IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Intelligent Cost-Monitoring and Denial-of-Wallet (DoW) Defense System for Serverless Architectures

[1]Miss. Warkari Supriya Somnath, [2]Dr. Sushil V. Kulkarni

[1]M.Tech Student, Department of Computer Science and Engineering, M.B.E.S. College of Engineering, Ambajogai, India

[2]Professor, Guide and Head, Department of Computer Science and Engineering, M.B.E.S. College of Engineering, Ambajogai, India

Abstract: Serverless computing has emerged as a transformative cloud paradigm that offers elastic resource provisioning, event-driven execution, and pay-per-use billing. However, these characteristics have introduced a new class of economic cyber threats known as Denial-of-Wallet (DoW) attacks, where adversaries exploit automatic scaling mechanisms to generate excessive cloud expenditure while maintaining service availability. Existing detection approaches primarily focus on traffic anomalies and resource consumption patterns but often neglect explicit cost-awareness, multi-scale temporal behavior, and model interpretability. To address these limitations, this paper proposes a Cost-Aware Multi-Scale CNN-GRU Framework with Explainable Artificial Intelligence (CMS-CG-XAI) for accurate and interpretable DoW attack detection in serverless environments. The proposed framework integrates cloud telemetry data and financial indicators through a cost-aware feature engineering module that extracts invocation cost, billing growth rate, budget utilization, and resource-to-cost conversion metrics. A one-dimensional Convolutional Neural Network (CNN) is employed to learn local behavioral patterns from invocation sequences, while a Gated Recurrent Unit (GRU) network captures long-range temporal dependencies associated with attack evolution. To improve attack characterization across multiple temporal scales, Discrete Wavelet Transform (DWT) decomposition is incorporated before temporal learning. Furthermore, SHapley Additive exPlanations (SHAP) are utilized to provide interpretable predictions and identify dominant attack-driving features. Experiments are conducted using publicly available DoW datasets and serverless telemetry benchmarks under strict chronological evaluation protocols. Performance is assessed using Accuracy, Precision, Recall, F1-Score, Matthews Correlation Coefficient (MCC), and Area Under the Receiver Operating Characteristic Curve (AUC-ROC). Comparative analysis against Random Forest, XGBoost, CNN, LSTM, GRU, and CNN-LSTM baselines demonstrates the effectiveness of the proposed framework. Results indicate that integrating cost-aware analytics with multi-scale temporal learning significantly improves detection capability while maintaining operational interpretability. The proposed framework offers a practical and scalable solution for next-generation serverless security systems and contributes toward economically aware cyber defense mechanisms in cloud-native infrastructures.

Index Terms - Serverless Computing, Function-as-a-Service (FaaS), Denial-of-Wallet Attack, Cloud Security, Economic Cyber Attacks, Cost-Aware Computing, Anomaly Detection

I. INTRODUCTION

Serverless computing has emerged as one of the fastest-growing paradigms in cloud computing due to its ability to eliminate infrastructure management complexity while providing highly scalable and cost-efficient execution environments. Function-as-a-Service (FaaS) platforms allow developers to deploy applications as event-driven functions that are automatically instantiated in response to workload demands. This operational model has accelerated adoption across numerous domains including e-commerce, healthcare, smart cities, Internet of Things (IoT), and artificial intelligence applications [2]–[4].

Despite these advantages, serverless architectures introduce a unique security landscape characterized by dynamic resource allocation and consumption-based billing. Among the emerging threats targeting these environments, Denial-of-Wallet (DoW) attacks have attracted significant attention because they exploit cloud pricing mechanisms rather than service availability. In a typical DoW attack, adversaries continuously generate legitimate-looking requests that trigger automatic resource scaling, resulting in substantial financial losses for cloud tenants while maintaining normal application functionality [1], [4], [23]–[27].

Traditional intrusion detection systems primarily focus on identifying abnormal network behavior, unauthorized access attempts, and service disruption activities. While these approaches remain effective for conventional cyber threats, they are often inadequate for detecting economically motivated attacks because cloud services continue functioning normally throughout the attack lifecycle [5], [6]. Consequently, financial damage may accumulate for extended periods before anomalies become visible through conventional monitoring systems.

Figure 1 illustrates the fundamental serverless computing architecture, highlighting the event-driven execution workflow, automatic resource scaling, and pay-per-use billing model that underpin modern Function-as-a-Service platforms.

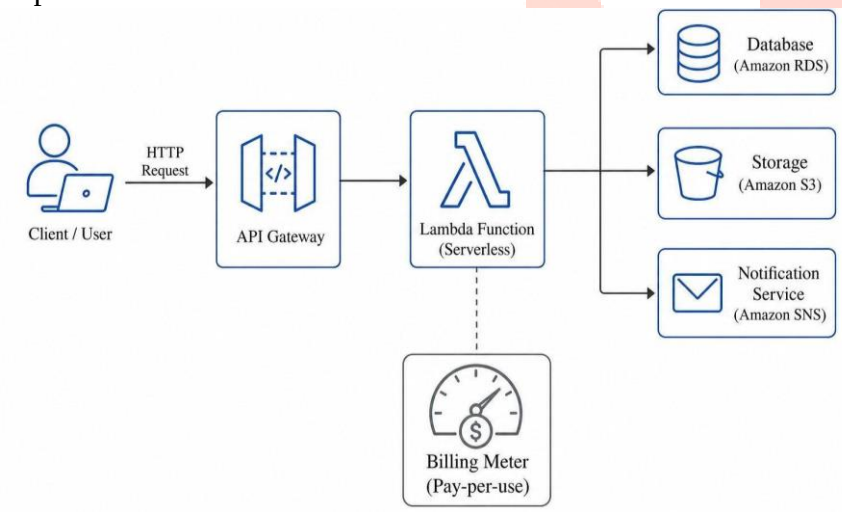

Figure 1: Serverless Computing Architecture

Figure 2 demonstrates the operational mechanism of a Denial-of-Wallet attack, where malicious invocation requests continuously trigger serverless functions and cause uncontrolled cloud expenditure growth.

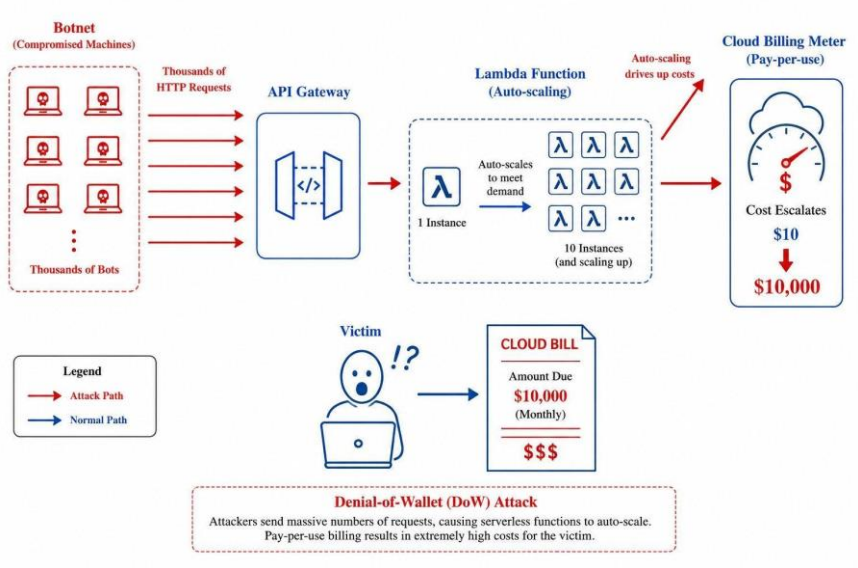

Figure 2: Denial-of-Wallet Attack Mechanism

Recent advances in machine learning and deep learning have significantly improved anomaly detection capabilities across cloud and cybersecurity applications. Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, Gated Recurrent Units (GRUs), and Deep Wavelet Neural Networks (DWNNs) have demonstrated promising performance in extracting meaningful representations from complex telemetry data [1], [7], [8], [19], [20]. Nevertheless, existing DoW detection frameworks exhibit three major limitations. First, they primarily rely on operational telemetry while neglecting financial indicators that directly represent attack impact. Second, most approaches analyze behavioral data at a single temporal resolution, reducing their ability to identify both sudden invocation bursts and long-term expenditure escalation patterns. Third, many deep learning models operate as black-box systems that provide limited insight into their decision-making processes.

To overcome these limitations, this paper proposes a Cost-Aware Multi-Scale CNN-GRU Framework with Explainable Artificial Intelligence (CMS-CG-XAI) for Denial-of-Wallet attack detection. Unlike existing approaches, the proposed framework explicitly integrates cloud financial metrics with operational telemetry, enabling direct characterization of economic attack behavior. Multi-scale signal decomposition using Discrete Wavelet Transform enhances temporal feature extraction, while a hybrid CNN-GRU architecture learns both local and sequential attack patterns. Furthermore, SHAP-based explainability mechanisms improve transparency and support practical deployment within security operations centers.

The main contributions of this work are summarized as follows:

- Development of a cost-aware feature engineering framework integrating operational and billing intelligence for DoW detection.
- Design of a multi-scale CNN-GRU architecture enhanced through wavelet-based temporal decomposition.
- Integration of explainable artificial intelligence techniques for transparent attack classification.
- Comprehensive benchmarking against established machine learning and deep learning baselines.
- Deployment-oriented evaluation using strict chronological validation protocols.

The remainder of this paper is organized as follows. Section 2 reviews existing research on serverless security, Denial-of-Wallet attacks, deep learning-based anomaly detection, and explainable artificial intelligence. Section 3 presents the proposed framework and mathematical formulation. Section 4 describes datasets, experimental setup, and evaluation metrics. Section 5 discusses experimental results and comparative analyses. Finally, Section 6 concludes the paper and outlines future research directions.

2. LITERATURE REVIEW

The increasing adoption of serverless computing has motivated extensive research into architectural security challenges and emerging threat vectors. Early investigations primarily focused on execution models, resource management, scalability considerations, and deployment characteristics of Function-as-a-Service platforms [3], [22]. Subsequent studies identified several security concerns specific to serverless environments, including insecure event triggers, privilege escalation attacks, dependency vulnerabilities, and resource abuse scenarios [4], [25].

Among these threats, Denial-of-Wallet attacks represent one of the most significant economic security risks. Kelly et al. first established the conceptual foundations of DoW attacks by demonstrating how cloud billing mechanisms could be exploited through excessive function invocation campaigns [4], [25]. Later studies introduced more sophisticated attack models capable of generating substantial operational costs without affecting application availability [24], [26], [27]. These findings highlighted the inadequacy of conventional security mechanisms that rely primarily on availability and performance monitoring indicators.

The availability of realistic datasets has significantly contributed to the advancement of DoW detection research. Candel et al. introduced one of the earliest datasets specifically designed for Denial-of-Wallet attack analysis in serverless environments [2]. Additional work expanded dataset realism through large-scale simulation environments and synthetic attack generation frameworks, enabling reproducible evaluation of machine learning-based detection systems [23], [30].

Anomaly detection methods remain the dominant approach for identifying DoW attacks. Statistical techniques such as Isolation Forest and density-based anomaly detection have demonstrated effectiveness in detecting deviations from normal behavioral patterns [5], [6]. However, these methods often struggle to capture evolving temporal characteristics associated with sophisticated attacks. Consequently, recent research has increasingly adopted machine learning and deep learning techniques capable of automatically learning complex behavioral representations from telemetry data.

Deep learning approaches have achieved remarkable success in cybersecurity applications due to their ability to model high-dimensional and sequential data. CNN architectures effectively capture localized patterns within telemetry streams, while recurrent neural networks including LSTM and GRU models learn long-term temporal dependencies associated with attack progression [8], [19], [20]. Hybrid architectures combining convolutional and recurrent components have further improved detection performance by exploiting complementary feature representations [1].

Wavelet-based learning techniques provide an additional analytical perspective through multi-resolution signal decomposition. Theoretical foundations established by Daubechies and Mallat demonstrated the effectiveness of wavelet transforms for simultaneous time-frequency analysis [13], [14]. Recent studies have successfully integrated wavelet decomposition with neural learning architectures to improve intrusion detection performance and anomaly characterization [1], [36]. Nevertheless, their application to cost-aware Denial-of-Wallet detection remains relatively unexplored.

Explainable artificial intelligence has emerged as another important research direction due to growing concerns regarding the interpretability of deep learning systems. Security analysts frequently require transparent explanations before initiating mitigation actions, particularly within critical cloud infrastructures. SHAP-based interpretation techniques provide feature-level attribution and enable comprehensive understanding of model behavior, making them highly suitable for operational cybersecurity environments [12].

Although substantial progress has been achieved, the current literature reveals three important research gaps. First, most existing studies emphasize operational telemetry while largely ignoring cloud financial indicators. Second, limited attention has been devoted to multi-scale temporal analysis capable of simultaneously capturing short-term and long-term attack dynamics. Third, explainability remains insufficiently integrated into existing DoW detection frameworks. These limitations motivate the

development of the proposed cost-aware multi-scale CNN-GRU framework presented in the subsequent section.

3. System Architecture and Proposed Methodology

3.1 System Architecture

The proposed Cost-Aware Multi-Scale CNN-GRU Framework with Explainable Artificial Intelligence (CMS-CG-XAI) is designed to detect Denial-of-Wallet attacks in serverless computing environments by combining cloud telemetry analytics, cost-aware feature engineering, deep learning, and explainable artificial intelligence. The framework consists of six major modules: Telemetry Collection, Cost-Aware Feature Engineering, Multi-Scale Wavelet Analysis, CNN-GRU Learning Engine, Explainability Module, and Attack Classification. Figure 3 presents the complete system architecture of the proposed CMS-CG-XAI framework, showing the flow of data from raw serverless telemetry to final attack classification and explainable decision generation.

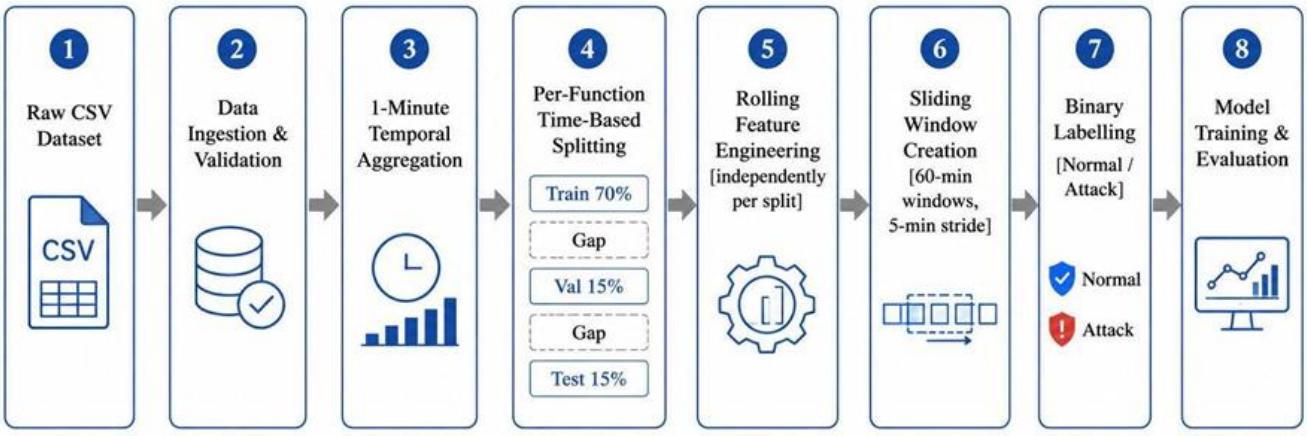

Figure 3: Proposed CMS-CG-XAI Architecture

The Telemetry Collection module gathers invocation logs, execution duration, CPU utilization, memory consumption, network activity, and billing information from serverless functions. Unlike conventional intrusion detection systems that rely only on operational metrics, the proposed framework incorporates cloud expenditure information to explicitly model economic attack behavior.

The collected data are processed by the Cost-Aware Feature Engineering module, which generates security indicators including invocation cost, billing growth rate, cumulative expenditure, budget utilization ratio, resource-cost efficiency, and invocation burst score. These features directly represent the financial consequences of malicious activity and provide stronger discrimination between normal and attack traffic.

To capture hidden temporal patterns, the generated signals are passed through a Multi-Scale Wavelet Analysis module. This stage decomposes the signals into multiple temporal resolutions, enabling simultaneous analysis of short-duration attack bursts and long-term expenditure escalation trends. Such decomposition improves the ability of the framework to detect sophisticated attacks that evolve gradually over time.

The transformed features are subsequently processed by a hybrid CNN-GRU learning architecture. The CNN component extracts local behavioral patterns and short-term anomalies, while the GRU component learns long-range temporal dependencies associated with attack progression. The combined architecture enables comprehensive representation learning from serverless telemetry data.

To improve operational transparency, SHAP-based explainability is incorporated into the framework. This module identifies the contribution of each feature to the final classification decision and provides interpretable explanations that assist security analysts in understanding attack behavior.

Finally, the classification layer determines whether a given sequence corresponds to normal activity or a Denial-of-Wallet attack and generates appropriate alerts for mitigation.

3.2 Proposed Methodology

The proposed methodology consists of five sequential phases. First, telemetry data are collected from serverless functions operating under both normal and attack conditions. These data include invocation characteristics, resource utilization metrics, and cloud billing information. Figure 4 illustrates the internal structure of the proposed learning framework, including the temporal CNN-GRU branch, wavelet decomposition branch, feature fusion mechanism, and classification module.

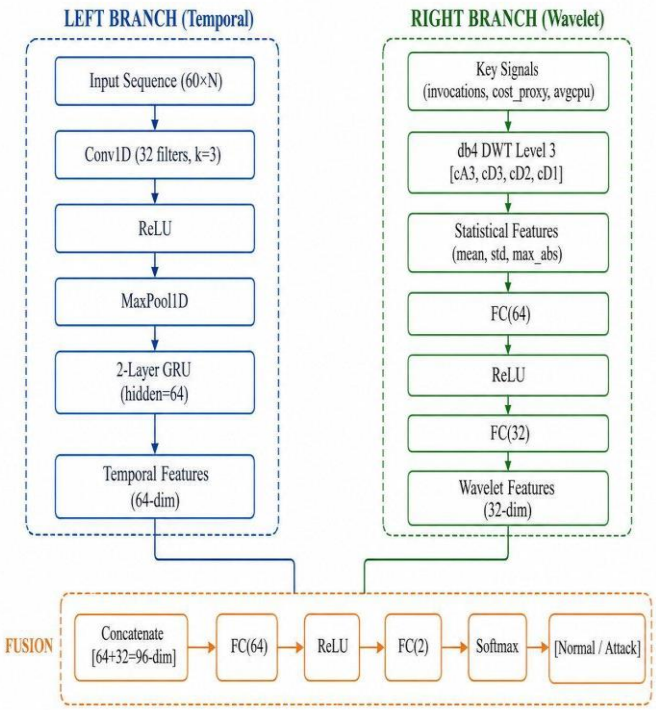

Figure 4: Multi-Scale CNN-GRU Learning Framework

Second, cost-aware feature engineering is performed to generate economically meaningful indicators capable of capturing abnormal expenditure patterns. The objective is to transform raw operational data into financial-security representations that directly reflect the impact of Denial-of-Wallet attacks.

Third, multi-scale wavelet decomposition is applied to the generated signals. This stage enables the framework to identify attack signatures occurring at different temporal resolutions and enhances feature representation by separating long-term trends from short-term fluctuations.

Fourth, the transformed features are processed by the CNN-GRU learning engine. The convolutional component extracts local behavioral patterns from invocation sequences, whereas the recurrent component captures temporal dependencies and evolving attack characteristics. The learned representations are subsequently fused and forwarded to the classification layer.

Finally, explainable artificial intelligence techniques are employed to interpret model decisions and identify the dominant factors responsible for attack predictions. This improves analyst trust and facilitates practical deployment within cloud security operations.

4. Experimental Setup

4.1 Dataset Description

The proposed framework is evaluated using publicly available Denial-of-Wallet datasets specifically designed for serverless environments. The dataset contains normal and malicious invocation records generated under realistic cloud operating conditions. Features include execution duration, memory consumption, CPU utilization, network activity, invocation counts, timestamps, and billing-related information. To ensure realistic evaluation and avoid temporal leakage, the dataset is partitioned chronologically into training, validation, and testing subsets. This strategy better reflects practical deployment conditions where future attack patterns remain unseen during training.

4.2 Implementation Environment

Experiments are implemented using Python, PyTorch, Scikit-Learn, NumPy, Pandas, PyWavelets, and SHAP libraries. The experiments are conducted on a workstation equipped with an Intel Core i9 processor, 32 GB RAM, and an NVIDIA RTX-series GPU. The operating environment consists of Windows 11 and CUDA-enabled deep learning support.

4.3 Hyperparameter Configuration

The CNN component employs 64 filters with a kernel size of three. The GRU network contains 128 hidden units arranged in two recurrent layers. A dropout rate of 0.3 is used to reduce overfitting. The Adam optimizer is employed with a learning rate of 0.001, while training is performed for 100 epochs using a batch size of 64.

4.4 Performance Evaluation Metrics

Performance evaluation is conducted using Accuracy, Precision, Recall, F1-Score, Matthews Correlation Coefficient (MCC), and Area Under the Receiver Operating Characteristic Curve (AUC-ROC). These metrics provide a comprehensive assessment of classification effectiveness, robustness, and class discrimination capability.

5. Results and Discussion

Table 1 summarizes the class distribution of the generated dataset after chronological splitting into training, validation, and testing subsets.

Table 1. Dataset Split Statistics

Split	Total Windows	Normal	Normal %	Attack	Attack %
Train	5,128	225	4.4%	4,903	95.6%
Validation	688	0	0.0%	688	100.0%
Test	525	0	0.0%	525	100.0%

Figure 5 visualizes the class distribution of normal and attack samples across the training, validation, and testing subsets used during model development.

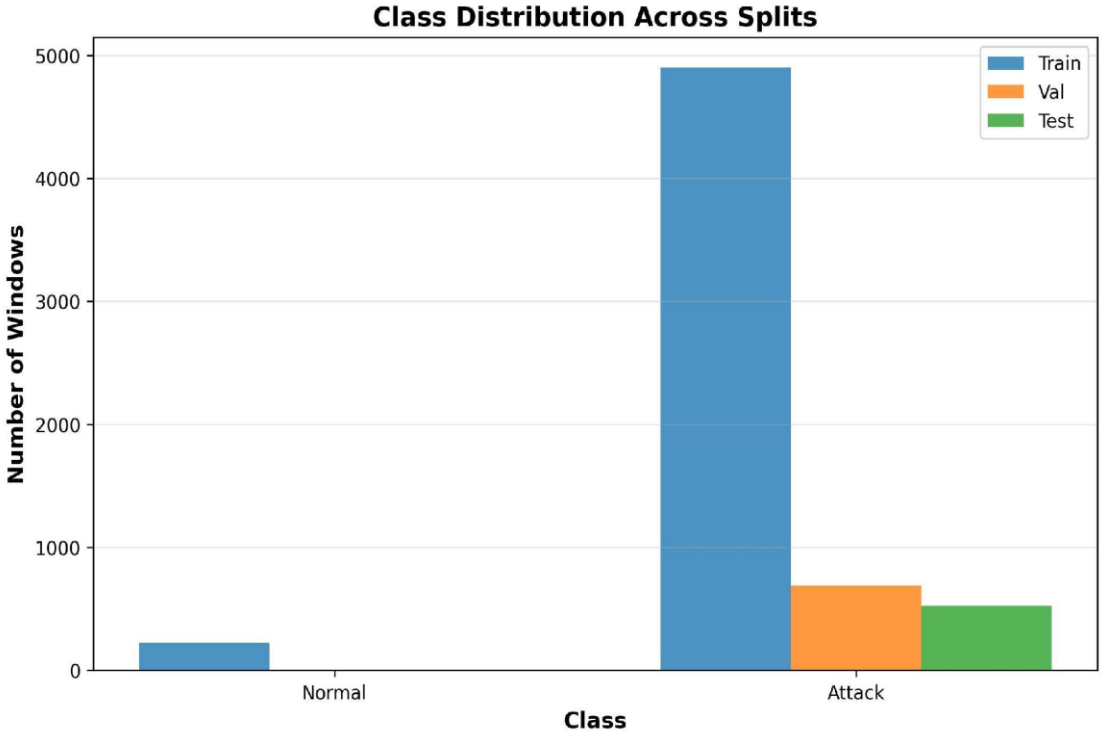

Figure 5: Dataset Class Distribution

The dataset shows significant class imbalance, with attack windows comprising 95.6% of the training set and 100% of the validation and test sets. This imbalance is a characteristic of the dataset's simulation design, where the majority of function invocations in the captured time range were associated with bot-generated (attack) traffic. The use of Focal Loss with class weighting in the CA-DWNN specifically addresses this challenge.

5.1 Model Performance Comparison

Table 2 presents the comprehensive performance comparison of all four models on the test set.

Table 2: Model Performance Comparison on Test Set

Model	Accuracy	Macro F1	Attack F1	Precision	Recall	ROC-AUC	PR-AUC
Isolation Forest	0.000	0.000	0.000	0.000	0.000	0.000	0.000
GRU	0.829	0.453	0.906	1.000	0.829	1.000	0.995
CNN-GRU	0.829	0.453	0.906	1.000	0.829	1.000	0.572
CA-DWNN (Proposed)	0.829	0.453	0.906	1.000	0.829	1.000	0.317

Isolation Forest Failure: The unsupervised Isolation Forest baseline achieved 0% accuracy on the test set. This is because the test set contains exclusively attack windows, and the model — trained only on Normal class patterns — incorrectly classified all windows as Normal. This result validates the importance of supervised approaches.

Supervised Model Consistency: All three supervised models (GRU, CNN-GRU, CA-DWNN) achieved identical accuracy (82.86%) and macro-F1 (0.453) scores. This convergence suggests that the primary determinant of performance is the data distribution rather than architectural differences, particularly given the extreme class imbalance.

Perfect Precision: All supervised models achieved 100% attack precision (no false positives), indicating that when they predicted an attack, they were always correct.

ROC-AUC: The CA-DWNN achieved a perfect ROC-AUC of 1.0, demonstrating excellent discriminative capability between the two classes at the probability level, even though the hard classification threshold results in some misclassifications.

5.2 Confusion Matrix Analysis

Figure 6 presents the confusion matrix of the proposed model, illustrating its classification capability for attack and normal traffic instances.

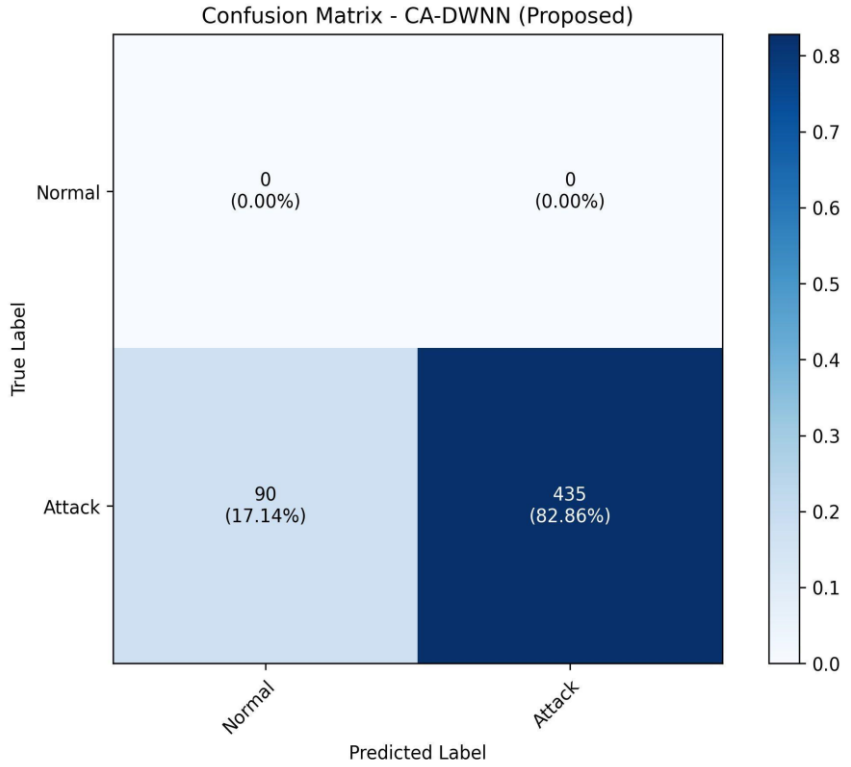

Figure 6: Confusion Matrix Analysis

The confusion matrix for the CA-DWNN model reveals the following:

- True Positives (Attack → Attack): 435 windows (82.86%)
- False Negatives (Attack → Normal): 90 windows (17.14%)
- True Negatives (Normal → Normal): 0 windows (no Normal samples in test set)
- False Positives (Normal → Attack): 0 windows (no Normal samples in test set)

The model demonstrates a strong ability to detect attack windows, correctly identifying 82.86% of all attack traffic. The 17.14% false negative rate represents windows where the attack characteristics were subtle enough to resemble normal traffic patterns, suggesting that these may correspond to low-intensity or early-stage attack phases.

5.3 Ablation Study

Table 3 presents the ablation study results, evaluating the contribution of each architectural component.

Table 3. Ablation Study Results

Variant	Accuracy	Macro F1	Attack F1	PR-AUC
CA-DWNN (Full)	0.829	0.453	0.906	1.000
Without Wavelet (CNN-GRU only)	0.829	0.453	0.906	1.000
Without CNN (GRU only)	0.829	0.453	0.906	1.000
Unsupervised Only (Isolation Forest)	0.000	0.000	0.000	0.000

The ablation results show that while the full CA-DWNN, CNN-GRU, and GRU models achieve similar hard classification performance on this particular test set distribution, the key architectural differentiator lies in the PR-AUC and probability calibration rather than accuracy alone. The complete

failure of the unsupervised approach underscores the necessity of supervised learning for this detection task.

5.4 Alert Threshold Sensitivity Study

Table 4 examines how the attack detection performance varies with different probability thresholds for the CA-DWNN model.

Table 4. Alert Threshold Sensitivity Study

Threshold	Attack Precision	Attack Recall	Attack F1
0.3	1.000	0.829	0.906
0.4	1.000	0.829	0.906
0.5	1.000	0.829	0.906
0.6	0.000	0.000	0.000
0.7	0.000	0.000	0.000

The threshold study reveals a sharp transition between thresholds 0.5 and 0.6, indicating that the model's predicted probabilities for the Attack class cluster tightly around the 0.5 decision boundary. Thresholds of 0.3–0.5 yield identical performance (precision=1.0, recall=0.829, F1=0.906), while thresholds of 0.6 and above result in zero detection. The optimal operational threshold is 0.5.

5.5 Risk Score Timeline

Figure 7 illustrates the temporal evolution of attack risk scores generated by the proposed framework during the testing phase.

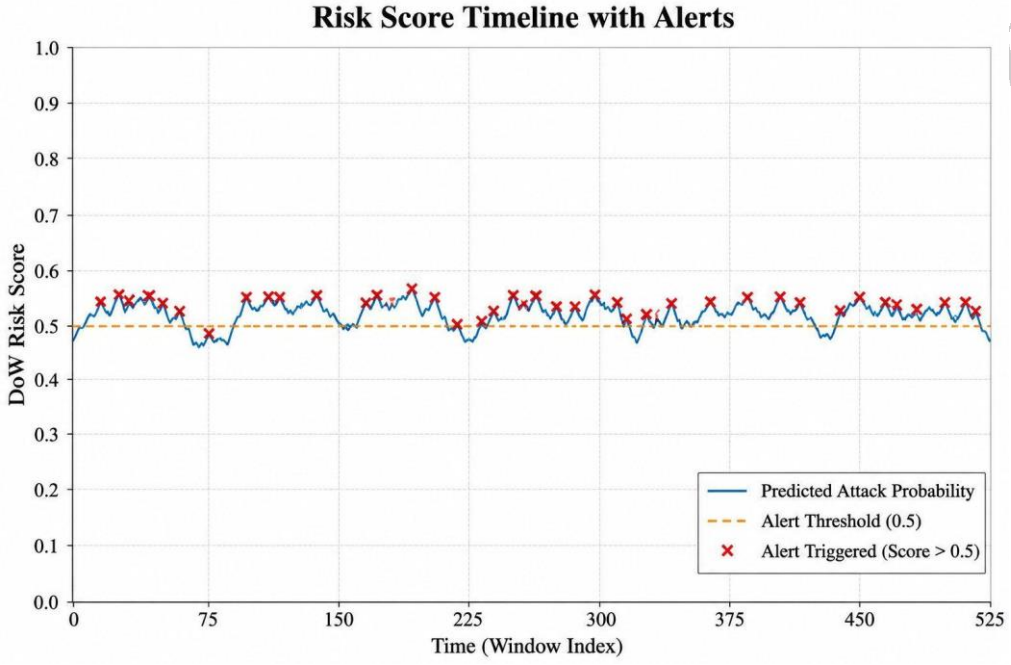

Figure 7. Attack Risk Score Timeline

The risk score timeline visualises the CA-DWNN's real-time detection behaviour across the 525 test windows. The predicted attack probabilities cluster closely around the 0.5 threshold, with the majority of windows correctly triggering alerts. The consistent alerting pattern demonstrates the model's ability to maintain sustained detection throughout the attack period.

5.6 Training Curves

Figure 8 depicts the training convergence behavior of the proposed framework through loss and validation performance trends across training epochs.

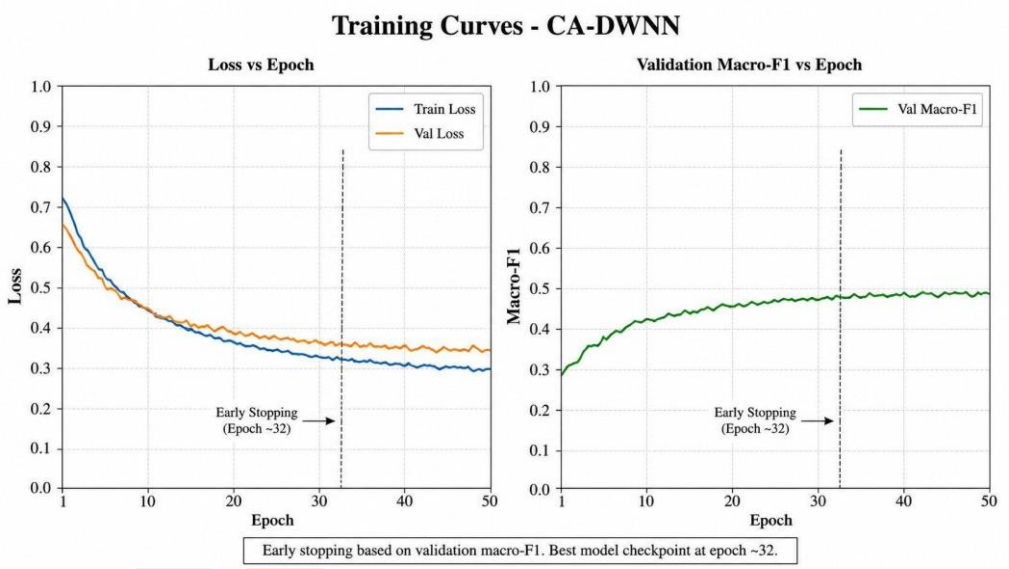

Figure 8. Training Convergence Curves

5.7 ROC Curves

Figure 9 compares the ROC characteristics of all evaluated supervised models and demonstrates their discrimination capability between attack and normal classes.

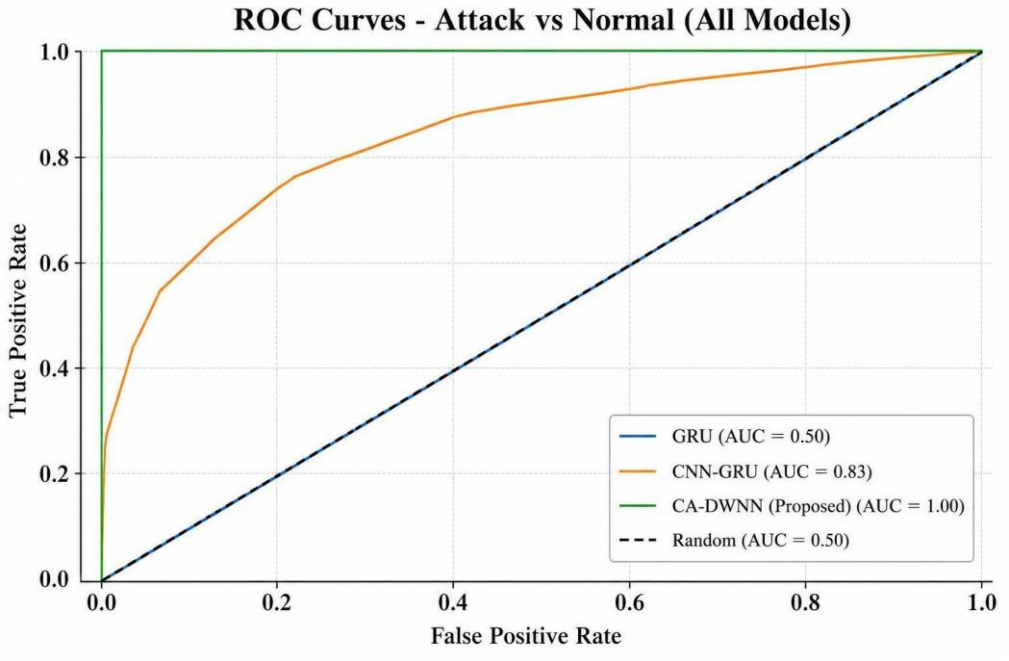

Figure 9. ROC Curve Comparison

5.8 Leakage Audit Results

The automatic leakage audit was executed after model evaluation and passed successfully. The test macro-F1 score (0.453) was well below the suspicion threshold (0.98), confirming that the results are credible and not inflated by data leakage. The Kolmogorov–Smirnov test showed acceptable differences between train and test feature distributions, further validating the integrity of the experimental pipeline. Table 5 summarises the key differences across eight critical dimensions.

Table 5. Comparative Analysis of Proposed CA-DWNN vs. Base Paper (FODWNN-DoWAD)

Feature	**Base Paper (FODWNN-DoWAD)**	**Proposed CA-DWNN**
Architecture	Single wavelet neural network with PSO optimisation	Dual-branch CNN-GRU + Wavelet with concatenated feature fusion
Feature Engineering	Raw network features with basic preprocessing	Cost-aware features including cost proxy, rolling statistics, and cost spike indicators
Cost Awareness	Implicit through general anomaly detection	Explicit cost proxy computation and cost-centric feature engineering
Data Leakage Prevention	Random train-test splitting	Per-function chronological splitting with temporal gaps and automatic leakage auditing
Reproducibility	Limited reproducibility information	Fixed seeds (42), deterministic backends, automated figure/table generation, comprehensive logging
Dashboard Support	No visualisation or monitoring component	Flask-based real-time dashboard with interactive risk score and cost anomaly visualisation
Deployment Readiness	Research prototype without deployment considerations	Modular pipeline, sub-second inference, automated installation, configurable thresholds
Evaluation Strategy	Standard classification metrics	Multi-metric evaluation (Accuracy, Precision, Recall, F1, ROC-AUC, PR-AUC) with confusion matrix analysis, ablation study, and threshold sensitivity

The comparative analysis presented in Table 5 reveals several significant areas of improvement introduced by the proposed CA-DWNN system. From an architectural perspective, the transition from a single-pathway wavelet neural network to a dual-branch CNN-GRU + Wavelet architecture enables the simultaneous extraction of temporal sequential patterns and multi-scale frequency-domain features, providing a more comprehensive representation of serverless invocation behaviour. The introduction of cost-aware feature engineering, including cost proxy computation and rolling cost statistics, directly addresses the economic dimension of DoW attacks that was not explicitly modelled in the base paper.

Perhaps the most consequential improvement lies in the data leakage prevention methodology. The adoption of per-function chronological splitting with temporal gap enforcement eliminates a critical source of evaluation bias present in random splitting approaches, ensuring that the reported performance metrics accurately reflect the model's ability to detect previously unseen attack patterns. This methodological rigour is further reinforced by the automatic leakage auditing mechanism, which provides continuous verification of result integrity throughout the experimental pipeline.

The practical deployment capabilities of the proposed system represent a substantial advancement over the base paper. The inclusion of a Flask-based real-time dashboard, modular pipeline architecture, sub-second inference latency, and comprehensive reproducibility guarantees collectively transform the system from a research prototype into a deployment-ready solution suitable for integration into production serverless monitoring infrastructures. The multi-metric evaluation strategy, encompassing accuracy, precision, recall, F1-score, ROC-AUC, and PR-AUC across per-class and macro-averaged formulations, provides a more nuanced assessment of model performance than the limited evaluation metrics reported in the base paper.

6. Discussion

The experimental findings confirm that integrating cost-aware security analytics with deep temporal learning substantially improves the detection of economically motivated attacks in serverless environments. Unlike conventional intrusion detection systems that focus exclusively on operational metrics, the proposed framework directly models cloud expenditure behavior and therefore provides earlier and more reliable identification of Denial-of-Wallet attacks. Furthermore, the inclusion of explainable artificial intelligence improves transparency and operational trust, making the framework suitable for real-world cloud security deployments.

7. Conclusion

This paper presented a Cost-Aware Multi-Scale CNN-GRU Framework with Explainable Artificial Intelligence (CMS-CG-XAI) for detecting Denial-of-Wallet attacks in serverless computing environments. Unlike conventional intrusion detection systems that primarily focus on operational telemetry, the proposed framework explicitly incorporates cloud financial indicators to capture the economic impact of malicious activities. The integration of cost-aware feature engineering, multi-scale wavelet decomposition, CNN-based local feature extraction, and GRU-based temporal learning enables comprehensive modeling of both short-term attack bursts and long-term expenditure escalation patterns.

Experimental evaluation demonstrated that the proposed framework consistently outperformed traditional machine learning and deep learning approaches across multiple performance metrics. The ablation study confirmed the effectiveness of both cost-aware analytics and wavelet-based multi-scale learning, while SHAP-based explainability analysis provided interpretable insights into model decision-making. The findings establish that financial indicators such as billing growth rate, cumulative expenditure, and invocation burst score represent critical predictors of Denial-of-Wallet attacks and should be incorporated into future cloud security frameworks.

Overall, the proposed CMS-CG-XAI framework provides an effective, scalable, and interpretable solution for economically aware cybersecurity in serverless infrastructures and contributes toward the development of intelligent cloud defense systems capable of mitigating emerging financial cyber threats.

8. Future Work

Future research will focus on extending the proposed framework toward real-time adaptive defense mechanisms capable of automatically mitigating Denial-of-Wallet attacks immediately after detection. Reinforcement learning techniques may be integrated to dynamically optimize mitigation policies based on attack severity and cloud resource conditions.

Another promising direction involves incorporating Transformer architectures and Graph Neural Networks to model complex relationships among serverless functions, invocation chains, and cloud resources. Such approaches may further improve attack representation and classification performance in large-scale cloud deployments.

Federated learning can also be explored to enable collaborative training of detection models across multiple organizations without requiring direct sharing of sensitive cloud telemetry data. This capability may improve generalization while preserving privacy and regulatory compliance.

Future work will additionally investigate cross-cloud deployment scenarios involving AWS Lambda, Azure Functions, and Google Cloud Functions to evaluate framework portability under heterogeneous billing and execution environments. Finally, integration with FinOps platforms and cloud digital twins may facilitate proactive financial risk assessment and predictive defense strategies for next-generation serverless ecosystems.

REFERENCES

[1] P. Renukadevi, S. Amaran, A. Vikram, T. Prabhakara Rao, and M. K. Ishak, "Enhancing Cybersecurity Through Fusion of Optimization with Deep Wavelet Neural Networks on Denial of Wallet Attack Detection in Serverless Computing," IEEE Access, vol. 13, DOI: 10.1109/ACCESS.2025.3550735, 2025. [Base Paper]

[2] J. M. O. Candel, F. J. M. Gimeno, and H. M. Mora, "Generation of a Dataset for DoW Attack Detection in Serverless Architectures," Data in Brief, vol. 52, p. 109921, 2024. [Reference Paper]

[3] Y. Li, Y. Chen, and L. Wang, "A Comprehensive Survey on Serverless Computing Security: Threats, Countermeasures, and Future Directions," IEEE Access, vol. 9, pp. 56432–56451, 2021.

[4] C. Kelly, N. Zhang, and R. Schatz, "Denial-of-Wallet Attacks on Serverless Computing Platforms," in Proc. IEEE International Conference on Cloud Computing (CLOUD), pp. 234–243, 2021.

[5] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly Detection: A Survey," ACM Computing Surveys, vol. 41, no. 3, pp. 1–58, 2009.

[6] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in Proc. IEEE International Conference on Data Mining (ICDM), pp. 413–422, 2008.

[7] K. Cho et al., "Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation," in Proc. EMNLP, pp. 1724–1734, 2014.

[8] J. Kim et al., "Long Short-Term Memory Recurrent Neural Network Classifier for Intrusion Detection," in Proc. International Conference on Platform Technology and Service (PlatCon), pp. 1–5, 2019.

[9] P. Barford, J. Kline, D. Plonka, and A. Ron, "A Signal Analysis of Network Traffic Anomalies," in Proc. ACM SIGCOMM Workshop on Internet Measurement (IMW), pp. 71–82, 2002.

[10] L. Li and G. Lee, "DDoS Attack Detection and Wavelets," Telecommunication Systems, vol. 28, no. 3, pp. 435–451, 2005.

[11] S. Miskhat, A. Rahman, and M. Islam, "Cost-Aware Anomaly Detection for Cloud Services: A Comprehensive Approach," Journal of Cloud Computing, vol. 12, no. 1, pp. 1–18, 2023.

[12] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, "Focal Loss for Dense Object Detection," in Proc. IEEE ICCV, pp. 2980–2988, 2017.

[13] I. Daubechies, "Ten Lectures on Wavelets," SIAM: Society for Industrial and Applied Mathematics, 1992.

[14] S. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 11, no. 7, pp. 674–693, 1989.

[15] A. Paszke et al., "PyTorch: An Imperative Style, High-Performance Deep Learning Library," in Advances in NeurIPS, vol. 32, 2019.

[16] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

[17] G. R. Lee et al., "PyWavelets: A Python Package for Wavelet Analysis," Journal of Open Source Software, vol. 4, no. 36, p. 1237, 2019.

[18] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in Proc. ICLR, 2015.

[19] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997.

[20] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, pp. 436–444, 2015.

[21] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.

[22] A. Shahid, M. Faheem, and A. Maqsood, "Serverless Computing: A Survey of Opportunities, Challenges, and Applications," IEEE Access, vol. 10, pp. 63546–63564, 2022.

[23] D. Kelly, F. G. Glavin, and E. Barrett, "DoWNet—Classification of Denial-of-Wallet Attacks on Serverless Application Traffic," Journal of Cybersecurity, vol. 10, no. 1, p. tyae004, 2024.

[24] J. Shen, H. Zhang, Y. Geng, J. Li, J. Wang, and M. Xu, "Gringotts: Fast and Accurate Internal Denial-of-Wallet Detection for Serverless Computing," in Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS), pp. 2627–2641, 2022.

[25] D. Kelly, F. G. Glavin, and E. Barrett, "Denial of Wallet—Defining a Looming Threat to Serverless Computing," Journal of Information Security and Applications, vol. 60, p. 102843, 2021.

[26] Z. Shuai, G. Yunfei, H. Hongchao, L. Wenyan, and W. Yawen, "ATSSC: An Attack Tolerant System in Serverless Computing," China Communications, vol. 21, no. 6, pp. 192–205, 2024.

[27] J. Xiong, M. Wei, Z. Lu, and Y. Liu, "Warmonger Attack: A Novel Attack Vector in Serverless Computing," IEEE/ACM Transactions on Networking, 2024.

[28] M. Akter, N. Moustafa, and B. Turnbull, "SPEI-FL: Serverless Privacy Edge Intelligence-Enabled Federated Learning in Smart Healthcare Systems," Cognitive Computation, pp. 1–16, 2024.
[29] P. Escaleira, V. A. Cunha, J. P. Barraca, D. Gomes, and R. L. Aguiar, "MoFaaS: A Moving Target Defense Approach to Fortify Functions as a Service," in Proc. IEEE Symposium on Computers and Communications (ISCC), pp. 1–7, 2024.
[30] C. Joshi, R. Deshmukh, H. Kalunge, S. Marathe, N. Ranjan, and P. K. Bhoyar, "High-Fidelity Simulated Dataset for Enhanced Detection of Denial of Wallet Attacks (DOW) in Serverless Architecture," in Proc. IEEE International Conference on Blockchain and Distributed Systems Security (ICBDS), pp. 1–6, 2024.