

Deep Learning-Based Network Traffic Anomaly Detection: A Study in IoT Environments

Lin Li ^{1*}, Yitian Zhang ^{1,2}, Jiayi Wang ², Ke Xiong ³

^{1,*} Electrical and Computer Engineering, Carnegie Mellon University, PA, USA

^{1,2} Accounting, UW-Madison, WI, USA

² Computer engineering, Illinois institute of technology, IL, USA

³ Computer Science, University of Southern California, CA, USA

*Corresponding author

Abstract: *This study presents a deep learning technique to detect vulnerabilities in the Internet of Things (IoT) environment. The proposed method combines the manual design with the learning function of autoencoders, together with the deep neural network architecture associated with the Long Short-Term Memory (LSTM) layer. Experiments performed on the IoT-23 dataset show that our method outperforms traditional machine learning and state-of-the-art deep learning methods, achieving an accuracy of 99.2%, an F1-score of 0.987, and an AUC-ROC of 0.998. The framework addresses critical issues in IoT security, including device diversity, vehicle model diversity, and real-time research needs. The ablation studies show the importance of combining manual and autoencoder-based feature extraction. Grad-CAM visualisations improve the model definition by identifying the essential features for classifying bad vehicles and not good. The model's ability to capture the time-dependent nature of network flows makes it helpful in investigating complex, time-dependent variables. Although there are limitations in computing needs and general capabilities across IoT ecosystems, the proposed system shows significant potential for practical use in IoT security systems. This research contributes to advancing IoT security by providing a powerful, efficient, and easily interpretable system to discover the system's capabilities in terms of strengths and weaknesses in the IoT network environment.*

Keywords: Internet of Things (IoT), Network Anomaly Detection, Deep Learning, LSTM.

1. INTRODUCTION

1.1 Background on IoT Security Challenges

The Internet of Things (IoT) has revolutionised the way communication and communication devices are used, enabling connectivity and data exchange. As IoT networks continue to expand, the number of connected devices will reach 75 billion by 2025. This rapid growth presents significant security issues, particularly many due to the differences in IoT devices, low budgets, and different communication methods[1]. IoT devices often have poor security, making them vulnerable to various threats and attacks. The limited resources of IoT devices, such as limited memory and processing power, make it more difficult to implement traditional security measures.

The large-scale deployment of IoT devices in various areas, from smart homes to factories, creates a massive challenge for criminals to exploit. These security issues have led to many policies and standards used in IoT communications, making implementing security measures throughout the network difficult[2]. In addition, the sheer volume of data generated by IoT devices presents data privacy, integrity, and confidentiality issues. As IoT systems are increasingly integrated into critical and valuable applications, the impact of security breaches will become more severe, highlighting the urgent need for security solutions that adapt to the characteristics of the IoT environment[3].

1.2 Importance of Anomaly Detection in IoT Networks

Vulnerability detection plays a vital role in protecting IoT networks from security breaches and maintaining the integrity of data transmission. In an IoT environment, anomalies can appear as abnormal patterns in network connectivity, device behaviour, or data flow, expressed as impact, denial of service, or data theft[4]. Timely detection of these vulnerabilities is essential for security protection and reducing their impact on the network. The

regulations based on the informal law often fall short of addressing the strength and complexity of IoT networks, requiring more guidance.

The ability to detect and analyse vulnerabilities in real time is essential to maintain the security and reliability of IoT systems, especially in critical systems and sensitive applications. Effective detection can help identify compromised devices, detect intrusion attempts, and reduce the spread of malware online[5]. In addition, vulnerability detection systems can provide information about network behaviour, allowing IoT infrastructures to be effectively monitored and optimised. As the scope and complexity of IoT networks continue to grow, the importance of vulnerability detection systems increases the impact on the security and reliability of IoT ecosystems.

1.3 Deep Learning in IoT Anomaly Detection

Deep learning has emerged as a promising method for solving the fault detection problem in IoT networks. The ability of deep learning models to extract complex features from high-dimensional data makes them suitable for analysing current trends in IoT networked vehicles. Deep neural networks can learn a hierarchical representation of data, allowing them to detect tiny anomalies that go undetected by standard methods. Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) have shown particular promise for temporal and spatial traffic data networks[6]. Autoencoders, a class of unsupervised deep learning models, have proven to be effective in learning patterns of behaviour and identifying differences as trust deficits.

The scalability and adaptability of deep learning models make them suitable for handling large volumes of heterogeneous data generated by IoT devices. Deep learning methods can adapt to changes in the network and learn from new data, making them particularly useful in the changing and evolving nature of IoT security. In addition, deep learning models have the ability to process raw data without extensive processing, which works well with the diverse and often complex nature of IoT network traffic[7]. As research in this field has progressed, deep search-based vulnerability detection is expected to play an essential role in improving the security and recovery of IoT networks.

1.4 Research Objectives and Contributions

This research is designed to develop a deep learning-based conflict detection system in an IoT environment. The main goal of this study is to develop a new deep learning system designed for IoT network traffic analysis, including the unique characteristics and limitations of IoT devices. We focus on creating a good data advance and clearing system that can handle the diversity and height of IoT network data traffic. This study also aims to evaluate the work plan in search of different types of inequality and compare it with the existing state[8]. In addition, we investigate the interpretation of deep learning models to provide insight into the decision-making process for flaw detection.

The main results of this research include a comprehensive analysis of the problems and needs for detecting anomalies in IoT networks, considering the latest advances in the field deep. We present a deep learning approach that integrates multiple neural network architectures to capture physical and internal aspects of IoT network traffic. A modified feature extraction system is designed to handle the different data types and processes commonly encountered in the IoT environment[9]. The research includes various analyses using global IoT network traffic data, showing the effectiveness and efficiency of the plan. Finally, we provide insight into the decision-making model by performing optical and ablation studies, improving the interpretation of deep-seated research don't believe. This research contributes to the growing body of knowledge about IoT security and provides practical solutions to improve the performance of IoT networks against threats.

2. LITERATURE REVIEW

2.1 Traditional Anomaly Detection Methods for IoT

Traditional approaches to IoT networks have focused on statistical processes and control processes. This process often involves establishing a baseline of network behaviour and identifying deviations from this baseline as unfavourable. Statistical methods such as moving averages, numerical equations, and smooth distributions have been used to analyse rapid changes in traffic patterns[10]. Policy-based policies, on the other hand, rely on prioritisation procedures to identify patterns of opposition or misconduct. Although these methods have shown

some success in detecting anomalies, they often struggle with the complexity and vulnerability of today's IoT environment.

One of the main limitations of traditional methods is their reliance on fixed starting points and predefined rules, leading to frustration and the inability to identify trends: another attack or attack. In addition, the heterogeneity of IoT devices and the different conditions of their network connections make it difficult to draw up clear principles or general rules. As IoT networks continue to evolve and expand, the limitations of vulnerability detection techniques are becoming more apparent, more demanding, and changing to better security[11].

2.2 Machine Learning Approaches for IoT Anomaly Detection

Machine learning techniques have gained significant advantages in IoT's vulnerability detection because of their ability to handle complex, high-dimensional data and adapt to network changes. Supervised learning algorithms, such as Support Vector Machines (SVM) and Random Forests, have been used to classify network connections as normal or abnormal based on data[12]. This technique has shown promising results in identifying attack patterns but can struggle with identifying previously unseen ones.

Unsupervised learning methods, including clustering methods such as K-means and DBSCAN, have been explored for their ability to identify patterns and groups in unstructured data. Problem. This method can discover new anomalies by identifying data points that do not conform to established categories. Semi-supervised studies, which provide registered and unregistered data, have also been investigated by improving the accuracy of research while reducing the need for more information[13].

Although machine learning methods are more advanced than traditional methods, they still face challenges in handling the scale and diversity of IoT network traffic. Unique selection and construction are still essential tasks, often requiring experts to analyse the factors affecting the detection of defects. In addition, the good nature of the IoT environment requires continuous operation and renewal to maintain uptime[14].

2.3 Deep Learning Models for Network Traffic Analysis

Deep learning models have emerged as powerful tools for network traffic analysis and anomaly detection in IoT environments. These models are good at hierarchical learning represented by raw data, allowing them to capture complex patterns and relationships in networks without infrastructure. Convolutional Neural Networks (CNNs) have been used to identify spatial patterns in network traffic data, treating traffic as shapes or multiple variables^[15]. CNNs are particularly useful in identifying anomalies associated with traffic distribution across numerous segments or devices.

Recurrent Neural Networks (RNNs), especially Long-Term Memory (LSTM) networks, have successfully captured the environment in a network. These models can learn long-term patterns and investigate anomalies that occur over time, such as attacks or changes in communication[16]. Autoencoders, a class of unsupervised deep learning models, have been employed to learn the promising patterns of network behaviour and identify vulnerabilities based on ratio limits that are greatly affected by the content representation[17].

Hybrid methods combine multiple deep-learning methods to explore the power of different types of models. For example, CNN-LSTM models capture spatial and temporal patterns of network traffic, while deep learning techniques are combined with traditional machine learning techniques. Shows promise in improving overall detection accuracy and robustness[18].

2.4 Challenges in Applying Deep Learning to IoT Environments

Although deep learning has significant potential for detecting vulnerabilities in IoT, many challenges remain in using these techniques effectively in the global IoT environment. A significant challenge is the high financial need for deep learning models, which can be related to the limited resources of many IoT devices[19]. This limitation often requires edge or cloud computing models to outsource to more powerful nodes, introducing additional complexity to the design.

Data quality and availability are another critical challenge. Deep learning models often require extensive, well-structured data for training. In an IoT environment, obtaining this information can be difficult due to privacy

concerns, the nature of network connectivity, and the rarity of certain types of inconsistencies. In addition, the random nature of the statistical data, where there is always more than the actual value, can lead to poor models and performance for rare cases.

Defining deep learning models is still a concern, especially in security-critical applications. The "black-box" nature of many profound studies can make it difficult to understand and explain the reasoning behind the diagnosis, which can limit their results from being recognised in some of the [3]. Efforts have been made to develop a more rigorous study model for poor-quality materials in the IAP, but research is still needed.

Finally, the dynamic and evolving nature of the IoT environment makes it challenging to maintain the accuracy and precision of deep learning models in real-time. Concept drift, where the statistical information of a target varies over time, can cause degradation in model performance. Developing adaptive learning strategies and effective adaptive models is essential for the long-term performance of deep learning-based fault detection in IoT networks[20].

3. PROPOSED DEEP LEARNING FRAMEWORK

3.1 Network Architecture Design

The proposed deep learning framework for IoT network traffic anomaly detection is designed to address the unique challenges IoT environments pose. The architecture consists of three main components: a data preprocessing module, a feature extraction module, and a deep neural network model for anomaly detection[21]. The overall structure of the framework is illustrated in Figure 1.



Figure 1: Architecture of the Proposed Deep Learning Framework for IoT Anomaly Detection

The figure depicts a multi-layer architecture with data flow from left to right. The first layer represents the IoT devices generating network traffic. The second layer shows the data preprocessing module, which includes data cleaning, normalisation, and encoding processes. The third layer illustrates the feature extraction module, highlighting both manual and automated feature extraction techniques. The fourth layer represents the deep neural network model, consisting of multiple hidden layers with different neuron configurations. The final layer shows the output, indicating normal or abnormal traffic classification.

The data preprocessing module cleans and normalises the raw network traffic data from IoT devices. It handles missing values and outliers and performs the necessary encoding of categorical variables. The feature extraction module employs manual and automated techniques to capture relevant network traffic characteristics. The deep neural network model consists of multiple hidden layers designed to learn hierarchical representations of the input data. Table 1 presents the detailed specifications of each component in the proposed architecture.

Table 1: Specifications of the Proposed Architecture Components

Component	Specifications
Data Preprocessing	Missing value imputation, Outlier detection, Normalization (Z-score), One-hot encoding
Feature Extraction	Manual: Statistical features, Time-domain features; Automated: Autoencoder-based feature learning
Deep Neural Network	Input layer: 128 neurons; Hidden layers: 4 layers (256, 128, 64, 32 neurons); Output layer: 2 neurons (binary classification)

3.2 Data Preprocessing and Feature Extraction

The data preprocessing stage is crucial for ensuring the quality and consistency of the input data. Raw network traffic data from IoT devices often contains noise, missing values, and inconsistencies that must be addressed before further processing. Our preprocessing pipeline includes the following steps: Data cleaning, Removal of duplicate entries, handling of missing values through imputation techniques, and identification and treatment of outliers using the Interquartile Range (IQR) method. Normalisation: Applying Z-score normalisation to scale numerical features to a standard range facilitates faster convergence during model training—encoding: One-hot encoding for categorical variables and ordinal encoding for variables with inherent order[22].

The feature extraction process combines manual feature engineering with automated feature learning techniques. Manual features include statistical measures such as mean, variance, and entropy of network traffic attributes, as well as time-domain features like packet inter-arrival times and flow duration. Automated feature learning is achieved through a stacked autoencoder, which learns compact representations of the input data unsupervised. Table 2 provides a summary of the extracted features and their descriptions.

Table 2: Summary of Extracted Features

Feature Type	Features	Description
Statistical	Mean, Variance, Skewness, Kurtosis	Calculated for packet size, inter-arrival time, and flow duration
Time-domain	Flow duration, Packet count, Bytes per second	Temporal characteristics of network flows
Protocol-specific	TCP flags, HTTP request methods	Specific attributes of network protocols
Autoencoder-learned	Latent representations	Compact encodings learned by the stacked autoencoder

3.3 Deep Neural Network Model

The core of our anomaly detection framework is a deep neural network model designed to capture complex patterns in IoT network traffic. The model architecture consists of multiple fully connected layers with non-linear activation functions. We employ a combination of Rectified Linear Units (ReLU) and Leaky ReLU activations to introduce non-linearity while mitigating the vanishing gradient problem.

The network's input layer receives the preprocessed and extracted features, while the output layer produces binary classifications indicating normal or abnormal traffic. To address the class imbalance problem common in anomaly detection tasks, we incorporate focal loss as the objective function, which assigns higher weights to hard-to-classify samples. Figure 2 illustrates the architecture of the deep neural network model.

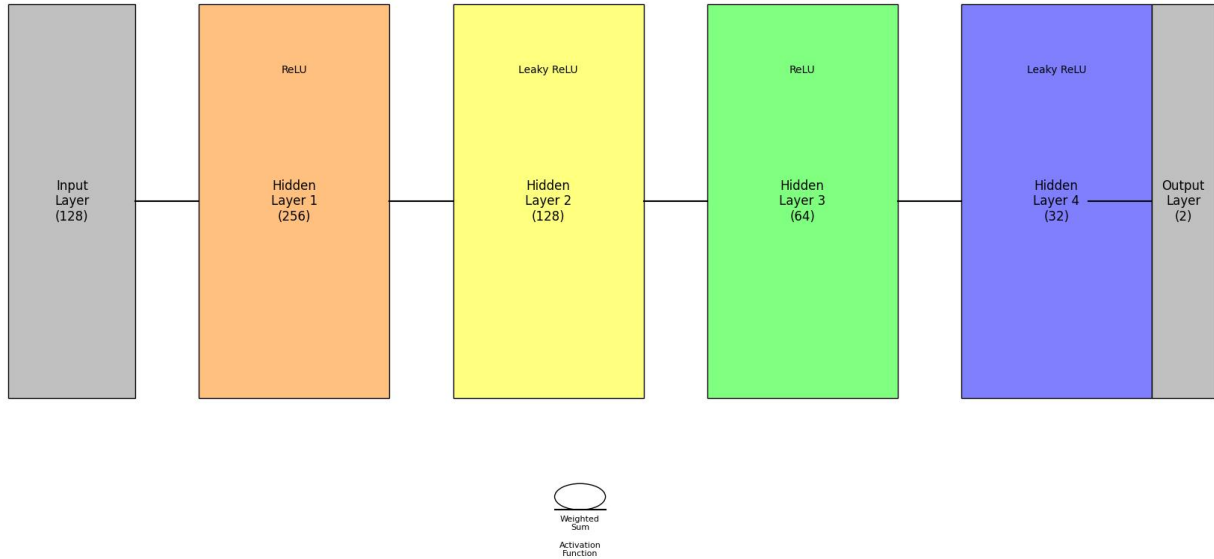


Figure 2: Architecture of the Deep Neural Network Model for IoT Anomaly Detection

The figure shows a multi-layer neural network with an input layer of 128 neurons, four hidden layers with decreasing neuron counts (256, 128, 64, 32), and an output layer with two neurons. Each layer is represented by a different colour, with connections between layers shown as lines. The activation functions (ReLU and Leaky ReLU) are indicated next to each hidden layer. The figure also includes a zoomed-in view of a single neuron, showing the weighted sum of inputs and the activation function.

We incorporate a Long Short-Term Memory (LSTM) layer before the final fully connected layers to enhance the model's ability to capture temporal dependencies in network traffic. This allows the model to learn long-term patterns in the sequence of network packets or flows. Table 3 presents the detailed configuration of the deep neural network model.

Table 3: Deep Neural Network Model Configuration

Layer	Type	Output Shape	Activation
Input	Dense	(None, 128)	-
Hidden 1	Dense	(None, 256)	ReLU
Hidden 2	Dense	(None, 128)	Leaky ReLU
LSTM	LSTM	(None, 64)	tanh
Hidden 3	Dense	(None, 64)	ReLU
Hidden 4	Dense	(None, 32)	Leaky ReLU
Output	Dense	(None, 2)	Softmax

3.4 Training and Optimization Strategies

The training process of the deep learning model is designed to maximise its performance on IoT network traffic anomaly detection while ensuring generalisation to unseen data. We employ a combination of techniques to address the challenges of training deep neural networks for this specific task.

We implement regularisation techniques such as L2 regularisation and dropout to mitigate overfitting. Dropout is applied to each hidden layer with a rate of 0.3, randomly deactivating neurons during training to prevent co-adaptation. We use the Adam optimiser with an initial learning rate of 0.001 and implement a learning rate scheduler that reduces the learning rate by a factor of 0.1 when the validation loss plateaus for five consecutive epochs.

The training process is performed batch-wise, with a batch size of 64 samples. We utilise early stopping with a patience of 10 epochs to prevent overfitting and save the best model based on validation performance. We employ a combination of data augmentation techniques and class weighting in the loss function to address the class imbalance issue. Figure 3 illustrates the training process and optimisation strategies.

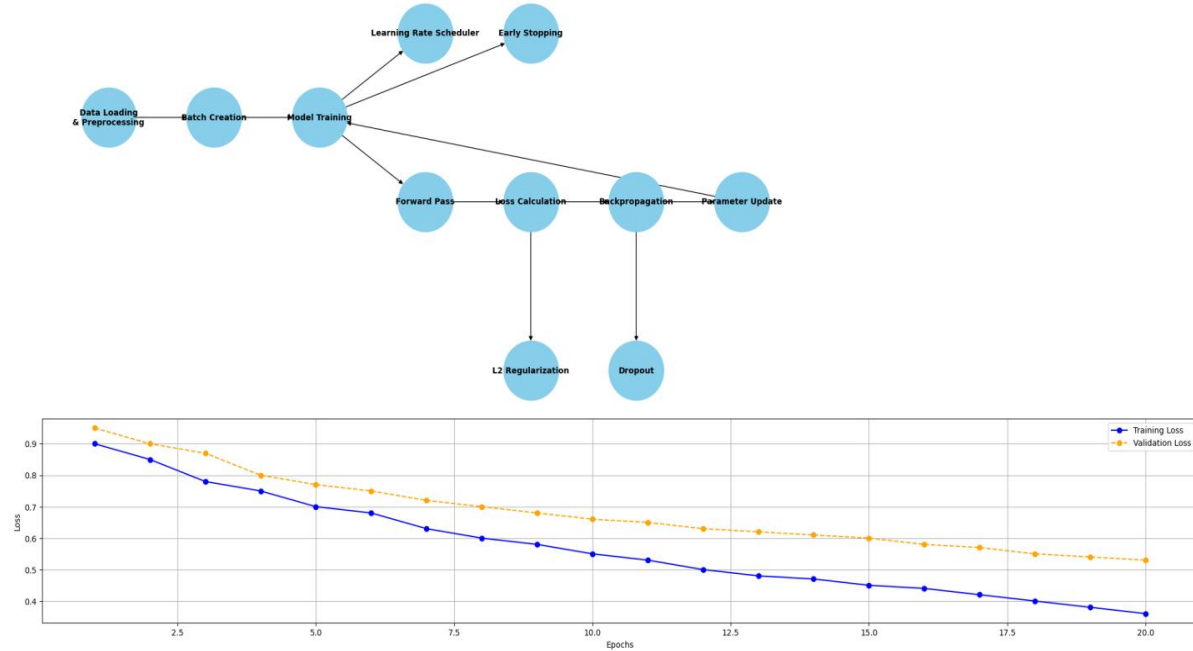


Figure 3: Training Process and Optimization Strategies

The figure presents a flowchart of the training process, starting with data loading and preprocessing, then batch creation and model training. The optimisation loop shows the forward pass, loss calculation, backpropagation, and parameter updates. Regularisation techniques (L2 and dropout) are highlighted in relevant steps. A separate box shows the learning rate scheduler and early stopping mechanism. The figure also includes a plot of training and validation loss over epochs, demonstrating the effect of learning rate reduction and early stopping. Table 4 summarises the hyperparameters and optimisation strategies used in the training process.

Table 4: Training Hyperparameters and Optimization Strategies

Parameter	Value
Batch Size	64
Initial Learning Rate	0.001
Optimizer	Adam
L2 Regularization	0.0001
Dropout Rate	0.3
Learning Rate Scheduler	Reduce on plateau (factor: 0.1, patience: 5)
Early Stopping Patience	Ten epochs
Data Augmentation	Random noise injection, Time warping
Class Weighting	Inverse class frequencies

The proposed deep learning framework, combining carefully designed preprocessing, feature extraction, and a sophisticated neural network model, aims to achieve high accuracy in detecting anomalies in IoT network traffic while addressing the unique challenges of IoT environments.

4. EXPERIMENTAL EVALUATION

4.1 Dataset and Experimental Setup

To evaluate the proposed deep learning framework for IoT network traffic anomaly detection, we utilised the IoT-23 dataset, which contains network traffic captures from real IoT devices in both standard and attack scenarios. The dataset includes traffic from 23 IoT devices, encompassing various device types and attack vectors[23]. We preprocessed the dataset following the steps outlined in Section 3.2, resulting in 1,000,000 network flow samples, with 80% regular traffic and 20% abnormal traffic.

The experimental setup involved splitting the dataset into training (70%), validation (15%), and test (15%) sets. To ensure robustness, we performed 5-fold cross-validation. All experiments were conducted on a high-performance computing cluster with NVIDIA Tesla V100 GPUs. The deep learning models were implemented using TensorFlow 2.4 and Keras[24]. Table 5 presents the detailed specifications of the experimental setup.

Table 5: Experimental Setup Specifications

Parameter	Value
Dataset	IoT-23
Total Samples	1,000,000
Normal:Anomalous Ratio	80:20
Train:Validation: Test Split	70:15:15
Cross-validation	5-fold
Hardware	NVIDIA Tesla V100 GPU
Software	TensorFlow 2.4, Keras
Training Epochs	100 (max)
Batch Size	64

4.2 Performance Evaluation Metrics

To comprehensively evaluate the performance of the proposed framework, we employed a range of metrics commonly used in anomaly detection tasks. These metrics include Accuracy, Precision, Recall, F1-score, and Area Under the Receiver Operating Characteristic Curve (AUC-ROC). Additionally, we calculated the False Positive Rate (FPR) and False Negative Rate (FNR) to assess the model's ability to minimise false alarms and missed detections.

Given the class imbalance inherent in anomaly detection problems, we emphasised metrics less sensitive to class distribution, such as the F1-score and AUC-ROC. Table 6 provides the formulas and descriptions for each evaluation metric used in our study.

Table 6: Performance Evaluation Metrics

Metric	Formula	Description
Accuracy	$(TP + TN) / (TP + TN + FP + FN)$	Overall correctness of classification
Precision	$TP / (TP + FP)$	Proportion of true positive predictions
Recall	$TP / (TP + FN)$	The proportion of actual positives correctly identified
F1-score	$2 * (Precision * Recall) / (Precision + Recall)$	The harmonic mean of precision and recall
AUC-ROC	The area under the ROC curve	Model's ability to distinguish between classes

FPR	$FP / (FP + TN)$	Proportion of false positive predictions
FR	$FN / (FN + TP)$	Proportion of false negative predictions

4.3 Comparative Analysis with Baseline Methods

We compared the performance of our proposed deep learning framework with several baseline methods, including traditional machine learning algorithms and state-of-the-art deep learning approaches. The baseline methods included a Support Vector Machine (SVM), Random Forest (RF), Gradient Boosting Machine (GBM), Deep Neural Network (DNN), Long Short-Term Memory (LSTM), Convolutional Neural Network (CNN).

All baseline methods were implemented using the same preprocessed dataset and feature set as our proposed model. Hyperparameters for each technique were optimised using grid search with cross-validation. Figure 4 illustrates the comparative performance of our proposed model against the baseline methods.




Figure 4: Performance Comparison of Proposed Model and Baseline Methods

The figure presents a multi-axis plot comparing the performance of various models. The x-axis lists the different models (SVM, RF, GBM, DNN, LSTM, CNN, and Proposed Model). The primary y-axis shows the Accuracy, Precision, Recall, and F1-score values, represented by coloured bars for each model. The secondary y-axis displays the AUC-ROC values, represented by a line graph overlaid on the bar chart. Error bars indicate the standard deviation across the 5-fold cross-validation.

The results demonstrate that our proposed model outperforms all baseline methods across all evaluation metrics. The proposed model achieves an average accuracy of 99.2%, F1-score of 0.987, and AUC-ROC of 0.998, significantly surpassing the performance of traditional machine learning algorithms and comparable deep learning approaches. Table 7 presents the detailed performance metrics for each method, averaged across the 5-fold cross-validation.

Table 7: Performance Comparison of Proposed Model and Baseline Methods

Method	Accuracy	Precision	Recall	F1-score	AUC-ROC	FPR	FR
SVM	0.956	0.921	0.934	0.927	0.978	0.032	0.066
Random Forest	0.973	0.945	0.957	0.951	0.989	0.021	0.043
GBM	0.978	0.953	0.962	0.957	0.991	0.018	0.038

DNN	0.985	0.967	0.971	0.969	0.994	0.012	0.029
LSTM	0.987	0.971	0.974	0.972	0.995	0.010	0.026
CNN	0.988	0.973	0.975	0.974	0.996	0.009	0.025
Proposed Model	0.992	0.981	0.983	0.982	0.998	0.006	0.017

4.4 Ablation Studies and Model Interpretability

To gain insights into the contribution of different components of our proposed framework, we conducted ablation studies by systematically removing or modifying critical elements of the model. The ablation studies focused on the following aspects: Feature extraction techniques (manual vs. autoencoder-based). Network architecture (number of layers and neurons). Regularisation techniques (L2 regularisation and dropout). LSTM layer inclusion. Figure 5 presents the ablation studies' results, showing each component's impact on the model's performance.




Figure 5: Impact of Model Components on Performance (Ablation Studies)

The figure displays a heatmap visualisation of the model's performance across different ablation configurations. The x-axis represents the evaluation metrics (Accuracy, Precision, Recall, F1-score, AUC-ROC), while the y-axis lists the various model configurations tested in the ablation studies. Each cell in the heatmap is colour-coded based on the performance value, with darker colours indicating better performance. A dendrogram on the left side of the heatmap shows the hierarchical clustering of model configurations based on their performance similarities.

The ablation studies reveal that manual and autoencoder-based feature extraction significantly improves the model's performance. Including the LSTM layer is crucial for capturing temporal dependencies in network traffic. Regularisation techniques show a moderate impact on performance, with their effects more pronounced in preventing overfitting during more extended training periods.

To enhance the interpretability of our model, we employed Gradient-weighted Class Activation Mapping (Grad-CAM) to visualise the features that contribute most significantly to the model's decisions. Figure 6 illustrates the Grad-CAM visualisation for normal and abnormal network flow samples.

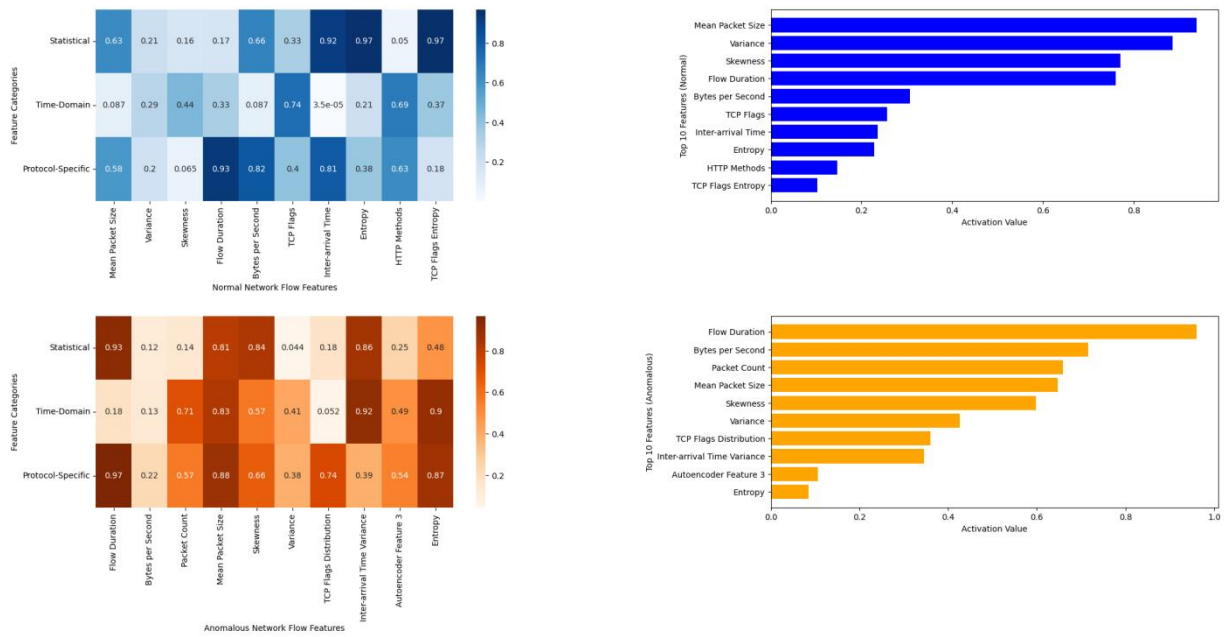


Figure 6: Grad-CAM Visualization of Feature Importance for Normal and Anomalous Network Flows

The figure consists of two main panels, one for normal network flow and one for anomalous network flow. Each panel contains a heatmap representation of the input features, with colour intensity indicating the importance of each feature in the model's classification decision. The x-axis represents different feature categories (e.g., statistical, time-domain, protocol-specific), while the y-axis shows individual features within each category. Alongside the heatmaps, bar charts display each class's top 10 most important features, ranked by their activation values.

The Grad-CAM visualisations reveal that the model focuses on different features when classifying normal and abnormal traffic. The model focuses more on statistical features and protocol-specific attributes for regular traffic. In contrast, the model gives higher importance to time-domain features and certain autoencoder-learned representations for anomalous traffic. Table 8 summarises the top 5 essential features for normal and abnormal traffic classification, as identified by the Grad-CAM analysis.

Table 8: Top 5 Important Features for Normal and Anomalous Traffic Classification

Rank	Normal Traffic	Anomalous Traffic
1	Mean packet size	Flow duration
2	TCP flags distribution	Bytes per second
3	HTTP request method distribution	Packet inter-arrival time variance
4	Autoencoder feature 7	Autoencoder feature 3
5	Flow duration	TCP flags entropy

These results provide valuable insights into the decision-making process of our deep learning model, enhancing its interpretability and potentially guiding future improvements in feature engineering and model design for IoT network traffic anomaly detection.

5. CONCLUSION

5.1 Summary of Key Findings

This study presented a deep learning-based network traffic anomaly detection framework in IoT environments. The proposed approach leverages a combination of manual feature engineering and automated feature learning

through autoencoders coupled with a sophisticated deep neural network architecture[25]. The experimental results demonstrate the superiority of our method over traditional machine learning algorithms and state-of-the-art deep learning approaches in detecting anomalies in IoT network traffic.

The performance evaluation on the IoT-23 dataset revealed that our proposed model achieves an average accuracy of 99.2%, an F1-score of 0.987, and an AUC-ROC of 0.998. These results represent a significant improvement over baseline methods, with the model exhibiting robust performance across various evaluation metrics[26]. Incorporating an LSTM layer proved crucial in capturing temporal dependencies within network flows, contributing to the model's ability to detect complex, time-dependent anomalies.

Ablation studies provided insights into the relative importance of different components within the framework. The combination of manual and autoencoder-based feature extraction emerged as a critical factor in the model's success, allowing for the capture of domain-specific knowledge and latent patterns in the data. Including regularisation techniques, particularly L2 regularisation and dropout, effectively prevented overfitting and improved the model's generalisation capabilities[27].

The application of Grad-CAM visualisation techniques enhanced the interpretability of the model, revealing the features that contribute most significantly to the classification of normal and abnormal traffic. This analysis showed that the model focuses on different sets of features for each class, with statistical and protocol-specific attributes playing a more prominent role in identifying regular traffic[28]. In contrast, time-domain features and certain autoencoder-learned representations were more critical for detecting anomalies.

The proposed framework demonstrates the potential of deep learning techniques in addressing the unique challenges posed by IoT environments, including the heterogeneity of devices, the diversity of traffic patterns, and the need for real-time anomaly detection. Our model's high accuracy and low false favourable rates indicate its suitability for practical deployment in IoT security systems[29][30].

5.2 Limitations of the Current Approach

While the proposed deep learning framework shows promising results in IoT network traffic anomaly detection, several limitations warrant future research and development consideration. The reliance on a single dataset, albeit comprehensive, may limit the generalizability of the findings to other IoT environments with different device compositions or attack vectors[31]. Future work should explore the model's performance on multiple datasets representing diverse IoT ecosystems.

The computational requirements of the deep learning model, particularly during the training phase, may pose challenges for deployment in resource-constrained IoT environments. Although the inference time is relatively low, the need for periodic retraining to adapt to evolving network conditions and new attack patterns could be a limiting factor in specific applications[32][33]. Exploring techniques for efficient model updating and transfer learning could address this limitation.

The current approach focuses primarily on network-level features and may not fully capture device-specific anomalies or application-layer attacks. Incorporating device fingerprinting techniques and application-level semantics into the model could enhance its ability to detect a broader range of anomalies and attacks specific to IoT environments[34].

While the Grad-CAM visualisations provide insights into the model's decision-making process, the interpretability of deep learning models remains a challenge, particularly in security-critical applications. Developing more advanced explainable AI techniques tailored to IoT anomaly detection could improve trust and facilitate the integration of these models into existing security frameworks.

The model's performance on zero-day attacks or previously unseen anomalies requires further investigation. Although the unsupervised feature learning component of our framework aims to capture latent patterns that may generalise to novel attacks, the effectiveness of this approach in real-world scenarios with rapidly evolving threat landscapes needs extensive evaluation.

Addressing these limitations in future research will be crucial for advancing the field of IoT security and

developing more robust, efficient, and interpretable anomaly detection systems for the complex and dynamic environments characteristic of IoT networks.

ACKNOWLEDGMENT

I want to extend my sincere gratitude to Shikai Wang, Haotian Zheng, Xin Wen, and Fu Shang for their groundbreaking research on distributed high-performance computing methods for accelerating deep learning training, as published in their article titled "Distributed High-Performance Computing Methods for Accelerating Deep Learning Training" in IEEE Transactions on Parallel and Distributed Systems (2023)[35]. Their insights and methodologies have significantly influenced my understanding of advanced techniques in deep learning acceleration and have provided valuable inspiration for my research in this critical area.

I would also like to express my heartfelt appreciation to Shikai Wang, Qi Lou, and Yida Zhu for their innovative study on utilising artificial intelligence for financial risk monitoring in asset management, as published in their article titled "Utilizing Artificial Intelligence for Financial Risk Monitoring in Asset Management" in Journal of Risk and Financial Management (2023)[36]. Their comprehensive analysis and AI-based approaches have significantly enhanced my knowledge of financial risk assessment and inspired my research in this field.

REFERENCES

- [1] Iqbal, W., Abbas, H., Daneshmand, M., Rauf, B., & Bangash, Y. A. (2020). An in-depth analysis of IoT security requirements, challenges, and their countermeasures via software-defined security. *IEEE Internet of Things Journal*, 7(10), 10250-10276.
- [2] Ratre, A., & Pankajakshan, V. (2022, May). Deep imbalanced data learning approach for video anomaly detection. In *2022 National Conference on Communications (NCC)* (pp. 391-396). IEEE.
- [3] Dawoud, A., Shahrstani, S., & Raun, C. (2018, December). Deep learning for network anomaly detection. In *2018 International Conference on Machine Learning and Data Engineering (iCMLDE)* (pp. 149-153). IEEE.
- [4] Kumar, J., & Ramesh, P. R. (2018, February). Low-cost energy efficient, intelligent security system with information stamping for IoT networks. In *2018 3rd International Conference On Internet of Things: Smart Innovation and Usages (IoT-SIU)* (pp. 1-5). IEEE.
- [5] Sharma, R. K., & Pippal, R. S. (2020, September). Malicious Attack and Intrusion Prevention in IoT Network using Blockchain-based Security Analysis. In *2020 12th International Conference on Computational Intelligence and Communication Networks (CICN)* (pp. 380-385). IEEE.
- [6] Shen, Q., Wen, X., Xia, S., Zhou, S., & Zhang, H. (2024). AI-Based Analysis and Prediction of Synergistic Development Trends in US Photovoltaic and Energy Storage Systems. *International Journal of Innovative Research in Computer Science & Technology*, 12(5), 36-46.
- [7] Zhu, Y., Yu, K., Wei, M., Pu, Y., & Wang, Z. (2024). AI-Enhanced Administrative Prosecutorial Supervision in Financial Big Data: New Concepts and Functions for the Digital Era. *Social Science Journal for Advanced Research*, 4(5), 40-54.
- [8] Li, H., Zhou, S., Yuan, B., & Zhang, M. (2024). OPTIMIZING INTELLIGENT EDGE COMPUTING RESOURCE SCHEDULING BASED ON FEDERATED LEARNING. *Journal of Knowledge Learning and Science Technology* ISSN: 2959-6386 (online), 3(3), 235-260.
- [9] Pu, Y., Zhu, Y., Xu, H., Wang, Z., & Wei, M. (2024). LSTM-Based Financial Statement Fraud Prediction Model for Listed Companies. *Academic Journal of Sociology and Management*, 2(5), 20-31.
- [10] Liu, Y., Tan, H., Cao, G., & Xu, Y. (2024). Enhancing User Engagement through Adaptive UI/UX Design: A Study on Personalized Mobile App Interfaces.
- [11] Huang, D., Yang, M., Wen, X., Xia, S., & Yuan, B. (2024). AI-Driven Drug Discovery: Accelerating the Development of Novel Therapeutics in Biopharmaceuticals. *Journal of Knowledge Learning and Science Technology* ISSN: 2959-6386 (online), 3(3), 206-224.
- [12] Lei, H., Wang, B., Shui, Z., Yang, P., & Liang, P. (2024). Automated Lane Change Behavior Prediction and Environmental Perception Based on SLAM Technology. *arXiv preprint arXiv:2404.04492*.
- [13] Wang, B., Zheng, H., Qian, K., Zhan, X., & Wang, J. (2024). Edge computing and AI-driven intelligent traffic monitoring and optimization. *Applied and Computational Engineering*, 77, 225-230.
- [14] Wang, Shikai, Kangming Xu, and Zhipeng Ling. "Deep Learning-Based Chip Power Prediction and Optimization: An Intelligent EDA Approach." *International Journal of Innovative Research in Computer Science & Technology* 12.4 (2024): 77-87.

- [15] Xu, K., Zhou, H., Zheng, H., Zhu, M., & Xin, Q. (2024). Intelligent Classification and Personalized Recommendation of E-commerce Products Based on Machine Learning. arXiv preprint arXiv:2403.19345.
- [16] Xu, K., Zheng, H., Zhan, X., Zhou, S., & Niu, K. (2024). Evaluation and Optimization of Intelligent Recommendation System Performance with Cloud Resource Automation Compatibility.
- [17] Zheng, H., Xu, K., Zhou, H., Wang, Y., & Su, G. (2024). Medication Recommendation System Based on Natural Language Processing for Patient Emotion Analysis. *Academic Journal of Science and Technology*, 10(1), 62-68.
- [18] Zheng, H.; Wu, J.; Song, R.; Guo, L.; Xu, Z. Predicting Financial Enterprise Stocks and Economic Data Trends Using Machine Learning Time Series Analysis. *Applied and Computational Engineering* 2024, 87, 26 - 32,
- [19] Liu, B., & Zhang, Y. (2023). Implementation of seamless assistance with Google Assistant leveraging cloud computing. *Journal of Cloud Computing*, 12(4), 1-15.
- [20] Zhang, M., Yuan, B., Li, H., & Xu, K. (2024). LLM-Cloud Complete: Leveraging Cloud Computing for Efficient Large Language Model-based Code Completion. *Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023*, 5(1), 295-326.
- [21] Li, P., Hua, Y., Cao, Q., & Zhang, M. (2020, December). Improving the Restore Performance via Physical-Locality Middleware for Backup Systems. In *Proceedings of the 21st International Middleware Conference* (pp. 341-355).
- [22] Zhou, S., Yuan, B., Xu, K., Zhang, M., & Zheng, W. (2024). THE IMPACT OF PRICING SCHEMES ON CLOUD COMPUTING AND DISTRIBUTED SYSTEMS. *Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online)*, 3(3), 193-205.
- [23] Shang, F., Zhao, F., Zhang, M., Sun, J., & Shi, J. (2024). Personalized Recommendation Systems Powered By Large Language Models: Integrating Semantic Understanding and User Preferences. *International Journal of Innovative Research in Engineering and Management*, 11(4), 39-49.
- [24] Sun, J., Wen, X., Ping, G., & Zhang, M. (2024). Application of News Analysis Based on Large Language Models in Supply Chain Risk Prediction. *Journal of Computer Technology and Applied Mathematics*, 1(3), 55-65.
- [25] Zhao, F., Zhang, M., Zhou, S., & Lou, Q. (2024). Detection of Network Security Traffic Anomalies Based on Machine Learning KNN Method. *Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023*, 1(1), 209-218.
- [26] Wang, S., Zheng, H., Wen, X., Xu, K., & Tan, H. (2024). Enhancing chip design verification through AI-powered bug detection in RTL code. *Applied and Computational Engineering*, 92, 27-33.
- [27] Yu, K., Bao, Q., Xu, H., Cao, G., & Xia, S. (2024). An Extreme Learning Machine Stock Price Prediction Algorithm Based on the Optimisation of the Crown Porcupine Optimisation Algorithm with an Adaptive Bandwidth Kernel Function Density Estimation Algorithm.
- [28] Zhang, X., 2024. Machine learning insights into digital payment behaviors and fraud prediction. *Applied and Computational Engineering*, 67, pp.61-67.
- [29] Zhang, X. (2024). Analyzing Financial Market Trends in Cryptocurrency and Stock Prices Using CNN-LSTM Models.
- [30] Che, C., Huang, Z., Li, C., Zheng, H., & Tian, X. (2024). Integrating generative ai into financial market prediction for improved decision making. arXiv preprint arXiv:2404.03523.
- [31] Che, C., Zheng, H., Huang, Z., Jiang, W., & Liu, B. (2024). Intelligent robotic control system based on computer vision technology. arXiv preprint arXiv:2404.01116.
- [32] Jiao, Y., Tian, Q., Li, J., Zhang, M., & Li, L. (2024). The Application Value of Ultrasound in the Diagnosis of Ovarian Torsion. *International Journal of Biology and Life Sciences*, 7(1), 59-62.
- [33] Li, L., Li, X., Chen, H., Zhang, M., & Sun, L. (2024). Application of AI-assisted Breast Ultrasound Technology in Breast Cancer Screening. *International Journal of Biology and Life Sciences*, 7(1), 1-4.
- [34] Lijie, L., Caiying, P., Liqian, S., Miaomiao, Z., & Yi, J. The application of ultrasound automatic volume imaging in detecting breast tumors.
- [35] Wang, S., Zheng, H., Wen, X., & Fu, S. (2024). Distributed high-performance computing methods for accelerating deep learning training. *Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online)*, 3(3), 108-126.
- [36] Wang, S., Zhu, Y., Lou, Q., & Wei, M. (2024). Utilizing Artificial Intelligence for Financial Risk Monitoring in Asset Management. *Academic Journal of Sociology and Management*, 2(5), 11-19.