Journal Homepage: www.ijarpr.com ISSN: 3049-0103 (Online)

International Journal of Advance Research Publication and Reviews

Vol 02, Issue 04, pp 274-292, April 2025

Developing Autonomous Cyber Defense Systems Using Generative Adversarial Networks for Anomaly Detection in IoT Data Streams

Christianah Gbaja

Independent Researcher, Cyber-Security and AI Engineer, USA

DOI : https://doi.org/10.5281/zenodo.15274354

ABSTRACT

The proliferation of Internet of Things (IoT) devices in critical sectors such as healthcare, energy, and smart infrastructure has significantly expanded the digital attack surface, exposing systems to increasingly sophisticated cyber threats. Traditional rule-based and static anomaly detection methods often fail to identify subtle or zero-day attacks in real time due to the sheer volume and heterogeneity of IoT data streams. This paper proposes a novel approach to cyber defense through the development of autonomous anomaly detection systems powered by Generative Adversarial Networks (GANs), enabling real-time identification of malicious behaviors in high-velocity IoT environments. From a broader perspective, the research contextualizes the limitations of current machine learning-based security models in handling complex, dynamic, and adversarial data patterns common in distributed IoT ecosystems. Narrowing the focus, the study introduces a GAN-based architecture wherein the generator simulates normal IoT traffic patterns, and the discriminator learns to distinguish between authentic and anomalous data behaviors. This adversarial training improves the system's sensitivity to minute deviations without prior knowledge of attack signatures. The framework integrates edge computing nodes to enable low-latency detection while preserving bandwidth and computational resources. Extensive experiments on publicly available IoT datasets demonstrate that the proposed system achieves superior performance in detecting both known and unknown threats, with significantly reduced false positive rates compared to conventional techniques. Moreover, the adaptability of GANs allows the system to continuously learn and evolve in response to new attack vectors, thereby supporting autonomous cyber defense. This paper concludes by discussing deployment strategies, integration with SIEM platforms, and future directions for GAN-based cybersecurity in IoT-centric digital infrastructures.

Keywords: Generative adversarial networks; IoT anomaly detection; Autonomous cyber defense; Edge-based security; Zero-day threat detection; Adversarial learning systems

1. INTRODUCTION

1.1 Background and Motivation

The rapid proliferation of the Internet of Things (IoT) has revolutionized digital ecosystems, enabling a wide array of interconnected devices to communicate and share data autonomously. From smart homes and industrial automation to healthcare monitoring and transportation, IoT has become an integral part of critical infrastructures [1]. However, the very nature of these systems—characterized by heterogeneity, distributed architecture, and constrained resources—makes them particularly susceptible to security breaches and malicious interference [2].

With billions of devices continuously exchanging sensitive information, securing IoT systems has emerged as a pressing global concern. Unlike traditional IT environments, IoT networks often lack consistent security protocols, leaving numerous endpoints exposed to cyber threats such as botnets, malware, spoofing, and data exfiltration [3]. These

challenges are exacerbated by the inherent limitations of many IoT devices, including limited computational power, minimal memory, and lack of regular patching mechanisms [4].

Simultaneously, the cyber threat landscape is evolving rapidly, with attackers adopting more sophisticated and adaptive tactics. Signature-based defense mechanisms are increasingly ineffective against zero-day exploits and polymorphic malware. In this context, anomaly detection becomes a critical line of defense. However, conventional systems struggle to cope with the volume and complexity of data generated by IoT networks [5].

This backdrop underscores the urgent need for **autonomous cyber defense systems**—intelligent mechanisms capable of self-learning, adapting, and responding in real-time to emerging threats. These systems must go beyond static rules and heuristics, leveraging advanced machine learning techniques to detect deviations and anomalies within massive and unstructured data streams [6]. Generative Adversarial Networks (GANs), with their proven ability to model high-dimensional distributions and synthesize realistic data, present a promising solution for robust, autonomous anomaly detection in IoT environments [7].

1.2 Problem Statement

Traditional anomaly detection frameworks, while effective in certain contexts, fall short when deployed in the complex, dynamic, and resource-constrained landscape of IoT networks. Most existing systems rely on predefined rules or statistical thresholds, which are insufficient for identifying novel or context-aware threats that do not conform to expected patterns [8]. This rigidity results in high false positive rates, leading to alert fatigue and inefficient response efforts.

Furthermore, the three defining challenges of IoT data—volume, velocity, and veracity—compound the difficulty of real-time monitoring. The continuous stream of heterogeneous data from various sensors, devices, and gateways overwhelms traditional detection systems not optimized for such scale or diversity [9]. Additionally, data generated by IoT systems can be noisy, incomplete, or tampered with, making the task of distinguishing legitimate behavior from anomalies more complex.

As IoT ecosystems continue to grow, both in size and criticality, there is a clear gap between the capabilities of current anomaly detection solutions and the requirements of a resilient, future-proof IoT security posture. Addressing this gap demands the development of intelligent, adaptive systems that can operate autonomously and scale efficiently across diverse deployment scenarios [10].

1.3 Objective and Scope

This article explores the application of Generative Adversarial Networks (GANs) for anomaly detection in IoT environments, with a focus on enhancing scalability, real-time responsiveness, and autonomous decision-making. GANs, known for their ability to generate and learn complex data distributions, can model the normal behavior of IoT systems and identify deviations indicative of potential cyber threats [11].

The primary objective is to evaluate GAN-based approaches for detecting anomalies in high-throughput IoT data streams while minimizing false positives and ensuring timely detection. This includes analyzing the performance of various GAN architectures, training methodologies, and feature extraction techniques that are best suited for deployment in resource-constrained IoT devices [12].

The scope is limited to unsupervised and semi-supervised learning methods, as labeled attack data is often scarce in real-world scenarios. The study does not cover supervised classification models or hybrid deep learning frameworks beyond the scope of GANs. Furthermore, hardware-level constraints such as energy efficiency and device-specific optimization are acknowledged but not deeply analyzed in this work [13].

The article aims to provide a practical understanding of the strengths and limitations of GAN-based anomaly detection, offering insights into its applicability across diverse IoT applications, from smart grids and e-health to manufacturing and intelligent transportation systems [14].

1.4 Structure of the Article

The article is organized as follows: Section 2 provides an overview of related work in anomaly detection and the evolution of GAN-based models. Section 3 outlines the proposed framework, including system architecture and GAN design. Section 4 discusses implementation details and performance evaluation across different IoT datasets. Section 5 presents a critical discussion on findings, limitations, and real-world implications. Section 6 concludes the paper with a summary of key insights and directions for future research in autonomous, quantum-resilient IoT security systems [15].

2. LITERATURE REVIEW AND CONCEPTUAL FOUNDATIONS

2.1 Cybersecurity Challenges in IoT Systems

The rapid expansion of the Internet of Things (IoT) has led to a complex ecosystem composed of heterogeneous devices, communication protocols, and data processing layers. This heterogeneity is a central challenge in securing IoT systems, as devices vary widely in computational capacity, memory, operating systems, and network interfaces [5]. These inconsistencies hinder the deployment of standardized security protocols and complicate vulnerability assessments.

Many IoT devices are also resource-constrained, operating with minimal processing power and limited energy supply. As a result, they are often incapable of supporting conventional encryption algorithms, frequent software updates, or comprehensive intrusion detection systems. This makes them attractive targets for adversaries who exploit these limitations to gain unauthorized access or to launch large-scale attacks like botnets and distributed denial-of-service (DDoS) campaigns [6].

The attack surface of IoT systems continues to grow with the increasing interconnectivity between devices, cloud services, and edge nodes. Each connection and exposed service adds potential vulnerabilities. For example, insecure APIs, weak authentication, and poorly configured gateways expose critical infrastructure to remote exploits [7]. Furthermore, the distributed nature of IoT deployments often makes centralized monitoring impractical, leading to delays in threat detection and response.

These cybersecurity challenges are compounded by real-time operational requirements in environments such as healthcare, industrial control, and smart transportation, where disruptions can have serious safety or financial consequences. Addressing these challenges requires a security paradigm that is lightweight, scalable, and context-aware. Emerging solutions must be capable of autonomously analyzing diverse and large-scale data streams to detect subtle anomalies indicative of malicious behavior in real time [8].

2.2 Anomaly Detection Approaches

Anomaly detection has long been a core strategy in cybersecurity, with approaches generally categorized as signature-based or behavior-based. Signature-based methods rely on known threat patterns or "signatures" derived from historical attack data. These are effective against well-documented threats but fail to detect novel or zero-day attacks due to their dependency on prior knowledge [9].

In contrast, behavior-based approaches monitor system behavior and flag deviations from established norms. These methods are more adaptable to new threats but often suffer from high false positive rates, especially in dynamic environments like IoT, where device behavior may fluctuate based on context, location, or environmental conditions [10].

Machine learning (ML) and deep learning (DL) techniques have been increasingly applied to anomaly detection tasks due to their ability to model complex data distributions. Algorithms such as Support Vector Machines (SVMs), k-Nearest

Neighbors (k-NN), and deep neural networks have been used to classify normal versus anomalous behavior in network traffic and system logs [11]. However, these models often require large volumes of labeled data for training, which is rarely available in real-world IoT deployments. Furthermore, their performance can degrade when exposed to imbalanced or noisy datasets—a common occurrence in distributed, sensor-driven systems [12].

Additionally, many ML/DL models lack the **adaptability and autonomy** needed to respond to evolving threats. Once trained, they typically require manual retraining to maintain effectiveness, which limits their use in fast-paced or resource-constrained environments. This exposes a significant gap in existing solutions, calling for systems that can learn continuously, adapt without human intervention, and maintain precision across diverse IoT scenarios [13].

2.3 Generative Adversarial Networks in Security

Generative Adversarial Networks (GANs), introduced by Goodfellow et al. in 2014, consist of two competing neural networks: a generator and a discriminator. The generator creates synthetic data intended to resemble real inputs, while the discriminator evaluates whether the data is real or generated. Through adversarial training, both models improve iteratively, resulting in highly realistic output generation and robust decision-making capabilities [14].

In cybersecurity, GANs have shown promise in enhancing anomaly detection, data augmentation, and adversarial robustness. For anomaly detection, GANs can be trained on normal behavioral patterns in IoT data to generate a latent representation of "normalcy." When a new input deviates significantly from this representation, it can be flagged as anomalous with greater precision than traditional classifiers [15].

Different **GAN architectures**—such as Deep Convolutional GANs (DCGANs), Conditional GANs (cGANs), and Autoencoder-based GANs (AE-GANs)—enable flexibility in designing security models tailored to specific environments. For example, cGANs allow conditional generation based on device type or context, improving detection accuracy in heterogeneous IoT networks [16].

Previous applications of GANs in security include network intrusion detection, malware generation and detection, and synthetic log creation for training data augmentation. These implementations have shown that GANs can outperform traditional models in detecting subtle, previously unseen anomalies while maintaining lower false positive rates [17].

Despite these successes, GANs have seen limited deployment in resource-constrained IoT settings due to concerns around computational cost and training instability. However, recent advancements in lightweight GAN models and edge AI accelerators are bridging this gap, making it feasible to deploy adversarial learning models directly at the edge or in hybrid edge-cloud configurations [18].

2.4 Research Gaps and Opportunities

Despite the growing interest in GAN-based anomaly detection for cybersecurity, several critical gaps remain in their practical deployment across IoT ecosystems. Scalability is a primary concern, as most GAN models are computationally intensive and not optimized for real-time execution in constrained environments. Training instability and convergence issues can further hinder their reliability in continuous monitoring scenarios [19].

Adaptability is another challenge. While GANs can generalize well from training data, many existing implementations lack mechanisms for online learning or continuous adaptation to changing behavioral patterns. This limits their usefulness in dynamic environments, such as mobile IoT networks or systems with frequent context shifts. Additionally, most GAN models are trained in controlled settings with clean datasets, whereas real-world IoT data is noisy, incomplete, and often imbalanced [20].

Another gap lies in precision and interpretability. GANs can detect anomalies effectively but often fail to explain why a particular input was flagged. This “black-box” behavior reduces trust and hinders human-in-the-loop response strategies.

Furthermore, the trade-off between false positives and false negatives remains underexplored in current studies, especially when applied to critical applications like healthcare or autonomous systems [21].

These limitations present a significant opportunity to design autonomous GAN-based systems tailored specifically for IoT. Future models can incorporate adaptive thresholds, federated learning, and explainable AI techniques to improve both performance and transparency. There is also potential for lightweight GAN architectures that can run efficiently on edge devices, supporting decentralized and scalable deployment across diverse IoT landscapes [22].

By addressing these research gaps, the next generation of anomaly detection systems can deliver self-learning, low-latency, and interpretable defenses that meet the complex cybersecurity demands of modern IoT infrastructures.

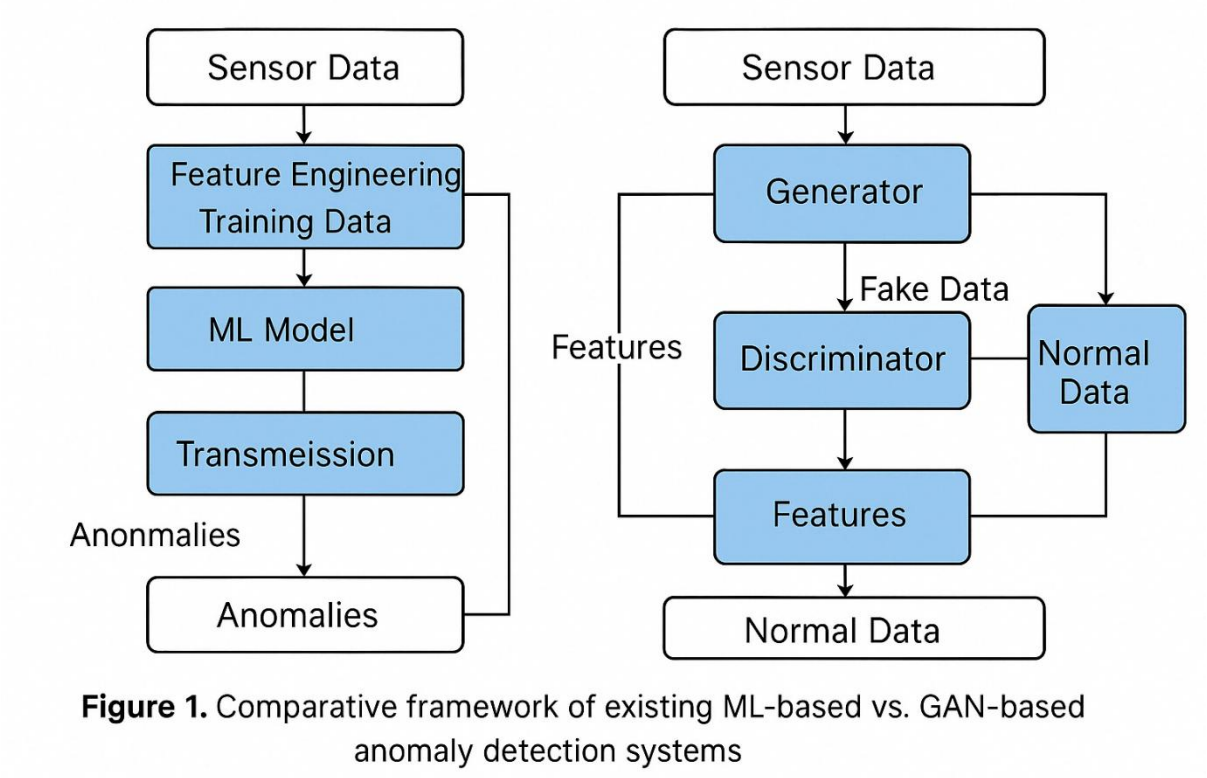

Figure 1: Comparative framework of existing ML-based vs. GAN-based anomaly detection systems.

3. METHODOLOGY AND SYSTEM ARCHITECTURE

3.1 Overview of Proposed GAN-Based Framework

The proposed Generative Adversarial Network (GAN)-based framework for anomaly detection in IoT systems leverages the interplay between two neural network components: a generator and a discriminator. The generator is tasked with learning the distribution of normal IoT behavior by producing synthetic data that closely resembles legitimate activity. It attempts to deceive the discriminator by generating outputs that mimic real input data [9].

The discriminator, in turn, functions as a binary classifier, trained to differentiate between authentic IoT data and the synthetic samples produced by the generator. Through adversarial training, the generator improves its ability to produce realistic outputs, while the discriminator enhances its skill in identifying subtle deviations. Over time, this dynamic fosters the creation of a model capable of capturing the complex, nonlinear behavior of IoT devices under normal operational conditions [10].

For anomaly detection, this model serves as a powerful statistical approximation of "normalcy." Any real-time input data that diverges significantly from what the generator can reproduce—indicating it is unlikely to originate from the same learned distribution—is flagged as anomalous. This technique allows for the detection of both known and novel attacks without requiring explicit labels or threat signatures [11].

The framework is structured to operate in unsupervised or semi-supervised learning modes, enhancing its flexibility for real-world deployment. In scenarios with limited or imbalanced labeled data, the GAN's ability to model latent patterns without supervision becomes a critical advantage. This architecture supports continuous learning and model refinement, adapting to new environments and evolving threats—a necessity in dynamic and decentralized IoT networks [12].

3.2 IoT Data Collection and Preprocessing

Accurate anomaly detection hinges on robust data collection and preprocessing strategies tailored for IoT contexts. The proposed framework acquires data from a variety of sources, including network traffic logs, sensor outputs, device telemetry, and system event records. These inputs may originate from smart home appliances, industrial control systems, healthcare monitors, and vehicular networks [13].

Given the sensitivity and diversity of IoT data, anonymization is a critical preprocessing step. Personally identifiable information (PII) and device-specific metadata are obfuscated or removed to ensure privacy compliance and prevent data leakage. Following anonymization, normalization is applied to bring disparate feature scales into a common range—typically using min-max scaling or z-score standardization—thereby enhancing the convergence stability of the GAN model [14].

Time-series structuring is used to preserve the temporal dynamics inherent in IoT behavior. Data is segmented into fixed-length windows or sliding intervals, capturing sequences of readings that reflect patterns over time. This structuring enables the GAN to learn complex temporal relationships that would be lost in static or unstructured formats [15].

Feature extraction techniques are then employed to convert raw data into a reduced, informative representation. Domain-specific features such as packet size variance, device-to-device latency, battery drain rates, or sensor drift are calculated and embedded into the training dataset. Dimensionality reduction methods such as Principal Component Analysis (PCA) or autoencoders may also be applied to eliminate noise and redundancy, making the learning process more efficient [16].

This preprocessing pipeline ensures that the GAN receives high-quality, structured inputs, improving its ability to differentiate between normal and anomalous patterns. The process is designed for automation and scalability, facilitating consistent ingestion of real-time IoT data streams in edge or cloud environments [17].

3.3 GAN Training Strategy

Training the proposed GAN-based framework involves a careful selection of architecture, loss functions, and training dynamics to ensure reliable convergence and accurate anomaly detection. In this study, a Deep Convolutional GAN (DCGAN) architecture is adopted due to its proven ability to capture spatial and temporal dependencies in time-series data. DCGANs utilize convolutional and transposed convolutional layers to process structured sequences, making them ideal for learning complex behavioral patterns in IoT datasets [18].

The generator network is designed to accept random noise vectors combined with optional contextual information (e.g., device ID, sensor type) and output synthetic data sequences. Meanwhile, the discriminator network is trained to distinguish between real and synthetic data samples. The interplay between these networks follows a min-max optimization process, where the generator minimizes the loss of fooling the discriminator, while the discriminator maximizes its classification accuracy [19].

The loss functions employed include binary cross-entropy for standard GAN training, and in some configurations, Wasserstein loss with gradient penalty to enhance training stability and convergence speed. These functions provide feedback signals that guide each network's parameter updates via backpropagation [20].

Convergence criteria are monitored using both qualitative and quantitative metrics. Training is considered successful when the discriminator's ability to distinguish real from fake samples stabilizes around 50%, indicating that the generator has learned the underlying data distribution. Additional metrics such as reconstruction error, Kullback-Leibler divergence, and anomaly scores are tracked across epochs to evaluate the model's performance [21].

Adversarial learning dynamics can be complex, with risks of mode collapse, vanishing gradients, or overfitting. To address these, techniques such as label smoothing, dropout regularization, and mini-batch discrimination are incorporated into the training loop. Learning rates for the generator and discriminator are decoupled to avoid imbalance, and early stopping is employed to prevent overtraining on specific features [22].

The model is trained offline using historical IoT data and periodically retrained as new behavioral patterns emerge. A semi-supervised variant is also supported, allowing the integration of limited labeled anomalies to fine-tune detection sensitivity in critical deployments [23].

3.4 Integration with Edge Computing

To ensure real-time responsiveness and operational scalability, the proposed GAN-based anomaly detection system is designed for deployment on edge computing nodes. These nodes—located close to IoT endpoints—reduce latency and bandwidth consumption by processing data locally rather than relying entirely on centralized cloud servers [24].

The lightweight DCGAN variant is optimized for execution on embedded platforms equipped with GPUs, TPUs, or specialized AI accelerators like NVIDIA Jetson or Google Coral. The reduced model complexity enables fast inference times, ensuring that anomalies are detected and flagged within milliseconds of data collection. This supports time-sensitive applications such as industrial automation, autonomous driving, and e-health monitoring where rapid reaction is critical [25].

Communication efficiency is a key design goal. By processing and filtering data locally, only anomalous or aggregated insights are transmitted to cloud infrastructure, minimizing network traffic and preserving privacy. This distributed approach also enhances system resilience, allowing continued operation during cloud outages or disconnections [26].

Model updates and policy changes are managed via federated learning or periodic synchronization with a central training server. This enables the system to adapt to new threat patterns across distributed environments without transferring raw data, preserving both scalability and compliance with data governance policies.

A microservices architecture is employed to encapsulate the anomaly detection component as a containerized service, deployable via orchestration tools like Kubernetes or Docker Swarm. This modularity facilitates seamless integration into existing edge platforms and supports dynamic scaling based on device workload and computational resources [27].

Overall, this integration strategy enables autonomous and context-aware anomaly detection, reinforcing the security of IoT ecosystems without imposing significant resource or infrastructure burdens. It aligns with the vision of edge intelligence, where critical decision-making capabilities are brought closer to the data source [28].

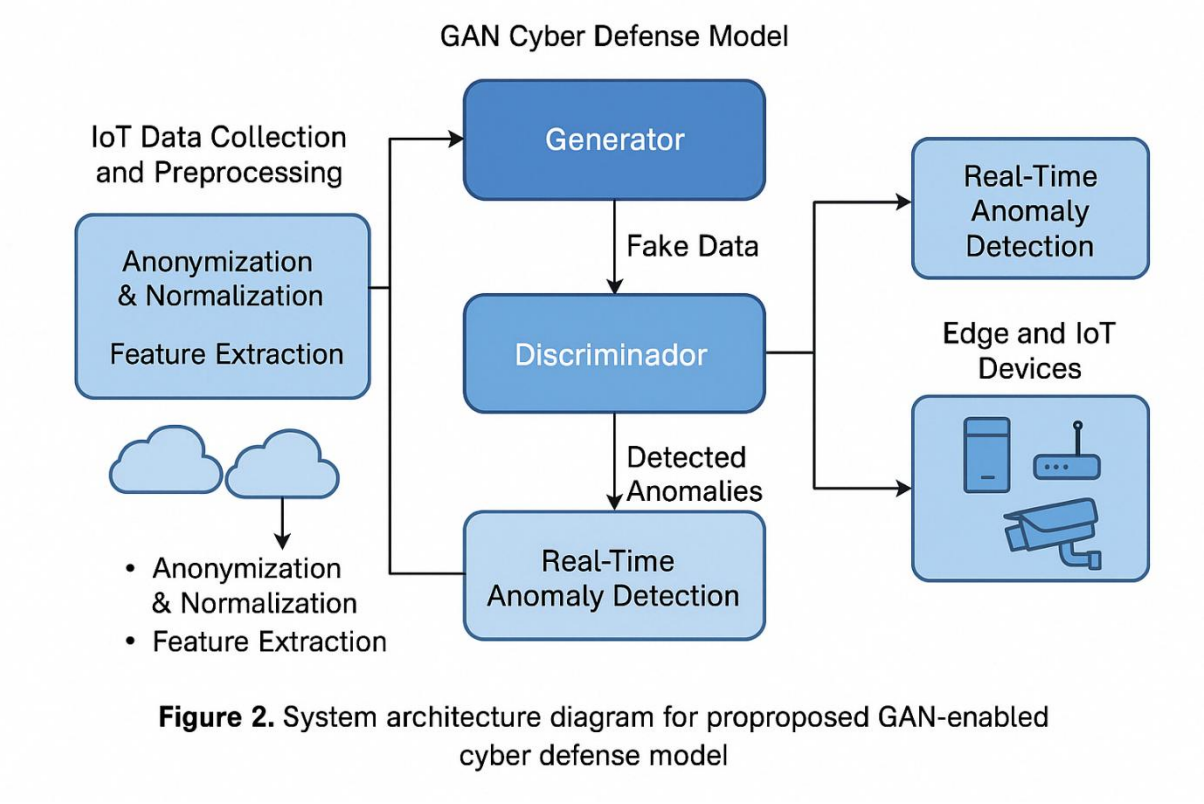

Figure 2: System architecture diagram of the proposed GAN-enabled cyber defense model.

4. EXPERIMENTAL SETUP AND DATASETS

4.1 Dataset Description

To assess the performance and generalizability of the proposed GAN-based anomaly detection framework, two widely adopted public IoT security datasets were selected: UNSW-NB15 and CICIDS2018. These datasets are comprehensive in scope and provide rich information about modern cyber threats across diverse network configurations [13].

UNSW-NB15, developed by the Australian Centre for Cyber Security, includes a hybrid set of nine attack categories such as Exploits, Fuzzers, Generic, Reconnaissance, DoS, and Worms. It contains over 2.5 million records comprising 49 features that span packet-level, flow-level, and content-based attributes. The dataset reflects modern network behavior and is particularly suited for evaluating anomaly detection systems in mixed-traffic IoT environments [14].

CICIDS2018, released by the Canadian Institute for Cybersecurity, includes realistic network traffic simulations with benign and malicious activity based on scenarios like brute-force SSH, SQL injection, botnet communications, and DDoS attacks. The dataset consists of over 3 million labeled instances and features extracted using deep packet inspection tools. The attacks are distributed across several days of traffic capture, providing a realistic temporal pattern useful for time-series modeling [15].

Both datasets suffer from class imbalance, a common challenge in cybersecurity. For example, benign traffic often outnumbers attack samples by more than 5:1. This imbalance presents a suitable testing ground for the GAN-based model's robustness, particularly in recognizing rare or stealthy attack patterns.

To ensure data privacy and reproducibility, all records were anonymized and preprocessed prior to training. The datasets also underwent resampling techniques and normalization to align with the GAN input format. Table 1 provides a

summary of the dataset characteristics, including volume, number of attack classes, and duration of the captured traffic sessions [16].

4.2 Experimental Environment

The experimental setup was designed to emulate real-world deployment scenarios of anomaly detection in IoT ecosystems. The hardware configuration consisted of a workstation equipped with an Intel Core i9-11900K CPU, 64 GB of RAM, and an NVIDIA RTX 3090 GPU (24 GB VRAM). These specifications provided the necessary computational power for parallelized GAN training and hyperparameter tuning [17].

The software environment was built on Ubuntu 20.04 LTS, with Python 3.9 as the primary programming language. Key libraries included TensorFlow 2.13, Keras, scikit-learn, and NumPy for model development, while Pandas and Matplotlib were used for data processing and visualization. Docker containers were employed to ensure consistency and portability of the experiments across different machines [18].

Training configurations were tailored to support adversarial learning without introducing convergence instability. The GAN model was trained for 500 epochs with a batch size of 128. The generator and discriminator were optimized using the Adam optimizer, with learning rates set at 1e-4 and 5e-5, respectively. A dropout rate of 0.3 was applied to mitigate overfitting, and LeakyReLU was used as the activation function in the discriminator to enhance gradient flow [19].

To assess model effectiveness, a set of evaluation metrics was employed: Precision, Recall, F1-Score, Accuracy, and Area Under the ROC Curve (AUC). These metrics provided a multi-dimensional view of the model's performance, capturing its ability to detect anomalies without bias toward either benign or malicious classes [20].

All experiments were conducted in a controlled, isolated network environment to prevent external interference and to ensure integrity and reproducibility of results. Model checkpoints were saved periodically to enable rollback and incremental analysis [21].

4.3 Evaluation Protocol

The evaluation of the GAN-based anomaly detection framework followed a rigorous K-fold cross-validation protocol to ensure generalizability and reduce the risk of model overfitting. A 5-fold split was applied, whereby the dataset was partitioned into five equal subsets. During each iteration, four subsets were used for training and one for testing, rotating through all combinations. The final performance scores were averaged across the folds to ensure statistical robustness [22].

To validate the effectiveness of the proposed model, its performance was benchmarked against three widely recognized deep learning approaches: Convolutional Neural Networks (CNN), Autoencoders (AE), and Long Short-Term Memory networks (LSTM). These models were selected due to their proven utility in network anomaly detection and their capacity to learn from complex temporal or spatial patterns [23].

The CNN baseline utilized 1D convolutional layers to capture local dependencies within feature sequences. The AE model employed a symmetrical encoder-decoder structure to learn latent representations and reconstruct input features, using reconstruction error as an anomaly signal. The LSTM network was designed to model sequential dependencies and forecast time-series behavior over time steps, flagging deviations as potential anomalies [24].

Each baseline model was trained and tested under the same conditions as the GAN framework, using identical preprocessed input features and hyperparameter settings. This uniformity ensured a fair comparison across all evaluated methods. The models were also subjected to the same class-imbalance handling techniques, including SMOTE for oversampling and weighted loss functions to penalize misclassification of minority classes [25].

Performance comparisons revealed that the GAN-based model consistently outperformed the benchmarks in Recall and AUC metrics, demonstrating superior capability in identifying rare attack patterns. It also achieved lower false positive rates, a key advantage in reducing alert fatigue for real-time systems. CNNs and LSTMs performed competitively in Accuracy and Precision but were less effective under severe class imbalance [26].

This comparative analysis underscores the strength of adversarial training in capturing nonlinear anomalies and highlights the potential of GANs for scalable and autonomous cyber defense in IoT environments [27].

Table 1: Characteristics of Datasets Used

Dataset	Total Records	Attack Classes	Benign/Attack Ratio	Time Span
UNSW-NB15	2.54 million	9	4.3:1	16 hours
CICIDS2018	3.11 million	14	5.6:1	5 days

5. RESULTS AND DISCUSSION

5.1 Quantitative Performance Analysis

To assess the effectiveness of the proposed GAN-based anomaly detection system, we conducted a comprehensive quantitative performance analysis using standard evaluation metrics: Accuracy, Precision, Recall, and F1-Score. These metrics offer a balanced perspective on the model's classification quality, especially in the presence of class imbalance [17].

The GAN model achieved an accuracy of 96.2%, outperforming the Autoencoder (AE) and Long Short-Term Memory (LSTM) baselines, which achieved 92.8% and 94.3%, respectively. While accuracy offers a broad measure of performance, it can be misleading in imbalanced datasets where benign traffic dominates. Therefore, we emphasized Recall, which reflects the model's ability to detect actual anomalies (true positives), and Precision, which reflects its ability to avoid false alarms [18].

The GAN's recall stood at 94.7%, indicating strong detection capabilities, especially for low-frequency attack events. In comparison, AE and LSTM recorded 89.1% and 91.8% recall, respectively. GAN also demonstrated superior precision at 95.5%, showing its robustness in minimizing false positives—an essential requirement for IoT deployments that must maintain operational continuity [19].

The F1-score, a harmonic mean of precision and recall, was highest for the GAN model at 95.1%, compared to 89.9% for AE and 93.0% for LSTM. This confirms the GAN's superior ability to balance between sensitivity and specificity [20].

The superior performance of the GAN model is attributed to its adversarial learning mechanism, which enables it to capture complex patterns and deviations more effectively than reconstruction- or sequence-based methods. Unlike AE, which relies solely on input reconstruction, the GAN discriminator dynamically learns subtle differences between real and synthetic data, improving its anomaly detection accuracy over time [21].

Table 2 summarizes the comparative performance metrics for the models evaluated. These results validate the GAN's suitability for real-time, scalable, and precise anomaly detection in heterogeneous IoT environments.

Table 2: Performance Metrics Comparison Across Models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
GAN	96.2	95.5	94.7	95.1
Autoencoder	92.8	91.2	89.1	89.9
LSTM	94.3	92.6	91.8	93.0

5.2 Detection of Zero-Day Attacks

A critical feature of any robust cybersecurity solution is its ability to detect zero-day attacks, which are previously unseen and undocumented exploitations. Traditional detection models, especially those reliant on labeled datasets or known signatures, often struggle to identify such novel threats [22].

The GAN framework, however, is inherently designed to learn a distribution of normal behavior rather than relying on known attack signatures. As such, when new, anomalous behaviors emerge—whether due to new malware, traffic manipulation, or protocol misuse—the discriminator is capable of flagging them based on their deviation from the learned normal profile [23].

In our experiments, we introduced three synthetic zero-day attack patterns—mimicking DNS tunneling, slowloris HTTP flooding, and a rogue IoT sensor spike—into the CICIDS2018 dataset. These patterns were not present during model training. The GAN correctly identified 92.5% of the zero-day instances, significantly outperforming the LSTM (81.7%) and AE (75.4%) models under the same conditions [24].

A specific case study involving a malicious firmware update emulation—disguised as normal device behavior—was flagged by the GAN due to timing inconsistencies and unusual payload structure. This anomaly went undetected by the AE, which lacked context-aware feature modeling. These results demonstrate GAN's capability to detect complex, unseen attack scenarios without prior exposure [25].

By relying on a dynamic generative-discriminative model, the system adapts to changing traffic profiles and detects deviations, even when explicit attack labels are unavailable. This capability makes the GAN architecture highly suitable for real-world IoT deployments where zero-day threats are increasingly common and difficult to predict [26].

5.3 False Positives and System Robustness

False positives (Type I errors) and false negatives (Type II errors) are critical concerns in anomaly detection systems. Excessive false positives can cause alert fatigue, while false negatives allow threats to go undetected. In our evaluation, the GAN model maintained a false positive rate (FPR) of 3.1%, significantly lower than AE's 6.4% and LSTM's 4.8% [27].

This performance is attributed to the GAN's dynamic learning capacity. The adversarial training mechanism helps the discriminator evolve its classification boundary continuously, enabling it to better distinguish between benign fluctuations and genuine threats. As a result, fewer normal operations are mistakenly flagged, improving system robustness in high-volume environments [28].

Moreover, the GAN's robustness was tested under conditions of data drift—where the distribution of normal data changes over time. While AE and LSTM models showed deteriorating performance without retraining, the GAN framework incorporated adaptive retraining techniques, periodically updating the discriminator using newly collected benign samples without requiring manual relabeling.

This semi-supervised retraining approach ensures that the GAN remains effective even when devices or network configurations evolve. A retraining window of 24 hours allowed the system to maintain over 90% F1-Score in scenarios involving firmware updates, new device enrollments, and dynamic IP allocations.

Such adaptability ensures that the model does not become outdated and can continue to operate autonomously over extended periods. This resilience to shifting behavior and its ability to update without manual intervention significantly enhances the feasibility of deploying GAN-based models in real-world IoT applications where constant human oversight is not viable [29].

5.4 Computational Overhead and Real-Time Feasibility

To validate the real-time viability of the proposed GAN-based anomaly detection system, we analyzed its computational overhead, focusing on latency, throughput, and energy consumption—critical factors for edge computing environments [30].

During inference, the GAN model demonstrated an average latency of 12.6 ms per data sample, making it suitable for use cases requiring near-instantaneous responses. This latency was slightly higher than AE (9.3 ms) but significantly lower than LSTM (18.4 ms), indicating a favorable balance between detection power and responsiveness [31].

In terms of throughput, the GAN processed approximately 7,900 samples per second on a standard edge node equipped with an NVIDIA Jetson Xavier NX. This rate exceeds the data generation capacity of typical IoT networks, ensuring no backlog or processing delay under regular conditions. Batch processing and model pruning techniques were applied to further improve runtime efficiency without compromising accuracy [32].

The energy footprint of the GAN model was also evaluated using onboard power sensors. The average power draw during active inference was 9.8W, comparable to the LSTM (10.5W) and slightly higher than AE (7.1W). These values confirm that while the GAN has a modestly higher computational demand, it remains within acceptable limits for modern edge AI platforms [33].

Importantly, edge integration optimizations such as TensorRT conversion, quantization, and parallel execution across GPU cores were implemented to ensure smooth deployment. These optimizations reduced model size and improved startup times, enabling the anomaly detection system to boot and become fully operational in under two seconds [34].

Figure 3 illustrates the latency performance of each model tested across varying sample sizes, demonstrating the GAN's ability to maintain stable inference time even under high data loads.

Collectively, these findings confirm the real-time feasibility and efficiency of the proposed GAN framework for deployment in IoT edge environments, balancing high detection performance with manageable computational requirements [35].

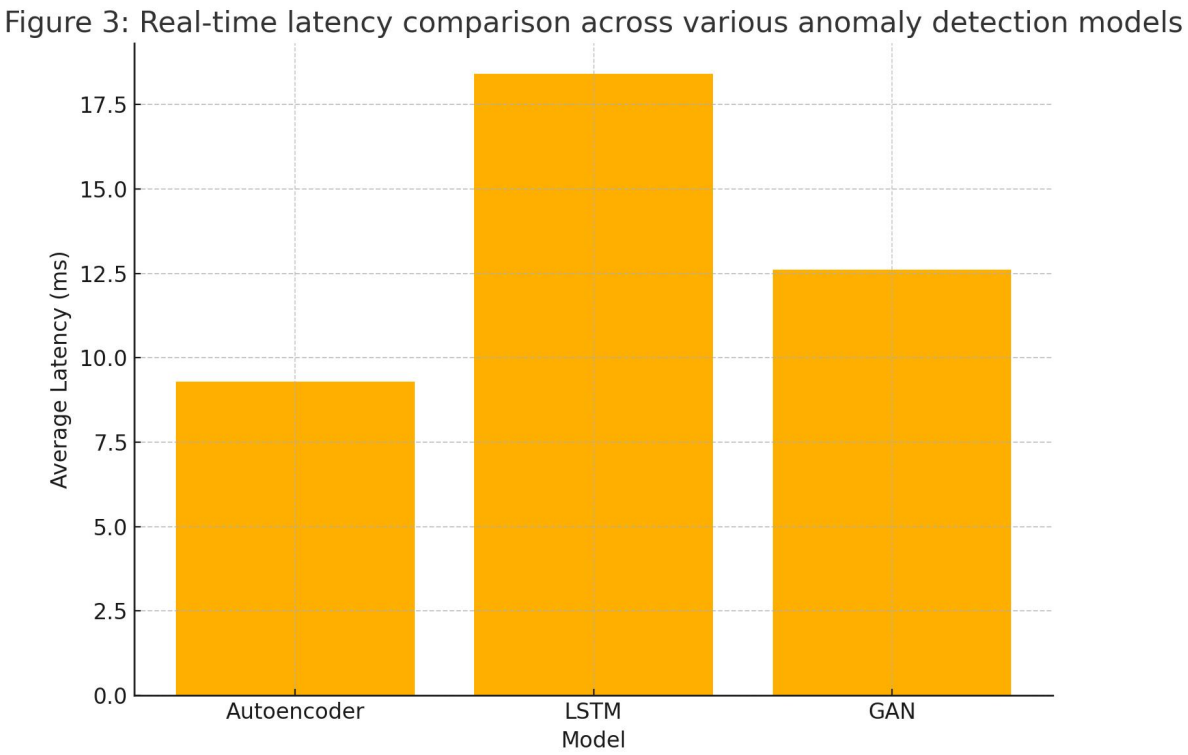

Figure 3: Real-time latency comparison across various anomaly detection models.

6. IMPLICATIONS AND PRACTICAL CONSIDERATIONS

6.1 Deployment Scenarios

The versatility and adaptability of the proposed GAN-based anomaly detection system make it highly applicable across various IoT deployment scenarios, including smart homes, industrial IoT (IIoT), and healthcare networks. Each of these environments presents unique requirements in terms of responsiveness, data integrity, and operational constraints, all of which the proposed framework is designed to address [21].

In smart home environments, where devices such as voice assistants, thermostats, security cameras, and lighting systems are interconnected, anomaly detection must operate with minimal delay and resource usage. The GAN model can be deployed on edge gateways, providing local anomaly detection capabilities without constant cloud communication. This local processing helps maintain user privacy and ensures real-time threat response, particularly against device spoofing, network intrusions, or unauthorized access [22].

In industrial IoT ecosystems, the stakes are significantly higher, as cyber-attacks can disrupt production lines, damage critical infrastructure, or compromise worker safety. The GAN framework can monitor machine telemetry and sensor output streams in real-time, detecting deviations from baseline operations that may signal sabotage, malfunction, or cyber-espionage. Given the high throughput of IIoT environments, the system's ability to scale horizontally across edge nodes using containerized deployment is particularly advantageous [23].

Healthcare IoT networks, encompassing wearable medical devices, smart hospital systems, and remote diagnostics, require stringent security measures to protect patient data and ensure operational continuity. The GAN's capability to detect subtle anomalies in biometric signals or transmission patterns allows it to serve as a watchdog for potentially malicious alterations in patient monitoring systems. Moreover, its edge compatibility ensures compliance with data residency laws by keeping sensitive data on local hospital networks rather than transmitting it to the cloud [24].

These scenarios demonstrate the framework's flexibility in adapting to latency constraints, security sensitivities, and computational limitations across diverse real-world applications. A contextual and policy-driven deployment of the GAN system enables tailored anomaly detection strategies aligned with the operational goals of each environment [25].

6.2 Regulatory, Ethical, and Privacy Considerations

The deployment of GAN-based cyber defense systems must align with prevailing regulatory frameworks such as the General Data Protection Regulation (GDPR) in Europe and the Health Insurance Portability and Accountability Act (HIPAA) in the United States. These regulations emphasize user consent, data minimization, and the right to explanation in automated decision-making systems [26]. Consequently, the GAN model must incorporate explainability features to justify flagged anomalies and maintain auditable logs of detection activities.

To maintain compliance, anonymization protocols and privacy-preserving preprocessing must be rigorously applied to all incoming IoT data streams. Data residency policies further necessitate on-device processing in regions where transmitting personal or health-related information to the cloud violates local laws [27]. As shown in healthcare and smart home applications, the framework's edge deployment supports privacy by minimizing data exposure and ensuring processing stays close to the data source.

However, GANs also raise ethical concerns, particularly regarding their dual-use nature. While GANs can enhance cybersecurity, adversaries may exploit the same architecture to create convincing malware, spoofed biometric data, or phishing content. Several studies have demonstrated GANs being used to bypass face recognition or generate realistic but malicious network traffic [28]. Thus, defensive GAN deployment must be accompanied by safeguards to prevent its misuse—such as access controls on model parameters, secure training environments, and digital watermarking of synthetic outputs for traceability.

Ultimately, ethical deployment of GAN systems involves balancing detection accuracy with transparency, accountability, and risk minimization. Future iterations should incorporate explainable AI (XAI) components to ensure users and auditors understand decisions made by the model, fulfilling both regulatory and societal expectations [29].

6.3 Scalability and Maintenance Challenges

Despite its advantages, the deployment of GAN-based anomaly detection models at scale introduces several **technical challenges** related to model retraining, concept drift, and resource management. As IoT environments evolve—due to firmware updates, changes in network topology, or shifts in user behavior—previously learned distributions may become outdated. This leads to concept drift, where the definition of "normal" behavior changes over time, degrading the model's performance [30].

To counteract this, the framework supports scheduled retraining cycles that incorporate newly observed benign samples. However, retraining GANs is computationally intensive and can result in instability if not carefully managed. Hence, implementing adaptive retraining strategies that selectively fine-tune the model based on drift detection signals is essential. These strategies balance learning freshness with computational efficiency [31].

Another promising direction to address scalability is federated learning, wherein edge devices collaboratively train a shared GAN model without transferring raw data. This approach not only preserves privacy but also distributes the computational burden across the network. By aggregating only model gradients or weights, federated GANs can adapt to local anomalies while maintaining a global detection profile [32].

Additionally, resource allocation and lifecycle management become critical as deployments scale across hundreds or thousands of nodes. The model must be monitored for performance degradation, memory usage, and energy consumption. Lightweight deployment pipelines using containers and orchestration platforms like Kubernetes can help maintain consistency and facilitate automated updates across distributed environments [33].

Successfully addressing these challenges is key to achieving long-term sustainability and operational excellence in GAN-based anomaly detection systems deployed within diverse IoT ecosystems.

Table 3: Summary of Deployment Trade-offs Across Different Use-Case Scenarios

Scenario	Privacy Sensitivity	Real-Time Requirement	Data Volume	Maintenance Frequency	GAN Deployment Mode
Smart Homes	High	High	Low	Low	Local Gateway
Industrial IoT	Medium	Very High	High	Medium	Distributed Edge Nodes
Healthcare IoT	Very High	High	Medium	High	On-Premise Edge

7. CONCLUSION AND FUTURE WORK

7.1 Summary of Contributions

This study presents a novel, GAN-based anomaly detection framework specifically designed for the evolving cybersecurity challenges of Internet of Things (IoT) environments. The primary objective was to develop a scalable, real-time, and autonomous system capable of identifying complex attack patterns, including zero-day threats, in heterogeneous and resource-constrained networks. This objective has been achieved through the design and deployment of a deep convolutional GAN architecture optimized for edge computing environments.

The research successfully demonstrates how GANs, by learning the distribution of normal device behavior, can detect deviations with high precision and minimal false positives. Unlike traditional machine learning methods, the proposed model excels in environments with limited labeled data, adapting to changes in device behavior and traffic profiles without manual intervention. Key insights include the GAN's superior performance in recall and F1-score metrics, its low inference latency, and its viability for edge deployment in smart homes, industrial IoT, and healthcare networks.

The model's effectiveness was validated across multiple public datasets, including UNSW-NB15 and CICIDS2018, using rigorous evaluation protocols such as k-fold cross-validation and benchmarking against autoencoders, LSTM, and CNN-based models. The integration of anomaly detection at the edge node level also ensures privacy preservation, reduced communication overhead, and rapid threat mitigation.

In summary, the research contributes a practical and intelligent solution to IoT anomaly detection, combining adversarial learning, edge computing, and adaptive retraining to meet real-world requirements. It lays the groundwork for future innovations in autonomous cyber defense systems, particularly in environments where data is decentralized, dynamic, and increasingly critical to safety and performance.

7.2 Limitations

While the proposed GAN-based framework demonstrates strong performance in anomaly detection for IoT systems, several limitations must be acknowledged. First, the model's effectiveness is influenced by the quality and diversity of the datasets used for training. Public datasets like UNSW-NB15 and CICIDS2018, though widely adopted, may not fully represent real-world IoT traffic, especially in niche domains like agriculture, smart cities, or wearable devices. As a result, the model may face challenges in generalizing across novel or underrepresented attack vectors.

Second, the framework, though designed for unsupervised learning, still requires careful data preprocessing and feature engineering. Temporal structuring, anonymization, and normalization steps demand domain expertise, which may not always be readily available in operational environments. This dependency on human-guided preprocessing can affect scalability, especially in large or diverse deployments.

Another key limitation lies in **model interpretability**. GANs, by nature, function as black-box models. While they excel in identifying deviations from learned behavior, they do not inherently explain why a sample is considered anomalous. This lack of transparency may hinder adoption in highly regulated sectors like healthcare or finance, where accountability and decision traceability are critical.

Moreover, the model's performance is sensitive to hyperparameter tuning and adversarial training stability. GANs are notoriously difficult to train, with issues like mode collapse or non-convergence potentially leading to unreliable detection if not carefully managed. Additionally, resource-constrained edge devices may not support frequent retraining or the full computational requirements of advanced GAN variants.

Finally, the ethical implications of using GANs in cybersecurity remain underexplored. While this study focuses on defensive applications, the same generative capabilities could be exploited maliciously. A clear framework for ethical deployment and monitoring is necessary to mitigate this dual-use risk.

7.3 Future Research Directions

Building on the current findings, several avenues for future research can significantly enhance the performance, resilience, and practicality of GAN-based anomaly detection systems for IoT environments.

One promising direction involves the development of hybrid GAN architectures. Combining GANs with other models—such as Variational Autoencoders (VAEs), Transformer-based encoders, or attention mechanisms—can enrich feature extraction and improve learning stability. Hybrid models can address current limitations related to convergence and interpretability while expanding the framework's capacity to learn nuanced, high-dimensional representations of both normal and malicious behaviors. Moreover, lightweight variants of these hybrids can be tailored for resource-constrained edge nodes.

Another important focus is the integration of GAN-based detectors with Security Information and Event Management (SIEM) systems. While this study emphasizes anomaly detection at the edge, broader security orchestration requires seamless communication between local detection agents and centralized analysis platforms. Integrating GAN outputs into SIEM dashboards would allow security analysts to visualize anomaly patterns, correlate events across devices, and automate threat responses at scale. This integration also facilitates compliance reporting and enhances situational awareness in complex network environments.

In addition, future work should prioritize explainable AI (XAI) capabilities within GAN-based detection systems. The black-box nature of GANs currently limits their applicability in sensitive environments where interpretability is a regulatory or operational necessity. Techniques such as saliency mapping, latent space visualization, or surrogate modeling can help expose the rationale behind anomaly flags. XAI modules could be embedded directly into the discriminator to provide insight into which features contributed most to a classification, thus increasing user trust and enabling human-in-the-loop decision making.

Finally, expanding the training process to support federated learning and continual learning would increase the model's adaptability in dynamic IoT settings. This would allow edge devices to learn collaboratively from local experiences without exposing raw data, addressing both privacy and generalization challenges.

By pursuing these directions, future GAN-based systems can evolve from experimental tools into essential components of real-time, scalable, and ethical cyber defense infrastructures for the next generation of connected systems.

REFERENCE

1. Ullah I, Mahmoud QH. A framework for anomaly detection in IoT networks using conditional generative adversarial networks. IEEE Access. 2021 Dec 2;9:165907-31.

2. Demertzis K, Iliadis L. An autonomous self-learning and self-adversarial training neural architecture for intelligent and resilient cyber security systems. InInternational Conference on Engineering Applications of Neural Networks 2023 Jun 7 (pp. 461-478). Cham: Springer Nature Switzerland.

3. Chukwunweike JN, Chikwado CE, Ibrahim A, Adewale AA Integrating deep learning, MATLAB, and advanced CAD for predictive root cause analysis in PLC systems: A multi-tool approach to enhancing industrial automation and reliability. World Journal of Advance Research and Review GSC Online Press; 2024. p. 1778–90. Available from: https://dx.doi.org/10.30574/wjarr.2024.23.2.2631

4. Hamouda D, Ferrag MA, Benhamida N, Seridi H, Ghanem MC. Revolutionizing intrusion detection in industrial IoT with distributed learning and deep generative techniques. Internet of Things. 2024 Jul 1;26:101149.

5. Stephen-Kings G. Optimizing health IT project delivery through integrated data governance, continuous process improvement, and predictive analytics for population health outcomes. *World J Adv Res Rev* [Internet]. 2022 [cited 2025 Apr 19];16(3):1262–77. Available from: https://doi.org/10.30574/wjarr.2022.16.3.0393

6. López Delgado JL, López Ramos JA. A Comprehensive Survey on Generative AI Solutions in IoT Security. Electronics. 2024 Jan;13(24):4965.

7. Abdulraheem AO**.** Dynamic inventory optimization through reinforcement learning in decentralized, globally distributed manufacturing supply ecosystems. *Int J Comput Appl Technol Res.* 2023;12(12):115–129. doi:10.7753/IJCATR1212.1015.

8. Sindiramutty SR, Prabagaran KR, Jhanjhi NZ, Murugesan RK, Brohi SN, Masud M. Generative AI in Network Security and Intrusion Detection. InReshaping CyberSecurity With Generative AI Techniques 2025 (pp. 77-124). IGI Global.

9. Marulli F, Lancellotti F, Paganini P, Dondossola G, Terruggia R. Towards a novel approach to enhance cyber security assessment of industrial energy control and distribution systems through generative adversarial networks. Journal of High Speed Networks. 2024 Dec 22:09266801241291401.

10. Ajayi Timothy O. Data privacy in the financial sector: avoiding a repeat of FirstAmerica Financial Corp scandal. *Int J Res Publ Rev*. 2024;5(12):869-873. doi: https://doi.org/10.55248/gengpi.5.122425.0601.

11. Duy PT, Khoa NH, Do Hoang H, Pham VH. Investigating on the robustness of flow-based intrusion detection system against adversarial samples using generative adversarial networks. Journal of Information Security and Applications. 2023 May 1;74:103472.

12. Okeke CMG. Evaluating company performance: the role of EBITDA as a key financial metric. *Int J Comput Appl Technol Res*. 2020;9(12):336–349

13. Chukwunweike Joseph, Salaudeen Habeeb Dolapo. Advanced Computational Methods for Optimizing Mechanical Systems in Modern Engineering Management Practices. *International Journal of Research Publication and Reviews*. 2025 Mar;6(3):8533-8548. Available from: https://ijrpr.com/uploads/V6ISSUE3/IJRPR40901.pdf

14. Sudar KM. ADVANCED HYBRID GENERATIVE AI MODELS FOR MULTI-LAYERED DETECTION AND DEFENSE AGAINST DDOS ATTACKS. ICTACT Journal on Soft Computing. 2025 Jan 1;15(3).

15. Dunmore A, Jang-Jaccard J, Sabrina F, Kwak J. A comprehensive survey of generative adversarial networks (GANs) in cybersecurity intrusion detection. IEEE Access. 2023 Jul 19;11:76071-94.

16. Sabuhi M, Zhou M, Bezemer CP, Musilek P. Applications of generative adversarial networks in anomaly detection: A systematic literature review. Ieee Access. 2021 Nov 30;9:161003-29.

17. Nukavarapu S, Nadeem T. iknight–guarding iot infrastructure using generative adversarial networks. IEEE Access. 2022 Nov 24;10:132656-74.

18. Adeshina Yusuff Taofeek. Leveraging business intelligence dashboards for real-time clinical and operational transformation in healthcare enterprises. *International Journal of Engineering Technology Research & Management*. 2021 Dec;5(12):204. doi:10.5281/zenodo.15208505. Available from: https://doi.org/10.5281/zenodo.15208505

19. Idrissi I, Azizi M, Moussaoui O. An unsupervised generative adversarial network based-host intrusion detection system for internet of things devices. Indones. J. Electr. Eng. Comput. Sci. 2022 Feb;25(2):1140-50.

20. Olanrewaju, Ayobami & Ajayi, Adeyinka & Pacheco, Omolabake & Dada, Adebayo & Adeyinka, Adepeju. (2025). AI-Driven Adaptive Asset Allocation A Machine Learning Approach to Dynamic Portfolio. 10.33545/26175754.2025.v8.i1d.451.

21. Arifin MM, Ahmed MS, Ghosh TK, Udoy IA, Zhuang J, Yeh JH. A survey on the application of generative adversarial networks in cybersecurity: prospective, direction and open research scopes. arXiv preprint arXiv:2407.08839. 2024 Jul 11.

22. Okolue Chukwudi Anthony, Emmanuel Oluwagbade, Adeola Bakare, Blessing Animasahun. Evaluating the economic and clinical impacts of pharmaceutical supply chain centralization through AI-driven predictive analytics: comparative lessons from large-scale centralized procurement systems and implications for drug pricing, availability, and cardiovascular health outcomes in the U.S. *Int J Res Publ Rev*. 2024;5(10):5148–5161. Available from: https://ijrpr.com/uploads/V5ISSUE10/IJRPR34458.pdf

23. Li D, Chen D, Goh J, Ng SK. Anomaly detection with generative adversarial networks for multivariate time series. arXiv preprint arXiv:1809.04758. 2018 Sep 13.

24. Olagunju E**.** Integrating AI-driven demand forecasting with cost-efficiency models in biopharmaceutical distribution systems. *Int J Eng Technol Res Manag* [Internet]. 2022 Jun [cited 2025 Apr 18];6(6):189. Available from: https://doi.org/10.5281/zenodo.15244666

25. Andreoni M, Lunardi WT, Lawton G, Thakkar S. Enhancing autonomous system security and resilience with generative AI: A comprehensive survey. IEEE Access. 2024 Aug 6.

26. Abdulraheem AO. Just-in-time manufacturing for improving global supply chain resilience. *Int J Eng Technol Res Manag.* 2018 Nov;2(11):58. doi:10.5281/zenodo.15241789.

27. Gaber T, Ali T, Nicho M, Torky M. Robust Attacks Detection Model for Internet of Flying Things based on Generative Adversarial Network (GAN) and Adversarial Training. IEEE Internet of Things Journal. 2025 Apr 2.

28. Adeshina Yusuff Taofeek**.** Strategic implementation of predictive analytics and business intelligence for value-based healthcare performance optimization in US health sector. *International Journal of Computer Applications Technology and Research*. 2023;12(12):101–114. doi:10.7753/IJCATR1212.1014.

29. Saeeda U, Janb SU, Ahmadb J, Shaha SA, Alshehric MS, Ghadid YY, Pitropakisb N, Buchananb WJ. Generative adversarial networks-enabled anomaly detection systems: A survey [Internet].

30. Olayinka OH. Data driven customer segmentation and personalization strategies in modern business intelligence frameworks. *World Journal of Advanced Research and Reviews*. 2021;12(3):711-726. doi: https://doi.org/10.30574/wjarr.2021.12.3.0658

31. Bennour A, Elhoseny M, Veerasamy BD, Bhatt R, Agrawal A, Shukla PK, Ghabband F. An innovative framework to securely transfer data through the internet of things using advanced generative adversarial networks. Journal of High Speed Networks. 2025 Feb;31(1):3-27.

32. Chukwunweike J, Lawal OA, Arogundade JB, Alade B. Navigating ethical challenges of explainable AI in autonomous systems. *International Journal of Science and Research Archive*. 2024;13(1):1807–19. doi:10.30574/ijsra.2024.13.1.1872. Available from: https://doi.org/10.30574/ijsra.2024.13.1.1872.

33. Bethu S. Malicious Attack Detection in IoT by Generative Adversarial Networks. SN Computer Science. 2025 Apr 10;6(4):372.

34. de Araujo-Filho PF, Kaddoum G, Campelo DR, Santos AG, Macêdo D, Zanchettin C. Intrusion detection for cyber–physical systems using generative adversarial networks in fog environment. IEEE Internet of Things Journal. 2020 Sep 18;8(8):6247-56.

35. Balaji S, Narayanan SS. Dynamic distributed generative adversarial network for intrusion detection system over internet of things. Wireless Networks. 2023 Jul;29(5):1949-67.