Informatics Bulletin, Faculty of Computers and Artificial Intelligence, Helwan University
Published Online *Vol 7 Issue 2, July 2025*
(https://fcihib.journals.ekb.eg)

A Literature Review on Anomaly Detection using Deep Learning Techniques

Saad Khalifa[1], Wael Mohamed[1], Mohamed Marie[1]
[1]Information System Dept., Faculty of Computers and Artificial Intelligence, Helwan University, Egypt
khalifasaad43@fci.helwan.edu.eg, waelmohamed@fci.helwan.edu.eg, dr.mmarie@fci.helwan.edu.eg

***Abstract*—Anomaly detection, which involves identifying irregular patterns that diverge from normal behavior, plays a vital role in various fields such as cybersecurity, healthcare, financial systems, and the Internet of Things (IoT). Recognizing anomalies is key to uncovering problems like fraudulent activities, system failures, or security intrusions. Traditional methods for anomaly detection, which typically rely on statistical or threshold-based techniques, are effective for low-dimensional or static data but struggle with high-dimensional, intricate, and dynamic datasets. As data complexity and volume have increased, machine learning (ML) and deep learning (DL) techniques have become essential for enhancing detection precision and adaptability. Identifying anomalies in networks is particularly important for bolstering cybersecurity, acting as a proactive approach to prevent or reduce cyber threats. With the rapid progress in Artificial Intelligence (AI), many deep learning-based approaches utilizing Autoencoders (AEs) have been created to improve network security. However, the performance of these advanced AE models varies widely, and they often lack a thorough framework for assessing critical performance metrics that impact detection accuracy.**

***Index Terms*— Anomaly Detection, Internet of Things (IoT), Traditional Detection Methods, Artificial Intelligence (AI), Cyber Threats, Cybersecurity, Deep Learning (DL), Machine Learning (ML), Autoencoders (AEs).**

I. Introduction

In the digital age, technologies like the Internet, smartphones, and robotics are now essential to everyday life. The fast-paced growth of IT and the reliance on real-time, data-driven decisions have boosted global data transmission. Yet, this progress has also brought complex cybersecurity risks, challenging developers, manufacturers, users, and security groups. Cybersecurity organizations strive to manage the rising volume, speed, and variety of data for real-time threat detection. Increased data exchange has amplified security issues such as malware, DDoS attacks, phishing, and APTs, endangering individuals, businesses, and critical systems [1]. Anomaly detection is crucial for spotting deviations from normal patterns. It is widely used in areas like financial fraud prevention and traffic monitoring to identify rare, impactful events. Detecting such anomalies, however, poses challenges due to their infrequent occurrence and subtle nature. This underscores the importance of advanced detection techniques to effectively identify and mitigate potential threats [2]. Anomaly detection is vital as even a few anomalies can lead to significant impacts, making it crucial in areas like cybersecurity for spotting network intrusions and social media for identifying fraud, such as Sybil accounts. Consequently, network anomaly detection has grown in importance. However, labeling anomalies is often laborious and time-consuming, leading most methods to rely on unsupervised techniques. Popular approaches include autoencoder-based models and matrix factorization, which detect anomalies without needing large amounts of labeled data [3]. In IoT and sensor networks, anomaly detection is key for spotting sensor failures, environmental shifts, or intrusions. With IoT devices generating constant data streams, real-time detection is crucial. Deep learning techniques, such as CNNs and RNNs, are effective as they analyze both spatial and temporal patterns. Hybrid methods that merge ML, DL, and statistical approaches have shown enhanced accuracy and efficiency, particularly in real-time edge computing scenarios [4]. ML and DL-based anomaly detection methods are typically divided into supervised, unsupervised, and semi-supervised approaches. Supervised techniques depend on labeled data but are hindered by the lack of labeled anomalies. Unsupervised methods, like clustering and density-based models, are more prevalent as they identify deviations by modeling normal behavior. Common approaches include k-means clustering and isolation forests, though they may face challenges with intricate data [5]. Deep learning models like autoencoders and RNNs provide significant progress in identifying anomalies in high-dimensional and sequential data. Autoencoders excel at reconstructing normal patterns, using reconstruction errors to flag anomalies. RNNs, particularly LSTMs, are adept at capturing temporal dependencies in time-series data, making them highly effective for detecting anomalies in sequences like network traffic or financial transactions [6]. Autoencoder (AE) models have become increasingly popular in deep learning for identifying anomalies in large-scale network traffic datasets. Their strength lies in efficiently learning and reconstructing data. During training, the AE reduces reconstruction loss, and the resulting

loss rate acts as a critical measure to classify network samples as normal or anomalous [7].

While several existing surveys have examined anomaly detection and its applications, many focus either broadly on machine learning techniques or are limited to specific application domains. In contrast, this review provides a comprehensive and up-to-date synthesis of deep learning-based approaches, with a focus on categorizing models based on their architectures (e.g., autoencoders, GANs, RNNs) and highlighting how each performs across various data types and domains. Furthermore, we go beyond summarization by offering a critical analysis of the strengths and limitations of each method, as well as discussing practical challenges such as interpretability, scalability, and training on imbalanced datasets. We also introduce a refined taxonomy that groups models not just by architecture, but also by their training paradigms (supervised, unsupervised, self-supervised), which have been underexplored in prior reviews. Lastly, we identify emerging trends and future directions to guide researchers in addressing current gaps in the field. These contributions position our work as a valuable and forward-looking resource for both newcomers and experienced researchers in anomaly detection.

Additionally, this paper is divided into three sections: Section II provides background on key concepts, such as machine learning and deep learning. Section III presents reviews of recent literature and highlights key findings. Finally, Section IV concludes the paper.

II. Background

This section illustrates the fundamental concepts, theories, and advancements in anomaly detection, emphasizing its significance in modern data analysis.

Anomaly detection, also known as outlier detection, is important for finding data patterns that deviate considerably from predicted behavior. These anomalies may signal fraudulent activity, cyber-attacks, medical abnormalities, or equipment problems in industrial systems. Historically, statistical techniques and rule-based systems were used for this goal, but these approaches frequently fail when dealing with high-dimensional and complicated datasets. Recent improvements in machine learning (ML) and deep learning (DL) have improved anomaly detection models by allowing them to process large volumes of data, learn complex patterns, and spot deviations more accurately. Unlike previous methods, ML and DL algorithms may adapt to new abnormalities without requiring considerable user intervention, making them particularly successful in dynamic contexts [8].

A. Theories and Frameworks in Anomaly Detection

A variety of theoretical models and computational frameworks, which can be broadly classified into statistical techniques, supervised learning, unsupervised learning, semi-supervised learning, and hybrid models, form the basis of anomaly detection. Depending on the type of data and the problem domain, each of these approaches has unique benefits and drawbacks.

1. Statistical Approaches

   Because classical statistical approaches can mathematically represent data distributions, they have been employed for a long time in anomaly identification. Normal data is assumed to follow a particular statistical distribution by methods like Principal Component Analysis (PCA), Z-score analysis, and Gaussian Mixture Models (GMM). Anomalies are defined as instances that substantially depart from the expected distribution.

   - Gaussian Mixture Models (GMM): GMM assumes that data comes from a combination of several Gaussian distributions. GMM gives each case a probability by estimating the parameters of these distributions, designating as anomalies those with low likelihood.
   - Z-score Analysis: By calculating the number of standard deviations a point deviates from the mean; this method standardizes data. Anomalies are identified when data points above a predetermined threshold, such as ±3 standard deviations. The equation of Z-score is shown in formula (1).

$$\text{Z-score Formula: } Z = \frac{x - \mu}{\sigma} \qquad (1)$$

Where:

- X= data point
- μ= mean of the dataset
- σ= standard deviation

Figure 1 illustrates the concept of detecting outliers using Z-scores in statistics.

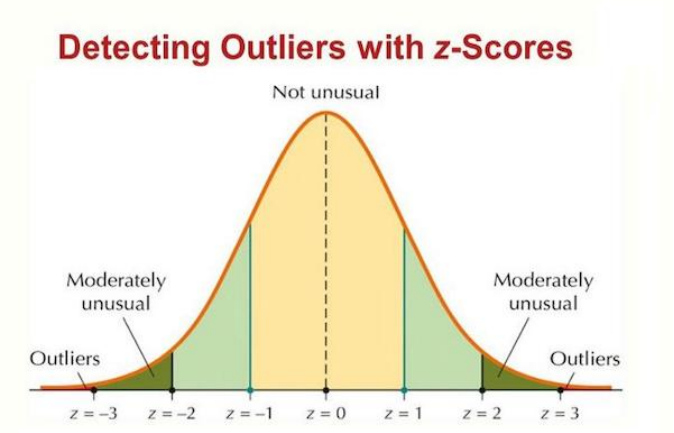

Fig. 1. Detecting outliers using Z-scores in statistics.

Figure 1 Z-score > 3 or Z-score < -3: These results show that the data point may be an outlier since it deviates from the mean by more than three standard deviations [9].

While statistical models perform well with low-dimensional and regularly distributed data, they struggle with high-dimensional datasets, multimodal distributions, and nonlinear patterns found in real-world applications such as cybersecurity and industrial monitoring [10].

2. Supervised Learning Approaches

Supervised machine learning algorithms have grown in favor of anomaly identification due to their predictive power. Support Vector Machines (SVM), Decision Trees, and Random Forests are trained on labeled datasets that contain both normal and anomalous samples to develop a decision boundary that distinguishes the two groups.

- Support Vector Machines (SVM): Hyperplanes are used by SVM to distinguish between examples that are normal and those that are aberrant. It can identify irregularities in non-linear datasets when paired with kernel techniques.
- Random Forests: To increase generality, this ensemble learning technique creates several decision trees, each trained on a distinct subset of data. Unusual occurrences are frequently classified inconsistently across trees, which facilitates their detection.

Figure 2 illustrates the detection of anomalies using random forests.

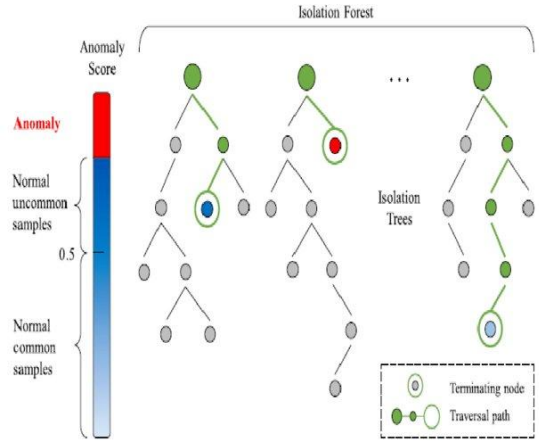

Fig. 2. Detecting anomalies using random forests.

The Isolation Trees in figure 2 demonstrate the Isolation Forest technique, in which anomalies (red) are separated in fewer splits than typical samples (blue). Gray nodes show intermediate steps, and green nodes show decision splits. Anomalies are found in sparser regions and are identified by shorter traversal paths, as indicated by the anomaly score scale on the left [11].

The primary drawback of supervised learning for anomaly detection is the scarcity of labeled anomaly data. Since anomalies are rare, acquiring sufficient labeled examples is challenging, and models trained on insufficient anomaly samples may fail to generalize well about unseen attacks or failures [10].

3. Unsupervised Learning Techniques

Unsupervised learning techniques are now the go-to option for anomaly identification because labeled data is difficult to obtain. By examining the data's natural structure, these methods identify abnormalities without depending on previously tagged examples.

- Clustering-Based Methods (e.g., k-Means, DBSCAN): According to these approaches, anomalies are located far from dense clusters formed by typical cases.
- k-Means Clustering: divides the data into k clusters, with outliers usually located far from the centroids.
- DBSCAN (Density-Based Spatial Clustering of Applications with Noise): points in low-density areas with few nearby data points are identified as anomalies.
- Density-Based Methods (e.g., Local Outlier Factor - LOF): Each data point's local density is calculated by LOF, which then contrasts it with its neighbors. Examples that are substantially less dense than their surroundings are called anomalies.

Although unsupervised learning methods are often successful, they can have significant false positive rates, particularly when the distribution of normal data is highly variable [9].

Figure 3 illustrates clustering-based anomaly detection in a two-dimensional dataset.

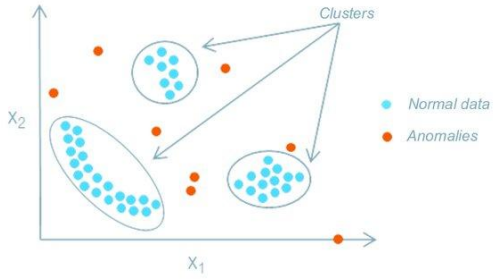

Fig. 3. Detecting anomalies using clustering based.

Clustering-based anomaly detection in a two-dimensional dataset is depicted in figure 3. Normal data is represented by the blue spots, which clump together to form clusters. The orange points indicate anomalies, or outliers, that diverge from these clusters and point to uncommon or infrequent events. This strategy is frequently employed in unsupervised learning techniques like DBSCAN and k-Means, which identify anomalies by their separation from dense clusters [12].

4. Semi-Supervised Learning Approaches

By combining a small quantity of labeled data with a larger pool of unlabeled data, semi-supervised learning fills the gap between supervised and unsupervised approaches. Given the scarcity of labeled anomalies, this method is especially helpful in anomaly detection.

- Self-Learning Techniques: Using reliable predictions from the unlabeled dataset, a model iteratively improves after being trained with labeled data.
- Autoencoders with Semi-Supervised Training: Neural networks known as autoencoders can

recognize deviations when given anomalous inputs and learn normal behavior from labeled normal data.

The quality and representativeness of the labeled subset determines how effective semi-supervised learning is, even though it outperforms fully unsupervised techniques in detection performance [10].

The Autoencoder, a neural network used for anomaly detection and dimensionality reduction, is depicted in figure 4. It is composed of an encoder that captures key information by compressing the input (x) into a lower-dimensional latent space (z), also referred to as the bottleneck. With the goal of minimizing reconstruction loss ($x \approx x'$), the decoder then uses this compressed representation to recover the input (x'). Because the reconstruction error is substantial when an input deviates greatly from taught patterns, autoencoders are useful for anomaly detection because they can identify data that cannot be reliably recreated. Autoencoders are a type of neural network used to learn a compressed representation of input data. They consist of two main parts: an encoder, which reduces the input data to a lower-dimensional latent space, and a decoder, which tries to reconstruct the original input from this compressed form. In anomaly detection, autoencoders are trained on normal data so they learn to reconstruct it well. When the model encounters an anomaly—something different from the normal patterns—it struggles to reconstruct it accurately, resulting in a high reconstruction error. This error can then be used to flag potential anomalies [13].

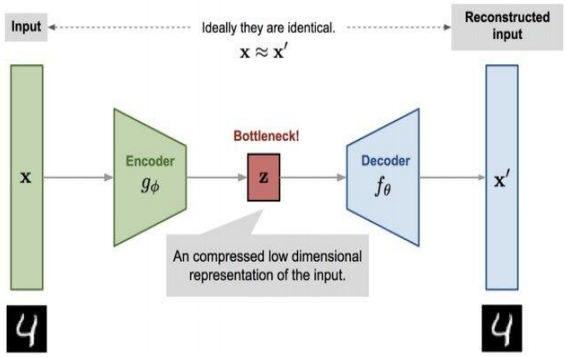

Fig. 4. Detecting anomalies using autoencoders.

5. Hybrid Models for Anomaly Detection

Hybrid approaches integrate various strategies to create more robust anomaly detection, hence overcoming the limits of individual techniques. To improve detection accuracy and lower false positives, these models use supervised, unsupervised, and deep learning techniques.

- Autoencoder + Isolation Forest: To further differentiate between normal and anomalous points, autoencoders effectively rebuild normal data, and the reconstruction mistakes they produce are then put into an isolation forest.
- Deep Learning + Statistical Methods: More efficient anomaly identification in complicated datasets is made possible by combining neural networks with conventional statistical outlier detection methods.

Unless they are effectively optimized, hybrid models are less appropriate for real-time applications since they frequently demand greater processing power and longer training periods [10].

Figure 5 depicts an anomaly detection framework combining an Autoencoder and HDBSCAN. QAR data is processed through a Time-Feature Attention Module, encoded into a latent space (L), and reconstructed by a Decoder. Reconstruction loss aids model training. HDBSCAN clusters latent representations, identifying normal points (blue) and outliers (orange) for anomaly detection [14].

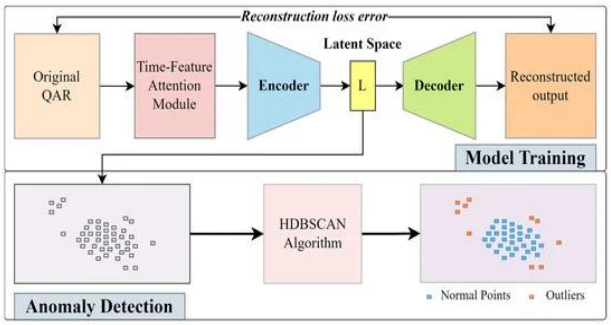

Fig. 5. Detecting anomalies using hybrid models.

Overall, these theoretical foundations provide the basis for modern anomaly detection techniques. The selection of an appropriate method depends on the nature of the data, the availability of labeled anomalies, computational constraints, and domain-specific requirements.

B. Historical Evolution of Anomaly Detection

Over time, the anomaly detection field has changed. Simple statistical tests that assumed data fit predetermined distributions were the foundation of early methods. These techniques were useful for small-scale datasets, but they were not flexible or scalable. An important change occurred with the introduction of machine learning algorithms in the early 2000s, which made it possible to automatically identify patterns in big, unstructured datasets. Among the earliest ML-based methods were Principal Component Analysis (PCA) and One-Class SVM. But as data complexity has increased, deep learning methods like Autoencoders, Generative Adversarial Networks (GANs) where GANs are composed of two neural networks that compete: a generator and a discriminator. The generator creates synthetic data samples that try to mimic the real data, while the discriminator tries to distinguish between real and generated samples. Through this adversarial process, the generator becomes better at producing realistic data. In anomaly detection, a GAN can be trained on normal data so that it learns the characteristics of non-anomalous patterns. When a new data sample is introduced, if the generator fails to produce a similar version or if the discriminator identifies it as “unrealistic,” the sample may be considered an anomaly, and Variational

Autoencoders (VAEs) have become more popular because of their capacity to extract high-dimensional features and generalize to abnormalities that are not visible [15].

C. The Role of Autoencoders in Anomaly Detection

Autoencoders, a specialized type of neural network, are designed to learn efficient representations of input data by compressing and reconstructing it. During training, they minimize the reconstruction error for normal instances. However, when an anomalous sample is introduced, the reconstruction error spikes, making it an effective anomaly detection tool. Variants like Sparse Autoencoders, Denoising Autoencoders, and Deep Autoencoding Gaussian Mixture Models (DAGMM) have further improved performance. These architectures are widely applied in cybersecurity, fraud detection, and medical diagnostics, where anomalies need to be accurately identified [16].

D. Variational Autoencoders (VAEs) for Anomaly Detection

Variational Autoencoders (VAEs), a more sophisticated type of autoencoder, add a probabilistic latent space representation to the basic framework. VAEs are very good at differentiating between normal and anomalous patterns because they model the distribution of normal data rather than just learning to recreate inputs. Their capacity to produce realistic samples also improves their ability to identify irregularities in cybersecurity applications, financial transactions, and medical imaging [17].

E. Variational Autoencoders (VAEs) for Anomaly Detection

For sequential data, Long Short-Term Memory (LSTM) networks, a variant of Recurrent Neural Networks (RNNs), have proven to be exceptionally powerful. These models capture long-range dependencies in time-series data, making them ideal for anomaly detection in financial transactions, predictive maintenance, and network intrusion detection. By learning temporal patterns, LSTMs differentiate between normal sequences and anomalous fluctuations, significantly improving detection accuracy [18].

F. CNNs for Anomaly Detection in Visual Data

Because Convolutional Neural Networks (CNNs) can extract spatial characteristics, they are quite effective in situations where there are abnormalities in pictures and video data. CNN-based models are extensively utilized in autonomous car safety, medical imaging for tumor identification, and manufacturing flaw detection. CNNs' capacity to identify small abnormalities in a variety of areas is further improved by sophisticated methods like Transfer Learning and GAN-based approaches [19].

G. Diffusion Models in Unsupervised Anomaly Detection

Diffusion models have been investigated recently for anomaly identification, particularly in fields with little labeled data. By learning to produce normal data distributions, these models are able to identify deviations and identify anomalous cases. Their use in medical imaging, cybersecurity, and finance has shown encouraging outcomes and provided a fresh viewpoint on outlier detection [20].

H. Bias and Fairness in Anomaly Detection Models

A critical issue in anomaly detection is bias, which can lead to unfair or discriminatory outcomes in applications like fraud detection and hiring decisions. If models are trained on biased datasets, they may unfairly classify certain groups as anomalies. Strategies like fairness-aware training, adversarial debiasing, and explainable AI (XAI) have been developed to address these ethical challenges, ensuring anomaly detection models remain equitable and trustworthy [21].

Having established the foundational concepts of anomaly detection and the motivation for leveraging deep learning, the subsequent section provides an in-depth examination of the key methodologies that have emerged in recent years. We explore how various deep learning architectures ranging from autoencoders and convolutional neural networks to recurrent models and generative frameworks have been employed to tackle the challenges of detecting anomalies across different domains. This literature review not only categorizes these methods based on their underlying architecture and application area but also critically assesses their performance, strengths, and limitations, offering insights into the current state of the field.

Traditional anomaly detection methods such as statistical models, distance-based techniques, and clustering algorithms face several persistent challenges. These include a strong reliance on assumptions about data distribution (e.g., normality), poor performance in high-dimensional or non-linear datasets, and limited adaptability to complex, dynamic data environments. Additionally, traditional models typically require manual feature engineering and often fail to generalize across different domains. They also struggle with imbalanced datasets, where anomalies are rare and difficult to learn. These limitations reduce their effectiveness in modern real-world applications, thus motivating the shift toward deep learning approaches, which can automatically learn complex patterns, handle unstructured data, and scale to high-dimensional problems more effectively. [9],[10],[11].

Compared to traditional anomaly detection techniques—such as statistical models, clustering algorithms, and distance-based approaches—deep learning models offer several substantial improvements. Traditional methods often rely on manual feature engineering and strong assumptions about data distribution, which limit their flexibility and effectiveness in complex, high-dimensional, or unstructured datasets. In contrast, deep learning models, such as autoencoders, CNNs, and RNNs, automatically learn relevant features from raw data, allowing for more accurate and robust detection of subtle or non-linear anomalies. Deep learning is also more adaptable to diverse data types, including time-series, images, text, and graphs, where traditional models often struggle. Additionally, deep models can scale better with large datasets and leverage transfer learning or self-supervised learning to improve performance in low-label settings. These capabilities make deep learning a powerful and increasingly essential approach to modern anomaly detection.[16],[17],[18],[19].

III. LITERATURE REVIEW

In recent years, a diverse array of deep learning techniques has been proposed to enhance the effectiveness of anomaly detection. This section reviews key contributions in the literature, organizing them based on the type of neural network architecture employed and the specific application domains addressed.

Guo et al. [22] proposed an energy-efficient anomaly detection method for IoT multivariate time series data using a Graph Neural Network (EGNN). They evaluated their model on four real-world datasets—SWaT, PSM, MSL, and SMAP spanning water treatment, server monitoring, and astronautical systems. EGNN employs graph attention-based forecasting and deviation scoring to identify anomalies while minimizing energy consumption. The experimental results demonstrated that EGNN reduces energy usage significantly compared to baseline methods, though it sometimes exhibits slightly lower F1-scores compared to state-of-the-art approaches like OmniAnomaly and GDN. However, its limitations include dependency on accurate graph structure learning, restricted evaluation across diverse IoT domains, and scalability challenges in large-scale deployments.

Contreras-Cruz et al. [23] introduced a Generative Adversarial Network (GAN)-based approach for anomaly detection in aerial images, utilizing the fast Anomaly Generative Adversarial Network (f-AnoGAN). The study employed two datasets which are urban and rural space image sets, where anomalies were human-made structures such as buildings and roads. The f-AnoGAN model outperformed other methods, achieving AUC scores of 0.99 and 0.92 for urban and rural datasets, respectively. However, the approach faces challenges in handling lower contrast anomalies, dependency on high-quality normal data, and computational cost during training.

Iqbal and Amin [24] explored deep learning-based approaches for time series forecasting and anomaly detection, focusing on credit card fraud detection. They utilized benchmark datasets, including the Numenta Anomaly Benchmark (NAB) corpus and a credit card fraud detection dataset, to evaluate various models such as LSTM, Autoencoder, GAN, Transformer, and ensemble methods. The results showed that the LSTM-Autoencoder, GAN, and ensemble models performed best, achieving high accuracy and AUC-ROC scores. However, limitations include the sensitivity to data imbalance, computational complexity, and the need for extensive hyperparameter tuning.

Khan and Haroon [25] proposed an unsupervised deep learning ensemble model for anomaly detection in static attributed social networks. They evaluated their approach on the BlogCatalog and Flickr datasets, which contain network structures and node attributes relevant to social media interactions. The ensemble model integrates Autoencoders (AEs), Variational Autoencoders (VAEs), and Generative Adversarial Networks (GANs), leveraging a novel weighted averaging mechanism to improve anomaly detection. Experimental results showed that their ensemble model outperformed individual baseline methods such as DOMINANT, VGAE, and EfficientGAN, achieving the highest AUC scores (0.8503 for BlogCatalog and 0.8418 for Flickr). However, limitations include the need for further evaluation of low-dimensional datasets and challenges in extending the model to dynamic network environments.

Adiban et al. [26] introduced STEP-GAN, a novel multi-generator Generative Adversarial Network (GAN) approach for anomaly detection, specifically targeting cybersecurity threats. The model was evaluated on the ICS (Industrial Control System) and UNSW-NB15 datasets, both highly imbalanced, containing significantly more normal samples than anomalies. STEP-GAN leverages multiple generators in a stepwise interaction with a discriminator to capture different data distribution modes, effectively mitigating the mode collapse issue common in GAN-based methods. Experimental results demonstrated that STEP-GAN outperformed state-of-the-art approaches in terms of accuracy and F-measure across both datasets. However, the method's limitations include potential sensitivity to shifts in the normal data distribution over time and challenges in defining optimal hyperparameters for different domains.

Li et al. [27] proposed a Controlled Graph Neural Network (ConGNN) with a denoising diffusion probabilistic model (DDPM) for anomaly detection in attributed networks. They evaluated their approach on five benchmark datasets, Cora, Citeseer, PubMed, Photo, and Computer, demonstrating its effectiveness in handling label scarcity. ConGNN generates augmented data by injecting reference node characteristics into source nodes, enhancing network anomaly detection. The model outperformed state-of-the-art baselines in AUC and precision metrics. However, limitations include sensitivity to the choice of reference nodes, computational overhead from the diffusion model, and challenges in extending the method to dynamic graphs.

Kopčan et al. [28] developed an anomaly detection framework using Adversarial Autoencoders (AAE) and Deep Convolutional Generative Adversarial Networks (DCGAN) for autonomous transportation systems. Their approach was tested on the MNIST, Fashion-MNIST, and CIFAR-10 image datasets to evaluate the models' ability to detect anomalies in visual data. The study introduced an optimal decision threshold using cumulative and reverse cumulative distribution functions, leading to anomaly detection errors of 0.08% for AAE and 1.89% for DCGAN on the MNIST dataset. However, the method's effectiveness decreases when applied to more complex datasets like CIFAR-10, limiting its generalizability to real-world autonomous transport scenarios.

Sevyeri and Fevens [29] introduced AD-CGAN, a Contrastive Generative Adversarial Network for anomaly detection, designed to address mode collapse and instability

issues in GAN-based models. They evaluated AD-CGAN on four benchmark datasets—CIFAR-10, FashionMNIST, MNIST, and CatsVsDogs using both one-vs-all and all-vs-one anomaly detection schemes. The model integrates contrastive learning with GANs and autoencoders, enhancing feature discrimination and reconstruction accuracy. Experimental results demonstrated that AD-CGAN outperformed state-of-the-art anomaly detection methods, achieving significant improvements in ROC-AUC scores. However, the approach faces challenges when handling multi-modal normal distributions, leading to reduced performance in complex datasets

Yu et al. proposed [30] a graph-based anomaly detection framework utilizing an attention mechanism to enhance feature extraction and anomaly classification. Their approach was evaluated on multiple datasets, including WebKB, Cora, Citeseer, and a power grid dataset, where anomalies were artificially injected for benchmarking. The method integrates deep neural networks with graph representation learning, incorporating a graph attention module and an optimal transport-based classifier to improve anomaly detection performance. Experimental results demonstrated superior recall rates compared to baseline methods, particularly in structured data domains. However, limitations include increased memory consumption and reduced efficiency on large-scale graphs.

Yang et al. [31] introduced AnoTrans, a Transformer-based Generative Adversarial Network (GAN) designed to enhance anomaly detection by capturing both long-range dependencies and local details in image data. They evaluated their model on four benchmark datasets: CIFAR10, STL10, LBOT, and MVTecAD demonstrating superior performance over state-of-the-art CNN-based methods. AnoTrans employs a U-Net-inspired generator with self-attention mechanisms and a novel skip attention connection (SAC) to improve feature representation. Experimental results showed that AnoTrans outperformed SAGAN by over 3% in AUC scores, particularly excelling in detecting subtle anomalies. However, the method's limitations include increased computational complexity and potential instability in GAN training.

Ning et al. proposed [32] MST-GNN, a Multi-Scale Temporal-Enhanced Graph Neural Network for anomaly detection in multivariate time series. The model was evaluated on three real-world datasets: MSL, SWaT, and WADI capturing anomalies in industrial and sensor-based environments. MST-GNN integrates shapelets learning, a recurrent-skip neural network, and raw time series data to enhance temporal feature representation, while a graph attention network captures dependencies among multivariate time series. Experimental results demonstrated superior F1-scores, and recall compared to baseline methods, particularly excelling in detecting subtle anomalies. However, limitations include computational complexity and sensitivity to hyperparameter tuning.

Hassan et al. [33] proposed a real-time anomaly detection framework for network traffic using Graph Neural Networks (GNNs) and Random Forest models. Their study utilized the Hornet dataset, which consists of network traffic data collected from honeypots deployed in eight global locations. The approach leverages GNNs for graph-based anomaly detection and Random Forest for feature-based classification, comparing their effectiveness. Experimental results showed that the Random Forest model outperformed GNNs, achieving near-perfect accuracy of 99%, while GNNs struggled with dynamic graphs and class imbalance. However, GNNs exhibited potential for capturing structural anomalies. Limitations include the need for further optimization of GNN-based approaches and scalability challenges.

Ounasser et al. [34] conducted a comparative study on unsupervised anomaly detection using generative models and autoencoders, evaluating various deep learning approaches across seven datasets, including KDDCup99, Credit Cards, and WDBC. The study focused on DAGMM, SO-GAAL, and MO-GAAL models, demonstrating their superiority over traditional machine learning methods such as Isolation Forest and LOF. DAGMM achieved up to a 14% improvement in F1-score, particularly excelling in high-dimensional and contaminated datasets. However, the models face challenges in mode collapse (for GANs), sensitivity to contamination rates, and computational resource requirements.

Sharma et al. [35] introduced Inspection-L, a Graph Neural Network (GNN)-based anomaly detection framework for fraudulent transaction identification in blockchain networks. They evaluated their model using the Elliptic dataset, a large-scale Bitcoin transaction dataset, leveraging self-supervised Deep Graph Infomax (DGI) and a Graph Isomorphism Network (GIN), combined with a supervised Random Forest classifier. Experimental results demonstrated that Inspection-L outperformed traditional machine learning models in detecting illicit transactions, achieving superior recall and F1-scores. However, the method faces limitations in scalability, computational cost, and potential sensitivity to evolving blockchain fraud tactics.

Chen et al. [36] proposed a Dual Auto-Encoder GAN-based anomaly detection model (DAGAN) for industrial control systems (ICS), addressing challenges such as the long-tailed distribution of data and the difficulty of obtaining abnormal samples. The DAGAN model employs an "encoder-decoder-encoder" architecture to learn the latent and marginal distributions of normal data without requiring any abnormal samples for training. A parameter-free dynamic strategy was introduced to robustly learn the marginal distribution, reducing the misjudgment of marginal samples. The model was evaluated on the DS2OS and SWaT datasets, achieving high accuracy and outperforming other state-of-the-art methods, including GANomaly and FenceGAN. The DAGAN model demonstrated superior performance in terms of accuracy, recall, and F1-score, particularly for hard-to-identify attacks. However, the model's performance could be further improved

for multi-class anomaly detection, which was identified as a future research direction.

Xu et al. [37] introduced TGAN-AD, a Transformer-based Generative Adversarial Network (GAN) model for anomaly detection in time series data. The model was evaluated using three public datasets: SWaT, WADI, and KDDCUP99. TGAN-AD employs a Transformer-based generator to capture contextual patterns in time series data and a discriminator to assist in detecting anomalies. Experimental results demonstrated that TGAN-AD outperformed state-of-the-art anomaly detection methods, achieving the highest Recall and F1-Score across all datasets. However, the study highlighted the model's sensitivity to hyperparameter tuning, particularly the sliding window size and Transformer layers, which could impact detection performance.

Lian et al. [38] introduced a digital twin-driven anomaly detection method for oil and gas stations using MTAD-GAN, combining knowledge graph attention and temporal Hawkes attention to enhance spatio-temporal correlations. Tested on datasets like KDD99, SWaT, and WADI, the method achieved significant improvements in accuracy, precision, F1 score, and AUC-ROC, outperforming traditional and deep learning approaches. However, challenges remain in handling imbalanced data and exploring complex time series relationships for broader industrial applications.

Daniel et al. [39] proposed AnomEn, a robust graph neural network encoder for anomaly detection in attributed networks. The model was tested on multiple datasets, including Twitter, Enron, and Amazon, for node anomaly detection, and PolitiFact and GossipCop for edge anomaly detection. AnomEn introduces a weighted aggregation mechanism that balances node features and neighborhood influences, addressing challenges of false positives and false negatives in Graph Neural Networks (GNNs). Experimental results showed that AnomEn outperformed existing methods, improving node anomaly detection by 5.63% and edge anomaly detection by 7.87%. However, the study noted that parameter tuning, particularly the weighting factor in aggregation, impacts performance.

Emane et al. [40] proposed an anomaly detection framework that integrates Graph Convolutional Networks (GCNs) with Density-Based Spatial Clustering of Applications with Noise (DBSCAN) for large-scale graph data. The model was tested on datasets such as YelpChi, Amazon, and ACM, where it demonstrated improved accuracy and robustness in identifying anomalous nodes. GCNs were used to generate expressive node embeddings, which were then clustered using DBSCAN, leveraging a novel heuristic for automatic hyperparameter tuning. The results showed that this approach outperformed traditional methods like K-means and OPTICS, as well as state-of-the-art models like RioGNN and CARE-GNN. However, the study noted that the model's performance is influenced by hyperparameter selection, particularly the number of GCN layers and DBSCAN's minimum points threshold.

Nakao et al. [41] proposed an unsupervised deep anomaly detection method for chest radiographs using a Variational Autoencoder-Generative Adversarial Network (VAE-GAN). The model was trained on 29,684 frontal chest radiographs from the Radiological Society of North America (RSNA) Pneumonia Detection Challenge dataset, using only normal images for training. The VAE component captured normal image distributions, while the GAN improved reconstruction quality. The model achieved an AUROC of 0.752 for detecting abnormal images, with higher performance for the "Opacity" class (0.838) compared to "No Opacity/Not Normal" (0.704). Although the method successfully detected various abnormalities, including lung masses and pleural effusion, its performance was limited by the inherent challenges of unsupervised learning, such as difficulty in precisely diagnosing specific conditions.

Xu et al. [42] introduced a two-stage anomaly detection model based on Generative Adversarial Networks (GANs) to enhance detection accuracy, particularly for small and positive samples. The approach was tested on the liver CT image dataset and the CIFAR10 public dataset. The first stage extracts multi-scale image features using convolutional neural networks, while the second stage employs an anomaly detection GAN with an Attention Gate mechanism to improve reconstruction quality. The proposed model achieved superior results, with an 8.8% improvement on the liver CT dataset and a 19.2% increase on CIFAR10 compared to the skip-GANomaly baseline. However, the study noted that the method's effectiveness depends on feature extraction quality, and further improvements in handling complex abnormalities are needed.

Ko et al. [43] proposed an anomaly detection method using Graph Neural Networks (GNNs) to analyze feature correlations in network data, focusing on real-time detection of Distributed Denial-of-Service (DDoS) attacks. The study utilized the Coburg Intrusion Detection Dataset (CIDDS) and KDDCup datasets, applying GNNs to trace feature interrelationships and identify anomaly signals. The method achieved high accuracy rates of 94.5% for KDDCup and 98.85% for CIDDS, demonstrating its effectiveness in detecting network anomalies. However, the study acknowledged limitations in handling real-time detection for highly dynamic network environments and the need for further optimization of feature selection to improve computational efficiency. The authors highlighted the potential of GNNs in enhancing real-time anomaly detection but emphasized the challenges in scaling the approach for broader network applications.

Luo et al. [44] proposed an anomaly detection model that combines Generative Adversarial Networks (GANs) with Convolutional Autoencoders (CAEs) to enhance feature extraction for time-series data. The model was evaluated on ECG and 2D gesture datasets, leveraging a modified version of the Unsupervised Anomaly Detection (USAD) architecture to

improve stability during adversarial training. Additionally, the study introduced an Exponential Weighted Moving Average (EWMA) method to smooth reconstruction errors and reduce false positives. Experimental results showed an improvement of 0.028% in AUROC, 0.233% in AUPRC, and 0.187% in F1-score compared to existing methods. However, the model requires manual threshold tuning for optimal detection, limiting its adaptability to real-time applications.

Park et al. [45] proposed an unsupervised anomaly detection method for breast cancer screening using StyleGAN2 to generate synthetic mammograms from 105,948 normal images. Evaluated on 50 cancer and 50 normal mammograms, the method achieved an AUC of 70.0%, sensitivity of 78.0%, and specificity of 52.0%. While the synthetic images showed high fidelity (FID score: 4.383), limitations included noise-like artifacts, reliance on craniocaudal views, and insufficient performance for clinical use. Future improvements, such as higher-resolution images and additional views, were suggested to enhance accuracy.

Benaddi et al. [46] proposed a hybrid anomaly detection model combining Distributional Reinforcement Learning (DRL) and Generative Adversarial Networks (GANs) to enhance cybersecurity in Industrial Internet of Things (IIoT) networks. The model was evaluated using the DS2OS dataset, which contains various IIoT attack types such as Denial of Service, malicious control, and data probing. The GAN component was used for data augmentation to address class imbalance, while DRL improved detection by modeling the probability distribution of anomalous events. Experimental results demonstrated that the DRL-GAN model outperformed standard DRL in both binary and multi-class classification, achieving higher accuracy, precision, recall, and F1-score. However, the study noted that the approach requires high computational resources and manual fine-tuning of GAN training for optimal results.

Duan et al. [47] proposed a log anomaly detection method, GAN-EDC, based on Generative Adversarial Networks (GANs), utilizing an Encoder-Decoder framework with Long Short-Term Memory (LSTM) as the generator and Convolutional Neural Networks (CNN) as the discriminator. The method was evaluated on real-world log datasets, including HDFS and BGL, achieving an average precision of 95% for detecting log point anomalies. The generator maps log keywords to templates, while the discriminator distinguishes between real and generated templates, with anomaly detection performed using Euclidean distance. GAN-EDC outperformed traditional methods like clustering, SVM, and decision trees in accuracy and efficiency. However, the study acknowledged limitations in handling large-scale datasets and the need for manual parameter tuning, particularly for the threshold k*k*. Future work suggested incorporating reinforcement learning for automated parameter optimization.

Deep learning (DL) has emerged as a powerful tool for anomaly detection across various domains, offering enhanced capabilities over traditional methods. In the realm of cybersecurity, DL models such as Long Short-Term Memory (LSTM) networks and hybrid architecture have demonstrated high accuracy in detecting network intrusions and malicious activities within Internet of Things (IoT) environments. For instance, a study achieved a detection accuracy of up to 99.9% using LSTM-based models for IoT security applications [48].

In healthcare, DL techniques have been instrumental in early disease detection and monitoring. A notable study applied DL-based models to breast ultrasonography, achieving significant improvements in anomaly detection accuracy. Similarly, in intensive care units, DL models have been utilized for automated anomaly detection in EEG signals, enhancing patient monitoring and care [49].

The industrial IoT sector has also benefited from DL-based anomaly detection. A real-time deep anomaly detection framework combining Convolutional Neural Networks (CNNs) and LSTM networks was developed to monitor multivariate time-series data, effectively identifying anomalies in industrial processes. Additionally, a memory-efficient DL model named TinyAD was proposed to facilitate anomaly detection on resource-constrained IIoT devices, demonstrating reduced memory consumption with negligible computational overhead [50].

In the financial domain, DL models have revolutionized anomaly detection by enabling real-time fraud detection and risk management. These models analyze transaction data to identify fraudulent activities and assess financial statements for inconsistencies, thereby aiding in the prevention of financial crimes [51].

This review goes beyond summarizing prior work by critically examining the strengths, limitations, and assumptions of existing deep learning-based anomaly detection models. Through this analysis, we identify specific gaps in literature such as the lack of interpretability, insufficient real-time evaluation, and limited generalizability across domains that remain underexplored. These gaps highlight the need for more robust, explainable, and adaptable approaches. The organization of this review, including our taxonomy and comparative analysis, is designed to underscore these research deficiencies and provide a foundation for future studies aimed at addressing them.

Table 1 presents a comprehensive comparison of the papers mentioned in the literature review section. The comparison made based on many aspects such as dataset, techniques, finding and limitations.

TABLE I
COMPREHENSIVE ANALYSIS OF LITERATURE REVIEW PAPERS

Paper	Dataset	Techniques	Results	Limitations
(Guo et al.) [22]	SWaT, PSM, MSL, SMAP	Graph Neural Network (EGNN)	Energy-efficient anomaly detection with reduced power consumption	Dependency on graph structure accuracy, scalability issues
(Contreras-Cruz et al.) [23]	Urban & rural aerial image datasets	f-AnoGAN (fast AnoGAN)	AUC 0.99 (urban), 0.92 (rural)	Struggles with low-contrast anomalies, high training cost
(Iqbal & Amin) [24]	NAB, credit card fraud dataset	LSTM, Autoencoder, GAN, Transformer, ensemble methods	LSTM-AE, GAN, and ensemble models achieved high accuracy GAN results are accuracy 1.000, precision 1.000, recall 1.000, F1 1.000 and R2 Score 0.978	Sensitivity to class imbalance, hyperparameter tuning complexity
(Khan & Haroon) [25]	BlogCatalog, Flickr	Ensemble (AE, VAE, GAN)	AUC 0.8503 (BlogCatalog), 0.8418 (Flickr)	Limited evaluation on low-dimensional data, not tested on dynamic networks
(Adiban et al.) [26]	ICS, UNSW-NB15	STEP-GAN (multi-generator GAN)	Higher accuracy & F-measure than baselines. Metrics on UNSW-NB15 Dataset Accuracy 97.24 and F1 0.9644	Sensitive to data shifts, complex hyperparameter tuning
(Li et al.) [27]	Cora, Citeseer, PubMed, Photo, Computer	ConGNN + DDPM	Higher AUC, precision in label-scarce settings. Metrics on computer, photo and cora datasets are AUC 0.983, 0.85 and 0.881 respectively.	Sensitivity to reference node choice, high computational cost
(Kopčan et al.) [28]	MNIST, Fashion-MNIST, CIFAR-10	Adversarial Autoencoder (AAE), DCGAN	Anomaly detection error 0.08% (AAE) and 1.89% (DCGAN).	Decreased performance on complex datasets like CIFAR-10
(Sevyeri & Fevens) [29]	CIFAR-10, FashionMNIST, MNIST, CatsVsDogs	AD-CGAN (Contrastive GAN)	Higher ROC-AUC than baselines. For CIFAR-10 dataset and MNIST AUC 89.8 and 94.6 respectively.	Struggles with multi-modal distributions
(Yu et al.) [30]	WebKB, Cora, Citeseer, power grid dataset	Graph-based attention mechanism	Higher recall than baselines. Recall@L = 20 on core dataset is 0.79	High memory usage, inefficiency in large-scale graphs
(Yang et al.) [31]	CIFAR-10, STL10, LBOT, MVTecAD	AnoTrans (Transformer-GAN)	3% higher AUC than SAGAN	High computational complexity, unstable GAN training
(Ning et al.) [32]	MSL, SWaT, WADI	MST-GNN (Multi-Scale Temporal GNN)	Best F1-score, recall on benchmark datasets	Computationally expensive, sensitive to hyperparameters
(Hassan et al.) [33]	Hornet (honeypot network traffic)	GNN + Random Forest	RF: 99% accuracy, GNN struggled with imbalance	GNN scalability issues, optimization needed
(Ounasser et al.) [34]	KDDCup99, Credit Cards, WDBC	DAGMM, SO-GAAL, MO-GAAL	DAGMM improved F1-score by 14%	Mode collapse in GANs, sensitivity to contamination rates
(Sharma et al.) [35]	Elliptic (Bitcoin transactions)	GNN (DGI, GIN) + RF	Higher recall & F1-score than ML models. Inspection-L AF+DNE precision 0.972. recall 0.72 and F1 0.8282	Scalability issues, sensitive to evolving fraud patterns
Chen et al. [36]	DS2OS, SWaT	Dual Auto-Encoder GAN (DAGAN)	Improved anomaly detection accuracy over five baseline models. the average values of the ACC, Rec and F1 are greater than 0.82 on SWaT dataset.	Effectiveness depends on accurate discriminator-based marginal sample selection
(Xu et al.) [37]	SWaT, WADI, KDDCUP99	TGAN-AD (Transformer-based GAN)	Highest Recall and F1-Score across all datasets. One SWaT data set F1 0.953, recall 0.99 and precision 91.8.	Sensitive to hyperparameter tuning (sliding window size, Transformer layers)
Paper 20 (Lian et al.) [38]	KDD99, SWaT, WADI, J10031, SKAB, DAMADICS, MSL, SMAP, SMD	MTAD-GAN (GAN + Digital Twin + Knowledge Graph + Temporal Hawkes Attention)	Accuracy improved by 2.6% over TenED algorithm, best precision, F1-score, AUC-ROC	Struggles with highly imbalanced data, needs further exploration of implicit time-series relationships

Paper	Dataset	Techniques	Results	Limitations
(Daniel et al.) [39]	Twitter, Enron, Amazon (node anomaly detection); PolitiFact, GossipCop (edge anomaly detection)	AnomEn (GNN-based encoder with weighted aggregation)	Improved node anomaly detection by 5.63%, and edge anomaly detection by 7.87%. For the Gossipcop dataset achieved the highest accuracy score of 0.9798 and AUC score of 0.9796.	Performance affected by parameter tuning, particularly weighting factor in aggregation
(Emane et al.) [40]	YelpChi, Amazon, ACM	GCN + DBSCAN (Graph Convolutional Networks + Clustering)	Outperformed K-means, OPTICS, RioGNN, and CARE-GNN in node anomaly detection. F1 on Amazon, Yelp and ACM datasets are 0.95,0.92 and 0.98 respectively.	Sensitive to hyperparameter selection (GCN layers, DBSCAN min points threshold)
(Nakao et al.) [41]	RSNA Pneumonia Detection Challenge (chest radiographs)	VAE-GAN (Variational Autoencoder + Generative Adversarial Network)	AUROC 0.752 (overall), 0.838 (Opacity class), 0.704 (No Opacity/Not Normal)	Limited by unsupervised learning, struggles with precise diagnosis of specific conditions
(Ko et al.) [43]	CIDDS, KDDCup	GNN-based anomaly detection for network traffic	94.5% accuracy (KDDCup), 98.85% accuracy (CIDDS)	Limited real-time performance in dynamic networks, needs optimization for feature selection
(Luo et al.) [44]	ECG, 2D gesture datasets	GAN + Convolutional Autoencoder (USAD-based) + EWMA	AUROC +0.028%, AUPRC +0.233%, F1-score +0.187% vs. baseline	Requires manual threshold tuning, limiting adaptability to real-time applications
(Park et al.) [45]	Mammography dataset (105,948 normal, 100 test cases)	StyleGAN2 (synthetic image generation for anomaly detection)	AUC 70.0%, Sensitivity 78.0%, Specificity 52.0%, FID 4.383	Limited to craniocaudal views, synthetic artifacts, insufficient clinical accuracy
(Benaddi et al.) [46]	DS2OS (IIoT cybersecurity dataset)	DRL-GAN (Distributional Reinforcement Learning + GAN)	Improved accuracy, precision, recall, F1-score over standard DRL. On normal DRL f1 score 99.22, precision 99.5 and accuracy 98.955	High computational cost, manual fine-tuning required for GAN training
(Duan et al.) [47]	HDFS, BGL (log datasets)	GAN-EDC (GAN + LSTM encoder-decoder + CNN discriminator)	In BGL dataset with Precision 96%, recall 89% and F-measure 92%.	Struggles with large-scale logs, manual parameter tuning needed

Despite the promising performance of deep learning techniques in anomaly detection, several limitations hinder their widespread applicability and generalization. One major challenge is the sensitivity to hyperparameter tuning. Many models, such as autoencoders and GANs, require careful calibration of parameters like learning rates, architecture depth, and threshold settings for anomaly scores, which can significantly impact performance. This makes replication and real-world deployment more difficult.

Another critical limitation is computational complexity. Deep learning models often involve high training and inference costs, making them less suitable for real-time or resource-constrained environments such as IoT devices or edge computing platforms. Furthermore, training deep models on high-dimensional data may require substantial GPU resources and time, which can limit accessibility for researchers or organizations with fewer computational resources.

Imbalanced data remains a persistent issue in anomaly detection, where anomalies are inherently rare. While techniques like oversampling or synthetic anomaly generation can help, they may also introduce noise or fail to represent realistic deviations. This imbalance can lead to biased models that overfit the majority class.

Moreover, lack of interpretability is a significant concern, particularly in safety-critical applications. Most deep learning models act as black boxes, offering little insight into the rationale behind anomaly predictions. This undermines trust and makes debugging or refining models difficult.

Finally, many existing approaches lack generalization across domains. A model trained in one type of data (e.g., network traffic) often fails when applied to another (e.g., medical records), highlighting the need for domain adaptation and transfer learning strategies.

Future research should focus on developing explainable and interpretable models to enhance trust, especially in critical areas like healthcare and finance. Approaches like self-supervised and few-shot learning can address the challenge of limited labeled anomaly data. Improving cross-domain generalization through transfer learning is also vital, as is creating lightweight models for real-time, edge-device deployment. Standardizing

benchmarks and evaluation metrics will help ensure fair comparisons. Additionally, future systems should provide actionable explanations alongside anomaly detection. Finally, ethical considerations such as privacy, fairness, and the impact of false detections must be addressed to ensure responsible use.

IV. CONCLUSION

In this review, we have provided a comprehensive overview of deep learning techniques applied to anomaly detection across various domains. We discussed the strengths and limitations of popular architectures such as autoencoders, generative adversarial networks (GANs), recurrent neural networks (RNNs), and convolutional neural networks (CNNs), while also highlighting benchmark datasets and evaluation metrics commonly used in the literature.

Looking forward, several promising research directions are emerging in this space. One key area is the development of explainable and interpretable deep learning models, which are crucial for real-world deployment, especially in sensitive domains like healthcare and finance. Another growing trend is the integration of multimodal data (e.g., combining images, text, and sensor data) to enhance anomaly detection performance in complex environments.

In addition, there is increasing interest in self-supervised and few-shot learning approaches, which aim to reduce the dependency on large, labeled datasets, an ongoing challenge in anomaly detection. The use of foundation models and pre-trained architectures tailored for anomaly detection is also gaining traction.

Moreover, the rise of edge computing calls for lightweight and efficient deep learning models capable of performing real-time anomaly detection with limited computational resources. Future research could explore model compression techniques and hardware-aware designs to address these constraints.

Lastly, collaborative and federated learning frameworks offer potential for privacy-preserving anomaly detection across distributed systems without the need to centralize data, a critical aspect for domains such as cybersecurity and medical diagnostics.

REFERENCES

:

[1] W. Hassan, S. E. Hosseini, and S. Pervez, “Real-Time Anomaly Detection in Network Traffic Using Graph Neural Networks and Random Forest,” *Lecture notes in computer science*, pp. 194–207, Jan. 2024, doi: https://doi.org/10.1007/978-3-031-60994-7_16.

[2] J. Kopčan, O. Škvarek, and M. Klimo, “Anomaly detection using Autoencoders and Deep Convolution Generative Adversarial Networks,” *Transportation Research Procedia*, vol. 55, pp. 1296–1303, 2021, doi: https://doi.org/10.1016/j.trpro.2021.07.113.

[3] X. Li, C. Xiao, Z. Feng, S. Pang, W. Tai, and F. Zhou, “Controlled graph neural networks with denoising diffusion for anomaly detection,” *Expert Systems with Applications*, vol. 237, p. 121533, Mar. 2024, doi: https://doi.org/10.1016/j.eswa.2023.121533.

[4] Y. Wang, X. Li, and J. Han, "Hybrid Models for Real-Time Anomaly Detection in Edge Computing Environments," *IEEE Transactions on Industrial Informatics*, vol. 19, no. 5, pp. 2307-2318, 2023.

[5] R. Chalapathy and S. Chawla, "Deep Learning for Anomaly Detection: A Survey," *ACM Computing Surveys*, vol. 55, no. 3, pp. 1-38, 2022.

[6] B. Liu, X. Zhang, and X. Chen, "An Improved Isolation Forest for High-Dimensional Data Anomaly Detection," *Knowledge-Based Systems*, vol. 220, 106951, 2021.

[7] “Improving Performance of Autoencoder-Based Network Anomaly Detection on NSL-KDD Dataset | IEEE Journals & Magazine | IEEE Xplore,” *ieeexplore.ieee.org*. https://ieeexplore.ieee.org/abstract/document/9552882

[8] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly Detection: A Survey," *ACM Computing Surveys (CSUR)*, vol. 41, no. 3, pp. 1-58, 2009.

[9] I. Alam, “Detecting Outliers Using Z-score — part -2 - Irshad Alam - Medium,” *Medium*, Oct. 02, 2024. https://medium.com/@irshadalamtech/detecting-outliers-using-z-score-part-2-85297f54ea84

[10] M. A. Pimentel, D. A. Clifton, L. Clifton, and L. Tarassenko, "A Review of Novelty Detection," *Signal Processing*, vol. 99, pp. 215-249, 2014.

[11] D. Mohamed, Ayman El-Kilany, and Hoda, “A Hybrid Model for Documents Representation,” vol. 12, no. 3, Jan. 2021, doi: https://doi.org/10.14569/ijacsa.2021.0120339.

[12] B. L. Bars, “Détection d’événements et inférence de structure pour des vecteurs sur graphes,” Jan. 29, 2021.

[13] Joeri Lenaerts, H. Pinson, and V. Ginis, “Data driven design of optical resonators,” Dec. 21, 2021.

[14] K. Qin, Q. Wang, B. Lu, H. Sun, and P. Shu, “Flight Anomaly Detection via a Deep Hybrid Model,” *Aerospace*, vol. 9, no. 6, pp. 329–329, Jun. 2022, doi: https://doi.org/10.3390/aerospace9060329.

[15] V. J. Hodge and J. Austin, "A Survey of Outlier Detection Methodologies," *Artificial Intelligence Review*, vol. 22, no. 2, pp. 85-126, 2004.

[16] M. Sakurada and T. Yairi, "Anomaly Detection Using Autoencoders with Nonlinear Dimensionality Reduction," in *Proceedings of the MLSDA 2014 2nd Workshop on Machine Learning for Sensory Data Analysis*, pp. 4-11, 2014.

[17] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *arXiv preprint arXiv:1312.6114*, 2013.

[18] P. Malhotra, L. Vig, G. Shroff, and P. Agarwal, "Long Short-Term Memory Networks for Anomaly Detection in Time Series," in *Proceedings of the 23rd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, pp. 89-94, 2015.

[19] M. Sabokrou, M. Fathy, M. Hoseini, and R. Klette, "Deep-CNN-Based Approach for Anomaly Detection in Video Surveillance," *IEEE Transactions on Image Processing*, vol. 27, no. 5, pp. 2515-2527, 2018.

[20] C. Ho and M. Welling, "Denoising Score Matching with Annealed Langevin Sampling," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 2726-2735, 2020.

[21] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A Survey on Bias and Fairness in Machine Learning," *ACM Computing Surveys (CSUR)*, vol. 54, no. 6, pp. 1-35, 2021.

[22] H. Guo, Z. Zhou, D. Zhao, and W. Gaaloul, "EGNN: Energy-efficient anomaly detection for IoT multivariate time series data using graph neural network," *Future Generation Computer Systems*, vol. 151, pp. 45-56, Feb. 2024, doi: 10.1016/j.future.2023.09.028.

[23] M. A. Contreras-Cruz, F. E. Correa-Tome, R. López-Padilla, and J.-P. Ramirez-Paredes, "Generative Adversarial Networks for anomaly detection in aerial images," *Computers & Electrical Engineering*, vol. 106, p. 108470, Mar. 2023, doi: 10.1016/j.compeleceng.2022.108470.

[24] A. Iqbal and R. Amin, "Time Series Forecasting and Anomaly Detection Using Deep Learning," *Computers & Chemical Engineering*, vol. 182, p. 108560, Mar. 2024, doi: 10.1016/j.compchemeng.2023.108560.

[25] W. Khan and M. Haroon, "An unsupervised deep learning ensemble model for anomaly detection in static attributed social networks," *International Journal of Cognitive Computing in Engineering*, vol. 3, pp. 153-160, Jun. 2022, doi: 10.1016/j.ijcce.2022.08.002.

[26] M. Adiban, S. M. Siniscalchi, and G. Salvi, "A step-by-step training method for multi generator GANs with application to anomaly detection and cybersecurity," *Neurocomputing*, vol. 537, pp. 296-308, Jun. 2023, doi: 10.1016/j.neucom.2023.03.056.

[27] X. Li, C. Xiao, Z. Feng, S. Pang, W. Tai, and F. Zhou, "Controlled graph neural networks with denoising diffusion for anomaly detection," *Expert Systems with Applications*, vol. 237, p. 121533, Mar. 2024, doi: 10.1016/j.eswa.2023.121533.

[28] J. Kopčan, O. Škvarek, and M. Klimo, "Anomaly detection using Autoencoders and Deep Convolution Generative Adversarial Networks," *Transportation Research Procedia*, vol. 55, pp. 1296-1303, 2021, doi: 10.1016/j.trpro.2021.07.113.

[29] L. R. Sevyeri and T. Fevens, "AD-CGAN: Contrastive Generative Adversarial Network for Anomaly Detection," *Lecture Notes in Computer Science*, pp. 322-334, Jan. 2022, doi: 10.1007/978-3-031-06427-2_27.

[30] Y. Yu, Z. Zhiyong, B. Jin, G. Wu, and C. Dong, "Graph-Based Anomaly Detection via Attention Mechanism," pp. 401-411, Jan. 2022, doi: 10.1007/978-3-031-13870-6_33.

[31] C. Yang *et al.*, "A Transformer-Based GAN for Anomaly Detection," *Lecture Notes in Computer Science*, pp. 345-357, 2022, doi: 10.1007/978-3-031-15931-2_29.

[32] Z. Ning, Z. Jiang, H. Miao, and L. Wang, "MST-GNN: A Multi-scale Temporal-Enhanced Graph Neural Network for Anomaly Detection in Multivariate Time Series," *Lecture Notes in Computer Science*, pp. 382-390, Jan. 2023, doi: 10.1007/978-3-031-25158-0_29.

[33] W. Hassan, S. E. Hosseini, and S. Pervez, "Real-Time Anomaly Detection in Network Traffic Using Graph Neural Networks and Random Forest," *Lecture Notes in Computer Science*, pp. 194-207, Jan. 2024, doi: 10.1007/978-3-031-60994-7_16.

[34] N. Ounasser, M. Rhanoui, M. Mikram, and B. El Asri, "Generative and Autoencoder Models for Large-Scale Multivariate Unsupervised Anomaly Detection," *Smart Innovation, Systems and Technologies*, pp. 45-58, Oct. 2021, doi: 10.1007/978-981-16-3637-0_4.

[35] A. Sharma, P. K. Singh, E. Podoplelova, V. Gavrilenko, A. Tselykh, and A. Bozhenyuk, "Graph Neural Network-Based Anomaly Detection in Blockchain Network," *Lecture Notes in Networks and Systems*, pp. 909-925, Jan. 2023, doi: 10.1007/978-981-99-1479-1_67.

[36] L. Chen, Y. Li, X. Deng, Z. Liu, M. Lv, and H. Zhang, "Dual Auto-Encoder GAN-Based Anomaly Detection for Industrial Control System," *Applied Sciences*, vol. 12, no. 10, p. 4986, May 2022, doi: 10.3390/app12104986.

[37] L. Xu *et al.*, "TGAN-AD: Transformer-Based GAN for Anomaly Detection of Time Series Data," *Applied sciences*, vol. 12, no. 16, pp. 8085–8085, Aug. 2022, doi: https://doi.org/10.3390/app12168085.

[38] Y. Lian, Y. Geng, and T. Tian, "Anomaly Detection Method for Multivariate Time Series Data of Oil and Gas Stations Based on Digital Twin and MTAD-GAN," *Applied Sciences*, vol. 13, no. 3, p. 1891, Jan. 2023, doi: 10.3390/app13031891.

[39] G. V. Daniel, K. Chandrasekaran, M. Venkatesan, and P. Prabhavathy, "Robust Graph Neural-Network-Based Encoder for Node and Edge Deep Anomaly Detection on Attributed Networks," *Electronics*, vol. 12, no. 6, p. 1501, Mar. 2023, doi: 10.3390/electronics12061501.

[40] C. Retiti *et al.*, "Anomaly Detection Based on GCNs and DBSCAN in a Large-Scale Graph," *Electronics*, vol. 13, no. 13, p. 2625, Jul. 2024, doi: 10.3390/electronics13132625.

[41] T. Nakao *et al.*, "Unsupervised Deep Anomaly Detection in Chest Radiographs," *Journal of Digital Imaging*, Feb. 2021, doi: 10.1007/s10278-020-00413-2.

[42] C. Xu, D. Ni, B. Wang, M. Wu, and H. Gan, "Two-stage anomaly detection for positive samples and small samples based on generative adversarial networks," *Multimedia Tools and Applications*, Jan. 2023, doi: 10.1007/s11042-022-14306-9.

[43] H. Ko, I. Praca, and S. G. Choi, "Anomaly detection analysis based on correlation of features in graph neural network," *Multimedia Tools and Applications*, Aug. 2023, doi: 10.1007/s11042-023-15635-z.

[44] X. Luo, Y. Jiang, E. Wang, and X. Men, "Anomaly detection by using a combination of generative adversarial networks and convolutional autoencoders," *EURASIP Journal on Advances in Signal Processing*, vol. 2022, no. 1, Nov. 2022, doi: 10.1186/s13634-022-00943-7.

[45] S. Park, K. H. Lee, B. Ko, and N. Kim, "Unsupervised anomaly detection with generative adversarial networks in mammography," *Scientific Reports*, vol. 13, no. 1, Feb. 2023, doi: 10.1038/s41598-023-29521-z.

[46] H. Benaddi, M. Jouhari, K. Ibrahimi, J. Ben Othman, and E. M. Amhoud, "Anomaly Detection in Industrial IoT Using Distributional Reinforcement Learning and Generative Adversarial Networks," *Sensors*, vol. 22, no. 21, p. 8085, Oct. 2022, doi: 10.3390/s22218085.

[47] X. Duan, S. Ying, W. Yuan, H. Cheng, and X. Yin, "A Generative Adversarial Networks for Log Anomaly Detection," *Computer Systems Science and Engineering*, vol. 37, no. 1, pp. 135-148, 2021, doi: 10.32604/csse.2021.014030.

[48] A. A. H. D. Almowsawi, "Deep Guard-IoT: A Systematic Review of AI-Based Anomaly Detection Frameworks for Next-Generation IoT Security (2020-2024)," *Wasit Journal for Pure Sciences*, vol. 22, no. 1, pp. 1–10, 2024.

[49] C. Yun et al., "A Study on the Effectiveness of Deep Learning-Based Anomaly Detection Methods for Breast Ultrasonography," *Sensors*, vol. 23, no. 5, p. 2864, 2023

[50] H. Nizam et al., "Real-Time Deep Anomaly Detection Framework for Multivariate Time-Series Data in Industrial IoT," *IEEE Sensors Journal*, vol. 22, no. 23, pp. 22812–22821, Dec. 2022.

[51] "Deep Learning Revolutionizes Financial Anomaly Detection," *PyQuant News*, 2023. [Online]. Available: https://www.pyquantnews.com/free-python-resources/deep-learning-revolutionizes-financial-anomaly-detection