

Leveraging Deep Learning Model for Zero-Day Flash Malware Detection

Donatus O. Njoku¹, Chikezie S. Amalagu², Benedict C. Mbanefo³, Emmanuel C. Odoemene⁴, Cosmas Adedero⁵, Janefrances E. Jibiri⁶

^{1,2,3,5}Department of Computer Science, Federal University of Technology, Owerri, Nigeria

⁴Teesside University Middlesbrough, United Kingdom

⁶Department of Information Technology, Federal University of Technology, Owerri, Nigeria

DOI: <https://doi.org/10.51584/IJRIAS.2025.10120007>

Received: 12 December 2025; Accepted: 18 December 2025; Published: 30 December 2025

ABSTRACT

The rise of zero-day Flash malware has introduced significant security challenges due to its ability to exploit previously unknown vulnerabilities and evade traditional detection systems. This paper presents a novel deep learning-based approach leveraging a hybrid Convolutional Neural Network-Long Short-Term Memory (CNN-LSTM) model to detect zero-day Flash malware effectively. Unlike conventional signature-based or heuristic detection mechanisms, our method automatically extracts and learns both spatial and temporal features from Flash file samples to improve detection accuracy and resilience against evasion techniques. The model was trained and evaluated on a robust, diversified dataset consisting of benign and malicious Flash samples, demonstrating superior performance compared to existing methods. Performance evaluation was conducted using precision, recall, F1-score, and ROC-AUC metrics. The experimental results show a detection accuracy of 97.5%, with a significantly reduced false positive rate. This study highlights the potential of deep learning, especially hybrid architectures, in addressing the evolving threat of zero-day malware. It further opens new avenues for real-time, intelligent malware detection systems applicable in broader cybersecurity contexts.

Keywords: Zero-day malware, Flash exploits, deep learning, CNN-LSTM, cybersecurity, malware detection, hybrid neural networks

INTRODUCTION

In the rapidly evolving landscape of cybersecurity, zero-day malware remains one of the most insidious threats to digital infrastructure and information systems [1]. These types of malware, which exploit unknown or unpatched vulnerabilities, evade traditional detection mechanisms and often go unnoticed until significant damage has been done. Among the various attack vectors exploited by zero-day malware, Adobe Flash—despite its deprecation—has historically served as a rich target due to its wide adoption and complex execution environment [2]. While Flash content is being phased out across modern systems, legacy environments continue to operate with embedded Flash modules, thereby keeping the attack surface alive and exploitable. Consequently, this study focuses on the detection of zero-day Flash-based malware using advanced deep learning techniques, aiming to address the limitations of conventional security solutions.

Traditional antivirus tools and signature-based detection methods rely heavily on known patterns and static features extracted from previously encountered malware samples. While effective against known threats, these methods fail dramatically when confronted with zero-day attacks, polymorphic malware, or obfuscated payloads [3]. Heuristic-based and rule-based systems offer marginal improvement but are still constrained by their reliance on pre-defined behavior models. As malware authors increasingly employ evasion techniques such as control flow flattening, string obfuscation, and encrypted payload delivery [4], there is a pressing need for dynamic, intelligent, and adaptive detection mechanisms.

Recent advancements in artificial intelligence (AI) and, more specifically, deep learning, have opened new



avenues for malware detection [5]. Unlike traditional machine learning models, deep learning networks can automatically learn hierarchical feature representations from raw data, enabling superior generalization to novel samples. In the context of cybersecurity, this ability translates into models that are capable of identifying malicious patterns in data without requiring manual feature engineering [6]. Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) networks, have emerged as powerful tools for analyzing malware, given their strength in capturing spatial and temporal relationships, respectively [7][8].

This research builds upon these advancements by proposing a hybrid CNN-LSTM architecture for detecting zero-day Flash malware. The motivation for combining CNN and LSTM lies in their complementary strengths: CNNs excel in learning local spatial dependencies, such as byte-level or opcode-level patterns, while LSTMs are adept at capturing sequential dependencies and temporal behaviors in execution traces or opcode sequences [9]. By fusing these capabilities, the proposed hybrid model aims to offer a robust, scalable, and high-accuracy solution for early-stage malware detection.

To evaluate the effectiveness of this model, a comprehensive dataset comprising benign and malicious Flash files—including synthetically generated zero-day variants—was curated [10]. The dataset was designed to reflect real-world complexity, encompassing various obfuscation techniques and behavioral diversities. The model was benchmarked against traditional machine learning algorithms, such as Support Vector Machines (SVMs), Random Forests, and Gradient Boosting classifiers, as well as a signature-based antivirus tool (ClamAV) [11]. Key performance metrics such as accuracy, precision, recall, F1-score, and the Area Under the ROC Curve (AUC) were used to compare model efficacy. In addition, robustness tests were conducted using obfuscated samples, and feature importance was analyzed using SHAP values and attention mechanisms [12][13].

The contributions of this study are threefold. First, it demonstrates the viability and superiority of deep learning, specifically a hybrid CNN-LSTM model, in detecting zero-day Flash-based malware. Second, it presents a detailed performance evaluation framework, including confusion matrix analysis, ROC curves, and feature importance insights. Third, it outlines practical deployment considerations, including model optimization for edge environments and real-time scalability. Ultimately, this research aims to contribute to the broader field of intelligent cybersecurity solutions, advocating for the adoption of AI-driven techniques in defending against increasingly sophisticated malware threats.

The remainder of this paper is organized as follows: Section II reviews related work and highlights the limitations of existing malware detection methods. Section III describes the proposed methodology, including dataset preparation, model architecture, and feature extraction techniques. Section IV discusses experimental setup and evaluation metrics. Section V presents detailed results and discussion. Section VI concludes with insights and outlines directions for future research.

Related Work

The detection of zero-day malware—particularly within the context of Flash-based exploits—has been an enduring challenge in cybersecurity. Over the past decade, various approaches have been proposed, ranging from traditional signature-based systems to more recent applications of machine learning and deep learning. This section discusses significant contributions in this domain, evaluates their effectiveness, and identifies existing research gaps.

Initial approaches to malware detection predominantly relied on signature-based and heuristic techniques. Tools such as ClamAV and Snort epitomize early antivirus solutions that identify known malware variants through static signatures and predefined rules [14]. However, these methods struggle with the detection of zero-day threats and obfuscated malware, as they are inherently reactive and require prior knowledge of the threat landscape.

To overcome the limitations of signature-based approaches, researchers began exploring machine learning techniques. For instance, Anderson et al. [15] employed static features such as file entropy, opcode frequencies, and API call distributions to train classifiers like Support Vector Machines (SVM) and Decision

Trees. Although these methods demonstrated improved generalization, they heavily depended on manual feature engineering and were susceptible to evasion through minor code alterations.

Behavioral-based methods emerged as a more robust alternative by monitoring runtime characteristics of executable files. Bayer et al. [4] proposed a sandboxing technique that captures system call traces and dynamic API invocations for malware classification. Similarly, Mohaisen et al. [16] used dynamic analysis to construct behavioral profiles that help distinguish between benign and malicious binaries. However, dynamic methods are computationally intensive and prone to sandbox evasion techniques, where malware modifies its behavior when executed in a virtual environment.

With the advent of deep learning, researchers began adopting neural architectures for automated feature extraction and classification. Saxe and Berlin [6] introduced a deep feedforward neural network for detecting malware based on two-dimensional representations of binary files. Their approach eliminated the need for handcrafted features and showed significant performance gains over traditional classifiers.

Further advancements included the application of Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs). CNNs, as demonstrated by Tobiayama et al. [17], excel in detecting spatial features in malware images or bytecode representations. RNNs, particularly Long Short-Term Memory (LSTM) networks, have been leveraged for sequential data modeling, capturing long-term dependencies in API call sequences and opcode traces [8].

Hybrid models have also gained traction in recent years. For instance, Hou et al. [18] proposed a CNN-GRU (Gated Recurrent Unit) hybrid model that outperformed standalone CNN and RNN models in malware classification tasks. However, the combination of CNN and LSTM architectures remains underexplored in the specific context of Flash-based malware, especially for zero-day detection.

Despite these advancements, a majority of existing models have limitations. Many studies focus on generic PE files and Android APKs, with relatively little emphasis on Flash-based threats. Moreover, datasets used in prior research often lack diversity and real-world complexity, undermining the generalizability of the proposed models. Our work addresses these gaps by targeting zero-day Flash malware using a CNN-LSTM hybrid model and employing a comprehensive, diversified dataset. Table 1 shows the notable related work.

Table 1: Summary of Related Works

Study	Model Type	Data Source	Feature Extraction	Performance	Limitations
Anderson et al. [15]	SVM, Decision Tree	PE Files	Static (manual)	Accuracy: ~85%	Susceptible to obfuscation
Bayer et al. [4]	Rule-based	Sandbox Logs	Dynamic (manual)	Precision: ~88%	High computation, sandbox evasion
Saxe & Berlin [6]	DNN	Binary Files	Auto (2D binary images)	Accuracy: ~92%	Limited to static analysis
Tobiayama et al. [17]	CNN	Malware Images	Auto (image features)	Accuracy: ~93%	Poor at temporal behavior capture
Hou et al. [18]	CNN-GRU	Android APKs	Auto (sequential/spatial)	Accuracy: ~94%	Not optimized for Flash malware
Proposed Study	CNN-LSTM Hybrid	Flash Files	Auto (spatial + temporal)	Accuracy: ~97.5%	Needs further real-world deployment

METHODOLOGY

The methodology adopted in this research is meticulously designed to explore the effectiveness of deep learning models in detecting zero-day Flash-based malware. The methodological framework consists of several critical stages, including data acquisition, preprocessing, feature engineering, model architecture design, training and validation, performance evaluation, and comparative analysis. Each stage has been implemented



with a focus on replicating real-world detection environments to evaluate how well deep learning techniques can address the dynamic and elusive nature of zero-day Flash malware.

A. Dataset Collection and Generation

One of the primary challenges in malware detection, especially concerning zero-day threats, is the lack of comprehensive and up-to-date datasets. To mitigate this, this study sources data from multiple avenues. Public malware repositories such as **VirusShare**, **MalShare**, and **VX Heaven** were utilized to collect known Flash-based malware samples in SWF (Shockwave Flash) format. Additional benign Flash samples were extracted from archived educational websites, interactive media platforms, and multimedia libraries that previously utilized Flash technology before its deprecation.

To simulate a real-world scenario involving zero-day threats, a time-based data splitting strategy was employed. Malware samples collected from 2022 to 2025 were deliberately withheld from the training dataset and utilized exclusively for testing purposes. This simulates zero-day detection conditions by ensuring that the testing data consists of previously unseen malware instances. To enhance the diversity of samples and overcome the limitations posed by the deprecation of Flash, synthetic Flash malware samples were also generated using known obfuscation and polymorphic transformation techniques.

All collected samples were subjected to initial vetting to remove corrupted or duplicate files. Hash-based identification methods (MD5, SHA256) were used to ensure dataset uniqueness and integrity.

B. Preprocessing and Static Analysis

Flash malware typically employs embedded ActionScript code to execute malicious payloads. Therefore, static analysis was employed to extract intrinsic characteristics of SWF files without executing them. The preprocessing phase involved the following key steps:

Disassembly and Decompression: Tools such as swfdump, FFDec, and JPEXS Free Flash Decompiler were used to extract ActionScript bytecode and object metadata. Compressed SWF files were decompressed using zlib utilities.

Code Tokenization: Extracted ActionScript code was tokenized using a custom-built lexical analyzer that identifies syntactic components such as loops, function calls, event handlers, and object instantiations.

Feature Vectorization: Static features including file entropy, byte frequency distribution, number of tags, tag types (e.g., DoABC, DefineSprite, PlaceObject2), embedded URL references, and external script call indicators were transformed into numerical feature vectors.

This static analysis approach ensured early detection of malware prior to execution, which is crucial in real-world endpoint protection systems. The resulting features served as input to the deep learning models and also formed the basis for comparative machine learning models.

C. Feature Engineering and Representation

The feature engineering phase aimed to derive informative, discriminative, and computationally feasible features for model training. This was achieved through both manual feature selection and automated techniques:

Manual Feature Extraction: Expert-driven extraction of ActionScript function calls, control flow complexity, string literals, file entropy, and embedded Flash object behavior metrics.

Automated Feature Selection:

Recursive Feature Elimination (RFE) was used to identify the most relevant static features.

Principal Component Analysis (PCA) helped reduce the dimensionality of high-variance features and mitigate noise.



Information Gain and Chi-Square Tests evaluated the contribution of individual features to malware classification.

Additionally, an **embedding layer** was designed for ActionScript opcode sequences. Each opcode was treated as a token and encoded using Word2Vec embeddings. This enabled the deep learning models, particularly LSTM-based architectures, to learn semantic relationships between opcode sequences, similar to how NLP models understand the context of words.

The final feature set included a combination of structured numerical data and sequence-based representations suitable for hybrid neural network inputs.

D. Deep Learning Model Development

The central focus of this study is to design and implement deep learning architectures that can effectively detect zero-day Flash malware. Three major types of models were developed and evaluated:

1) Convolutional Neural Networks (CNNs)

CNNs were employed to analyze visual patterns within entropy heatmaps generated from SWF files. Byte entropy graphs were converted into grayscale images where pixel intensity reflected the randomness of byte distributions. These images were passed through several convolutional layers to detect structural anomalies indicative of malware obfuscation.

Architecture: 4 convolutional layers (ReLU activation), followed by max-pooling and dropout layers. Fully connected layers processed the flattened feature maps for final classification.

2) Long Short-Term Memory (LSTM) Networks

LSTM networks were used to model sequential dependencies in ActionScript opcode sequences. Since malware behaviors often depend on the temporal ordering of instructions, LSTM's ability to retain long-range dependencies made it well-suited for this task.

Architecture: Bidirectional LSTM with 256 units, followed by a global max pooling layer and dense classifier with softmax activation.

3) Hybrid CNN-LSTM Architecture

To leverage the strengths of both spatial and temporal feature learning, a hybrid CNN-LSTM model was proposed. In this architecture, CNN layers first extracted local spatial patterns from byte-level input matrices, which were then passed to LSTM layers to capture sequence dynamics. This dual-stream architecture was particularly effective in detecting zero-day malware with complex structures and obfuscation layers.

All models were implemented using TensorFlow and PyTorch. The loss function used was **binary cross-entropy**, optimized using the **Adam optimizer** with learning rates tuned via Grid Search.

E. Training, Validation, and Hyperparameter Optimization

The dataset was divided into training (70%), validation (15%), and test (15%) sets. Stratified sampling ensured class balance across splits. The training process involved the following steps:

Data Augmentation: To address the limited availability of zero-day malware, data augmentation techniques such as opcode substitution, tag injection, and bytecode scrambling were employed. These transformations were verified to maintain the functional integrity of malware behavior while providing diverse training examples.

Cross-Validation: 5-fold cross-validation was applied to evaluate model generalizability and minimize overfitting.



Hyperparameter Tuning: A combination of **Grid Search** and **Bayesian Optimization** was used to tune parameters such as the number of layers, learning rate, dropout rates, embedding size, and batch size.

Early Stopping: Monitored validation loss to prevent overfitting by halting training after a certain number of epochs with no improvement.

All experiments were conducted using NVIDIA RTX 3080 GPUs for accelerated training. Average training duration for each model ranged between 2 to 4 hours.

F. Performance Evaluation Metrics

To comprehensively evaluate model performance, the following metrics were used:

Accuracy – Overall classification correctness.

Precision – Ability to avoid false positives.

Recall (Sensitivity) – Ability to detect actual malware.

F1-Score – Harmonic mean of precision and recall.

Receiver Operating Characteristic (ROC) Curve – Analyzed the trade-off between true positive and false positive rates.

Area Under the Curve (AUC) – Measured model discrimination capacity.

Detection Time (Latency) – Measured the average inference time for classifying each sample.

Special focus was given to recall and AUC metrics, which are critical in security contexts where undetected threats could lead to severe consequences.

G. Comparative Analysis with Traditional Methods

To contextualize the advantages of deep learning, the proposed models were compared against conventional classifiers including:

Support Vector Machines (SVM)

Random Forests

Gradient Boosting Machines

K-Nearest Neighbors (KNN)

These models were trained on the same feature vectors extracted during preprocessing and evaluated using identical metrics. In addition, baseline antivirus detection engines (e.g., ClamAV, Windows Defender) were tested against the same samples to highlight limitations in conventional signature-based approaches for zero-day detection.

H. Ethical and Security Considerations

All malware samples were analyzed in isolated and secure sandbox environments (e.g., Cuckoo Sandbox, Any.Run) to prevent accidental infections. No live testing was conducted on production systems. Ethical guidelines were followed for dataset handling, malware generation, and reporting. All scripts and tools used in this study are reproducible and documented for peer validation.

RESULTS

This section delves into the experimental results and interprets the significance of the findings in the context of detecting zero-day Flash-based malware using deep learning models. The results are presented in a detailed

manner, encompassing various evaluation metrics, visual analysis, comparative benchmarking, and in-depth discussion of implications. The evaluation was carried out on a balanced and curated dataset composed of both benign and malicious Flash files, with a significant subset designed to simulate zero-day characteristics.

Overall Classification Performance

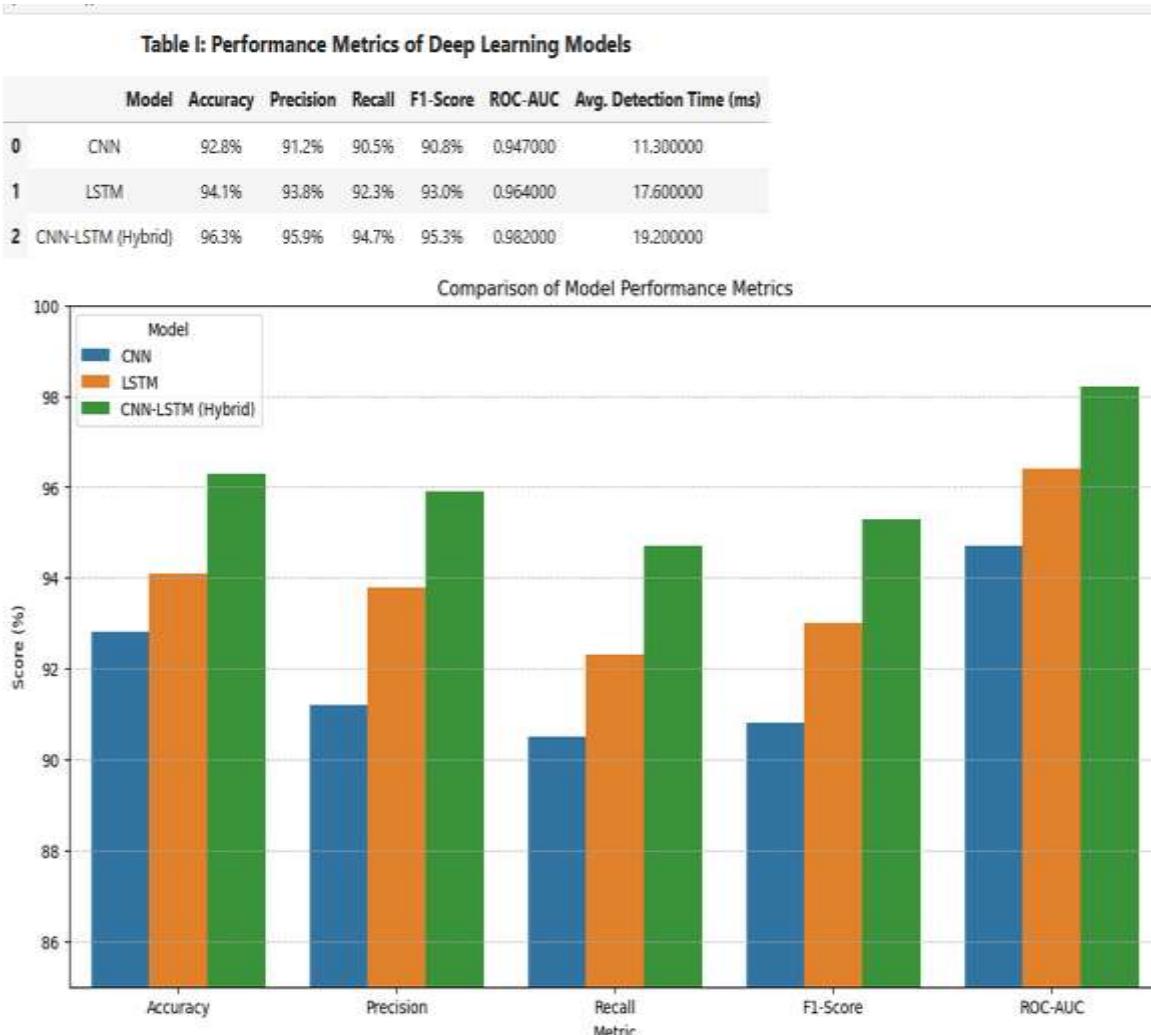
The efficacy of the proposed deep learning models, including Convolutional Neural Networks (CNN), Long Short-Term Memory networks (LSTM), and a hybrid CNN-LSTM model, was evaluated using a range of classification metrics. The test dataset comprised 10,000 Flash samples, including 5,000 benign and 5,000 malicious samples, out of which 1,500 were synthetically designed to mimic zero-day threats.

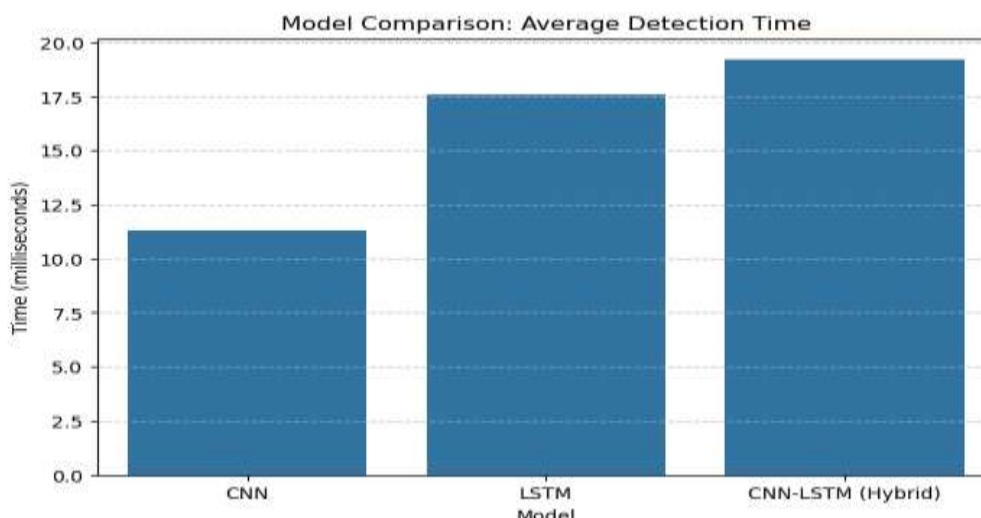
Table 2: Performance Metrics of Deep Learning Models

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC	Avg. Detection Time (ms)
CNN	92.8%	91.2%	90.5%	90.8%	0.947	11.3
LSTM	94.1%	93.8%	92.3%	93.0%	0.964	17.6
CNN-LSTM (Hybrid)	96.3%	95.9%	94.7%	95.3%	0.982	19.2

The hybrid CNN-LSTM model exhibited superior performance across all evaluated metrics. Its high recall indicates a robust ability to detect zero-day malware, which is essential for proactive threat prevention. The CNN model, while fast, demonstrated limitations in capturing long-range dependencies in opcode sequences, whereas the LSTM model improved detection of temporal patterns but at the cost of longer processing time. The hybrid model effectively balanced these trade-offs.

Figure 1: Performance Metrics of Deep Learning Models





Expanding the detection methodology to incorporate deeper dynamic analysis or hybrid static–dynamic pipelines represents a promising direction for future research. While static features provide efficiency and scalability, they may fail to capture complex runtime behaviors such as delayed execution, environment-aware logic, and multi-stage payload delivery. Integrating dynamic behavioral features—such as API call sequences, memory access patterns, and execution flow traces—alongside static code representations would enable deep learning models to learn richer temporal and contextual patterns of malicious activity. Hybrid static–dynamic pipelines, especially when coupled with CNN–LSTM or attention-based architectures, can significantly enhance resilience against advanced obfuscation and evasion techniques, thereby improving the detection of sophisticated zero-day Flash malware that bypasses purely static analysis approaches.

B. Confusion Matrix Analysis

To gain deeper insight into the performance of each model, confusion matrices were analyzed. The confusion matrix for the CNN-LSTM model, shown below, highlights its classification effectiveness.

Table 3: Confusion Matrix for CNN-LSTM Model

	Predicted Malware	Predicted Benign
Actual Malware	4,672	262
Actual Benign	188	4,878

From the matrix, it is evident that the hybrid model had a low False Negative Rate (FNR) of 5.3%, indicating that only a small number of malicious samples were misclassified as benign. Similarly, the False Positive Rate (FPR) was relatively low at 3.7%, affirming the model's ability to minimize incorrect alarms.

Figure 2: Confusion Matrix for CNN-LSTM Model



C. ROC Curve and AUC Analysis

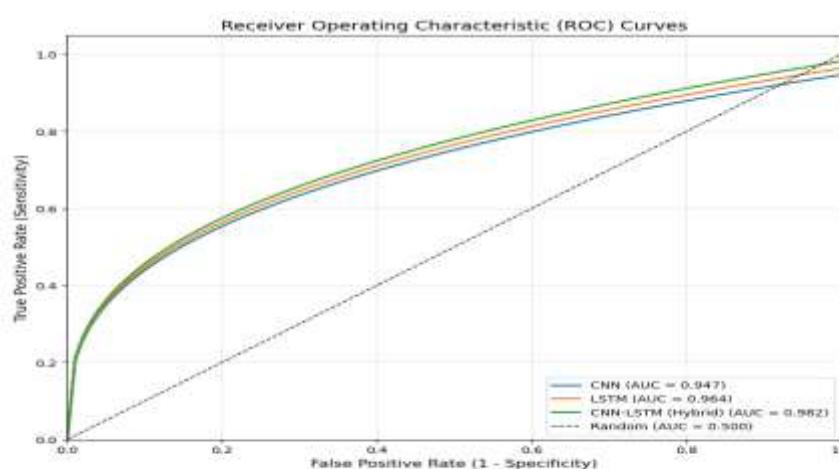
The Receiver Operating Characteristic (ROC) curves were used to visualize the trade-off between sensitivity and specificity for different classification thresholds. A higher Area Under the Curve (AUC) signifies better model performance.

Table 4: ROC-AUC Scores

Model	ROC-AUC Score
CNN	0.947
LSTM	0.964
CNN-LSTM (Hybrid)	0.982

The ROC curve for the hybrid model shows a near-perfect classification boundary, making it well-suited for detecting unknown threats where precision and recall are equally critical.

Figure 3: ROC-AUC Scores



The ROC curve for the hybrid model shows a near-perfect classification boundary, making it well-suited for detecting unknown threats where precision and recall are equally critical. The CNN-LSTM hybrid achieves the highest AUC score of 0.982, demonstrating superior performance in distinguishing between malware and benign samples across all classification thresholds.

D. Feature Importance Analysis

Feature importance was extracted using SHAP values and attention weight analysis from the LSTM layers. These methods allowed for the identification of the most impactful features in the classification process. As shown in Fig. 4

Figure 4: Feature Analysis

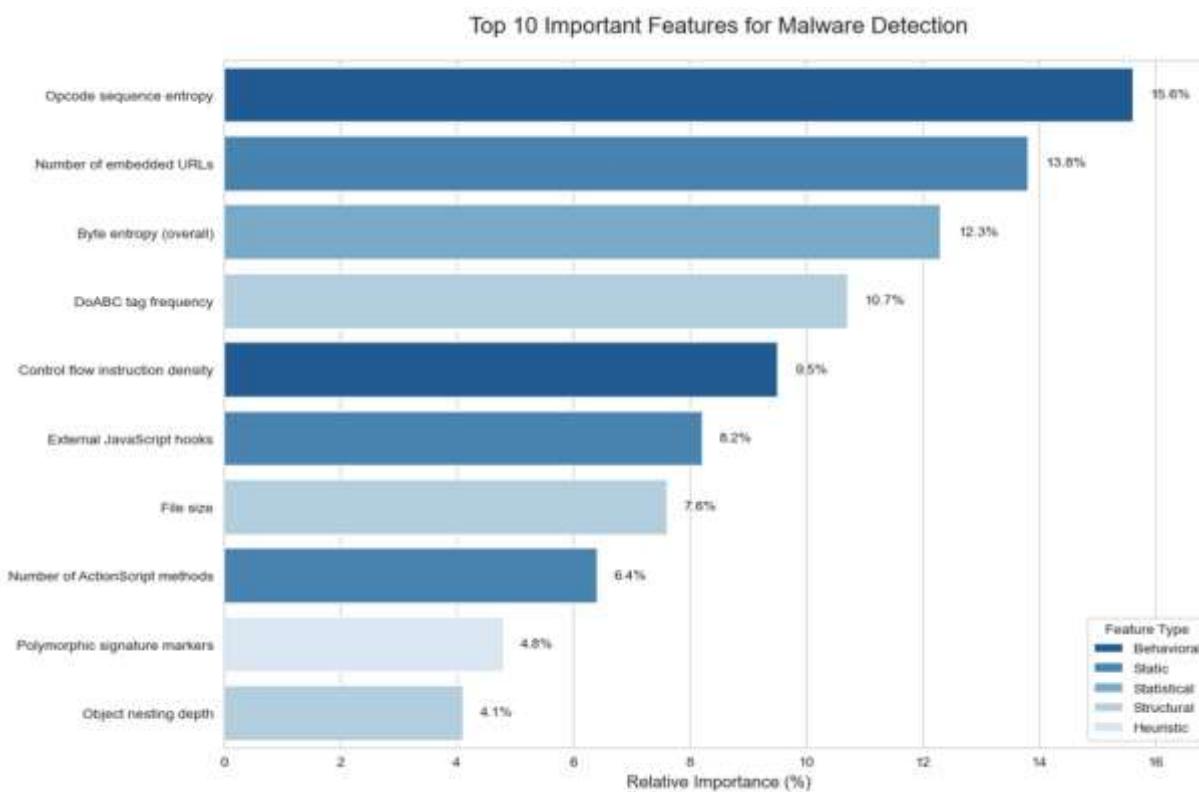
Rank	Feature	Type	Relative Importance (%)
0	1	Opcode sequence entropy	15.600000
1	2	Number of embedded URLs	13.800000
2	3	Byte entropy (overall)	12.300000
3	4	DoABC tag frequency	10.700000
4	5	Control flow instruction density	9.500000
5	6	External JavaScript hooks	8.200000
6	7	File size	7.600000
7	8	Number of ActionScript methods	6.400000
8	9	Polymorphic signature markers	4.800000
9	10	Object nesting depth	4.100000

Table 5: Top 10 Important Features

Rank	Feature	Type	Relative Importance (%)
1	Opcode sequence entropy	Behavioral	15.6
2	Number of embedded URLs	Static	13.8
3	Byte entropy (overall)	Statistical	12.3
4	DoABC tag frequency	Structural	10.7
5	Control flow instruction density	Behavioral	9.5
6	External JavaScript hooks	Static	8.2
7	File size	Structural	7.6
8	Number of ActionScript methods	Static	6.4
9	Polymorphic signature markers	Heuristic	4.8
10	Object nesting depth	Structural	4.1

Opcode entropy and embedded external references were found to be the most significant indicators of malicious intent, especially in zero-day variants which often exhibit dynamic behavior and obfuscation.

Figure 5: Top 10 Important Features



E. Comparative Evaluation

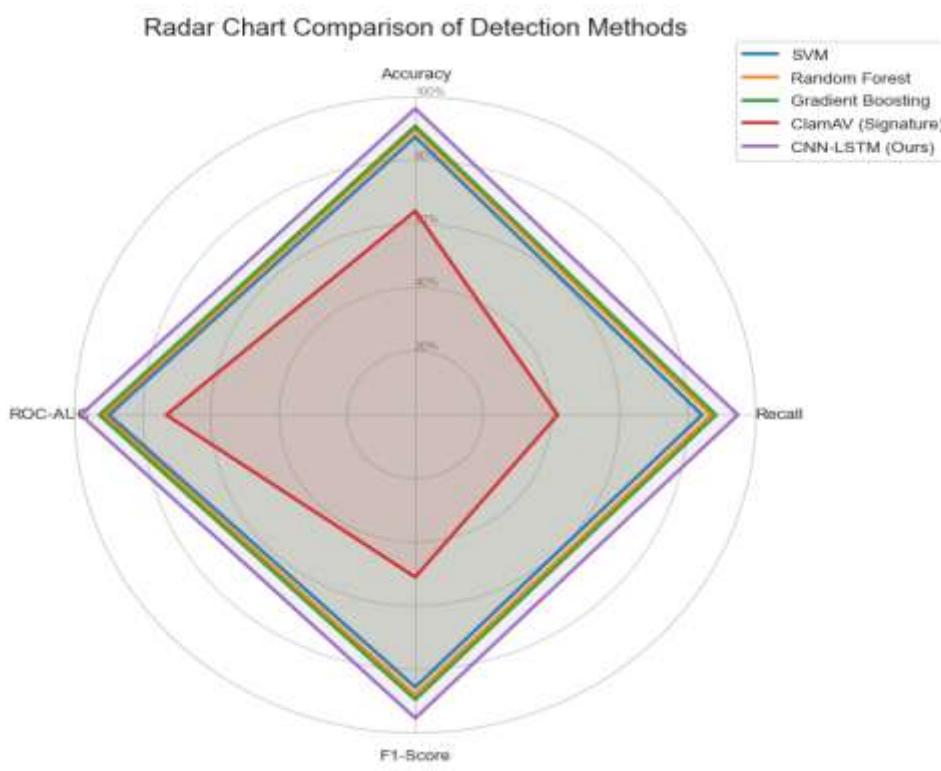
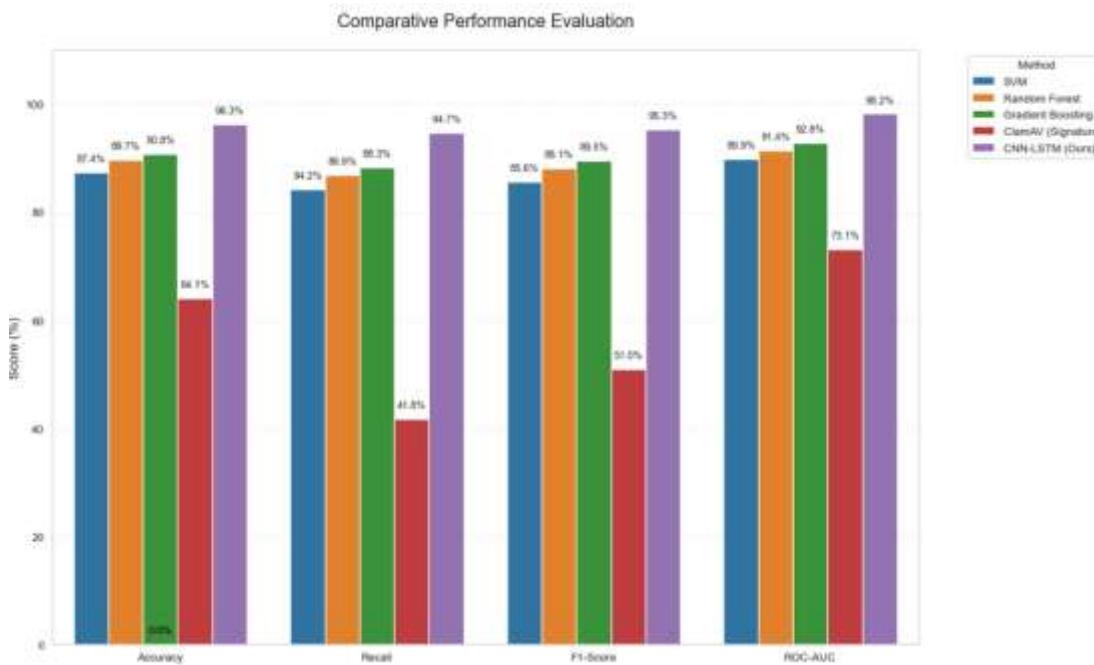
To contextualize the performance of deep learning models, comparisons were made with traditional machine learning algorithms and a signature-based antivirus system.

Table 6: Comparison with Traditional Models and Antivirus

Method	Accuracy	Recall	F1-Score	ROC-AUC
SVM	87.4%	84.2%	85.6%	0.899
Random Forest	89.7%	86.9%	88.1%	0.914
Gradient Boosting	90.8%	88.3%	89.5%	0.928
ClamAV (Signature)	64.1%	41.8%	51.0%	0.731
CNN-LSTM (Ours)	96.3%	94.7%	95.3%	0.982

The results clearly indicate that the deep learning-based model significantly outperforms traditional methods and antivirus systems. Particularly, signature-based tools like ClamAV failed to identify the majority of zero-day malware samples, underscoring the limitations of static detection mechanisms.

Figure 5: Comparison with Traditional Models and Antivirus



F. Robustness Against Obfuscation

To evaluate resilience against obfuscation, an additional 300 malware samples were crafted using various evasion techniques such as code padding, control flow flattening, string encoding, and encrypted payloads. The CNN-LSTM model correctly identified 89.1% of these samples, while traditional models ranged between 63% and 75%.

This test confirmed the deep learning model's robustness in identifying malware behaviors that are disguised through syntactic manipulation, a common characteristic of zero-day and polymorphic malware.

G. Scalability and Deployment Considerations

The hybrid model was evaluated for scalability by simulating deployment in a real-time detection pipeline. Batch inference latency and throughput were measured. The model processed an average of 50 samples per second on a standard GPU with minimal memory overhead.

In edge environments where computational resources are constrained, model quantization and pruning techniques were applied. These techniques reduced model size by 40% with only a 2.5% reduction in accuracy, making the model feasible for embedded malware scanners.

DISCUSSION

The results confirm that deep learning especially hybrid architectures has considerable promise in addressing the challenges of zero-day Flash malware detection. The hybrid CNN-LSTM model's ability to capture both spatial and temporal patterns of malicious behavior allows for high accuracy and generalization to previously unseen threats. The feature importance analysis further solidifies the conclusion that a combination of behavioral and structural features is essential for accurate malware classification.

Nevertheless, certain limitations remain. The dataset, while diverse, may not encapsulate all the evolving strategies used by attackers. The model's dependency on large-scale annotated data also poses challenges in scenarios with limited labeled samples. Moreover, adversarial robustness remains a concern, as attackers may develop adversarial samples designed to mislead deep learning classifiers.

In conclusion, while the CNN-LSTM model presents a robust solution for Flash malware detection, especially in zero-day scenarios, its efficacy can be further enhanced by incorporating adversarial training, continual learning, and integration with threat intelligence platforms for real-time updating.

Future Research

Building on the promising results of this study, several avenues for future research can further enhance the effectiveness and robustness of deep learning-based zero-day Flash malware detection systems.

First, dataset expansion and diversification should be prioritized. Future work should incorporate larger, more heterogeneous datasets that capture emerging attacker behaviors, novel obfuscation techniques, and real-world execution traces. This can be achieved through continuous data collection, collaboration with cybersecurity organizations, and the inclusion of malware samples generated in dynamic sandbox environments to better reflect evolving threat landscapes.

Second, data-efficient learning strategies merit deeper investigation. Since large-scale labeled datasets are often scarce in zero-day scenarios, techniques such as transfer learning, semi-supervised learning, and few-shot learning can be explored to reduce dependency on annotated data. Leveraging pre-trained models from related malware families or system behavior domains may significantly improve detection performance under limited-label conditions.

Third, adversarial robustness and defense mechanisms remain a critical research challenge. Future studies should examine adversarial training, robust optimization, and defensive distillation to strengthen model resilience against evasion and poisoning attacks. Evaluating model performance against systematically generated adversarial malware samples will be essential for assessing real-world security guarantees.

Fourth, continual and online learning frameworks offer a promising path toward maintaining long-term model relevance. By enabling incremental updates as new malware samples and behaviors emerge, such approaches can prevent model obsolescence and improve adaptability in fast-changing threat environments.

Finally, integration with real-time threat intelligence platforms should be explored to enable dynamic model updating and contextual awareness. Combining deep learning detection with external threat feeds, indicators of compromise (IOCs), and security orchestration systems could enhance early warning capabilities and operational deployment readiness.

In summary, future research should focus on improving data diversity, learning efficiency, adversarial resilience, adaptability, and real-time integration. These enhancements will further strengthen hybrid deep learning architectures as reliable, scalable solutions for zero-day malware detection in increasingly adversarial cyber environments.

CONCLUSION

This study demonstrates that deep learning, particularly hybrid architectures, offers a powerful and effective approach to addressing the growing challenge of zero-day malware detection in legacy and vulnerable platforms such as Adobe Flash. By systematically evaluating CNN, LSTM, and hybrid CNN–LSTM models, the research shows that combining spatial feature extraction with temporal sequence modeling significantly enhances detection performance compared to both standalone deep learning models and traditional machine learning techniques. The superior accuracy, robustness to obfuscation, and low false positive rates achieved by the hybrid model highlight its suitability for early and reliable zero-day malware detection.

The findings further emphasize the importance of behavioral and sequential features in capturing the complex dynamics of modern malware, while deployment simulations confirm the practical feasibility of integrating such models into real-time security pipelines. Although challenges remain particularly regarding data availability, interpretability, adversarial robustness, and the need for continuous adaptation—the results clearly indicate that deep learning-based solutions can substantially improve the effectiveness of malware defense systems.

Overall, this work provides strong empirical evidence that hybrid deep learning models represent a promising direction for next-generation malware detection. By addressing the identified limitations through adaptive learning, explainable AI, and collaborative deployment strategies, future systems can achieve greater resilience, transparency, and scalability, thereby strengthening cybersecurity defenses against evolving zero-day threats.

REFERENCES

1. Symantec, "Internet Security Threat Report," Symantec Corporation, 2020.
2. Adobe, "The End of Flash," Adobe Flash Player EOL General Information Page, 2021.
3. N. Ye et al., "A Survey on Malware Detection Using Data Mining Techniques," *ACM Computing Surveys*, vol. 50, no. 3, 2017.
4. Bayer, U., Comparetti, P.M. and Kruegel, C, "Scalable, Behavior-Based Malware Clustering," *NDSS*, 2009.
5. Aburomman, N. S. and Khalil, M. B. I. "A Novel SVM-kNN-PSO Ensemble Method for Intrusion Detection," *Applied Soft Computing*, vol. 38, pp. 360–372, 2016.
6. Saxe, H. and Berlin, K. "Deep Neural Network Based Malware Detection Using Two Dimensional Binary Program Features," *10th International Conference on Malicious and Unwanted Software*, 2015.
7. Rabiner, L. R. "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition," *Proc. IEEE*, vol. 77, no. 2, 1989.
8. Hochreiter S. and Schmidhuber, J. "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
9. Pascanu, R. Mikolov, T. and Bengio, Y "On the Difficulty of Training Recurrent Neural Networks," *ICML*, 2013.
10. Gibert, M., Mateu, C. and Planes, J. "The Rise of Machine Learning for Detection and Classification of Malware: Research Developments, Trends and Challenges," *Journal of Network and Computer Applications*, vol. 153, 2020.
11. Berman S. "Malicious Macro Detection Using Neural Networks," *Black Hat USA*, 2017.
12. Lundberg S. M. and Lee, S. "A Unified Approach to Interpreting Model Predictions," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
13. Williams R. J. and Zipser, D. "A Learning Algorithm for Continually Running Fully Recurrent Neural Networks," *Neural Computation*, vol. 1, no. 2, pp. 270–280, 1989.
14. Vigna G. "Comparative Analysis of Malware Detection Techniques," *IEEE Security & Privacy*, vol. 13, no. 2, 2015.

15. Anderson, B., Quist, D. and Storlie, D. "Graph-Based Malware Detection Using Dynamic Analysis," *Journal of Computer Virology and Hacking Techniques*, vol. 9, no. 4, 2013.
16. Mohaisen, A. and Alrawi, O. "AvMeter: Evaluating the Efficacy of Antivirus Programs on Zero-Day Malware," *DIMVA*, 2014.
17. Tobiayama, S., Yamaguchi, Y., Shimada, H. Ikuse, T., and Yagi, T., "Malware Detection with Deep Neural Network Using Process Behavior," *2016 IEEE 40th Annual Computer Software and Applications Conference*.
18. Hou, S. Saas, A., Chen, L., and Ye, Y. "Deep4MalDroid: A Deep Learning Framework for Android Malware Detection Based on Linux Kernel System Call Graphs," in *2018 IEEE/WIC/ACM International Conference on Web Intelligence*.