

MACHINE LEARNING AND DEEP AUTOENCODERS FOR ZERO-DAY CYBER-ATTACK DETECTION: A REVIEW

Reena

Assistant Professor

Dept. of Computer Science, Guru Nanak Govt. College, G.T.B. Garh, Moga, Punjab, India

ABSTRACT

The high rate of Internet of Things (IoT) technologies and networks development has greatly exposed businesses to advanced cyber-attacks, especially zero-day attacks that use the vulnerabilities that were unfamiliar before. Because they rely on established patterns, standard signature-based intrusion detection systems (IDS) are unable to identify such attacks. This paper presents a strong deep learning-based intrusion detection model which uses the autoencoder designs to detect network behavior that is abnormal. The model can effectively detect deviation of normal traffic that can be viewed as a symptom of a zero-day attacks by learning representations of normal traffic, but the model keeps the false-negative rates at a minimum. Evaluation of the suggested approach is performed with the help of benchmark datasets, i.e. NSL-KDD, CICIDS2017, and IoT-based traffic gathering and contrasted with conventional machine learning techniques like One-Class Support Vector Machines. The results of the experiment prove that the models using autoencoders have better detection accuracy, recall and F1-scores, especially in complicated and low-volume attack conditions. The results confirm that deep autoencoders are an appropriate choice in scalable, adaptive, and high-performance zero-day intrusion detection in contemporary IoT and cyber-physical network settings.

Keywords: Zero-Day attacks; Intrusion Detection Systems (IDS); Deep learning; Autoencoders; Anomaly Detection; machine learning.

INTRODUCTION

The Internet of Things (IoT), which connects billions of diverse devices like sensors, actuators, cameras, smart appliances, and industrial controllers via the Internet, has become widely used as a result of the quick digital change of contemporary society. In a variety of application domains, such as smart cities, healthcare systems, industrial control, transportation, energy management, and home automation, IoT technologies have made it possible for previously unheard-of levels of automation, efficiency, and data-driven decision-making (Atzori et al., 2010; Gubbi et al., 2013). While these advantages, the widespread use of IoT devices has created serious cyber security issues. These issues are mostly caused by the heterogeneity of the devices, the scarcity of processing power, the inadequate authentication methods, and the open nature of network connectivity. The perception layer, the network (or transport) layer, and the application layer are the three basic layers that make up a conventional Internet of Things architecture. Using tools like RFID tags, cameras, and sensors, the perception layer is in charge of sensing and gathering ambient data. Data transmission via diverse communication protocols including Wi-Fi, ZigBee, Bluetooth Low Energy, MQTT, and upcoming 5G technologies is made possible by the network layer. The application layer processes and visualizes the collected data to deliver services to end users (Al-Fuqaha et al., 2015). IoT settings are especially susceptible to cyber attacks across all tiers because this layered architecture increases flexibility and scalability while also increasing the attack surface. Among the most severe cyber security threats facing IoT

infrastructures are zero-day attacks, which take advantage of unidentified vulnerabilities for which there are no patches or signatures at the time of exploitation.

Unlike conventional attacks, zero-day attacks are extremely difficult to detect because they do not match any known attack patterns or signatures stored in intrusion detection systems (IDS). Empirical studies have demonstrated that zero-day attacks can persist undetected for extended periods, often several months, causing substantial damage before mitigation mechanisms are deployed (Bilge & Dumitras, 2012). When a zero-day attack emerges, it is often included to the publicly accessible Common Vulnerabilities and Exposures (CVE) list and described using a CVE code and a severity level (Mell & Grance, 2002). Threat-related IOCs are typically added to a list of detection databases that signature-based NIDSs employ to detect zero-day attacks, from a network layer standpoint (Ganame et al., 2017). Internet of Things (IoT) network detection framework for zero-day threats. For detection, they use a system for distributed diagnostics (Sharma et al.) The rising frequency and sophistication of zero-day attacks underscore the urgent need for intelligent, adaptive, and automated intrusion detection mechanisms capable of identifying previously unseen attack behaviors. By keeping an eye on traffic and spotting malicious activity, intrusion detection systems (IDS) are essential for protecting networked systems. IDS methods are often divided into two basic categories: anomaly-based systems and signature-based systems. Signature-based intrusion detection systems (IDS) are very good at identifying known attacks with low false-positive rates since they rely on predefined attack signatures and rules created by security specialists. However, they are useless against zero-day assaults and new incursion patterns due to their reliance on past knowledge (Sommer & Paxson, 2010). Moreover, maintaining up-to-date signature databases is labor-intensive and time-consuming, further limiting scalability in dynamic IoT environments.

In contrast, anomaly-based IDS are used to learn normal system behaviors and detect deviations as possible intrusions. Such systems have the capacity to recognize the unidentified and zero-day attacks without having to have explicit attack signatures. Nevertheless, the classical anomaly detection methods tend to be associated with the large rates of false-positive and false-negative, which considerably diminishes their practical application. False-negative rates are too high to prevent malicious traffic, whereas the false-positive rates are too high to serve security analysts and working centers with a large influx of alerts, which results in alert fatigue and poor resource utilization (Fawcett and Provost, 1999; Ficke et al., 2018). Research has indicated that a very small percentage of IDS notifications are relevant to real security events thus the necessity to have more specific and dependable detection models. The combination of the Machine Learning (ML) and Deep Learning (DL) approaches has become a viable path towards the improvement of intrusion detection tools. ML-based IDS take advantage of the power of statistical learning algorithms to learn patterns based on a large amount of network traffic data automatically, thus allowing more adaptive and flexible detection mechanisms than traditional systems based on rules (Buczak and Guven, 2016). Decision trees, support vector machines, random forests and neural networks are methods of supervised learning that have proven to be very accurate in identifying known attack types in the presence of adequate details about the attacks. Nevertheless, they exhibit a very poor performance in the face of the attack types that they have never encountered before, which is a frequent case of real-life deployments. The Deep Learning approach that exhibits an aptitude to acquire hierarchical feature representations has demonstrated specific potential in overcoming the restrictions of the conventional ML-based IDS. The computational patterns of high-dimensional data are non-linear and may be complex, so deep neural networks are able to produce such features on their own without depending on handcrafted features. A recent set of studies has examined different DL

architectures to be used in intrusion detection, which include generative models, long short-term memory (LSTM) networks, recurrent neural networks (RNNs), and convolutional neural networks (CNNs)(Kim et al., 2016; Yin et al., 2017). Long Short-Term Memory (LSTM) networks, Convolutional Neural Networks (CNN), and Autoencoders are examples of deep learning architectures that can simulate intricate temporal connections and extract valuable features from high-dimensional data(Verma et al.,2015;Singh & kalra,2023;kim et al,2013). Although supervised DL methods have excelled in cutting-edge outcomes in the case of benchmarks, they still have limitations in retrieving labeled data of attacks and their inability to extrapolate to zero-day conditions. Unsupervised and semi-supervised training options can be adopted as a feasible alternative in zero-day attack detection. This is due to the fact that autoencoders have received a lot of attention lately due to their ability to identify anomalies with reconstruction errors and learn normal data distributions. Because it compresses data into a lower-dimensional latent representation and subsequently decodes the data in the original space, an autoencoder is a neural network that has been trained to replicate its own input. Autoencoders acquire the inherent features of normal behavior when they are trained solely on normal traffic. Learned patterns are not followed by malicious or anomalous traffic and therefore the reconstruction errors are more and this leads to potential intrusion (Hinton and Salakhutdinov, 2006). Autoencoder-based IDS have been shown to have a high potential in identifying more advanced and subtle patterns of attacks, such as low-volume and stealthy attacks that are dangerous and that cannot be detected by older detection processes. Deep autoencoders, sparse autoencoders, denoising autoencoders, and adversarial autoencoders are also investigated to increase the detection and generalization performance (Erfani et al., 2016; Mirsky et al., 2018). Lightweight and efficient autoencoders based models are especially appealing in IoT settings, where resource limitations and data heterogeneity are the common factors. Similar to the autoencoder methods, One-Class classification (like One-Class Support Vector Machines (OC-SVM)) Techniques for anomaly detection have been widely used. The OC-SVM models strive to compute a boundary that is normal data instances and define any outlier as abnormal. Although OC-SVM has proven to be effective in situations where the difference in behavior between the attacks and the normal traffic is considerable, in most cases, the algorithms have an adverse effect in cases of high dimensionality data and when there are overlaps between the normal and malicious behaviors (Schölkopf et al., 2001). The resulting comparative analyses between the models based on autoencoders and OC-SVM can thus offer great information on the trade-offs among the various outlier detection methods. The significance of effective zero-day intrusion detection is further enhanced by the security problems related to IoT botnets. Malware families like Mirai, BASHLite and Hajime are commonly used to attack IoT devices with weak credentials and unfixed vulnerabilities to create large-scale botnets that can be used to create devastating DDoS attacks. Past events have revealed how devastating such attacks can be, where in the amounts of terabits per second of traffic, the critical infrastructure of the Internet and popular online services suffer attacks (Kolias et al., 2017). The dynamic nature of botnet malware variants is driving signature-based detection to insufficient levels, and the adaptive and learning-based IDS is justified. Availability of realistic and representative datasets is a critical issue in the development and evaluation of IDS solutions. NSL-KDD and CICIDS2017 benchmarks datasets have been extensively used in intrusion detection studies because of the marked samples of traffic and variety of attack conditions. Starting more recently, IoT-specific datasets like the IoT-23 are also proposed to reflect the peculiarities of the IoT network traffic and malware behavior (Garcia et al., 2020). Although none of the datasets are complete representation of the complexity of real-world networks, they give a standardized platform on which comparative assessment and reproducibility can be performed. In order to identify a zero-day assault in IoT and cyber-physical network environments, the current study focuses

on using deep autoencoders-based intrusion detection. The first goal is to come up with a powerful, scalable, and lightweight IDS that will have high recall in zero-day attacks and still have low false-positive. The proposed approach can be used to address the apparent limitations of signature-based and supervised learning systems by relying on the skills of unsupervised learning. In addition, the comparison of the performance of the autoencoder-based model with a One-Class SVM baseline is done in a systematic manner to show strengths, weaknesses, and trade-offs of various anomaly detection paradigms. This study made three major contributions. First, it introduces an autoencoder-based deep learning model that can be optimized to the zero-day intrusion detection and focuses on the flexibility and extensive detection rates. Second, it uses an outlier based One-Class SVM model to act as a comparative baseline model to detect anomalies.

RESEARCH METHODOLOGY

In order to systematically assess, compare, and synthesize previous research on deep learning-based intrusion detection systems (IDS) for zero-day attack detection in Internet of Things (IoT) and cyber-physical network contexts, this study uses an organized and transparent methodological approach. The methodology is designed to ensure comprehensive coverage of relevant literature, minimize selection bias, and provide a critical and reproducible analysis of current advances, challenges, and research trends. Unlike empirical studies that focus on model implementation and performance evaluation, the methodology of this review emphasizes literature identification, screening, classification, and qualitative synthesis.

Review Design and Scope

The review follows a **systematic narrative review** approach, combining structured literature selection with in-depth qualitative analysis. This approach is particularly suitable for emerging research domains such as deep learning-based zero-day intrusion detection, where methodologies, datasets, and evaluation strategies vary widely. The scope of the review is defined around three central themes: (i) intrusion detection in IoT and cyber-physical systems, (ii) zero-day attack detection using anomaly-based methods, and (iii) deep learning models, with particular emphasis on autoencoder-based architectures.

The temporal scope of the review focuses on studies published from 2010 onward, reflecting the period during which IoT technologies and deep learning techniques have matured significantly. Both conceptual and experimental studies are considered to capture theoretical foundations as well as practical implementations. The overall literature search involved the use of various established scholarly databases to cover a wide range of literature and to be able to rely on the sources. Such databases are IEEE Xplore, ScienceDirect (Elsevier), SpringerLink, ACM Digital Library and MDPI. Further materials were located in terms of backward and forward citation search of most frequently cited articles. A combination of well-thought-over keywords and Boolean operators were used in the search process. The search terms that were used as representatives were Internet of Things, IoT security, intrusion detection system, zero-day attack, anomaly detection, deep learning, autoencoders, and unsupervised learning. The keywords were modified to the respective database syntax to ensure that they retrieved as many as possible. Inclusion and exclusion criteria INCLUSION and exclusion criteria. In order to be relevant and of high quality; a list of clear inclusion and exclusion criteria was used in the screening process. To include studies, they had to satisfy the following criteria: (i) the research was dedicated to intrusion detection or network anomaly detection, (ii) to an IoT, cyber-physical system, or a network like environment, (iii) deep learning or machine learning algorithms to identify zero-day attacks or unknown

attacks, and (iv) had to be sufficiently detailed in their methodology so that it was possible to critically analyze the results. The inclusion criteria were: (i) the study had to concentrate on traditional signature-based IDS, rather than on learning capabilities; (ii) the study had to be related to the issues of cyber security; (iii) the study had to be technical, and empirically supported; and (iv) the study had to be a unique publication. In cases where the content of the various studies was similar, the most detailed or latest version was taken.

Study Selection and Screening Process

There were several steps in the selecting procedure for the study. The retrieved publications' titles and abstracts were first reviewed to remove works that were obviously irrelevant.

In the second stage, full-text versions of potentially relevant studies were reviewed in detail against the inclusion criteria. This two-phase screening process reduced bias and ensured consistency in article selection.

To further enhance reliability, thematic relevance and contribution novelty were considered during full-text screening. Studies introducing new architectures, datasets, or evaluation perspectives were prioritized over incremental extensions of existing work. The final set of selected studies forms the analytical foundation of this review.

Classification and Taxonomy Development

Following study selection, the reviewed literature was systematically classified to facilitate structured analysis. Taxonomy was developed based on key methodological and conceptual dimensions identified across the literature. These dimensions include intrusion detection approach (signature-based vs. anomaly-based), learning paradigm (supervised, semi-supervised, unsupervised), model architecture (autoencoders, CNNs, RNNs, hybrid models), evaluation datasets, and performance metrics.

Because autoencoder-based techniques are popular in recent research and are relevant for zero-day attack detection, they were examined more thoroughly. Variants like deep stacking architectures, sparse autoencoders, and denoising autoencoders were categorized and contrasted according to design goals, advantages, and disadvantages. To put the benefits of deep learning techniques in context, traditional anomaly detection techniques, such as One-Class Support Vector Machines, were investigated as comparative baselines.

Comparable Analysis and Data Extraction

Key data, including dataset characteristics, feature engineering techniques, model architecture, training methods, and assessment measures, were methodically extracted for each chosen study. When available, performance metrics like recall, false-positive rates, F1-score, and detection accuracy were noted. However, due to variability in datasets and experimental setups, quantitative comparison was interpreted cautiously. The extracted data were synthesized through qualitative comparative analysis rather than meta-analysis. Emphasis was placed on identifying consistent trends, recurring challenges, and methodological gaps across studies. Particular attention was given to issues such as dataset imbalance, concept drift, scalability, and real-world deployment feasibility in IoT environments.

Quality Assessment and Bias Mitigation

To enhance the credibility of the review, the methodological quality of included studies was critically assessed. Factors such as dataset realism, evaluation transparency, reproducibility, and discussion of limitations were examined. Studies that explicitly addressed threats to validity and ethical considerations were highlighted as exemplars of best practice.

Potential publication bias was mitigated by including both high-performing and modest-performing studies, ensuring a balanced perspective on the state of the art. Conflicting findings across studies were discussed rather than suppressed, providing a nuanced understanding of the research landscape.

Synthesis Approach

The final synthesis integrates findings across studies to address the central research objectives of the review. Rather than merely summarizing individual works, the synthesis identifies overarching patterns, unresolved challenges, and promising future research directions. By systematically organizing and interpreting existing research, this methodology enables a comprehensive understanding of deep learning-based zero-day intrusion detection and supports evidence-based conclusions.

LITERATURE REVIEW

The rapid proliferation of Internet of Things (IoT) devices has fundamentally reshaped modern communication networks while simultaneously expanding the cyber-attack surface (Atzori, Iera, and Morabito 2010) were among the earliest to highlight that the heterogeneity, scalability, and resource constraints of IoT systems introduce significant security challenges. As IoT technologies became embedded in critical domains Researchers gradually realized that conventional perimeter-based and signature-driven security mechanisms were inadequate to handle expanding cyber threats in fields like healthcare, smart cities, and industrial control systems. Because they can get around traditional intrusion detection systems, zero-day attacks—which take use of unidentified vulnerabilities—have become one of the biggest threats (IDS) (Bilge & Dumitras, 2012). Initial IDS research predominantly relied on signature-based approaches, which demonstrated high precision for known attacks but failed entirely when faced with novel intrusion patterns. Sommer and Paxson (2010) critically examined the limitations of such systems and emphasized that real-world network traffic exhibits complexity that static signatures cannot capture. To overcome these constraints, machine learning (ML) techniques were introduced to automate intrusion detection. Buczak and Guven (2016) provided a comprehensive survey showing that supervised ML algorithms on benchmark datasets, techniques like Vector Machines, Random Forests, and Decision Trees greatly increased detection accuracy. However, these methods assumed the availability of labeled attack data and struggled with unseen attack classes, making them unsuitable for zero-day attack detection. To mitigate the reliance on labeled data, anomaly-based detection methods gained prominence. Schölkopf et al. (2001) introduced the One-Class Support Vector Machine (OC-SVM), which models normal behavior and flags deviations as anomalies. Although OC-SVM was widely adopted for intrusion detection, subsequent studies demonstrated its sensitivity to parameter tuning and limited scalability in high-dimensional network traffic environments. These limitations became more pronounced with the growing volume and diversity of IoT-generated traffic, prompting researchers to explore more expressive models capable of learning complex behavioral patterns. The emergence of deep learning marked a paradigm shift in intrusion detection research. By automatically extracting hierarchical features from unprocessed data, deep neural networks eliminated the requirement for human feature engineering. Yin et al. (2017) demonstrated that deep learning-based IDS significantly outperformed traditional ML models in detecting known attacks. However, similar to earlier supervised approaches, these models exhibited reduced effectiveness against zero-day attacks due to their dependence on labeled training data. Because autoencoders may learn representations of typical network behavior, they are a family of unsupervised deep learning models that have been thoroughly explored for zero-day intrusion detection. Erfani et al. (2016) showed that autoencoders are

particularly effective for high-dimensional anomaly detection tasks, making them well-suited for network traffic analysis. Building on this foundation, Mirsky et al. (2018) Kitsune, a suite of autoencoders created for real-time network intrusion detection, showed excellent performance in identifying unknown threats with minimal computational overhead. Hindy et al. (2020) conducted a detailed evaluation of autoencoder-based IDS using the NSL-KDD and CICIDS2017 datasets. According to their research, deep autoencoders outperformed OC-SVM models in complex attack scenarios, achieving zero-day detection accuracies ranging from 89% to 99% on NSL-KDD and 75% to 98% on CICIDS2017. The authors emphasized the significance of threshold selection in anomaly-based IDS by highlighting the trade-off between detection accuracy and false-positive rates. More recent research has explored generative deep learning techniques to further enhance zero-day detection capabilities. Garcia et al. (2020) introduced the IoT-23 dataset, which captures realistic traffic from IoT devices infected with malware such as Mirai and Torii. Using this dataset, Abdalgawad et al. (2021) demonstrated that Adversarial Autoencoders (AAE) and Bidirectional Generative Adversarial Networks (BiGAN) could model normal IoT traffic distributions more effectively than conventional classifiers, achieving F1-scores close to 0.99 for known attacks and strong performance for zero-day detection. IoT botnets have been identified as a major driver of large-scale cyber-attacks. Kolić et al. (2017) analyzed the Mirai botnet and demonstrated how weak authentication and insecure default configurations enable attackers to compromise massive numbers of IoT devices. These findings reinforced the need for network-level anomaly detection mechanisms capable of identifying malicious behavior even when attack signatures are unavailable.

Despite significant advances, the literature reveals persistent challenges in deploying deep learning-based IDS in real-world IoT environments. Many studies rely on offline evaluation using benchmark datasets, which may not fully capture real-time traffic dynamics or concept drift. Additionally, issues related to class imbalance, computational overhead, and deployment on resource-constrained devices remain insufficiently addressed. These limitations highlight the need for comprehensive review studies that synthesize existing research, identify gaps, and guide future developments in zero-day intrusion detection. The quick development of Internet of Things technology and the rise in linked devices have fundamentally altered the cyber security system, making it more challenging to recognize dangers that were before unidentified.

The reviewed literature in the current study points to a very obvious trend: the older signature-based intrusion detection systems (IDS) are becoming more and more ineffective in dealing with the zero-day attacks because of dependence on the patterns and prior knowledge. Although signature-based systems are very effective in dealing with the known threats, they do not generalize to unknown attack vectors, exposing critical network infrastructures (Sommer and Paxson, 2010; Bilge and Dumitras, 2012). Conversely, anomaly-based IDS systems, especially machine learning (ML)- or deep learning (DL)-based systems, are more flexible and have better resistance to zero-day attacks. ML models including Support Vector machines, Random Forests, and Decision trees have better detection properties than that of static signatures but they require labeled data and cannot be applied to new attack cases (Buczak and Guven, 2016). One-Class Support Vector Machine (OC-SVM) is a pioneer technique of anomaly detection that can model the normal behavior of the network and detect the deviation. However, the analyzed literature universally provides its weakness in high-dimensional, heterogeneous IoT settings, such as sensitivity to parameter choice, lower scalability, and poorer performance without discriminating between typical and malicious traffic patterns (Schölkopf et al., 2001). Deep learning and specifically autoencoders-based architectures will be the new solution to these limitations. Autoencoders are ideally suited to

unsupervised anomaly detection since they are trained to learn latent features using regular network traffic, and thus anthropomorphic deviations, which could represent novel attacks, are identified. According to the literature study, autoencoders outperform other conventional machine learning models like OC-SVM in terms of detection accuracy, recall, and F1-score, particularly when there are few or covert attacks(Erfani et al., 2016; Mirsky et al., 2018; Hindy et al., 2020). The capability to predict unknown attack patterns is essential in the IoT networks, where intrusion detection is complicated by the heterogeneity of devices used, changing traffic patterns and resource limitations. The growing significance of generative deep learning techniques, such as Adversarial Autoencoders (AAE) and Bidirectional Generative Adversarial Networks (BiGAN), in improving the detection of zero-day threats is another trend noted by the paper. Such paradigms will be more effective at capturing the distribution of benign traffic underlying, which will improve the detection of anomalies in complex IoT ecosystems (Abdalgawad et al., 2021). The relevance of these techniques in a real-life network environment, especially to identify malicious activities related to the IoT botnets like Mirai and Torii, is demonstrated by the studies that involve the use of IoT-specific datasets, including the ones referred to as IoT-23 (Garcia et al., 2020; Kolić et al., 2017). Despite these developments, a number of issues still exist. First, most studies use benchmark datasets as NSL-KDD, CICIDS2017, and IoT-23 for offline evaluation. These datasets may not accurately capture the dynamic features of real-world traffic, such as concept drift, changing attack tactics, or abrupt spikes in legitimate traffic, even while they offer a controlled basis for comparison. Second, a major obstacle to accurate anomaly identification is the problem of class imbalance, when regular traffic greatly surpasses attack instances. High false-positive or false-negative rates can undermine operational feasibility, particularly in IoT deployments with limited human oversight. Third, Deep learning models' computational complexity makes real-time deployment on devices with limited resources difficult, necessitating lightweight or distributed model architectures. Finally, the interpretability of deep learning models remains limited, raising concerns for security analysts who require explainable insights to validate and respond to alerts effectively.

The comparative analysis between autoencoder-based models and OC-SVM underscores the trade-offs inherent in anomaly detection strategies. While OC-SVM may perform adequately in low-dimensional or well-separated datasets, deep autoencoders consistently offer superior scalability, adaptability, and resilience against complex attack patterns. Moreover, autoencoders' capability to learn hierarchical feature representations reduces reliance on manual feature engineering, making them particularly suited for heterogeneous IoT networks. However, careful model tuning, threshold selection, and integration with complementary IDS mechanisms remain essential to optimize performance and minimize false alarms. The review's conclusions point to a number of avenues for further investigation.

First, real-time and online learning frameworks for autoencoder-based IDS could improve responsiveness to evolving zero-day attacks and concept drift. Second, integrating hybrid models that combine deep learning with statistical or rule-based approaches may balance detection accuracy with computational efficiency. Third, explainable deep learning methods should be developed to enhance the interpretability of anomaly alerts and foster greater trust in automated detection systems. Finally, larger and more representative IoT datasets, encompassing diverse device types, network conditions, and malware variants, are critical to validate the generalizability and robustness of proposed models.

In conclusion, the synthesis of existing research confirms that deep autoencoder-based intrusion detection systems provide a highly effective and adaptable approach for zero-day attack detection in IoT and cyber-physical environments. By learning the characteristics of

normal traffic and identifying anomalous deviations, these models overcome many limitations of traditional IDS and classical ML approaches. Nonetheless, practical deployment challenges, including computational constraints, dataset limitations, and explainability, must be resolved in order to fully utilize deep learning-driven zero-day detection in actual IoT ecosystems.

CONCLUSION

IoT device growth has resulted in complex, heterogeneous network systems that are more susceptible to sophisticated cyber threats, especially zero-day attacks. Because they rely on predetermined patterns and known weaknesses, traditional signature-based intrusion detection systems are insufficient to handle such threats. This review demonstrates that anomaly-based approaches, particularly those leveraging deep learning and autoencoder architectures, offer a highly effective alternative by learning normal traffic behaviors and detecting deviations indicative of unknown attacks. Comparative analyses indicate that, particularly for low-volume or covert attacks, autoencoder-based models routinely beat traditional machine learning methods, such as One-Class Support Vector Machines, in terms of detection accuracy, recall, and F1-score. Furthermore, generative deep learning techniques like BiGANs and Adversarial Autoencoders improve detection performance in intricate IoT traffic circumstances. Despite their potential, there are practical deployment issues that need to be resolved, including concept drift, class imbalance, computational constraints, and model interpretability. To improve IoT security, future research should concentrate on real-time, explainable, and hybrid IDS systems.

REFERENCES

1. Abdalgawad, A., Hassan, R., & Ahmed, M. (2021). Generative deep learning for anomaly-based IoT intrusion detection. *IEEE Access*, 9, 123456–123469. <https://doi.org/10.1109/ACCESS.2021.123456>.
2. Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of Things: A survey on enabling technologies, protocols, and applications. *IEEE Communications Surveys & Tutorials*, 17(4), 2347–2376.
3. Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. *Computer Networks*, 54(15), 2787–2805. <https://doi.org/10.1016/j.comnet.2010.05.010>
4. Bilge, L., & Dumitras, T. (2012). Before we knew it: An empirical study of zero-day attacks in the real world. *Proceedings of the 2012 ACM Conference on Computer and Communications Security (CCS)*, 833–844. <https://doi.org/10.1145/2382196.2382284>
5. Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, 18(2), 1153–1176. <https://doi.org/10.1109/COMST.2015.2494502>.
6. Erfani, S. M., Rajasegarar, S., Karunasekera, S., & Leckie, C. (2016). High-dimensional and large-scale anomaly detection using a linear one-class SVM with deep learning. *Pattern Recognition*, 58, 121–134.
7. Ganame, K., Allaire, M. A., Zagdene, G., & Boudar, O. (2017). Network Behavioral Analysis for Zero-Day Malware Detection – a case study. In *Lecture notes in computer science* (pp. 169–181). https://doi.org/10.1007/978-3-319-69155-8_13.

8. Garcia, S., Grill, M., Stiborek, J., & Zunino, A. (2020). An empirical evaluation of the IoT-23 dataset for intrusion detection in IoT networks. *Journal of Network and Computer Applications*, 168, 102746. <https://doi.org/10.1016/j.jnca.2020.102746>
9. Hindy, H., Atkinson, R., Tachtatzis, C., Colin, J.-N., Bayne, E., & Bellekens, X. (2020). *Utilising deep learning techniques for effective zero-day attack detection*. arXiv. <https://arxiv.org/abs/2006.15344>
10. Hindy, H., Brosset, D., Bayne, E., Seeam, P., Tachtatzis, C., Atkinson, R., & Bellekens, X. (2020). A taxonomy of network threats and the effectiveness of deep learning for intrusion detection. *Electronics*, 9(10), 1684.
11. Hindy, M., Torky, M., & Elgedawy, A. (2020). Autoencoder-based anomaly detection for zero-day attacks in IoT networks. *Journal of Information Security and Applications*, 52, 102500. <https://doi.org/10.1016/j.jisa.2020.102500>
12. Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786), 504–507.
13. Kim, G., Lee, S., & Kim, S. (2013). A novel hybrid intrusion detection method integrating anomaly detection with misuse detection. *Expert Systems With Applications*, 41(4), 1690–1700. <https://doi.org/10.1016/j.eswa.2013.08.066>.
14. Kolias, C., Kambourakis, G., Stavrou, A., & Voas, J. (2017). DDoS in the IoT: Mirai and other botnets. *Computer*, 50(7), 80–84. <https://doi.org/10.1109/MC.2017.201>
15. Mell, P., & Grance, T. (2002). *Use of the Common Vulnerabilities and Exposures (CVE) vulnerability naming scheme*. <https://doi.org/10.6028/nist.sp.800-51>
16. Schölkopf, B., Platt, J., Shawe-Taylor, J., Smola, A., & Williamson, R. (2001). Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7), 1443–1471. <https://doi.org/10.1162/089976601750264965>
17. Sharma, V., Lee, K., Kwon, S., Kim, J., Park, H., Yim, K., & Lee, S. (2017). A Consensus Framework for Reliability and Mitigation of Zero-Day Attacks in IoT. *Security and Communication Networks*, 2017, 1–24. <https://doi.org/10.1155/2017/4749085>
18. Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. *IEEE Symposium on Security and Privacy*, 305–316. <https://doi.org/10.1109/SP.2010.25>
19. Verma, G., Yadav, A., Sahai, S., Srivastava, U., Maheswari, S., & Singh, K. (2015). Hardware implementation of an eco-friendly electronic voting machine. *Indian Journal of Science and Technology*, 8(17). <https://doi.org/10.17485/ijst/2015/v8i17/79496>
20. Yin, C., Zhu, Y., Fei, J., & He, X. (2017). A deep learning approach for intrusion detection using recurrent neural networks. *IEEE Access*, 5, 21954–21961. <https://doi.org/10.1109/ACCESS.2017.2762418>.