Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Hybrid Intrusion Detection System for IoT and Cloud Environments

Dipo Dunsin

d.dunsin@uwtsd.ac.uk

University of Wales Trinity Saint David https://orcid.org/0009-0009-7376-0477

Raju Molla

University of Wales Trinity Saint David

Alireza Esfahani

University of West London

Research Article

Keywords: Intrusion Detection System, IoT Security, Cloud Security, Machine Learning, XGBoost, Random Forest, LSTM Autoen-coder, Anomaly Detection

Posted Date: June 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10144004/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Molla, Esfahania and Dunsin. *Cybersecurity* (2026) XX:XX
https://doi.org/xx.xxxx/xxxx

Cybersecurity

RESEARCH Open Access

Hybrid Intrusion Detection System for IoT and Cloud Environments

Raju Molla[1], Alireza Esfahania[2], Dipo Dunsin[1*]

Abstract

The rapid growth of Internet of Things (IoT) devices and cloud-based network infrastructures has greatly increased the complexity of detecting malicious network traffic in heterogeneous environments. Conventional supervised intrusion detection systems usually perform well on labelled attack data but are often limited in effectiveness on emerging or unseen attack patterns due to their dependence on predefined attack knowledge. The study explores whether the combination of supervised classification and anomaly-aware sequence modelling can improve the performance of intrusion detection in combined IoT and enterprise network environments. To solve this problem, a hybrid intrusion detection framework is proposed using feature-aligned CICIDS2017 and TON_IoT datasets. A common traffic-related attribute, the destination port, the flow duration, the packet statistics, byte statistics and a domain indicator, were used to build a common cross-domain feature space. The proposed framework combines Random Forest and XGBoost as supervised learning models and an LSTM Autoencoder-based anomaly detection component. The LSTM Autoencoder was trained on normal traffic sequences only and used reconstruction-error analysis for anomalous behaviour detection. The models were assessed on numerous performance metrics including accuracy, precision, recall, F1 score, false positive rate, false negative rate, ROC-AUC analysis, and precision-recall analysis. Experimental results show the proposed hybrid framework has achieved good overall intrusion detection performance, i.e., 99.45% accuracy, 99.75% precision, 99.64% recall and 99.70% F1 score under standard evaluation settings. The standalone LSTM autoencoder obtained high precision, but the performance of recall was relatively low, which indicated that anomaly detection alone was not sufficient to reliably classify the intrusion. We also conduct an unseen-attack experiment to further evaluate the generalisation capability, where we exclude some attack families in the training and evaluate the models on previously unseen attack traffic. The results showed that the hybrid fusion approach improved the sensitivity and recall of the attack detection but at the cost of additional false positives and decreased overall precision. The results show that the hybrid intrusion detection frameworks based on the combination of supervised learning and anomaly-aware sequence modelling can improve the detection robustness in heterogeneous cloud and IoT environments. The paper also discusses the trade-offs between recall improvement and false alarm control in anomaly-aware intrusion detection systems. This work presents a feature-aligned cross-domain intrusion detection framework and a realistic evaluation of supervised, anomaly-based, and hybrid detection strategies for modern cybersecurity settings.

Keywords Intrusion Detection System, IoT Security, Cloud Security, Machine Learning, XGBoost, Random Forest, LSTM Autoencoder, Anomaly Detection

Raju Molla, Alireza Esfahania and Dipo Dunsin have contributed to this work.

*Correspondence:
Dipo Dunsin
d.dunsin@uwtsd.ac.uk

[1] Department of Applied Computing IICL, University of Wales Trinity Saint David, E14 4HA, UK
[2] School of Computing and Engineering, University of West London, London, UK

© The Author(s) 2026. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as appropriate credit is given to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit

http://creativecommons.org/licenses/by/4.0/.

Introduction

The rapid development in Internet of Things (IoT) and cloud computing technologies has radically transformed the modern communication infrastructures and digital services (Katta et al., 2025). IoT devices are progressively being deployed in smart homes, healthcare systems, industrial automation, transportation systems and smart city environments, where they continuously generate vast quantities of heterogeneous data. At the same time, cloud computing systems offer scalable storage, processing resources and network services to process and manage data dense situations. The integration of IoT and cloud technologies gives improved connectivity, automation, and operational efficiency, but also poses serious cybersecurity challenges because of distributed architectures, heterogeneous communication protocols, resource-constrained devices, and limited built-in security mechanisms (Moustafa, 2020; Buczak and Guven, 2016; Hindy et al., 2018; Sedjelmaci et al., 2017).

As current network infrastructures become more interconnected, cyberattacks against IoT and cloud environments are increasingly common and complex. Malicious actors commonly target vulnerable devices, poor authentication mechanisms and network communication protocols to carry out attacks such as Distributed Denial of Service (DDoS), infiltration, reconnaissance, botnet propagation and unauthorized access. Such threats jeopardize the confidentiality of the data, the integrity of the system and the availability of the service (Kulkarni et al., 2026). A high level of cybersecurity protection is demanded by modern network environments (Buczak and Guven, 2016; Hindy et al., 2018).

The intrusion detection system or IDS is required to protect the network infrastructure against any suspicious or malicious activities by analyzing the network traffic (Bijone, 2016). The conventional IDS methods can be divided into signature based detection system and anomaly based detection system (Khraisat et al., 2019). The network traffic is matched with the known attack signatures to identify the attack via the signature based IDS approaches. These systems are quite effective at detecting established attack patterns, but generally cannot identify new or emerging threats, since new attack signatures must be manually updated (Axelsson, 2000). Alternatively, anomaly-based IDS methods detect deviations from the regular traffic behavior and can identify suspicious activity, not only based on predetermined signatures. However, anomaly-based systems usually have higher false positive rates and inconsistent detection performance in complex network environments (Chandola et al., 2009; Buczak and Guven, 2016).

The field of cybersecurity research has seen widespread application of machine learning (ML) techniques to improve intrusion detection capacity. Supervised learning methods such as Random Forest and XGBoost have shown good classification performance on structured data of network traffic. In the proposed work, Random Forest is chosen because of its robustness to overfitting and efficiency in high dimensional traffic features (Breiman, 2001). Likewise, XGBoost provides scalable gradient boosting optimization and high predictive performance on large scale intrusion detection datasets (Chen and Guestrin, 2016). These models can achieve great accuracy on benchmark datasets such as CICIDS2017 which contains actual corporate style network traffic and different attack scenarios (Sharafaldin et al., 2018). However, supervised learning methods mainly rely on labeled attack data and may experience diminished generalization capabilities when traffic patterns differ substantially in heterogeneous contexts.

The ability of deep learning techniques to simulate complex nonlinear and sequential traffic patterns has been extensively studied for intrusion detection studies Ahmad et al. (2021). In particular, Long Short-Term Memory (LSTM) networks are appropriate to studying the temporal dependencies in sequential network traffic data (Hochreiter and Schmidhuber, 1997; Kim et al., 2016; Yin et al., 2017). Traditional machine learning approaches usually have difficulties in describing sequential behaviors and latent traffic patterns that can be captured by LSTM based models. Besides, the autoencoder-based anomaly detection models learn the low-dimensional representations of typical traffic patterns, and identify the anomalous samples with respect to the analysis of the reconstruction error (Hinton and Salakhutdinov, 2006; Shone et al., 2018a; Mirsky et al., 2018). These algorithms are anomaly-aware and can assist detect abnormal traffic patterns without the requirement for explicit attack signatures.

Although machine learning and deep learning approaches have made progress, there are still several limits in the current research of IDS. Models of supervised learning can provide excellent precision and classification accuracy for known attack classes . But they need labeled training data. However, anomaly-based models detect anomalies from typical traffic patterns, but they are sensitive to thresholds, have low recall values and high false positive rates (Shone et al., 2018a; Mirsky et al., 2018). Moreover, most of the IDS research explore a single dataset for the models investigation, which makes them less universal in different IoT and enterprise environments (Ring et al., 2019). The modern infrastructures are likely to incorporate the cloud systems, enterprise networks, and IoT ecosystems. Therefore, the intrusion detection frameworks must be evaluated in many traffic domains for robustness and practical application.

Recently, hybrid framework of Intrusion Detection Systems (IDS) has been explored by combining supervised learning and anomaly detection techniques to take the advantage of their complementary benefits (Chandrashekhar et al., 2026; Kim et al., 2014). Chandrashekhar et al. (Chandrashekhar et al., 2026) proposed a new threat detection framework for cloud-IoT environment that integrates autoencoder-based anomaly detection with deep learning classification models. The detection performance of the proposed system was good, but the system heavily depended on one dataset and an intensive computation architecture, which might limit the scalability and applicability across domains.

Another major issue in IDS research is dataset heterogeneity. The CICIDS2017 dataset delivers realistic enterprise-style network traffic with the diverse attack categories (Sharafaldin et al., 2018), and the TON_IoT dataset captures realistic IoT and IIoT network traffic behaviors (Moustafa, 2020). These datasets are very distinct in terms of feature structure, traffic characteristics and attack behaviors. Therefore, it is vital to build effective feature alignment and cross-domain dataset integration methodologies to develop a unified intrusion detection framework that can work in heterogeneous environments (Ring et al., 2019).

To address these problems, this research provides a hybrid intrusion detection system, a unified cross-domain environment that combines supervised machine learning and sequence-based anomaly detection. The proposed system integrates Random Forest and XGBoost for supervised attack categorization with an LSTM Autoencoder for modelling anomaly-aware sequences. Also, we merge the feature-aligned CICIDS2017 and TON_IoT datasets into a single dataset containing both business and IoT traffic scenarios.

The experimental results demonstrate that the proposed hybrid framework delivers good performance of intrusion detection on heterogeneous cloud and IoT traffic datasets. And the unseen-attack evaluation reveals that fusion of supervised classification and anomaly-aware sequence modeling can improve generalization ability under previously unseen attack situations, although this has trade-offs between detection sensitivity and false positive rates.

Research Questions

This paper investigates the potential of a hybrid intrusion detection framework that combines supervised machine learning with anomaly aware sequence modelling to improve the performance of intrusion detection in heterogeneous cloud and IoT settings.

This study also addresses the following research questions:

RQ1: Does a hybrid XGBoost–LSTM Autoencoder model improve intrusion detection performance?

RQ2: Can combining CICIDS2017 and TON_IoT datasets improve model generalization across different network environments?

RQ3: Does anomaly-aware sequence modelling help detect previously unseen attacks?

Hypothesis Statement

The paper hypothesises that the use of supervised machine learning and sequence-based anomaly detection in a unified hybrid intrusion detection framework can improve the robustness, generalisation ability and detection effectiveness of the intrusion detection system in heterogeneous cloud and IoT environments.

More specifically, this paper offers:

> A hybrid intrusion detection framework combining XGBoost-based supervised classification with LSTM Autoencoder-based anomaly-aware sequence modelling will achieve superior overall intrusion detection performance across feature-aligned CICIDS2017 and TON_IoT datasets compared to standalone anomaly detection approaches.

The proposed system is expected to exploit the high classification power of supervised learning for detecting known attack patterns while boosting the sensitivity to anomalous and previously unreported attacks using sequential reconstruction-error analysis. In addition, the integration of feature-aligned multi-domain datasets is anticipated to improve the model's generalisation capabilities across different network contexts, such as business, cloud, and IoT traffic cases.

Research Contributions

This paper contributes to the field of intrusion detection by offering a hybrid intrusion detection framework that combines supervised machine learning and anomaly-aware sequence modelling for heterogeneous cloud and IoT settings. The suggested framework combines the XGBoost based supervised classification and LSTM Autoencoder based anomaly detection method to enhance the detection ability of known and unseen cyber threats.

The other most important contribution of this research is the development of a unified cross-domain intrusion detection dataset through the integration of the CICIDS2017 and TON_IoT datasets. Since these datasets originate from different network domains and contain inconsistent feature structures, a feature-alignment strategy was designed to construct a compact and interpretable shared feature space suitable for both supervised classification and sequence-based anomaly detection.

This also provide a comprehensive empirical evaluation of several intrusion detection approaches, including Random Forest, XGBoost, LSTM Autoencoder, and the proposed hybrid fusion model. Several performance metrics such as accuracy, precision, recall, F1-score, false positive rate, false negative rate, ROC-AUC analysis, and unseen-attack evaluation were used to compare the performance. The assessment provides a broader picture on the advantages and disadvantages of supervised, anomaly-based and hybrid intrusion detection techniques.

Moreover, the study emphasizes the real-world limitations of applying isolated anomaly detection models in sequence-based intrusion detection contexts. Experimental results show that the LSTM Autoencoder has good anomaly-awareness and high precision, but its detection ability is limited by low recall and reconstruction-threshold sensitivity. The results indicate the significance of combining the anomaly detection and supervised learning techniques for balanced intrusion detection performance.

Finally, this work contributes to the ongoing discussion of cross-domain intrusion detection by studying the generalization ability of intrusion detection systems in heterogeneous enterprise and IoT traffic environments. The proposed framework combines multi-domain datasets and tests previously unseen attack scenarios providing a more realistic evaluation of the intrusion detection performance in modern cyber security environments.

Literature Review

The growing complexity, volume and diversity of cyber attacks on enterprise, cloud and Internet of Things (IoT) environments have made Intrusion Detection Systems (IDS) an essential part of cybersecurity architectures (Alakhras et al., 2025). The rapid integration of IoT devices with cloud computing platforms resulted in highly heterogeneous and distributed network architectures, which greatly augment the attack surface of modern systems (Hindy et al.,

2018; Sedjelmaci et al., 2017). Therefore, traditional security mechanisms are often unable to detect complex and evolving attack behaviors in dynamic network environment.

Therefore, the existing research in IDS has been increasingly focused on machine learning and deep learning techniques that can automatically identify malicious traffic patterns from the network data (Hamdi et al., 2026). Despite the significant progress made, there are still some existing challenges that limit the existing intrusion detection approaches, such as dependence on labelled data, inability to generalise over heterogeneous datasets, computational complexity and reduced effectiveness against previously unseen attacks (Ring et al., 2019; Buczak and Guven, 2016). This section critically reviews the evolution of intrusion detection approaches, evaluates recent machine learning and deep learning-based IDS frameworks, discusses limitations in the existing research, and highlights the research gaps that motivate the proposed study.

Traditional Intrusion Detection Approaches

The early intrusion detection systems mostly adopted signature-based detection techniques, which match network traffic against predefined attack signatures or known malicious patterns (Axelsson, 2000). Signature-based IDS approaches are very efficient in detecting already known attacks and usually produce low false positive rates, because detection decisions are based on explicitly specified attack characteristics (Ahmed et al., 2025). As a result, these systems have been widely adopted in enterprise cybersecurity environments.

A major disadvantage of signature-based IDS is that they are unable to detect new or evolving attacks (Iyer, 2021). Since these systems are based entirely on known attack signatures, new threats can only be identified after the corresponding signatures have been manually generated and updated (Buczak and Guven, 2016). Such dependency seriously limits adaptability and long-term effectiveness in fast changing IoT and cloud environments where attack behaviours are constantly changing.

In order to overcome these limitations the anomaly based intrusion detection methods were introduced. Anomaly detection approaches create a baseline representation of normal network behaviour and any deviation from this baseline is classified as a potential intrusion (Chandola et al., 2009). This enables the detection of zero-day attacks and new malicious activities not defined by signature.

All in all, anomaly based IDS approaches have their own challenges. In heterogeneous and highly dynamic environments, such as IoT ecosystems, communication patterns are very diverse between devices and services (Sedjelmaci et al., 2017), thus modelling normal behaviour is still a challenge. This often results in high false positive rates in anomaly-based systems, affecting operational reliability and increasing alert fatigue for security analysts. Anomaly detection systems can detect never-before-seen attacks, but they have a hard time to keep their detection performance stable in different traffic environments.

Machine Learning-Based Intrusion Detection

Traditional Intrusion Detection System (IDS) approaches suffer from constraints, thus, we exploited machine learning (ML) techniques that can understand complex traffic patterns from data without human intervention (Buczak and Guven, 2016). ML-based intrusion detection systems improve the flexibility by discovering statistical correlations between traffic features and attack behaviours without the requirement for manually crafted signatures.

Random Forest is one of the most popular supervised machine learning algorithms for intrusion detection, because it is robust, scalable and resistant to overfitting (Breiman, 2001). Random forest is an ensemble learning technique that uses many decision trees to make the final prediction. According to many studies, Random Forest has high classification accuracy and low false positive rates for benchmark intrusion detection datasets (Buczak and Guven, 2016). It is capable to manage high-dimensional traffic features and noisy datasets, which makes it especially applicable to network security.

XGBOOST also draws a lot of attention due to its gradient boosting optimization framework and high predictive performance (Chen and Guestrin, 2016). XGBoost builds decision trees one after the other . Each tree tries to minimise the errors of previous trees . This is done by iterative optimisation . XGBoost outperforms traditional ensemble methods in terms of better scalability, efficient parallel computation and better classification performance on large-scale datasets. Thus, many of the recent IDS studies have used XGBoost as a primary classification model for known attack pattern detection.

Supervised machine learning models are successful, but they still heavily rely on labelled training data. In real-world cybersecurity settings, it is expensive and time-consuming to build large-scale accurately labelled datasets of attacks. Moreover, labelled datasets rarely cover the full range of attack behaviours in the real-world. Thus, supervised IDS models often suffer from limited generalisation ability when deployed in environments that include novel or evolving attacks (Buczak and Guven, 2016). While Random Forest and XGBoost models usually attain very high benchmark accuracy, those results might not be representative of the real-world deployment conditions where traffic distributions are very different from training datasets.

Another important limitation is that many ML-based IDS frameworks are evaluated on a single benchmark dataset with static train-test splits. Such an evaluation strategy is susceptible to dataset dependency, where models end up learning dataset-specific traits rather than generalizable intrusion patterns (Ring et al., 2019). Thus, good performance on benchmark datasets does not necessarily translate into good performance in heterogeneous IoT and enterprise environments.

Deep Learning-Based Intrusion Detection

Deep learning (DL) techniques have been increasingly used in intrusion detection research to overcome the limitation of feature engineering in traditional machine learning approaches (Goodfellow et al., 2016). Deep learning models can automatically learn hierarchical feature representations

and complex non-linear relationships from the network traffic data.

One of the most common deep learning designs for IDS is the Long Short-Term Memory (LSTM) network (Hochreiter and Schmidhuber, 1997). LSTM is a recurrent neural network architecture designed explicitly to model the temporal dependencies of sequential data. LSTM models are well suited for identifying multi-stage assaults and temporal irregularities in traffic, since the network data is sequential and time-dependent.

Kim et al. (Kim et al., 2016) show that LSTM-based intrusion detection models can well simulate the temporal traffic characteristics and improve the detection accuracy. Yin et al. (Yin et al., 2017) also showed good performance in intrusion detection with recurrent neural network designs trained on benchmark network datasets. However, the studies have mostly evaluated their models on data sets from a single domain and have not explored the generalization across different IoT and corporate systems. Hence the robustness of these techniques on cross domain deployment settings are not well explored.

Autoencoder based methods have also gained immense popularity for anomaly detection tasks (Shone et al., 2018a), apart from sequence modelling. Autoencoders are trained to learn the compact representation of normal traffic behaviour using a reconstruction model of encoding-decoding. Typically, anomalous samples result in higher reconstruction errors when trained on benign traffic, which allows the detection of suspicious behaviour without the need for labelled attack data (Chandola et al., 2009).

Shone et al. (Shone et al., 2018a) proved the effectiveness of deep Autoencoder architectures for automatic feature extraction and anomaly detection for intrusion detection systems. Similarly, kitsune framework using ensemble of lightweight autoencoders for online network intrusion detection in resource-constrained scenarios was proposed by Mirsky et al. (Mirsky et al., 2018). Kitsune showed good anomaly detection capacity and suitability for IoT based systems.

However, current Autoencoder-based IDS approaches have a few significant limitations. The choice of the reconstruction-error threshold significantly affects the detection performance, and minor changes in the threshold can result in unstable false positive and false negative rates (Shone et al., 2018a). Furthermore, stand-alone anomaly detection models often have a lower recall performance when the attack patterns are similar to the normal traffic behaviour. While anomaly detection enhances sensitivity to unknown attacks, it can also increase false alarm rates, which in turn decreases operational reliability in practical deployment scenarios.

Deep learning models also suffer from computational issues. Architectures using recurrent networks, attention mechanisms or multi-stage ensemble pipelines frequently require huge amount of computing and big scale training data (Potluri and Diedrich, 2016; Aldweesh et al., 2020). However, the complexity may limit the feasibility of deployment in real-time IoT settings where computational resources and latency are constrained.

Hybrid Intrusion Detection Systems

Due to the limitations of individual methods, recent intrusion detection studies focus more on the hybrid IDS frameworks that combine supervised learning and anomaly detection techniques (Kim et al., 2014). The hybrid architectures are designed to take advantage of the strong classification ability of supervised models, while incorporating anomaly detection to identify previously unseen attacks.

Supervised classifiers such as Random Forest and XGBoost are often able to identify known attack categories with high precision and low false positive rates (Breiman, 2001; Chen and Guestrin, 2016). Instead, the anomaly detection methods such as autoencoders and LSTM networks enhance sensitivity to unknown traffic behaviour by detecting anomalies from learned normal patterns (Shone et al., 2018a; Mirsky et al., 2018). These complementary capabilities are combined by hybrid IDS frameworks to balance detection accuracy and generalisation capability.

Chandrashekhar et al. (Chandrashekhar et al., 2026) proposed a hybrid deep learning framework based on Autoencoder anomaly detection with CNN, BiLSTM, and attention for cloud–IoT enterprise environments. Their framework had high performance in intrusion detection and was more adaptive to the dynamic cyber threats. But the proposed architecture is very computationally complex, due to its multi-stage deep learning pipeline, even with high benchmark accuracy. Besides, the study mainly used one dataset, which made it difficult to evaluate the generalisation ability in different domains.

Similarly, there are a number of recent studies that explore the ensemble based IDS architecture, combining multiple classifiers to improve robustness and reduce misclassification rates (Hindy et al., 2018). Many of these approaches demonstrate strong detection performances but are often highly benchmark-dependent and computationally demanding. In practice, very complex hybrid frameworks can be difficult to deploy in resource constrained IoT systems where latency, memory consumption and computational overhead represent critical operational constraints.

A key trade-off in hybrid IDS research is the trade-off between recall and false positive control. Enhanced anomaly detection usually improves sensitivity to novel attacks but may also increase false alarms. Designing lightweight hybrid frameworks that can balance precision, recall, computational efficiency and anomaly-awareness is still an open research challenge.

Dataset Challenges in Intrusion Detection

The effectiveness of intrusion detection systems (IDSs) is highly reliant on the quality and variety of datasets used for training and evaluation (Hosseini et al., 2025). Most of the existing studies are based on a single dataset, which limits the generalization capability of the developed models. The CICIDS dataset is one of the most popular benchmark datasets used in intrusion detection research. It provides realistic network traffic data with a variety of attack scenarios, and is suitable for supervised learning approaches (Sharafaldin et al., 2018). But it is mainly an enterprise or cloud network environment. On the other hand, the TON IoT dataset characterizes the features of IoT en-

vironments, e.g., device level interactions and lightweight communication patterns (Moustafa, 2020). This dataset considers different kinds of attacks on IoT systems such as denial-of-service and data exfiltration attacks. The heterogeneity of datasets is a major challenge for multi-domain intrusion detection. Differences in the feature representation, distribution, and labeling schemes obstruct the integration process. Ring et al. (Ring et al., 2019) point out the importance of using different datasets to increase the robustness and reliability of IDS models.

Table 1 shows that the recent research on intrusion detection has been mainly focused on three major areas: realistic dataset development, deep learning-based anomaly detection and hybrid intrusion detection frameworks. Studies like CICIDS2017 and TON_IoT have significantly contributed to the progress of IDS evaluation by providing realistic modern attack scenarios and heterogeneous traffic environments. Each dataset is focused on different operational domains, CICIDS2017 mostly represents enterprise network traffic and TON_IoT represents IoT and IIoT environments. Current deep learning methods like Autoencoder and recurrent neural network models have demonstrated good ability in learning complex traffic patterns and detecting anomalous behaviour. Nevertheless, other studies also mention several limitations, including high computational complexity, sensitivity to the choice of a reconstruction-error threshold, lower recall performance in stand-alone anomaly detection models, and limited generalisation due to evaluation on a single dataset. Furthermore, lightweight anomaly detection frameworks like Kitsune show the effectiveness of Autoencoder-based online intrusion detection for novel attacks. Nevertheless, the performance of these techniques is dependent on the traffic characteristics and reconstruction behaviour. These limitations demand more robust and generalisable intrusion detection frameworks. These research gaps have inspired the proposed study that introduces a hybrid intrusion detection framework that integrates supervised learning using XGBoost and anomaly-aware sequence modelling using an LSTM Autoencoder. Moreover, this work combines the feature-aligned CICIDS2017 and TON_IoT datasets to allow cross-domain intrusion detection in both enterprise and IoT environments.

Research Gap Analysis

Despite significant advances in intrusion detection research, several important limitations remain unresolved in existing IDS frameworks.

First, many current intrusion detection approaches rely heavily on supervised machine learning models trained using labelled benchmark datasets (Breiman, 2001; Chen and Guestrin, 2016; Buczak and Guven, 2016). Although these models frequently achieve high classification accuracy for known attack categories, their effectiveness remains strongly dependent on the availability of representative labelled attack data. In practical cybersecurity environments, attack behaviours continuously evolve, reducing the ability of supervised models to generalise effectively to previously unseen threats.

Second, anomaly-based deep learning techniques improve anomaly-awareness by modelling normal traffic behaviour (Shone et al., 2018a; Mirsky et al., 2018). However, stand-alone anomaly detection algorithms are susceptible to high false positive rates, sensitivity to the choice of the reconstruction-error threshold and poor recall performance. Recent research also suggests that models trained only on anomalies may struggle to differentiate subtle malicious behaviors from regular traffic in high dynamic situations.

Another significant limitation is the dominance of single-dataset evaluation. Most of the IDS research assess their models on either CICIDS2017 or TON_IoT separately (Sharafaldin et al., 2018; Moustafa, 2020). However, the communication behaviors and feature distributions of business and IoT traffic are quite different. Therefore, models trained on a single dataset may learn dataset specific behaviour rather than generalizable incursion patterns.

Moreover, some of the modern hybrid IDS frameworks rely on computationally expensive deep learning architectures such as CNNs, BiLSTM layers, attention mechanisms or multi-stage ensemble pipelines (Chandrashekhar et al., 2026). These techniques in many cases perform well on benchmarks, but their computational demands might prevent their scalability and feasibility of implementation in resource-constrained IoT settings.

Therefore, there remains a need for a lightweight hybrid intrusion detection framework capable of integrating supervised classification with anomaly-aware sequence modelling while supporting cross-domain generalisation across heterogeneous enterprise and IoT traffic environments. This research addresses this gap by combining XGBoost-based supervised learning with an LSTM Autoencoder-based anomaly detection component using feature-aligned CICIDS2017 and TON_IoT datasets. In addition, the proposed framework evaluates unseen-attack detection capability to investigate the balance between precision, recall, anomaly-awareness, and cross-domain robustness.

Summary of Literature Review

As can be seen from literature, machine learning and deep learning algorithms are very successful in intrusion detection. Traditional IDS techniques are good at detecting known attacks but they are limited in detecting emerging threats. Supervised machine learning methods like Random Forest and XGBoost enhance the classification accuracy of known attack types. Deep learning methods like LSTM networks and Autoencoders provide better sequence modelling and anomaly detection capabilities.

However, current IDS research remains constrained by several unresolved limitations, including dependence on labelled datasets, reduced robustness against previously unseen attacks, computational complexity, and limited cross-domain generalisation capability. Furthermore, many existing studies rely on single-dataset evaluation and highly complex architectures that may not generalise effectively to heterogeneous real-world environments.

To address these limitations, this paper proposes a lightweight hybrid intrusion detection system that combines XGBoost based supervised classification with LSTM Autoencoder based anomaly-aware sequence modelling. The proposed framework utilizes the feature aligned CICIDS2017 and TON_IoT datasets to improve the cross-

Table 1: Comparative Analysis of Related Intrusion Detection Studies

Reference	Dataset / Environment	Approach Used	Key Contributions	Research Gap / Limitation	Connection to Proposed Study
Chandrashekhar et al. (Chandrashekhar et al., 2026)	CSE-CIC-IDS2018 for Cloud–IoT enterprise systems	Hybrid deep learning model combining Autoencoder, CNN, BiLSTM, and Attention mechanism	Achieved high intrusion detection accuracy through multi-stage anomaly detection and attack classification	High computational complexity and dependency on a single dataset reduce generalisation capability	The proposed work employs a lightweight hybrid IDS using XGBoost and LSTM Autoencoder with integrated multi-domain datasets for improved generalisation
Sharafaldin et al. (Sharafaldin et al., 2018)	CICIDS2017	Intrusion traffic generation and benchmark dataset development	Introduced realistic benign and malicious traffic with modern attack scenarios and extensive flow-based features	Focused mainly on enterprise and conventional network traffic without IoT telemetry behaviour	CICIDS2017 is utilised in this research as the cloud and enterprise network component of the integrated IDS dataset
Moustafa (Moustafa, 2020)	TON_IoT dataset for IoT and IIoT environments	IoT/IIoT telemetry and cybersecurity dataset generation	Captured realistic IoT telemetry, operating system logs, and network attack behaviours from heterogeneous IoT environments	Dataset structure differs from traditional network datasets, creating feature compatibility challenges	This research combines TON_IoT with CICIDS2017 to construct a cross-domain IDS framework supporting both IoT and enterprise environments
Shone et al. (Shone et al., 2018a)	Network intrusion detection systems	Deep Autoencoder-based anomaly detection model	Demonstrated the effectiveness of deep autoencoders in automatic feature extraction and anomaly detection	Reconstruction-based detection may produce lower recall and depends heavily on threshold selection	Similar findings are observed in this study, where standalone LSTM Autoencoder models exhibit limitations in recall performance
Mirsky et al. (Mirsky et al., 2018)	Online and real-time network intrusion detection	Kitsune lightweight ensemble Autoencoder framework	Proposed an efficient anomaly-based IDS capable of detecting previously unseen attacks in real time	Detection performance is sensitive to traffic distribution and reconstruction characteristics	The proposed framework incorporates an LSTM Autoencoder as an anomaly-awareness component integrated with XGBoost rather than using anomaly detection alone

domain intrusion detection ability without sacrificing the trade-off between the precision, recall, computational efficiency and anomaly-awareness.

Background and Problem Definition

The increasing reliance on digital infrastructure has made network security a critical concern for organizations, industries, and individuals. Modern networks are continuously evolving with the emergence of cloud computing and Internet of Things (IoT) devices, which increases the difficulty of protecting these environments significantly. Intrusion Detection Systems (IDS) are one of the major defense system against cyber threats. It is important for detecting malicious activity in network traffic. However, the use of classic IDS techniques in heterogeneous and dynamic network environments is affected by several problems.

Background of Intrusion Detection Systems

Intrusion Detection Systems are designed to monitor network traffic and system activities to detect unauthorized access, misuse, or anomalies (Depren et al., 2005). IDS can be deployed in various forms, including network-based IDS (NIDS), which analyze network traffic, and host-based IDS (HIDS), which monitor activities on individual devices (Efe and Abacı, 2022). Both processes are intended to detect potential security violations and to alert for further review.

As discussed in previous sections, IDS techniques are broadly categorized into signature-based and anomaly-based detection methods. Signature-based systems rely on predefined patterns of known attacks and are effective in identifying previously encountered threats with high accuracy (Li et al., 2019). However, these systems are inherently limited in their ability to detect new or unknown attack types, as they require continuous updates of signature databases (Axelsson, 2000).

Anomaly based systems however try to predict normal behavior and look for deviations as potential dangers. This allows them to find unknown or zero-day attacks that are not solved by signature-based methods. However, anomaly detection methods are often prone to higher false positive rates, because of the difficulty in accurately defining normal behavior in complex and dynamic network environments (Buczak and Guven, 2016).

Challenges in Modern Intrusion Detection

With the emergence of IoT and cloud computing, intrusion detection has become more challenging due to several key factors:

Data Heterogeneity Modern network environments consist of various data sources such as cloud servers, IoT devices, and Edge computing platforms. These sources generate data having different formats, structures and features. For example, cloud network traffic is typically high-volume and structured, while IoT traffic is lightweight, dynamic, and often erratic. This variability makes it challenging to build a single intrusion detection model that performs well in multiple domains (Ring et al., 2019).

Class Imbalance Intrusion detection datasets are often highly imbalanced because the number of normal traffic samples is significantly larger than that of malicious samples. Such imbalance could lead to biased model predictions as the model is going to tend to predict the majority class (regular traffic) and miss the minority class (attacks). This leads to accurate models, but poor detection

of real intrusions, in the sense of recall and false negative rates (Wang et al., 2017).

Detection of Unknown Attacks One of the most important problems in intrusion detection is the ability to detect zero-day attacks, attacks that have never been seen before. Supervised machine learning models are useful for detecting known attack patterns, and they use labeled data. However, they do not generalize well to new types of attacks that are not present in the training data. This limitation emphasizes the necessity for more adaptive and sophisticated detection mechanisms (Buczak and Guven, 2016).

Feature Inconsistency Across Datasets Many datasets may contain different feature representations, naming conventions and data formats. For instance, features like flow length, packet counts and byte statistics may have different representations in CICIDS and TON IoT datasets. This discrepancy makes it difficult to combine datasets directly for training a unified model requiring additional preprocessing and feature alignment techniques (Sharafaldin et al., 2018; Moustafa, 2020).

Computational Complexity Advanced intrusion detection models based on deep learning require significant computational resources for training and inference. For example, the LSTM networks have complex operations and sequential data processing, which are not suitable for real-time deployment in resource-limited scenarios such as IoT devices (Hochreiter and Schmidhuber, 1997).

Limitations of Existing Approaches

There have been many advances in the field of intrusion detection but the existing solutions suffer from a number of key drawbacks that prevent them from being useful in the modern heterogeneous network environment. Traditional supervised learning models, such as Random Forest and XGBoost, have shown excellent classification performance and detection accuracy on benchmark intrusion detection datasets. However, they depend heavily on labeled training data and learn mostly predefined assault patterns. Thus, their generalization to unseen or emerging cyber risks is limited (Breiman, 2001; Chen and Guestrin, 2016). This dependence on labeled data is a limiting factor for supervised intrusion detection systems in practical cybersecurity environments with constantly changing attack technique patterns, and limits their adaptability and long-term stability.

Anomaly detection approaches based on autoencoders and sequential deep learning models address some of these limitations by learning representations of normal traffic behavior and identifying deviations as potential intrusions (Chandola et al., 2009; Shone et al., 2018a). Although these methods improve anomaly-awareness and offer the potential to detect previously unseen attacks, their standalone detection performance is often constrained by reconstruction-threshold sensitivity, lower recall, and difficulty distinguishing subtle malicious behaviors from legitimate network traffic. Consequently, anomaly-based systems may generate unstable detection performance and increased false negative rates in complex traffic environments.

Deep learning based intrusion detection methods are computationally expensive. Architectures based on recurrent neural networks, sequence modelling or multi-stage deep learning pipelines often require high computing resources, large-scale training datasets and long training times (Goodfellow et al., 2016). Such criteria may present a challenge to their feasibility in real-time intrusion detection scenarios, particularly in resource-constrained IoT systems, where latency and processing efficiency are critical operational factors.

Another major weakness of the existent intrusion detection research is the frequent use of single-dataset evaluation approaches (Issa et al., 2024). Many research evaluate intrusion detection models on a single benchmark dataset, which may lead the model to learn dataset-specific properties instead of generalized intrusion patterns (Ring et al., 2019). The models trained on a single dataset typically exhibit poor cross-domain generalization ability when deployed in different real-world environments, because the traffic behaviors, feature distributions and communication patterns are substantially different in corporate and IoT contexts.

These constraints indicate the need for more adaptive and resilient intrusion detection frameworks, that can combine supervised classification, anomaly-aware sequence modeling, and multi-domain learning to boost detection ability, generalization performance, and practical applicability across modern cloud and IoT network infrastructures.

Problem Definition

The growth of cloud infrastructures and Internet of Things (IoT) environments pose major challenges for modern intrusion detection systems. However, existing intrusion detection techniques often fail to simultaneously provide high detection accuracy, anomaly-awareness, cross-domain generalization, and computing efficiency in diverse network environments (Qin et al., 2025). Supervised learning models work well for the type of attack that they are trained on Hossain et al. (2025) But they can find it difficult to spot new or unexpected risks. On the other hand, the performance of anomaly detection methods can be erratic, missing of true attacks and the results are sensitive to the threshold settings. Moreover, the heterogeneity of the traffic features in the enterprise and IoT environments makes it difficult to design a unified and generalizable intrusion detection system.

To this end, the present study addresses the problem of designing a hybrid intrusion detection system that can combine supervised attack categorization with anomaly-aware sequence modelling in the context of heterogeneous cloud and IoT network environments.

This research aims to investigate if the combination of XGBoost-based supervised classification and LSTM Autoencoder-based anomaly detection can improve the performance of intrusion detection and support cross-domain generalization using the feature-aligned CICIDS2017 and TON_IoT datasets, specifically. The proposed framework is designed to increase the detection of both known and unknown threats, and achieve high precision and minimize the number of false negatives.

Another key part of the problem is how to handle heterogeneous multi-domain datasets with inconsistent feature

structures and traffic characteristics. The study also aims to address the challenge of constructing a unified feature representation that can be used for both supervised classification and sequence-based anomaly detection due to the very different communication behaviors, feature distributions, and attack patterns of enterprise and IoT traffic.

The paper also discusses the effect of class imbalance and dependency of datasets on intrusion detection performance. The proposed methodology aims at improving the model robustness and generalization ability by learning from the integrated multi-domain traffic data instead of learning from a single benchmark dataset.

Finally, the research addresses the practical problem of trade-off between detection capability and processing efficiency. Deep learning methods provide a powerful sequential modelling capability but the computational complexity may limit its usability for real-time applications in resource constrained contexts. Therefore, the proposed hybrid framework aims to provide efficient intrusion detection performance with a relatively lightweight and scalable architecture suitable for modern cybersecurity applications.

Proposed Solution Overview

To overcome the limitations of existing intrusion detection techniques, this paper proposes a hybrid intrusion detection system based on supervised machine learning and anomaly-aware sequence modeling for heterogeneous cloud and IoT environments. The proposed framework aims to combine the high classification power of supervised learning models with the anomaly detection ability of deep learning-based sequence analysis to improve the overall performance of intrusion detection.

The supervised part of the framework contains Random Forest and XGBoost classifiers trained on labeled network traffic data (Breiman, 2001; Chen and Guestrin, 2016). The models are trained to learn the discriminative patterns of known attack categories and normal traffic behavior. These ensemble learning algorithms have demonstrated the powerful classification ability and robustness on the structured network traffic data, leading to very excellent detection performance for the attack patterns previously reported.

In addition to the supervised learning part (Hochreiter and Schmidhuber, 1997; Chandola et al., 2009), an LSTM Autoencoder based anomaly detection model is used for the detection of aberrant or unreported attack behaviors. The LSTM Autoencoder is trained only with normal traffic sequences to learn the temporal communication patterns and sequential behavior of the traffic . During the inference phase, the reconstruction-error analysis is used to detect the anomalous traffic sequences, where larger reconstruction errors are considered as potential anomalies or malicious activities.

Another important aspect of the proposed system is the fusion of heterogeneous network traffic datasets . In this work, we combine the CICIDS2017 (Sharafaldin et al., 2018) and TON_IoT (Moustafa, 2020) datasets to create a unified cross-domain intrusion detection environment that contains enterprise/cloud and IoT traffic scenarios. The different feature structure and traffic characteristics of the two datasets motivate us to design a feature-alignment technique to obtain a compact and consistent shared feature representation for both supervised classification and sequence-based anomaly detection.

The proposed framework also uses a hybrid fusion mechanism to fuse the output of the supervised classifier with the anomaly detection component by using a logical OR-based decision method. The purpose of this fusion strategy is to increase the sensitivity of attack detection, so that either the supervised classification model or the anomaly-aware sequence modelling component can detect the malicious traffic.

The proposed system incorporates supervised learning, anomaly detection, feature-aligned multi-domain datasets and hybrid decision fusion into a single architecture, with the aim of enhancing the robustness of intrusion detection, improving the generalization capability over heterogeneous environments, and providing a more realistic and scalable solution for modern cybersecurity applications.

Summary

In this section, the background and major issues of intrusion detection in IoT and cloud systems are discussed. It discussed the main shortcomings of current techniques, such as the incapability to identify unknown threats, the reliance on labeled data, the heterogeneity of data, and the restrictions of computing.

The problem statement emphasizes the need of a hybrid and multi-domain intrusion detection system to overcome these challenges. Our proposed strategy employs machine learning and deep learning techniques, and dataset alignment tactics to enhance detection performance and generalization capability.

Dataset Description

The efficiency of an intrusion detection system (IDS) (Hossain et al., 2025) is highly dependent on the quality, diversity and realism of the dataset used for model training and evaluation. Datasets are the basis of intrusion detection study, which learns network behavior patterns and detects malicious actions. Thus, the selection of appropriate datasets is an important step towards developing a reliable and generalizable IDS. In this paper we use two popular datasets, the CICIDS dataset and the TON IoT dataset. These datasets are chosen to represent different network domains i.e. cloud-based and IoT-based environments to facilitate the development of a multi-domain intrusion detection framework.

Recent work on intrusion detection has moved away from traditional datasets, such as the KDD Cup 99 dataset, due to many disadvantages such as duplicate records, unrealistic distribution of traffic, and outmoded attack scenarios (Tavallaee et al., 2009). With the change of network environment, there is a need for the dataset which can represent the real environment accurately. Recent datasets such as CICIDS and TON IoT alleviate these limitations with the inclusion of realistic traffic patterns, multiple attack types, and rich feature representations . Such datasets are used to increase the reliability of the findings of the experiments and to improve the generalization capability of the intrusion

detection models (Ring et al., 2019).

CICIDS Dataset

The CICIDS dataset introduced by Sharafaldin et al. (Sharafaldin et al., 2018) is one of the most popular benchmark datasets used in intrusion detection research. It is primarily intended to emulate real world network traffic in enterprise and cloud environments. The dataset contains a large number of network flow records, each representing a communication session between two endpoints. The flows have been extracted from real network scenarios, including typical user activities such as web browsing, email communication and file transfers.

One of the major advantages of the CICIDS dataset is that it is well endowed with a rich set of features. It has a wide set of features that describe the behavior of a network such as flow duration, packet counts of forward and backward flows, byte statistics and timing information. These attributes provide accurate information about the traffic patterns and are very useful for machine learning based research.

The data collection includes not only normal traffic but also other attack scenarios such as distributed denial-of-service attacks, brute force attacks, web-based attacks, penetration and botnet activity. This diversity enables models to learn multiple patterns related to malicious activity. Therefore, the CICIDS dataset is especially appropriate for supervised learning algorithms, which require labeled data to perform classification tasks.

The CICIDS dataset has some advantages, but it is mainly based on cloud or enterprise network context. However, this is not enough to describe the nature of IoT systems, whose traffic patterns are more dynamic and less structured. This constraint highlights the necessity of incorporating more datasets from IoT contexts.

TON IoT Dataset

Besides the cloud-based representation provided by CICIDS, this research uses the TON IoT dataset suggested from Moustafa (Moustafa, 2020). The dataset is meant to model the behavior of the IoT-based systems. It provides a realistic representation of the IoT network traffic.

The TON IoT dataset includes telemetry data generated by multiple IoT devices such as sensors, smart appliances and industrial control systems (Alsaedi et al., 2020). It records the interactions among devices and network services that typify the dynamic and heterogeneous nature of IoT environments. The IoT traffic is more lightweight, less structured and very variable than classic network statistics, which presents an even bigger challenge for intrusion detection models.

The dataset contains both numerical and categorical attributes describing different aspects of network communication and device behavior (Gad et al., 2021). like connection state, protocol type, packet counts, and byte level statistics. These features can be a useful information to detect anomalous activity in IoT systems.

The TON IoT dataset contains a variety of attack scenarios relevant to IoT systems, such as denial of service attacks, scanning, password attacks and data exfiltration. These attacks allow for a more detailed test of the intrusion models in the context of IoT.

In this study, we used the TON IoT dataset which improves the generalizability of the proposed methodology in other domains. The inclusion of data from both the cloud and the IoT offers a wider variety of patterns to the model, making it more robust and applicable to real-world scenarios.

Challenges in Multi-Domain Dataset Usage

The use of several datasets provides advantages in terms of diversity and realism, but also presents several obstacles to be dealt with in order to allow for good model training.

One of the main issues is the inconsistency of the feature representation among datasets . The CICIDS and TON IoT datasets are created with different techniques and tools, thus there are differences in feature names, formats and definitions. For example, different naming conventions or units may be used to represent the same network attributes, which makes direct integration more difficult.

Another problem is the distribution of the data. The datasets differ in traffic volume, feature scale and attack frequency. Such differences can lead to biases in the training of a model, which may favor patterns from one dataset over the other. This problem is more relevant for the fusion of cloud and IoT data as the traffic characteristics of them are somewhat different.

Another important challenge in the intrusion detection datasets is the class imbalance. The CICIDS and TON IoT datasets contain more normal traffic samples than malicious samples. This imbalance may affect the model performance and thus results in lower detection rates of attack instances. This problem should be resolved to get reliable intrusion detection.

Dataset Integration Approach

The present study tackles the challenges of merging multiple datasets by following a structured integration approach. The first step is to identify a set of similar features available in both datasets that are relevant to describe the network behavior. Such features typically include communication endpoint features, traffic volume and temporal features.

After identifying the key features, the datasets are processed to bring them into a common representation. This approach consists of the alignment of the names and formats of the features in such a way that equivalent attributes of both data sets are represented in a unified way. The data types are also standardized to match the model training.

Once the features are aligned, the data sets can be combined by assigning consistent labels. Both datasets are multi-class at the beginning however they are reduced to binary classification problems. In this form, normal traffic is tagged as a label and all attack types are clustered under a malicious label. This simplification helps the model to focus on the main task of distinguishing between benign and malicious traffic.

After preprocessing, the datasets are combined into one dataset. This joint dataset consists of samples from cloud and IoT contexts, which enables the model to learn from

Table 2: Dataset Statistics Throughout the Preprocessing Pipeline

Dataset Stage	Samples	Features
CICIDS2017 original dataset	2,830,743	79
TON_IoT original dataset	211,043	44
CICIDS2017 after feature selection	2,830,743	7
TON_IoT after feature selection	211,043	7
Merged dataset before duplicate removal	3,041,786	8
Dataset after duplicate removal	2,338,695	8
Final cleaned dataset	2,338,647	8
Training set	1,870,917	7
Testing set	467,730	7

heterogeneous traffic patterns. The combination of several datasets enhances the generalization capability of the intrusion detection system, thus making it more suitable for real world deployment.

Summary

In this section we present the two complementary datasets, CICIDS and TON IoT, as the foundation of the proposed intrusion detection methodology. The CICIDS dataset has a detailed description of the cloud-based network traffic. The TON IoT dataset describes the unique features of IoT environments. By mixing different datasets it is possible to build multi-domain IDS able to cope with diverse and heterogeneous network scenarios.

However, there are some challenges like inconsistency of features, heterogeneity of data and class imbalance when using multiple datasets. This is achieved by a systematic integration method. The present work provides a solid data set base for the next steps of pre-processing and modeling on feature alignment, label standardization and dataset merging into one structure.

Data Preprocessing

Data preprocessing is an important step in design of effective intrusion detection system (IDS). The quality and uniformity of the input data are critical factors in the performance of machine learning and deep learning models. Raw network traffic statistics can have errors such as missing values, duplicate records, and fluctuations in feature size, which can greatly affect the accuracy and generalizability of models. Thus, we need a systematic preparation pipeline that cleans the raw data and makes it consistent and ready for the model.

In this work, the CICIDS and TON IoT datasets were preprocessed before being merged into a combined dataset. The datasets are collected from diverse scenarios and possess heterogeneous features, thus, preprocessing is necessary to guarantee compatibility and improve data quality.

Table 2 summarises the dataset size at each stage of the preprocessing pipeline. The CICIDS2017 dataset initially contained 2,830,743 samples with 79 features, while the TON_IoT dataset contained 211,043 samples with 44 features. To construct a unified cross-domain dataset, six common traffic-related features and one label column were selected from each dataset. A domain indicator was then added before merging, resulting in an 8-column integrated dataset.

In total, the data collection consisted of 3,041,786 samples. 2.338,695 samples after filtering invalid negative duration values). During the cleaning process, duplicates were removed from the dataset. The input variables were separated from the label column before training. The machine learning and deep learning models had 7 input features.

Data Cleaning and Handling Missing Values

Network traffic datasets have often missing or undefined values due to insufficient packet captures or errors in feature extraction. Such inconsistencies, if left unaddressed, can corrupt statistical properties and damage model training.

Some features in the CICIDS dataset such as flow bytes per second and flow packets per second can result in undefined results in the calculation when divided by zero. To overcome this problem, zero is used to replace the missing and infinite values. This strategy preserves the numerical stability and the overall structure of the data collection. Similarly, the TON IoT dataset is checked for inconsistencies and any invalid entries are corrected to ensure data integrity.

Removal of Duplicate Records

Duplicate records are common in intrusion detection data sets, because of recurring network flows or artefacts of data collection. The presence of duplicate data may bias the model by over-representing some patterns, which may lead to overfitting and decrease the generalization capacity.

To deal with this, duplicate records are detected and removed from both datasets. This technique ensures that each instance in the dataset corresponds to a unique network event. This increases the diversity of the dataset, making it more representative of real-world traffic situations, thereby improving the robustness of the trained models.

Label Transformation and Standardization

The CICIDS and TON IoT datasets have different types of network traffic, different types of assaults and typical behaviour. But these categories are not consistent across datasets, each dataset has its own labelling scheme.

All labels are transformed to a binary form to remain compatible, which makes the classification easier. In this form normal traffic is assigned a value of 0 and all sort of malicious traffic is assigned a value of 1. With this modification, the model can focus on the difference of benign and harmful actions, which is the basic goal of intrusion detection systems.

Feature Scaling and Normalization

The numerical range of characteristics in a network dataset is often very diverse. Packet counts can be in the hundreds, but byte related features can be in the thousands or millions. Such discrepancies can be harmful for the model performance especially for gradient-based optimization based techniques.

To solve this problem, all values of features are normalized to a common range from 0 to 1 using Min-Max normalization. This transformation is defined by the following equation:

$$X_{\text{norm}} = \frac{X - X_{\min}}{X_{\max} - X_{\min}} \tag{1}$$

where X represents the original feature value, $X_{\min}$ and $X_{\max}$ denote the minimum and maximum values of the feature, and X_{norm} is the normalized value.

This normalization strategy assures that all features are contributing equally during training of the model and no feature with greater magnitudes dominates the learning process . Also, it improves the convergence speed and stability, particularly in deep learning models like LSTM networks.

Class Distribution Analysis

Before training the model, we need to check the class distribution in the data set. Most intrusion detection datasets are imbalanced, i.e., the number of normal traffic samples is usually larger than the number of attack samples. Such an imbalance can lead to biased prediction of the model, towards the majority class.

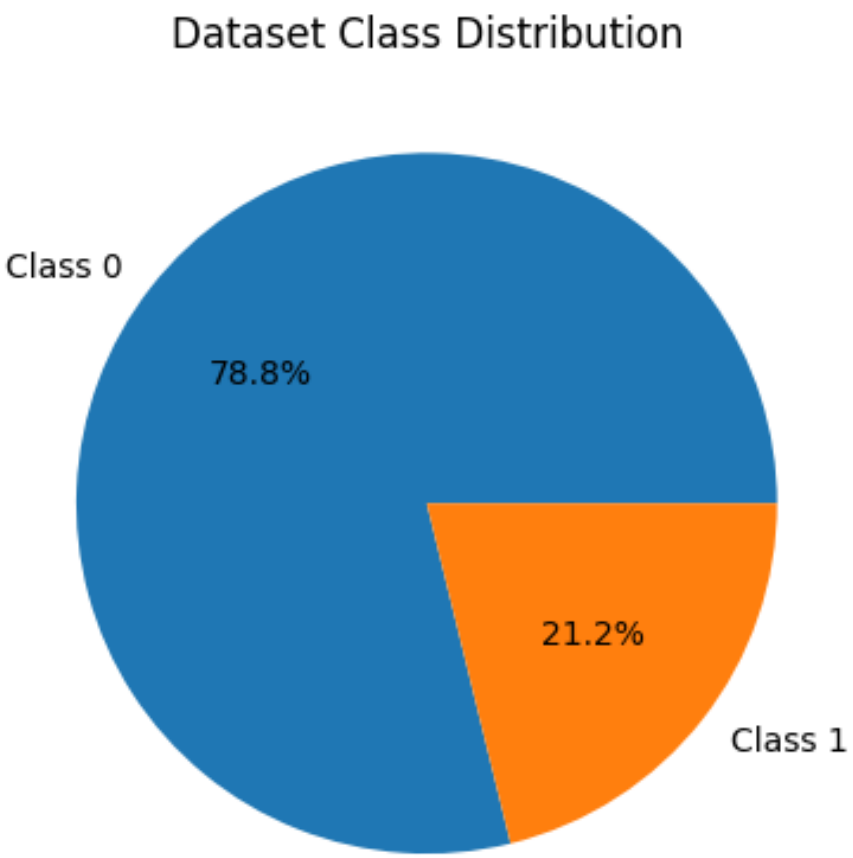


Figure 1: Class Distribution Before Preprocessing

As illustrated in Figure 1, the dataset is heavily skewed toward normal traffic. This imbalance highlights the importance of careful evaluation and dataset handling during model development.

Dataset Splitting

Then the data is pre-processed and split between testing and training data to test the model on unseen data. In this study, we employed 80:20 split where 80% of the data were used for training and 20% for testing.

In this technique, stratified sampling is used to ensure that both subsets have the same class distribution. This means that the test set is a true representation of the real world and that the evaluation of the model is not biased.

Preprocessing Pipeline Overview

The preprocessing approach can be succinctly described as a sequential transformation of raw data into a structured dataset appropriate for model training.

The pipeline follows this sequence:

$$\text{Raw Dataset} \rightarrow \text{Data Cleaning} \rightarrow \text{Missing Value Handling} \rightarrow \text{Duplicate Removal} \rightarrow \text{Label Transformation} \rightarrow \text{Normalization} \rightarrow \text{Train/Test Split}$$

This pipeline ensures that the dataset is consistent, noise-free, and ready for subsequent analysis and modeling.

Summary

In this section, a complete data preparation methodology is introduced to prepare the CICIDS and TON IoT datasets for intrusion detection. Pre-processing steps are executed, include dealing with missing values, removal of duplicates, label modification, feature normalization and separating the dataset to training and test sets.

These procedures are important to improve data quality and to support interoperability across diverse data sources. The obtained clean and consistent dataset boosts the ability of the proposed system to detect essential patterns and its accurate detection performance in the following stages of the job.

Exploratory Data Analysis (EDA)

Exploratory Data Analysis (EDA) is an important step before applying machine learning and deep learning models to understand the structure, properties and underlying trends of the dataset . EDA plays a crucial role in identifying the relation between features, anomaly detection and class imbalance analysis which in turn directly impacts the model performance for intrusion detection system. Both the CICIDS and TON_IoT datasets offer useful insights into traffic distribution, feature importance and attack behaviors.

Overview of Data Distribution

As a first step of Exploratory Data Analysis (EDA), we analyze the statistical distribution of crucial traffic-related aspects in the integrated datasets. Knowing how the features are distributed helps us detect skewness, outliers, and changes in traffic behavior between domains.

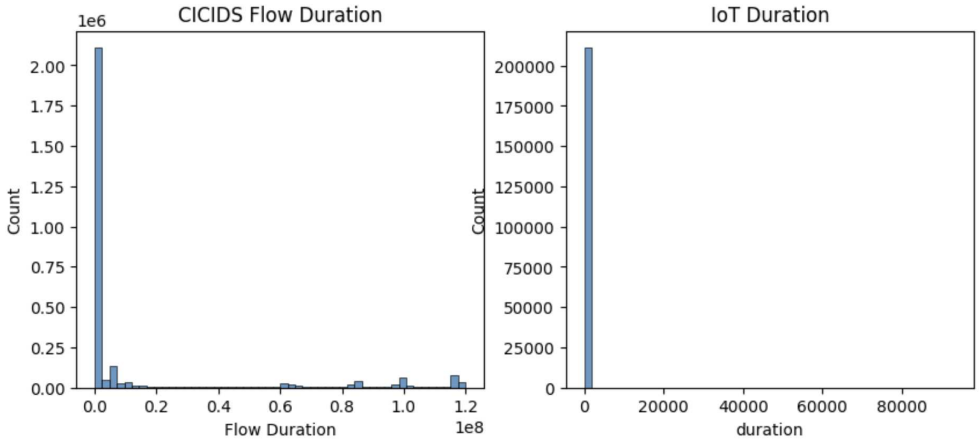


Figure 2: Distribution of Flow Duration Values in CICIDS2017 and TON_IoT Datasets Before Preprocessing

Figure 2 shows the distribution of flow duration values in the CICIDS2017 and TON_IoT datasets before any preprocessing. The distributions are strongly skewed with a large concentration of samples near lower duration values and some extreme values towards higher durations. This suggests that both datasets contain many outliers and irregular traffic behavior.

Furthermore, the picture shows the difference between the standard network traffic and the IoT traffic , emphasizing the heterogeneity of the combined dataset. If not handled, such skewed distributions can adversely affect machine learning performance. Hence, preprocessing steps such as logarithmic transformation and normalization are applied in the later stages to reduce the effect of extreme values and to enhance the stability of the model.

Feature Distribution Analysis

The distribution of distinct attributes can be useful to recognize patterns and find anomalies in network traffics . Flow length, packet counts and byte statistics of the communication behavior are valuable insights.

Attack traffic usually shows anomalous patterns, e.g., very high packet rates or inconsistent flow durations. We can learn the feature distributions of normal and malicious traffic, and find out which features distinguish them.

The analysis of the flow duration feature reveals that the flow duration of malicious traffic is generally shorter than that of normal traffic. Such activity suggests fast or bursty communication patterns which are often indicative of harmful operations such as scanning, flooding and denial of ser-

vice attacks.

In addition, the packet and byte related features showed significant differences among benign and malicious traffic samples, indicating the need of feature level analysis to build intrusion detection models. These facts support the choice of statistical approaches for pre-processing and normalization in order to boost the robustness and the generalization capability of the proposed models.

Correlation Analysis Between Features

Correlation analysis is used to see the links between features. Highly correlated features may lead to redundancy while weakly correlated features may provide supplementary information for classification.

The Pearson correlation coefficient is a measure of the linear relationship between two variables. It is defined as:

$$r = \frac{\sum(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum(X_i - \bar{X})^2}\sqrt{\sum(Y_i - \bar{Y})^2}} \quad (2)$$

where r represents the correlation coefficient, and X and Y denote the variables under consideration.

A correlation matrix is typically used to visualize relationships among features. Strong positive or negative correlations indicate dependencies that may influence model behavior. Eliminating redundant features based on correlation analysis can enhance model efficiency and reduce overfitting.

Detection of Outliers

Outliers are data points that are far from the normal distribution. The points are usually malicious or anomalous actions in the case of intrusion detection systems. Outlier detection is an important problem to enhance the performance of anomaly detection and robustness of machine learning models.

During the exploratory investigation, extreme values and anomalous traffic patterns were observed in various indicators such as flow length, packet counts and byte statistics . These anomalous observations are significantly different from most benign traffic samples and could be indicative of suspicious or malicious network activity.

The existence of outliers emphasizes the importance of using anomaly detection techniques to the proposed intrusion detection architecture. Autoencoders (Chandola et al., 2009; Shone et al., 2018a) and similar models are very effective in learning normal traffic and identifying anomalies based on reconstruction error. Moreover, some pre-processing methods were adopted, such as logarithmic transformation and normalization, to limit the influence of extreme values and improve the stability of the model during training process.

Feature Importance Insights

Feature importance analysis to find the most important qualities for intrusion detection. Tree-based models like Random Forest and XGBoost have a natural way to measure the relevance of features.

The most important features to distinguish between normal and malicious traffic are usually packet counts, byte statistics and flow duration (Breiman, 2001; Chen and Guestrin, 2016). These attributes capture important aspects of the communication activity and provide useful information for identifying suspicious network activities.

Trained model analysis showed that flow duration and bytes related characteristics significantly contributed to the classification performance. Packet based features also played an important role in distinguishing malicious traffic patterns from benign network behavior. The influential traits identified will enable dimensionality reduction, increase model interpretability and help to design more efficient intrusion detection systems.

Besides, the feature importance understanding helps to interpret the model decision making process and understand which network traffic features are the most relevant to the detection of cyber threats.

Summary

Exploratory Data Analysis allows us to understand the structure and properties of the dataset. The research points out the main issues including class imbalance, feature variability and existence of outliers and highlights the features that have considerable impact on intrusion detection.

These findings are a good basis for feature engineering and model building. The results of the EDA can inform the choice of appropriate preprocessing techniques and model selection approaches to improve the performance of the proposed intrusion detection system.

Feature Engineering

Feature engineering was performed to construct a unified and consistent representation of network traffic across the CICIDS2017 and TON_IoT datasets. Since the two datasets contain different feature structures, naming conventions, and traffic characteristics, direct merging was not possible. Therefore, this study adopted a manual feature-alignment strategy by selecting semantically similar numerical traffic features from both datasets.

The purpose of this stage was not to perform high-dimensional feature optimization, but to create a compact and interpretable cross-domain feature space suitable for both supervised classification and sequence-based anomaly detection.

Feature Alignment and Selection

A common feature space was constructed by selecting numerical traffic attributes that were available or semantically comparable in both CICIDS2017 and TON_IoT. The final selected input features were:

- *dst_port*: destination port number,
- *duration*: duration of the network flow or connection,
- *src_pkts*: number of source/forward packets,
- *dst_pkts*: number of destination/backward packets,

- *src_bytes*: number of source/forward bytes,
- *dst_bytes*: number of destination/backward bytes,
- *domain*: dataset origin indicator.

For CICIDS2017, the corresponding features were mapped from destination port, flow duration, total forward packets, total backward packets, total length of forward packets, and total length of backward packets. For TON_IoT, equivalent features were selected from destination port, duration, source packets, destination packets, source bytes, and destination bytes.

A binary *domain* feature was introduced to indicate the source of each sample, where 0 represents CICIDS2017 traffic and 1 represents TON_IoT traffic. This feature allows the model to retain information about dataset origin while learning from a unified feature representation.

Feature Transformation

Network traffic features such as duration, packet counts, and byte counts are often highly skewed and may contain extreme values. To reduce the influence of large outliers and improve numerical stability, logarithmic transformation was applied using `log1p()` to the following features:

- *duration*
- *src_pkts*
- *dst_pkts*
- *src_bytes*
- *dst_bytes*

The transformation is defined as:

$$X' = \log(1 + X) \tag{3}$$

where X represents the original feature value and X' represents the transformed value. This transformation preserves zero values while reducing the scale of large numerical values.

Feature Scaling

The selected characteristics were scaled into a similar range from 0 to 1 using Min-Max normalization after logarithmic transformation. This phase is especially critical for the LSTM Autoencoder, since neural network models are sensitive to changes in feature scale.

The Min-Max normalization is defined as:

$$X_{\text{norm}} = \frac{X - X_{\min}}{X_{\max} - X_{\min}} \tag{4}$$

where X represents the original feature value, $X_{\min}$ and $X_{\max}$ represent the minimum and maximum values of the feature, and X_{norm} represents the normalized value.

To avoid data leakage, the scaler was fitted only on the training set and then applied to the testing set.

Feature Reduction Strategy

This study did not use any dimensionality reduction methods such as Principal Component Analysis (PCA). Rather, feature reduction was done by manual feature alignment. We have kept only those features which are semantically consistent in both CICIDS2017 and TON_IoT.

This approach reduced the original CICIDS2017 feature space from 79 features and TON_IoT feature space from 44 features to a compact set of seven input features . This sacrificed the richness of the original datasets, but increased cross-domain compatibility, interpretability and computational efficiency.

Summary

Feature engineering was performed to obtain a compact, interpretable and cross-domain feature representation. The last set of features included six aligned numerical traffic features and one domain indication. Skewness was reduced through logarithmic transformation and the features were normalized using Min-Max scaling.

This feature engineering approach fulfills the main objective of this study, which is to evaluate supervised, anomaly-based, and hybrid intrusion detection models on feature aligned CICIDS2017 and TON_IoT traffic data.

Proposed Methodology

This section describes the proposed hybrid intrusion detection framework for improving detection performance in cloud and IoT environments. The methodology is a combination of supervised machine learning models and deep learning-based anomaly detection to overcome the shortcomings of the two approaches. The proposed system combines different methodologies and utilizes heterogeneous datasets in order to achieve high accuracy for the detection of known assaults and to be able to detect previously unreported threats.

Overview of the Proposed Framework

The proposed methodology is based on a hybrid architecture using Random Forest and XGBoost as supervised learning components and an autoencoder based on LSTM for anomaly detection. The solution is based on a unified dataset created by combining CICIDS and TON_IoT datasets to expose different traffic patterns.

The workflow starts with pre-processed data which is followed by feature engineering to obtain an enhanced feature set. This processed data is then used to train models, both supervised and unsupervised. The outputs of these models are combined to reach the final classification decision.

The motivation of the hybrid method is that the supervised models are good to detect the known attack patterns while the anomaly detection models are better to detect the unknown or zero-day assaults. The combination of both methodologies allows to find a trade-off between the generality and the accuracy of the system.

Supervised Learning Component

The supervised learning component consists of Random Forest and XGBoost classifiers trained on labeled data to classify network traffic as normal or malicious.

Random Forest constructs multiple decision trees and aggregates their predictions. The final prediction is given by:

$$\hat{y} = \frac{1}{N} \sum_{i=1}^{N} T_i(x) \tag{5}$$

where $T_i(x)$ represents the prediction of the i-th decision tree and N is the total number of trees.

XGBoost builds trees sequentially, where each new tree minimizes the error of the previous ones. The objective function is defined as:

$$\mathcal{L}(\phi) = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{k} \Omega(f_k) \tag{6}$$

where l is the loss function and Ω is a regularization term that controls model complexity.

These models are effective in detecting known attack patterns and achieve high classification accuracy when trained on labeled datasets (Breiman, 2001; Chen and Guestrin, 2016).

Anomaly Detection Component

To detect unknown attacks, an LSTM-based autoencoder is incorporated. The autoencoder is trained using normal traffic data to learn its underlying patterns. During testing, deviations from these patterns result in higher reconstruction errors, indicating anomalies.

The reconstruction error is computed as:

$$E = \|X - \hat{X}\|^2 \tag{7}$$

where X is the original input and $\hat{X}$ is the reconstructed output.

A threshold is defined such that if the reconstruction error exceeds this threshold, the sample is classified as anomalous. The use of LSTM networks enables the model to capture temporal dependencies in network traffic, making it effective for detecting sequential attack patterns (Hochreiter and Schmidhuber, 1997; Chandola et al., 2009; Shone et al., 2018a).

Hybrid Decision Mechanism

The final decision is obtained by combining the outputs of the supervised and anomaly detection components. This ensures effective detection of both known and unknown attacks.

The decision function is defined as:

$$D(x) = \begin{cases} 1, & \text{if } S(x) = 1 \text{ or } A(x) > \theta \\ 0, & \text{otherwise} \end{cases} \tag{8}$$

where $S(x)$ represents the supervised model prediction, $A(x)$ is the anomaly score, and θ is the anomaly threshold. [h] Hybrid Fusion Decision Mechanism [1] each traffic sequence x Obtain XGBoost prediction $S(x)$ Reconstruct sequence using LSTM Autoencoder Compute reconstruction error $A(x)$ $S(x) = 1$ **or** $A(x) > \theta$ Classify x as Attack Classify x as Normal

Model Training Process

The training is divided into different supervised and unsupervised phases. Supervised models are trained on labeled data, whereas the LSTM autoencoder is trained just on regular traffic data.

The hyperparameters such as number of trees in Random Forest, learning rate in XGBoost, number of LSTM layers are modified to obtain the optimum results. We employ cross-validation to ensure that the model generalizes to unseen data.

System Workflow Visualization

The overall workflow of the proposed framework is illustrated in Fig. 3.

The framework follows a sequential pipeline from data preprocessing and feature engineering to model training and decision fusion.

Performance Illustration

The effectiveness of the hybrid approach is demonstrated by comparing it with individual models.

As shown in Fig. 4, the hybrid model achieves higher accuracy compared to individual models, highlighting the advantage of combining multiple techniques.

Unseen-Attack Evaluation

An unseen-attack experiment was conducted to evaluate the generalization capability of the proposed intrusion detection framework. For this evaluation we excluded some attack families from the training and only considered them for the testing. The experimental setting simulates new attack scenarios never met before and enables a more realistic evaluation of anomaly awareness and generalization performance.

Many denial of service related attacks, DDoS, DoS Hulk, DoS GoldenEye, DoS Slowhttptest, DoS Slowloris and DoS traffic were included in Held out attack categories. In the training phase, the supervised models are trained with only normal traffic and the remaining types of attacks.

The evaluation is done by comparing the performance of XGBoost, LSTM Autoencoder and hybrid fusion model on unseen-attack condition using Accuracy, Precision, Recall and F1-score metrics.

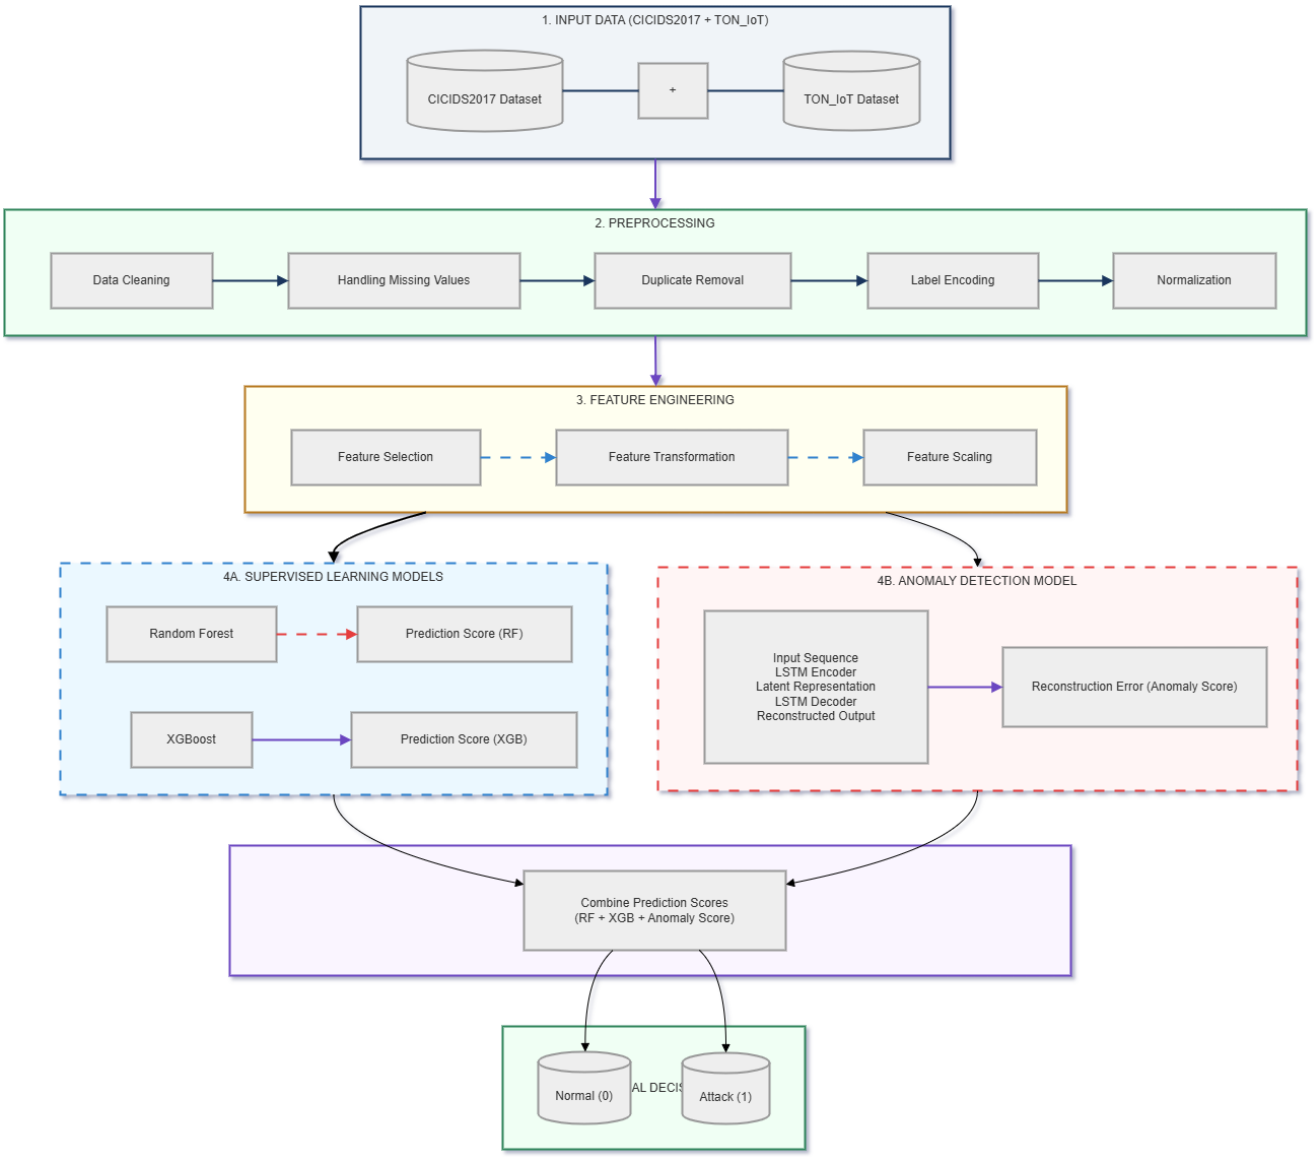


Figure 3: Proposed Hybrid Intrusion Detection Framework

Summary

In this part we present a hybrid intrusion detection system by combining supervised learning model with deep learning based anomaly detection. Known assault detection is performed with Random Forest and XGBoost and unknown threat detection is performed with an LSTM autoencoder.

The hybrid decision mechanism takes advantage of the complementing characteristics of various models and thus increases the detection power. Besides, the multi-domain datasets use allows better generalization for the system to be used in real-world cloud and IoT scenarios.

Model Design

The model design describes the architecture and working structure of the proposed hybrid intrusion detection system. It shows the interaction of the system components, the flow of data in the model and the process of generation of forecast. The proposed approach combines supervised learning models with a deep learning based anomaly detection method to build a robust and scalable intrusion detection system.

Architectural Overview

In this paper, we propose a hybrid architecture model which combines classification based detection and anomaly based detection. In the system, the network traffic data is processed in a sequential pipeline from the input data to the final decision, i.e. normal or malicious traffic.

The architecture consists of three main components: a supervision learning module, an anomaly detection module and a fusion mechanism. The supervised module includes Random Forest and XGBoost classifiers that are trained on labeled data to detect known attack patterns. We use an LSTM Autoencoder to learn the normal behavior of traffic and implement the anomaly detection module by detecting deviation as anomaly.

The outputs of these components are fused using a decision fusion approach to get the final classification result. The hybrid design allows the system to successfully detect known

Model	Accuracy	Precision	Recall	F1-score
Random Forest	0.9958	0.9902	0.9901	0.9901
XGBoost	0.996	0.9939	0.9871	0.9905
LSTM AE	0.2039	0.993	0.1235	0.2197
Hybrid	0.9945	0.9975	0.9964	0.997

Figure 4: Performance Comparison of Models

and unknown attacks.

Supervised Model Design

The supervised component contains two ensemble learning algorithms, Random Forest and XGBoost. We selected these models because they have demonstrated good performance on classification tasks and can handle high-dimensional network data (Breiman, 2001; Chen and Guestrin, 2016).

A random forest is a collection of decision trees . Each tree is trained on a random subset of data. The final prediction is made by majority voting of all trees. This promotes robustness and avoids overfitting.

XGBoost is a gradient boosting framework that constructs trees one by one. Each tree is trained to reduce the inaccuracy of previous trees and thus the overall accuracy is improved. The model also employs regularization techniques to control complexity and improve generalization.

Both models are trained on the merged CICIDS and TON_IoT dataset, to learn the cloud and IoT traffic patterns.

LSTM Autoencoder Design

The anomaly detection component of the proposed intrusion detection system is based on an LSTM Autoencoder designed to learn temporal patterns in sequential network traffic data. Long Short-Term Memory (LSTM) networks are suitable for modelling sequential dependencies and temporal behaviour in network traffic, which is important for identifying anomalous communication patterns (Hochreiter and Schmidhuber, 1997; Chandola et al., 2009; Shone et al., 2018a).

The model was trained exclusively on normal traffic sequences so that it could learn the characteristics of benign behaviour. During inference, sequences that produce high reconstruction error are considered anomalous and potentially malicious.

For the LSTM Autoencoder, tabular traffic records were transformed into fixed-length sequences with a sequence length of 10. Each sequence contained seven input features. A sequence was labelled as malicious if at least one sample within the sequence corresponded to attack traffic.

The encoder section contains 2 LSTM layers. The first encoder was an LSTM with 64 hidden units and `return_sequences=True` followed by a second LSTM with 32 hidden units which provided the latent bottleneck representation. The original sequence was decoded by the decoder using `RepeatVector` layer and two decoder LSTM layers with 32 and 64 hidden units respectively. A last `TimeDistributed Dense` layer reconstructed the feature values at each time step.

The model is trained with Mean Squared Error (MSE) reconstruction loss and Adam optimizer with learning rate of 0.001. I trained for 10 epochs, with a batch size of 64, and early terminated based on monitoring of validation loss.

For inference, anomaly scores were computed as the mean squared reconstruction error between the original and reconstructed sequences. The threshold for reconstruction-error 0.001014 was tuned empirically to distinguish between regular and aberrant traffic patterns.

Fusion Mechanism Design

The outputs of the supervised and the anomalous detection parts are combined using the fusion mechanism. This stage is important for the trade-off between detection accuracy and generalization capabilities .

Supervised models predict classification labels and thus classify the samples as normal or malicious. The LSTM Autoencoder outputs a reconstruction error, which is used as the anomaly score. These results are combined by a rule-based judgment mechanism.

If the trained models classify a sample as malicious AND/OR if the anomaly score is higher than a predetermined threshold, the sample is marked as an attack. Otherwise the sample is taken in the normal way. This approach enables efficient detection of known and unknown threats.

Data Flow Through the Model

The data has a structured process. First, the input data are pre-processed and transformed into normalized feature set. The processed data is then fed in parallel to the supervised models and to the LSTM Autoencoder.

The supervised models output classification probabilities as per the learned patterns and the autoencoder outputs reconstruction errors. Then the fusion layer aggregates these results for the final deduction.

The parallel structure of the processing improves the efficiency and guarantees the execution of various detection approaches in simultaneously, which raises the overall performance of the system.

Model Implementation Considerations

Several implementation considerations are addressed to optimize model performance. Choosing the right hyperparameters is crucial to achieve the best possible results. The parameters that are changed for Random Forest are number of trees and max depth. For XGBoost, the learning rate, number of estimators and regularization parameters are tuned carefully .

The architectural parameters of the LSTM Autoencoder (number of layers, hidden units and sequence length) are determined by the features of the dataset. The threshold for anomaly detection is determined empirically based on the distribution of the reconstruction error.

The model is implemented with standard machine learning and deep learning tools, ensuring repeatability, scalability and ease of deployment.

Summary

The model design presented in this section outlines a hybrid intrusion detection framework that integrates supervised learning and anomaly detection techniques. The combination of Random Forest, XGBoost, and LSTM Autoencoder enables effective detection of both known and unknown attacks.

The fusion mechanism plays a critical role in combining outputs from different components, resulting in improved detection performance. The structured data flow and carefully designed architecture make the proposed model suitable for deployment in complex network environments, including cloud and IoT systems.

Experimental Setup and Evaluation

Dataset Configuration

The experiments are performed on the combined dataset of CICIDS and TON_IoT datasets. The datasets are extracted from cloud based and IoT based network environments, respectively, and thus contain a wide range of traffic patterns and attack scenarios (Sharafaldin et al., 2018; Moustafa, 2020). Both datasets are pre-processed and feature engineered prior to training as described in previous sections.

The integrated dataset is divided into training and testing sub-groups with the ratio of 80:20. Stratified sampling is used to guarantee that each subset contains the same ratio between normal and malicious samples. This method decreases bias during training and evaluation.

Implementation Environment

The proposed model is implemented in Python which is a flexible and efficient platform for machine and deep learning applications. For data processing and model development, widely used libraries are used.

The supervised models, Random Forest and XGBoost are implemented in the standard machine learning libraries and the LSTM Autoencoder is built in a deep learning framework. The tests are conducted on a system with sufficient computation resources to support large-scale network data and training operations.

The existing frameworks allow for reproducibility and deployment in real-world applications.

Model Configuration

The configuration of each component in the proposed hybrid intrusion detection system was selected empirically based on detection performance and computational efficiency.

The Random Forest classifier was implemented using 100 decision trees with `random_state=42` and `n_jobs=-1` to enable parallel processing.

The XGBoost classifier was configured using:

- `n_estimators=100`
- `max_depth=6`
- `learning_rate=0.1`
- `subsample=0.8`
- `colsample_bytree=0.8`

To detect anomalies this paper used an LSTM Autoencoder with overlapping network traffic windows of length 10. The encoder was composed of LSTM layers with 64 and 32 hidden units, and the decoder was composed of LSTM layers with 32 and 64 hidden units with a `TimeDistributed Dense` output layer.

The autoencoder was trained with Adam optimizer with learning rate 0.001 and the reconstruction loss function was Mean Square Error (MSE). Training was conducted for 10 epochs with batch size of 64 and early stopping based on monitoring of validation loss.

The last hybrid model was a hybridization of XGBoost sequence level prediction and LSTM Autoencoder anomaly prediction with logical OR fusion.

Training Procedure

The training procedure is divided into two parts, corresponding to the supervised and unsupervised parts.

In the first phase, supervised models are trained on labeled data to classify the network traffic into normal and harmful classes. During training, validation techniques are used to ensure that the learning is correct and does not overfit.

In the second stage, only regular traffic samples are trained with the LSTM Autoencoder. The objective is to reconstruct the input data with minimal error. The reconstruction error is later used as an anomaly score.

The loss function used for training the autoencoder is the Mean Squared Error (MSE), defined as:

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(X_i - \hat{X}_i)^2 \tag{9}$$

where X_i represents the original input and $\hat{X}_i$ represents the reconstructed output.

Unseen Attack Evaluation Setup

The unseen-attack experiment was designed to evaluate whether the proposed framework could detect attack families that were not observed during supervised training. The attack labels were first standardized by converting them to lowercase and removing extra spaces. The excluded attack families were DDoS, DoS Hulk, DoS GoldenEye, DoS Slowhttptest, DoS Slowloris, and DoS traffic.

The training set was constructed using all normal samples and all attack samples except the held-out unseen attack families. Therefore, the supervised model was not exposed to these attacks during training. The test set was constructed using normal samples together with only the excluded unseen attack samples. This ensured that the evaluation simulated a zero-day or previously unseen attack scenario.

To reduce data leakage, Min-Max scaling was fitted only on the training data and then applied to the unseen test data. The domain feature was retained so that both CICIDS2017 and TON_IoT traffic sources could be represented during evaluation. This helped control domain-specific bias by preserving cross-domain traffic information during testing.

Evaluation Metrics

To evaluate the performance of the proposed intrusion detection system, several standard metrics are employed. These metrics provide a comprehensive assessment, particularly in imbalanced datasets.

Accuracy is defined as:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{10}$$

Precision is defined as:

$$\text{Precision} = \frac{TP}{TP + FP} \tag{11}$$

Recall is defined as:

$$\text{Recall} = \frac{TP}{TP + FN} \tag{12}$$

The F1-score is defined as:

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{13}$$

These metrics are widely used in intrusion detection research and provide a balanced evaluation of model performance (Buczak and Guven, 2016).

Experimental Results Overview

The performance of the proposed hybrid model is compared with individual models to demonstrate its effectiveness.

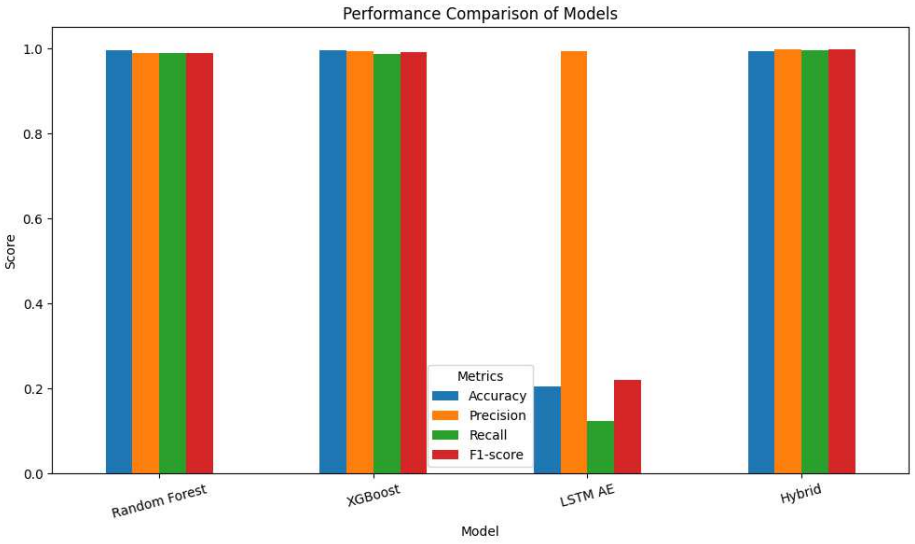


Figure 5: Performance Comparison of Models

As shown in Fig. 5, As shown in Fig. 6, the hybrid model achieved improved recall and F1-score performance compared to standalone anomaly detection while maintaining strong overall classification performance. The improvement is particularly evident in terms of F1-score, which reflects a better balance between precision and recall.

Summary

In this section, the experimental setup for the assessment of the proposed hybrid intrusion detection system is described, including dataset design, implementation environment, model parameters, training methodology, and evaluation metrics.

A mixed dataset displays different traffic pattern. The hybrid architecture uses supervised classification and anomaly-aware sequence modelling to improve robustness in heterogeneous network environments. The experimental results show the feasibility of the proposed approach and give a suitable foundation for further study.

Results

This section illustrates the experimental results of the proposed hybrid intrusion detection system. Performance of individual models, Random Forest, XGBoost and LSTM Autoencoder are considered and compared. Finally, the results of the hybrid model are revealed.

Random Forest Performance

The Random Forest model is evaluated as a baseline supervised classifier. The results show strong performance in detecting network intrusions.

The confusion matrix is:

$$\begin{bmatrix} 367817 & 971 \\ 980 & 97962 \end{bmatrix}$$

The model performs well, with only a small number of false positives and false negatives.

XGBoost Performance

The XGBoost model achieves slightly improved performance compared to Random Forest.

The confusion matrix is:

$$\begin{bmatrix} 368189 & 599 \\ 1273 & 97669 \end{bmatrix}$$

The results indicate improved precision but slightly reduced recall.

Sequence Construction

For the LSTM Autoencoder, tabular traffic records were transformed into fixed-length sequential windows with a sequence length of 10. The sequence construction process enabled the model to learn temporal traffic behaviour rather than isolated observations.

A sequence was labelled as malicious if at least one sample within the sequence corresponded to attack traffic. This security-oriented labelling strategy prioritizes the detection of suspicious temporal windows and reduces the risk of overlooking malicious activities embedded within predominantly normal traffic sequences.

However, this approach also alters the class distribution compared to the original row-level dataset by increasing the proportion of malicious sequences. Consequently, sequence-level evaluation metrics are not directly comparable to row-level classification metrics produced by Random Forest and XGBoost models.

LSTM Autoencoder Performance

The LSTM Autoencoder was evaluated as an anomaly detection model using reconstruction error. The model achieved very high precision but low recall, indicating that it was conservative in identifying anomalous traffic.

The confusion matrix is:

$$\begin{bmatrix} 42927 & 367 \\ 372008 & 52419 \end{bmatrix}$$

Although the LSTM Autoencoder produced very few false positives, it missed a large number of attack sequences. This indicates that the selected reconstruction-error threshold was highly conservative. Therefore, the LSTM Autoencoder should not be interpreted as a strong standalone classifier, but rather as a complementary anomaly-awareness component within the hybrid intrusion detection framework.

Hybrid Model Performance

The hybrid model combines the XGBoost classifier with the LSTM Autoencoder using an OR-based fusion mechanism. A sequence is classified as malicious if either the supervised model or the anomaly detection model identifies it as an attack.

The confusion matrix is:

$$\begin{bmatrix} 42236 & 1058 \\ 1514 & 422913 \end{bmatrix}$$

In general, the hybrid model performed very well with high precision, good recall and very low false negative rate. This indicates the fusion method improves the attack detection sensitivity with a low false alarm rate. In intrusion detection, the reduction of false negatives is very important, since false negatives are just missed attacks.

Comparative Analysis of Results

The comparison shows that supervised models perform well for known attacks, while the LSTM Autoencoder detects anomalies with high precision. The hybrid model effectively combines these strengths.

Unseen-Attack Evaluation

To further evaluate the generalization capability of the proposed intrusion detection framework, an unseen-attack experiment was conducted by excluding selected attack families during training and evaluating the models on previously unseen attack traffic. This experiment provides a more realistic assessment of anomaly-awareness and robustness under previously unseen attack conditions.

Table 10 presents the performance of XGBoost, the LSTM Autoencoder, and the hybrid fusion model under unseen-attack evaluation settings.

Table 10: Unseen-Attack Evaluation Results

Model	Accuracy	Precision	Recall	F1-score
XGBoost	99.19%	64.63%	96.37%	77.37%
LSTM Autoencoder	96.80%	29.92%	91.90%	45.14%
Hybrid Model	96.21%	27.16%	97.41%	42.47%

Table 11: Per-Attack-Type Detection Performance

Attack Type	Support	Detected	Missed	Detection Rate (%)
backdoor	87	82	5	94.25
bot	92	89	3	96.74
ddos	5886	5808	78	98.67
dos	609	602	7	98.85
dos goldeneye	438	433	5	98.86
dos hulk	7049	6943	106	98.50
dos slowhttptest	244	241	3	98.77
dos slowloris	194	193	1	99.48
ftp-patator	162	158	4	97.53
injection	837	825	12	98.57
mitm	47	46	1	97.87
password	615	607	8	98.70
portscan	3752	3707	45	98.80
ransomware	3	3	0	100.00
scanning	225	224	1	99.56
ssh-patator	131	127	4	96.95
web attack - brute force	65	64	1	98.46
web attack - sql injection	4	4	0	100.00
web attack - xss	28	28	0	100.00
xss	343	339	4	98.83

Per-Attack-Type Detection Performance

To provide a more detailed attack-wise evaluation, the detection capability of the XGBoost classifier was analysed using the original attack labels preserved from both the CICIDS2017 and TON_IoT datasets. Since the implemented intrusion detection system operates as a binary classifier, the analysis reports the detection rate for each attack category rather than multi-class precision or recall. The detection rate represents the percentage of attack samples correctly identified as malicious.

The detection performance of the proposed framework for each attack type is presented in Table 11. The results show consistent high detection performance for various intrusion scenarios. The detection rates for high-volume attacks such as DDoS, DoS Hulk, DoS GoldenEye, Slowloris and PortScan were over 98%. This demonstrates the effectiveness of the proposed framework in identifying malicious behavior that results in high traffic.

Similarly, the reconnaissance and scanning related attacks, PortScan and Scanning, also achieved good performance, which indicates that the selected network-flow features are effective to characterize the features of the malicious traffic at both the business and IoT environments. Other types of attack, such as Backdoor, Bot, FTP-Patator and SSH-Patator, which are more stealthy, had lower detection rates, but the detection capability was still higher than 94% in all cases.

Furthermore, several attack categories including ransomware and web attack SQL injection achieved perfect detection rates. However, these categories contained very limited support samples; therefore, the results should be interpreted cautiously. Overall, the findings demonstrate that the proposed hybrid intrusion detection framework maintains strong generalization capability across heterogeneous attack behaviours and diverse network environments.

Addressing the Research Questions

RQ1: Does a hybrid XGBoost–LSTM Autoencoder model improve intrusion detection performance?

Yes. The mixed model performed very well in the tests with a 99.45% accuracy rate, a 99.75% precision rate, a 99.644% recall rate and a 99.70% F1-score. The hybrid method performed better than the standalone models in attack discovery and detection accuracy.

RQ2: Can combining CICIDS2017 and TON_IoT datasets improve model generalization across different network environments?

The results indicate that the robustness of the proposed framework is improved by combining the CICIDS2017 and TON_IoT datasets. The feature-aligned dataset exposed the model to both enterprise and IoT traffic patterns enabling effective detection across a variety of attack categories and heterogeneous network environments.

RQ3: Does anomaly-aware sequence modelling help detect previously unseen attacks?

Yes. The unseen-attack evaluation revealed that the LSTM Autoencoder provided additional anomaly-awareness to the framework. The standalone autoencoder had limited recall, however, in combination with XGBoost, it was able to detect previously unseen attacks and increase the overall attack detection sensitivity.

Summary

The experimental results indicate that the proposed hybrid intrusion detection system combines supervised classification and anomaly-based analysis to achieve good overall detection performance. The method improves the robustness of detection against known attack types and unusual traffic characteristics. Separate models are limited . The hybrid approach is more balanced and more robust .

Results show that the combination of supervised learning with anomaly detection improves the detection performance and system reliability in complex network environments.

Evaluation Metrics

To assess the performance of an intrusion detection system, paper need appropriate metrics that can show the ability of the system to differentiate between normal and malicious data. The study fully analyzes the model performance by using many evaluation indicators, in particular, the class imbalance context. However, accuracy alone is not enough because it can be misleading when the dataset contains dominating class (Wang et al., 2017; Buczak and Guven, 2016).

The evaluation is done by using the confusion matrix that summarizes the classification performance in terms of true positives, true negatives, false positives and false negatives .

Confusion Matrix

The confusion matrix is defined as:

$$\begin{bmatrix} TN & FP \\ FN & TP \end{bmatrix}$$

where TN represents true negatives, FP represents false positives, FN represents false negatives, and TP represents true positives.

In intrusion detection, false negatives are particularly critical, as they correspond to undetected attacks.

Accuracy

Accuracy measures the overall proportion of correctly classified instances:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

Although accuracy is commonly used, it may not be reliable in imbalanced datasets (Ashfaq et al., 2017).

Precision

Precision measures the proportion of correctly predicted positive instances:

$$\text{Precision} = \frac{TP}{TP + FP}$$

High precision indicates fewer false positives.

Recall

Recall, also known as detection rate, measures the proportion of actual positive instances correctly identified:

$$\text{Recall} = \frac{TP}{TP + FN}$$

Recall is critical in intrusion detection, as it reflects the ability to detect attacks.

F1-Score

The F1-score is the harmonic mean of precision and recall:

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

It is particularly useful for imbalanced datasets.

False Positive Rate and False Negative Rate

The False Positive Rate (FPR) and False Negative Rate (FNR) provide additional insights:

$$FPR = \frac{FP}{FP + TN}$$

$$FNR = \frac{FN}{FN + TP}$$

A low FPR reduces false alarms, while a low FNR ensures fewer missed attacks.

ROC Curve and AUC

The Receiver Operating Characteristic (ROC) curve evaluates the trade-off between the True Positive Rate and False Positive Rate at different thresholds. The Area Under the Curve (AUC) summarizes overall performance.

A higher AUC indicates better performance, although ROC curves may be less informative for highly imbalanced datasets (Buczak and Guven, 2016).

Precision-Recall Curve

The Precision-Recall (PR) curve illustrates the trade-off between precision and recall.

PR curves are more suitable for intrusion detection tasks, as they focus on the minority (attack) class.

Summary

Various evaluation measures are used to thoroughly evaluate the model performance. Accuracy is a global metric, but precision and recall give a deeper insight into detection capability. F1-score is a weighted harmonic mean between these metrics, and FPR and FNR show trade-offs between false alarms and missed detections.

In this study, the use of ROC and Precision-Recall analysis provides a good evaluation framework in consistent with intrusion detection in imbalanced and real world scenarios.

Comparative Analysis

This section presents a comprehensive comparative analysis of the proposed hybrid intrusion detection system against individual models and existing approaches in the literature. The objective is to evaluate the effectiveness of the hybrid framework in terms of detection capability, robustness, and suitability for real-world deployment across heterogeneous environments.

The comparison is conducted at two levels: internal comparison among implemented models and external comparison with existing approaches reported in the literature.

Comparison with Baseline Models

The results in Section show that both Random Forest and XGBoost have very high accuracy and F1-scores, above 99%. These models understand complicated associations from labeled data and perform very well in the detection of known assault patterns.

The Random Forest has a slightly greater recall than XGBoost, which means it detects more attack instances. On the other hand, XGBoost has higher precision indicating that it has fewer false positives. This trade-off illustrates the benefits of ensemble-based supervised learning methods (Breiman, 2001; Chen and Guestrin, 2016).

Both models perform well . They require labels , however , and are most useful for recognizing established attack patterns . Their performance can suffer while facing undetected or novel threats (Buczak and Guven, 2016).

The LSTM Autoencoder is an anomaly detection model trained solely on normal traffic data. It is high precision but much lower recall, which means that it is cautious in spotting anomalies. This is consistent with the results of previous investigations (Shone et al., 2018a; Mirsky et al., 2018).

To overcome these constraints, the proposed hybrid model combines the supervised and the anomaly detection methodologies . The hybrid model enhances the recall with high precision while combining XGBoost with the LSTM Autoencoder, which helps to detect known and undiscovered threats.

Performance Comparison Across Models

On closer inspection of the two approaches, the hybrid technique has advantages. RF and XGBoost have good accuracy but are restricted to known attack patterns. LSTM Auto-Encoder LSTM Autoencoder is good for anomaly detection, but has low recall rate.

The hybrid model is well balanced in all measures with high accuracy, precision and recall. The increase in recall shows that it can identify a higher percentage of malicious activity.

Table 12: Comparison of Model Performance

Model	Accuracy	Precision	Recall	F1-score	FPR	FNR
Random Forest	99.58%	99.02%	99.01%	99.01%	0.26%	0.99%
XGBoost	99.60%	99.39%	98.71%	99.05%	0.16%	1.29%
LSTM Autoencoder	20.39%	99.30%	12.35%	21.97%	0.85%	87.65%
Hybrid Model	99.45%	99.75%	99.64%	99.70%	2.44%	0.36%

This comparison demonstrates that the hybrid model maintains the strengths of supervised learning while enhancing detection capability through anomaly detection.

Comparison with Existing Studies

To further evaluate the effectiveness of the proposed approach, comparisons are made with existing intrusion detection systems.

Chandrashekhar et al. (Chandrashekhar et al., 2026) proposed a hybrid framework integrating Autoencoder (AE) and CNN–LSTM models for cloud–IoT environments, achieving high classification accuracy. However, their framework relied heavily on computationally intensive deep learning architectures and was evaluated primarily on a single-domain dataset, which may limit scalability and generalization across heterogeneous environments.

XGBoost-based intrusion detection models reported in the

literature achieve high accuracy and precision (Chen and Guestrin, 2016); however, these approaches depend mainly on supervised learning and therefore exhibit limited capability in detecting previously unseen or zero-day attacks.

Autoencoder-based approaches (Shone et al., 2018a; Mirsky et al., 2018) demonstrate strong anomaly detection capability by learning normal traffic behavior. Nevertheless, these methods often suffer from low recall and sensitivity to reconstruction threshold selection, which may result in missed attack detections.

Table 13: Comparison with Existing Studies

Study	Approach	Strength	Limitation
Chandrashekhar et al.	CNN–LSTM + AE	High accuracy	High computational complexity
XGBoost-based IDS	Supervised ML	High precision	Cannot detect unknown attacks
Autoencoder-based IDS	Unsupervised DL	Detects anomalies	Low recall
Proposed Method	Hybrid ML + AE	Balanced performance	Slight increase in false positives

The comparison highlights the advantages of the proposed approach, particularly in terms of multi-domain dataset integration, improved generalization across cloud and IoT environments, and balanced detection performance for both known and unknown cyber threats. Furthermore, the proposed framework achieves a higher recall and F1-score while reducing dependency on computationally expensive deep learning architectures.

Impact of Hybrid Fusion

The hybrid fusion strategy plays a crucial role in improving overall model performance. The system combines results of supervised and anomaly detection algorithms, exploiting complimentary strengths.

The logical OR-based fusion:

$$\hat{y}_{Hybrid} = \hat{y}_{XGB} \vee \hat{y}_{AE}$$

ensures that a sample is classified as malicious if either model detects it as an attack. This significantly reduces false negatives, which is critical in intrusion detection.

Although this may increase false positives slightly, the trade-off is acceptable, as missing an attack is more critical than raising a false alarm.

Discussion of Generalization and Practical Applicability

One important feature of the proposed method is that it is applicable to a wide variety of networks. The model learns different traffic patterns and attack behaviors by merging the datasets of cloud and IoT domains.

Feature alignment allows consistent representation on datasets and thus reduces over-fitting and increases scalability.

This hybrid method can be applied in practice to network monitoring systems in real time. The hybrid technique is a combination of the supervised classification and anomaly based detection suitable to implementation in a dynamic and expanding network system where the traffic pattern changes over time.

Summary

The comparison investigation reveals that the hybrid intrusion detection system works better than the individual models and it addresses the primary shortcomings of the existing methodologies. The system integrates supervised learning and anomaly detection to increase the detection ability.

The results indicate that the combination of several models and datasets led to a more robust, scalable and effective intrusion detection solution for modern network security applications.

Discussion

In this section the experimental results are explained in detail. The behavior of individual models and the proposed hybrid framework is discussed. It also presents the advantages, limitations and practical implications of the proposed approach in modern intrusion detection systems.

The main point is that the supervised learning models, especially Random Forest and XGBoost, performed consistently well. Both models achieve accuracy and F1-scores above 99% and demonstrate their ability to learn patterns from labeled network traffic data. This performance is attributed to the ability of these methods to capture complex nonlinear relationships and interactions among features, which are common in network traffic datasets (Breiman, 2001; Chen and Guestrin, 2016). Also, tree based ensemble algorithms are noise tolerant and can efficiently handle high dimensional data which is suitable for intrusion detection applications.

Supervised models, although they perform very well, have a major limitation: they require labeled data. These models are trained on predefined attack patterns and have limited capability to detect previously undetected or zero day attacks. This constraint is well known in the study on intrusion detection and has motivated the usage of anomaly detection approaches (Buczak and Guven, 2016). Here we solve this problem by adding an LSTM Autoencoder to the system.

The performance profile of the LSTM Autoencoder is different from supervised models. It yields very high precision but very low recall, i.e., the detected anomalies are very accurate but many assault cases are missed. This behaviour is consistent with previous work where autoencoder based models tend to be conservative in anomaly detection (Shone et al., 2018b; Mirsky et al., 2018).

The low recall can be explained by several factors. Firstly, the model is trained only with typical traffic, which restricts

the ability to detect tiny anomalies that mimic normal patterns. Second, the threshold of the reconstruction error is important to the classification performance. A high threshold may cause assaults to be missed (false negatives). A lower threshold increases sensitivity, but may pay the price in more false positives. This trade-off poses a difficulty to the optimization of anomaly detection methods.

However, the LSTM Autoencoder can be embedded into the hybrid architecture to greatly improve the detection performance. The hybrid model significantly improves the recall to about 99.66%, indicating that the majority of harmful behaviors may be caught successfully. This shows the power of the combination of the supervised and anomaly detection methodologies.

The fusion strategy plays a crucial role in this improvement. By applying a logical OR operation:

$$\hat{y}_{Hybrid} = \hat{y}_{XGB} \vee \hat{y}_{AE}$$

The system will consider any instance reported by either model as harmful. This emphasizes detecting ability, something very important in cybersecurity when missing an assault is more dangerous than raising a false alarm. The unseen-attack evaluation indicated significant trade-offs between supervised and anomaly-based intrusion detection techniques. The XGBoost model was still able to obtain good recall performance on the attack families that were not seen previously. The LSTM Autoencoder was high in anomaly sensitivity at the cost of significantly reduced precision due to more false positives.

Also found that the OR-based hybrid fusion significantly enhanced the recall by integrating the supervised and anomaly-based predictions. However, the improved recall brought more false alarms and lowered the overall precision. The results indicate that anomaly-aware hybrid intrusion detection systems can improve the sensitivity to new threats, but also bring issues in controlling the false positive rates.

Another significant addition of this study is the integration of diverse datasets from different sectors. The integration of CICIDS and TON_IoT datasets makes the system more exposed to various traffic patterns and attack situations, therefore enhancing its generalization capacity. This overcomes a major problem of many previous studies relying on a single dataset (Ring et al., 2019). This connection is additionally supported by feature alignment which ensures consistent depiction of features.

The results further highlight the significance of preprocessing and feature engineering. Steps like duplication removal, standardization and feature alignment greatly increase data quality and model performance. These measurements are especially important in the multi-domain setting, because inconsistency among datasets can alter the results.

However, there are certain restrictions to the good performance. The LSTM Autoencoder alone has a low recall therefore there is a room for improvement. Techniques such as adaptive thresholding, innovative architectures, or attention techniques can boost the anomaly detection performance. Furthermore, the study uses a small number of datasets and more evaluation on more datasets could enhance generalization claims.

Another problem is the computational cost of deep learning models . Supervised models are reasonably efficient, but the LSTM Autoencoder demands a lot of CPU power, especially when working with sequence data. This may be challenging for real-time deployment in resource constrained environment.

From the practical point of view, the suggested hybrid intrusion detection system has various advantages. It is suitable for dynamic and changing network environment due to its capacity to identify both known and unknown assaults. It is more applicable in real life applications, because of integration of cloud and IoT data sets. Moreover, the modular design allows for flexibility in future integration of additional models or datasets.

The discussion in general shows that the suggested hybrid technique effectively tackles major issues in intrusion detection including unknown threat detection and multi-domain data integration. Despite restrictions, the framework provides a robust and scalable solution for modern network security applications.

Conclusion and Future Work

The hybrid intrusion detection framework is proposed in this study to overcome the issues in modern network settings especially cloud and IoT systems. This paper propose a supervised machine learning model combined with a deep learning based anomaly detection technique for efficient identification of known and novel cyber threats. In this paper, propose a new system that combines Random Forest, XGBoost and an LSTM Autoencoder in one framework. The system exploits the strengths of each model to produce a strong and balanced performance.

A main contribution of this work is the integration of heterogeneous datasets, CICIDS and TON_IoT, through a feature alignment technique. This allows to develop a multi-domain intrusion detection system that can work in different network contexts. We differ from many previous works that use a single dataset by demonstrating the possibility and effectiveness of mixing datasets from diverse domains to improve generalization.

The experimental results indicate that the supervised models, particularly XGBoost, are highly accurate and precise in detecting known attack patterns. But LSTM Autoencoder learns the typical behavior and detects the anomalies through reconstruction error and overcomes their restriction of identifying hidden threats. The anomaly detection model itself has a low recall. But when it is put into the hybrid framework, the overall detection ability is greatly enhanced.

The proposed hybrid model provides good results on all the evaluation metrics like accuracy, precision, recall and F1-score. The improvement in recall in particular shows the ability of the system to detect a significant amount of the malicious activity, thereby reducing the probability of unreported attacks. This highlights the effectiveness of the hybrid fusion approach, which exploits the output of both supervised and anomaly detection models to improve detection performance.

This work also shows the importance of data preparation, feature engineering and multi-domain learning for intrusion detection in order to enhance the performance. The results show that consistent model performance depends on careful

handling of data quality, feature alignment, and class imbalance. These are relevant aspects to the design of realistic intrusion detection systems that can be applied in real-world scenarios.

The results are promising, but there are a number of limitations. Moreover, the capability of the LSTM Autoencoder to identify tiny anomalies also depends on the choice of the threshold and the sequence length. Furthermore, the computational complexity of deep learning models can prevent their application in real-time applications in resource-constrained environments. Moreover, the evaluation is done on a few datasets that may not cover all possible network configurations.

Future work can further overcome these limitations by considering more advanced deep learning architectures such as attention-based models or transformer networks which can increase the performance of anomaly detection. Research can be done on adaptive thresholding approaches to change detection thresholds adaptively on the fly based on the data distribution. More experiments on other datasets and real-world network traffic using the proposed methodology would further help to establish its generalizability.

Another possible alternative is to build lightweight models for real-time intrusion detection in IoT systems with low computational power. The system may also be able to improve its performance on an ongoing basis by applying online learning algorithms to adapt to changing patterns of assault.

In summary, the proposed hybrid intrusion detection system provides an efficient and holistic solution for detection of cyber threats in complex and heterogeneous network environments. The framework consists of multiple models and datasets, which makes a good trade-off between accuracy and generalization, which is a promising direction for modern network security applications.

Data Availability Statement

The datasets supporting the findings of this study were obtained from publicly available sources. The *CICIDS2017* dataset was collected from the Canadian Institute for Cybersecurity (CIC), University of New Brunswick, and is publicly available at: https://www.unb.ca/cic/datasets/index.html

The *TON_IoT* dataset was obtained from the University of New South Wales (UNSW) Cyber Range Lab and is publicly accessible at: https://research.unsw.edu.au/projects/toniot-datasets

Conflicts of interest/Competing interests

The authors declare that they have no relevant financial or non-financial interests to disclose.

References

Ahmad, Z., Shahid Khan, A., Wai Shiang, C., Abdullah, J., and Ahmad, F. (2021). Network intrusion detection system: A systematic study of machine learning and deep learning approaches. *Transactions on Emerging Telecommunications Technologies*, 32(1):e4150.

Ahmed, U., Nazir, M., Sarwar, A., Ali, T., Aggoune, E.-H. M., Shahzad, T., and Khan, M. A. (2025). Signature-based intrusion detection using machine learning and deep learning approaches empowered with fuzzy clustering. *Scientific Reports*, 15(1):1726.

Alakhras, T. M. O., Ghanem, W. A. H. M., Yunus, F., Ghaleb, S. A. A., and Otair, M. (2025). A survey of intrusion detection systems in iot: Machine learning and feature selection approaches. *International Journal of Applied Mathematics*, 38(10s):828–884. Received August 09, 2025.

Aldweesh, A., Derhab, A., and Emam, A. (2020). Deep learning approaches for anomaly-based intrusion detection systems. *IEEE Access*.

Alsaedi, A., Moustafa, N., Tari, Z., Mahmood, A., and Anwar, A. (2020). Ton_iot telemetry dataset: A new generation dataset of iot and iiot for data-driven intrusion detection systems. *Ieee Access*, 8:165130–165150.

Ashfaq, R. A. R., Wang, X.-Z., Huang, J. Z., Abbas, H., and He, Y.-L. (2017). Fuzziness based semi-supervised learning approach for intrusion detection system. *Information sciences*, 378:484–497.

Axelsson, S. (2000). Intrusion detection systems: A survey. Technical report.

Bijone, M. (2016). A survey on secure network: intrusion detection & prevention approaches. *American Journal of Information Systems*, 4(3):69–88.

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.

Buczak, A. L. and Guven, E. (2016). A survey of data mining and machine learning for cybersecurity intrusion detection. *IEEE Communications Surveys & Tutorials*.

Chandola, V., Banerjee, A., and Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*.

Chandrashekhar, A. S., Patil, S. R., and Kounte, M. R. (2026). An intelligence framework for threat detection and response in cloud–iot-assisted enterprise environments. *Bulletin of Electrical Engineering and Informatics*, 15(1):546–556.

Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of KDD*.

Depren, O., Topallar, M., Anarim, E., and Ciliz, M. K. (2005). An intelligent intrusion detection system (ids) for anomaly and misuse detection in computer networks. *Expert systems with Applications*, 29(4):713–722.

Efe, A. and Abacı, İ. N. (2022). Comparison of the host based intrusion detection systems and network based intrusion detection systems.

Gad, A. R., Nashat, A. A., and Barkat, T. M. (2021). Intrusion detection system using machine learning for vehicular ad hoc networks based on ton-iot dataset. *IEEE access*, 9:142206–142217.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.

Hamdi, D. A., Bahaa-Eldin, A. M., and Sobh, M. A. (2026). A systematic literature review on deep learning-driven ids for software-defined wsn. *Journal on Wireless Communications and Networking*.

Hindy, H., Brosset, D., Bayne, E., Seeam, A., Tachtatzis, C., Atkinson, R., and Bellekens, X. (2018). A taxonomy of network threats and the effect of current datasets on intrusion detection systems. *arXiv preprint arXiv:1806.03517*.

Hinton, G. E. and Salakhutdinov, R. R. (2006). Reducing dimensionality with neural networks. *Science*.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*.

Hossain, S., Senouci, S.-M., Brik, B., and Boualouache, A. (2025). A privacy-preserving self-supervised learning-based intrusion detection system for 5g-v2x networks. *Ad Hoc Networks*, 166:103674.

Hosseini, S., Ebrahimi, A., Mosavi, M., and Shahhoseini, H. (2025). A novel hybrid cnn-cbam-gru method for intrusion detection in modern networks. *Results in Engineering*, 28:107103.

Issa, M. M., Aljanabi, M., and Muhialdeen, H. M. (2024). Systematic literature review on intrusion detection systems: Research trends, algorithms, methods, datasets, and limitations. *Journal of Intelligent Systems*, 33(1).

Iyer, K. I. (2021). From signatures to behavior: Evolving strategies for next-generation intrusion detection. *European Journal of Advances in Engineering and Technology*, 8(6):165–171.

Katta, V. K., Gadhiya, R. B., Habib, I., and Patavardhan, P. P. (2025). Anomaly detection in iot networks using ai machine learning and statistical models. *International Journal of Applied Mathematics*, 38(11s):2416–2429.

Khraisat, A., Gondal, I., Vamplew, P., and Kamruzzaman, J. (2019). Survey of intrusion detection systems: techniques, datasets and challenges. *Cybersecurity*, 2(1):1–22.

Kim, G., Lee, S., and Kim, S. (2014). A novel hybrid intrusion detection method integrating anomaly detection with misuse detection. *Expert Systems with Applications*, 41(4):1690–1700.

Kim, J., Kim, J., Thu, H. L. T., and Kim, H. (2016). Long short term memory recurrent neural network classifier for intrusion detection. In *Proceedings of ICNC*.

Kulkarni, T., Khadilkar, P., Kumar, B. R., and Jaiswal, R. (2026). A comparative study of different ids approaches and their performance. In Senjyu, T., So-In, C., and Joshi, A., editors, *Smart Trends in Computing and Communications*, pages 149–158, Singapore. Springer Nature Singapore.

Li, W., Tug, S., Meng, W., and Wang, Y. (2019). Designing collaborative blockchained signature-based intrusion detection in iot environments. *Future Generation Computer Systems*, 96:481–489.

Mirsky, Y., Doitshman, T., Elovici, Y., and Shabtai, A. (2018). Kitsune: an ensemble of autoencoders for online network intrusion detection. *arXiv preprint arXiv:1802.09089*.

Moustafa, N. (2020). Ton_iot datasets: A realistic dataset for iot intrusion detection. *IEEE Access*, 8:165130–165150.

Potluri, S. and Diedrich, C. (2016). Accelerated deep neural networks for enhanced intrusion detection. *IEEE*.

Qin, Z., Li, L., Luo, Q., Yu, X., and Zhang, Q. (2025). Enhancing intrusion detection performance using gcn-lof: A hybrid graph-based anomaly detection approach. *Computer Networks*, page 111606.

Ring, M., Wunderlich, S., Scheuring, D., Landes, D., and Hotho, A. (2019). A survey of network-based intrusion detection data sets. *Computers & security*, 86:147–167.

Sedjelmaci, H., Senouci, S. M., and Taleb, T. (2017). An efficient intrusion detection framework for iot. *IEEE*.

Sharafaldin, I., Lashkari, A. H., and Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. In *Proceedings of ICISSP*.

Shone, N., Ngoc, T. N., Phai, V. D., and Shi, Q. (2018a). A deep learning approach to network intrusion detection. *IEEE transactions on emerging topics in computational intelligence*, 2(1):41–50.

Shone, S., Ngoc, T. N., Phai, V. D., and Shi, Q. (2018b). A deep learning approach to network intrusion detection using autoencoders. *IEEE Access*.

Tavallaee, M., Bagheri, E., Lu, W., and Ghorbani, A. A. (2009). A detailed analysis of the kdd cup 99 data set. In *2009 IEEE symposium on computational intelligence for security and defense applications*, pages 1–6. Ieee.

Wang, W., Sheng, Y., Wang, J., Zeng, X., Ye, X., Huang, Y., and Zhu, M. (2017). Hast-ids: Learning hierarchical spatial-temporal features using deep neural networks to improve intrusion detection. *IEEE access*, 6:1792–1806.

Yin, Y., Zhu, J., Fei, J., and He, X. (2017). A deep learning approach for intrusion detection using recurrent neural networks. *IEEE Access*, 5.