were relatively limited, indicating that the model does not generate excessive dip alerts. However, a substantial number of false negatives were observed, meaning that many actual dip events were missed.

This outcome highlights the core challenge of rare-event prediction in cryptocurrency markets. Increasing sensitivity to capture more dip events would likely result in a higher false-positive rate, while maintaining conservative thresholds limits recall. The observed classification pattern suggests that the model prioritizes stability over aggressive dip detection.

3.5 Error Analysis

An analysis of classification errors provides insight into the limitations of the proposed ensemble model under realistic market conditions. As reflected in the confusion matrix (Table 3), the model exhibits a substantial number of false negatives, corresponding to dip events that were not detected in advance. These missed events are primarily associated with abrupt price declines that occur without strong precursor signals in the engineered features, highlighting the inherent difficulty of short-horizon dip prediction in highly volatile markets where "flash crashes" can occur near-instantaneously.

False positives, while present, occur less frequently and typically correspond to periods of elevated market volatility


without subsequent sharp declines. This behavior suggests that the model adopts a conservative prediction strategy, prioritizing stability over aggressive dip signaling. Such conservatism is reflected in the precision–recall trade-off (Figure 1), where increases in recall are accompanied by a gradual rise in false-positive predictions as the decision threshold is lowered.

Further inspection of probability outputs indicates that predicted dip likelihoods tend to increase shortly before actual dip events, even when the decision threshold is not crossed. This pattern implies that the model effectively captures weak early warning signals such as subtle momentum shifts or BTC–ETH decoupling but struggles to amplify them into high-confidence predictions under conditions of severe class imbalance. These findings underscore the challenge of distinguishing between standard market noise and true financial anomalies using price-action data alone.

3.6 Precision-Recall Curve Interpretation

The precision–recall curve provides a detailed visualization of the trade-off between precision and recall across varying probability thresholds.

Figure 1: Precision-Recall Curve of the Proposed Ensemble Model




High precision is observed only at very low recall levels, indicating that when the model predicts a dip with high confidence, it is often correct, but such predictions occur infrequently. As recall increases, precision declines steadily, reflecting the growing difficulty of distinguishing true dip events from normal market noise.

This curve shape is characteristic of highly imbalanced financial datasets and confirms that dip detection remains challenging when relying primarily on historical price-based information.

3.7 Sample Prediction Behavior

To further examine model behavior, a sample prediction run was generated using the proposed ensemble model on the final hold-out dataset. For each timestamp, the model produced a continuous probability score representing the estimated likelihood that a dip would occur within the next six hours. These probability outputs were plotted against time to visualize the model's response to changing market conditions.

Figure 2: Probability-Over-Time Plot of Dip Predictions on the Hold-Out Dataset



The probability outputs remained low for most timestamps, with occasional increases during periods of heightened volatility or momentum reversal. Some probability spikes occurred shortly before actual dip events, suggesting that the model captures weak precursors such as rising volatility or negative momentum indicators. However, many dips occurred without strong probability signals, and some elevated probabilities did not correspond to subsequent dips.

This behavior is consistent with the confusion matrix results and highlights the difficulty of anticipating abrupt market movements using price-derived features alone.

3.8 Summary of Results

Overall, the experimental results indicate that the proposed ensemble model learns limited but meaningful patterns associated with short-term Bitcoin price dips. While performance exceeds random classification, the system remains constrained by strong class imbalance, market noise, and the absence of external information such as news or sentiment indicators.

The results demonstrate that combining LightGBM and Bidirectional LSTM models provides modest improvements over individual approaches. However, further enhancements in feature diversity, temporal modeling, and external signal integration are required to achieve robust early-warning performance suitable for real trading environments.

CONCLUSION

This study proposed an ensemble prediction framework combining Light Gradient Boosting Machine and Bidirectional Long Short-Term Memory models for short-term Bitcoin price dip detection. A dip event was

defined as a price decline of at least two percent within a six-hour horizon, reflecting a realistic and practically relevant early-warning scenario in highly volatile cryptocurrency markets.

Experimental results on a chronological hold-out dataset demonstrate that the proposed ensemble model is capable of learning non-random patterns associated with dip events, achieving performance above random classification. The integration of feature-based learning through LightGBM and temporal modeling through BiLSTM provided modest improvements over individual model approaches, particularly in terms of recall-oriented evaluation metrics. These findings suggest that hybrid ensemble strategies offer advantages when dealing with noisy and non-stationary financial time-series data.

However, the results also highlight the inherent difficulty of short-term dip prediction using price-based indicators alone. Severe class imbalance, rapid market regime shifts, and the absence of external information such as macroeconomic signals, order-flow dynamics, or sentiment data significantly limited predictive performance. The model exhibited conservative behavior, prioritizing stability and low false-positive rates at the expense of recall, resulting in a substantial number of missed dip events.

Despite these limitations, the proposed framework establishes a solid baseline for early-warning dip detection in cryptocurrency markets. The methodology demonstrates how ensemble learning and imbalance-aware evaluation can be applied systematically to rare-event financial prediction problems. Future work may extend this approach by incorporating richer data sources, advanced sequence models, and economic-value-based evaluation to improve practical applicability.

In conclusion, while robust early detection of short-term Bitcoin price dips remains a challenging task, this study provides empirical evidence that ensemble learning frameworks can contribute meaningful predictive signals. The findings offer insights for researchers and practitioners seeking to develop more reliable risk-aware decision-support systems in highly volatile digital asset markets.

References

- [1] Ahmed, L.J. et al. 2021. Biomedical Image Processing with Improved SPIHT Algorithm and optimized Curvelet Transform Technique. *2022 8th International Conference on Advanced Computing and Communication Systems (ICACCS)*. (Mar. 2021), 1596–1602.
DOI:<https://doi.org/10.1109/icaccs51430.2021.9441832>.
- [2] Aydin, Ü. and Aydin, Ö. 2020. Herd Behavior in Crypto Asset Market and Effect of Financial Information on Herd Behavior. *INTERNATIONAL JOURNAL OF ECONOMICS AND FINANCE STUDIES*. 12, 2 (2020), 2020–1309.
DOI:<https://doi.org/10.34109/ijefs.202012221>.
- [3] Bakshi, S. and Bhar, K.K. 2020. Simulation of Tidal Morpho-dynamics in the Hooghly Estuary using CMS Flow and Artificial Neural Network Models. *Procedia Computer Science*. 167, (2020), 459–467.
DOI:<https://doi.org/10.1016/j.procs.2020.03.255>.
- [4] Balachandra Pattanaik et al. 2021. Design of Smart Inverter for Distribution system using PV-STATCOM. *2021 International Conference on Intelligent Technologies (CONIT)*. (Jun. 2021), 1–5.
DOI:<https://doi.org/10.1109/conit51480.2021.9498527>.
- [5] Bouteska, A. et al. 2024. Cryptocurrency price forecasting – A comparative analysis of ensemble learning and deep learning methods. *International Review of Financial Analysis*. 92, (Mar. 2024), 103055.
DOI:<https://doi.org/10.1016/j.irfa.2023.103055>.
- [6] Chen, Z. et al. 2020. Bitcoin price prediction using machine learning: An approach to sample dimension engineering. *Journal of Computational and Applied Mathematics*. 365, (Feb. 2020), 112395.
DOI:<https://doi.org/10.1016/j.cam.2019.112395>.
- [7] Davis, D.D. and Korenok, O. 2022. Policy experiments and financial regulation: Using laboratory methods to evaluate responses to the 2007–2009 financial crisis. *Journal of Economic Surveys*. (Jun. 2022).
DOI:<https://doi.org/10.1111/joes.12517>.
- [8] Hamayel, M.J. and Owda, A.Y. 2021. A Novel Cryptocurrency Price Prediction Model Using GRU, LSTM and bi-LSTM Machine Learning Algorithms. *AI*. 2, 4 (Oct. 2021), 477–496.
DOI:<https://doi.org/10.3390/ai2040030>.
- [9] Ji, S. et al. 2019. A Comparative Study of Bitcoin Price Prediction Using Deep Learning. *Mathematics*. 7, 10 (Sep. 2019), 898.
DOI:<https://doi.org/10.3390/math7100898>.
- [10] Jiang, W. 2021. Applications of deep learning in stock market prediction: Recent progress. *Expert Systems with Applications*. 184, (Dec. 2021), 115537.
DOI:<https://doi.org/10.1016/j.eswa.2021.115537>.
- [11] Kim, M. and Lee, I. 2022. Human-guided auto-labeling for network traffic data: The GELM approach. *Neural Networks*. 152, (Aug. 2022), 510–526.
DOI:<https://doi.org/10.1016/j.neunet.2022.05.007>.
- [12] Lee, M.-C. 2024. Bitcoin Trend Prediction with Attention-Based Deep Learning Models and Technical

- Indicators. *Systems*. 12, 11 (Nov. 2024), 498. DOI:<https://doi.org/10.3390/systems12110498>.
- [13] Li, Y. et al. 2024. Hybrid-LLM-GNN: Integrating Large Language Models and Graph Neural Networks for Enhanced Materials Property Prediction. *Digital Discovery*. (Dec. 2024). DOI:<https://doi.org/10.1039/d4dd00199k>.
- [14] Mudassir, M. et al. 2020. Time-series forecasting of Bitcoin prices using high-dimensional features: a machine learning approach. *Neural Computing and Applications*. (Jul. 2020). DOI:<https://doi.org/10.1007/s00521-020-05129-6>.
- [15] Musikawan, P. et al. 2019. Parallelized Metaheuristic-Ensemble of Heterogeneous Feedforward Neural Networks for Regression Problems. *IEEE Access*. 7, (2019), 26909–26932. DOI:<https://doi.org/10.1109/access.2019.2900563>.
- [16] Nabipour, M. et al. 2020. Deep Learning for Stock Market Prediction. *Entropy*. 22, 8 (Jul. 2020), 840. DOI:<https://doi.org/10.3390/e22080840>.
- [17] O'Leary, D.E. 2020. ISAFM Paper of the Year for 2012. *Intelligent Systems in Accounting Finance & Management*. 27, 1 (Jan. 2020), 40–40. DOI:<https://doi.org/10.1002/isaf.1469>.
- [18] Omole, O. and Enke, D. 2024. Deep learning for Bitcoin price direction prediction: models and trading strategies empirically compared. *Financial Innovation*. 10, 1 (Aug. 2024). DOI:<https://doi.org/10.1186/s40854-024-00643-1>.
- [19] Patel, M.M. et al. 2020. A Deep Learning-based Cryptocurrency Price Prediction Scheme for Financial Institutions. *Journal of Information Security and Applications*. 55, (Dec. 2020), 102583. DOI:<https://doi.org/10.1016/j.jisa.2020.102583>.
- [20] Pintelas, E. et al. 2020. Investigating the Problem of Cryptocurrency Price Prediction: A Deep Learning Approach. *Artificial Intelligence Applications and Innovations*. 584, (May 2020), 99–110. DOI:https://doi.org/10.1007/978-3-030-49186-4_9.
- [21] Scharnowski, S. and Shi, Y. 2021. Intraday Herding and Attention Around the Clock. *SSRN Electronic Journal*. (2021). DOI:<https://doi.org/10.2139/ssrn.3936733>.
- [22] Tianyu Ray Li, Anup S. Chamrajnagar, Xander R. Fong, Nicholas R. Rizik, and Feng Fu (2018). Sentiment-Based Prediction of Alternative Cryptocurrency Price Fluctuations Using Gradient Boosting Tree Model: 2018. <https://arxiv.org/abs/1805.00558>.
- [23] Soumitra Singh Kachhwaha et al. 2020. Analysis of Smart Front-End Structure in Frontal Crash Mitigation. *Lecture notes in mechanical engineering*. (Jan. 2020), 631–638. DOI:https://doi.org/10.1007/978-981-15-2696-1_60.
- [24] Sun, X. et al. 2018. A novel cryptocurrency price trend forecasting model based on LightGBM. *Finance Research Letters*. (Dec. 2018). DOI:<https://doi.org/10.1016/j.frl.2018.12.032>.
- [25] Tepe, G. et al. 2021. FinTech Companies: A Bibliometric Analysis. *International Journal of Financial Studies*. 10, 1 (Dec. 2021), 2. DOI:<https://doi.org/10.3390/ijfs10010002>.
- [26] Urquhart, A. 2016. The inefficiency of Bitcoin. *Economics Letters*. 148, (Nov. 2016), 80–82. DOI:<https://doi.org/10.1016/j.econlet.2016.09.019>.
- [27] Wang, C. et al. 2022. Stock market index prediction using deep Transformer model. *Expert Systems with Applications*. 208, (Dec. 2022), 118128. DOI:<https://doi.org/10.1016/j.eswa.2022.118128>.
- [28] Wang, D. et al. 2022. Corporate finance risk prediction based on LightGBM. *Information Sciences*. 602, (Jul. 2022), 259–268. DOI:<https://doi.org/10.1016/j.ins.2022.04.058>.
- [29] Wolk, K. 2019. Advanced social media sentiment analysis for short-term cryptocurrency price prediction. *Expert Systems*. 37, 2 (Nov. 2019). DOI:<https://doi.org/10.1111/exsy.12493>.
- [30] Wu, J. et al. 2024. Review of deep learning models for crypto price prediction: implementation and evaluation. *arXiv (Cornell University)*. (May 2024). DOI:<https://doi.org/10.48550/arxiv.2405.11431>.
- [31] Yamamoto, Y. et al. 2021. Attack Detection and Defense System Using an Unknown Input Observer for Cooperative Adaptive Cruise Control Systems. *IEEE Access*. 9, (Jan. 2021), 148810–148820. DOI:<https://doi.org/10.1109/access.2021.3124547>.
- [32] Katsiampa, P. (2017). Volatility estimation for Bitcoin: A comparison of GARCH models. *Economics Letters*, 158, 3–6.
- [33] Smales, L. A. (2019). Bitcoin as a safe haven: Is it even worth considering? *Finance Research Letters*, 30, 385–393.

Adaptive Shot Allocation for Quantum Observable with Statistical Guarantees

Marc Jermaine Pontiveros
University of the Philippines
mpontiveros@up.edu.ph

Henry N. Adorna
University of the Philippines
hnadorna@up.edu.ph

ABSTRACT

Estimating the expectation value of a quantum observable requires running the circuit many times, and the right number of shots is rarely known in advance. Fixed budgets oversample easy observables and may fall short of the required confidence for harder ones. We derive adaptive stopping rules from Hoeffding and empirical Bernstein inequalities that halt sampling once the confidence radius drops below a user-specified tolerance ε , with a $(1 - \delta)$ coverage guarantee.

The analysis is organized around a three-level hierarchy separating the ideal observable value, the device-level expectation, and the sample-mean estimator. The stopping rules control the sampling component independently of hardware bias. Simulations show 3–25 $\times$ shot reductions over fixed- n Hoeffding baselines, with larger savings in low-variance settings. This framework is designed to integrate with measurement routines in hybrid quantum-classical pipelines easily.

CCS CONCEPTS

• **Theory of computation** $\rightarrow$ **Quantum computation theory**; • **Computing methodologies** $\rightarrow$ *Statistical learning theory*; Sequential decision making; • **Mathematics of computing** $\rightarrow$ *Probability and statistics*;

KEYWORDS

quantum computing, adaptive sampling, concentration inequalities, quantum measurement, shot efficiency

1 INTRODUCTION

Quantum computing has long promised to solve some computational problems that are intractable for classical computers. Given the recent advances in both quantum hardware and algorithm design, the field has progressed significantly over the past decade, and the first claims of quantum advantage are now emerging. In the Noisy Intermediate-Scale Quantum (NISQ)-era processors, the remaining dominant cost comes from the repeated execution of circuits required to estimate expectation values with sufficiently small sampling error. Minimizing quantum measurement resources is therefore central to any path toward quantum advantage [9, 10].

One of the classes of problems where quantum advantage is believed to manifest is estimating the expectation values of quantum observables [9]. To estimate the expectation value of a quantum observable, one repeatedly executes a circuit—taking many *shots*—and averages the outcomes. A standard approach chooses a fixed number of shots *a priori*; however, this can lead to inefficient oversampling when the observable is easy to estimate, or insufficient statistical confidence when the observable is hard to estimate [11, 13].

The difficulty is that many shot-saving heuristics are not statistically valid under adaptive stopping: naively monitoring an evolving confidence interval and halting when it becomes small enough generally invalidates fixed- n guarantees. Recent work in sequential analysis provides tools to construct *anytime-valid* confidence intervals—confidence sequences—that remain correct even when the stopping time is itself data-dependent [7, 8]. These techniques leverage concentration inequalities, including Hoeffding-type worst-case bounds [6] and variance-sensitive Empirical Bernstein bounds [2, 12], to enable adaptive sampling with rigorous coverage guarantees.

To frame our analysis, we distinguish between three specific components:

- (1) **Ideal observable value**: the expected value assuming a perfect, noiseless device model;
- (2) **Device-level expected value**: the expectation value actually induced by the noisy hardware;
- (3) **Sample-mean estimator**: the empirical average derived from a finite number of measurement shots.

Breaking down the problem in this way allows us to isolate the finite-sampling component. This component is distinct from hardware bias, as it can be quantified via concentration inequalities and controlled independently. While hardware errors are typically handled using error mitigation techniques, the methods we discuss here focus specifically on managing the errors arising from sampling. To this end, we provide stopping rules that guarantee

$$|\hat{\mu}_\varepsilon - \mu| \leq \varepsilon$$

with probability of at least $1 - \delta$. Here, μ represents the device-level mean, $\varepsilon > 0$ is the desired error tolerance, and $\delta \in (0, 1)$ is the acceptable failure probability. We elaborate on the (ε, δ) framework in Section 3.

The difficulty of estimating an observable depends heavily on its measurement variance. Consider Pauli observables with outcomes in $\{\pm 1\}$; their variance follows $\sigma^2 = 1 - \mu^2$. Consequently, the variance is small whenever $|\mu|$ is close to 1 (indicating high certainty) and peaks when $\mu \approx 0$. This variance profile is common in practice—for instance, during optimization, where some Hamiltonian terms may become nearly deterministic while others continue to exhibit high variance. Relying on a single, worst-case budget for shots is therefore often overly conservative.

Hoeffding’s inequality yields a confidence radius that is agnostic to variance:

$$r_n^{\text{Hoeff}} \asymp \sqrt{\frac{\log(1/\delta)}{n}},$$

resulting in a worst-case sample complexity of order

$$n \asymp \frac{\log(1/\delta)}{\varepsilon^2}$$

to achieve an accuracy of ϵ . In contrast, Empirical Bernstein bounds incorporate an estimate of the empirical variance, producing radii of the form

$$r_n^{\text{EBS}} \approx \sqrt{\frac{\hat{\sigma}_n^2 \log(1/\delta)}{n}} + \frac{\log(1/\delta)}{n}.$$

This difference leads to two analytically distinct regimes:

First is the **low-variance regime** ($\sigma^2 \ll 1$): Here, empirical Bernstein concentrates much faster than Hoeffding. It allows the process to stop substantially earlier, with the leading term scaling as

$$n \asymp \frac{\sigma^2 \log(1/\delta)}{\epsilon^2}$$

(ignoring lower-order additive terms).

Next is the **high-variance regime** ($\sigma^2 = \Theta(1)$): In this case, empirical Bernstein essentially recovers the behavior of Hoeffding, up to constant and logarithmic factors. While adaptivity does not hurt here, it may still offer advantages through sharper constants and data-dependent variance estimates.

Defining these regimes makes it explicit *when* variance-adaptive stopping offers provable benefits. It motivates a unified framework that treats Hoeffding as a conservative baseline while enabling variance-sensitive early stopping whenever the data justifies it.

This paper addresses fixed-shot inefficiency by deriving stopping rules from sequential analysis and concentration inequalities. Our primary contributions are:

- **Bounded-mean formalization:** We cast quantum observable measurement as estimation of a bounded random variable and derive explicit $(1 - \delta)$ finite-sample guarantees.
- **Adaptive stopping rules:** We derive stopping criteria from Hoeffding and empirical Bernstein inequalities that meet prescribed (ϵ, δ) targets.
- **Anytime-valid confidence:** We extend the rules via union bounds and confidence sequences so the (ϵ, δ) guarantee holds regardless of when sampling stops.
- **Regime analysis:** We characterize when variance-adaptive stopping provably beats worst-case baselines and derive the sample-complexity gap between regimes.
- **Three-level error hierarchy:** We separate ideal observable values, device-level expectations, and sample-mean error—a decomposition absent from prior EBS work [4]. This makes precise that the stopping rules control *sampling* uncertainty only, independent of hardware bias, a distinction practitioners routinely conflate.

As error mitigation [17] reduces hardware bias, sampling error becomes the dominant bottlenecks—making rigorous statistical control of shot counts essential.

2 RELATED WORK

Adaptive shot management and sample-efficient quantum measurement span several distinct research threads. We position our work relative to the most closely related.

Empirical Bernstein Stopping. Gresch et al. [4] apply the Empirical Bernstein stopping (EBS) rule to many-body energy estimation, combining operator grouping with variance-adaptive sampling. Our work is complementary. We add the three-level hierarchy

(ideal/device/sample) that Gresch et al. omit, give unified coverage proofs across all three stopping strategies, and release an open-source Qiskit library with EstimatorV2 integration.

Complementary Methods. Measurement grouping [3, 18] and shadow-based methods (e.g., LBCS [5], adaptive estimation [16]) partition operators or exploit randomized measurements to reduce circuit executions; our rules apply *within* each measurement type and are composable with all such grouping strategies. Bayesian adaptive estimators terminate when posterior uncertainty falls below a threshold but require a prior and non-trivial per-sample updates; our frequentist approach costs $O(1)$ per sample via Welford’s algorithm and requires no prior. The Rosalin optimizer [1] allocates shots across gradient components; our stopping rules provide a complementary per-observable *coverage* guarantee ($|\hat{\mu}_\tau - \mu| \leq \epsilon$ with probability $1 - \delta$). More sophisticated confidence sequences (e.g., stitched or mixture-based Bernstein CS [8]) offer tighter constants than the π^2 union-bound schedule but at the cost of implementation complexity; we defer this comparison to future work.

3 BACKGROUND AND PRELIMINARIES

We review bounded-mean estimation, the concentration inequalities we rely on, and the construction of anytime-valid confidence sequences.

3.1 Finite-Shot Quantum Measurement

In quantum computing, the expectation value of an observable O with respect to a quantum state ρ is given by

$$\langle O \rangle = \text{Tr}(\rho O).$$

Since quantum measurements are inherently probabilistic, estimating $\langle O \rangle$ requires sampling: executing the quantum circuit multiple times and aggregating measurement outcomes. For observables with binary outcomes (e.g., Pauli operators with eigenvalues ± 1), each shot yields a random variable $X_i \in [-1, 1]$, and the empirical mean after n shots is

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n X_i.$$

This estimator converges to the true expectation value $\mu = \mathbb{E}[X]$, but quantifying how close $\hat{\mu}_n$ is to μ requires probabilistic bounds.

3.2 Statistical Guarantees: Error Tolerance and Confidence

Before introducing concentration inequalities, we formally define the two parameters that characterize statistical estimation guarantees:

Error Tolerance (ϵ). The parameter $\epsilon > 0$ specifies the maximum acceptable deviation between the estimated mean $\hat{\mu}_n$ and the true mean μ . We seek to guarantee that

$$|\hat{\mu}_n - \mu| \leq \epsilon$$

with high probability. Smaller ϵ demands higher precision but requires more samples. For example, $\epsilon = 0.02$ means we want our estimate to be within ± 0.02 of the true value.

Failure Probability (δ). The parameter $\delta \in (0, 1)$ specifies the maximum acceptable probability that our guarantee fails. We aim to ensure

$$\Pr(|\hat{\mu}_n - \mu| > \epsilon) \leq \delta,$$

or equivalently, the guarantee holds with *confidence* at least $1 - \delta$. Smaller δ provides stronger confidence but requires more samples. For example, $\delta = 0.01$ corresponds to 99% confidence that our estimate lies within $\pm\epsilon$ of the true value.

The Estimation Problem. Given user-specified tolerance ϵ and confidence level $1 - \delta$, our goal is to determine the minimum number of shots n such that

$$\Pr(|\hat{\mu}_n - \mu| \leq \epsilon) \geq 1 - \delta.$$

Concentration inequalities provide formulas for computing such n , with different bounds exploiting different properties of the random variables (e.g., range, variance).

Practical Interpretation. In quantum algorithms, typical choices are $\epsilon \in [0.01, 0.05]$ (1–5% relative error for observables in $[-1, 1]$) and $\delta \in [0.001, 0.01]$ (99–99.9% confidence).

3.3 Concentration Inequalities

With the (ϵ, δ) framework established, we turn to concentration inequalities, which give explicit formulas for the required n . These inequalities provide probabilistic upper bounds on the deviation $|\hat{\mu}_n - \mu|$ for bounded random variables.

Hoeffding Bound. For independent $X_i \in [a, b]$, Hoeffding’s inequality gives:

$$\Pr(|\hat{\mu}_n - \mu| \geq \epsilon) \leq 2 \exp\left(-\frac{2n\epsilon^2}{(b-a)^2}\right).$$

This yields a confidence interval of half-width:

$$r_n^{\text{Hoeff}} = (b-a) \sqrt{\frac{\log(2/\delta)}{2n}} \quad \text{with probability } \geq 1 - \delta.$$

For $X_i \in [-1, 1]$ (i.e., $b - a = 2$), this specializes to $r_n^{\text{Hoeff}} = \sqrt{2 \log(2/\delta)/n}$.

Empirical Bernstein Bound. If the variance of X_i is small, tighter bounds are possible. The Empirical Bernstein bound [2, 12] uses the sample variance $\hat{\sigma}_n^2$ to adapt the confidence radius:

$$|\hat{\mu}_n - \mu| \leq \sqrt{\frac{2\hat{\sigma}_n^2 \log(3/\delta)}{n}} + \frac{3R \log(3/\delta)}{n},$$

where $R = b - a$ is the range of the bounded variable. This bound shrinks when variance is low, improving sample efficiency over Hoeffding in many practical settings.¹

¹The Maurer–Pontil [12] form uses $\frac{7R \log(2/\delta)}{3(n-1)}$ for the additive term. Both forms are valid; at the SWAP anchor ($n = 8,557$, $R = 2$), they differ by less than 0.3% of the radius. We adopt $\frac{3R \log(3/\delta)}{n}$ consistently, as it provides slightly tighter coverage at moderate n (the regime relevant to quantum workflows).

3.4 Anytime-Valid Confidence Sequences

Standard confidence intervals are fixed- n statements: they are valid only after a predetermined number of measurements. However, adaptive procedures require guarantees that remain valid under *sequential stopping*—that is, whenever the sampling process halts based on intermediate results.

Anytime-valid confidence sequences address this need. A sequence of intervals $C_1, C_2, \dots$ satisfies:

$$\Pr(\forall n \geq 1 : \mu \in C_n) \geq 1 - \delta.$$

This property enables algorithms to stop early without compromising statistical guarantees.

We adopt techniques from the literature on time-uniform concentration bounds [8] and use them to derive stopping criteria for quantum measurement processes.

Throughout this paper, $X_1, X_2, \dots$ denote i.i.d. samples with $X_i \in [-1, 1]$, representing outcomes of quantum measurements. The goal is to adaptively determine the minimum number τ of shots such that $|\hat{\mu}_\tau - \mu| \leq \epsilon$ with probability at least $1 - \delta$, comparing fixed- n , variance-agnostic, and variance-adaptive approaches.

4 METHODOLOGY: ADAPTIVE STOPPING RULES

Each stopping rule here guarantees at least $1 - \delta$ confidence on the estimation error $|\hat{\mu}_n - \mu|$. The idea is to halt once the confidence radius falls below ϵ . We derive rules from two classes of concentration bounds and discuss their properties in sequential settings.

4.1 Hoeffding-Based Stopping Rule

Starting with the classical Hoeffding bound, for observations $X_i \in [-1, 1]$ (range $b - a = 2$), we define a confidence radius at time n as

$$r_n^{\text{Hoeff}} = \sqrt{\frac{2 \log(2/\delta)}{n}}.$$

A simple stopping criterion is:

$$\tau^{\text{Hoeff}}(\epsilon, \delta) = \min \{n \geq 1 : r_n^{\text{Hoeff}} \leq \epsilon\}.$$

This rule is variance-agnostic and conservative, since it does not exploit low-variance scenarios where tighter bounds are possible.

4.2 Empirical Bernstein Stopping Rule (EBS)

Let $\hat{\sigma}_n^2$ be the biased empirical variance. The empirical Bernstein confidence radius is:

$$r_n^{\text{EBS}} = \sqrt{\frac{2\hat{\sigma}_n^2 \log(3/\delta)}{n}} + \frac{3R \log(3/\delta)}{n},$$

where $R = b - a$ is the range of the bounded random variable (consistent with Section 3). The EBS stopping rule is then:

$$\tau^{\text{EBS}}(\epsilon, \delta) = \min \{n \geq n_{\min} : r_n^{\text{EBS}} \leq \epsilon\},$$

where $n_{\min}$ is the minimum sample size. We set $n_{\min} = 10$ to ensure the empirical variance estimate is stable; fewer than ~ 10 samples yield unreliable variance estimates. This rule is variance-adaptive and often requires fewer samples when the underlying variance is small.

In practice, we implement EBS using *geometric checkpoints* $n_k = \lceil \beta^k \rceil$ rather than evaluating at every sample. At each checkpoint, we:

- (1) Sample a batch of $\Delta n = n_k - n_{k-1}$ shots
- (2) Update running statistics using Welford's online algorithm
- (3) Compute the EB radius r_k^{EBS} with per-checkpoint delta $\delta_k = \delta/K$
- (4) Stop if $r_k^{\text{EBS}} < \varepsilon$, otherwise continue to next checkpoint
- (5) Enforce a Hoeffding cap: never exceed $n_{\max}$ from worst-case bound

This geometric checkpoint approach reduces computational overhead to K evaluations instead of n , which suits batched EstimatorV2 execution well.

4.3 Anytime-Valid Stopping via Confidence Sequences

Both previous bounds assume a fixed n . For early stopping to remain valid, the intervals must hold simultaneously over time—a *confidence sequence*. Two constructions are worth distinguishing:

Geometric Checkpoints. Rather than monitoring at every sample n , we adopt a geometric checkpoint strategy: evaluating the stopping criterion at $n_k = \lceil \beta^k \rceil$ for $k = k_0, k_0 + 1, \dots, K$, where $\beta > 1$ (default 1.1). We allocate the failure probability uniformly across K checkpoints: $\delta_k = \delta/K$. This approach reduces computational overhead compared to per-sample monitoring while maintaining the union bound guarantee:

$$\Pr(\exists k : \text{violation at } n_k) \leq \sum_{k=1}^K \delta_k = \delta.$$

π^2 Union Bound (Anytime). For finer-grained control, we also implement per-sample monitoring using:

$$\delta_n = \frac{6\delta}{\pi^2 n^2}, \quad \text{so that } \sum_{n=1}^{\infty} \delta_n = \frac{6\delta}{\pi^2} \cdot \frac{\pi^2}{6} = \delta.$$

Replacing δ with δ_n in the confidence radius gives an anytime-valid bound. For the standard Hoeffding bound with $\log(2/\delta)$, substituting δ_n yields $\log(2/\delta_n) = \log(\pi^2 n^2 / (3\delta))$, giving:

$$r_n^{\text{any-H}} = \sqrt{\frac{2 \log(\pi^2 n^2 / (3\delta))}{n}}.$$

Similar modifications apply to EBS, yielding a stopping time:

$$\tau^{\text{any-EBS}}(\varepsilon, \delta) = \min \{n \geq n_{\min} : r_n^{\text{EBS}}(\delta_n) \leq \varepsilon\}.$$

This provides anytime-valid guarantees at every n , enabling stopping at arbitrary sample sizes. However, this approach is more conservative (requires more samples) than geometric checkpoints because the log terms $\log(3/\delta_n) = \log(\pi^2 n^2 / (2\delta))$ grow with n , whereas uniform allocation uses fixed $\log(3K/\delta)$ across all checkpoints.

4.4 Algorithmic Implementation

Our companion library [14] implements all three strategies. **Hoeffding** is variance-agnostic and always provides correct coverage, making it the safe default for worst-case settings. **EBS with geometric checkpoints** is variance-adaptive and most efficient for

batched EstimatorV2 workflows, evaluating the stopping criterion at K checkpoints. **Anytime EBS** (π^2 union bound) provides validity at every n , best for streaming workflows, but is more conservative than geometric checkpoints due to the growing log terms in δ_n .

Efficiency tradeoff. Geometric checkpoint EBS achieves 3–25× shot savings over Hoeffding depending on the variance regime ($\sigma^2 \lesssim 0.05$ for the high end; up to 25× at $\sigma^2 \approx 0$), while maintaining validity at K checkpoints. Anytime EBS provides validity at every n but uses more samples than geometric checkpoint due to the more conservative π^2 union-bound schedule. For most quantum workflows with batched execution, geometric checkpoint EBS is recommended.

Both stopping rules maintain a running mean $\hat{\mu}_n$ and variance $\hat{\sigma}_n^2$ via Welford's online algorithm [19], requiring $O(1)$ memory regardless of sample count. A Hoeffding cap $n_{\max} = \lceil (b-a)^2 \log(2/\delta) / (2\varepsilon^2) \rceil$ ensures worst-case guarantees.

Algorithm 1 Geometric Checkpoint EBS Stopping

```

1: Input:  $\varepsilon_{\text{stat}}, \delta, a, b, \beta > 1, n_{\min}$ 
2: Compute Hoeffding cap:  $n_{\max} \leftarrow \lceil (b-a)^2 \log(2/\delta) / (2\varepsilon^2) \rceil$ 
3: Generate checkpoints:  $\mathcal{N} \leftarrow \{n_k = \lceil \beta^k \rceil : n_{\min} \leq n_k \leq n_{\max}\}$ 
4: Compute  $K \leftarrow |\mathcal{N}|$  and allocate  $\delta_k \leftarrow \delta/K$  for all  $k$ 
5: Initialize stats  $\leftarrow \text{RunningStats}()$ 
6: for  $k = 1, 2, \dots, K$  do
7:   Sample  $\Delta n \leftarrow n_k - n_{k-1}$  new observations
8:   for each sample  $x$  in batch do
9:     stats.update( $x$ )
10:  end for
11:  Compute EB radius:  $r_k \leftarrow \sqrt{\frac{2\hat{\sigma}_{n_k}^2 \log(3/\delta_k)}{n_k}} + \frac{3R \log(3/\delta_k)}{n_k}$ 
    $\triangleright \hat{\sigma}_{n_k}^2$ : empirical variance (Welford's)
12:  if  $r_k < \varepsilon_{\text{stat}}$  then
13:    return  $(\hat{\mu}_{n_k}, n_k, r_k, \text{"ebs"})$ 
14:  end if
15: end for
16: return  $(\hat{\mu}_{n_{\max}}, n_{\max}, r_{n_{\max}}, \text{"cap"})$ 

```

5 THEORETICAL GUARANTEES

Throughout, we take $X_i \in [-1, 1]$ (bounded observations), $\mu = \mathbb{E}[X_i]$ (true mean), and $\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n X_i$ (empirical estimate). The theorems below establish high-probability bounds on estimation error and sample complexity for each stopping rule from Section 4.

5.1 Correctness Guarantees

For a stopping rule τ , we define correctness as:

$$\mathbb{P}(|\hat{\mu}_\tau - \mu| \leq r_\tau) \geq 1 - \delta.$$

THEOREM 5.1 (HOEFFDING STOPPING GUARANTEE). *Let $\tau^{\text{Hoeff}}(\varepsilon, \delta)$ be defined as in Section 4. Then,*

$$\mathbb{P}(|\hat{\mu}_{\tau^{\text{Hoeff}}} - \mu| \leq \varepsilon) \geq 1 - \delta.$$

THEOREM 5.2 (EMPIRICAL BERNSTEIN GUARANTEE (GEOMETRIC CHECKPOINT VARIANT)). *Let $\tau^{\text{EBS-geom}}(\varepsilon, \delta)$ be the stopping time of*

Algorithm 2 Anytime π^2 Union Bound EBS

```

1: Input:  $\varepsilon_{\text{stat}}, \delta, a, b, n_{\min}$ 
2: Compute Hoeffding cap:  $n_{\max} \leftarrow \lceil (b-a)^2 \log(2/\delta) / (2\varepsilon^2) \rceil$ 
3: Initialize stats  $\leftarrow$  RunningStats()
4: for  $n = 1, 2, \dots, n_{\max}$  do
5:   Sample one observation  $x$ 
6:   stats.update( $x$ )
7:   if  $n \geq n_{\min}$  then
8:     Compute per-sample budget:  $\delta_n \leftarrow 6\delta / (\pi^2 n^2)$ 
9:     Compute EB radius:  $r_n \leftarrow \sqrt{\frac{2\hat{\sigma}_n^2 \log(3/\delta_n)}{n}} + \frac{3R \log(3/\delta_n)}{n}$ 
       $\triangleright \hat{\sigma}_n^2$ : online variance (Welford's); distinct from  $\delta_n$ 
10:    if  $r_n < \varepsilon_{\text{stat}}$  then
11:      return  $(\hat{\mu}_n, n, r_n, \text{"ebs"})$ 
12:    end if
13:  end if
14: end for
15: return  $(\hat{\mu}_{n_{\max}}, n_{\max}, r_{n_{\max}}, \text{"cap"})$ 

```

Algorithm 1, which halts only at the K predetermined checkpoints $n_1 < \dots < n_K$ with per-checkpoint budget $\delta_k = \delta/K$. Then,

$$\mathbb{P}(|\hat{\mu}_{\tau^{\text{EBS-geom}}} - \mu| \leq \varepsilon) \geq 1 - \delta.$$

PROOF SKETCHES. **Theorem 5.1.** The radius $r_n^{\text{Hoeff}} = \sqrt{2 \log(2/\delta)/n}$ is a deterministic, decreasing function of n alone; it does not depend on the data. The stopping time $\tau^{\text{Hoeff}} = \lceil 2 \log(2/\delta)/\varepsilon^2 \rceil$ is therefore a fixed constant. Applying the Hoeffding bound at that fixed n :

$$\Pr(|\hat{\mu}_{\tau^{\text{Hoeff}}} - \mu| > \varepsilon) \leq 2 \exp\left(-\frac{2 \tau^{\text{Hoeff}} \varepsilon^2}{(b-a)^2}\right) \leq 2 \exp\left(-\log \frac{2}{\delta}\right) = \delta.$$

Theorem 5.2. Algorithm 1 evaluates the stopping criterion only at the K checkpoints $n_1 < \dots < n_K$ chosen *before* any data is observed, with per-checkpoint budget $\delta_k = \delta/K$. For each fixed n_k , the Empirical Bernstein bound [2] gives

$$\Pr(|\hat{\mu}_{n_k} - \mu| > r_k) \leq \delta_k.$$

A union bound over the K predetermined checkpoints yields

$$\Pr(\exists k : |\hat{\mu}_{n_k} - \mu| > r_k) \leq \sum_{k=1}^K \delta_k = \delta.$$

Hence, with probability $\geq 1 - \delta$, the bound holds at *every* checkpoint simultaneously. The algorithm stops at the first k for which $r_k < \varepsilon$; at that checkpoint $|\hat{\mu}_\tau - \mu| \leq r_\tau < \varepsilon$, completing the proof. Note that this argument is valid precisely because the checkpoints are predetermined: a naïve per-sample rule monitoring r_n^{EBS} with a fixed δ at every n would be invalidated by optional stopping and requires a time-uniform construction (Theorem 5.3).

Theorem 5.3. Setting $\delta_n = 6\delta/(\pi^2 n^2)$ and applying the Bernstein bound at each fixed $n \geq 2$:

$$\Pr(|\hat{\mu}_n - \mu| > r_n^{\text{EBS}}(\delta_n)) \leq \delta_n.$$

A union bound over all $n \geq 2$ gives

$$\Pr(\exists n \geq 2 : |\hat{\mu}_n - \mu| > r_n^{\text{EBS}}(\delta_n)) \leq \sum_{n=2}^{\infty} \delta_n \leq \frac{6\delta}{\pi^2} \sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{6\delta}{\pi^2} \cdot \frac{\pi^2}{6} = \delta,$$

where the last step uses the Basel identity $\sum_{n=1}^{\infty} 1/n^2 = \pi^2/6$. Hence $\mu \in C_n$ simultaneously for all $n \geq 2$ with probability $\geq 1 - \delta$. $\square$

5.2 Anytime Validity

Let C_n denote the confidence interval at time n . A confidence sequence satisfies:

$$\mathbb{P}(\forall n \geq 1, \mu \in C_n) \geq 1 - \delta.$$

THEOREM 5.3 (ANYTIME VALIDITY OF UNION-BOUND ADJUSTED EBS). Let $\delta_n = \frac{6\delta}{\pi^2 n^2}$. Then,

$$\mathbb{P}(\forall n \geq 2, |\hat{\mu}_n - \mu| \leq r_n^{\text{EBS}}(\delta_n)) \geq 1 - \delta.$$

5.3 Sample Complexity Bounds

We characterize the number of samples required to achieve accuracy ε under each rule.

PROPOSITION 5.4 (HOEFFDING SAMPLE COMPLEXITY). For $X_i \in [-1, 1]$ (range $b - a = 2$), the stopping time satisfies

$$\tau^{\text{Hoeff}}(\varepsilon, \delta) \leq \left\lceil \frac{2 \log(2/\delta)}{\varepsilon^2} \right\rceil.$$

PROOF. The rule stops at the smallest n for which $r_n^{\text{Hoeff}} \leq \varepsilon$. Substituting $r_n^{\text{Hoeff}} = \sqrt{2 \log(2/\delta)/n}$ and solving:

$$\sqrt{\frac{2 \log(2/\delta)}{n}} \leq \varepsilon \iff n \geq \frac{2 \log(2/\delta)}{\varepsilon^2}.$$

The stopping time is therefore $\tau^{\text{Hoeff}} = \lceil 2 \log(2/\delta)/\varepsilon^2 \rceil$, which equals the Hoeffding cap $n_{\max}$ enforced by the algorithm. For $\varepsilon = 0.02$ and $\delta = 0.01$ this gives $n_{\max} = \lceil 2 \log(200)/0.0004 \rceil = 26,492$ shots. $\square$

PROPOSITION 5.5 (EMPIRICAL BERNSTEIN SAMPLE COMPLEXITY). Let σ^2 be the true variance. Then asymptotically,

$$\mathbb{E}[\tau^{\text{EBS}}(\varepsilon, \delta)] = O\left(\frac{\sigma^2 \log(1/\delta)}{\varepsilon^2} + \frac{\log(1/\delta)}{\varepsilon}\right).$$

PROOF. By the strong law of large numbers, $\hat{\sigma}_n^2 \rightarrow \sigma^2$ almost surely, so for large n the EBS radius satisfies

$$r_n^{\text{EBS}} \approx \sqrt{\frac{2\sigma^2 \log(3/\delta)}{n}} + \frac{3R \log(3/\delta)}{n}.$$

Setting the variance term to $\varepsilon/2$ and solving gives $n \asymp 2\sigma^2 \log(3/\delta)/\varepsilon^2$; setting the additive term to $\varepsilon/2$ gives $n \asymp 6R \log(3/\delta)/\varepsilon$. The stopping time is dominated by the larger of the two, yielding $\mathbb{E}[\tau^{\text{EBS}}] = O(\sigma^2 \log(1/\delta)/\varepsilon^2 + \log(1/\delta)/\varepsilon)$. When $\sigma^2 \rightarrow 0$, the quadratic-in- $(1/\varepsilon)$ term vanishes and only the linear additive term survives, explaining the dramatic shot savings observed in the zero-variance Bell-state and near-optimal QAOA regimes. $\square$

This shows that EBS yields better sample efficiency when variance is low, aligning with practical needs in shallow circuits or expectation values close to extremal values ± 1 .

To make the regime differences concrete, Table 1 reports sufficient shot counts for representative (ε, δ) pairs under a Hoeffding baseline and Empirical Bernstein bounds in low-variance and high-variance regimes.

(ϵ, δ)	Hoeffding n_H	EBS (low var) $\sigma^2 = 0.01$	EBS (high var) $\sigma^2 = 1$
$(0.05, 10^{-2})$	4239	686	4565
$(0.02, 10^{-2})$	26492	1713	28520
$(0.05, 10^{-3})$	6081	962	6407

Table 1: Sample complexity comparison for bounded observations $X \in [-1, 1]$. Hoeffding uses $n_H = \lceil (b-a)^2 \log(2/\delta)/(2\epsilon^2) \rceil$ where $b-a=2$. Empirical Bernstein (EBS) uses the sufficient condition $n_{\text{EBS}} = \lceil \max \{ 2\sigma^2 \log(3/\delta)/\epsilon^2, 3R \log(3/\delta)/\epsilon \} + 1 \rceil$ where $R = b-a=2$. Values computed using the companion library [14].

Comparison and Insights. Hoeffding is simple and worst-case optimal; EBS exploits empirical variance and stops earlier when variance is low. Both support sequential use with appropriate union-bound adjustments. The guarantees establish that adaptive stopping rests on formal sample-complexity bounds, not heuristic reasoning.

6 INSTANTIATIONS AND SIMULATIONS

Section 5 predicts 3–25 $\times$ shot savings depending on variance. Here we check whether that holds on the observables practitioners actually use—Bell states, SWAP-test fidelity, and QAOA MaxCut—and connect the results to the companion library.

6.1 Simulation Setup

To validate the theoretical bounds, we implement both stopping rules compatible with the Qiskit ecosystem. All code, including the adaptive shot planning logic, is open-source and available as a companion library at [14].

We simulate Bell pair measurements ($Z \otimes Z$), single-qubit Pauli measurements (Z, X), parametric depolarizing noise at rates $p \in [0, 0.50]$, and realistic device noise via Qiskit’s FakeMontrealV2 backend [15] (real IBM Montreal calibration data: T_1/T_2 , gate errors, readout errors). For each setup we report stopping time τ , confidence interval width at stopping, and empirical coverage probability across independent trials.

6.2 Analytical Examples: Bell State and Noisy Circuit

Bell state ($\sigma^2 = 0$). Consider the Bell state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ measured under $Z \otimes Z$. The true expectation is $\mu = 1$ with zero variance (deterministic output). With $\widehat{\sigma}^2 = 0$ and $R = 2$, the EBS radius reduces to:

$$r_n^{\text{EBS}} = \frac{6 \log(3/\delta)}{n},$$

whereas the Hoeffding radius is $r_n^{\text{Hoeff}} = \sqrt{2 \log(2/\delta)/n}$. Since $6 \log(3/\delta)/n$ falls below ϵ much faster than $\sqrt{2 \log(2/\delta)/n}$, EBS achieves far greater shot efficiency in the zero-variance limit.

Noisy circuit ($\sigma^2 = 1 - \mu^2$). Under depolarizing noise with strength p , the same observable acquires mean $\mu = (1-p)^2$ and variance $\sigma^2 = 1 - \mu^2$. The EBS radius becomes:

$$r_n^{\text{EBS}} = \sqrt{\frac{2\widehat{\sigma}_n^2 \log(3/\delta)}{n}} + \frac{6 \log(3/\delta)}{n}.$$

As $p \rightarrow 0$, $\sigma^2 \rightarrow 0$ and EBS recovers near-zero-variance savings; as p increases EBS degrades gracefully toward the Hoeffding cap, as quantified in Figure 1.

6.3 SWAP Test Fidelity Estimation

To make the comparison concrete, consider estimating the fidelity between two single-qubit states $|\psi\rangle = R_y(\theta_1)|0\rangle$ and $|\phi\rangle = R_y(\theta_2)|0\rangle$ using a 3-qubit SWAP test.

Setup. Consider $\theta_1 = 0.3$ and $\theta_2 = 0.8$. The ideal overlap squared is:

$$F_{\text{ideal}} = |\langle\psi|\phi\rangle|^2 = \cos^2\left(\frac{\theta_1 - \theta_2}{2}\right) \approx 0.9388.$$

The SWAP test circuit produces measurements of the ancilla qubit in the computational basis, with outcomes $\{+1, -1\}$ and $\mathbb{E}[Z_{\text{anc}}] = F$. The variance for this measurement is:

$$\sigma^2 = 1 - F^2 \approx 1 - 0.9388^2 \approx 0.1187.$$

Hoeffding Planning. For parameters $\epsilon = 0.02$ and $\delta = 0.01$, Hoeffding’s formula gives:

$$n_H = \left\lceil \frac{(1 - (-1))^2 \log(2/0.01)}{2(0.02)^2} \right\rceil = \left\lceil \frac{4 \times 5.298}{0.0008} \right\rceil = 26,492 \text{ shots.}$$

Empirical Bernstein (Geometric Checkpoint). Using $\beta = 1.1$ and $n_{\min} = 10$, the algorithm generates $K = 83$ geometric checkpoints with uniform delta allocation $\delta_k = \delta/K \approx 0.00012$. With true variance $\sigma^2 \approx 0.1187$, at $n = 8,557$ (the actual stopping checkpoint):

$$\epsilon_{8557} \approx 0.0191 < 0.02.$$

Thus, geometric checkpoint EBS stops at $n = 8,557$ shots, achieving 3.1 $\times$ reduction compared to Hoeffding. These values are computed directly from the companion library [14]. The classical overhead is negligible: $K = 83$ checkpoint evaluations, each requiring $O(1)$ Welford updates (~ 3 arithmetic operations per sample). At a typical QPU rate of 100 shots/s, the 3.1 $\times$ shot reduction yields a proportional 3.1 $\times$ wall-clock saving, with < 1 ms total classical overhead—less than 0.01% of total job time.

Anytime π^2 EBS. For the anytime approach, the per-sample delta $\delta_n = 6\delta/(\pi^2 n^2)$ shrinks rapidly, causing $\log(3/\delta_n) \propto \log(n^2)$ to grow with n . At moderate variance ($\sigma^2 \approx 0.12$), this conservatism means the radius does not drop below $\epsilon = 0.02$ until the Hoeffding cap, so anytime EBS reaches $n = 26,492$. Geometric checkpoint EBS is preferred for batched QPU execution; anytime EBS is suited to streaming workflows where stopping at any checkpoint must remain valid.

Figure 1 shows stopping time across noise levels $p \in [0, 0.50]$ ($\epsilon = 0.02$, $\delta = 0.01$), with three distinct behaviors visible on a log scale. Hoeffding (fixed n) is flat at 26,492 shots regardless of noise. EBS-geom starts at $\sim 1,052$ shots at $p = 0$ ($25\times$ savings) and rises to plateau at the Hoeffding cap around $p \approx 0.20$ ($\sigma^2 \approx 0.59$), demonstrating graceful degradation. Anytime-EBS diverges above




Figure 1: Stopping time τ vs. depolarizing noise p for Hoeffding (fixed n), EBS-geom, and Anytime-EBS. $\epsilon = 0.02$, $\delta = 0.01$; shaded bands show $\pm 1\sigma$ over 100 trials.

Table 2: Mean stopping times, empirical coverage, and shot-reduction ratio for $\epsilon = 0.02$, $\delta = 0.01$ across noise levels. Results averaged over 100 trials.

p	σ^2	τ_{Hoeff}	τ_{EBS}	Coverage	Reduction
0.00	0.000	26492	1052	100%	25.18×
0.01	0.039	26492	3887	100%	6.82×
0.05	0.185	26492	11789	100%	2.25×
0.10	0.344	26492	20197	100%	1.31×

the Hoeffding cap beyond $p \approx 0.03$ due to its conservative π^2 union-bound schedule, confirming that EBS-geom is the recommended choice for batched QPU workflows. Table 2 reports mean τ , empirical coverage, and shot-reduction ratio for selected noise levels (100 trials each). EBS-geom achieves 100% empirical coverage across all noise levels, consistent with the $(1 - \delta)$ guarantee (at $\delta = 0.01$, the expected failure count is ≤ 1 per 100 trials).

6.4 Realistic Noise Validation: FakeMontrealV2

We validate under realistic noise using AerSimulator seeded from FakeMontrealV2 calibration data (27-qubit IBM Montreal device). A 50,000-shot pilot estimates $F_{\text{device}} = 0.9404$ (vs. $F_{\text{ideal}} = 0.9388$; bias = $0.0016 < \epsilon$). Both stopping rules target the device expectation within $\epsilon = 0.02$. Table 3 shows EBS-geom achieving 1.92× shot reduction while meeting the accuracy target; the small bias is fully within the ϵ_{stat} budget.

6.5 QAOA MaxCut: Near-Optimal Parameter Regime

We further evaluate on the quantum approximate optimization algorithm (QAOA) applied to 3-node MaxCut, using optimal circuit

Table 3: Results on FakeMontrealV2 (IBM Montreal calibration data, 27-qubit device). $F_{\text{ideal}} = 0.9388$, $F_{\text{device}} = 0.9404$ (50,000-shot pilot). $\epsilon = 0.02$, $\delta = 0.01$.

Method	Shots τ	$\hat{F}$	$ \hat{F} - F_{\text{dev}} $	$\leq \epsilon$?
Hoeffding (fixed n)	26492	0.9408	0.0005	✓
EBS-geom	13781	0.9390	0.0014	✓
Ideal (noiseless ref.)	—	0.9388	0.0016	—

Table 4: QAOA MaxCut energy estimation: $\gamma = 0.783$, $\beta = 0.438$, $\langle ZZ \rangle_{\text{ideal}} = 0.9836$, $\sigma^2 = 0.032$. $\epsilon = 0.02$, $\delta = 0.01$. Averaged over 200 trials.

Method	τ (mean)	Coverage (%)	Reduction
Hoeffding	26492	100.0	1.00×
EBS-geom	3481	100.0	7.61×
Anytime-EBS	6671	100.0	3.97×

parameters $\gamma = 0.783$, $\beta_{\text{QAOA}} = 0.438$ (distinct from the EBS geometric ratio $\beta_{\text{EBS}} = 1.1$). At this operating point, $\langle ZZ \rangle_{\text{ideal}} = 0.9836$ and $\sigma^2 = 0.032$ (low-variance, high-fidelity regime), making it ideal for adaptive stopping. Table 4 shows EBS-geom achieving 7.6× shot reduction over Hoeffding with 100% empirical coverage across 200 trials, consistent with the $(1 - \delta) = 0.99$ guarantee.

6.6 Limitations and Scope of Guarantees

The guarantees in Section 5 ensure statistical validity under adaptive sampling; EBS is especially effective when measurement variance is low, a common scenario for circuits near their ideal parameter values (e.g., VQE or QAOA settings). These guarantees apply exclusively to the *sampling error* component: the deviation between $\hat{\mu}_\tau$ and the device-level expectation μ_{device} . Specifically,

$$|\hat{\mu}_\tau - \mu_{\text{device}}| \leq \epsilon \quad \text{with probability } \geq 1 - \delta.$$

Achieving accuracy with respect to the *ideal* observable value μ_{ideal} additionally requires bounding the hardware bias $|\mu_{\text{device}} - \mu_{\text{ideal}}|$, which must be addressed through error mitigation [17]. Our framework is a *complement* to error mitigation: the two address additive components of the total error budget ($\epsilon_{\text{total}} = \epsilon_{\text{bias}} + \epsilon_{\text{stat}}$).

The guarantees further assume i.i.d. measurement outcomes. Real hardware may exhibit non-stationary noise (drift, crosstalk) that violates this [20]; the assumption holds approximately for short experimental runs on stable devices. Finally, FakeMontrealV2 uses real calibration data but does not capture all time-dependent phenomena; full validation on live IBM Quantum systems is an important next step.

7 CONCLUSION

We presented a framework for adaptive shot allocation in quantum measurement using Hoeffding and Empirical Bernstein stopping rules with provable $(1 - \delta)$ coverage guarantees. The central contribution is the three-level error hierarchy (ideal observable $\rightarrow$ device-level expectation $\rightarrow$ sample-mean estimator), which makes

precise that stopping rules govern only the sampling error component, independently of hardware bias. This decomposition is absent from prior EBS work [4] and clarifies a distinction practitioners routinely conflate. Classical overhead is $O(1)$ per checkpoint via Welford’s online algorithm [19], requiring $K = 83$ variance updates for the default $\beta_{\text{EBS}} = 1.1$ schedule—negligible relative to QPU execution cost.

Empirical results confirm the theoretical predictions across four operating points. QAOA near-optimal ($\sigma^2 = 0.032$) yields $7.6\times$ savings; SWAP-test fidelity ($\sigma^2 \approx 0.12$) yields $3.1\times$; near-zero variance ($p=0$) yields up to $25\times$. Realistic noise validation on FakeMontrealV2—using real IBM Montreal calibration data for T_1/T_2 times, gate errors, and readout errors—yields $1.92\times$ shot reduction, with empirical coverage consistent with the $(1-\delta)$ guarantee across all 10 noise levels (100 trials each). EBS-geom degrades gracefully to the Hoeffding cap as variance approaches the worst-case regime ($\sigma^2 \approx 0.59$ at $p \approx 0.20$), confirming that adaptivity never undermines correctness.

All results are reproducible via the companion Qiskit library [14], which exposes backend-agnostic EstimatorV2-compatible stopping rules as a drop-in complement to any existing circuit workflow. Future directions include: hardware benchmarking on live IBM Quantum systems to validate the simulation methodology under real device noise; joint coverage control over multiple observables for VQE Hamiltonians [11]; composition with measurement grouping [3, 18] to compound savings across circuit families; and drift-robust stopping for non-stationary devices [20] via stitched or mixture-Bernstein confidence sequences [8].

ACKNOWLEDGMENTS

This work was supported in part by the Qiskit Advocate Program and the Qiskit Advocate Mentorship Program (QAMP), which provided the mentorship structure and research environment that facilitated the development of the library described in this paper.

The author also acknowledges Tharmashastha SAPV for valuable discussions and comments that informed the formulation of the results and the direction of the library’s implementation.

REFERENCES

- [1] Andrew Arrasmith, Lukasz Cincio, Andrew T Sornborger, Wojciech H Zurek, and Patrick J Coles. 2020. Operator sampling for shot-frugal optimization in variational algorithms. *arXiv preprint arXiv:2004.06252* (2020).
- [2] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. 2009. Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. In *Proceedings of the 20th International Conference on Algorithmic Learning Theory*. 437–442.
- [3] Matthew X Burns, Chenxu Liu, Samuel Stein, Bo Peng, Karol Kowalski, and Ang Li. 2024. GALIC: Hybrid Multi-Qubitwise Pauli Grouping for Quantum Computing Measurement. *arXiv preprint arXiv:2409.00576* (2024).
- [4] Jonas Gresch, Marius Tepe, and Martin Kliesch. 2025. Adaptive and provably accurate estimation of quantum expectation values using the empirical Bernstein stopping rule. *arXiv preprint arXiv:2502.01730* (2025). arXiv:2502.01730v1.
- [5] Charles Hadfield, Sergey Bravyi, Rudy Raymond, and Antonio Mezzacapo. 2022. Measurements of quantum Hamiltonians with locally-biased classical shadows. *Communications in Mathematical Physics* 391, 3 (2022), 951–967. arXiv:2105.01110.
- [6] Wassily Hoeffding. 1963. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* 58, 301 (1963), 13–30.
- [7] Steven R. Howard, Aaditya Ramdas, Odalric-Ambrym Maillard, and Dino Sejdinovic. 2022. Sequential estimation of bounded means: Is the empirical Bernstein bound best possible? *arXiv preprint arXiv:2202.06861* (2022).
- [8] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet S Sekhon. 2021. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys* 18 (2021), 1–29.
- [9] IBM Quantum. 2023. Announcing the Quantum Development Roadmap through 2025. <https://www.ibm.com/quantum/blog/qdc-2025>. Accessed January 2026.
- [10] IBM Quantum. 2024. IBM Enters the Era of Quantum Utility. <https://www.ibm.com/quantum/blog/quantum-advantage-era>. Accessed January 2026.
- [11] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. 2017. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* 549, 7671 (2017), 242–246.
- [12] Andreas Maurer and Massimiliano Pontil. 2009. Empirical Bernstein bounds and sample variance penalization. *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)* (2009).
- [13] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. 2014. A variational eigenvalue solver on a quantum processor. *Nature communications* 5 (2014), 4213.
- [14] Marc Jermaine Pontiveros. 2026. *qamp-shotplanner: Statistical Shot Planning for Quantum Circuits*. <https://github.com/marcjermaine-pontiveros/QAMP-2025---Shot-Budgeting-for-Expectation-Values-of-Quantum-Circuits> Qiskit Advocate Mentorship Program (QAMP) 2025, Issue #53.
- [15] Qiskit Development Team. 2024. *Qiskit: An Open-Source Framework for Quantum Computing*. <https://qiskit.org/>
- [16] Ariel Shlosberg, Andrew Jena, Priyanka Mukhopadhyay, Tobias Haug, Alice Reardon-Smith, Connie Sheratt, Ryan Babbush, and Sergey Bravyi. 2023. Adaptive estimation of quantum observables. *Quantum* 7 (2023), 1160. arXiv:2110.15339.
- [17] Kristan Temme, Sergey Bravyi, and Jay M Gambetta. 2017. Error mitigation for short-depth quantum circuits. *Physical Review Letters* 119, 18 (2017), 180509.
- [18] Vladyslav Verteletskyi, Tzu-Ching Yen, and Artur F Izmaylov. 2020. Measurement optimization in the variational quantum eigensolver using a minimum clique cover. *The Journal of Chemical Physics* 152, 12, 124114.
- [19] BP Welford. 1962. Note on a method for calculating corrected sums of squares and products. *Technometrics* 4, 3 (1962), 419–420.
- [20] Leonardo Zambrano. 2026. Quantum tomography for non-iid sources. *arXiv preprint arXiv:2602.22057* (2026).

Neuro-Symbolic AI and Autoformalization: Restoring Determinism in the Age of Probabilistic Computing

Junel Alje B. Isanan
Computational Science Research Center
College of Science
University of the Philippines Diliman
Diliman, Quezon City, Philippines
jbisanan@up.edu.ph

ABSTRACT

The transition from classical deterministic computing to probabilistic artificial intelligence has exposed a critical *verification gap*, evidenced by the production of syntactically valid yet semantically incorrect code and proofs. This paper surveys the growing role of **Neuro-Symbolic AI** in addressing this gap through a specific interface: *autoformalization*, or the automated translation of informal natural language specifications into machine-verifiable formal logic. Through a systematic synthesis of LEANDJO, FORMALALIGN, ALPHAGEOMETRY, and BALDUR, we identify three recurring architectural mechanisms that characterize modern neuro-symbolic pipelines. First, *semantic projection* constrains ambiguous natural language inputs into well-formed formal representations. Second, *satisfiability gating* enforces logical correctness by subjecting candidate outputs to symbolic consistency checks. Third, *error-guided refinement* leverages diagnostic feedback from formal solvers to iteratively restrict the probabilistic search space of the neural component. Our analysis suggests that this emerging solver-aided paradigm offers meaningful labor amortization by automating substantial portions of the formalization process, while remaining explicitly bounded by fundamental computational constraints such as undecidability and NP-completeness. We observe that the neuro-symbolic architecture currently represents the most structurally coherent approach to high-assurance verification: it preserves the semantic flexibility required for natural language processing while enforcing the logical soundness guarantees that purely neural systems cannot provide.

CCS CONCEPTS

• **Theory of computation** → **Logic and verification**; • **Computing methodologies** → **Artificial intelligence**; *Neural networks*; • **Software and its engineering** → **Formal software verification**;

KEYWORDS

Theorem Proving, Formal Verification, Autoformalization, Neuro-Symbolic AI, Large Language Models

1 INTRODUCTION

Modern AI systems built on large language models (LLMs) have demonstrated remarkable capacity for generating syntactically plausible outputs across a wide range of tasks. However, this fluency masks a critical *verification gap*: the inability to determine whether a generated artifact — a proof, a formal specification, a piece of safety-critical code — is semantically correct rather than merely

plausible [5, 24]. This gap is not incidental. It is structural, arising from the fundamentally probabilistic nature of neural token prediction, which optimizes for linguistic likelihood rather than logical soundness [5].

Formal verification offers a principled solution: symbolic systems such as interactive proof assistants and SMT solvers provide deterministic correctness guarantees by construction. Yet purely symbolic approaches face their own limitations. They require formal specifications as input, demand years of specialized expertise to operate, and cannot process the ambiguous, context-dependent nature of natural language specifications [28, 31]. The result is a structural incompatibility: the systems best suited to *generating* formal content cannot *verify* it, and the systems capable of *verifying* it cannot easily *consume* informal human intent.

Autoformalization — the automated translation of natural language specifications into machine-verifiable formal representations — has emerged as the critical interface between these two paradigms [36]. Recent years have seen a proliferation of neuro-symbolic frameworks that attempt this translation, including LEANDJO [37], FORMALALIGN [27], ALPHAGEOMETRY [32], and BALDUR [12]. Despite their differences in domain, scale, and design, these systems share a common architectural commitment: constraining probabilistic neural generation through deterministic symbolic verification.

However, the existing literature lacks a systematic mechanistic account of *how* this constraint is operationalized across systems. Prior surveys of neuro-symbolic AI tend toward broad taxonomies of integration strategies [13, 20] without identifying the specific architectural mechanisms that make formal verification tractable in practice. This paper addresses that gap directly.

1.1 Contributions

This paper’s primary contribution is a **mechanistic taxonomy** of neuro-symbolic autoformalization: the identification of three recurring architectural mechanisms — semantic projection, satisfiability gating, and error-guided refinement — that appear consistently across systems with otherwise heterogeneous tasks and designs. In support of this taxonomy, the paper additionally provides:

- (1) A **structured framework comparison** classifying four contemporary systems — LEANDJO, FORMALALIGN, ALPHAGEOMETRY, and BALDUR — according to task type and mechanism instantiation rather than performance ranking.

- (2) A **computational constraint analysis** situating neuro-symbolic verification within the theoretical limits of undecidability and NP-completeness, clarifying what these systems can and cannot guarantee.
- (3) A **motivated theoretical argument**, grounded in the surveyed evidence, that neuro-symbolic integration represents the most structurally coherent currently available architecture for high-assurance autoformalization.

1.2 The Neuro-Symbolic Paradigm

The tension between neural and symbolic computation can be understood through Kahneman’s Dual-Process Theory [19]. Neural systems (System 1) offer high semantic flexibility and broad linguistic coverage but lack formal correctness guarantees. Symbolic systems (System 2) provide mathematical rigor and deterministic soundness but are constrained by the *semantic bottleneck* — an architectural inability to resolve the ambiguity inherent in natural language. Table 1 summarizes these complementary limitations.

Table 1: Comparison of Neural and Symbolic Computation

Property	Symbolic	Neural
Computation Model	Discrete	Continuous
Interpretability	Explicit	Opaque
Correctness	Formally verifiable	No guarantees
Ambiguity Handling	Brittle	Robust
Scalability	Expert-dependent	Data-driven

Neuro-symbolic integration addresses both limitations simultaneously by coupling System 1 generation with System 2 verification [9, 13]. The cost of this integration is computational: reconciling neural inference with symbolic reasoning frequently involves NP-hard satisfiability problems [34]. The central question this paper investigates is therefore not whether neuro-symbolic integration is desirable — the theoretical motivation is clear — but *what specific architectural mechanisms make it tractable and effective for autoformalization in practice*.

2 METHODOLOGY

2.1 Autoformalization: Theoretical Framework

The process of methodically and automatically converting informal mathematical or computational claims expressed in natural language into accurate, machine-verifiable formal language representations such as Lean, Coq, Isabelle/HOL, or F* is known as autoformalization [36]. This translation process constitutes the critical interface between automated machine verification systems and human specification intent.

Autoformalization can be defined formally as a mapping:

$$f : \mathcal{L}_{\text{natural}} \times C \rightarrow \mathcal{L}_{\text{formal}} \quad (1)$$

where C denotes contextual knowledge structures (such as type hierarchies, proof corpora, and mathematical libraries), $\mathcal{L}_{\text{natural}}$ denotes natural language proposition spaces, and $\mathcal{L}_{\text{formal}}$ denotes formal specification languages [37, 39].

2.2 Review Methodology

2.2.1 Search Scope and Strategy. The primary literature search was conducted across Google Scholar, the ACM Digital Library, arXiv (cs.AI, cs.LO, cs.PL), and Semantic Scholar, using the following query terms:

- *autoformalization, formal verification, neuro-symbolic AI*
- *LLM theorem proving, semantic parsing formal language*
- *interactive proof assistant, neural symbolic integration*

Search results were supplemented using Gemini Deep Research as an exploratory tool for identifying additional candidate works, followed by manual verification of all identified sources against their original publications.

2.2.2 Inclusion Criteria. A framework was included in the comparative analysis if it satisfied all of the following conditions:

- (1) It combines a neural generation component with a symbolic verification component in an integrated pipeline
- (2) Its symbolic component produces or verifies outputs in a formal language amenable to automated reasoning (e.g., Lean 4, Isabelle/HOL, Coq, SMT-LIB)
- (3) It targets the direction of *neural-to-symbolic* translation — that is, neural generation constrained by symbolic verification — rather than the converse
- (4) It reports empirical evaluation on at least one established benchmark

2.2.3 Exclusion Criteria. Frameworks were excluded if they met any of the following conditions:

- (1) Systems that produce formal specifications expressible entirely in natural language without requiring logical axioms, type-theoretic constructs, or machine-verifiable proof terms
- (2) Systems where symbolic components serve only as parsers or knowledge retrieval mechanisms rather than as verification oracles
- (3) Neuro-symbolic systems that follow the *symbolic-to-neural* direction, where symbolic representations are used to guide or constrain neural learning, as this falls outside the autoformalization scope defined in Section 2

2.2.4 Selection Outcome. The search and screening process identified twelve candidate frameworks satisfying the neural-to-symbolic integration criterion. Of these, four were selected for detailed comparative analysis on the basis of meeting all inclusion criteria and representing meaningfully distinct task types: tactic-level proof search (LEANDojo), statement alignment evaluation (FORMALALIGN), domain-constrained geometry solving (ALPHAGEOMETRY), and whole-proof generation with repair (BALDUR). All four have publicly available benchmark evaluations against established formal reasoning datasets.

2.2.5 Source Verification. All cited claims were independently verified against their primary sources. Reported metrics, architectural descriptions, and benchmark results were cross-checked directly against the original publications rather than secondary summaries.

2.3 The Two-Stage Pipeline Architecture

Modern neuro-symbolic autoformalization frameworks implement a standard two-stage processing pipeline:

Stage 1: Neural Semantic Parsing (System 1)

- **Input:** Natural language propositions with contextual information
- **Procedure:** An LLM generates candidate formal translations via probabilistic inference
- **Output:** Syntactically plausible but potentially semantically incorrect formal code
- **Mechanisms:** Statistical pattern matching, probabilistic token generation, and premise retrieval from knowledge databases

Stage 2: Symbolic Verification (System 2)

- **Input:** Candidate formal specifications from Stage 1
- **Method:** An interactive proof assistant checks type accuracy and logical consistency
- **Output:** A binary validation decision accompanied by structured error diagnostics for iterative refinement
- **Mechanisms:** SMT solving, type checking, proof verification, and counterexample generation

Stage 2 offers deterministic soundness guarantees relative to the formal system’s axioms: if the proof assistant validates a specification, that specification is assured to be logically consistent within the axiom system [29]. Correctness of the specification itself with respect to the original informal intent remains the responsibility of the system designer.




Figure 1: Error-guided refinement loop in a two-stage autoformalization pipeline

2.4 Analytical Framework: Mechanisms of Deterministic Verification

We identify three core operational mechanisms through which neuro-symbolic systems bridge the gap between probabilistic generation and deterministic logical verification. These mechanisms are presented as observed patterns across the surveyed literature, not as theoretically derived requirements.

2.4.1 Mechanism 1: Semantic Projection. The surveyed literature shows that successful frameworks treat the target formal language

as an active constraint system rather than merely an output format. In this architectural view, the target logic — such as SMT-LIB or Lean 4 type theory — acts as a controlled bottleneck [36] that forces the neural model to resolve semantic ambiguity before proof synthesis.

For example, the natural language statement “The function is monotonically increasing” admits multiple interpretations. We observe that surveyed frameworks address this by compelling the neural model to collapse the ambiguity into a precise, machine-parseable definition:

```
def Monotone (f: R -> R) =
  forall a b, a <= b -> f a <= f b
```

We argue that this mechanism is architecturally significant because it converts an unbounded semantic space into a finite, decidable search space before any symbolic reasoning begins.

2.4.2 Mechanism 2: Satisfiability Gating. The surveyed frameworks consistently prioritize logical soundness as their primary optimization criterion. The literature shows this is achieved through a gatekeeping pattern in which a symbolic solver — such as an interactive proof assistant or Z3 SMT solver — serves as the authoritative arbiter of logical validity [4].

When a neural component incorrectly asserts that “For all integers $n > 0$, $2n + 1$ is even,” the symbolic verification component produces a concrete counterexample (e.g., $n = 1$, yielding $2(1) + 1 = 3$, which is odd) [10]. We observe that all four surveyed frameworks implement this pattern: no output that fails symbolic verification propagates to the final result, regardless of the neural component’s confidence.

2.4.3 Mechanism 3: Error-Guided Refinement. The surveyed literature shows a consistent shift from one-shot generation toward iterative refinement driven by structured error feedback. Rather than treating verification failures as terminal, modern systems use diagnostic outputs from symbolic components as deterministic feedback signals to guide neural regeneration [27, 37].

We observe that this mechanism is what transforms the two-stage pipeline from a static filter into a dynamic closed-loop system. When the symbolic engine returns a structured error such as “unknown identifier ‘sum_first_n_odds’”, this diagnostic is injected into the neural model’s context window, progressively narrowing the probabilistic search space toward provably correct solutions.

2.5 Evaluation Dimensions

We assess current frameworks across six dimensions identified as critical for closing the verification gap:

- **System 1 (Neural Component):** The probabilistic generation engine used for semantic translation
- **System 2 (Symbolic Component):** The deterministic verification oracle providing soundness guarantees relative to the formal axiom system
- **Feedback Mechanism:** Error-guided refinement protocols implemented architecturally
- **Retrieval Augmentation:** Knowledge databases used to ground and constrain neural generation
- **Soundness Guarantee:** The degree of formal correctness assurance provided, and its conditions

- **Empirical Performance:** Benchmark results from the original publications, reported for reference rather than cross-framework comparison

3 RESULTS

3.1 Framework Task and Mechanism Map

Table 2 maps the surveyed frameworks according to task type, architectural components, and the primary verification mechanism each instantiates. Because the frameworks address distinct tasks under different evaluation protocols, benchmark figures are included for reference only and must not be interpreted as a performance ranking.

A Note on Framework Heterogeneity. The surveyed frameworks address related but meaningfully distinct tasks. LeanDojo targets tactic-level proof search within a retrieval-augmented pipeline; FormalAlign evaluates and improves statement-level alignment between informal and formal representations; AlphaGeometry solves competition-level geometry problems through domain-constrained construction; and Baldur generates whole proofs with post-hoc symbolic repair. We do not treat these as directly comparable implementations of a single task, and we make no cross-framework performance claims. Rather, the contribution of this survey lies in demonstrating that the same three architectural mechanisms recur consistently across this heterogeneous landscape — a finding that suggests their generality as organizing principles for neuro-symbolic autoformalization independent of domain or task specification.

3.2 Framework-Specific Findings

3.2.1 Retrieval-Augmented Theorem Proving: LeanDojo. Yang et al. [37] introduce LeanDojo, a framework that grounds LLM outputs in retrieved mathematical premises from Lean’s mathlib library. By prioritizing relevant theorems and lemmas, the retrieval mechanism reduces neural hallucinations and limits exploration of unpromising proof paths.

Surveyed finding: The Lean 4 type-theoretic kernel enforces syntactic and semantic correctness, ensuring only valid proof steps are accepted. Retrieval-augmented guidance, combined with this strict kernel, improves proof search efficiency and the synthesis of nontrivial formal proofs.

3.2.2 Statement Alignment Evaluation: FormalAlign. Lu et al. [27] present FormalAlign, a framework for evaluating and improving alignment between informal mathematical statements and their formal representations. The system uses structured feedback from alignment evaluation metrics to iteratively improve autoformalization performance.

Surveyed finding: Incorporating alignment feedback into model training improves formalization accuracy. While the system does not claim token-level error correction, structured feedback allows more reliable translation from informal to formal statements.

3.2.3 Domain-Constrained Geometry Solving: AlphaGeometry. Trinh et al. [32] describe AlphaGeometry, which combines neural generation of geometric constructions with a symbolic deduction

engine applying Wu’s method. By restricting generation to geometrically valid constructions, AlphaGeometry improves reasoning efficiency in Euclidean geometry.

Surveyed finding: Strong domain-specific symbolic constraints reduce the effective search space, enabling solutions to competition-level geometry problems. Neural construction proposals verified via symbolic reasoning enhance reliability and accuracy in proof synthesis.

3.2.4 Whole-Proof Generation with Repair: Baldur. First et al. [12] introduce Baldur, a framework that generates entire proofs using LLMs and applies post-hoc symbolic repair via Isabelle/HOL. This approach treats proof construction as a global optimization problem rather than incremental tactic prediction.

Surveyed finding: Post-hoc symbolic repair allows recovery from initial neural generation errors. Deterministic verification via the proof kernel ensures global logical consistency without requiring token-level interleaving of neural generation and verification.

3.3 Cross-Framework Synthesis

A repeating architectural pattern emerges across all surveyed frameworks. The surveyed literature shows that each implementation maintains a fundamental division of computational responsibilities, despite significant variation in target domains, system scale, and learning objectives:

- Neural components produce candidate hypotheses, constructions, or proof structures through statistical pattern matching
- Symbolic components enforce mathematical soundness and logical correctness through deterministic formal verification

We observe that this division is not incidental: no surveyed framework attempts to achieve formal soundness through neural generation alone. We argue that this consistent pattern supports the interpretation that solver-aided verification functions as an architectural requirement in this domain rather than an optional enhancement.

3.4 Demonstration of Mechanisms in Practice

The following synthesized example concretely illustrates how the three mechanisms operate in sequence:

Natural Language Specification: “Prove that for all positive integers n , the sum of the first n odd numbers equals n^2 .”

Pure Neural Output (Hallucination):

```
theorem sum_odd_numbers (n: Nat):
  sum_first_n_odds n = n^2 := by
  induction n with
  | zero => rfl
  | succ n ih =>
    calc sum_first_n_odds (n+1)
      = sum_first_n_odds n + (2*n+1) := by simp
      = n^2 + (2*n+1)                 := by rw [ih]
      = (n+1)^2                       := by ring
```

The LLM exhibits premise hallucination: it assumes the existence of `sum_first_n_odds` without defining it. This is a documented failure mode of pure neural systems [5].

Table 2: Neuro-Symbolic Autoformalization Frameworks: Task and Mechanism Map

Framework	Task Type	Neural (System 1)	Symbolic (System 2)	Primary Mechanism	Benchmark (Reference Only)
LeanDojo [37]	Tactic-level proof search	ReProver (Encoder-Decoder)	Lean 4 Kernel	Error-Guided Refinement	26–51% Pass@1
FormalAlign [27]	Statement alignment evaluation	CodeT5+ (Fine-tuned)	Lean 4 Type Checker	Satisfiability Gating	66–99% Alignment Accuracy
AlphaGeometry [32]	Domain-constrained geometry solving	Custom Transformer	Algebraic Solver (Wu’s Method)	Semantic Projection	25/30 IMO-AG-30
Baldur [12]	Whole-proof generation with repair	PaLM-2 (Minerva)	Isabelle Proof Kernel	Error-Guided Refinement	47.5% miniF2F

Note: This table is a task-and-mechanism map, not a performance comparison. LeanDojo reports Pass@1 on tactic prediction splits; FormalAlign reports alignment accuracy rather than end-to-end proving success; AlphaGeometry is evaluated on the IMO-AG-30 geometry benchmark; Baldur reports theorem-proving success on miniF2F. These metrics measure different things under different protocols and cannot be meaningfully compared across rows.

Neuro-Symbolic Resolution:

- (1) **Mechanism 1 (Semantic Projection):** Neural output is constrained to valid Lean 4 type-theoretic syntax before proof synthesis begins.
- (2) **Mechanism 2 (Satisfiability Gating):** The Lean type checker rejects the proof and returns the diagnostic “unknown identifier ‘sum_first_n_odds’”.
- (3) **Mechanism 3 (Error-Guided Refinement):** The structured error is injected into the neural context, triggering regeneration with the missing definition supplied:

```
def sum_first_n_odds: Nat -> Nat
| 0     => 0
| n + 1 => sum_first_n_odds n + (2 * n + 1)

theorem sum_odd_numbers (n: Nat):
  sum_first_n_odds n = n^2 := by
  induction n with
  | zero => rfl
  | succ n ih =>
    calc sum_first_n_odds (n+1)
      = sum_first_n_odds n + (2*n+1) := by rfl
      = n^2 + (2*n+1)                := by rw [ih]
      = (n+1)^2                      := by ring
```

This example illustrates how iterative neuro-symbolic refinement converts a plausible but formally invalid neural output into a machine-checked proof.

3.5 Fundamental Computational Limitations

The surveyed literature consistently acknowledges that neuro-symbolic systems inherit foundational constraints from theoretical computer science. These are not engineering limitations to be overcome but mathematical boundaries within which all such systems must operate.

3.5.1 Undecidability Constraints. It is a proven mathematical result that first-order logic with equality is undecidable: no algorithm can determine the validity of arbitrary logical statements in finite time [7]. The surveyed frameworks exhibit this limitation in practice:

- **LeanDojo:** The literature reports that performance deteriorates on problems with long chains of logical dependencies. Premise retrieval can recommend tactics that fail to close proof goals, potentially leading to non-terminating search behaviour [37].

- **AlphaGeometry:** The system enforces explicit constructive depth limits on auxiliary point introduction to preserve tractability. Proofs requiring constructions beyond these limits cannot be solved [32].

3.5.2 Intractability Constraints. Even within decidable logical fragments, the surveyed systems encounter severe complexity barriers:

- **Complexity-Theoretic Barriers:** Satisfiability checking in many SMT theories is NP-complete; combinations of theories may induce higher complexity classes [3].
- **LeanDojo:** Global proof planning may trigger combinatorial explosion in the tactic search space, leading to timeouts on difficult problems [2].
- **Symbolic Equivalence:** Logical equivalence checking for deeply nested expressions exhibits exponential worst-case complexity ($O(2^n)$), requiring heuristic approximation in practice [23].
- **FormalAlign:** The combinatorial alignment space between natural language and formal code is computationally unmanageable at scale; beam search pruning trades global optimality for tractable execution [27].

3.6 Comparison with Classical Verification Methodologies

Table 3 synthesizes the surveyed neuro-symbolic systems relative to traditional manual formal methods and pure neural approaches. **Key Takeaway:** Neuro-Symbolic AI offers labour amortisation rather than a computational breakthrough. Verification complexity remains bounded by NP-completeness and undecidability, but the approach automates the formalization process, saving substantial expert effort. The value proposition is making the theoretically feasible practically achievable, not circumventing fundamental limits.

4 DISCUSSION

The preceding results section is descriptive: it reports what the surveyed frameworks do, how they are constructed, and what computational limits they face. The observations in this section go beyond those findings. They represent our analytical interpretation of the surveyed evidence and constitute the paper’s theoretical argument. We signal this distinction explicitly: claims prefaced with *the surveyed literature shows* or *we observe* are grounded directly in reported findings; claims prefaced with *we argue* are interpretive positions constructed from those findings.

Table 3: Verification Paradigm Comparison

Criterion	Characterization
<i>Manual Formal Methods</i> [21, 22]	
Human Effort	Months to years of expert effort
Computational Cost	Moderate (automated solver time)
Coverage	Limited by human cognitive capacity
Soundness	Formal guarantees when successfully verified
<i>Pure Neural Approaches</i> [6, 33]	
Human Effort	Minimal (inference only)
Computational Cost	Low (single forward pass)
Coverage	High breadth, statistically unreliable
Soundness	No formal guarantees
<i>Neuro-Symbolic Hybrid</i> [27, 37]	
Human Effort	Minimal to moderate (review only)
Computational Cost	Moderate to high (iterative solver calls)
Coverage	High breadth with formal verification
Soundness	Conditional: logically sound relative to the formal specification; correctness of the specification with respect to original informal intent remains the designer’s responsibility

4.1 On the Recurrence of Three Mechanisms

The surveyed literature shows that semantic projection, satisfiability gating, and error-guided refinement appear in all four surveyed frameworks despite their heterogeneous tasks and designs. We observe that this recurrence is not a coincidence of selection: the frameworks were chosen to represent maximally distinct task types, yet all three mechanisms are present in each.

We argue that this recurrence reflects a structural property of the autoformalization problem rather than a design fashion. Any system that must translate ambiguous natural language into verifiable formal representations faces the same three-part challenge: collapsing semantic ambiguity before reasoning begins (semantic projection), enforcing logical correctness rather than statistical plausibility (satisfiability gating), and recovering from the inevitable generation failures that arise when operating near the boundary of decidability (error-guided refinement). These mechanisms are, we argue, natural responses to the structure of the problem itself.

4.2 On Theoretical Well-Motivatedness

The following presents a structured theoretical argument for why neuro-symbolic integration is well-motivated for autoformalization. This argument does not constitute a formal proof of necessity. It is an analytical position, constructed from the convergent findings of the reviewed literature, that we present for critical evaluation.

- (1) *The Semantic Bottleneck (surveyed)*: The literature consistently shows that pure symbolic systems cannot process natural language effectively without learning-based mechanisms, owing to the inherent ambiguity of informal discourse [11, 18, 36].

- (2) *The Verification Gap (surveyed)*: Pure neural systems lack grounded logical entailment mechanisms and are intrinsically prone to stochastic hallucinations. The surveyed literature shows they cannot provide formal verification guarantees [5, 16, 17].
- (3) *Logical Requirement of Symbolic Verification (proven)*: It is an established result in formal methods that any computational system offering formal correctness guarantees must incorporate symbolic verification to satisfy soundness and decidability requirements [8, 15, 22]. This is not an empirical claim but a logical consequence of the definitions of soundness and completeness.
- (4) *Statistical Requirement for Natural Language Processing (surveyed)*: The literature shows that statistical learning mechanisms are necessary to handle context-dependence, pragmatic inference, and linguistic ambiguity in natural language [6, 30, 35].
- (5) *Synthesis (argued)*: Taken together, these four observations support the following position: neuro-symbolic integration is not merely a practical engineering convenience but a theoretically well-motivated architecture for automated formalization. A system that must both process natural language and provide formal correctness guarantees requires, by the preceding observations, both statistical and symbolic components. We present this as a reasoned argument grounded in surveyed evidence and established formal methods results, not as a formal proof of logical necessity [13, 20, 24, 38].

4.3 On Computational Limits as Design Constraints

We observe that every surveyed framework implements explicit mechanisms — depth limits, beam pruning, timeout heuristics — to manage the NP-completeness and undecidability barriers documented in Section 3.5. We argue that this is the correct engineering response: not to circumvent these limits, which is mathematically impossible, but to operate tractably within them by using neural heuristics to guide symbolic search away from computationally intractable regions. The value proposition of neuro-symbolic AI, on this view, is precisely that the neural component functions as a learned approximation of the intractable symbolic search, trading worst-case optimality for practical tractability.

5 CONCLUSION

Both neural and symbolic paradigms possess distinct and complementary advantages. Neural systems offer broad linguistic coverage and semantic flexibility but lack formal correctness guarantees. Symbolic systems provide mathematical certainty and logical rigor but struggle with the ambiguity and scale of natural language. Our review finds that Neuro-Symbolic AI, operationalized through autoformalization, effectively bridges these approaches: neural components handle semantic interpretation while symbolic components enforce strict logical verification.

The surveyed systems — LEANDOJO, FORMALALIGN, BALDUR, and ALPHAGEOMETRY — instantiate three key mechanisms across heterogeneous tasks: **Semantic Projection, Satisfiability Gating,**

and **Error-Guided Refinement**. The recurrence of these mechanisms across systems with distinct tasks, domains, and designs constitutes the primary finding of this survey. Although neuro-symbolic systems inevitably operate within the fundamental limits of undecidability and NP-completeness, we argue that they represent a theoretically well-motivated architecture for autoformalization — one that achieves a balance between formal verifiability and semantic flexibility that neither paradigm can offer in isolation.

5.1 Prospects for Further Research

Toward Fully Automated Formal Development: The long-term goal is AI systems that accept high-level natural language specifications, produce formally verified implementations, and improve continuously through structured feedback. Progress depends on advances in specialised neural architectures such as Aristotle [1] and Goedel-Prover-V2 [25], improved symbolic reasoning techniques such as TacticToe [14] and Deep Network Guided Proof Search [26], and richer human-AI collaborative interfaces.

Benchmark Standardisation: Standardised evaluation metrics are currently absent from the autoformalization domain. While miniF2F [40] and domain-specific assessments [38] offer preliminary frameworks, comprehensive cross-domain benchmarks with standardised protocols remain a critical research priority.

Cross-Domain Transfer Learning: Modern systems exhibit strong domain specialisation but limited transfer. Future work should address multilingual autoformalization, transfer between heterogeneous formal systems (e.g., Lean to Coq), and adaptation of mathematical reasoning models to software verification contexts.

DECLARATION OF GENERATIVE AI IN THE WRITING PROCESS

During the preparation of this work, the authors used ChatGPT, Gemini, and Claude to assist with technical drafting, literature synthesis, and $\text{\LaTeX}$ manuscript formatting. These tools were employed to improve structural organization and refine the conversion of theoretical frameworks into formal academic prose. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the intellectual content, technical accuracy, and scientific integrity of the work.

ACKNOWLEDGEMENT

The author gratefully acknowledges the University of the Philippines – Intelligent Systems Center XAI Team for funding this research. The author also extends appreciation to the Computational Science Research Center for their institutional and technical support. Special thanks are given to Dr. Henry Adorna and Dr. Alfonso Labao for their valuable inputs and insights that contributed to the development of this work.

REFERENCES

- [1] Tudor Achim et al. 2025. Aristotle: IMO-Level Automated Theorem Proving. *arXiv preprint arXiv:2510.01346* (2025). Large-scale transformer-based system for mathematical theorem proving.
- [2] Albert Atserias and Moritz Müller. 2021. Proof Search and the Structure of Solution Spaces. *IMTech Research Note* (2021). Discusses NP-hardness of proof search and combinatorial explosion aspects.
- [3] Clark Barrett, Roberto Sebastiani, Sanjit A Seshia, and Cesare Tinelli. 2010. Satisfiability Modulo Theories. In *Handbook of Satisfiability*. IOS Press, 825–885.
- [4] Sam Bayless, Stefano Buliani, Darion Cassel, Byron Cook, et al. 2025. A Neurosymbolic Approach to Natural Language Formalization and Verification. *arXiv preprint arXiv:2511.09008* (2025).
- [5] Emily M Bender et al. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *EACL*.
- [6] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, Vol. 33. 1877–1901.
- [7] Alonzo Church. 1936. A note on the Entscheidungsproblem. *The Journal of Symbolic Logic* 1, 1 (1936), 40–41.
- [8] Edmund M Clarke, Orna Grumberg, and Doron Peled. 1999. *Model Checking*. MIT press.
- [9] Brandon C. Colelough and William Regli. 2025. Neuro-Symbolic AI in 2024: A Systematic Review. *arXiv preprint arXiv:2501.05435* (2025).
- [10] Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 337–340.
- [11] Edsger W. Dijkstra. 1978. On the foolishness of “natural language programming”. *Selected Writings on Computing: A Personal Perspective* (1978), 332–335. Originally circulated as EWD667.
- [12] Emily First, Markus N. Rabe, Talia Ringer, and Yuriy Brun. 2023. Baldur: Whole-Proof Generation and Repair with Large Language Models. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE 2023)* (San Francisco, CA, USA).
- [13] Artur d’Avila Garcez, Marco Gori, Luis C. Lamb, Luciano Serafini, Michael Spranger, and Son N. Tran. 2019. Neural-Symbolic Computing: An Effective Methodology for Principled Integration of Machine Learning and Reasoning. *Journal of Applied Logics* 6, 4 (2019), 611–632.
- [14] Thibault Gauthier, Cezary Kaliszyk, and Josef Urban. 2020. TacticToe: Learning to Prove with Tactics. *Journal of Automated Reasoning* 64 (2020), 1–37.
- [15] C. A. R. Hoare. 1969. An Axiomatic Basis for Computer Programming. *Commun. ACM* 12, 10 (1969), 576–580.
- [16] Xiaowei Huang, Daniel Kroening, Wenjie Ruan, James Sharp, Youcheng Sun, Emese Thamo, Min Wu, and Xingyu Yi. 2020. A survey of safety and trustworthiness of deep neural networks. *Computer Science Review* 37 (2020), 100270.
- [17] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of Hallucination in Natural Language Generation. *Comput. Surveys* 55, 12 (2023), 1–38.
- [18] Daniel Jurafsky and James H Martin. 2023. *Speech and language processing*. Pearson.
- [19] Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- [20] Henry A. Kautz. 2022. The Third AI Summer: AAAI Robert S. Engelmore Memorial Lecture. *AI Magazine* 43, 1 (2022), 105–125. doi:10.1002/aaai.12036
- [21] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Phillip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. 2009. seL4: Formal Verification of an OS Kernel. In *Proceedings of the 22nd ACM Symposium on Operating Systems Principles (SOSP)*. 207–220.
- [22] Xavier Leroy. 2009. Formal Verification of a Realistic Compiler. *Commun. ACM* 52, 7 (2009), 107–115.
- [23] Zenan Li, Yifan Wu, Zhaoxu Li, Xinming Wei, Fan Yang, Xian Zhang, and Xiaoxing Ma. 2024. Autoformalize Mathematical Statements by Symbolic Equivalence and Semantic Consistency. In *Advances in Neural Information Processing Systems*, Vol. 37.
- [24] Xixun Lin, Yucheng Ning, Jingwen Zhang, Yan Dong, Yilong Liu, Yongxuan Wu, Xiaohua Qi, Nan Sun, Yanmin Shang, Pengfei Cao, Lixin Zou, Xu Chen, Chuan Zhou, Jia Wu, Shirui Pan, Bin Wang, Yanan Cao, Kai Chen, Songlin Hu, and Li Guo. 2025. LLM-based Agents Suffer from Hallucinations: A Survey of Taxonomy, Methods, and Directions. *arXiv preprint arXiv:2509.18970* (2025). <https://arxiv.org/abs/2509.18970>
- [25] Yong Lin, Shange Tang, Bohan Lyu, Ziran Yang, Jui-Hui Chung, Haoyu Zhao, Lai Jiang, Yihan Geng, Jiawei Ge, Jingruo Sun, Jiayun Wu, Jiri Gesi, Ximing Lu, David Acuna, Kaiyu Yang, Hongzhou Lin, Yejin Choi, Danqi Chen, Sanjeev Arora, and Chi Jin. 2025. Goedel-Prover-V2: Scaling Formal Theorem Proving with Scaffolded Data Synthesis and Self-Correction. *arXiv preprint arXiv:2508.03613* (2025).
- [26] Sarah Loos, Geoffrey Irving, Christian Szegedy, and Cezary Kaliszyk. 2017. Deep Network Guided Proof Search. In *International Conference on Logic for Programming, Artificial Intelligence and Reasoning*.
- [27] Jialin Lu, Yixuan Wan, Yinghui Huang, Jing Xiong, Zhengying Liu, and Zhijiang Guo. 2024. FormalAlign: Automated Alignment Evaluation for Autoformalization. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [28] Gary Marcus. 2020. The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence. *arXiv preprint arXiv:2002.06177* (2020).
- [29] Lawrence C. Paulson. 1994. *Isabelle: A Generic Theorem Prover*. Springer.

- [30] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. In *OpenAI Preprint*. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf.
- [31] Stuart Russell and Peter Norvig. 2010. *Artificial Intelligence: A Modern Approach*. Prentice Hall.
- [32] Trieu H Trinh, Yuhuai Wu, Quoc V Le, He He, and Thang Luong. 2024. Solving Olympiad Geometry without Human Demonstrations. In *Nature*, Vol. 625. 476–482.
- [33] Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, and Subbarao Kambhampati. 2023. On the Planning Abilities of Large Language Models: A Critical Investigation. In *Advances in Neural Information Processing Systems*, Vol. 36.
- [34] Frank van Harmelen and Annette Teije. 2019. A Boxology of Design Patterns for Hybrid Learning and Reasoning Systems. In *Semantic Technology*. Springer, 3–13.
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Advances in Neural Information Processing Systems*, Vol. 30. 5998–6008.
- [36] Yuhuai Wu, Albert Q. Jiang, Wenda Li, Markus N. Rabe, Charles Staats, Christian Stein, and Christian Szegedy. 2022. Autoformalization with Large Language Models. *arXiv preprint arXiv:2205.12615* (2022).
- [37] Kaiyu Yang et al. 2023. LeanDojo: Theorem Proving with Retrieval-Augmented Language Models. In *Advances in Neural Information Processing Systems*, Vol. 36.
- [38] Lan Zhang, Marco Valentino, and André Freitas. 2025. Autoformalization in the Wild: Assessing LLMs on Real-World Mathematical Definitions. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. 1720–1738.
- [39] Yedi Zhang, Yufan Cai, Xinyue Zuo, Xiaokun Luan, Kailong Wang, Zhe Hou, Yifan Zhang, Zhiyuan Wei, Meng Sun, Jun Sun, Jing Sun, and Jin Song Dong. 2024. The Fusion of Large Language Models and Formal Methods for Trustworthy AI Agents: A Roadmap. *arXiv preprint arXiv:2412.06512* (2024).
- [40] Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. 2022. miniF2F: A Cross-System Benchmark for Formal Olympiad-Level Mathematics. In *International Conference on Learning Representations*.



April 23–25, 2026
University of the Immaculate Conception, Davao City

Author Index

Abrenica, F.R., 156	Adorna, H.N., 187	Afable, J.V., Jr., 34
Agarap, A.F., 10, 93, 125	Aguinaldo, M.A., 10	Antoque, J.T., 47
Ascaño, C.I.E., 55	Azcarraga, A., 125	Barria, R., 179
Benavides, I.M., 10	Bernabe, M.A.S., 71	Binangbang, D.C., 63
Bitayo, B.C., 63	Bulio, N., 3	Buño, K.C., 85, 171
Cabarle, F.G., 77, 85, 132, 140, 171	Calabia, B.N.F., 18	Cardaño, M.S., 3
Cayabyab, G.S.S., 55	Chio, M.A., 171	Consunji, B., 163
Datiles, G., 34	De la Cruz, R.T.A., 85, 171	Deja, J.A., 148
Delos Reyes, J., 148	Dy, S., 148	Fernandez, R.A., 148
Floreta, Z.O., 63	Gabud, R.S., 18	Gamilla, K.L.Q., 55
Go, D.J.L., 26	Go, J.R., 93	Guirnela, E.J., 109
Heffron, J.S., 18	Isanan, J.A.B., 195	Jakosalem, C.M., 156
Jimenez, Z.B., 132	Ko, D.C.L., 85, 171	Lanting, K.N.P., 171
Lasala, K.C.C., 26	Lipat, J., 77	Lucero, S.R., 47
Luzada, J.S., 77	Madrid, V.R.M., 71	Madrigal, R., 10
Manlangit, V., ??	Manlises, M.M., 26	Marqueses, L.B., 93
Martinez, M.K., 93	Martínez-del-Amor, M.Á., 77	Monserrat, T.-J.K., 148
Muli, L., 163	Naraval, V.D.D., 85	Omictin, J.V., 63
Orellana-Martín, D., 77	Paglinawan-Muñoz, N., 3, 41	Paglinawan-Muñoz, N.M., 47
Pardillo, J.A.S., 63	Pepito, G.B., 117	Peña, C., 109
Peña, C.F., 179	Pilapil, M.D.M., 179	Pillodar, F.C.O., 47
Pintor, R.R.H., 117	Pontiveros, M.J., 187	Ramirez, R.M., II, 140
Ramos, T.J.E., 101	Regonia, P.R.R., 18	Rosillosa, K.A.R., ??
Roxas, R.E.O., 101	Sales, R.M., 148	Sarmiento, J.K.P., 93
Silva, L.T., 41	Soriano, J., 34	Tumampos, M.V., 109
Uy, O.L., 148	Uy, R.L.L., 26	Valdez, A.R., 34
Venturina, S.A., 10	Vergara, J., 163	Villar, J.J., 34
Wang, L.J.P., 117	Zapico, J., 163	



April 23–25, 2026
University of the Immaculate Conception, Davao City

Institution Index

Ateneo de Manila University, 163
Cebu Institute of Technology – University, 63
De La Salle University, 26, 93, 125, 148
De La Salle University – Dasmariñas, 55
Globe Fintech Innovations Inc., 10
South Philippine Adventist College, 3, 41, 47
Universidad de Sevilla, 77
University of San Carlos, 109, 117, 156, 179
University of the Philippines, 187
University of the Philippines Diliman, 18, 34, 77, 85, 132, 140, 171, 195
University of the Philippines Los Baños, 71, 101