

Deploying Deep Learning Models for Adaptive Malware Classification and Behavioral Threat Analysis in Evolving Cyber Landscapes

Marcel Okoebor¹

Publication Date: 2025/10/30

Abstract

The synthesised findings of five recent studies on the implementation of deep learning (DL) models to achieve adaptive malware classification and behavioural threat analysis in dynamic cyber environments are presented in this empirical review. The reviewed research shows that DL methods, convolutional neural networks (CNNs), recurrent neural networks (RNNs), transformers, and hybrid approaches are significantly more efficient tools in terms of accuracy, feature extraction, and adaptation to the complexity of malware behaviors than traditional machine learning methods. Hybrid models that merge supervised, unsupervised, and federated learning are even better at improving the scalability and privacy preservation. Despite these developments, several challenges persist, including high computational costs, limited access to high-quality labeled data, susceptibility to adversarial attacks, the need for explainability, and the incorporation of these systems into real-world settings. Research gaps were noted in creating lightweight IoT and edge computing systems, adversarially robust DL models, explainable AI, and continual learning systems. The article review generally confirms the instrumental role of DL in developing adaptive cybersecurity strategies, emphasizing the need for innovative, interdisciplinary, and practical implementations that match the sophistication and practicality required in a dynamic digital environment.

Keywords: Deep Learning, Malware Classification, Cybersecurity, Federated Learning, Adversarial Robustness, Explainable AI, IoT Security, Edge Computing.

I. INTRODUCTION

The issue of high-speed cyber threats increases the challenge of digital security systems, which have to face new and advanced methods of detection by malicious actors (Okoli et al., 2024). In more detail, the malware has become more flexible, exploiting such distributed systems as cloud, edge, and IoT (Gulatas et al., 2023). Zero-day attacks and polymorphic malware cannot always be detected using conventional signature-based and heuristic techniques. Actually, the industry demands more dynamic and intelligent methods of threat analysis (Kaur & Singh, 2014). Deeper learning has come to be recognised as an innovative mechanism in cybersecurity with improved functions in malware classification and behavioural threat analysis (Okoli et al., 2024). Its potential to auto learn complex patterns in high-dimensional data makes it well applicable to inherent subtle anomalies in static feature sets, dynamic feature sets, as well as hybrid feature sets. In contrast to classic machine learning, deep learning is less dependent on manual features and more flexible in terms of recognizing patterns and adaptability--in particular, it can facilitate real-time adaptations to new adversarial

strategies and maintain resilience. As a result, the use of CNNs, RNNs, transformers, and hybrids has become the subject of increased research (Alomar, Aysel & Cai, 2024).

The purpose of this review is to examine empirical evidence from recent studies on the deployment of deep learning models for adaptive malware classification and behavioural threat analysis. The scope covers five selected studies published between 2024 and 2025, which collectively investigate diverse methodologies, datasets, and applications. These studies encompass hybrid learning models, adversarial resilience, federated approaches, and large-scale empirical testing in distributed systems. By synthesising their findings, the review highlights performance outcomes, key contributions, limitations, and research gaps, thereby providing insight into the evolving role of deep learning in strengthening cybersecurity.

II. REVIEW METHODOLOGY

This review draws on five peer-reviewed studies published between 2024 and 2025, sourced from reputable

journals and research repositories such as the *World Journal of Advanced Research and Reviews*, *Frontiers in Physics*, and ResearchGate. The selection process focused on works that presented empirical investigations or systematic analyses of deep learning models for malware classification and behavioural threat detection. The inclusion criteria required studies to (i) fall within the 2024–2025 timeframe, (ii) employ or review deep learning approaches such as CNNs, RNNs, hybrid and ensemble methods, and (iii) report experimental results or detailed

evaluations using benchmark or real-world datasets. Studies that were purely conceptual without reference to empirical evidence, or outside the scope of cybersecurity threat detection, were excluded. The resulting corpus reflects a balance of experimental, hybrid, and review-oriented research, thereby providing a comprehensive foundation for synthesising empirical findings in evolving cyber landscapes.

➤ *Summary of Characteristics of the Selected Studies*

Table 1 Summary of Characteristics of the Selected Studies

#	Criterion	Study 1	Study 2	Study 3	Study 4	Study 5
1	Bibliographic Details	Chukwuani et al. (2025). Machine Learning Techniques for Real-Time Malware Classification and Threat Detection in Distributed Systems. <i>World Journal of Advanced Research and Reviews</i> .	Karki et al. (n.d). Cybersecurity with Deep Learning: Malware Detection through Real-Time Pattern Recognition... ResearchGate.	Sharma (2024). Machine Learning-Driven Approaches for Contemporary Cybersecurity... Nuvern Machine Learning Reviews.	Redhu, A., Choudhary, P., Srinivasan, K., Das, T.K. (2024). Deep learning-powered malware detection in cyberspace: a contemporary review. <i>Front. Phys.</i> 12:1349463.	Okafor, M.O. (2024). Deep learning in cybersecurity: Enhancing threat detection and response. <i>WJARR</i> , 24(03), 1116–1132.
2	Objective / Research Problem	Real-time malware classification & threat detection in distributed systems (IoT, cloud, edge).	Improve malware detection with DL; address evasion, real-time pattern recognition, collaborative intelligence.	Survey ML in intrusion detection, malware classification, and incident response.	Review DL models for malware detection; identify strengths, limitations, future directions.	Explore DL models (CNN, RNN, hybrids) in real-time detection; address adversarial & scalability issues.
3	Dataset Used	CICIDS, EMBER, custom ICS logs. Features: hybrid (static, dynamic, behavioral).	Not explicitly named; grayscale images, byte sequences, API logs.	Not specified; references public datasets and real-world logs.	Not specified; references Drebin, NSL-KDD, IoT-23, etc.	CICIDS2017, CERT, PhishTank.
4	Deep Learning Model / Methodology	Hybrid: RF, SVM, GBM + CNN, LSTM, Transformer. Supervised+unsupervised; federated learning.	CNN (images), RNN/LSTM (sequences), hybrid; adversarial training.	CNN, RNN, autoencoders, ensembles, RL. Supervised+unsupervised+RL.	RNNs, DAEs, LSTMs, DNNs, DBNs, CNNs, Generative Models, DRL, Attention models.	Hybrid CNN-RNN; adversarial training; gradient masking.
5	Evaluation Metrics & Results	Accuracy>96%, AUC=0.984, low FPR.	CNN up to 97.8%; LSTM 93.4%. DL > ML.	Not quantified; qualitative improvements (fewer false positives, faster response).	Acc, Prec, Rec, F1, AUC, Evasion. CNN: 94–99%, Hybrid: up to 99.9%.	Acc, Prec, Rec, F1, AUC. Hybrid CNN-RNN: 96.8% vs CNN 95.6% (CICIDS2017).
6	Key Findings / Contributions	Hybrid models + federated learning enable privacy-preserving, scalable detection.	DL > ML; hybrid & transfer improve robustness.	ML enhances accuracy, reduces manual effort, integrates threat intelligence.	DL excels in feature extraction, adaptability, real-time detection; hybrids improve robustness.	Hybrid CNN-RNN captures spatial & temporal features; adversarial training improves resilience.

7	Limitations Identified	Labeling inconsistencies, concept drift, adversarial evasion, high computational cost.	Adversarial vulnerability, computational overhead, legacy integration, explainability gaps.	General survey—no specific limitations discussed.	Data imbalance, interpretability, adversarial vulnerability, scalability, deployment challenges.	High computational cost, interpretability issues, adversarial attacks, scalability at IoT edge, privacy concerns.
---	------------------------	--	---	---	--	---

➤ *Empirical Evidence and Comparative Synthesis of Deep Learning for Malware Classification and Adaptive Threat Detection*

Examples in the five selected works show empirically that the field of deep learning has moved beyond hype to actual tools of use in the malware classification process when the model architecture is adjusted appropriately based on feature modality and constraints of deployment.

First, there are the approaches to static analysis, where models take a target artefact, such as bytecode, opcodes or PE-header characteristics. Karki et al. (n.d) indicate comparatively high accuracies of convolutional neural networks when either malware binaries are mapped as grayscale images or the byte sequences are utilized as such whereby the CNN can seek accuracies up to 97.8 percent equivalent to those of sequence-based LSTMs which can reach 93.4 percent with identical input. This is supported by Redhu et al. (2024), which summarises numerous empirical studies reporting 94-99% by CNNs on careful curated static datasets (e.g. Drebin-style datasets). Although Sharma (2024) considers much wider range of evidence, the qualitative evidence credits faster response times and fewer false positives to deep models in which representations learn directly on underneath the static artefacts, limiting the use of the derailing handcrafted features. This lack of sensitivity to additions and packaging comes at the cost of brittleness to additions and packaging, obscurity and concept drift: Static pipelines are thus good at throughput and reproducibility and poor at being sensitive to changes in additions and packaging.

"Moving to dynamic/behavioural analysis deep sequence models capture time-based consistent patterns in API (application programming interface) calls, system calls, and the run-time interaction graph. Hukwuani et al. (2025) combine behavioural telemetry via CICIDS, EMBER and bespoke ICS logs, testing LSTMs, Transformers and classical baselines; they achieve prediction accuracy of over 96% and AUC = 0.984 with low false-positive rates, indicating good discriminative power in conditions where temporal context is used. Karki et al. (2022) also note the effectiveness of the RNN/LSTM models in the context of API sequences and encourage the use of adversarial training to strengthen behaviour-based detectors against the ejection-based on feedback. The results tabulated in Okafor (n.d), reveal that temporal modelling does indeed matter in practice: the hybrid CNN-RNN model achieves a higher result (96.8%) compared to a pure CNN (95.6%) on CICIDS2017, when taken to

indicate the incremental usefulness of capturing both the spatial and the temporal footprints of malicious activity. Sharma (2024) additionally notes that behavioural models tend to minimize false positives and indeed enjoy greater flexibility and resilience towards zero-day behaviours compared to purely static pipelines; but Karki et al. (n.d) and Okafor (2024) warn that they present higher computational and integration costs as well as decreased resilience to conflicting or incompletely captured runtime events.

The most consistent empirical indication throughout the corpus is in hybrid and ensemble models, a combination of static and dynamic features and / or deep and classical learners. Hukwuani et al. (2025) merge RF, SVM, GBM with CNNs, LSTMs, and Transformers in both supervised and unsupervised environments and find that feature-based fusion is more precise-recall-balanced, whereas federated learning maintains privacy under distributed environments. Reviewing hybrids that drive overall scores to the highest end of the scale-up to 99.9 percent-Redhu et al. (2024) comment on how diversity (both in features and architecture) helps diminish overfitting and enhance resistance to prevalent obfuscation tactics. Okafor (2024) augments this by incorporating adversarial training and gradient masking (with inference in sparse settings), and this hybrid of CNN-RNN architecture actually improves detection without incurring an excessive latency cost in a strongly equipped environment. The findings are that ensembles and hybrids are always superior to single-modality models when tested on mixed or heterogeneous telemetry but introduce a complexity factor that casts doubts on the explainability and operational cost.

The issue of non-stationarity of the malware ecosystems is addressed explicitly in the empirical work concerning adaptive and evolving threat detection. Concept drift in distributed/IoT/cloud/edge systems can be addressed by use of federated learning, enabling to continuously update the model with local data, yet not to violate privacy: Hukwuani et al. (2025). Karki et al. (n.d) place more emphasis on transfer learning, giving evidence that pre-training on historic bigger language corpora, and fine-tuning on latest samples, enhances generalisation to new families and different forms. Redhu et al. (2024) also point to the need of continual learning as the direction to sustain the performance of the prolonged drift against the adversarial pressure. Longitudinal testing is skewed: CICIDS2017, CERT and PhishTank (Okafor, 2024) provide base diversity, but neither Chukwuani et al (2025)

custom logging of ICS logs nor use of multi-domain telemetry can hope to match changed distributions faced in production. Chukwuani et al. (2025) and Okafor, (2025) are most concrete when it comes to fielding, framing the challenge of real-time detection in distributed systems, with Chukwuani et al. (2025) highlighting the need to do so in a privacy-preserving way and Okafor, (2025) considering limitations at both the IoT edge (compute, energy, latency) and requirements of adversarial resilience.

A cross-analysis of the different datasets, metrics, strengths, and limitations refines these cues. Datasets include benchmark corpora (CICIDS, EMBER, Drebin, NSL-KDD, IoT-23) and operational logs (ICS, CERT, PhishTank), although Karki et al. (n.d), Sharma (2024), and Redhu et al. (2024) occasionally cite them in non-specific terms, rather than reporting specific sample sizes. In cases where the metrics are available, we find Accuracy > 96% and AUC = 0.984 on hybrid behavioural pipelines (Chukwuani et al., 2025), CNN = 0.94-0.99 in a static setting and hybrids up to 0.999 (Redhu et al., 2024), and CNN-RNN 0.968 compared with CNN 0.956 on CICIDS2017 (Okafor, 2024). As observed by Sharma (2023), Karki et al. (n.d), the response is quicker and the false positives are reduced, which corresponds with reports on DL > ML baselines. Overall, the most convincing empirical performances come through (i) high accuracy of both static and behavioural features, (ii) enhanced adaptability with the use of temporal/sequential context, (iii) robustness through hybridisation and adversarial training, and (iv) privacy-aware scaling using federated learning in distributed environments.

In counterbalance, there are certain constraints to external validity and operationalisation recurrently. Generalisation is challenged by bias and imbalance in the datasets (Redhu et al., 2024) and inconsistencies in labelling (Chukwuani et al., 2025) and a number of studies report adversarial vulnerability after defensive training (Karki et al., n.d; Redhu et al., 2024; Okafor, 2024). Compute overhead and latency also throws a spanner in the works in the IoT/edge setting of deployment (Chukwuani et al., 2025; Okafor, 2024), and legacy compatibility and explainability gaps (Karki et al., n.d) hinder adoption in the enterprise setting as well as post-incident follow-up investigations. Concept drift is an unsolved threat: although federated or transfer learning can alleviate short-term robustness (Chukwuani et al., 2025; Karki et al., n.d), there are few or no continuing, updated benchmarks, and Sharma (2024) does not provide quantified durability statements, in spite of her overall positive survey position. Privacy and governance issues, finally, become evident in issues of the privacy of behavioural telemetry and how to preserve privacy on the pathways through which models are updated.

Summarising the evidence, three applicable results are presented. First, select the model that matches modality: CNNs or autoencoders on high-throughput, static screening; LSTMs/Transformers where sequential behaviour is available and time restrictions allow. Second,

implement hybrid/ensemble pipelines in production to prevent obfuscation and tooling changes; the quantifiable improvements in AUC, F1 and false-positive control make the extra complexity worth the effort in high-risk circumstances. Third, plan to evolve: implement federated and transfer learning now, and consider future continual learning and routine drift audits over rolling windows of data like CICIDS/CERT and, where possible, specialized logs (e.g. ICS). In these papers, the inch in the empirical performance of deep learning is apparent, but its long-term lead lies in data governance, adversarial testing, and lifecycle engineering nearly as much as in design decision.

III. CHALLENGES AND RESEARCH GAPS IN EMPIRICAL STUDIES

Through the discussed papers, there are some prevailing challenges that limit the complete implementation of deep learning (DL) in the sub-domains of malware classification and adaptive threat detection. Prohibitive computational costs preclude use in IoT and edge settings, and limited availability and finicky annotations prevent generalisation. Behavioural models tend to have very high false positive rate and all DL models are susceptible to adversarial evasion. Real-world deployment also has the problems of scalability, explainability gaps and inability to integrate with legacy systems. Based on such limitations, it is clear that there are research gaps. Future efforts are needed to enhance explainable AI to improve trust and adoption and to mitigate concept drift and privacy constraints, through continual and federated learning. Compact DL models are required in resource-constrained IoT/edge environments and adversarially robust models should be given priority. The absence of longitudinal benchmarks is a major impediment and it is demonstrated that it is important to have dynamic datasets at hand, which more closely match malware activity.

IV. RECOMMENDATIONS

Research in the future directions must focus on realistic deployment strategies, namely developing lightweight and energy-efficient DL models suitable to be deployed in the IoT and distributed systems. One should aim at explainable AI that will increase transparency assisting in the adoption of enterprise and regulator. Solving privacy issues through practices like federated learning will be critical in ensuring data-sharing is reduced across sectors. Moreover, researchers should develop standardised regularly changing benchmark datasets which would reflect the malware evolution in time. Lastly, collaboration across disciplines between academic establishments, industry, and cybersecurity agencies will be important to hike resilience towards adversarial attack and make sure that the models put forward are operationally viable in different cyber environments.

V. CONCLUSION

The practice shows that deep learning has been introduced into the sphere of malware detection and

behavioural threat detection, showing higher accuracy rates and greater adaptability and resilience than traditional machine learning. CNNs are among the best at static feature analysis, RNNs and Transformers are best at capturing temporal behaviour, hybrid ensembles are generally best at providing robustness against obfuscation and adversarial approaches. However, the data may remain limited, lack scalability, and not be adversarially resistant especially in distributed IoT and edge setups. The review shows that the potential of deep learning does not only lie in the architecture design but continuous adaptation, adversarial tests and lifecycle integration and these aspects define the future of proactive, adaptive and trustworthy cybersecurity solutions.

REFERENCES

- [1]. Alomar, K., Aysel, H. I., & Cai, X. (2024). RNNs, CNNs and transformers in human action recognition: A survey and a hybrid model. *arXiv preprint arXiv:2407.06162*.
- [2]. Chukwuani, E. N., Odunsi, O. R., & Ikemefuna, C. D. (2025). Machine learning techniques for real-time malware classification and threat detection in distributed systems.
- [3]. Gulatas, I., Kilinc, H. H., Zaim, A. H., & Aydin, M. A. (2023). Malware threat on edge/fog computing environments from Internet of Things devices perspective. *Ieee Access*, 11, 33584-33606.
- [4]. Karki, N., Maharjan, P., Dhungana, R., Shrestha, A., Sharma, B. P., Gurung, D., & Tamang, K. Cybersecurity with Deep Learning: Malware Detection through Real-Time Pattern Recognition, Evasion Resistance, and Collaborative Threat Intelligence in Modern Digital Ecosystems.
- [5]. Kaur, R., & Singh, M. (2014). A survey on zero-day polymorphic worm detection techniques. *IEEE Communications Surveys & Tutorials*, 16(3), 1520-1549.
- [6]. Okafor, M. O. (2024). Deep learning in cybersecurity: Enhancing threat detection and response. *World Journal of Advanced Research and Reviews*, 24(3), 1116–1132.
- [7]. Okoli, U. I., Obi, O. C., Adewusi, A. O., & Abrahams, T. O. (2024). Machine learning in cybersecurity: A review of threat detection and defense mechanisms. *World Journal of Advanced Research and Reviews*, 21(1), 2286-2295.
- [8]. Okoli, U. I., Obi, O. C., Adewusi, A. O., & Abrahams, T. O. (2024). Machine learning in cybersecurity: A review of threat detection and defense mechanisms. *World Journal of Advanced Research and Reviews*, 21(1), 2286-2295.
- [9]. Redhu, A., Choudhary, P., Srinivasan, K., & Das, T. K. (2024). Deep learning-powered malware detection in cyberspace: A contemporary review. *Frontiers in Physics*, 12, 1349463.
- [10]. Sharma, R. (2024). Machine learning-driven approaches for contemporary cybersecurity: Intrusion detection, malware classification, and incident response. *Nuvern Machine Learning Reviews*, 18(1), 77–95.