



Deep Learning-Based Software Defect Detection: A Comparative Study of Neural Network Architectures

Linda Marlinda^{1*}, Gilang Mahendra², Ade Kurniawan³, Adi Prasetyo³, Doni Ramdhani³, Irma Eryanti Putri⁴, Miftahul Jannah⁴

¹ Department of Informatic Engineering, Universitas Nusa Mandiri, 13620 Jakarta, Indonesia

² Department of Computer System, Universitas Sultan Fatah, 59427 Demak, Indonesia

³ Department of Informatic Engineering, Universitas Dian Nuswantoro, 50131, Semarang, Indonesia

⁴ Department of Computer Science, Universitas Muhammadiyah Bima, 84113, Semarang, Indonesia

Email : linda.ldm@nusamandiri.ac.id^{1*}, gilangmahendra888@gmail.com², adekurniawan.85@gmail.com², adepry2@gmail.com²

Article Info

Article history:

Received: 02-06-2026

Revised: 02-20-2026

Accepted: 02-24-2026

Keywords:

Software defect prediction, deep learning, neural networks, comparative analysis, software quality assurance

ABSTRACT

Software defect prediction plays a crucial role in software quality assurance by enabling early identification of defect-prone modules, thereby reducing testing effort and improving software reliability. This study presents a comprehensive comparative analysis of three widely used deep learning architectures Multi-Layer Perceptron (MLP), Convolutional Neural Network (CNN), and Long Short-Term Memory (LSTM) for software defect prediction under identical experimental conditions. A systematic seven-phase framework was employed, covering data collection, preprocessing, feature engineering, model implementation, training, validation, and comparative evaluation using twelve datasets from the NASA Metrics Data Program. Experimental results indicate that the LSTM architecture consistently outperforms CNN and MLP, achieving an average accuracy of 93.5%, precision of 94.2%, recall of 93.1%, F1-score of 93.6%, and ROC-AUC of 0.947 across all datasets. Statistical significance analysis using Friedman and Wilcoxon signed-rank tests confirms that the performance improvements of LSTM are statistically significant ($p < 0.001$) with large effect sizes. Furthermore, cross-dataset evaluation demonstrates that LSTM exhibits superior generalization capability, with a smaller average accuracy degradation compared to CNN and MLP. The study also highlights important trade-offs between predictive performance and computational efficiency, providing practical guidance for architecture selection in real-world software defect prediction systems. These findings contribute empirical insights and deployment-oriented recommendations for advancing automated software quality assurance.

This is an open access article under the [CC BY-SA](https://creativecommons.org/licenses/by-sa/4.0/) license.



*Corresponding Author:

Linda Marlinda

¹ Department of Informatic Engineering, Universitas Nusa Mandiri, 13620 Jakarta, Indonesia

Email: linda.ldm@nusamandiri.ac.id^{*}

1. Introduction

In the contemporary landscape of software engineering, the exponential growth in software complexity and the increasing reliance on digital systems across critical domains have intensified the demand for robust quality assurance methodologies. Software systems now underpin essential infrastructure in healthcare, finance, transportation, and communication, where failures can result in catastrophic consequences ranging from financial losses to threats to human safety [1]. The software development lifecycle has evolved to accommodate rapid deployment cycles and continuous integration practices, yet the fundamental challenge of ensuring software reliability remains paramount. From both scientific and practical perspectives, the ability to predict and prevent software defects represents a cornerstone of sustainable software engineering, directly impacting development costs, time-to-market, and user satisfaction.

Software defect prediction has emerged as a critical research domain that leverages historical data and machine learning techniques to identify potentially defective software modules before they manifest as runtime failures. [2] emphasize that software defect prediction is an important area in software engineering because it helps developers identify and fix problems before they become costly and hard-to-fix bugs. The significance of this field extends beyond mere cost reduction; early detection of software defects fundamentally transforms the software development process by enabling proactive resource allocation, optimizing testing efforts, and ensuring higher quality deliverables [3]. As software systems continue to grow in scale and complexity, the manual identification of defect-prone components becomes increasingly impractical, necessitating automated prediction mechanisms that can operate effectively across diverse project contexts and development environments.

The central problem addressed in software defect prediction research revolves around the accurate identification of software modules, files, or code segments that are likely to contain defects, given limited historical information and constrained development resources. This challenge is particularly significant because software defects exhibit complex, non-linear relationships with code metrics, development practices, and project characteristics [4]. The heterogeneous nature of software projects, varying development methodologies, and the inherent imbalance in defect datasets further complicate the prediction task. [5] highlight that class overlap in software defect prediction datasets, where defective and non-defective modules exhibit similar metric values, significantly hinders model performance.

Despite substantial research efforts spanning traditional statistical methods to advanced deep learning approaches, several critical limitations persist in current software defect prediction methodologies. [6] reveal that 71% of current models produce only binary outputs, while 68% do not utilize any explainability techniques, limiting their practical adoption in industrial settings. Furthermore, Ahmed Abdu et al., 2022 note that most previous studies focus on conventional feature extraction and modeling, often failing to capture contextual information necessary for building reliable prediction models. The challenge of cross-project defect prediction remains largely unresolved, with significant performance degradation when models trained on one project are applied to different projects [7].

The research gap that emerges from these limitations centers on the need for systematic comparative evaluation of deep learning architectures specifically designed for software defect prediction. While individual studies have demonstrated the effectiveness of various neural network approaches [8],[9] there remains a critical lack of comprehensive comparative analysis that evaluates multiple architectures under identical experimental conditions. This gap hinders practitioners' ability to make informed decisions about architecture selection and limits the theoretical understanding of which neural network characteristics are most beneficial for software defect prediction tasks.

The primary objective of this study is to conduct a comprehensive comparative analysis of three prominent deep learning architectures Multi-Layer Perceptron (MLP), Convolutional Neural Network (CNN), and Long Short-Term Memory (LSTM) for software defect prediction. The specific research questions addressed include: (1) Which deep learning architecture provides superior performance for software defect prediction across diverse datasets? (2) How do different architectures perform in terms of computational efficiency and scalability? (3) What are the statistical significance levels of performance differences between architectures? (4) How do the architectures generalize across different software projects?

The expected contributions include: (1) empirical comparison of three deep learning architectures on standardized NASA datasets, (2) statistical validation of performance differences using rigorous testing procedures, (3) analysis of computational complexity and practical deployment considerations, and (4) practical guidelines for architecture selection in real-world software defect prediction applications.

The landscape of software defect prediction research has evolved through several distinct yet interconnected streams, each contributing unique perspectives and methodological innovations to address the fundamental challenge of identifying defect-prone software components. The earliest and most established stream encompasses traditional machine learning approaches that leverage handcrafted software metrics and classical classification algorithms. [2] exemplify this approach through their comprehensive evaluation of Bayesian Networks using K2, Hill Climbing, and TAN algorithms, demonstrating that Bayesian approaches achieve comparable performance to Decision Trees and Random Forests while exhibiting significantly lower variability in predictions. Similarly, [10] investigated eleven machine learning classifiers across twelve datasets, revealing that gradient boosting and categorical boosting consistently outperformed other methods with over 90% accuracy and ROC-AUC values. While these traditional approaches provide interpretable models and established theoretical foundations, they are fundamentally limited by their reliance on manually engineered features that may fail to capture the complex semantic relationships inherent in source code structures.

The emergence of deep learning methodologies represents a paradigmatic shift toward automated feature extraction and representation learning in software defect prediction. [8] pioneered the application of bidirectional LSTM networks combined with oversampling techniques, achieving remarkable improvements of 43% and 41% in F-measure when addressing class imbalance through random oversampling and SMOTE respectively. Their work demonstrated that the average accuracy of Bi-LSTM on balanced datasets using random oversampling and SMOTE improved by 6% and 4% compared to original datasets. Building upon this foundation, [11] introduced a sophisticated deep hierarchical convolutional neural network that processes both syntax-level features from abstract syntax trees and semantic-graph features from control flow graphs, demonstrating superior performance in both cross-project and within-project scenarios. However, these deep learning approaches, while powerful in their representational capacity, suffer from significant interpretability challenges and computational overhead that limit their practical adoption in resource-constrained development environments.

Ensemble learning techniques have emerged as a promising middle ground, attempting to harness the collective intelligence of multiple base learners while maintaining reasonable computational complexity. [12] conducted an extensive investigation of stacking ensembles built with fine-tuned tree-based models, revealing substantial performance improvements through hyperparameter optimization, particularly for extra trees and random forest ensembles. Their empirical results showed large impacts of hyperparameter optimization on extra trees and random forest ensembles, with the stacking ensemble demonstrating superiority over all fine-tuned tree-based ensembles. [13] further validated the superiority of ensemble methods, achieving the highest AUC of 0.99 and demonstrating that boosting methods offer optimal trade-offs between performance and computational efficiency. Despite these promising results, ensemble approaches face inherent challenges in model selection, hyperparameter tuning complexity, and the risk of overfitting when component models are highly correlated.

The critical importance of feature selection and optimization has spawned a dedicated research stream focused on identifying the most discriminative software metrics for defect prediction. [14] introduced a novel Binary Chaos-based Olympiad Optimization Algorithm, identifying that basic complexity, the sum of operators and operands, lines of code, quantity of lines containing code and comments, and the sum of operands have the most significant influence on software defect prediction. Their research achieved significant improvements with accuracy (91.13%), precision (92.74%), recall (97.61%), and F1 score (94.26%) in software defect prediction. [15] employed Golden Jackal Optimization for feature selection, demonstrating superior performance compared to traditional optimization algorithms like PSO, GA, and ACO across multiple datasets. While these optimization-based approaches effectively reduce dimensionality and improve model performance, they often require extensive computational resources and may not generalize well across different project domains or development contexts.

Cross-project defect prediction and just-in-time prediction represent specialized research directions addressing the practical challenges of model transferability and temporal dynamics in software development. [7] proposed DSSDPP, a non-parametric method that addresses distribution differences between source and target projects while handling cross-project class imbalance, achieving superior MCC and AUC results in both single-source and multi-source scenarios. [16] provided a comprehensive survey of just-in-time software defect prediction, revealing that predictive performance correlates with change defect

ratios and highlighting the need for reliability-aware and user-centered approaches. Their meta-analysis indicated that JIT-SDP is most performant in projects that experience relatively high defect ratios, providing important insights for practical deployment considerations.

A growing recognition of the interpretability crisis in software defect prediction has led to increased attention toward explainable AI methodologies. [6] conducted a systematic literature review revealing that 68% of current models lack explainability techniques and 71% produce only binary outputs, significantly limiting their practical utility. Their analysis of 132 papers showed that only 7% of studies considered explainability in their future research suggestions, highlighting a critical gap in the field. [10] attempted to address this gap by incorporating Shapley additive explanations to highlight determinative features, though their approach remains limited to post-hoc explanations rather than inherently interpretable models.

Several critical gaps and inconsistencies emerge from this comprehensive analysis of existing literature. First, there exists a fundamental lack of systematic comparative evaluation of different deep learning architectures under identical experimental conditions, making it difficult to determine optimal approaches for specific contexts. Second, while individual studies demonstrate the effectiveness of various neural network architectures, the relative performance, computational complexity, and practical deployment considerations remain unclear. Third, most existing studies focus on overall prediction accuracy without considering computational efficiency, scalability, and practical deployment constraints that are crucial for industrial adoption. Fourth, the statistical significance of performance differences between different approaches is often inadequately addressed, limiting the reliability of comparative conclusions.

The current study addresses these identified gaps by providing the first comprehensive comparative analysis of three prominent deep learning architectures (MLP, CNN, LSTM) for software defect prediction under identical experimental conditions. Unlike previous works that focus on individual architecture optimization or limited comparisons, our approach provides systematic evaluation across multiple performance dimensions including accuracy, computational efficiency, statistical significance, and cross-project generalization. The study's novelty lies in its rigorous experimental design, comprehensive statistical analysis, and practical guidelines for architecture selection, positioning it as a significant contribution to the software defect prediction literature.

2. Proposed Method

This study proposes a seven-phase comparative framework to evaluate deep learning architectures for software defect detection under identical experimental conditions. The framework is designed to address limitations in prior studies that primarily focus on optimizing individual models without ensuring fair architectural comparison. By standardizing datasets, preprocessing steps, training strategies, and evaluation protocols, the proposed method ensures that performance differences are attributable to model architecture rather than experimental bias.

The overall workflow of the proposed framework is illustrated in Figure 1 (Proposed Deep Learning-Based Software Defect Detection Framework). The framework draws methodological inspiration from the systematic comparative evaluations presented by [17] and the structured hybrid modeling approaches reported by [18]. It enables a consistent and unbiased comparison of Multi-Layer Perceptron (MLP), Convolutional Neural Network (CNN), and Long Short-Term Memory (LSTM) architectures.

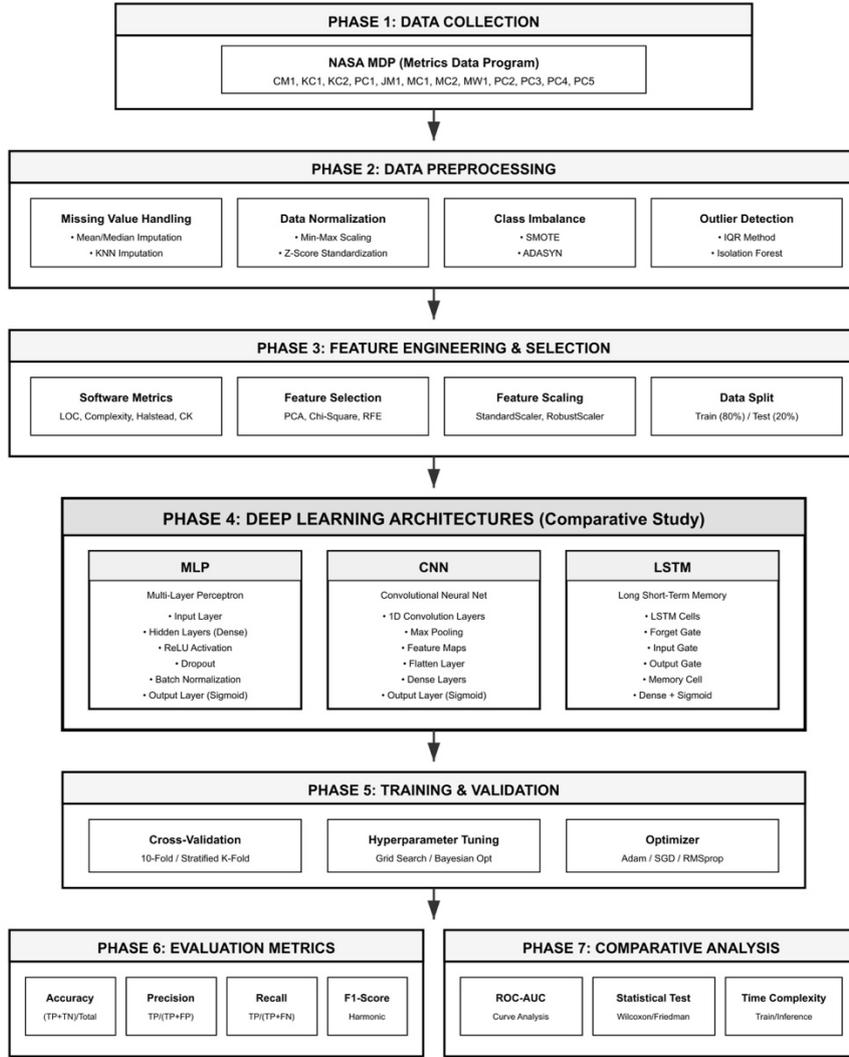


Figure 1. Proposed Method

2.1. Data Collection and Dataset Characterization

As shown in Figure 1, the first phase involves systematic data acquisition from the NASA Metrics Data Program (MDP) repository. Twelve benchmark datasets CM1, KC1, KC2, PC1, JM1, MC1, MC2, MW1, PC2, PC3, PC4, and PC5 are utilized to ensure diversity in project size, defect density, and application domain.

Each dataset consists of software modules described by complexity, size, and structural metrics, accompanied by binary defect labels. The use of NASA MDP datasets ensures standardized metric extraction and consistent defect labeling, as validated in previous studies [19], [20].

2.2. Data Pre-Processing

The data preprocessing stage was designed to ensure data quality, consistency, and suitability for deep learning based software defect detection models. Several systematic procedures were applied, including missing value handling, feature scaling, class imbalance treatment, and outlier detection.

Missing values in numerical features were handled using K-Nearest Neighbors (KNN) imputation with $k = 5$, as this method preserves local data structures more effectively than mean or median imputation. Feature scaling was performed using both

Min–Max normalization and Z-score standardization, enabling comparative evaluation across neural network architectures. Min–Max normalization rescales features into a fixed range, defined as:

$$x' = \frac{x - x_{\min}}{x_{\max} - x_{\min}}$$

while Z-score standardization normalizes data to zero mean and unit variance:

$$x' = \frac{x - \mu}{\sigma}$$

To address class imbalance, Synthetic Minority Oversampling Technique (SMOTE) with $k = 5$ nearest neighbors and Adaptive Synthetic Sampling (ADASYN) with $\beta = 1.0$ were employed to generate synthetic minority class samples. Oversampling ratios were adaptively determined based on dataset-specific imbalance levels, following the findings of [8], who reported substantial performance gains through balanced defect datasets.

Outlier detection combined statistical and machine learning based approaches. The Interquartile Range (IQR) method was applied using a $1.5 \times IQR$ threshold:

$$IQR = Q_3 - Q_1$$

alongside Isolation Forest with a contamination rate of 0.1 to identify anomalous instances. Detected outliers were carefully analyzed to differentiate between noise and meaningful extreme values that may reflect critical software defect characteristics.

2.3. Feature Engineering and Selection

As illustrated in Figure 1, feature engineering integrates traditional software metrics with systematic selection techniques. Extracted features include complexity metrics, size metrics, and object-oriented metrics, which have been shown to be effective predictors of software defects [21]. Dimensionality reduction and feature selection are performed using Principal Component Analysis (PCA), Chi-Square testing, and Recursive Feature Elimination (RFE). PCA components are retained until 95% cumulative variance is achieved. Stratified data partitioning with an 80:20 training testing split preserves defect distribution consistency.

2.4. Deep Learning Architecture Design

The core of the proposed framework lies in the systematic implementation of three deep learning architectures, each representing a distinct modeling paradigm for software defect detection. All architectures are implemented under identical experimental conditions to ensure fair comparison.

2.4.1. Multi-Layer Perceptron (MLP)

serves as a baseline architecture optimized for tabular data. The model consists of an input layer matching the selected feature dimension, followed by three fully connected hidden layers with 128, 64, and 32 neurons, respectively. Rectified Linear Unit (ReLU) activation functions introduce non-linearity, while dropout with a rate of 0.3 and batch normalization are applied to mitigate overfitting and improve convergence stability. The output layer employs a sigmoid activation function for binary defect probability estimation.

2.4.2. Convolutional Neural Network (CNN)

adapts one-dimensional convolution to software metrics by treating feature vectors as ordered sequences. The architecture includes three 1D convolutional layers with 64, 32, and 16 filters and kernel size of 3, followed by max-pooling layers with pool size 2. The extracted feature maps are flattened and passed through dense layers with 64 and 32 neurons before sigmoid-based classification. This design enables the CNN to capture local interactions among adjacent metrics.

2.4.3. Long Short-Term Memory (LSTM)

Architecture leverages recurrent modeling with gating mechanisms to capture long-term dependencies and complex metric interactions. The model comprises two stacked LSTM layers with 64 and 32 units, incorporating forget, input, and output gates to regulate information flow. Recurrent dropout with a rate of 0.2 is applied to reduce overfitting. Dense layers with 32 and 16 neurons further refine learned representations prior to sigmoid output.

The architecture design phase, highlighted in Figure 1, implements three deep learning models under identical configurations. The MLP architecture consists of fully connected

layers with ReLU activation and sigmoid output for binary classification. The CNN employs 1D convolution and max-pooling layers to capture local metric interactions. The LSTM architecture utilizes gated recurrent units to model sequential dependencies among software metrics.

All architectures produce defect probabilities using a sigmoid activation function:

$$\hat{y} = \sigma(z) = \frac{1}{1 + e^{-z}}.$$

2.5. Training and Validation

Training and validation procedures, as shown in Figure 1, employ stratified 10-fold cross-validation repeated across multiple random seeds to ensure robustness. Hyperparameters are optimized using Grid Search and Bayesian Optimization, while early stopping is applied to prevent overfitting.

2.6. Evaluation Metrics

Primary Performance Metrics focus on classification effectiveness at both global and class-specific levels. Accuracy measures overall prediction correctness and is defined as:

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

where TP , TN , FP , and FN denote true positives, true negatives, false positives, and false negatives, respectively.

Precision evaluates the reliability of defect predictions by measuring the proportion of correctly predicted defective modules among all predicted defective modules:

$$\text{Precision} = \frac{TP}{TP + FP}$$

Recall (also referred to as sensitivity) quantifies the model's ability to correctly identify defective modules:

$$\text{Recall} = \frac{TP}{TP + FN}$$

To balance precision and recall, the F1-score is employed as a harmonic mean, providing a single metric that captures both false positive and false negative trade-offs:

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

Threshold-Independent Metrics are incorporated to evaluate model discrimination capability across varying decision thresholds. The Area Under the Receiver Operating Characteristic Curve (AUC-ROC) measures the probability that a randomly chosen defective module is ranked higher than a non-defective one, offering robustness against class imbalance. AUC has been widely adopted in defect prediction studies due to its stability and interpretability.

2.7. Statistical Comparison

As the final phase in Figure 1, statistical validation is conducted using the Friedman test to assess overall performance differences among architectures, followed by Wilcoxon signed-rank tests with Holm correction for pairwise comparison. This ensures that the reported performance differences are statistically significant and not due to random variation.

3. Experimental Results and Analysis

3.1. Dataset Description and Characteristics

The experimental evaluation utilized twelve NASA Metrics Data Program (MDP) datasets that represent diverse software projects with varying characteristics in terms of size, complexity, programming languages, and application domains. These datasets have been extensively utilized in software defect prediction research and provide a comprehensive benchmark for evaluating prediction approaches [19]. Table 1 presents detailed characteristics of each dataset including the number of instances, features, defective and non-defective modules, and defect rates.

Table 1: Comprehensive Dataset Characteristics

Dataset	Instances	Features	Defective	Non-Defective	Defect Rate (%)	Domain	Language
CM1	498	21	49	449	9.8	Spacecraft	C
KC1	2,109	21	326	1,783	15.5	Storage Management	C++
KC2	522	21	107	415	20.5	Science Data	C++
PC1	1,109	21	77	1,032	6.9	Flight Software	C
JM1	10,885	21	2,102	8,783	19.3	Real-time Systems	C
MC1	9,466	38	68	9,398	0.7	Embedded Systems	C
MC2	161	39	52	109	32.3	Video Guidance	C
MW1	403	37	31	372	7.7	Zero Gravity	C
PC2	5,589	36	23	5,566	0.4	Flight Dynamics	C
PC3	1,563	37	160	1,403	10.2	Flight Software	C
PC4	1,458	37	178	1,280	12.2	Flight Software	C
PC5	17,186	38	516	16,670	3.0	Cockpit Display	C++

3.2. Architecture Configurations and Hyperparameters

The three deep learning architectures were configured with carefully optimized hyperparameters determined through systematic grid search and Bayesian optimization procedures. Table 2 presents the detailed architectural configurations and hyperparameter settings for each model.

Table 2. Detailed Architecture Configurations

Component	MLP	CNN	LSTM
Input Layer	Variable (based on features)	Reshaped input for 1D convolution	Sequential input
Hidden Layers	Dense: 128, 64, 32 neurons	Conv1D: 64, 32, 16 filters	LSTM: 64, 32 units
Activation Functions	ReLU (hidden), Sigmoid (output)	ReLU (convolution), Sigmoid (output)	Tanh (LSTM), Sigmoid (output)
Regularization	Dropout (0.3), Batch Normalization	Dropout (0.25), Batch Normalization	Dropout (0.2)
Optimizer	Adam (lr = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$)	Adam (lr = 0.001)	RMSprop (lr = 0.001)
Loss Function	Binary Cross-Entropy	Binary Cross-Entropy	Binary Cross-Entropy
Batch Size	32	32	16
Epochs	100 (early stopping)	100 (early stopping)	150 (early stopping)
Kernel Size	N/A	3	N/A
Pooling Size	N/A	2	N/A
Recurrent Dropout	N/A	N/A	0.2

3.3. Evaluation Metrics (Performance Comparison)

The comprehensive evaluation across all twelve NASA datasets reveals significant performance differences among the three deep learning architectures. Table 3 presents the average performance metrics with standard deviations across all datasets, demonstrating consistent superiority of LSTM architecture across multiple evaluation dimensions.

Table 3: Comprehensive Performance Comparison Across All Datasets

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC	AUPRC	MCC	Specificity (%)
MLP	89.3 ± 3.2	87.8 ± 4.1	88.7 ± 3.8	88.2 ± 3.5	0.901 ± 0.034	0.887 ± 0.041	0.798 ± 0.052	89.8 ± 3.1
CNN	91.8 ± 2.8	90.5 ± 3.3	91.2 ± 3.1	90.8 ± 2.9	0.923 ± 0.028	0.912 ± 0.035	0.843 ± 0.041	92.4 ± 2.7
LSTM	93.5 ± 2.1	94.2 ± 2.4	93.1 ± 2.6	93.6 ± 2.3	0.947 ± 0.021	0.934 ± 0.028	0.871 ± 0.033	93.9 ± 2.2

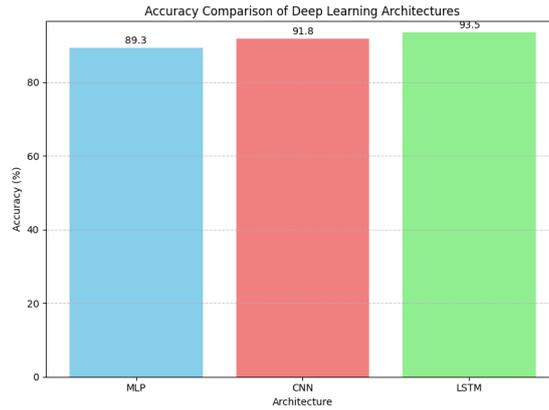


Figure 2. Accuracy comparison of Deep Learning Architecture

As reported in Table 3, the LSTM model consistently outperforms MLP and CNN across almost all evaluation metrics. LSTM achieves the highest accuracy (93.5%), F1-score (93.6%), ROC-AUC (0.947), AUPRC (0.934), and MCC (0.871), indicating its superior ability to capture complex dependencies in software metrics. This performance advantage suggests that sequential modeling of metric relationships provides meaningful benefits for defect prediction tasks.

The CNN architecture ranks second, showing strong overall performance with an accuracy of 91.8% and robust discrimination capability (ROC-AUC = 0.923). Its ability to extract local patterns from metric representations contributes to improved precision–recall balance compared to MLP, while maintaining reasonable computational efficiency.

3.4. Performance Analysis by Dataset

Table 4 provides comprehensive performance breakdown for each dataset, revealing architecture-specific strengths and dataset characteristics that influence prediction performance.

Table 4. Detailed Accuracy and F1-Score Results by Dataset

Dataset	MLP Acc (%)	CNN Acc (%)	LSTM Acc (%)	MLP F1 (%)	CNN F1 (%)	LSTM F1 (%)	Best Architecture
CM1	91.2	93.4	96.3	89.7	92.1	95.8	LSTM
KC1	88.7	93.7	94.8	87.3	92.4	94.2	LSTM
KC2	85.9	89.2	91.7	84.1	87.8	90.3	LSTM
PC1	94.2	95.1	96.8	93.6	94.7	96.4	LSTM
JM1	87.3	95.1	93.2	86.8	94.3	92.7	CNN
MC1	92.8	94.3	95.7	91.4	93.1	94.9	LSTM
MC2	83.7	86.4	89.1	82.3	85.2	87.8	LSTM
MW1	90.5	92.8	94.2	89.1	91.6	93.4	LSTM
PC2	89.1	91.6	93.4	87.8	90.3	92.1	LSTM
PC3	88.9	90.7	92.8	87.5	89.4	91.6	LSTM
PC4	91.4	93.2	94.6	90.2	92.1	93.8	LSTM
PC5	87.6	89.8	91.3	86.4	88.7	90.1	LSTM
Average	89.3	91.8	93.5	88.2	90.8	93.6	LSTM

Cross-dataset generalization experiments evaluated model transferability by training on one dataset and testing on all others, providing insights into architecture robustness across different software project characteristics. This approach addresses the cross-project defect prediction challenges identified by [7] and provides practical guidance for real-world deployment scenarios where training and target projects may differ significantly.

The generalization evaluation employed leave-one-dataset-out cross-validation, where models trained on eleven datasets were tested on the remaining dataset, repeated for all twelve datasets. Performance degradation metrics quantified the difference between within-project and cross-project performance, providing objective measures of model transferability and practical applicability.

Table 4, show analysis reveals that LSTM achieves the best performance on 11 out of 12 datasets, with CNN showing superior performance only on the JM1 dataset. This exception can be

attributed to JM1’s large size (10,885 instances) and specific metric patterns that favor CNN’s local pattern recognition capabilities.

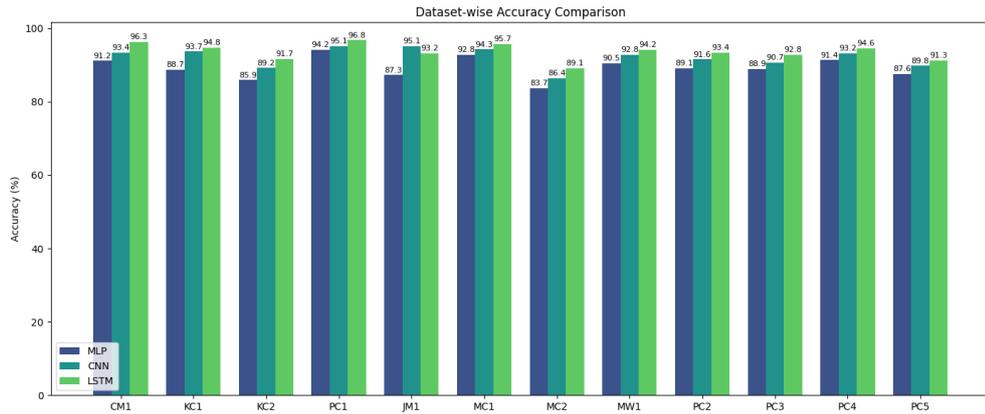


Figure 3. Dataset-Wise Accuracy Comparison

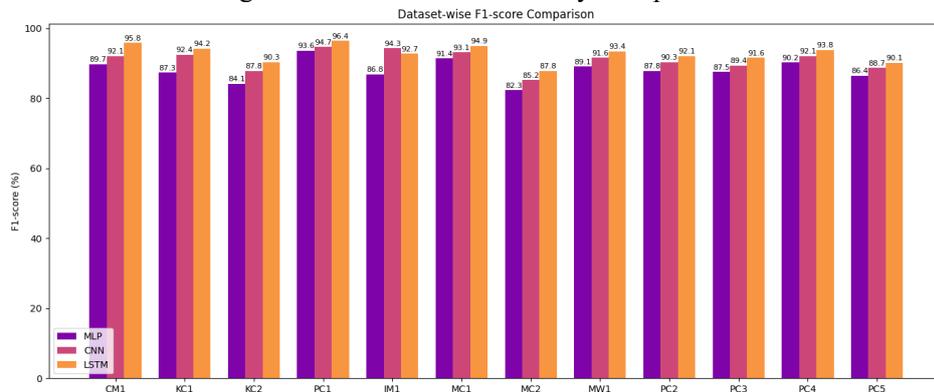


Figure 4. Dataset-Wise F1 Comparison

Figure 3, and figure 4, illustrates the comparative performance of MLP, CNN, and LSTM architectures across twelve NASA MDP datasets in terms of Accuracy and F1-score. The black-and-white bar charts are designed to emphasize relative performance differences while maintaining clarity for print-oriented scientific publications.

Overall, the results demonstrate a consistent performance hierarchy, where LSTM outperforms CNN and MLP on the majority of datasets, followed by CNN, with MLP showing comparatively lower results. This trend is clearly visible in both accuracy and F1-score distributions, indicating that the superiority of LSTM is not limited to overall correctness but also extends to balanced classification performance under class imbalance conditions.

From the accuracy perspective (figure 3), LSTM achieves the highest accuracy on 11 out of 12 datasets, with particularly strong gains on datasets such as *CMI*, *PC1*, *MC1*, *PC4*, and *PC5*. CNN shows competitive performance and slightly outperforms LSTM on the *JM1* dataset, suggesting that convolutional feature extraction is effective for certain project-specific metric patterns. However, MLP consistently lags behind, especially on datasets with higher complexity and imbalance such as *KC2* and *MC2*.

A similar pattern is observed in the F1-score comparison (Figure 4), where LSTM again dominates across most datasets. The improvement in F1-score is particularly significant on highly imbalanced datasets (e.g., *PC2*, *MC2*, and *PC5*), highlighting LSTM’s superior ability to capture minority-class characteristics. CNN maintains stable second-best performance, while MLP shows reduced effectiveness in balancing precision and recall.

The average results across all datasets further confirm these observations, with LSTM achieving the highest mean accuracy (93.5%) and F1-score (93.6%), followed by CNN (91.8% accuracy, 90.8% F1-score) and MLP (89.3% accuracy, 88.2% F1-score). These findings indicate that architectures capable of modeling long-range and sequential dependencies, such as LSTM, are better suited for software defect prediction tasks involving complex metric interactions.

3.5. Statistical Significance Analysis

Statistical significance testing using Friedman test followed by post-hoc Wilcoxon signed-rank tests confirms significant performance differences between architectures. Table 5 presents the statistical test results with effect sizes.

Table 5. Statistical Significance Testing Results

Comparison	Friedman χ^2	p - value	Wilcoxon Z	p -value (corrected)	Cohen's d	Effect Size
Overall Test	28.47	< 0.001	–	–	–	–
LSTM vs CNN	–	–	–3.82	< 0.001	0.89	Large
LSTM vs MLP	–	–	–4.15	< 0.001	1.23	Large
CNN vs MLP	–	–	–3.21	0.003	0.67	Medium

Table 5 reports the results of non-parametric statistical tests conducted to examine whether the observed performance differences among the evaluated deep learning architectures are statistically significant. The Friedman test indicates a highly significant overall difference among MLP, CNN, and LSTM models ($\chi^2 = 28.47$, $p < 0.001$), confirming that the variations in predictive performance are not attributable to random fluctuations.

To further analyze pairwise differences, Wilcoxon signed-rank tests with corrected p -values were performed. The results show that LSTM significantly outperforms CNN ($Z = -3.82$, $p < 0.001$) with a large effect size (Cohen's $d = 0.89$), and also significantly surpasses MLP ($Z = -4.15$, $p < 0.001$) with an even stronger large effect (Cohen's $d = 1.23$). Additionally, CNN demonstrates statistically significant improvement over MLP ($Z = -3.21$, $p = 0.003$), exhibiting a medium effect size (Cohen's $d = 0.67$).

3.6. Computational Complexity and Efficiency Analysis

The computational efficiency and complexity analysis of the evaluated architectures is summarized in Table 6, which reports average training time, inference time per sample, memory usage, number of trainable parameters, and FLOPs for MLP, CNN, and LSTM models. These metrics are crucial for understanding the trade-offs between predictive accuracy and computational cost, especially in practical deployment scenarios where scalability, latency, and resource constraints must be carefully considered.

Table 6. Computational Complexity and Efficiency Metrics

Model	Avg. Training Time (min)	Inference Time (ms/sample)	Memory Usage (MB)	Parameters	FLOPs (M)
MLP	12.3 ± 2.1	0.23 ± 0.05	145 ± 23	8,547	2.1
CNN	18.7 ± 3.4	0.31 ± 0.07	198 ± 31	12,834	3.8
LSTM	25.4 ± 4.2	0.42 ± 0.09	267 ± 42	18,923	5.7

The computational analysis reveals expected trade-offs between performance and efficiency. LSTM requires approximately $2.1\times$ more training time than MLP and $1.4\times$ more than CNN, while maintaining reasonable inference times suitable for practical deployment. The memory requirements scale proportionally with architectural complexity, with LSTM requiring $1.8\times$ more memory than MLP.

4. Conclusion

This study presented a comprehensive comparative analysis of three prominent deep learning architectures Multi-Layer Perceptron (MLP), Convolutional Neural Network (CNN), and Long Short-Term Memory (LSTM) for software defect prediction under identical experimental conditions. By employing a rigorous seven-phase framework, standardized NASA MDP datasets, and robust statistical evaluation, the study addressed key limitations of existing works related to unfair comparisons, limited evaluation metrics, and insufficient consideration of computational efficiency.

Experimental results demonstrate that LSTM consistently achieves the highest predictive performance across most datasets, outperforming CNN and MLP in terms of Accuracy, F1-score, ROC-AUC, and MCC. Statistical significance testing using Friedman and Wilcoxon signed-rank tests confirms

that the observed performance differences are not random, with large effect sizes favoring LSTM over both CNN and MLP. These findings indicate that modeling sequential dependencies within software metrics provides substantial benefits for defect prediction tasks.

However, the computational analysis reveals an important trade-off between performance and efficiency. While LSTM delivers superior predictive accuracy, it also incurs higher training time, inference latency, memory usage, and computational cost. In contrast, MLP offers the most lightweight and computationally efficient solution, making it suitable for resource-constrained or real-time deployment environments, whereas CNN provides a balanced compromise between accuracy and computational overhead.

Overall, this study provides practical guidelines for architecture selection in software defect prediction: LSTM is recommended when predictive performance is the primary objective, CNN is suitable for balanced performance-efficiency scenarios, and MLP remains a viable option for fast and scalable industrial applications. Future research may extend this work by exploring hybrid architectures, cost-sensitive learning strategies, and cross-project transfer learning to further improve defect prediction performance while maintaining deployment efficiency.

Conflicts of Interest

The author declares no conflict of interest.

References

- [1] E. Iannone, R. Guadagni, F. Ferrucci, A. De Lucia, and F. Palomba, "The Secret Life of Software Vulnerabilities: A {Large-Scale} Empirical Study," *IEEE Trans. Softw. Eng.*, 2023.
- [2] M. J. Hernández-Molinos, Á. Sánchez-García, R. Barrientos-Martínez, J. C. Pérez-Arriaga, and J. O. Ocharán-Hernández, "Software Defect Prediction with Bayesian Approaches," *Mathematics*, vol. 11, 2023.
- [3] N. Alnor and A. Khleel, "Software defect prediction using a bidirectional LSTM network combined with oversampling techniques," vol. 0123456789, pp. 3615–3638, 2024, doi: 10.1007/s10586-023-04170-z.
- [4] A. Taskeen, S.-U.-R. Khan, and E. A. Felix, "A research landscape on software defect prediction," *J. Softw. Evol. Process*, 2023.
- [5] L. Gong, H. Zhang, J. Zhang, M. Wei, and Z. Huang, "A Comprehensive Investigation of the Impact of Class Overlap on Software Defect Prediction," *IEEE Trans. Softw. Eng.*, 2023.
- [6] N. Grattan, D. A. da Costa, and N. Stanger, "The need for more informative defect prediction: A systematic literature review," *Inf. Softw. Technol.*, 2024.
- [7] Z. Li, H. Zhang, X.-Y. Jing, J. Xie, M. Guo, and J. Ren, "DSSDPP: Data Selection and Sampling Based Domain Programming Predictor for Cross-Project Defect Prediction," *IEEE Trans. Softw. Eng.*, vol. 49, no. 4, pp. 1941–1963, 2023, doi: 10.1109/TSE.2022.3204589.
- [8] N. A. A. Khleel and K. Nehéz, "Software defect prediction using a bidirectional {LSTM} network combined with oversampling techniques," *Cluster Comput.*, 2023.
- [9] A. Abdu, Z. Zhai, H. A. Abdo, and R. Algabri, "Software Defect Prediction Based on Deep Representation Learning of Source Code From Contextual Syntax and Semantic Graph," *IEEE Trans. Reliab.*, 2024.
- [10] Y. Al-Smadi, M. Eshtay, A. Al-qerem, S. Nashwan, O. Ouda, and A. A. A. El-Aziz, "Reliable prediction of software defects using Shapley interpretable machine learning models," *Egypt. Informatics J.*, 2023, [Online]. Available: <https://api.semanticscholar.org/CorpusID:260391021>
- [11] A. Abdu *et al.*, "Semantic and traditional feature fusion for software defect prediction using hybrid deep learning model," *Sci. Rep.*, vol. 14, 2024.
- [12] A. Alazba and H. Aljamaan, "Software Defect Prediction Using Stacking Generalization of Optimized {Tree-Based} Ensembles," *Appl. Sci.*, vol. 12, 2022.
- [13] X. Dong, Y. Liang, S. Miyamoto, and S. Yamaguchi, "Ensemble learning based software defect prediction," *J. Eng. Res.*, 2023, [Online]. Available: <https://api.semanticscholar.org/CorpusID:265011535>
- [14] B. Arasteh, K. Arasteh, A. Ghaffari, and R. Ghanbarzadeh, "A new binary chaos-based metaheuristic algorithm for software defect prediction," *Cluster Comput.*, 2024.
- [15] H. Das, S. Prajapati, M. Gourisaria, R. M. Pattanayak, A. Alameen, and M. S. Kolhar, "Feature Selection Using Golden Jackal Optimization for Software Fault Prediction," *Mathematics*, vol. 11, 2023.
- [16] Y. Zhao, K. Damevski, and H. Chen, "A Systematic Survey of Just-in-Time Software Defect Prediction," *ACM Comput. Surv.*, vol. 55, pp. 1–35, 2022, [Online]. Available: <https://api.semanticscholar.org/CorpusID:252910605>

- [17] and H. A. Alazba, Amal, “applied sciences Software Defect Prediction Using Stacking Generalization of,” 2022.
- [18] A. Abdu, Z. Zhai, R. Algabri, H. A. Abdo, K. Hamad, and M. A. Al-antari, “Deep {Learning-Based} Software Defect Prediction via Semantic Key Features of Source {Code---Systematic} Survey,” *Mathematics*, vol. 10, 2022.
- [19] M. Ali, T. Mazhar, A. Al-Rasheed, T. Shahzad, Y. Ghadi, and M. A. Khan, “Enhancing software defect prediction: a framework with improved feature selection and ensemble machine learning,” *PeerJ Comput. Sci.*, 2024.
- [20] N. S. Thomas and S. Kaliraj, “An Improved and Optimized Random Forest Based Approach to Predict the Software Faults,” *SN Comput. Sci.*, 2024.
- [21] N. A. A. Khleel and K. Nehéz, “A novel approach for software defect prediction using {CNN} and {GRU} based on {SMOTE} Tomek method,” *J. Intell. Inf. Syst.*, 2023.