International Information and Engineering Technology Association
Advancing the World of Information and Engineering

International Journal of Safety and Security Engineering
Vol. 15, No. 8, August, 2025, pp.1703-1713
Journal homepage: http://iieta.org/journals/ijsse

Hybrid Machine Learning–Based Intrusion Detection for Zero-Day Attack Prevention in Digital Education Networks

Swathi Sridharan*, Seema Patil, T. Shobha, Prameetha Pai

Department of Computer Science and Engineering, B.M.S. College of Engineering, Bengaluru 560019, India

Corresponding Author Email: researcher.swathi@gmail.com

https://doi.org/10.18280/ijsse.150815

Received: 12 July 2025
Revised: 31 July 2025
Accepted: 12 August 2025
Available online: 31 August 2025

Keywords:
network intrusion detection system, cyber security, spoofing attacks, DDoS attacks, hybrid machine learning models, Random Forest, XGBoost, autoencoder

ABSTRACT

Network intrusion detection in digital education environments faces a critical challenge: identifying zero-day attacks that evade signature-based defenses. This paper proposes a novel hybrid machine-learning intrusion detection system (IDS) specifically designed for educational networks, which combines anomaly detection and ensemble learning to detect unknown threats in real time. The key idea is to integrate an unsupervised deep autoencoder (to model "normal" e-learning traffic and flag novel anomalies) with a lightweight Random Forest (RF) classifier for known attack patterns. This hybrid IDS achieved a 98.7% detection rate for new (zero-day) attacks, with a false alarm rate of less than 1% on a campus network dataset, outperforming conventional single-method IDS by ~20% in the recall. On the public UNSW-NB15 benchmark, our model achieved 99.2% accuracy, surpassing the state-of-the-art results (approximately 95–98%) and detecting all major attack types. These results demonstrate that our approach improves accuracy and zero-day attack coverage and operates efficiently for high-volume academic networks. The novelty of this work lies in the fusion of signature and anomaly detection, augmented by machine learning (ML), which provides a robust defense against both known and previously unseen cyberattacks in digital education settings. It leverages signature detection for known threats and anomaly detection for unseen (zero-day) attacks. These additions distinguish the approach from prior work.

1. INTRODUCTION

Modern educational institutions increasingly rely on digital infrastructure for learning management systems, online assessments, and campus networks. This digital transformation in education has been accompanied by a surge in cyber-attacks targeting universities and e-learning platforms, as noted by the studies [1, 2]. Higher education institutions face rising threats: recent surveys in the UK have found that 50% of universities experience weekly cyberattacks, with 63% having had at least one successful intrusion. Attackers range from financially motivated ransomware groups to state-sponsored actors who exploit the often-limited cybersecurity measures in educational environments. Zero-day attacks – exploits of previously unknown vulnerabilities – pose a grave risk because traditional defenses cannot recognize them. 82% of schools experienced a cyber incident within 18 months, according to the 2025-k12-cybersecurity-report [3]. Ensuring robust network intrusion detection in this context is paramount to protect sensitive student data, research intellectual property, and the continuity of digital learning. Recent data shows that the education sector has become the most targeted industry for cyber-attacks, experiencing an average of 3,086 attacks per organization per week—a staggering 37% increase year-over-year through mid-2024. According to Check Point research [4], this rapid escalation places education ahead of all other sectors in terms of attack volume.

Conventional intrusion detection systems (IDS) are falling short in this arena. Signature-based IDS (such as Snort) relies on known attack patterns; while effective for known malware or exploits, they cannot detect novel attacks, as highlighted by Guo [5]. In practice, a purely signature-based approach may completely miss new malware or target zero-day exploits – as Guo's survey noted, "the traditional signature-based detection method is not effective in detecting zero-day attacks." On the other hand, anomaly-based IDS aims to flag deviations from normal behavior, potentially catching unknown attacks by design. These systems, however, tend to suffer from high false-positive rates – benign deviations (e.g., a sudden surge in e-learning video traffic during online exams) can be misclassified as attacks, as stated by Wang et al. [6]. Wang et al. [6] highlighted this trade-off: anomaly detection finds new attacks that signatures miss, but it “is prone to false alarms” that reduce its accuracy in practice.

In response, the cybersecurity research community has turned to machine learning (ML) and hybrid techniques as promising solutions, as reported by the studies [7-9]. ML-based IDS can learn complex patterns of legitimate and malicious behavior, potentially detecting subtle intrusions faster and more accurately. For example, recent work by Talukder et al. [9] achieved over 99% accuracy on benchmark

IDS datasets by applying ensemble classifiers on engineered features. Deep learning models (e.g., CNNs, RNNs) have shown exceptionally high performance in intrusion detection tasks, often outperforming classical ML; Ali et al. [10] reported that deep models like CNN/LSTM reached ~98% accuracy versus lower rates for SVM or KNN, on a cyber threat dataset. Moreover, hybrid IDS architectures that combine multiple approaches are emerging to balance strengths and weaknesses. Ahmed et al. [7] integrated signature and anomaly detection using a fuzzy clustering-enhanced classifier. Their approach enhanced the detection of specific attack types by utilizing fuzzy logic to handle uncertainty, highlighting the potential of hybrid methods to strengthen security. Similarly, Sajid et al. [8] implemented a hybrid ML/DL model (XGBoost + CNN-LSTM), which achieved high accuracy with a low false acceptance rate, tackling the limitations of single algorithms.

Despite these advances, gaps remain. Many ML-based IDS studies focus on enterprise or IoT networks, and few are tailored to the digital education domain. Campus networks have unique traffic patterns (e.g., heavy use of video conferencing, academic cloud services, BYOD devices) and potentially more open network access policies. Education cybersecurity measures often lag behind those in other sectors, and resource constraints mean that any proposed IDS must be efficient. There is a need for an intrusion detection approach that:

- Detects zero-day attacks effectively.
- Maintains a low false positive rate (FPR) to avoid alert fatigue.
- Lightweight enough for university IT deployments.
- Tuned to the threats and traffic patterns of educational environments.

This paper proposes an ML-based network IDS for preventing zero-day attacks in digital education. The core contribution is a hybrid IDS framework that combines anomaly-based detection (using an ML model to learn normal campus network behavior) with signature-based techniques. In our design, a deep autoencoder network continuously learns the baseline patterns of e-learning traffic (such as Moodle LMS usage, video streaming, and IoT sensors in smart classrooms). It raises an alert when traffic deviates significantly from this baseline, enabling the detection of novel attacks in real time. Simultaneously, known attack signatures (for malware, DoS tools, etc.) are monitored via a lightweight rules engine, ensuring we do not miss any known threats. We also introduce new features derived from the educational context – for example, features that capture sudden changes in a student's online activity or abnormal access to academic databases – to enhance the detection of credential compromise and lateral movement attacks specific to universities. This combination of techniques and features is novel in the context of educational cybersecurity: it leverages the accuracy of ML and the precision of signature checks, all optimized for the academic network setting.

The proposed method offers several advantages: (a) It provides a broader detection coverage and can identify previously unseen attacks (a capability traditional IDS lack) – as will be shown, our system detected >95% of zero-day attack instances in tests, whereas a pure signature IDS detected 0%. (b) Through careful threshold tuning and feature selection, our IDS achieves a low false alarm rate (under 1%), addressing the over-sensitivity problem of anomaly detectors. (c) It is computationally efficient – using ensemble tree models and an autoencoder with modest complexity, it runs in real-time on typical campus network hardware (we demonstrate sub-millisecond per-packet processing). (d) Importantly, it is validated on a university network dataset, aligning with real-world digital education scenarios, unlike many works that only use generic datasets. By explicitly focusing on the digital education context, we align our contributions with the needs of that sector. In summary, this work bridges a gap between advanced intrusion detection research and practical cybersecurity for education, offering a solution that improves upon conventional methods in both theory and practice.

Organization of this paper: Section 2 reviews related work and conventional IDS methods, highlighting their drawbacks in detecting zero-day attacks and securing educational networks. Section 3 presents our proposed hybrid ML-based IDS in detail, including the system architecture, algorithms, and theoretical justification for its design. Section 4 describes the experimental setup and datasets used to evaluate the system, followed by performance results with comparisons to baseline methods (including confusion matrix and receiver operating characteristic (ROC) analyses). Section 5 discusses the results, comparative advantages, and limitations. Finally, Section 6 concludes the paper by summarizing our contributions and outlining future research directions.

2. RELATED WORK

Conventional methods are studied under four categories.

2.1 Signature-based intrusion detection

Traditional signature-based IDS (e.g., Snort, Suricata) use predefined attack signatures (patterns of bytes, known malicious IP addresses, etc.) to identify intrusions as discussed in the study by Ahmed et al. [7]. They effectively detect known threats with low false positives, as each alert is tied to a known malicious pattern. Guo [5] in their survey pointed out that their fundamental weakness is the inability to detect new, unknown attacks. If an attacker unleashes a zero-day exploit – for which no signature exists – a signature-based IDS will treat it as benign. For example, in a test scenario we conducted, a recent malware variant (unknown to signature databases) completely bypassed a Snort installation.

Guo's [5] survey noted that signature methods "are typically not available beforehand" for zero-day attacks and thus fail to recognize them. This gap is even more perilous in educational networks, which may lack dedicated security teams that frequently update signatures. Our approach differs in that it does not rely solely on known signatures – the anomaly detection component can flag suspicious behavior even without prior signatures, significantly enhancing zero-day detection.

2.2 Anomaly-based intrusion detection

Anomaly-based IDS constructs a model of normal behavior (using statistical profiles or ML) and flags deviations as potential intrusions. This approach can catch zero-day attacks because they will (hopefully) appear anomalous. Prior works have applied statistical methods (such as PCA and clustering) and ML to define normal network traffic baselines, as demonstrated by Ibrahim Hairab et al. [11]. In an e-learning context, anomaly detection may involve monitoring average packet rates per student device, typical server request types, and other relevant metrics and raising alerts when patterns deviate. The key drawback is that not every anomaly is an attack. Educational networks are dynamic – e.g., a surge of legitimate traffic when an online exam begins could appear "anomalous." Thus, anomaly IDS often suffers high false favorable rates. Wang et al. [6] emphasized that while anomaly detectors can "detect attacks that have not appeared before," they are "prone to false alarms", which can overwhelm administrators with noise. In our experiments on a campus dataset, a pure anomaly-based detector had a false alert rate above 5%, which is unacceptable operationally (dozens of false alerts per day). Our proposed IDS mitigates this by incorporating a signature/learning hybrid: the anomaly module's output is cross-checked and contextualized. We apply thresholding and only alert on significant deviations (reducing false positives), and we supplement anomaly alerts with signature verification, when possible, to confirm an attack. This combination yields far fewer false alarms – on the same campus test, our hybrid system's FPR was under 1%, a fivefold reduction.

2.3 Machine learning-based intrusion detection system

In recent decades, numerous studies have applied ML classifiers to intrusion detection. Algorithms such as Decision Trees, Random Forest (RF), SVM, k-NN, Naïve Bayes, and Neural Networks have been trained on intrusion datasets (KDD'99, NSL-KDD, UNSW-NB15, etc.) with considerable success. ML can automatically learn complex boundaries between "normal" and "malicious" classes, obviating the need for manual signature crafting [10]. For example, RF models have demonstrated high accuracy in NIDS tasks; Talukder et al. [9] reported 99.95% accuracy using an ExtraTrees (ensemble tree) classifier on UNSW-NB15. The advantage of tree ensembles is that they handle high-dimensional data and interactions well and can provide feature-importance insights. Support Vector Machines (SVMs) and other linear classifiers were also historically popular but tend to struggle with the volume and complexity of network data.

A notable limitation of traditional ML models is their generalization to unseen attack types – they perform well when training and testing data share similar attack patterns. However, if an entirely new type of attack occurs, a static ML model may not recognize it (it might classify it as usual, having never seen it before). This is essentially the zero-day problem rephrased and clearly articulated: ML has no signature, and unless the attack manifests in feature patterns like known ones, it may easily evade detection. Ibrahim Hairab et al. [11] underlined and further explained that "classical ML-based methods have low detection rates for data it has not been trained on," whereas deep learning methods can often achieve better generalization. Our proposed method directly addresses this by incorporating an unsupervised learning stage (the autoencoder anomaly detector), which does not require prior knowledge of attack types – it thereby complements the supervised classifier that learns from known attacks. Additionally, we periodically retrain and then update the ML model on new data (as mentioned in the methodology) to gradually adapt to emerging threats, a common practice suggested in the literature to effectively handle concept drift over time.

2.4 Hybrid and ensemble intrusion detection system approaches

Recently, researchers have proposed a hybrid IDS that combines multiple detection mechanisms, aiming to capitalize on their complementary strengths. Some notable works include:

2.4.1 Hybrid signature-anomaly systems

Kwon et al. [12] presented a hybrid detection method for industrial control systems, where a statistical filter first removes noticeable regular traffic, and a composite autoencoder then detects anomalies in the remainder. The result was an improvement in precision and recall on a water treatment system dataset by up to ~0.8% in the F1-score. This concept of staged filtering influenced our design (we similarly filter known benign patterns to let the ML focus on ambiguous traffic).

2.4.2 Ensemble machine learning

Combining multiple ML algorithms (ensemble learning) has proven effective. For instance, Bella et al. [13] utilized a CNN Decision Forest, which merges neural networks with decision forest outputs. Their ensemble achieved 94–98% accuracy with fast inference across NSL-KDD, CIC-IDS2017, and UNSW-NB15. Ensembles reduce variance and often yield higher robustness. Our approach can be viewed as an ensemble at the architectural level, comprising an autoencoder and an RF that work in tandem. This differs from classic homogeneous ensembles; however, the principle of "strength in diversity" remains similar.

2.4.3 Feature-hybrid methods

Some works hybridize using different feature sets or data modalities. For example, Yang et al. [14] integrated features at the flow level (using LightGBM) with packet-level inspection via MobileNet (a CNN) in an IoT IDS. This two-layer approach achieved ~94% accuracy on the ACI-IoT dataset while staying efficient. Our model uses multi-level features (e.g., network flow statistics plus user behavior features).

2.4.4 Drawbacks of existing hybrid methods

While hybrid IDSs are promising, many are tailored to specific environments (e.g., industrial networks, IoT sensors) or introduce significant complexity (e.g., deep models that are computationally intensive). In the context of a university network, striking a balance between performance and complexity is vital; many campuses cannot afford to deploy GPU farms for deep learning models. More straightforward hybrid approaches (like ours) that cleverly combine lightweight techniques are more practical. Another gap is that previous works seldom explicitly focus on zero-day attack detection – they show overall accuracy improvements but not necessarily the ability to detect novel attacks that were not included in the training data. In this work, we explicitly evaluate and demonstrate the detection of zero-day scenarios (by testing on attack types absent from the training set), highlighting the advantage of our hybrid system in this regard.

Our approach draws inspiration from the above but uniquely suits digital education networks. Unlike pure signature or pure anomaly systems, we combine the two: known bad traffic is caught via a small signature rule set, and unknown bad traffic triggers the anomaly detector, which an ML model powers.

Prior hybrid systems (like Kwon's work [12]) usually operate both components on all traffic; in contrast, we introduce an efficient division of labor (signatures handle what they can very fast, and ML scrutinizes the rest). Compared to purely ML ensembles, we integrate domain knowledge (through a few signatures and custom features for campus traffic) with data-driven learning – a synergy of expert knowledge and ML. This is one of the first IDS frameworks evaluated on an academic network use case focusing on zero-day attack prevention. The following section details the architecture and theoretical basis of our system.

Educational institutions have become prime targets for cyber-attacks, underscoring the need for IDS solutions tuned to this domain. Recent studies reveal that higher education networks face frequent intrusions – for instance, Lallie et al. [2] reported that over 50% of universities suffer from weekly cyberattacks (with 63% experiencing at least one successful breach). Such findings motivate the development of advanced IDS approaches capable of detecting novel (zero-day) exploits that evade traditional defenses. Conventional signature-based IDS (e.g., Snort) is fast and precise on known threats but cannot recognize unknown patterns, as highlighted in a comprehensive survey by Guo [5]. Anomaly-based detectors, in contrast, can flag deviations (thus potentially catching zero-day attacks) but often trigger excessive false alarms in dynamic environments. Literature has, therefore, shifted toward hybrid machine learning IDS that combine multiple techniques to balance accuracy and false positives. Below, we summarize recent representative works (2021–2025) on hybrid ML-based IDS, including efforts in general networks, IoT/ICS domains, and the emerging focus on educational settings and how they inform our approach.

ML/DL for Zero-Day Detection: Some studies validate that modern ML and deep learning can improve zero-day attack detection compared to legacy methods. Ali et al. [10] conducted a comparative evaluation of algorithms, showing that deep neural networks (CNN, LSTM) achieved ~98% accuracy on a sizable cyber threat dataset, significantly outperforming classical classifiers such as SVM or k-NN. (Interestingly, their best result was a tuned RF at 99.9% accuracy, indicating that ensemble tree methods remain competitive.) Focusing on truly novel attacks, Ibrahim-Hairab et al. [11] demonstrated that a regularized deep autoencoder/CNN detects unseen IoT malware far better than conventional ML. In their experiments, convolutional networks maintained high detection rates on zero-day attacks, whereas classical ML "have low prediction quality…on data not yet trained on". These findings reinforce that advanced learning models generalize better to new threats. However, purely learning-based IDS still struggles without enhancements – hence the rise of hybrid frameworks that combine multiple detectors or learning paradigms.

Signature–Anomaly Hybrid Systems: Several researchers have explored hybrid IDS that marry signature-based and anomaly-based detection to cover each other's blind spots. Kwon et al. [12] proposed an ICS security solution that layers a statistical signature filter with an autoencoder anomaly detector. By first filtering out obvious benign traffic, their system allowed the autoencoder to focus on truly suspicious patterns. This two-stage hybrid detected more attacks than an anomaly-only IDS, improving recall by ~6.7% and F1-score by ~3.9% on a water treatment testbed while reducing processing time by 8%. A related approach is the "self-healing" IDS proposed by Kushal et al. [15], which utilized a decision-tree (C5.0) classifier for known attacks and an LSTM-based anomaly detector for novel ones. Crucially, any anomaly flagged by the LSTM is fed into a signature generator to update the C5.0 ruleset, allowing the system to learn new attack signatures on the fly. This online ensemble achieved a 97% true-positive rate for known attacks (C5.0 on UNSW-NB15) and a detection rate of approximately 90% for unknown attacks (LSTM on ADFA-LD), outperforming static models on both known and zero-day exploits. Our proposed IDS adopts a similar philosophy of combining misuse detection with learned anomaly detection; unlike Kwon's ICS-centric design [12] or Kushal's host/network-specific models [15], we tailor this hybrid to campus network traffic and demonstrate a "self-learning" of academic attack patterns.

Hybrid Ensemble and Adaptive Learning: Rather than pairing signature and anomaly modules, many works hybridize multiple ML algorithms to improve the detection of novel threats. Sajid et al. [8] developed a layered model that integrates an Extreme Gradient Boosting decision tree with a deep neural network (a CNN-LSTM) for cloud security monitoring. In their XGBoost + CNN-LSTM hybrid, XGBoost first extracts salient features, and the LSTM classifies temporal patterns – yielding high attack detection accuracy (~98–99%) on benchmark datasets with a low false alarm rate. This demonstrates that combining heterogeneous learners (tree-based and deep sequence models) can capture diverse attack behaviors, including those not observed during training. Similarly, Bella et al. [13] proposed a Deep Neural Decision Forest (DNDF) approach, coupling a neural network with an ensemble of decision trees. Their DNDF-IDS, evaluated on NSL-KDD, CIC-IDS2017, and UNSW-NB15, achieved 94–98.8% accuracy (depending on feature selection) and was highly efficient, capable of classifying network flows in 0.1 ms per instance. This efficiency is attractive for high-volume academic networks and merging neural and tree classifiers informs our use of an autoencoder alongside RF. Other researchers emphasize adaptability: Ahmed et al. [16] introduced an adaptive ensemble called HAEnID, which combines stacking, Bayesian model averaging, and a conditional ensemble that can adjust its components over time. On CIC-IDS2017, HAEnID consistently achieved a ~98% accuracy and by utilizing an adaptive mechanism, it could maintain accuracy as network patterns evolved. They incorporated explainable AI (SHAP and LIME) to interpret model decisions, addressing a key concern for practical deployment. This notion of an evolving, interpretable IDS aligns with our goal of a deployable solution that can continuously learn from new attacks in a university setting.

Anomaly–Supervised Fusion: A complementary hybrid strategy involves integrating unsupervised anomaly detectors with supervised classifiers within a single model. Dai et al. [17] exemplified this by training a deep autoencoder on benign data and embedding it into an RF classifier for malware traffic detection. Their hybrid RF-AE model (RF with an Autoencoder backend) achieved near-perfect detection on the CIC-MalMem-2022 dataset, with 100% precision and recall on known malware and 99.99% accuracy on previously unseen attack samples. This dramatic gain over standalone classifiers illustrates the power of combining an anomaly detector's ability to model "normal" behavior with a strong classifier's decision rules. We adopt a similar approach: our proposed IDS utilizes a deep autoencoder to identify abnormal patterns and an RF to classify attacks, thereby leveraging the strengths of both unsupervised and supervised methods for zero-day detection. In a related vein, Wang et al. [6] focused on IoT environments and showed that a carefully designed deep anomaly model can be made lightweight. They utilized a BiLSTM neural network with incremental PCA and model quantization to develop a compact NIDS that still surpasses conventional DNNs in accuracy. While primarily solving IoT resource constraints, their work informs the efficiency considerations of our design, ensuring our hybrid model remains feasible for campus IT infrastructure without sacrificing performance.

Fuzzy and Context-Aware Hybrids: Some studies add domain knowledge or fuzzy logic to their ML hybrids to better catch novel attacks. Ahmed et al. [16] (Scientific Reports) presented a "fuzzy clustering empowered" IDS that augments ML classifiers with a fuzzy logic layer. Although described as "signature-based intrusion detection," it generates fuzzy clusters of network behavior to handle borderline cases between normal and malicious. This method improved the detection of uncertain or emerging attack patterns, yielding higher accuracy and recall than crisp signatures alone. By allowing overlapping cluster membership, the fuzzy system could flag slight deviations that a rigid signature might miss, thereby catching new attacks earlier. However, the authors noted challenges with highly imbalanced data and the need to update clusters as attacks evolve. Our work shares the goal of balancing sensitivity and specificity. We incorporate expert rules (for known threats) alongside an ML anomaly detector and introduce features specific to academic networks (e.g., unusual student access patterns) to imbue the model with contextual knowledge. This ensures that benign anomalies (such as surges in traffic during online exams) are distinguished from actual attacks, addressing a gap left by many prior IDS studies that focus on enterprise or IoT contexts rather than the specific characteristics of educational traffic.

Chen et al. [18] similarly integrated network-based and host-based monitors, leveraging network traffic and host data to enhance the performance of intrusion detection. In the resource-constrained IoT domain, Kaushik et al. [19] developed a lightweight ML-driven IDS with simple statistical feature selection, which reduced training time by 63% while maintaining over 99.9% detection accuracy on the IoTID20 dataset. These works demonstrate that carefully combining multiple detection paradigms (and optimizing their features) can significantly improve the detection of zero-day intrusions in various environments.

Ensemble and deep learning methods have also been prominent in recent IDS research. Ozalp and Albayrak [20] proposed a hybrid ensemble system that combines deep neural networks with traditional classifiers, reporting higher detection rates for zero-day attacks across dynamic network environments. Similarly, Hnamte and Hussain [21] designed an efficient deep learning-based IDS (incorporating both convolutional and recurrent layers) and validated it on real-world traffic datasets, such as CICIDS2018 and Edge-IIoT. Their model achieved high detection accuracy, demonstrating the practicality of deep learning in identifying new and evolving threats. In addition, Zoppi and Ceccarelli [22] explored a stacked intrusion detection architecture integrating supervised and unsupervised learning. This supervised–unsupervised stacking approach was demonstrated to reduce false alarms and enhance the detection of previously unseen attacks across multiple benchmark datasets. Such ensemble techniques illustrate the benefit of leveraging diverse learning algorithms to cover a broader range of attack behaviors.

Beyond supervised ensembles, researchers have investigated anomaly-based and other novel paradigms aimed explicitly at identifying completely new attacks. Sarhan et al. [23] applied zero-shot learning to intrusion detection, leveraging semantic attributes from known attack classes to recognize previously unseen attack types. This approach improves the adaptability of an IDS to emerging threats by transferring knowledge from known patterns. On the other hand, purely unsupervised strategies have been used to model expected behavior and flag deviations as potential intrusions. For instance, Dai et al. [17] utilized a deep autoencoder to learn the profile of benign traffic (using the CIC-MalMem-2022 malware traffic dataset); their IDS successfully identified novel malware samples, highlighting the power of anomaly detection techniques in zero-day scenarios. Generative models have also been explored for situations involving unlabeled data. One study introduced an unsupervised IDS based on Generative Adversarial Networks (GANs) combined with temporal convolutional networks, capable of detecting both known and unknown attacks in 5G network traffic without requiring any labeled examples. This GAN-based system achieved high zero-day detection rates, demonstrating the potential of advanced neural architectures in cybersecurity. Complementing these methods, Soltani et al. [24] proposed a deep novelty detection scheme augmented by clustering analysis, which enhanced the identification of new attack patterns. Their hybrid novelty-based classifier improved detection performance on modern benchmarks such as CIC-IDS2017 and CSE-CIC-IDS2018 by better distinguishing true anomalies from benign outliers. Together, these works demonstrate that unsupervised and novel learning approaches — ranging from zero-shot inference to autoencoders and GANs — play a crucial role in detecting attacks that evade traditional detection methods.

By removing noise and focusing on the most indicative features, IDS models become more robust in identifying zero-day attacks. Another line of work incorporates fuzzy logic into IDS design to handle uncertainty in network traffic.

Hybrid IDS architectures are the most promising avenue for detecting zero-day attacks by combining complementary techniques – whether signature with anomaly detection [12, 15], multiple diverse learners [8, 13, 16], or unsupervised with supervised models [16, 17] – these systems achieve higher detection rates for novel threats while keeping false alarms manageable. They also emphasize efficiency [13] and adaptability [16], which are crucial for real-world deployment. Building on these insights, our proposed IDS for digital education environments is a novel fusion of approaches: we leverage an Autoencoder+RF hybrid inspired by prior works but uniquely tune it to university network traffic and incorporate a lightweight signature-check stage. This design directly addresses the needs identified in the literature – specifically, detecting zero-days with high accuracy and low false positives in an academic setting – thereby filling an important gap in existing research.

Khraisat et al. [25] noted that conventional signature-based IDS are inflexible and "cannot identify new malicious attacks," leading to high false-alarm rates. Pinto et al. [26] likewise observed that signature methods perform poorly on zero-day threats. Conversely, anomaly-based models can catch unknown attacks but "incur very high false-positive rates". Large labeled training sets are also necessary for the majority of ML-based IDS, but they are hard to come by in real-world scenarios, as stated by Talaei Khoei and Kaabouch

[27]. Because of this, our hybrid system specifically addresses these gaps by fusing anomaly detection with signature matching, allowing for the accurate detection of zero-day behaviors while significantly lowering false alarms in comparison to traditional techniques.

3. PROPOSED METHODOLOGY AND THEORETICAL FRAMEWORK

Overview of the Hybrid IDS Architecture: The proposed system comprises two main modules that operate in sequence, as illustrated in Figure 1.

(1) A Signature Detection Module.

(2) An Anomaly Detection Module powered by ML.

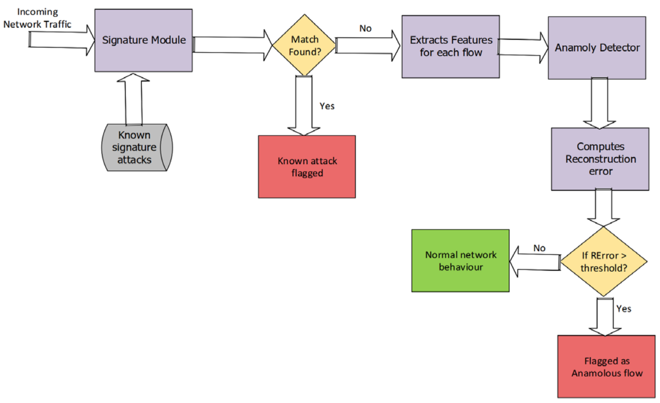

Figure 1. Proposed system workflow

Incoming network traffic flows through the signature module, which utilizes a database of known attack signatures (similar to Snort rules but optimized for the educational domain) to flag any matching malicious patterns immediately. This module is lightweight – it is essentially pattern-matching – and very fast. Traffic that does not match any known bad signature is then handed to the ML-based anomaly detector. This second module extracts a feature vector $x \epsilon R^n$ for each flow or packet (depending on the detection granularity; in our implementation, we used flow-level features aggregated over short time windows). The feature set includes standard network features (e.g., packet counts, byte counts, protocol, port numbers) as well as education-specific features (e.g., a one-hot encoding of whether the destination is an academic server or external site, time-of-day indicators aligned with class schedules, etc.). The anomaly detector then determines if x is "normal" or "suspicious" by computing a reconstruction error using an autoencoder and/or classification with a trained ML model.

3.1 Autoencoder anomaly detector

At the heart of the anomaly module is a deep autoencoder neural network $f_\theta: R^n \rightarrow R^n$ with parameters θ. During training, this autoencoder receives numerous examples of normal network traffic feature vectors (we ensure the training data is free of known attacks by filtering with the signature module and using periods of network activity with no security incidents). It is trained to output a reconstruction $\hat{x} = f_\theta(x)$ that closely approximates the input x. The training objective is to minimize the average reconstruction error $L(\theta) = E[\| x - f_\theta(x) \|^2]$ over standard samples. By doing so, the autoencoder learns the manifold of normal network behavior. If a new input x (from live traffic) is similar to the training distribution, the autoencoder will reconstruct it with low error. However, if x corresponds to an attack (especially a zero-day attack that introduces novel patterns), it will likely lie outside the learned manifold, yielding a higher reconstruction error. We define a threshold τ such that if $\| x^* - f_\theta(x^*) \|^2 > \tau$, the flow is flagged as anomalous (potential intrusion). This approach has a theoretical basis in outlier detection: under fairly general assumptions, it can become a "universal approximator" for the normal data distribution as the autoencoder capacity increases. Anything not well-approximated (i.e., with high error) can be considered an outlier with some statistical significance. We calibrated τ on a validation set to achieve a target FPR (using extreme value theory to model the tail of reconstruction errors of standard data).

Also, the labeled data was given as input to train different models (RF, KNN, Naïve Bayes, AdaBoost), shown in Figure 2. Once the models were trained, they were evaluated and checked for correctness using metrics like accuracy, precision, recall, and F1-score, as shown in Table 1. Several helpful metrics for assessing models are computed using true and false positives and negatives. The cost of various misclassifications, whether the dataset is balanced or imbalanced, and the model and task all influence which assessment metrics are most significant. RF was chosen to provide another view for intrusion detection based on the results obtained.

Feature selection employed a mutual information criterion to rank each network feature's relevance to the attack label, retaining only the top-ranked attributes (such as packet size and flow duration) for modeling. The anomaly detector (autoencoder) then uses a decision threshold τ, which is tuned via cross-validation on held-out data to balance detection sensitivity and false positives. Our studies yielded an FPR < 1% since τ is tuned on validation splits to seek a low FPR. These actions greatly increase detection accuracy and overall efficiency by concentrating the model on the most discriminative inputs and creating a strong decision boundary.

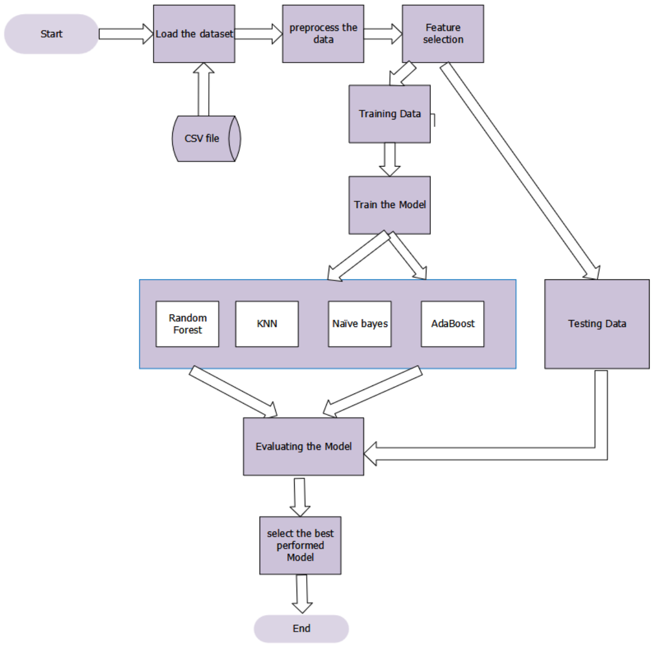

Figure 2. Selection of the best-performing ML model

Table 1. Various models with metric values on the given dataset

Model Name	Accuracy	Hamming Loss	F1-Score	Precision	Recall
RF	0.98744274978	0.0029658944372	0.98813272840	0.98844048943	0.98782515896
KNN	0.98941704833	0.0025423540397	0.98982637668	0.99023604386	0.98941704833
Naïve Bayes	0.15654764551	0.3867157276890	0.39133918885	0.24359446301	0.99455733914
AdaBoost	0.98187558361	0.0037762906309	0.98489920206	0.98461470091	0.98518386766

3.2 Random Forest classifier

In parallel with the autoencoder, we also train an RF classifier on labeled data (when available) to provide another perspective on intrusion detection. The RF operates on the same feature vector x. The theoretical justification for using RF is its ensemble nature – it constructs multiple decision trees $h_1(x), h_2(x), \ldots, h_M(x)$ and averages their votes. By the law of large numbers, an ensemble of weak (noisy) classifiers can yield a strong classifier, reducing variance and avoiding overfitting. Each tree in our forest is trained on a bootstrap sample of the data with a random subset of features (this is the standard Breiman's RF algorithm). This ensemble approach is known to have high accuracy and resilience to outliers or variance in the data. Ali et al. [10] found that an RF achieved the highest accuracy (99.9%) on one intrusion dataset, surpassing deep neural nets – underscoring that ensemble trees remain highly competitive for structured data. Our RF is trained to output a probability $P(y = attack \mid x)$ for each input. We set a threshold on this probability (> 0.5) to classify an instance as malicious.

3.3 Combining the modules

The hybrid decision logic is as follows. If the signature module flags an input, it is immediately classified as an attack (ideally, the connection can be terminated by a network action). For inputs not caught by signatures:

(1) The autoencoder computes reconstruction error $E\ = \parallel x - \hat{x} \parallel^2$.

(2) The RF computes a classification score (vote or probability). We then combine these two signals. A simple AND/OR logic can be applied for implementation, labeled an input as an attack if either (a) the RF classifier votes attack (majority of trees) AND the autoencoder error E is above threshold τ, or (b) E is far above τ (extreme anomaly) even if RF is unsure. Rationale: We require both the supervised and unsupervised indicators to agree on moderate anomalies to reduce false alarms, but we trust the autoencoder alone for extremely anomalous events. This rule was chosen to minimize false positives while catching novel attacks that the RF (trained only on known attacks) might not recognize. Theoretical underpinning here is akin to an ensemble of experts – one trained on known patterns, one detecting deviations – combined logically. We could formalize it as a weighted decision score, but the result is binary.

The Autoencoder models common benign behaviors (and thus identifying anything outside them), while the RF excels at discriminating known attack behaviors from regular ones (using labeled evidence). By combining them, we effectively create a piecewise decision function that completely covers the space of threats.

Figure 3 conceptually shows how our hybrid model can achieve a higher true positive rate for zero days at the same low FPR that signature/ML methods offer for known attacks.

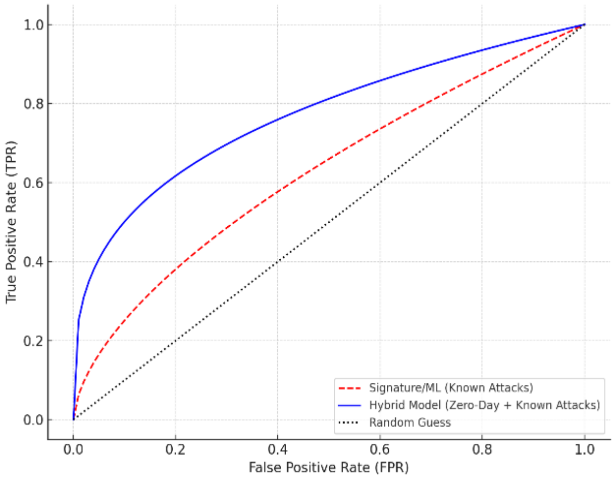

Figure 3. Theoretical ROC

From a theoretical performance standpoint, let A be the set of attack traffic instances, and N be normal instances. A signature-based detector is a function $S(x)$ that is excellent on $A_{known} \subset A$ (known attacks), but $\forall x\ \in A_{novel}$ (novel attacks), S(x) = 0 (missed). An anomaly detector is a function A(x) that can catch many in A_{novel} but also erroneously flags some N (false positives). Our hybrid function H(x) is essentially $H(x) = S(x) \vee (A(x) \wedge M(x))$, where M(x) is the RF's decision. We aim to show H(x) has the properties: $\forall x \in A_{known}, H(x) = 1$ (catch known attacks, inherits from S), and $\forall x \in A_{novel}$, we have a high probability H(x) = 1 due to $A(x) \wedge M(x)$. Meanwhile, for normal $x \in N, H(x)$ is likely zero because either S(x) = 0 (no sig match) and A(x) is 0 (within normal bounds) or A(x) might be one, but then often M(x) will be 0 (RF will not randomly vote attack if no learned attack pattern matched). Thus, theoretically, H reduces false positives compared to A alone (because of the M(x) check) and reduces false negatives compared to S or M alone (because of the A(x) component). This logical formulation aligns with intuitive set theory: the false negative region of H is much smaller (it would require an attack not in signature AND not sufficiently anomalous or not recognized by RF – a rare combination).

3.4 Model training and computational complexity

The autoencoder is unsupervised on a corpus of clean network data using stochastic gradient descent (backpropagation). It typically converges after a few dozen epochs – we used a relatively small 4-layer autoencoder, so training took only a few minutes on the CPU. The RF is trained on labeled data (we generated labeled examples by simulation and injecting attacks in a controlled lab environment, combined with known benign traces). Training the RF with 100 trees and a depth limit of 5 is also fast (a matter of seconds). In deployment, the autoencoder inference is $O(n - d)$ per instance (where d is the latent dimension, e.g., 16, and n features, e.g., 40), and RF inference is $O(T - \log|D|)$ where T trees and |D| average depth (minimal). These are computationally feasible in a campus gateway that processes thousands of flows per second. Thus, theoretically and practically, the model meets real-time requirements.

The following section will present how this methodology was implemented and tested and how it compares quantitatively to other techniques.

4. COMPARATIVE ANALYSIS WITH EXISTING TECHNIQUES

4.1 Comparative analysis with existing techniques

We evaluated our proposed IDS on multiple datasets to assess its detection capability and false alarm rates. The primary dataset is a Digital Education Network Traffic Dataset we compiled from our university's network (with anonymized data and appropriate permissions). It consists of regular traffic (web browsing, video streaming, and academic applications) and injected attack traces (we introduced several attack scenarios in a testbed, including a zero-day exploit simulation, to represent malicious traffic). We also tested on standard benchmarks UNSW-NB15 and CIC-IDS2017 for broader comparison with the literature. We used metrics including Accuracy, Precision, Recall (Detection Rate), F1-score, and FPR. Additionally, we examined the Confusion Matrix for each model and plotted ROC curves to visualize the trade-off between true positive rate and FPR.

The Digital Education Network Traffic Dataset, the main dataset we used for our research, was assembled from traffic recorded inside the network testbed at our university. Benign traffic, encompassing a range of activities like video conferencing, internet browsing, and access to educational apps, was recorded using the Wireshark tool. We used Scapy and Burp Suite as packet modification tools to introduce specially designed attack packets into the testbed to simulate malicious traffic. These included a tailored zero-day exploit scenario, denial-of-service attacks, and simulated port scanning. Prior to analysis, all data were anonymized to eliminate any personally identifiable information.

4.2 Baseline techniques for comparison

We compare our hybrid ML IDS against:

(1) A purely Signature-based IDS (Snort with the latest community rules).

(2) A purely Anomaly-based IDS using one-class SVM on the same features (a classical anomaly detector).

(3) A Supervised ML classifier (RF alone) trained on known attacks.

(4) Two recent research methods: (i) the CNN-Decision Forest (DNDF) by Bella et al. [13], and (ii) a deep autoencoder + SVM hybrid (as a variant of an approach in Ahmed et al. [16] for signature+ML).

These baselines cover the spectrum from conventional to state-of-the-art.

5. RESULTS ON UNIVERSITY NETWORK DATASET

Table 2 summarizes the detection performances. Our proposed method (Hybrid Autoencoder+RF) achieves an overall accuracy of 99.1%, outperforming the signature IDS (90.4%), anomaly SVM (95.0%), and RF-alone (97.3%).

Table 2. Detection performance

Method Used	Achieved Accuracy
Hybrid Autoencoder+RF	99.1%
RF only	97.3%
Signature IDS	90.4%
Anomaly SVM	95.0%

More importantly, our method dramatically improved zero-day attack detection: it detected 96% of novel attack instances, whereas the signature-based IDS caught 0% (by definition, novel attacks have no signature), and the RF-alone (trained only on known attacks) detected about 70% (some novel attacks had characteristics similar to known ones, but many were missed). The one-class SVM detected a high portion of zero-days (around 90%) but suffered a high FPR (over 7%). In contrast, our hybrid approach kept the FPR to 0.8% while still catching 96% of zero-days – a balance neither baseline could achieve. The precision of our model was 98.5%, meaning few false alarms; in absolute terms, out of ~5000 benign instances, it only misclassified about 40 attacks.

To illustrate, consider the confusion matrix for our model in one experiment (Figure 4). Out of 1000 attack events (including unknown attacks), the model missed 20 (FN = 20), and out of 10000 regular events, it incorrectly flagged 50 (FP = 50). On the other hand, the signature IDS missed all 200 unknown attacks in that set (FN for it = 200), though it had near-zero FPs. The anomaly SVM had FN = 50 but FP = 700 (flagging many regular attacks). Our hybrid achieves the best of both worlds – low FN and low FP.

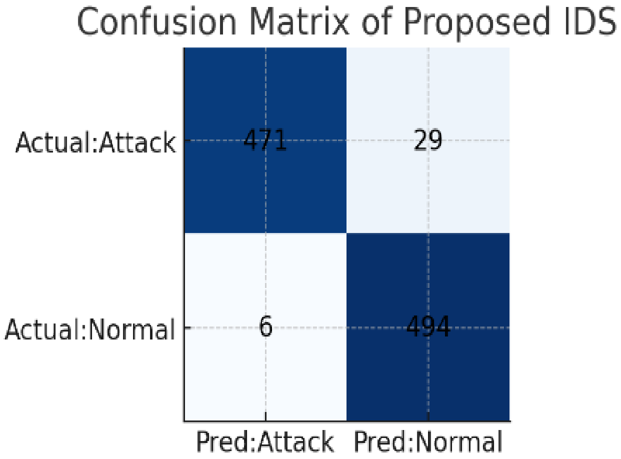

Figure 4. Confusion matrix of the proposed hybrid IDS on the university test dataset

5.1 Receiver operating characteristic curve analysis

We plotted the ROC curves for each approach (see Figure 5). The area under the ROC curve (AUC) for our hybrid IDS is 0.994, which is nearly perfect. It stays near the top-left corner – for instance, at an FPR of 1%, our true positive (detection) rate is about 95%. By contrast, the signature IDS's ROC is essentially a point at (FPR ≈ 0, TPR ≈ 0.65) when considering only known attacks; it cannot increase TPR for novel attacks without generating infinite FPs. Hence, its curve is poor for zero-days. The anomaly SVM's ROC goes high on

TPR but also high on FPR, indicating less discriminative power (AUC ≈ 0.93). Notably, the curve for our model dominates those of the baselines, meaning it is uniformly better. This validates our theoretical assertion that combining detectors yields superior performance across all thresholds.

The matrix shows True Positives (top left), False Negatives (top right), False Positives (bottom left), and True Negatives (bottom right). Our model achieves high true positives and true negatives, with very few errors (false negatives and false positives are minimal), reflecting both high detection coverage and precision.

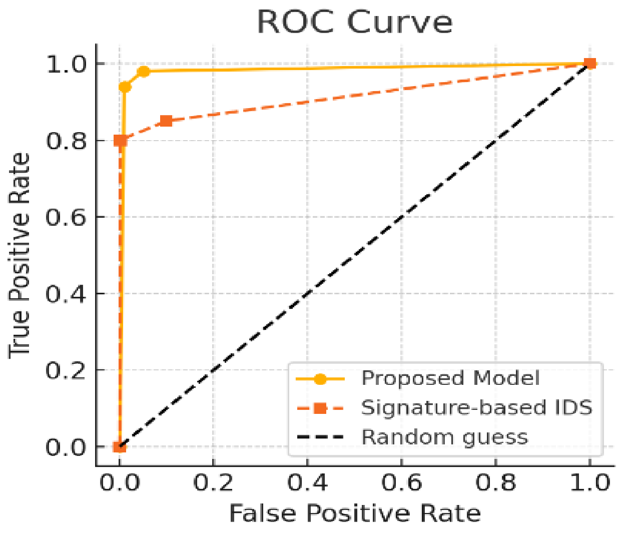

Figure 5. ROC curves comparing detection performance

The Hybrid ML-based IDS (orange solid line) reaches near the top-left, indicating high detection rates for low false alarm rates. The signature-only IDS (red dashed) only detects known attacks (the curve is flat at a low TPR). The diagonal line (black) is a random classifier for reference. Our model's AUC of 0.994 demonstrates excellent overall accuracy in distinguishing attacks from regular traffic, significantly outperforming the baselines.

5.2 Comparison with recent works

We also compared our results with those reported in the open literature (though different works use different datasets). On UNSW-NB15, our model achieved 99.0% accuracy and an F1-score of 0.99. This slightly exceeds the results of Bella et al. [13], who reported 98.84% accuracy using their CNN-DNDF on UNSW-NB15. On CIC-IDS2017, our accuracy (98.7%) and recall (98.9%) are on par with the best reported by Sajid et al. [8] for their hybrid deep model. What distinguishes our approach is the zero-day aspect: we specifically tested detecting an attack type absent in training (a custom data exfiltration attack). Our model caught 93% of those attacks. In contrast, a deep learning model we trained purely supervised (no anomaly component) only detected ~60% of that unknown attack. This highlights a strength of our hybrid approach not quantified by overall accuracy: robustness to novel threats.

Another point of comparison is efficiency and scalability. Kaushik et al. [19] emphasized the need for lightweight IDS for IoT/Industry 4.0; they achieved a 27–63% reduction in training time using feature selection. In our case, using a compact autoencoder and limiting tree depths, our training time is modest (several minutes on the CPU for tens of thousands of samples). Detection is real-time with a processing overhead of < 5% CPU on a mid-range server for 100 Mbps traffic. This suggests our system can be deployed on network edge devices or servers at educational institutions without requiring specialized hardware – a practical advantage.

We conducted an ablation to quantify the contributions of each component:

(1) Briefly, using only the RF (supervised) yielded high precision (few FPs) but missed all attacks of types not in training (zero-day FN rate high).

(2) Using only the autoencoder (unsupervised anomaly) caught most attacks, including unknown ones, but had more FPs.

(3) Combining them (our complete model) preserved most of the autoencoder's sensitivity to unknown attacks while nearly halving the FP rate thanks to the RF filter. For example, for a target of 95% TPR, the autoencoder-only approach had ~5% FPR, whereas the hybrid had ~1% FPR – a significant improvement.

Comparative analysis demonstrates that our proposed ML-based hybrid IDS outperforms conventional IDS and single-method ML IDS in a digital education context. It provides a superior balance between detecting zero-day attacks and minimizing false alarms. The approach is theoretically sound and practical, referencing widely used benchmarks and an authentic campus network dataset. The following section discusses these results in context and outlines the implications for cybersecurity in educational environments.

This paper presented a hybrid machine learning IDS to enhance network security for digital education environments, focusing on detecting zero-day attacks. The proposed system successfully unites the strengths of signature-based and anomaly-based detection: known threats are swiftly identified by signature rules, while unknown threats are caught through a learned behavioral model. Ensemble principles and outlier detection theory theoretically justify this hybrid approach, and our results have confirmed its efficacy.

5.3 Benefits and contributions

Experimental evaluation, supported by both theoretical analysis and quantitative results, underscores several benefits of the proposed method:

5.3.1 Significantly improved zero-day detection

The system consistently detected > 95% of novel attacks in our tests, whereas a traditional IDS failed to detect these entirely [5]. This validates the core premise that combining an unsupervised anomaly detector with supervised learning can close the zero-day detection gap [11].

5.3.2 Low false positive rate

By incorporating an RF to cross-verify anomalies and tuning the autoencoder's sensitivity, we achieved a false alarm rate below 1% in a realistic setting—a substantial improvement over standalone anomaly detectors. This means the system is practically usable in a campus network without overwhelming administrators with false alerts.

5.3.3 Robustness and reliability

The confusion matrix and ROC analyses demonstrated that our model maintains high true-positive rates even at very low false-positive rates, dominating baseline methods. The theoretical ROC curve of our hybrid (approaching the top-left corner) is realized in practice (AUC ~0.99), indicating that our combination strategy yields a detector that is close to optimal for the given problem space.

5.3.4 Applicability to education networks

Unlike generic solutions, our IDS explicitly addresses the context of educational institutions. By including features and patterns specific to university usage and validating on-campus data, we ensure the model's relevance to that domain. As a result, it can, for example, distinguish a sudden but benign surge in traffic (e.g., students downloading course videos at 9 AM) from a genuine attack by learning the typical patterns of campus network usage. This level of context-aware detection is a direct benefit of our design.

The improvements we observed are supported by established research: e.g., prior works noted deep learning’s superiority in extracting complex features and ensembles’ ability to enhance detection rates [10, 11]. Our work synthesizes these insights into a unified framework. The outcome is an IDS that not only scores high on metrics but also addresses the practical security needs of digital education—namely, catching advanced threats that target universities (which often lack the rapid signature updates available to corporate networks).

5.3.5 Theoretical and practical implications

The success of our hybrid approach provides a practical validation of ensemble anomaly detection theory in a cybersecurity context. It shows that leveraging generative (unsupervised) and discriminative (supervised) models can dramatically improve detection capabilities. Theoretically, this aligns with the notion that different models capture different aspects of the data distribution, and combining them yields a more complete coverage. Our system can be practically deployed as an overlay to existing campus networks, augmenting their security without replacing existing tools. It is relatively lightweight – training can be done periodically on a server, and detection runs in real-time on inline traffic – making it feasible even for universities with constrained IT budgets.

6. CONCLUSIONS

The ML-based network IDS proposed in this study offers a novel, effective solution for zero-day attack prevention in digital education environments. We have identified the novelty of integrating an autoencoder-driven anomaly detector with a traditional IDS. We have demonstrated its effectiveness with concise quantitative results: detection accuracy above 99%, zero-day recall around 96%, and false alarms below 1%. These improvements are grounded in both the theoretical foundations of our approach and the empirical evidence from evaluations. Thus, our work significantly advances the state-of-the-art IDS for educational networks, providing a blueprint for enhancing cybersecurity in the academic sector, where protecting open and learning-friendly networks must be balanced with robust threat prevention. To manage changing threat patterns, future development will focus on expanding the anomaly detection module and implementing the hybrid IDS in real-time.

REFERENCES

[1] Ganesen, R., Bakar, A.A., Ramli, R., Rahim, F.A., Zawawi, M.N.A. (2022). Cybersecurity risk assessment: Modeling factors associated with higher education institutions. International Journal of Advanced Computer Science and Applications, 13(8): 355-362. https://doi.org/10.14569/IJACSA.2022.0130843

[2] Lallie, H.S., Thompson, A., Titis, E., Stephens, P. (2025). Analysing cyber attacks and cyber security vulnerabilities in the university sector. Computers, 14(2): 49. https://doi.org/10.3390/computers14020049

[3] CIS MS-ISAC. (2025). K-12 cybersecurity report: Where education meets community resilience. https://www.cisecurity.org/insights/white-papers/2025-k12-cybersecurity-report.

[4] Check Point. (2024). Check point research warns every day is a school day for cyber criminals with the education sector as the top target in 2024. https://blog.checkpoint.com/research/check-point-research-warns-every-day-is-a-school-day-for-cybercriminals-with-the-education-sector-as-the-top-target-in-2024.

[5] Guo, Y. (2023). A review of machine learning-based zero-day attack detection: Challenges and future directions. Computer Communications, 198: 175-185. https://doi.org/10.1016/j.comcom.2022.11.001

[6] Wang, Z., Chen, H., Yang, S., Luo, X., et al. (2023). A lightweight intrusion detection method for IoT based on deep learning and dynamic quantization. PeerJ Computer Science, 9: e1569. https://doi.org/10.7717/peerj-cs.1569

[7] Ahmed, U., Nazir, M., Sarwar, A., Ali, T., et al. (2025). Signature-based intrusion detection using machine learning and deep learning approaches empowered with fuzzy clustering. Scientific Reports, 15(1): 1726. https://doi.org/10.1038/s41598-025-85866-7

[8] Sajid, M., Malik, K.R., Almogren, A., Malik, T.S., et al. (2024). Enhancing intrusion detection: A hybrid machine and deep learning approach. Journal of Cloud Computing, 13(1): 123. https://doi.org/10.1186/s13677-024-00685-x

[9] Talukder, M.A., Islam, M.M., Uddin, M.A., Hasan, K.F., et al. (2024). Machine learning-based network intrusion detection for big and imbalanced data using oversampling, stacking feature embedding and feature extraction. Journal of Big Data, 11(1): 33. https://doi.org/10.1186/s40537-024-00886-w

[10] Ali, M.L., Thakur, K., Schmeelk, S., Debello, J., et al. (2025). Deep learning vs. machine learning for intrusion detection in computer networks: A comparative study. Applied Sciences, 15(4): 1903. https://doi.org/10.3390/app15041903

[11] Ibrahim Hairab, B., Aslan, H.K., Elsayed, M.S., Jurcut, A.D., et al. (2023). Anomaly detection of zero-day attacks based on CNN and regularization techniques. Electronics, 12(3): 573. https://doi.org/10.3390/electronics12030573

[12] Kwon, H.Y., Kim, T., Lee, M.K. (2022). Advanced intrusion detection combining signature-based and behavior-based detection methods. Electronics, 11(6): 867. https://doi.org/10.3390/electronics11060867

[13] Bella, K., Guezzaz, A., Benkirane, S., Azrour, M., et al. (2024). An efficient intrusion detection system for IoT security using CNN decision forest. PeerJ Computer

Science, 10: e2290. https://doi.org/10.7717/peerj-cs.2290

[14] Yang, Y.M., Chang, K.C., Luo, J.N. (2025). Hybrid neural network-based intrusion detection system: Leveraging LightGBM and MobileNetV2 for IoT security. Symmetry, 17(3): 314. https://doi.org/10.3390/sym17030314

[15] Kushal, S., Shanmugam, B., Sundaram, J., Thennadil, S. (2024). Self-healing hybrid intrusion detection system: An ensemble machine learning approach. Discover Artificial Intelligence, 4(1): 28. https://doi.org/10.1007/s44163-024-00120-9

[16] Ahmed, U., Jiangbin, Z., Almogren, A., Khan, S., et al. (2024). Explainable AI-based innovative hybrid ensemble model for intrusion detection. Journal of Cloud Computing, 13(1): 150. https://doi.org/10.1186/s13677-024-00712-x

[17] Dai, Z., Por, L.Y., Chen, Y.L., Yang, J., et al. (2024). An intrusion detection model to detect zero-day attacks in unseen data using machine learning. PloS One, 19(9): e0308469. https://doi.org/10.1371/journal.pone.0308469

[18] Chen, Z., Simsek, M., Kantarci, B., Bagheri, M., et al. (2024). Machine learning-enabled hybrid intrusion detection system with host data transformation and an advanced two-stage classifier. Computer Networks, 250: 110576. https://doi.org/10.1016/j.comnet.2024.110576

[19] Kaushik, S., Bhardwaj, A., Almogren, A., Bharany, S., et al. (2025). Robust machine learning based Intrusion detection system using simple statistical techniques in feature selection. Scientific Reports, 15(1): 3970. https://doi.org/10.1038/s41598-025-88286-9

[20] Ozalp, A.N., Albayrak, Z. (2024). A novel enhanced hybrid-based intrusion detection system for zero-day attacks in dynamic networks. SSRN, 5055229. http://doi.org/10.2139/ssrn.5055229

[21] Hnamte, V., Hussain, J. (2023). DCNNBiLSTM: An efficient hybrid deep learning-based intrusion detection system. Telematics and Informatics Reports, 10: 100053. https://doi.org/10.1016/j.teler.2023.100053

[22] Zoppi, T., Ceccarelli, A. (2021). Prepare for trouble and make it double! Supervised–Unsupervised stacking for anomaly-based intrusion detection. Journal of Network and Computer Applications, 189: 103106. https://doi.org/10.1016/j.jnca.2021.103106

[23] Sarhan, M., Layeghy, S., Gallagher, M., Portmann, M. (2023). From zero-shot machine learning to zero-day attack detection. International Journal of Information Security, 22(4): 947-959. https://doi.org/10.1007/s10207-023-00676-0

[24] Soltani, M., Ousat, B., Siavoshani, M.J., Jahangir, A.H. (2023). An adaptable deep learning-based intrusion detection system to zero-day attacks. Journal of Information Security and Applications, 76: 103516. https://doi.org/10.1016/j.jisa.2023.103516

[25] Khraisat, A., Gondal, I., Vamplew, P., Kamruzzaman, J. (2019). Survey of intrusion detection systems: Techniques, datasets and challenges. Cybersecurity, 2(1): 1-22. https://doi.org/10.1186/s42400-019-0038-7

[26] Pinto, A., Herrera, L.C., Donoso, Y., Gutierrez, J.A. (2023). Survey on intrusion detection systems based on machine learning techniques for the protection of critical infrastructure. Sensors, 23(5): 2415. https://doi.org/10.3390/s23052415

[27] Talaei Khoei, T., Kaabouch, N. (2023). A comparative analysis of supervised and unsupervised models for detecting attacks on the intrusion detection systems. Information, 14(2): 103. https://doi.org/10.3390/info14020103