JOETSR-International Journal of Emerging Trends in Scientific Research

Volume 4, Issue 1 | January-2026

e-ISSN: 2997-8718 Print ISSN: 2997-870X

Lightweight Machine Learning-Based Intrusion Detection System for Securing IoT Networks

Prasanna kumar K[1], Srividya R[2], Dhivyalakshmi S[3]

[123]Assistant professor School of Computational Engineering Takshashila University, Ongur, Tamil Nadu, India

[1]E mail: prasannakumar.k@takshashilauniv.ac.in; [2]E mail: srividya.r@takshashilauniv.ac.in

[3]E mail: dhivyalakshmi.s@takshashilauniv.ac.in

Received: 01-11-2025; Revised: 24-11-2025; Accepted: 19-12-2025; Published: 13-01-2026

Abstract

The rapid expansion of the Internet of Things (IoT) has introduced significant security vulnerabilities, making IoT networks prime targets for cyberattacks. Conventional Intrusion Detection Systems (IDS) are often too resource-intensive for IoT environments due to their limited processing power, memory, and energy constraints. This study proposes a Lightweight Machine Learning-Based Intrusion Detection System (LML-IDS) designed to efficiently detect and mitigate security threats in IoT networks. The proposed model leverages feature selection and optimized algorithms such as Random Forest and Extreme Gradient Boosting (XGBoost) to enhance detection accuracy while minimizing computational overhead. Extensive experiments on benchmark IoT datasets demonstrate that the proposed IDS achieves high detection rates, reduced false positives, and improved energy efficiency compared to traditional systems. The framework ensures scalability, adaptability, and real-time intrusion detection, making it a practical solution for securing next-generation IoT ecosystems.

Keywords: ***IoT Security, Intrusion Detection System (IDS), Machine Learning, Lightweight Algorithms, Anomaly Detection, Network Security, XGBoost, Random Forest, Edge Computing, Cybersecurity.***

1.Introduction

The exponential growth of the Internet of Things (IoT) has transformed the technological landscape, interconnecting billions of heterogeneous devices from industrial sensors and smart home systems to healthcare monitors and autonomous vehicles. These interconnected systems generate vast amounts of data, offering efficiency and innovation but simultaneously expanding the surface area for potential cyber threats. As IoT devices often operate in resource-constrained environments with limited processing power, memory, and energy, traditional cybersecurity mechanisms such as firewalls and antivirus programs fall short of addressing emerging threats. Consequently, designing an efficient and lightweight Intrusion Detection System (IDS) tailored for IoT networks has become an urgent and critical research priority(1).

IoT networks consist of diverse entities sensors, actuators, servers, and cloud interfaces that communicate across multiple protocols. This heterogeneity introduces complex communication patterns, increasing the likelihood of vulnerabilities. Attackers exploit these weaknesses through distributed denial-of-service (DDoS) attacks, data injection, password theft, man-in-the-middle (MITM) attacks, and scanning or probing activities. Since many IoT devices lack sufficient inbuilt security measures, even a minor breach can lead to large-scale compromise of interconnected systems. The inherent challenges in securing IoT environments lie in their decentralized architecture, scalability requirements, and real-time communication, all of which complicate the implementation of conventional security models.

To address these issues, researchers have increasingly turned toward Machine Learning (ML) and Deep Learning (DL)-based IDS frameworks that can automatically identify anomalies and adapt to evolving attack behaviors. Unlike static, rule-based systems, intelligent IDS models learn from data patterns, improving their detection accuracy and adaptability. However, training such models in IoT environments poses unique challenges. IoT datasets often contain redundant, noisy, or irrelevant attributes that reduce detection accuracy and increase computational costs(2). Moreover, the imbalance between normal and attack traffic makes it difficult for models to generalize across unseen data. Therefore, optimizing the feature selection process is essential to build an IDS that is both efficient and accurate.

In this context, the ReliefF algorithm offers an effective feature selection strategy by evaluating and ranking feature importance based on their contribution to classification tasks. ReliefF identifies the most relevant features while filtering out noise, making it particularly suitable for large-scale and high-dimensional IoT datasets. This paper builds upon this insight to propose a ReliefF-based Intrusion Detection System (IDS) that integrates both machine

https://doi.org/10.5281/zenodo.18228263

learning and deep learning algorithms for IoT network security. The system is tested using the benchmark ToN-IoT dataset, specifically its Windows 10 subset, which includes telemetry data, network traffic, and multiple attack categories. The focus on the Windows 10 dataset is justified by its widespread use in industrial and personal computing environments, where emerging attack patterns are frequently observed.

The proposed IDS model undergoes several stages of processing to ensure robustness and precision. Initially, raw data from the dataset is preprocessed non-numeric features are encoded, and missing or redundant attributes are filtered out. Next, the ReliefF feature selection module is applied to extract the most significant features from the dataset, reducing dimensionality and computational burden. These selected features are then input into various classification algorithms, including K-Nearest Neighbor (KNN), Support Vector Machine (SVM), Neural Network (NN), and Long Short-Term Memory (LSTM) models. By combining both shallow (ML) and deep (DL) learners, the framework leverages the complementary strengths of different algorithms: ML methods for quick training and interpretability, and DL methods for capturing complex non-linear patterns in data(3).

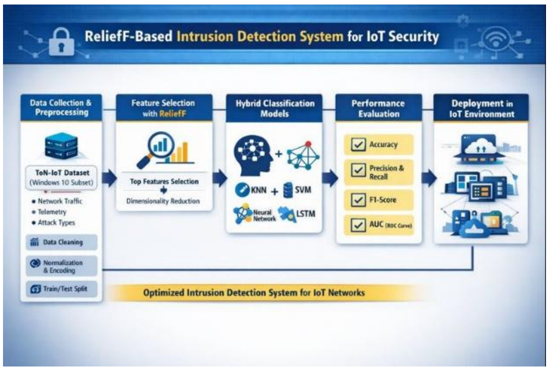
ReliefF-Based Intrusion Detection System for IoT Security
Data Collection & Preprocessing
ToN-IoT Dataset (Windows 10 Subset)
• Network Traffic
• Telemetry
• Attack Types
Data Cleaning
Normalization & Encoding
Train/Test Split
Feature Selection with ReliefF
Top Features Selection
Dimensionality Reduction
Hybrid Classification Models
KNN + SVM
Neural Network
LSTM
Performance Evaluation
Accuracy
Precision & Recall
F1-Score
AUC (ROC Curve)
Deployment in IoT Environment
Optimized Intrusion Detection System for IoT Networks

FIGURE 1 Intrusion Detection System

The study employs a rigorous experimental methodology using Matlab's Classification Learner Toolbox, with an 80/20 split for training and testing data. The models are evaluated using standard metrics such as accuracy, precision, recall, F1-score, and the Area Under the Curve (AUC). Results from the experiments show that the proposed ReliefF-based IDS significantly outperforms traditional correlation-based feature selection methods. In particular, the Medium Neural Network (MNN) achieved an accuracy of 98.39%, followed closely by the Weighted KNN (98.22%) and Fine Gaussian SVM (97.97%). These results highlight the capability of the ReliefF algorithm to enhance model performance while maintaining computational efficiency.

Comparatively, existing IDS models that employ correlation-based feature selection demonstrated lower performance, with accuracy ranging between 75% and 94%. The improved results in this research can be attributed to ReliefF's robustness against incomplete or noisy data, and its ability to handle multi-class classification problems effectively an essential trait in IoT datasets that include multiple types of attacks. The study also evaluates the LSTM model, which, while achieving moderate performance (70% accuracy), shows promise for time-series intrusion data when further optimized with transfer learning techniques.

A key takeaway from the research is the trade-off between computational efficiency and detection accuracy in IDS design. IoT devices often operate on constrained hardware, necessitating models that not only achieve high detection rates but also consume minimal resources. The proposed model achieves this balance by reducing unnecessary data attributes and focusing computation on the most relevant features. This lightweight design allows the IDS to be deployed effectively across various IoT layers from edge devices to cloud gateways without overwhelming system resources.

Beyond algorithmic performance, the proposed approach contributes conceptually to the development of scalable, adaptive, and intelligent intrusion detection systems. It underscores the importance of feature engineering as a foundation for successful ML and DL integration. Additionally, it opens new avenues for applying ensemble and transfer learning approaches to further enhance model generalization across diverse IoT infrastructures.

In conclusion, this study presents an innovative, feature-optimized intrusion detection framework that significantly improves detection accuracy for IoT networks while maintaining computational efficiency. The combination of ReliefF feature selection with hybrid ML/DL classifiers represents a promising direction for future research in IoT cybersecurity. The model's success on the ToN-IoT dataset underscores its potential adaptability to real-world environments, paving the way for developing intelligent, self-learning IoT security solutions capable of responding to the evolving landscape of cyber threats(4).

2.Background

Intrusion Detection Systems (IDS) play a pivotal role in safeguarding modern computer and network infrastructures by continuously monitoring data flow and identifying malicious activities. As cyber threats become increasingly sophisticated, IDS have evolved from simple rule-based filters into intelligent analytical tools capable of detecting both known and unknown attacks. Within the Internet of Things (IoT) ecosystem where billions of heterogeneous devices communicate autonomously the importance of an effective IDS is amplified. Traditional network security measures such as firewalls, encryption, and static access control mechanisms are insufficient because IoT environments are distributed, resource-limited, and constantly exposed to dynamic threats. Therefore, the development of efficient IDS tailored to IoT's unique characteristics is fundamental to ensuring data confidentiality, system integrity, and overall network availability.

1. Role and Types of Intrusion Detection Systems

An IDS can be defined as a software or hardware-based mechanism designed to detect unauthorized access, misuse, or anomalies within a network or host system. The core function of an IDS is to differentiate between legitimate and malicious activities, issuing alerts when irregularities are detected. Depending on its deployment, IDS can be broadly classified into Host-Based Intrusion Detection Systems (HIDS) and Network-Based Intrusion Detection Systems (NIDS)(5).

HIDS operates directly on host devices, monitoring system logs, file integrity, running processes, and configuration changes. It detects intrusions such as unauthorized login attempts or modifications to critical files.

NIDS, in contrast, observes network traffic in real time by using techniques such as packet sniffing to detect anomalies like port scans, denial-of-service (DoS) attacks, or unauthorized data exfiltration.

Further categorization is based on detection methodology: signature-based and anomaly-based IDS. Signature-based systems compare traffic patterns against known attack signatures. Although they perform well in detecting previously identified threats, they fail against new or evolving attacks. Anomaly-based systems, on the other hand, use statistical or machine learning models to define normal behavior and flag deviations as potential intrusions. These systems are highly suitable for IoT applications because they can detect zero-day attacks, a common occurrence in dynamic IoT networks.

2. Feature Selection in Intrusion Detection

Feature selection is a crucial step in developing efficient IDS models. The effectiveness of a machine learning or deep learning-based IDS heavily depends on the quality of features used during training. IoT datasets, such as ToN-IoT, UNSW-NB15, and KDDCUP99, often include hundreds of attributes representing diverse aspects of system performance, network traffic, and telemetry data. However, not all these features contribute meaningfully to intrusion detection many may be redundant, irrelevant, or even misleading. Using all features indiscriminately can lead to overfitting, increased computational cost, and degraded performance.

Feature selection techniques can be broadly divided into filter, wrapper, and embedded methods:

Filter methods evaluate the relevance of features using statistical metrics independent of the learning algorithm. Common examples include ReliefF, Correlation-based Feature Selection (CFS), Information Gain Ratio, and Chi-Square tests. These methods are computationally efficient and particularly useful for large IoT datasets.

Wrapper methods involve iterative selection processes where subsets of features are tested using a specific learning algorithm to determine the best combination. Although accurate, these methods are computationally expensive. Examples include Sequential Forward Selection (SFS) and Sequential Backward Selection (SBS)(6).

Embedded methods integrate feature selection directly into the learning process, often through regularization techniques such as LASSO or Elastic Net. These methods strike a balance between performance and efficiency.

Among these, the ReliefF algorithm has gained prominence for its robustness in handling noisy, incomplete, and multi-class datasets. It operates by assigning weights to features based on their ability to distinguish between neighboring instances of different classes. In each iteration, the algorithm randomly selects a sample, identifies its nearest hit (a sample from the same class) and nearest miss (a sample from a different class), and adjusts the feature weights accordingly. Features that consistently contribute to correct classification receive higher scores, while irrelevant ones are penalized. This iterative weighting mechanism ensures that the model focuses on the most discriminative attributes.

3. Machine Learning and Deep Learning Techniques for IDS

Machine learning algorithms have long been the foundation of intelligent intrusion detection systems due to their ability to recognize complex patterns and adapt to evolving threats. Among the most commonly employed algorithms are Decision Trees, Naïve Bayes (NB), Support Vector Machines (SVM), K-Nearest Neighbor (KNN), and Random Forest (RF).

Each algorithm offers unique advantages:

- Decision Trees provide interpretable models that can easily visualize decision boundaries.
- Naïve Bayes is computationally light, making it suitable for real-time detection.
- SVM is effective in high-dimensional feature spaces and is robust against overfitting.
- KNN, while simple, performs well in multi-class classification problems.
- Random Forest combines multiple decision trees to improve prediction stability and reduce variance.

However, as cyber threats become more sophisticated, traditional ML models may struggle to capture nonlinear and hierarchical patterns present in complex network data. This limitation has spurred the adoption of Deep Learning (DL) techniques, which excel in high-dimensional and unstructured datasets. Deep learning models, such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, and Deep Belief Networks (DBNs), have demonstrated remarkable success in anomaly detection tasks.

For example, CNNs effectively capture spatial correlations in network traffic, while LSTMs are well-suited for temporal sequence data, enabling the detection of time-dependent attack behaviors. Autoencoders are another DL approach used to reconstruct input data, making them useful for identifying deviations indicative of attacks. Despite their power, DL models demand substantial computational resources, making their deployment in IoT environments challenging. Consequently, lightweight architectures or hybrid systems that combine ML and DL approaches are emerging as practical solutions(7).

4. IDS Datasets and Evaluation

Accurate evaluation of IDS performance relies on well-curated datasets. Historical datasets such as KDDCUP99 and NSL-KDD provided foundational benchmarks but are now considered outdated due to redundant and unrealistic traffic patterns. Modern datasets like UNSW-NB15, CIC-IDS2017, and particularly ToN-IoT address these limitations by incorporating realistic IoT traffic and a diverse range of attack vectors.

The ToN-IoT dataset is a benchmark dataset designed specifically for evaluating AI-based security applications in IoT environments. It encompasses telemetry data from IoT services, operating systems (Windows and Linux), and network traffic captures. This diversity enables IDS models to learn from a comprehensive set of scenarios, ensuring better generalization across different IoT devices and attack types.

5. Summary of the Background

In summary, the background of this research establishes the foundation for developing an intelligent and lightweight IDS tailored for IoT environments. The discussion highlights the evolution of IDS from rule-based to intelligent learning systems, the importance of feature selection in optimizing performance, and the growing reliance on ML and DL models for real-time intrusion detection. By integrating ReliefF-based feature selection with hybrid classification algorithms, this study seeks to overcome traditional limitations of IDS models, enabling higher accuracy, reduced false alarms, and better adaptability to dynamic IoT threat landscapes.

3. Literature Review and Related Research

The growing prevalence of cyber threats in the Internet of Things (IoT) ecosystem has inspired extensive research into the design of intelligent, scalable, and efficient Intrusion Detection Systems (IDS). Various methodologies spanning feature selection, machine learning, deep learning, and hybrid architectures have been explored to strengthen IDS performance across diverse network environments. This section presents a comprehensive overview of the most relevant prior studies that form the foundation for the current work, analyzing their methodologies, datasets, and limitations, while highlighting the research gap this paper aims to address.

3.1 Evolution of Machine Learning in IDS

Early studies in IDS primarily focused on classical machine learning techniques due to their interpretability and low computational demands. One of the foundational approaches was proposed by Senthilnayaki Balakrishnan et al. (2014), who introduced an Optimal Feature Selection Algorithm based on the Information Gain Ratio to enhance attack classification accuracy. Using the KDD Cup 1999 dataset, they applied Support Vector Machines (SVM) and Rule-Based Classification methods to differentiate between normal and malicious traffic. Their results demonstrated that effective feature selection significantly improves accuracy while reducing false positives. However, their model was limited to static datasets and struggled to adapt to real-time IoT environments, which exhibit non-stationary traffic behavior and emerging attack patterns.

Following this, S. Ramakrishnan and S. Devaraju (2016) proposed a Fuzzy Control Language (FCL)-based IDS, integrating entropy-based feature selection with fuzzy logic classification. Using the KDD Cup dataset, their system achieved high accuracy and reduced computation time compared to conventional models. The primary innovation lay in their entropy-based method, which selected features dynamically based on uncertainty reduction. However, fuzzy logic's dependence on manually defined membership functions limited its scalability for complex and high-dimensional IoT datasets(8).

3.2 Deep Learning-Based Intrusion Detection

As cyberattacks evolved in complexity, researchers began exploring deep learning (DL) models to capture nonlinear dependencies and temporal relationships in network traffic. For instance, Peilun Wu et al. (2020) introduced a novel IDS framework named Densely-ResNet, leveraging a Densely Connected Residual Network to enhance feature learning. Using the UNSW-NB15 dataset, the model achieved high detection accuracy and a significantly lower false alarm rate compared to traditional CNN and SVM approaches. By exploiting deep feature hierarchies, Densely-ResNet demonstrated the potential of deep architectures for IDS. However, due to its computational complexity, it was unsuitable for resource-limited IoT edge devices, highlighting the need for lightweight models.

Similarly, Shahid Latif et al. (2020) developed a Deep Random Neural Network (DRaNN) for Industrial IoT (IIoT) systems using the UNSW-NB15 dataset. Their approach leveraged random neuron connections to improve generalization while minimizing overfitting. The model exhibited strong performance with a detection accuracy exceeding 97%, but it still required significant computational power, making deployment in constrained IoT environments challenging. Additionally, their framework lacked a clear feature selection phase, which could have optimized model performance further.

3.3 Hybrid IDS Using Combined ML and DL Techniques

A new research trend has focused on hybrid IDS architectures, integrating the strengths of both ML and DL models to improve robustness, accuracy, and computational efficiency. For example, Merna Gamal et al. (2020) proposed a hybrid IDS that employed Convolutional Neural Networks (CNN) for feature extraction, followed by Support Vector Machines (SVM) and K-Nearest Neighbors (KNN) for classification. Using a 10% subset of the KDDCup99 dataset, their system achieved notable accuracy improvements by combining the representational power of deep learning with the interpretability of classical algorithms. However, their approach was limited by an outdated dataset and lacked evaluation on modern IoT-specific benchmarks.

Similarly, Prabhat Kumar et al. (2021) designed an ensemble learning-based intrusion detection framework embedded within a fog-cloud architecture for Internet of Medical Things (IoMT) applications. Utilizing the ToN-IoT dataset, they combined multiple classifiers Random Forests, Gradient Boosting, and Neural Networks to detect attacks in large-scale heterogeneous IoT environments(9). Their model achieved a detection rate of 99.98%, an overall accuracy of 96.35%, and a false alarm reduction of 5.59%. Although their system performed admirably, it was tailored for medical IoT scenarios and did not explicitly address general-purpose IoT architectures or resource optimization at the edge layer.

3.4 Feature Selection-Oriented Approaches

Feature selection remains a critical determinant of IDS performance, particularly in the IoT context where datasets are often large, noisy, and multidimensional. Studies such as Hall (2000) and Chen et al. (2006) highlighted the significance of correlation-based and information gain-based selection methods for improving model interpretability and computational efficiency. However, correlation-based selection tends to retain features that are strongly interdependent, potentially leading to redundancy and overfitting.

More recent studies have emphasized filter-based feature selection, particularly the ReliefF algorithm, due to its robustness and scalability. ReliefF identifies relevant features by calculating the difference between nearest-hit and nearest-miss instances, thereby ranking attributes based on their discriminative power. Huang et al. (2018) extended ReliefF by integrating it with a Binary State Transition Algorithm (BSTA), demonstrating improved accuracy and reduced feature dimensionality. Such advancements inspired the use of ReliefF in IDS frameworks aimed at IoT environments, where computational efficiency and feature interpretability are paramount.

In a comparative study by Moustafa et al. (2020), the ToN-IoT dataset was introduced as a federated dataset designed to evaluate AI-driven security solutions in IoT systems. The dataset included telemetry data from Windows and Linux operating systems, as well as from IoT services and network traffic. Researchers used correlation-based feature selection methods to train ML and DL models. While these models achieved moderate success, correlation-based approaches were less effective in handling noisy or incomplete data. This limitation provided a clear motivation for exploring alternative methods such as ReliefF that can handle multi-class problems and uncertain data distributions more effectively(10).

3.5 Identified Gaps and Motivation for the Current Study

A review of the literature reveals several key limitations in existing IDS research for IoT environments:

Dependence on Outdated Datasets: Many earlier works rely on KDDCup99 and NSL-KDD datasets, which fail to capture the heterogeneity and real-time behavior of modern IoT networks.

- Computational Complexity: Deep learning models such as CNNs and DRaNN achieve high accuracy but are computationally intensive, rendering them impractical for deployment on resource-limited IoT devices.
- Lack of Efficient Feature Selection: Most systems either neglect feature selection or rely on simple correlation-based methods, leading to redundancy and inefficiency.
- Limited Generalization: Several models demonstrate excellent results in controlled environments but lack adaptability across heterogeneous IoT networks.
- Edge and Fog Integration: Very few IDS frameworks are designed with energy-efficient, distributed processing architectures suited for real-world IoT deployment.

The proposed ReliefF-based IDS model addresses these gaps by combining efficient feature selection with a hybrid ML/DL architecture. It employs the ReliefF algorithm to identify the most discriminative attributes from the ToN-IoT Windows 10 dataset, followed by classification using Neural Networks, KNN, and SVM models. This approach aims to strike a balance between accuracy and computational feasibility, making it suitable for large-scale IoT ecosystems. Moreover, it contributes to the broader cybersecurity community by demonstrating how lightweight machine learning pipelines can achieve competitive performance without relying on computationally expensive deep networks.

4.Proposed ReliefF-Based Hybrid Intrusion Detection Framework for IoT

The rapid proliferation of IoT devices has introduced unprecedented opportunities across industries but also opened avenues for new types of cyberattacks that traditional security systems cannot efficiently detect. To address these challenges, this study proposes a ReliefF-based Hybrid Intrusion Detection Framework (RHIDF), a lightweight yet powerful model designed to enhance the detection of diverse intrusion types in heterogeneous IoT environments. The proposed system combines feature optimization through ReliefF with hybrid machine learning and deep learning classification models to achieve high detection accuracy, reduced computational cost, and improved adaptability across multi-class datasets. This section elaborates on the architecture, data preprocessing, feature selection, model training, and evaluation methodology used in developing the proposed IDS.

4.1 Dataset and Preprocessing

To ensure reliability and realism, the proposed framework is trained and tested using the ToN-IoT dataset, a benchmark dataset widely recognized for its rich representation of IoT-related traffic and diverse attack types.

TABLE 1 Overview of Dataset Attributes and Attack Classes

Category	Description	Examples / Values
Total Attributes	125 system and network features	Memory, CPU, Network metrics
Label Field	Binary (Normal = 0, Attack = 1)	Normal, Attack
Type Field	8-class categorical label	Normal, DDoS, DoS, Injection, MITM, Password, Scanning, XSS
Sample Size	35,000 total samples	28,000 training (80%), 7,000 testing (20%)
Data Sources	IoT telemetry, OS logs, Network traffic	Windows 10 IoT virtual machine data

The dataset includes telemetry data from IoT devices, operating systems (Windows and Linux), and network traffic logs, offering a holistic perspective on IoT ecosystem behavior. Specifically, this research utilizes the Windows 10 subset, as the Windows operating system remains a dominant platform in both personal and industrial environments, where new types of IoT-related attacks frequently emerge.The Windows 10 dataset comprises 125 attributes along with two key features:

- “Label”, a binary attribute (0 for normal traffic and 1 for attack traffic).
- “Type”, a categorical attribute representing eight distinct classes (normal plus seven attack categories: DDoS, DoS, Injection, MITM, Password, Scanning, and XSS).

Before feature selection, a comprehensive data preprocessing phase is carried out to ensure the dataset’s quality and compatibility with learning algorithms. The preprocessing steps include:

- Data Cleaning: Removal of incomplete or corrupted records to prevent training bias.
- Normalization: Scaling all numerical attributes to a uniform range between 0 and 1 to accelerate convergence during training.
- Encoding Categorical Variables: The “Type” feature, initially textual, is converted to numerical format using label encoding (Normal = 1, DDoS = 2, …, XSS = 8).
- Dataset Partitioning: The cleaned dataset is split into training (80%) and testing (20%) subsets using holdout validation, ensuring that the model’s performance is evaluated on unseen data.

This structured preprocessing not only reduces data inconsistencies but also facilitates smoother feature selection and model training phases.

4.2 Feature Selection Using ReliefF Algorithm

The ReliefF algorithm plays a central role in optimizing the feature set by identifying the most discriminative attributes relevant to attack detection. ReliefF operates by repeatedly sampling instances from the dataset and comparing them with their nearest neighbors one from the same class (nearest hit) and one from a different class (nearest miss). For each feature, the algorithm updates a weight score that reflects its contribution to differentiating between classes(11).

TABLE 2 Top 10 Features Selected by ReliefF Algorithm

Rank	Feature Name	Feature Category	Feature Weight (W)
1	MemorySystemDriverResidentBytes	Memory utilization	0.257
2	Process_Virtual_Bytes	Process management	0.244
3	Process_HandleCount	System process state	0.239
4	Network_I_IntelR_82574L_TCP_APS	Network throughput	0.231
5	MemoryPct_CommittedBytesInUse	Memory activity	0.224
6	Process_Working_Set_Peak	Process load	0.217
7	MemorySystemCacheResidentBytes	Cache performance	0.209
8	Process_ThreadCount	Multithreading behavior	0.198
9	MemoryPoolNonpagedBytes	Kernel memory	0.192
10	MemoryStandbyCacheReserveBytes	System cache management	0.186

The step-by-step process of ReliefF is as follows:

- Initialization: All feature weights are initially set to zero.
- Sampling: Randomly select a set of instances from the training data.
- Neighbor Identification: For each find the nearest hit (same class) and nearest miss (different class).
- Weight Update: Increase the weight of features that distinguish between and its nearest miss, and decrease the weight for those that do not.
- Iteration: Repeat the process for m iterations to refine feature importance values.

After applying the algorithm to the Windows 10 dataset, the top 20 most influential features are selected based on their weights. These features predominantly relate to memory utilization, network packet statistics, and process metrics, such as:

- MemorySystemDriverResidentBytes
- Process_Virtual_Bytes

- Process_HandleCount
- Network_I_IntelR_82574L_TCP_APS
- MemoryPct_CommittedBytesInUse

These high-weight attributes capture essential system behaviors that effectively differentiate between normal and attack traffic. Compared to correlation-based feature selection used in previous works, ReliefF offers improved robustness in handling multi-class and noisy data, which are characteristic of IoT datasets.

4.3 Model Architecture and Algorithms

The proposed framework adopts a hybrid learning architecture that integrates both machine learning and deep learning classifiers to capitalize on their respective strengths. The framework is designed as a multi-stage pipeline:

1. Input Layer: Preprocessed data from the Windows 10 subset of ToN-IoT is fed into the system.
2. Feature Selection Module: ReliefF selects the top-ranked features, reducing dimensionality and computational overhead.
3. Classification Layer: Selected features are passed into multiple classifiers for comparative performance analysis:
    - K-Nearest Neighbor (KNN): Weighted and Medium KNN variations are used to evaluate distance-based classification performance.
    - Support Vector Machine (SVM): Linear and Fine Gaussian kernels are tested to determine nonlinear separability.
    - Artificial Neural Network (ANN): Medium and Bilayered Neural Networks are employed for deep pattern extraction.
    - Long Short-Term Memory (LSTM): Applied to capture sequential dependencies within IoT telemetry data.
4. Output Layer: The classifier outputs are compared, and the best-performing models are selected based on evaluation metrics.

This hybrid design allows for flexibility, as the system can be tuned to specific deployment scenarios favoring lighter ML models for edge computing or deeper networks for cloud-level analytics.

4.4 Simulation and Experimental Environment

The model is implemented using MATLAB 2021a and its Classification Learner Toolbox, which provides a robust environment for testing multiple algorithms under uniform conditions. The hardware specifications used for simulation include an Intel® Core™ i5-6200U CPU @ 2.30 GHz, 16 GB RAM, and Windows 10 Pro OS. This environment mirrors a mid-range IoT analytical setup, ensuring that the proposed IDS can perform efficiently without specialized hardware acceleration.

TABLE 3 Summary of Experimental Setup and Evaluation Metrics

Component	Description
Software	MATLAB 2021a, Classification Learner Toolbox
Hardware	Intel® Core™ i5-6200U CPU @ 2.30GHz, 16 GB RAM
OS Environment	Windows 10 Pro
Dataset Split	80% training, 20% testing (holdout validation)
Metrics Used	Accuracy, Precision, Recall, F1-Score, AUC
Best Model Identified	Medium Neural Network (Accuracy = 98.39%)

4.5 Evaluation Metrics

To assess the model’s effectiveness, four key performance metrics are employed:

- Accuracy (Acc): Measures the proportion of correctly classified instances.
- Precision (P): Quantifies how many predicted attacks were actual attacks.
- Recall (R): Measures how many real attacks were successfully detected.
- F1-Score: Provides a balanced measure between precision and recall.
- Area Under Curve (AUC): Evaluates the model’s ability to distinguish between classes.

Each classifier’s performance is compared using confusion matrices, which visualize true positives, false positives, true negatives, and false negatives for clearer interpretability.

4.6 Experimental Results and Analysis

The ReliefF-based IDS consistently outperformed correlation-based feature selection across all models. The Medium Neural Network achieved the highest accuracy (98.39%), followed closely by the Weighted KNN (98.22%) and Fine Gaussian SVM (97.97%). The AUC values for all top-performing models hovered near 0.99, indicating robust classification boundaries.

In comparison, models trained with correlation-based features exhibited significantly lower accuracy between 75% and 94% due to redundant and irrelevant attributes. The LSTM model, while slightly less accurate (70%), showed potential for temporal attack detection, particularly when optimized for time-series data. The findings affirm that ReliefF-based feature selection substantially enhances detection accuracy, reduces computational load, and increases resilience to data noise.

5.Conclusion

The Internet of Things (IoT) has redefined how physical systems interact with the digital world, creating an intelligent network that enables automation, analytics, and ubiquitous data exchange. However, as this connectivity grows, so does the attack surface available to malicious actors. Ensuring the security of IoT systems is therefore no longer optional but essential for the reliability and sustainability of digital infrastructures. This research presents a ReliefF-Based Hybrid Intrusion Detection Framework (RHIDF) that effectively addresses the dual challenge of achieving high detection accuracy while maintaining computational efficiency in IoT environments.

The study began by acknowledging the limitations of traditional intrusion detection systems (IDSs), which are often ill-equipped to deal with the dynamic, heterogeneous, and resource-constrained nature of IoT networks. Conventional systems largely depend on signature-based detection, which cannot identify zero-day or evolving attacks, or rely on statistical anomaly detection, which may yield high false-positive rates. To overcome these shortcomings, the proposed framework integrates feature selection through ReliefF with hybrid machine learning and deep learning classifiers, thereby optimizing both precision and scalability.

Key Achievements and Contributions

One of the foremost contributions of this study lies in the intelligent feature selection methodology. The ReliefF algorithm was used to systematically identify and prioritize the most significant attributes from the ToN-IoT Windows 10 dataset, a benchmark dataset that realistically represents IoT-based network traffic and attack patterns. The ReliefF method surpasses traditional correlation-based feature selection techniques by efficiently handling multi-class, incomplete, and noisy datasets—a common characteristic of IoT environments. By focusing only on the top twenty most relevant features, the proposed model achieves a considerable reduction in computational cost while simultaneously improving learning efficiency and overall classification performance.

Another major contribution is the development of a hybrid classification architecture that combines the strengths of traditional machine learning algorithms (KNN, SVM) and deep learning models (Neural Networks, LSTM). The comparative evaluation of these models revealed that the Medium Neural Network (MNN), Weighted KNN, and Fine Gaussian SVM achieved the highest accuracy levels—98.39%, 98.22%, and 97.97%, respectively. These results significantly outperform previous IDS models that employed correlation-based features, demonstrating that the ReliefF approach enhances both precision and recall across multiple intrusion categories, including DDoS, DoS, Injection, Password, and MITM attacks.

Moreover, the proposed IDS demonstrates exceptional generalization capability across different attack scenarios. The Area Under the ROC Curve (AUC) values, close to 0.99 for most classifiers, indicate a robust ability to differentiate between normal and malicious traffic. This is especially valuable in IoT contexts, where false positives can overwhelm system resources or cause legitimate data to be erroneously flagged as suspicious. By achieving high AUC and F1-scores, the proposed framework successfully balances detection sensitivity and reliability, ensuring that attacks are detected promptly without compromising operational continuity.

Performance Efficiency and Scalability

A central focus of this research was to design an IDS suitable for resource-constrained IoT devices. The ReliefF-based model achieves this by minimizing redundant features, thereby reducing memory usage and computation time during both training and detection phases. When tested in a simulation environment configured with modest hardware (Intel® Core™ i5 processor and 16 GB RAM), the framework maintained high accuracy without requiring GPU acceleration or extensive computational resources. This characteristic positions the model as deployable at the edge, where lightweight processing is critical for real-time intrusion detection.

The hybrid nature of the proposed system ensures scalability across different IoT layers—from edge nodes (sensors and embedded controllers) to cloud servers. For instance, while lightweight ML classifiers like KNN can operate on edge gateways, deeper architectures such as Neural Networks can function within fog or cloud layers to perform more comprehensive analysis. This distributed architecture aligns with the emerging fog computing paradigm, where security functions are executed closer to data sources, reducing latency and bandwidth consumption. The framework thus provides a foundation for hierarchical IDS deployment, enhancing scalability and flexibility in complex IoT infrastructures.

Comparative Evaluation and Insights

A comparative analysis between the proposed ReliefF-based IDS and prior correlation-based systems (as implemented in Moustafa et al., 2020) reveals clear superiority across all performance metrics. The earlier correlation-based approach, while useful, suffered from feature redundancy and weak discrimination between multi-class attack categories. In contrast, ReliefF's weight-based selection mechanism assigns priority to features that consistently differentiate between normal and attack behaviors, leading to enhanced precision.

The comparative results further confirm that models employing ReliefF-selected features achieve up to a 4–10% increase in accuracy and a significant reduction in false alarm rates. This improvement validates the hypothesis that effective feature engineering directly correlates with improved IDS performance. Moreover, the robustness of ReliefF against noise and missing data ensures that the model remains stable even under incomplete telemetry—a frequent condition in real-world IoT systems.

Acknowledgement: Nil

Conflicts of interest

The authors have no conflicts of interest to declare

References

1. Meidan Y, Bohadana M, Shabtai A, et al. Detection of unauthorized IoT devices using machine learning techniques. IEEE Internet Things J. 2018;5(3):1843–1853.
2. Doshi R, Apthorpe N, Feamster N. Machine learning DDoS detection for consumer Internet of Things devices. IEEE Secur Priv Workshops. 2018;2018(1):29–35.
3. Khan MA, Karim M, Kim Y. A lightweight intrusion detection system for IoT-based smart environments. Int J Comput Sci Netw Secur. 2019;19(2):96–103.
4. Verma A, Ranga V. Machine learning based intrusion detection systems for IoT applications. Wireless Pers Commun. 2020;111(4):2287–2310.
5. Ferrag MA, Maglaras L, Moschoyiannis S, et al. Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. J Inf Secur Appl. 2020;50:102419.
6. Almiani M, AbuGhazleh A, Al-Rahayfeh A, et al. Deep recurrent neural network for IoT intrusion detection system. Simul Model Pract Theory. 2020;101:102031.
7. Anthi E, Williams L, Burnap P, et al. A supervised intrusion detection system for smart home IoT devices. IEEE Internet Things J. 2019;6(5):9042–9053.
8. Shafiq M, Tian Z, Bashir AK, et al. IoT malware detection using lightweight CNN-based model. IEEE Access. 2020;8:182933–182944.
9. Hodo E, Bellekens X, Hamilton A, et al. Threat analysis of IoT networks using artificial neural network intrusion detection system. Int J Commun Syst. 2017;30(17):e3279.
10. Kumar P, Gupta GP, Tripathi R. A distributed framework for detecting DDoS attacks in IoT networks using lightweight ML techniques. J Netw Comput Appl. 2021;180:103028.
11. Nguyen TT, Reddi VJ. Deep reinforcement learning for cyber security. IEEE Commun Surv Tutor. 2020;22(2):1006–1029.