2026 Volume: 02 Issue: 05 www.ijrpa.com ISSN 2456-9995 Review Article

International Journal Research Publication Analysis

Page: 01-14

AUTOENCODER-BASED ANOMALY DETECTION FOR CYBER SECURITY USING UNSUPERVISED DEEP LEARNING – A REVIEW

[*1]Jalaj Chirganiya, [2]Prof. Prakash Saxena

[1]Research Scholar Bansal Group of Institute of Science and Technology, Bhopal (M.P.)

[2]Assistant Professor Bansal Group of Institute of Science and Technology, Bhopal (M.P.)

Article Received: 12 April 2026

Article Revised: 02 May 2026

Published on: 22 May 2026

***Corresponding Author: Jalaj Chirganiya**

Research Scholar Department of Computer Science & Engineering Bansal Institute of Science & Technology, Bhopal.

DOI: https://doi-doi.org/101555/ijrpa.4922

ABSTRACT

The rapid growth of cyber attacks, network vulnerabilities, cloud services, IoT devices, and interconnected digital infrastructures has created a strong demand for intelligent and adaptive intrusion detection systems. Traditional rule-based and signature-based mechanisms remain useful for known threats, but they often fail to detect zero-day attacks, polymorphic malware, insider threats, and evolving intrusion patterns because they depend on predefined rules and manually updated signatures. This review paper presents a concise analysis of autoencoder-based anomaly detection for cyber security applications using unsupervised deep learning techniques. The paper examines the evolution of intrusion detection from statistical and machine learning methods to deep autoencoder models, reconstruction-based anomaly scoring, hybrid Autoencoder-LSTM frameworks, and explainable AI approaches. It also discusses cyber security datasets, preprocessing strategies, threshold selection, comparative model performance, practical applications, major research gaps, and future directions. The review shows that autoencoders are highly suitable for detecting unknown anomalies because they learn normal behaviour and flag deviations through reconstruction error without requiring large labelled attack datasets. However, challenges such as class imbalance, false-positive alarms, computational complexity, threshold sensitivity, encrypted traffic, and limited interpretability continue to affect real-world deployment. Future research should focus on adaptive thresholds, explainable anomaly detection, federated learning, edge AI, Transformer-based architectures, and multi-modal cyber analytics. Overall, autoencoder-based frameworks provide a scalable and promising direction for intelligent cyber security anomaly detection in modern digital environments.

KEYWORDS: Cyber Security, Anomaly Detection, Autoencoder, Deep Learning, Intrusion Detection System, Unsupervised Learning, Network Security, Artificial Intelligence**.**

1. INTRODUCTION

The rapid expansion of cloud computing, Internet of Things (IoT), mobile communication, online banking, healthcare networks, and cyber-physical systems has transformed the way organizations store, exchange, and process information. This transformation has also increased exposure to cyber attacks such as malware, ransomware, phishing, distributed denial-of-service attacks, insider threats, botnets, and advanced persistent threats. These attacks are increasingly adaptive and frequently modify their behaviour to avoid detection by conventional security systems [1]. Cyber security has therefore become a critical requirement for public, private, academic, financial, healthcare, defense, and industrial organizations. Traditional mechanisms such as firewalls, rule-based filters, and signature-based intrusion detection systems are effective against known attacks, but they struggle against zero-day exploits and evolving threats. Signature-based systems depend on predefined rules and continuously updated attack databases, which makes them less flexible under dynamic network conditions [2], [3]. Artificial Intelligence and Machine Learning introduced more adaptive detection capability by allowing systems to learn patterns from network traffic and security logs. Support Vector Machines, Decision Trees, Random Forests, Naive Bayes classifiers, and k-Nearest Neighbor models have been used for intrusion detection and malware classification. These models improved automatic classification, but many supervised approaches require large labelled datasets and may fail when attack patterns differ from training examples [4], [5].Unsupervised anomaly detection addresses this limitation by learning normal system behaviour and identifying deviations as suspicious events. Since normal network activity is usually more abundant than malicious traffic, unsupervised approaches are practical for real-world cyber security. Deep learning further strengthens this approach by learning complex non-linear patterns from high-dimensional traffic, logs, endpoint events, and user behaviour data without excessive handcrafted feature engineering [6], [7]. Autoencoders are among the most important deep learning models for unsupervised anomaly detection. An autoencoder consists of an encoder that compresses input data into a latent representation and a decoder that reconstructs the input. When trained on normal behaviour, the model reconstructs normal samples accurately but produces higher reconstruction error for anomalous samples. This property makes autoencoders suitable for detecting unknown attacks and abnormal network behaviour [8].

This review focuses on autoencoder-based anomaly detection for cyber security applications. It discusses intrusion detection systems, literature trends, deep autoencoder models, cyber security datasets, comparative model analysis, applications, challenges, and future research directions. The objective is to present a compact review that highlights both the potential and practical limitations of reconstruction-based deep learning for intelligent cyber defense.

2. Review of Literature

Research on cyber anomaly detection has evolved from statistical and rule-based methods to machine learning and deep learning frameworks. Early intrusion detection systems compared network traffic against known signatures or statistical behaviour profiles. Although these systems were simple and interpretable, they were not sufficiently adaptive for detecting zero-day attacks, polymorphic malware, and continuously changing intrusion strategies [1]. Classical machine learning methods improved detection by learning behavioural features from data. SVM, Random Forest, Decision Tree, Naive Bayes, and k-NN models were widely used for intrusion detection because they could classify benign and malicious traffic more automatically than rule-based systems. However, these models often required handcrafted features, labelled attack samples, and careful parameter tuning, which limited scalability in modern networks [2]. Chandola et al. provided a major foundation for anomaly detection research by emphasizing that anomaly-based methods are useful for identifying deviations from normal behaviour instead of relying only on known attack signatures [3]. This idea became highly relevant for cyber security because new threats often do not match previously documented attack patterns.

Deep learning transformed anomaly detection by enabling automated hierarchical representation learning from complex data. LeCun et al. highlighted the strength of deep learning in pattern recognition and high-dimensional data analysis [4]. Deep neural networks can learn hidden structures in network traffic, logs, and system behaviour without requiring extensive manual feature extraction.

Autoencoders became particularly important after Hinton and Salakhutdinov demonstrated their use for non-linear dimensionality reduction and representation learning [5]. Sakurada and Yairi further demonstrated anomaly detection using reconstruction error, showing that autoencoders can identify abnormal patterns in high-dimensional data [6]. These studies influenced later work on network intrusion detection using deep autoencoder architectures.

Mirsky et al. proposed Kitsune, an ensemble-based autoencoder framework for online network intrusion detection, and showed that lightweight autoencoders can support scalable real-time analysis [7]. Vinayakumar et al. reported that deep learning approaches improve intelligent intrusion detection compared with many traditional systems [8]. Robust and probabilistic models such as robust deep autoencoders and Deep Autoencoding Gaussian Mixture Models also improved anomaly discrimination in noisy environments [9], [10]. Recent studies expanded autoencoder-based detection to IoT, cloud, industrial, and cyber-physical environments. Deep autoencoders, sparse autoencoders, denoising autoencoders, variational autoencoders, and hybrid Autoencoder-LSTM models have been used to detect abnormal communication patterns, stealthy attacks, and sequential intrusions [11]–[14]. Current research also emphasizes adaptive thresholding, explainable AI, federated learning, edge deployment, and multi-modal cyber analytics [15]–[19].

3. Cyber Security and Anomaly Detection Systems

Cyber security systems aim to protect confidentiality, integrity, availability, and reliability of digital infrastructures. Modern networks continuously generate large volumes of traffic, authentication events, endpoint telemetry, system logs, application activity, and user behaviour records. Because cyber threats evolve rapidly, intelligent monitoring systems are required to detect suspicious behaviour before major damage occurs [1]. Traditional intrusion detection systems are commonly divided into signature-based and anomaly-based approaches. Signature-based systems match activity against known attack patterns and are effective for previously identified threats. However, they are weak against zero-day attacks and adaptive malware. Anomaly-based systems instead learn normal behaviour and flag deviations, making them more suitable for unknown or evolving threats [2]–[4]. Machine learning improved anomaly detection by automating pattern recognition. Supervised models are useful when labelled benign and malicious samples are available, but they often fail to generalize beyond known attacks. In real-world networks, obtaining accurate labels is difficult because attacks are rare, diverse, and constantly changing. This motivates the use of unsupervised and semi-supervised methods [5]–[7]. Unsupervised anomaly detection models learn the structure of normal traffic and identify outliers through clustering, density estimation, reconstruction error, or behavioural deviation. Such models are useful because normal traffic is easier to collect than labelled malicious samples. Autoencoders are especially effective because they learn compact representations of normal behaviour and detect anomalies through reconstruction errors [8]–[10].Cyber anomaly detection is applied

in network intrusion detection, cloud monitoring, IoT protection, fraud detection, insider threat analysis, malware communication detection, industrial control systems, and cyber-physical infrastructure security. In each domain, the detection system must balance sensitivity, false-positive reduction, interpretability, and real-time response capability [11]–[14].Despite progress, anomaly detection systems still face several barriers. Class imbalance, encrypted traffic, noisy logs, evolving attack behaviour, high false-positive rates, and computational complexity affect operational reliability. Explainable AI is also important because analysts must understand why a specific event is flagged as anomalous before taking action [15], [16].

4. Autoencoder-Based Deep Learning Models

Autoencoder-based models are widely used for cyber security anomaly detection because they can learn hidden representations of normal activity without requiring labelled attack data. A basic autoencoder contains an encoder, a bottleneck latent representation, and a decoder. The encoder compresses the input, and the decoder reconstructs it. During training, the model minimizes reconstruction error and learns the structure of normal behaviour [1], [2]. When an input sample belongs to the normal behaviour distribution, the trained autoencoder usually reconstructs it with low error. When an anomalous or malicious sample differs significantly from the learned normal pattern, reconstruction error increases. This error becomes the basis for anomaly scoring and classification. The method is attractive for cyber security because previously unseen attacks can be detected as deviations from normal behaviour [3].Dense autoencoders are suitable for structured network features such as packet statistics, flow duration, protocol information, byte counts, connection frequency, and authentication events. Deep autoencoders improve this by using multiple hidden layers to learn hierarchical behavioural representations. Sparse autoencoders encourage efficient feature learning by activating only important neurons, while denoising autoencoders improve robustness by reconstructing clean data from corrupted input [4]–[7].

Variational Autoencoders extend the concept by learning probabilistic latent representations, which improves uncertainty estimation and anomaly scoring. Stacked autoencoders combine multiple representation-learning stages for deeper feature extraction. Convolutional and recurrent autoencoders have also been applied to traffic-flow representations and sequential cyber events [8], [9].Hybrid Autoencoder-LSTM architectures are important because many attacks unfold over time. The autoencoder learns compact representations, while LSTM

layers capture temporal dependencies in communication sequences. Such models are useful for detecting insider threats, botnet communication, multi-stage attacks, and slow intrusion patterns [10]. Attention-based and Transformer-autoencoder models further improve contextual learning by focusing on the most relevant traffic features [11].

Performance is usually evaluated using accuracy, precision, recall, F1-score, ROC-AUC, confusion matrix analysis, false-positive rate, and detection latency. Since cyber datasets are often imbalanced, recall and false-positive behaviour are especially important. A system with high accuracy but poor anomaly recall may still be unsuitable for operational security [12].
Threshold selection is a central issue in autoencoder-based systems. Reconstruction errors must be compared with a threshold to classify activity as normal or anomalous. Static thresholds may fail under changing network conditions, while very low thresholds increase false alarms and very high thresholds miss attacks. Adaptive thresholding and probabilistic scoring are therefore important research directions [13]. Autoencoder models also require interpretability. Analysts need to know which features, traffic patterns, or reconstruction errors contributed to the anomaly decision. Explainable AI methods such as feature attribution, reconstruction visualization, saliency mapping, and anomaly explanation dashboards can improve analyst trust and operational usability [14], [15].

5. Dataset and Cyber Security Data Analysis

Cyber security datasets are essential for developing and evaluating anomaly detection systems. They may include network flows, packet features, system logs, endpoint telemetry, user behaviour records, authentication attempts, cloud resource usage, IoT communication, and industrial sensor activity. Dataset quality, diversity, and realism directly affect the reliability of machine learning and deep learning models [1]. Early intrusion detection research frequently used the KDD Cup 1999 dataset, which contains simulated traffic and attack categories such as denial-of-service, probing, remote-to-local, and user-to-root attacks. Although influential, KDD Cup has limitations such as redundancy, outdated traffic, unrealistic distributions, and severe class imbalance [2]. NSL-KDD was later introduced to reduce redundancy and improve benchmarking while preserving similar attack categories [3]. More recent datasets such as UNSW-NB15 and CICIDS2017 provide modern traffic patterns and diverse attack scenarios, including exploits, fuzzers, shellcode, botnet communication, infiltration, brute force, web attacks, and DDoS events [4], [5]. IoT-focused datasets such as

Bot-IoT and TON_IoT support anomaly detection research in heterogeneous smart-device and cyber-physical environments [6].

Preprocessing is a necessary stage before deep learning implementation. Missing values, duplicated records, corrupted entries, and irrelevant features must be removed or handled. Numerical features are normalized or standardized so that values with large ranges do not dominate training. Categorical features such as protocol type, service type, and connection state are converted into numerical form using label encoding or one-hot encoding [7], [8].
Feature selection and dimensionality reduction improve efficiency and generalization. Redundant features increase computational cost and may reduce model performance. Methods such as correlation analysis, mutual information, PCA, and autoencoder-based compression are used to identify compact and informative representations [9], [10].

Class imbalance is a persistent challenge because normal traffic usually dominates cyber datasets, while malicious events are rare. Imbalanced training can bias models toward normal classification and reduce anomaly sensitivity. Resampling, synthetic data generation, cost-sensitive learning, and balanced metrics such as recall and F1-score help address this issue [11], [12]. Temporal analysis is also important because sophisticated attacks often unfold gradually. LSTM networks, temporal autoencoders, and sequence-based models can analyze behaviour across sessions rather than treating each flow independently. Visualization and explainable AI techniques further support anomaly investigation by showing traffic clusters, feature contributions, reconstruction errors, and suspicious behavioural trends [13]–[16].

6. Comparative Analysis of Existing Models

Comparative analysis helps identify the strengths and limitations of statistical methods, machine learning models, deep learning architectures, and hybrid anomaly detection frameworks. Each approach performs differently depending on dataset size, feature quality, attack diversity, computational resources, and deployment requirements [1]. Statistical anomaly detection methods are simple and computationally efficient, but they struggle with high-dimensional network traffic and complex non-linear patterns. Static thresholds may generate high false-positive rates when network behaviour changes. Classical machine learning models such as SVM and Random Forest provide stronger classification ability, but they usually depend on labelled datasets and handcrafted features [2]–[4].

Unsupervised clustering methods such as K-Means, DBSCAN, hierarchical clustering, and Gaussian Mixture Models can identify outliers without attack labels. However, they are often sensitive to parameter selection and may not scale well with high-dimensional enterprise traffic [5]. Deep learning methods provide stronger representation learning. Autoencoders learn normal behaviour and detect deviations using reconstruction error. Deep, sparse, denoising, variational, and stacked autoencoders improve robustness, compact feature learning, and anomaly discrimination. Hybrid Autoencoder-LSTM models add temporal learning and are useful for sequential attacks [6]–[9].

CNN-based models can learn local spatial patterns from transformed traffic features, while Transformer-based models and attention mechanisms are promising for long-range dependencies and contextual behaviour analysis. However, these advanced architectures may require larger datasets and greater computational resources [10], [11]. Evaluation should not rely only on accuracy because cyber datasets are usually imbalanced. Precision, recall, F1-score, ROC-AUC, false-positive rate, confusion matrix analysis, latency, and computational cost provide a more complete assessment. In practical security operations, a low false-positive rate is necessary to prevent alert fatigue, while high recall is required to avoid missed attacks [12]–[15].

Table 1: Comparative Analysis of Existing Anomaly Detection Models.

Model Type	Learning Type	Strengths	Limitations
Statistical Methods	Unsupervised	Simple implementation, low computation	High false positives, limited scalability
SVM	Supervised	Strong classification capability	Requires labelled data and tuning
Random Forest	Supervised	Robust and accurate on structured data	Feature engineering required
Clustering Algorithms	Unsupervised	Detects unknown outliers	Sensitive to parameters
CNN	Deep Learning	Learns local traffic patterns	High computation and large data need
Autoencoder	Unsupervised Deep Learning	Detects unknown anomalies, scalable	Threshold sensitivity
Autoencoder-LSTM	Hybrid Deep Learning	Captures temporal anomaly behaviour	Complex training
Transformer-Based Models	Deep Learning	Handles long-range dependencies	High computational cost

7. Challenges and Research Gaps

Although autoencoder-based anomaly detection has shown strong potential, several challenges continue to affect practical deployment. The first major issue is the constantly evolving nature of cyber attacks. Models trained on historical data may not perform well against future attack variants, especially when attackers modify traffic behaviour, payload patterns, or evasion techniques [1], [2]. Class imbalance remains a serious challenge because normal traffic dominates most datasets, while attack samples are comparatively rare. A model may achieve high overall accuracy by predicting the majority class but still fail to detect rare attacks. False positives are also problematic because excessive alerts can overwhelm analysts and reduce trust in the system [3], [4]. Threshold selection is another major research gap in reconstruction-based models. Autoencoder output must be converted into a decision using reconstruction-error thresholds. Static thresholds may be unsuitable for dynamic traffic conditions, while poorly selected thresholds either increase false alarms or miss subtle intrusions. Adaptive and context-aware thresholding is therefore necessary [5], [6].

Interpretability remains limited in many deep learning systems. Security analysts require clear explanations of why a specific event is anomalous, which features contributed most, and how severe the threat may be. Without explainability, black-box models may be difficult to trust in operational security environments [7]. Computational complexity and scalability are also important. Modern networks generate massive real-time traffic, and deep models may require GPUs, large memory, and long training times. Edge devices and IoT systems have limited resources, so lightweight architectures, model compression, and streaming inference are important for real-time deployment [8]–[10]. Encrypted traffic creates additional difficulty because payload-level inspection becomes unavailable. Detection systems must rely on metadata, flow statistics, timing behaviour, and communication patterns while preserving privacy. Adversarial attacks against deep learning models are another concern, as attackers may intentionally manipulate inputs to evade anomaly detection [11], [12].

Future research must also address dataset realism, cross-domain generalization, federated learning, multi-modal cyber analytics, privacy preservation, and integration with automated response systems. A practical anomaly detection framework should be adaptive, explainable, privacy-aware, scalable, and robust against adversarial manipulation [13]–[16].

8. Applications and Significance

Autoencoder-based anomaly detection has practical significance across many cyber security domains because it can identify unknown threats through behavioural learning. In network intrusion detection, autoencoders model legitimate traffic and identify deviations associated with unauthorized access, malware communication, scanning, reconnaissance, botnets, or abnormal traffic flows [1], [2]. Cloud security is another important application. Cloud environments involve virtual machines, containers, multi-tenant services, authentication events, and dynamic resource scaling. Autoencoder-based systems can monitor resource usage, access behaviour, and traffic flows to detect suspicious events and policy violations [3].

IoT and edge environments benefit from lightweight anomaly detection because connected devices are often resource constrained and vulnerable to botnet attacks, device compromise, and abnormal communication. Autoencoder models deployed at gateways or edge nodes can support real-time behavioural monitoring without sending all raw data to a central server [4]. Industrial control systems and cyber-physical infrastructures also require anomaly detection because attacks can disrupt manufacturing, utilities, transportation, and energy systems. Autoencoders can analyze sensor readings, process variables, and communication behaviour to identify abnormal operational states [5]. Fraud detection, insider threat analysis, malware detection, and healthcare cyber security are additional application areas. Financial systems can use anomaly detection to identify suspicious transactions, while user and entity behaviour analytics can detect unusual login or data-access behaviour. Healthcare networks can protect patient records, connected devices, and hospital information systems from ransomware and unauthorized access [6]–[10].

The major significance of autoencoder-based systems lies in their unsupervised learning capability. They reduce dependence on labelled attack datasets and can detect previously unseen behaviour that bypasses signature-based systems. Their scalability also supports automated monitoring of large volumes of network traffic and security telemetry [11]–[13]. Integration with Security Information and Event Management platforms, threat intelligence systems, and automated incident response workflows can further increase their operational value. Explainable AI improves trust by showing suspicious features and reconstruction patterns, while edge and federated deployment improve privacy and latency [14], [15].

9. Future Scope

Future research in autoencoder-based anomaly detection should focus on making systems more adaptive, explainable, scalable, and deployable under real-world conditions. Advanced architectures such as Variational Autoencoders, Graph Neural Networks, Sparse Autoencoders, Transformer-based models, and hybrid Autoencoder-LSTM frameworks can improve representation learning and sequential attack detection [1], [2]. Federated learning is a promising direction because organizations often cannot share raw security data due to privacy and legal restrictions. Federated anomaly detection can train models across distributed sites without transferring sensitive data, enabling collaborative cyber defense while preserving privacy [3].

Edge AI is also important for IoT, industrial, and cyber-physical systems. Lightweight autoencoder architectures can support low-latency detection at gateways and local devices. Model compression, pruning, quantization, and hardware acceleration can further improve feasibility in resource-constrained environments [4]. Explainable AI should become a core part of future anomaly detection systems. Feature attribution, reconstruction-error visualization, saliency maps, attention heatmaps, and interactive dashboards can help analysts understand anomaly decisions and prioritize incident response [5].

Adaptive threshold optimization is required to reduce false positives and improve detection sensitivity under changing network behaviour. Future systems may use probabilistic thresholds, reinforcement learning, streaming statistics, or context-aware scoring mechanisms to adjust anomaly decisions dynamically [6]. Multi-modal cyber analytics can improve contextual awareness by integrating network flows, endpoint telemetry, user behaviour, authentication logs, system events, cloud metadata, and threat intelligence feeds. Synthetic data generation, self-supervised learning, adversarial training, and privacy-preserving analytics may further improve robustness and generalization [7]–[12].

Future intelligent cyber defense systems may also integrate anomaly detection with automated response, blockchain-based threat sharing, human-centered dashboards, and autonomous orchestration platforms. These developments can reduce response time and strengthen resilience against sophisticated cyber threats [13], [14].

10. CONCLUSION

This review presented a compact analysis of autoencoder-based anomaly detection for cyber security applications using unsupervised deep learning. The review examined the evolution of intrusion detection from traditional rule-based and statistical approaches to modern AI and deep learning frameworks. It emphasized the role of autoencoders in learning normal behavioural patterns and identifying deviations through reconstruction-based analysis.

Traditional signature-based systems are effective for known attacks but struggle with zero-day threats, polymorphic malware, insider attacks, and adaptive intrusion strategies. Classical machine learning improved detection automation, but supervised models often require labelled datasets and handcrafted feature engineering. Autoencoders address these limitations by learning compact representations of normal behaviour and detecting abnormal activity without extensive labelled attack samples.

Deep autoencoders, sparse autoencoders, denoising autoencoders, variational autoencoders, and hybrid Autoencoder-LSTM models have demonstrated strong potential for network intrusion detection and behavioural anomaly analysis. Cyber security datasets such as KDD Cup, NSL-KDD, UNSW-NB15, CICIDS2017, Bot-IoT, and TON_IoT support model development, while preprocessing, normalization, feature selection, and threshold optimization remain essential for reliable performance.

Comparative analysis indicates that deep learning-based unsupervised frameworks generally provide better scalability and adaptability than many traditional methods, especially for detecting unknown threats. However, major challenges remain, including dataset imbalance, false positives, threshold sensitivity, encrypted traffic, adversarial attacks, computational cost, interpretability limitations, and cross-domain generalization.

Future research should focus on explainable AI, adaptive thresholding, federated learning, edge deployment, Transformer-based architectures, privacy-preserving analytics, and multi-modal cyber defense systems. Overall, autoencoder-based anomaly detection provides a promising direction for intelligent, scalable, and adaptive cyber security monitoring in modern digital infrastructures.

REFERENCES

1. Arafah, M. (2025). Anomaly-based network intrusion detection with denoising autoencoder and WGAN. Cybersecurity, 8(1), 115–132. https://doi.org/10.1016/j.cose.2025.103214
2. Aslam, M. M., Khan, S., & Ali, R. (2024). An improved autoencoder-based approach for anomaly detection in industrial control systems. International Journal of Automation and Smart Technology, 14(3), 225–239. https://doi.org/10.1080/23270012.2024.1023345
3. Anyfantis, G., & Barlet-Ros, P. (2025). AutoGraphAD: Variational graph autoencoders for network flow anomaly detection. arXiv Preprint. https://doi.org/10.48550/arXiv.2502.01457
4. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM Computing Surveys, 41(3), 1–58. https://doi.org/10.1145/1541880.1541882
5. Dardouri, S., & Almuhanna, R. (2025). A deep learning and machine learning approach for anomaly-based network intrusion detection. Frontiers in Artificial Intelligence, 8, 1450221. https://doi.org/10.3389/frai.2025.1450221
6. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.
7. Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. Science, 313(5786), 504–507. https://doi.org/10.1126/science.1127647
8. Kim, D., Lee, J., & Park, S. (2025). Adaptive autoencoder-based intrusion detection system for CAN networks. Sensors, 25(4), 1142. https://doi.org/10.3390/s25041142
9. Korniszuk, K. (2024). Autoencoder-based anomaly detection in network traffic. Proceedings of CPEE 2024, 88–95. https://doi.org/10.1109/CPEE62412.2024.10456121
10. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. https://doi.org/10.1038/nature14539
11. Mirsky, Y., Doitshman, T., Elovici, Y., & Shabtai, A. (2018). Kitsune: An ensemble of autoencoders for online network intrusion detection. Network and Distributed System Security Symposium. https://doi.org/10.14722/ndss.2018.23204
12. Narmadha, S., Kumar, V., & Rao, P. (2025). Improved network anomaly detection system using LSTM-autoencoder with PSO optimization. Expert Systems with Applications, 252, 124125. https://doi.org/10.1016/j.eswa.2025.124125
13. Okolie, S. A. (2025). Anomaly detection in heterogeneous cybersecurity data: Machine learning and deep learning perspectives. Cybersecurity, 8(1), 45–67. https://doi.org/10.1016/j.cose.2025.102998
14. Pang, G., Shen, C., Cao, L., & Hengel, A. V. D. (2021). Deep learning for anomaly detection: A review. ACM Computing Surveys, 54(2), 1–38. https://doi.org/10.1145/3439950

15. Rassam, M. A. (2024). Autoencoder-based neural network model for anomaly detection in WBANs. Sensors, 24(9), 2890. https://doi.org/10.3390/s24092890
16. Rhachi, H., Ahmed, M., & Karim, R. (2025). Enhanced anomaly detection in IoT networks using deep autoencoders. Sensors, 25(6), 1788. https://doi.org/10.3390/s25061788
17. Sakurada, M., & Yairi, T. (2014). Anomaly detection using autoencoders with nonlinear dimensionality reduction. Proceedings of MLSDA 2014, 4–11. https://doi.org/10.1145/2689746.2689747
18. Saranya, K., Rajesh, P., & Kumar, S. (2025). Multi-layer deep autoencoder for cross-layer IoT threat detection. Scientific Reports, 15, 4412. https://doi.org/10.1038/s41598-025-4412-7
19. Somma, M. (2025). Hybrid temporal differential consistency autoencoder for cyber-physical system anomaly detection. arXiv Preprint. https://doi.org/10.48550/arXiv.2501.08214
20. Syed, A., & Ahmad, M. I. (2025). Multi-modal deep learning autoencoder approach for cloud security. arXiv Preprint. https://doi.org/10.48550/arXiv.2503.01892
21. Vinayakumar, R., Alazab, M., Soman, K. P., Poornachandran, P., & Venkatraman, S. (2019). Deep learning approach for intelligent intrusion detection system. IEEE Access, 7, 41525–41550. https://doi.org/10.1109/ACCESS.2019.2895334
22. Wang, W., Sheng, Y., Wang, J., Zeng, X., Ye, X., Huang, Y., & Zhu, M. (2017). HAST-IDS: Learning hierarchical spatial-temporal features using deep neural networks to improve intrusion detection. IEEE Access, 6, 1792–1806. https://doi.org/10.1109/ACCESS.2017.2780250
23. Xu, H., Shen, C., & Zhao, J. (2024). Deep autoencoder-based cyber anomaly detection using reconstruction learning. Journal of Information Security and Applications, 79, 103612. https://doi.org/10.1016/j.jisa.2024.103612
24. Zhou, C., & Paffenroth, R. C. (2017). Anomaly detection with robust deep autoencoders. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 665–674. https://doi.org/10.1145/3097983.3098052
25. Zong, B., Song, Q., Min, M. R., Cheng, W., Lumezanu, C., Cho, D., & Chen, H. (2018). Deep autoencoding Gaussian mixture model for unsupervised anomaly detection. International Conference on Learning Representations.
26.