Available online at www.HEFJournal.org

HUMAN, EARTH AND FUTURE Ital Publication

ISSN: 2785-2997

Journal of Human, Earth, and Future

Vol. 7, No. 2, June, 2026

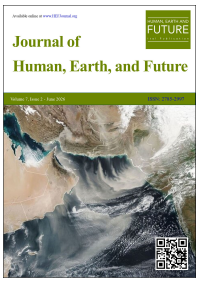

A Distributed Generative–Probabilistic Framework for Scalable Intrusion Detection in Smart Farming IoT Networks

Manikandan Thirumalaisamy [1, 2], Sumendra Yogarayan [1*], Siti Fatimah Abdul Razak [1], Md. Shohel Sayeed [1], Ramesh Shunmugam [2]

[1] *Faculty of Information Science and Technology, Multimedia University, Melaka 75450, Malaysia.*

[2] *Department of CSBS, Rajalakshmi Engineering College, Tamil Nadu 602105, India.*

Received 26 October 2025; Revised 22 April 2026; Accepted 27 April 2026; Published 01 June 2026

Abstract

Smart farming environments rely heavily on interconnected IoT–fog–cloud infrastructures, making them increasingly vulnerable to sophisticated cyber threats. The objective of this study is to develop a robust, scalable, and interpretable intrusion detection framework tailored for heterogeneous and resource-constrained agricultural IoT systems. The proposed method integrates a multi-stage generative–probabilistic pipeline: Principal Component Analysis (PCA) and Isolation Forest (iForest) operate at the fog layer for lightweight dimensionality reduction and early anomaly isolation, while the cloud layer employs an Intrusion Detection System Generative Adversarial Network (IDSGAN) model for advanced adversarial feature learning and Gaussian Mixture Models (GMM) for probabilistic attack-type clustering. Experimental analysis conducted on CIC IoT 2023 and CIC DIAD IoT 2024 datasets demonstrates strong detection performance, achieving Accuracy between 98.47–99.15%, F1-scores up to 0.996, and AUC values of 0.981–0.987, outperforming baseline IDS models. Clustering metrics including Silhouette (99.01%), ARI (98.97%), and NMI (98.91%) confirm highly coherent attack-grouping across major categories such as DDoS, Mirai, Spoofing, and Web-based intrusions. The novelty of this work lies in its distributed architecture combining edge-efficient processing with cloud-level generative learning, enabling low-latency, high-throughput, and interpretable detection suitable for real-world smart farming ecosystems. This framework thus offers a scalable and high-performing solution for securing agricultural IoT infrastructures.

Keywords: Distributed IoT Security; GAN; GMM; Isolation Forest; Fog-Cloud Computing; Anomaly Detection.

1. Introduction

The rapid proliferation of Internet of Things (IoT) devices in precision agriculture has transformed modern farming into a highly interconnected cyber–physical ecosystem, where soil sensors, irrigation controllers, unmanned aerial vehicles (UAVs), and cloud-based analytics systems collaborate to improve crop yield and resource efficiency [1, 2]. These IoT-enabled agricultural-computing networks continuously monitor soil nutrients, crop health, irrigation schedules, and environmental conditions, enabling data-driven and sustainable farm management practices [3]. However, the integration of heterogeneous and resource-limited devices substantially expands the attack surface. Compromised irrigation controllers may trigger crop damage, manipulated soil-sensor readings can misguide yield predictions, and distributed denial-of-service (DDoS) attacks on fog gateways can disrupt real-time decision-making

* Corresponding author: sumendra@mmu.edu.my

https://doi.org/10.28991/HEF-2026-07-02-05

processes [4, 5]. Unlike conventional IT infrastructures, agricultural IoT systems operate under stringent constraints—limited computational power, intermittent wireless connectivity, and strict latency requirements in control operations—which significantly undermine the effectiveness of traditional centralized intrusion detection systems (IDS) [6]. Furthermore, IoT network traffic is high-dimensional, heterogeneous, noisy, and non-stationary, while modern attacks increasingly exhibit sparse, probabilistic, and adaptive behaviors that often evade signature-based or shallow statistical detection methods [7, 8].

Existing intrusion detection efforts in smart agriculture predominantly rely on isolated machine learning models or shallow anomaly detection techniques, which struggle to address the distributed, resource-constrained, and probabilistic nature of cyberattacks in agricultural-computing environments [9, 10]. Lightweight IDS solutions deployed directly on IoT devices often compromise detection accuracy, resulting in undetected stealthy or low-frequency attacks, whereas cloud-centric deep learning models introduce significant latency and communication burdens, limiting their suitability for real-time response requirements [11, 12]. Moreover, most existing approaches do not integrate key functionalities—such as feature reduction, anomaly isolation, generative modeling, and probabilistic clustering—within a unified detection pipeline. This lack of cohesion leads to IDS frameworks that are either computationally infeasible for edge-level deployment or insufficiently adaptable to evolving attack behaviors [13]. Therefore, a clear research gap exists for an end-to-end distributed intrusion detection architecture that is resource-aware, capable of generative attack modeling, probabilistic in its decision-making, and inherently scalable across IoT, fog, and cloud layers.

Despite the notable progress in IoT security and machine learning–based IDS solutions, several critical gaps remain unaddressed. First, limitations in distributed adaptability: most existing IDS frameworks are either lightweight but constrained when deployed solely at the IoT device level, or accurate yet centralized, resulting in high latency and bandwidth consumption. Only a few works achieve truly distributed coordination across IoT, fog, and cloud layers while meeting real-world constraints on energy, computation, and response time [14]. Second, insufficient generative robustness: current anomaly detection techniques predominantly depend on supervised learning or shallow statistical models, making them vulnerable to adaptive and evasive cyberattacks. While generative methods such as GAN-based IDS have demonstrated potential, their adoption in agricultural IoT remains limited, and they are seldom integrated into a multi-stage, hybrid intrusion detection pipeline [15]. Third, the lack of probabilistic attack characterization: many IDS models can detect anomalies but fail to probabilistically categorize or cluster them under uncertainty, restricting interpretability, threat prioritization, and actionable mitigation in operational farming environments [16]. These limitations collectively highlight the necessity for a unified generative–probabilistic IDS that supports distributed feature reduction, robust anomaly isolation, adversarial modeling, and probabilistic clustering tailored for smart agriculture. As shown in Figure 1, modern smart farming ecosystems comprise heterogeneous IoT devices, fog nodes, and cloud platforms, all of which face diverse cyber threats such as data manipulation, DDoS attacks, unauthorized access, and disruptions caused by resource constraints in real-time decision-making.

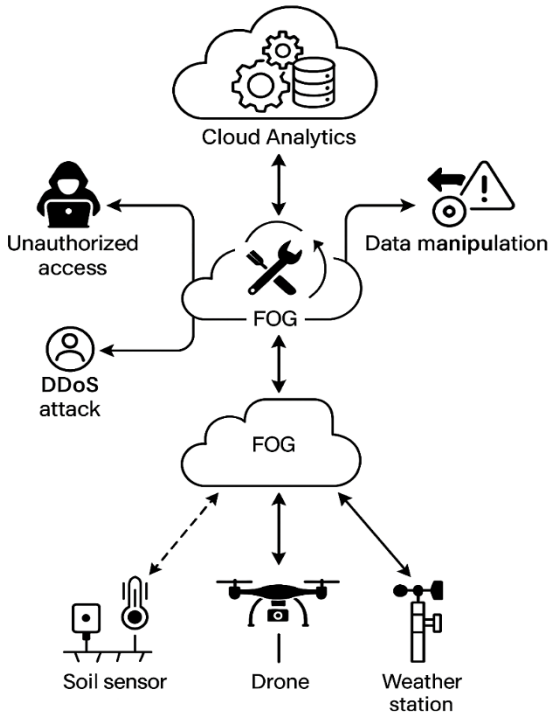

Figure 1. Smart farming IoT ecosystem with key devices data flows, and cybersecurity challenges

To address these challenges, this study introduces a distributed, multi-stage intrusion detection pipeline tailored for smart farming IoT environments. At the IoT device layer, lightweight statistical feature extraction minimizes redundancy, reduces communication overhead, and preserves key behavioral characteristics essential for early anomaly filtering. Fog nodes subsequently apply Principal Component Analysis (PCA) for compact representation of high-dimensional data and utilize Isolation Forest (iForest) for low-latency, energy-efficient anomaly isolation near the data source. Suspicious instances are then forwarded to the cloud layer, where an Intrusion Detection System Generative Adversarial Network (IDSGAN) performs advanced generative modeling, capturing complex and adversarial attack behaviors while enhancing detection robustness through synthetic augmentation.

Finally, a Gaussian Mixture Model (GMM) provides probabilistic clustering of anomalies into well-defined attack families, improving interpretability, uncertainty-aware classification, and prioritization of mitigation strategies. The key objectives are to develop a resource-aware IDS that operates collaboratively across IoT, fog, and cloud layers; incorporate generative intelligence for detecting adaptive threats; and enable probabilistic, interpretable clustering aligned with situational awareness needs in agriculture. By unifying distributed processing with generative–probabilistic modeling, the proposed framework achieves scalable, robust, and operationally practical intrusion detection for smart farming networks.

The key objectives of the proposed pipeline are as follows:

- Develop a resource-aware, distributed IDS that operates cohesively across IoT devices, fog nodes, and cloud platforms.
- Incorporate generative modeling to effectively capture adversarial, probabilistic, and evolving attack patterns.
- Enable probabilistic clustering to support interpretable classification and adaptive threat response.
- Demonstrate scalability and effectiveness using benchmark intrusion datasets tailored to agricultural IoT environments.

By integrating generative adversarial learning with probabilistic clustering in a distributed architecture, the proposed solution ensures robust intrusion detection and interpretable diagnosis, aligning security mechanisms with the operational constraints and real-time demands of smart agriculture.

The main contributions of this study are summarized as follows:

- A novel generative–probabilistic intrusion detection pipeline that distributes computational tasks across IoT, fog, and cloud layers, balancing resource efficiency with high detection fidelity.
- A hybrid integration of PCA-based dimensionality reduction, Isolation Forest-based anomaly isolation, and GAN-driven generative modeling, enabling adaptive detection against sophisticated and evolving cyber threats.
- A probabilistic attack characterization module using Gaussian Mixture Models (GMM), providing interpretable clustering and supporting risk-prioritized mitigation strategies.
- Comprehensive experimental validation on agricultural IoT intrusion datasets, demonstrating superior accuracy, reduced false alarms, and improved scalability compared with established IDS baselines.

Collectively, these contributions advance the state of the art in IoT intrusion detection by unifying distributed processing, generative modeling, and probabilistic analytics into a coherent pipeline optimized for smart farming networks.

The remainder of this paper is organized as follows: Section 2 reviews related work on IoT intrusion detection, hybrid architectures, and GAN-based security models. Section 3 details the proposed distributed, multi-stage detection framework. Section 4 presents the results and analytical evaluations. Section 5 discusses implications, limitations, and avenues for future research. Finally, the paper concludes with a summary of contributions and practical relevance.

2. Related Works

The rapid integration of IoT devices in smart agriculture has created highly interconnected cyber-physical ecosystems, where soil sensors, irrigation controllers, drones, and cloud analytics collaboratively optimize crop yield and resource utilization. This reliance on heterogeneous and resource-constrained devices, however, introduces a significantly expanded attack surface. Consequently, effective intrusion detection in these environments has become critical, prompting extensive research into machine learning (ML), deep learning (DL), and hybrid multi-stage frameworks tailored for IoT-based networks.

Recent studies highlight the importance of feature optimization and dimensionality reduction in enhancing detection efficiency. İleri (2025) [17] proposed a hybrid filter-based feature selection approach that integrates Chi-Square, Mutual Information, and Mean Absolute Deviation, demonstrating reduced computational overhead while maintaining high accuracy in smart farming networks. Zidi et al. (2024) [18] leveraged a downsized kernel partial least square (DKPLS) method combined with Kernel Extreme Learning Machines (KELM) to achieve near-perfect binary and multi-class detection in industrial IoT datasets. Similarly, Aburasain (2024) [19] employed Enhanced Black Widow Optimization for feature selection alongside a hybrid DCNN–DRNN model, achieving high accuracy with minimal false alarms in IoT-based smart farming. These approaches underscore the value of tailored feature selection methods for resource-constrained environments.

Multi-stage and hierarchical detection pipelines have been increasingly explored to address both binary and fine-grained multi-class classification. Ling et al. (2025) [20] introduced a two-stage framework using an Integrated Shared Feature Technique (ISFT) for feature reduction, followed by OC-SVM for normal vs. abnormal traffic detection and Decision Tree classifiers for attack-type classification, achieving over 98% accuracy. Polat et al. (2024) [21] and Alsagri (2025) [22] employed similar multi-stage pipelines combining deep learning for feature extraction with fast classifiers to detect complex DDoS and credit card fraud attacks, highlighting efficiency gains from staged processing.

Federated and distributed learning frameworks have been proposed to improve privacy, scalability, and resilience in distributed networks. Praharaj et al. (2025) [23] and Althunayyan et al. (2024) [24] demonstrated hierarchical federated learning systems that preserve local data privacy while collaboratively training anomaly detection models. Distributed GAN-based architectures, as introduced by Poongodi & Hamdi (2023) [25]and Yang et al. (2025) [26], generate synthetic adversarial traffic to improve robustness against rare and evolving attacks. GAN-LSTM hybrid models [27] further leverage temporal dependencies to detect sophisticated cyber-physical threats. Table 1 summarizes recent IDS approaches in smart farming IoT, showing methods, datasets, performance, and limitations, while highlighting gaps in real-world deployment and generalizability.

Despite these advances, significant gaps persist in the context of smart agriculture IoT. Existing solutions often remain either lightweight but limited to IoT devices or centralized and computationally intensive, hindering true distributed adaptability across IoT, fog, and cloud layers. Many models lack integration of generative adversarial learning and probabilistic clustering, leaving them vulnerable to adaptive threats and limiting interpretability. Furthermore, most IDS frameworks are dataset-specific, with limited validation under heterogeneous, resource-constrained, and dynamic agricultural environments. These observations underscore the need for an end-to-end, generative–probabilistic IDS pipeline that combines distributed feature reduction, adaptive anomaly detection, adversarial augmentation, and probabilistic classification to ensure robust, interpretable, and scalable intrusion detection in smart farming IoT networks.

In summary, while existing research has demonstrated significant progress through feature optimization, multi-stage architectures, federated learning, and GAN-based augmentation, a clear gap remains in developing an end-to-end distributed IDS that jointly leverages lightweight IoT-side preprocessing, fog-level anomaly isolation, cloud-driven generative adversarial modeling, and probabilistic clustering for interpretable diagnosis. These limitations in adaptability, interpretability, and cross-layer scalability directly motivate the proposed generative–probabilistic intrusion detection pipeline introduced in this study (see Table 1).

Table 1. Summary of Recent IDS Approaches in Smart Farming and IoT Environments

Ref	Methodology	Dataset	Performance	Key Contributions	Limitations
[17]	Hybrid filter-based feature selection (CS, MI, MAD), evaluated with Shallow & Deep ANN classifiers	Smart-Farm-IDS (172,800 samples, 21 features)	Shallow ANN: 7.6% reduced testing time; Deep ANN: 12.1%	Domain-specific dataset; hybrid feature selection; comparative ANN analysis	Limited to ANN models; dataset-specific; dependent on selected filters
[18]	Two-stage: DKPLS for dimensionality reduction + KELM classifier for binary & multi-class detection	X-IIOTID industrial IoT dataset	99.92% binary accuracy; 99.99% nine-class detection	Computationally efficient IDS; reduced false alarms; suitable for dynamic agriculture environments	Specifics on real-time adaptation not elaborated
[19]	Three-phase: AMMN preprocessing + EBWO feature selection + Hybrid Deep Learning (DCNN + DRNN)	BoT-IoT dataset (46 features, multi-class)	Accuracy 98.35%; F1-score 82.79%; Recall 80.95%; Precision 84.85%	EBWO-HDL model; combines advanced feature selection & hybrid deep learning	Single dataset; generalizability concerns
[23]	FedTDLR: Federated Transfer Learning + dynamic gradient compression; pre-trained models fine-tuned per farm	Two custom smart farming datasets	Outperformed traditional FL; MobileNetV2 efficient memory; EfficientNet-B0 highest accuracy 96.00%	FedTDLR framework; benchmark datasets; gradient compression for resource optimization	Validated on testbed-generated data; limited Non-IID scenario exploration
[28]	ML & DL models: MLP, Naïve Bayes, SVM, NN-RF hybrid; DL architectures for soil classification	Dry Beans dataset (13,611 images); Soil Type dataset (144 images)	MobileNetV2: accuracy 0.97, recall 0.97; SVM: 0.93; NN-RF hybrid: 92%	Demonstrated hybrid ML/DL effectiveness in IoT-enabled agriculture	Limited generalizability across datasets
[29]	GA-optimized BPNN; three-phase: filter-based feature selection + GA optimization + BPNN deployment on Fog nodes	UNSW-NB15, ToN_IoT	Execution time reduced 16.35%-37.07%; enhanced accuracy	GA-optimized BPNN; filter-based feature selection; Fog deployment	Generalizability to diverse attacks not addressed
[30]	Fog-edge architecture + Federated Learning; decentralized SVM	NSL-KDD, CICIDS2017	NSL-KDD: 98% accuracy; Recall/F1 improved 13.19%; CICIDS2017 Recall improved 7.57%	SVM-FL architecture; fog-edge deployment; real-time responsiveness	Optimization for constrained devices; adversarial attack resilience; data drift
[31]	Multi-stage: SAE for dimensionality reduction + CatBoost feature selection + Hybrid ensemble (Transformer, CNN, LSTM) with AGWO	NSL-KDD, UNSW-NB15, AWID	NSL-KDD: 99.7% accuracy, 0.996 F1-score; AWID: 99.9% accuracy	Combines feature selection, dimensionality reduction & ensemble deep learning in distributed setup	Requires real-world deployment validation; interpretability
[32]	GAN with self-attention: generator creates adversarial flows; discriminator learns from BlackBox IDS outputs	CICIDS2017	Adversarial flows reduced IDS detection rate by 15.93%; lower precision/recall/F1	Self-attention GAN for adversarial sample generation; strengthen IDS resilience	Experimental gap vs. real-world dynamic attacks
[25]	Multilevel distributed GAN with federated learning; peer-to-peer threat monitoring	KDD, SWAT	Accuracy 98.92%; F1-score 97.00%; Precision 99.00%; Recall 98.92%	Decentralized GAN architecture; high privacy; reduced communication overhead	Metrics & quantitative results not detailed; generalizability unknown
[26]	CE-GAN for synthetic minority sample generation + Nash equilibrium ensemble (RF, ET, GB)	NSL-KDD, UNSW-NB15	Improved classification of minority classes; robust multi-class IDS accuracy 99.71%	Integration of CE-GAN augmentation & game-theoretic ensemble	Data imbalance issue, evaluated on two datasets, and complexity of the model is relatively high
[27]	GAN generates rare attacks; LSTM for temporal dependencies; hybrid anomaly scoring	SWAT, WADI	SWAT: 99.99% accuracy; WADI: 98.12%; near-perfect recall/F1	Adversarial sample generation; temporal feature learning; superior detection vs. LSTM, CNN-LSTM, AE	High computational/resource demands
[24]	ANN for signature-based detection + LSTM Autoencoder for anomaly detection; hierarchical federated learning	Car-Hacking Dataset (4 attack types, 9 features)	F1-score 0.99 (seen); F1>0.95, DR 99.99% (unseen); FAR 0.016%	Multi-stage detection; LSTM-AE anomaly detection; H-FL framework	Evaluated on single dataset; generalizability limited
[22]	Two-stage: iForest anomaly detection + XGBoost classification	Kaggle Credit Card Fraud Detection Dataset (284,807 transactions, 492 frauds)	Accuracy 99.98%; Recall 95.8%; F1-score 93.6%	Two-stage architecture; improved minority fraud detection; handles high dimensionality & class imbalance	Implicit limitations from traditional methods
[21]	1D-CNN for feature extraction + Decision Tree for classification	IoTID20 dataset (46 features, 13 DDoS classes)	Accuracy 99.99%; F1-score 99.99%; Precision 99.98%; Recall 99.99%; minimal loss	Deep learning for feature engineering; fast classifier; multi-class DDoS detection	Validation needed on broader attacks
[33]	Two-phase: XGBoost/CNN for feature extraction + LSTM for classification	CIC IDS 2017, UNSW NB15, NSL KDD, WSN DS	CNN-LSTM: 98.55% accuracy (CIC IDS 2017 binary); outperformed other models	ML/DL integration for feature extraction & classification; high detection & low FAR	Difficulty detecting unseen attacks
[20]	ISFT for feature selection + two-stage classification OC-SVM, MHDEGAT	UNSW-NB15	accuracy of 96.99%, precision of 97.11%, recall of 96.99%, and F1 score of 96.93%	hierarchical detection GRL-HEGAT, and multi-classification intrusion detection model for MHDE GAT	Problem of low-sample attacks and the detection of unknown attacks

3. Material and Methods

This section presents the design and implementation of the proposed distributed, multi-stage generative–probabilistic intrusion detection framework for smart farming IoT systems. It begins by outlining the dataset and its adaptation to agricultural IoT contexts, followed by a detailed description of the distributed architecture across IoT devices, fog nodes, and cloud. Subsequently, each algorithmic stage of the pipeline—Principal Component Analysis (PCA), Isolation Forest (iForest), Generative Adversarial Networks (GAN), and Gaussian Mixture Models (GMM)—is described with mathematical formalism. Finally, the experimental design, evaluation metrics, and implementation details are presented.

3.1. Proposed IoT–Fog–Cloud Distributed Architecture for Generative–Probabilistic Intrusion Detection

The proposed intrusion detection framework is structured as a three-layer distributed architecture that follows the IoT–fog–cloud paradigm, as shown in Figure 2. At the IoT device layer, lightweight feature extraction and preprocessing are performed to generate compact statistical summaries such as packet counts, averages, and variances. This approach minimizes bandwidth usage and reduces energy consumption compared to transmitting raw traffic. The processed features are then transmitted to the fog layer, where Principal Component Analysis (PCA) reduces dimensionality and Isolation Forest (iForest) performs preliminary anomaly scoring. This stage ensures low-latency and resource-efficient anomaly filtering close to the data source, enabling timely responses without overwhelming local devices. Suspicious traffic flagged at the fog layer is subsequently forwarded to the cloud, where computationally intensive analysis is carried out. At this stage, Generative Adversarial Networks (GANs) are employed to refine anomaly detection by modeling complex adversarial traffic distributions, while Gaussian Mixture Models (GMMs) cluster anomalies into probabilistic groups corresponding to distinct attack families. In this way, the distributed architecture balances efficiency and robustness by combining early filtering at the edge with deep generative and probabilistic modeling in the cloud, ensuring both scalability and interpretability in securing smart farming IoT environments. This architecture ensures low-latency early filtering at the edge while enabling deep generative and probabilistic modeling at the cloud, where resources are abundant. The framework is also designed to handle Non-Independent and Non-Identically Distributed (non-IID) traffic data across heterogeneous farms. Local PCA and iForest processing preserves unique traffic patterns at each fog node, while the cloud-level IDSGAN and GMM models learn generalized representations, enabling robust detection across diverse datasets. This design ensures that both local variations and global patterns are effectively captured, maintaining high detection performance in distributed smart farming environments.

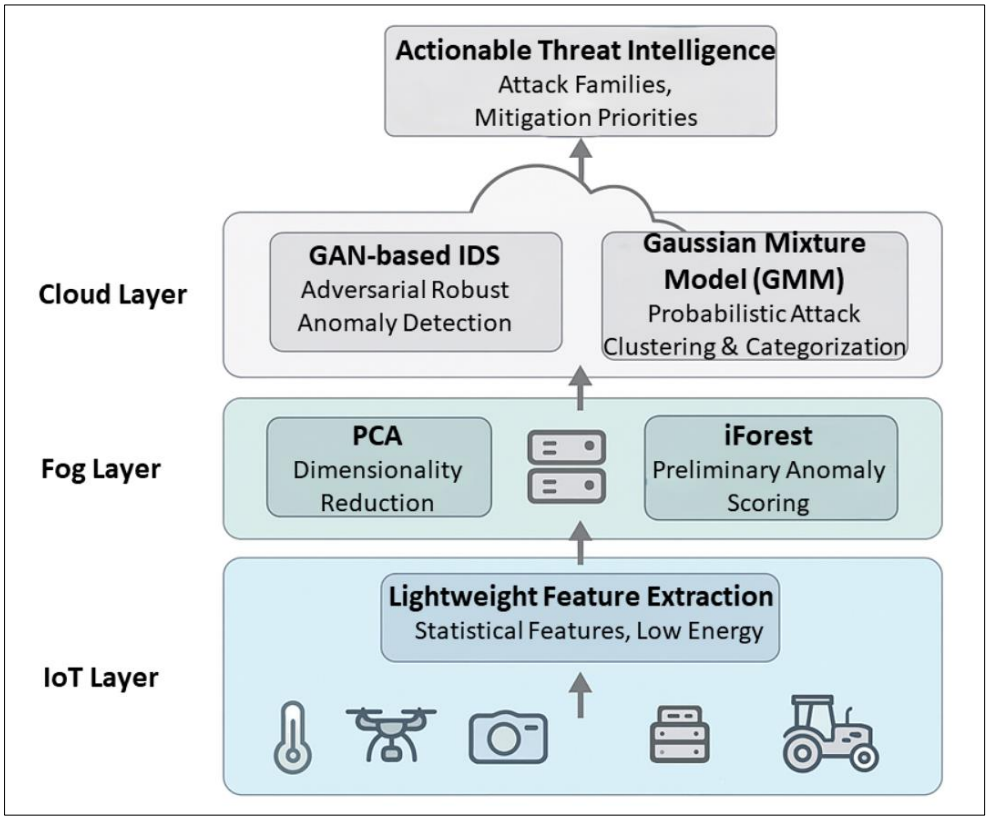

Figure 2. Proposed IoT–fog–cloud distributed architecture for generative–probabilistic intrusion detection in smart farming

3.2. Principal Component Analysis (PCA) for Dimensionality Reduction

The CIC IoT 2023 and CIC DIAD IoT 2024 datasets contain 39 and 83 flow-level attributes, many of which are highly correlated. Directly processing this high-dimensional feature space increases both computational cost and communication overhead, particularly in resource-constrained fog environments. To address this, Principal Component Analysis (PCA) is applied at the fog layer to generate compact feature representations while retaining most of the variance in the data [34].

Let the dataset be $D \in \mathrm{R}^{n\times d}$, where n is the number of samples and d the feature dimensionality. PCA identifies the principal components by solving the eigenvalue problem:

$$\Sigma v_i = \lambda_i v_i \tag{1}$$

where $\Sigma = \frac{1}{n} D^{\top} D$ is the covariance matrix, v_i are the eigenvectors, and λ_i the corresponding eigenvalues. The top-k eigenvectors are used to form the transformation matrix V_k. The dimensionality-reduced representation is then given by:

$$Z = DV_k, Z \in \mathrm{R}^{n\times k}. \tag{2}$$

To ensure minimal information loss, k is chosen such that:

$$\frac{\sum_{i=1}^{k} \lambda_i}{\sum_{j=1}^{d} \lambda_j} \geq 95\%. \tag{3}$$

This criterion guarantees that at least 95% of the original variance is preserved. By applying PCA at the fog layer, the system achieves an efficient balance between data fidelity and computational feasibility, enabling faster anomaly detection in the subsequent iForest stage while reducing communication load between IoT devices, fog nodes, and the cloud.

3.3. Isolation Forest (iForest) for Anomaly Scoring

At the fog layer, anomaly detection is carried out using the Isolation Forest (iForest) algorithm, which is well-suited for high-dimensional, large-scale data in resource-constrained environments. iForest operates on the principle that anomalies are "few and different," making them easier to isolate compared to normal data points [35].

An iForest consists of an ensemble of isolation trees (iTrees). Each iTree partitions the data recursively by randomly selecting a feature and a corresponding split value. For a given sample z, the path length $h(z)$ denotes the number of splits required to isolate it. Since anomalies typically lie in sparse regions of the feature space, they are isolated with fewer splits, resulting in shorter path lengths.

The anomaly score of a sample z is computed as:

$$s(z,n) = 2^{-\frac{E[h(z)]}{c(n)}} \tag{4}$$

where, $E[h(z)]$ is the expected path length of $x$across the ensemble of iTrees, n is the number of samples, and $c(n)$ represents the average path length of unsuccessful searches in a binary tree.

A threshold τ is defined such that:

$$s(z,n) \geq \tau \Rightarrow z \text{ is flagged as an anomaly.}$$

Flows exceeding the threshold are transmitted from the fog layer to the cloud for deeper adversarial analysis and probabilistic clustering. This design ensures that the fog layer acts as a computationally efficient filter, reducing bandwidth and cloud processing overhead while providing rapid anomaly scoring close to the data source (Figure 3).

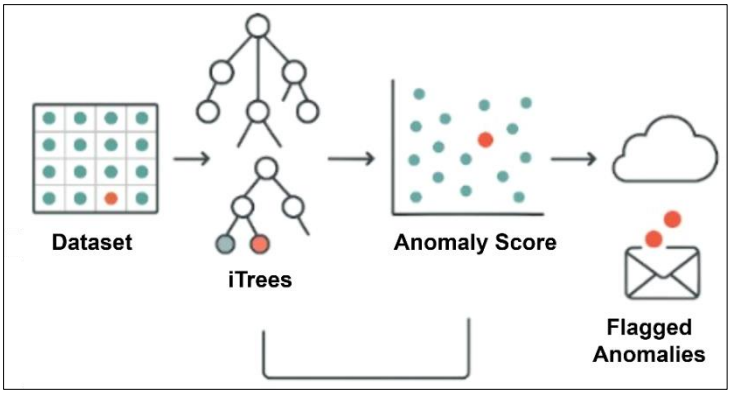

Figure 3. Isolation Forest workflow for anomaly detection, using recursive data partitioning to isolate anomalies at fog layer

3.4. IDSGAN-Based Intrusion Detection

At the cloud layer, advanced anomaly detection is performed using IDSGAN (Intrusion Detection System Generative Adversarial Network), a GAN variant tailored for intrusion detection [36]. The framework consists of a generator (G) and a discriminator (D) trained in a minimax game:

$$\min_{Gr} \max_{Dr} V(Dr, Gr) = \mathrm{E}_{z\sim p_{\mathrm{data}}(z)}[\log Dr(z)] + \mathbb{E}_{z\sim p_z(z)}[\log(1 - Dr(Gr(z)))] \quad (5)$$

- The generator $Gr(z)$maps random noise $z$to synthetic traffic samples, approximating the distribution of benign IoT network flows.
- The discriminator $Dr(z)$distinguishes real benign samples from generator outputs, effectively learning the normal traffic patterns.

For anomaly detection, the reconstruction loss is employed. Given a test sample z, the generator reconstructs it as $\hat{z}$, and the anomaly score is computed as:

$$A(z) = \alpha \parallel \mathrm{z} - \hat{z} \parallel_2 + (1 - \alpha)(1 - Dr(z)) \quad (6)$$

where, α balances the contribution of reconstruction error and discriminator confidence. Samples with high $A(z)$ values are flagged as potential attacks.

By modeling complex, non-linear traffic distributions and generating adversarial variants, IDSGAN enhances robustness against adaptive intrusions, making it well-suited for detecting stealthy attacks in smart farming IoT networks. Its integration (Figure 4) within the cloud layer ensures resource-intensive processing is offloaded from IoT and fog layers, maintaining low-latency and energy-efficient operations at the edge.

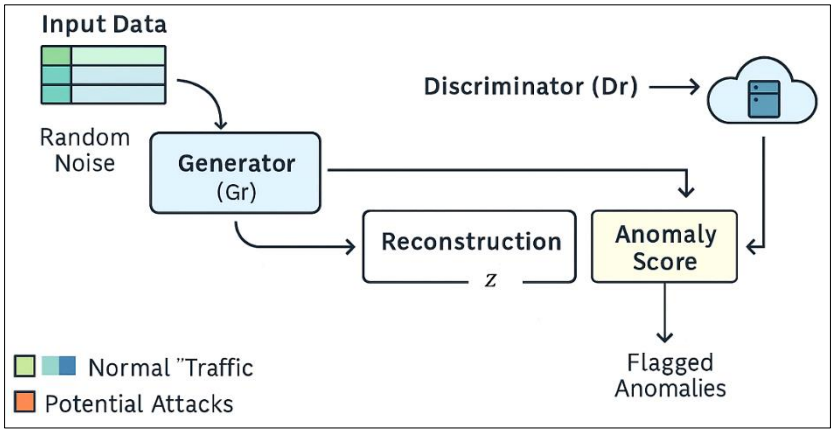

Figure 4. IDSGAN-based cloud IDS, showing adversarial generation, reconstruction, and anomaly scoring of IoT network flows

3.5. Gaussian Mixture Models (GMM) for Probabilistic Clustering

In the final stage, Gaussian Mixture Models (GMMs) cluster detected anomalies into interpretable attack categories. GMM assumes that anomaly data arises from a mixture of multiple Gaussian distributions, each representing a distinct attack type [37]. This probabilistic modeling accommodates heterogeneous and overlapping anomaly patterns typical of smart farming IoT networks [38].

Each Gaussian component is defined by its mean (μ_k), covariance (Σ_k), and mixture weight (π_k). The overall probability of a data point $x$is given by:

$$p(z) = \sum_{k=1}^{K} \pi_k \mathcal{N}(z \mid \mu_k, \Sigma_k) \quad (8)$$

Parameters are estimated using the Expectation–Maximization (EM) algorithm, which iteratively computes the probability of each anomaly belonging to each cluster and updates the component parameters until convergence.

The resulting clusters correspond to distinct attack families, such as flooding, injection, or advanced persistent threats, enabling prioritized response and actionable insights. By providing probabilistic characterization, GMM enhances interpretability and aligns with threat intelligence frameworks like MITRE ATT&CK, supporting informed mitigation in distributed smart farming environments.

4. Results and Analysis

This section presents a detailed evaluation of the proposed multi-stage generative–probabilistic intrusion detection framework, highlighting its effectiveness, scalability, and interpretability across multiple performance dimensions.

The analysis begins with an overview of the datasets utilized to validate the framework, followed by the application of PCA for dimensionality reduction to ensure computational efficiency. The detection performance is then systematically assessed using standard evaluation metrics, supported by visual representations such as confusion matrices and ROC curves. Comparative analysis with baseline techniques and state-of-the-art models is provided to demonstrating the superior performance of the proposed approach. Additionally, the deployment feasibility of the framework is examined within a distributed IoT–Fog–Cloud environment, emphasizing its adaptability and real-world applicability. Finally, clustering-based interpretability using GMM is explored to map detected anomalies to representative attack categories, thereby enhancing explainability and facilitating actionable cyber threat insights.

4.1. Dataset Description

To evaluate the effectiveness and generalizability of the proposed intrusion detection framework, two comprehensive benchmark datasets—CIC IoT 2023 [39] and CIC DIAD IoT 2024 [40]—were utilized. These datasets capture realistic IoT network traffic encompassing both benign and malicious scenarios, providing a diverse range of attack categories essential for robust anomaly detection and adversarial behavior analysis in distributed IoT environments. A detailed summary of these datasets, including the number of samples, benign instances, feature dimensions, and feature types, is presented in Table 2.

Table 2. Summary of Benchmark Datasets

Dataset	Year	Attack Categories	Total Samples	Benign Samples	Features	Feature Type
CIC IoT 2023	2023	DDoS, DoS, Reconnaissance, Web-Based, Mirai, Spoofing, Brute Force, Benign	48,788,784	1,098,191	39	Packet & Header-level (protocol flags, counts, rates, min/max/avg/std)
CIC DIAD IoT 2024	2024	DDoS, DoS, Mirai, Reconnaissance, Spoofing, Web-Based, Brute Force, Benign	12,169,400	398,330	83	Flow-based (packet lengths, IATs, segment sizes, bulk rates, active/idle stats, flags)

The datasets encompass diverse and evolving threat vectors, each serving a distinct role in evaluating the framework's capabilities. DDoS and DoS floods (e.g., SYN, UDP, ICMP, Slowloris) assess scalability under high-volume network stress; Reconnaissance scans (port, OS, vulnerability) evaluate sensitivity to stealthy probing activities; Web-based exploits (such as SQL Injection, Cross-Site Scripting, and Command Injection) target application-layer vulnerabilities; Botnet (Mirai) traffic introduces IoT-specific adversarial patterns; Spoofing attacks (ARP, DNS) simulate identity manipulation scenarios; Brute-force attempts represent persistence-driven intrusions; and benign traffic establishes a baseline for ensuring minimal false alarms. This diverse composition ensures that the framework is tested against a comprehensive spectrum of real-world threats, enabling accurate assessment of detection robustness, generalization, and scalability across heterogeneous IoT ecosystems.

4.2. Hyperparameter Configuration

The hyperparameters for each module in the proposed multi-stage intrusion detection pipeline were carefully tuned to balance detection accuracy, computational efficiency, and stability across distributed layers. PCA components were selected to retain ≥95% variance, while iForest used 150 estimators with a 0.05 contamination rate for effective anomaly isolation at the fog layer. At the cloud layer, IDSGAN was optimized using Adam (lr = 0.0002) over 89 epochs for stable adversarial training, and GMM employed five Gaussian components with full covariance to cluster anomalies into attack families. The complete configuration is summarized in Table 3, ensuring reproducibility and consistency across experiments.

Table 3. Hyperparameter configuration of the proposed multi-stage generative–probabilistic IDS

Algorithm	Key Parameters	Values/Settings
PCA	Components (k)	25 (≥95% variance)
	SVD Solver	full
Isolation Forest	n_estimators	150
	Contamination	0.05
	Max Samples	256
IDSGAN	lr (Adam)	0.0002
	Epochs	89
	Batch Size	1024
	α (Anomaly Score)	0.6
GMM	Components (K)	5
	Covariance Type	full

4.3. Performance Analysis of PCA on Feature Reduction

To ensure efficient representation and scalability, Principal Component Analysis (PCA) was applied to the encoder's latent features to evaluate its capability in reducing dimensionality while preserving critical variance. Table 4 reports the variance explained by the top three principal components, which collectively capture approximately 88.8% of the total variance. This indicates that PCA provides a compact yet highly informative representation, minimizing information loss while significantly reducing the feature space.

Table 4. Variance captured by PCA components

Principal Component	Variance Explained (%)	Cumulative Variance (%)
PC 1	62.4	62.4
PC 2	17.8	80.2
PC 3	8.6	88.8

To visualize the discriminative power of the reduced representation, the encoder's latent outputs were projected into a three-dimensional PCA space. Figure 5 illustrates clear separation between benign (sky blue) and attack (crimson) samples. Benign traffic forms dense, compact clusters, while attack instances occupy more dispersed and distinct regions, reflecting inherent behavioral variations.

This separation confirms two critical aspects: (i)Variance Preservation – The first three components capture the majority of relevant information, ensuring that essential patterns remain intact. (ii) Structural Interpretability – The spatial distinction in the latent space enables anomaly detection models to more effectively discriminate between normal and malicious traffic. Overall, the PCA analysis validates its role as an effective dimensionality reduction technique, enhancing scalability, noise resilience, and downstream detection accuracy in the proposed multi-stage pipeline.

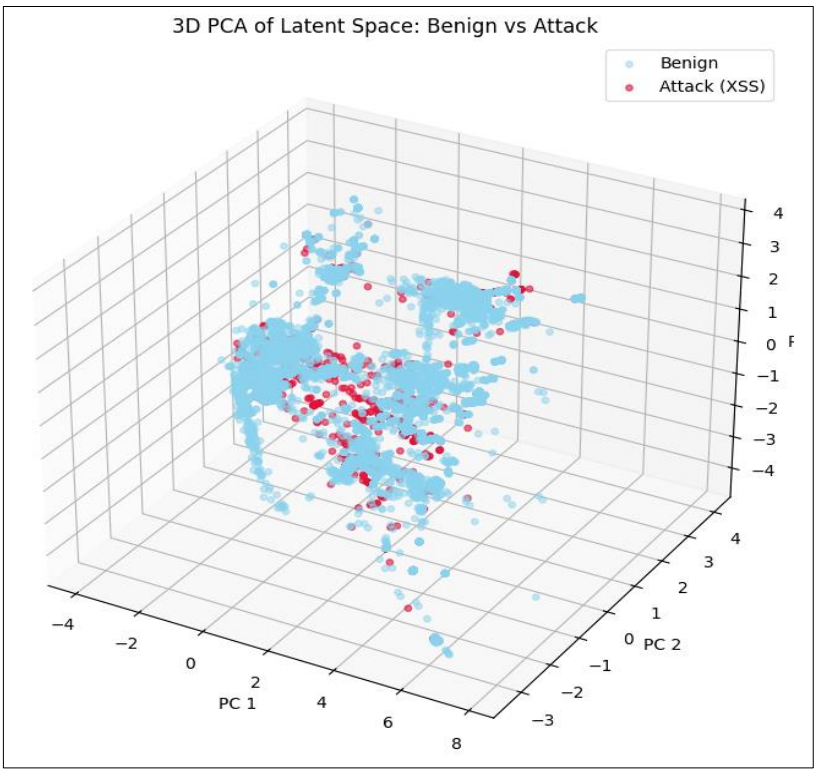

Figure 5. 3D-PCA visualization of the encoder's latent space, showing benign and attack sample distributions

To further interpret these findings, it is important to emphasize that the strong separability in the PCA-transformed space directly supports downstream anomaly detection. The compact clustering of benign traffic indicates stable operational behavior in smart farming IoT systems, whereas the dispersed attack points highlight variability typical of malicious flows such as DDoS bursts, spoofing attempts, and reconnaissance scans. This structural contrast in the low-dimensional space confirms that PCA not only reduces dimensionality but also enhances the discriminability of latent features. Such improved representation is essential in distributed environments, where fog nodes must execute rapid and lightweight analysis without compromising accuracy.

4.4. Detection Performance on the Benchmarked Datasets

The proposed multi-stage intrusion detection pipeline demonstrates consistently high performance across both benchmark datasets—CIC IoT 2023 and CIC DIAD IoT 2024. Evaluation metrics, summarized in Table 5, indicate excellent classification capability, with high accuracy, precision, recall, and F1-score. For a detailed view, Figure 6 presents the binary confusion matrices: 6(a) for CIC DIAD IoT 2024 and 6(b) for CIC IoT 2023, illustrating the number of correctly and incorrectly classified benign and attack samples.

Table 5. Detection Metrics of Proposed Pipeline on CIC DIoT 2023 and CIC DIoT 2024

Metric	CIC IoT DIAD 2024	CIC IoT 2023
Accuracy	98.47%	99.15%
Precision	99.94%	99.97%
Recall	98.48%	99.16%
F1-score	99.20%	99.56%
Specificity	98.20%	98.60%
AUC-ROC	0.981	0.987

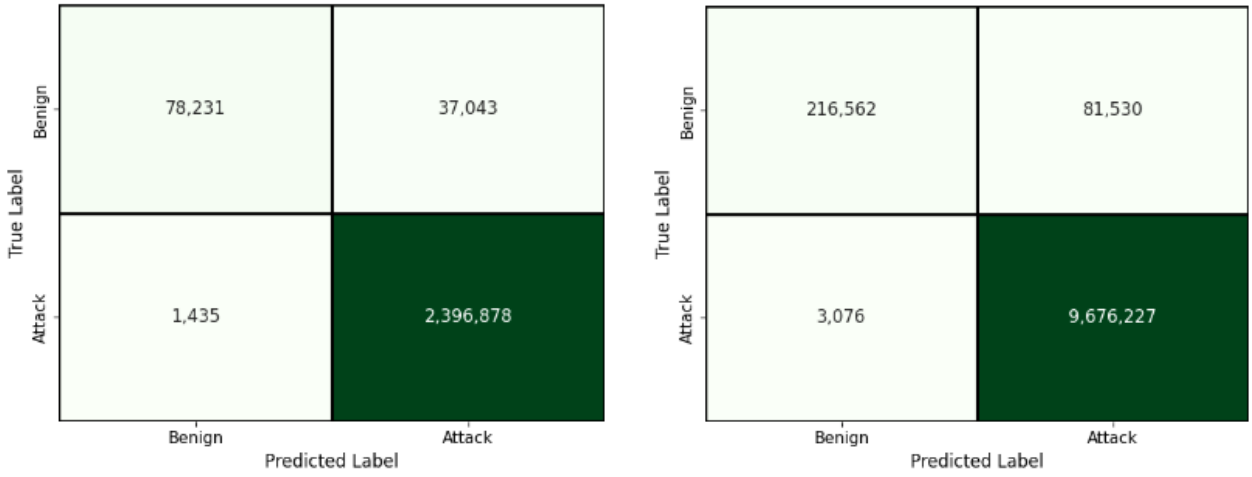

(a) (b)

Figure 6. Binary Confusion Matrices of the Proposed Framework on CIC IoT 2023 and CIC DIAD IoT 2024

Confusion matrices show the true positives, true negatives, false positives, and false negatives for the proposed pipeline on (a) CIC IoT 2023 and (b) CIC DIAD IoT 2024 datasets. High values along the diagonal indicate strong classification accuracy, demonstrating effective separation between benign and malicious traffic flows. The results confirm that the proposed framework maintains robust detection capabilities across heterogeneous IoT traffic. The CIC IoT 2023 dataset shows slightly higher overall metrics due to its larger volume and balanced feature diversity, enabling more comprehensive training. Meanwhile, CIC DIAD IoT 2024 still delivers strong generalization despite its higher attack variety and flow-level complexity, validating the adaptability and resilience of the proposed pipeline across evolving IoT threat landscapes.

These results collectively demonstrate that the proposed pipeline remains highly resilient across datasets with different traffic patterns, attack densities, and feature distributions. The consistently high recall values confirm that the system can reliably identify diverse attacks with minimal missed detections, which is crucial in smart-farming scenarios where even a single unmitigated intrusion may compromise yield, sensors, or automated irrigation systems. Similarly, the high precision indicates that the pipeline does not generate unnecessary alerts, reducing operational workload for farm administrators. These interpretations validate the robustness of the multi-stage architecture in real-world IoT deployments.

4.5. Comparative Evaluation with Baseline IDS Models

The proposed distributed framework was benchmarked against widely adopted IDS models: Random Forest, LSTM-based IDS, and Autoencoder-only IDS. Table 6 presents the comparative results on both CIC IoT 2023 and CIC DIAD IoT 2024, demonstrating that the proposed approach consistently outperforms baseline methods across datasets.

Table 6. Comparison of Baseline IDS Models and the Proposed Framework on CIC IoT 2023 and CIC DIAD IoT 2024

MODEL	CIC IoT 2023					CIC DIAD IoT 2024				
	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	AUC-ROC	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	AUC-ROC
Random Forest	92.13	90.87	91.24	91.09	0.927	91.54	90.14	90.71	90.43	0.921
LSTM-based IDS	94.35	93.57	94.07	93.71	0.948	93.82	93.09	93.34	93.14	0.943
Autoencoder-only IDS	95.09	94.24	94.84	94.52	0.954	94.62	93.91	94.22	94.08	0.949
Proposed Framework	99.15	99.97	99.16	99.56	0.987	98.47	99.94	98.48	99.20	0.981

The consistent improvements across both datasets confirm the robustness and generalizability of the proposed framework. By outperforming traditional and deep-learning baselines, it demonstrates strong adaptability to heterogeneous IoT environments, ensuring reliable performance in real-world intrusion detection scenarios. The ROC curves (Figure 7) further highlight the framework's superiority: Figure 7(a) for CIC IoT 2023 and Figure 7(b) for CIC DIAD IoT 2024. For both datasets, the ROC curves of the proposed model consistently remain closer to the top-left corner compared to the baselines, indicating a higher true positive rate and lower false positive rate. AUC values above 0.97 reinforce its capability to accurately discriminate between benign and malicious traffic with minimal trade-offs.

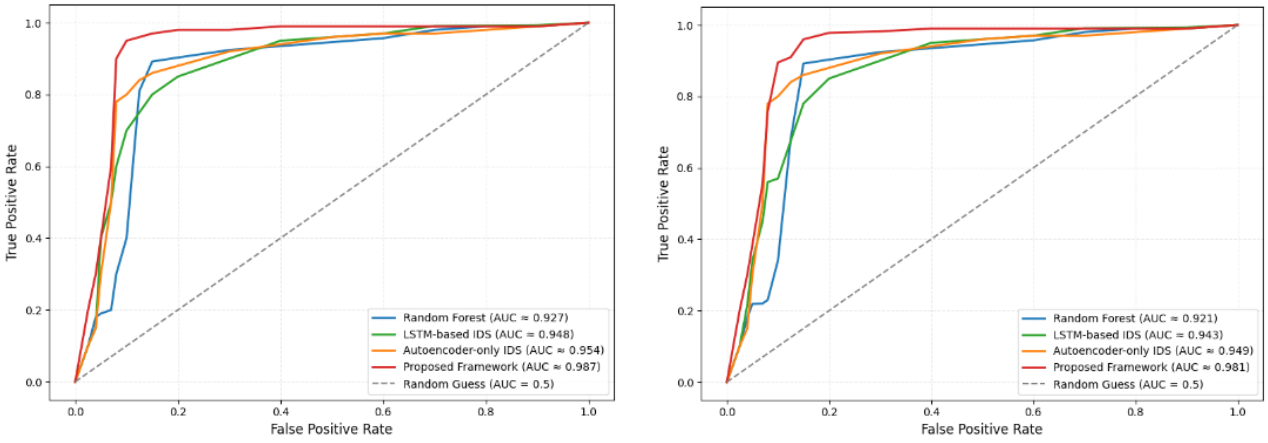

(a) (b)

Figure 7. ROC Curves for Proposed Framework vs Baseline Models on CIC IoT 2023 and CIC DIAD IoT 2024

The comparative findings further highlight the effectiveness of combining generative modeling with probabilistic clustering. Traditional models like Random Forest and LSTM-based IDS struggle to generalize across evolving IoT behaviors, whereas the proposed framework maintains significantly higher F1-scores across both datasets. This performance gain can be attributed to IDSGAN's ability to learn complex attack distributions and produce synthetic adversarial variations, which strengthens the model's resilience to unseen threats. Additionally, the ROC curves illustrate the advantages of the multi-stage design by consistently achieving higher true-positive rates at lower false-positive thresholds.

4.6. Distributed Deployment Performance and Scalability Evaluation

The computational efficiency of the proposed IDS was systematically evaluated across the IoT, fog, and cloud layers to assess both scalability and resource utilization. In this distributed deployment, IoT devices perform lightweight feature extraction, the fog layer executes dimensionality reduction and preliminary anomaly scoring using PCA combined with Isolation Forest, and the cloud layer handles resource-intensive generative modeling and probabilistic clustering with GAN and GMM.

Figure 8 illustrates the performance metrics across the three layers, including latency, energy consumption, and CPU utilization. IoT devices exhibit minimal latency (4.2 ms), low energy consumption (0.8 J), and modest CPU utilization (11%), demonstrating efficient local processing with minimal edge overhead. Fog nodes, responsible for PCA and anomaly scoring, incur moderate computational overhead (latency 12.6 ms, energy 2.1 J, CPU 23%). Cloud servers, executing complex generative and clustering algorithms, experience higher latency (45.4 ms), elevated energy usage (6.5 J), and increased CPU utilization (38%), consistent with the demands of large-scale probabilistic analysis.

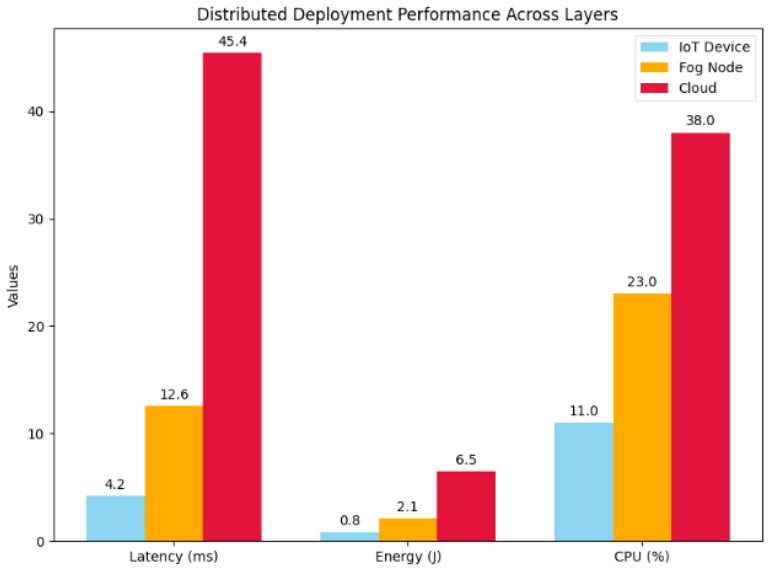

Figure 8. Distributed Deployment Performance Across IoT, Fog, and Cloud Layers

Throughput per layer further highlights the advantages of this distributed design: IoT devices process approximately 450 flows per second, fog nodes handle around 1,200 flows per second, and cloud servers manage approximately 3,500 flows per second. These results confirm that computational load is effectively balanced across layers, enabling real-time anomaly detection while optimizing resource allocation.

The distributed deployment demonstrates that the proposed framework achieves a favorable trade-off between efficiency and detection capability. Lightweight edge processing minimizes local resource consumption, fog-layer analysis ensures rapid anomaly scoring, and cloud-level generative modeling supports scalable, high-fidelity probabilistic detection. Collectively, these findings underscore the framework's suitability for real-world smart farming IoT environments where both performance and scalability are critical.

The observed computational distribution confirms that the system is well-suited for real-time agricultural IoT infrastructures. IoT devices remain minimally burdened, extending battery life and supporting long-term field deployments. Fog nodes effectively balance overhead and detection responsiveness, acting as an intermediate decision layer to filter anomalies before cloud transmission. The cloud layer's higher resource usage is justified, as it handles complex modeling tasks that benefit from centralization. Overall, the layered analysis demonstrates that the proposed design is not only scalable but also operationally feasible for large-scale, sensor-intensive smart-farming systems.

4.7. Attack-Type Detection and Interpretability

In the final stage of the proposed multi-stage intrusion detection framework, Gaussian Mixture Models (GMMs) are employed to cluster the detected anomalies into interpretable attack families. This stage provides insight into the nature of detected anomalies, supporting actionable threat intelligence mapping and enhancing overall interpretability. The clustering performance is evaluated using standard metrics, including Precision, Recall, Specificity, F1-score, Accuracy, and False Negative Rate (FNR) for each attack type, along with clustering quality metrics: Silhouette Score, Adjusted Rand Index (ARI), Normalized Mutual Information (NMI), and Davies–Bouldin Index (DBI).

Seven representative attack types—BruteForce, DDoS, DoS, Mirai, Recon, Spoofing, and Web-Based attacks—were selected from the CIC IoT 2023 and CIC DIAD IoT 2024 datasets. These attacks represent a range of behaviors, including volumetric attacks, integrity-based manipulations, reconnaissance activity, and command-level intrusions. Table 7 summarizes the detection metrics for each attack type, highlighting the framework's high accuracy and low false negative rates across all categories.

Table 7. GMM Clustering Performance Metrics for Representative Attack Types

Metric	BruteForce	DDoS	DoS	Mirai	Recon	Spoofing	Web-Based
Precision	0.9631	0.9980	0.9990	0.9752	0.9837	0.9585	0.9349
Recall	0.9724	0.9868	0.9896	0.9934	0.9908	0.9927	0.9881
Specificity	0.9997	0.9823	0.9810	0.9889	0.9818	0.9830	0.9980
F1-Score	0.9677	0.9924	0.9943	0.9842	0.9872	0.9753	0.9607
Accuracy	0.9994	0.9864	0.9892	0.9903	0.9865	0.9858	0.9978
False Negative Rate (FNR)	0.0276	0.0132	0.0104	0.0066	0.0092	0.0073	0.0119

To quantitatively evaluate the quality of the anomaly clusters, four standard clustering metrics were computed (Table 8): Silhouette Score, Adjusted Rand Index (ARI), Normalized Mutual Information (NMI), and Davies–Bouldin Index (DBI). High Silhouette, ARI, and NMI values, together with a low DBI, indicate that the GMM effectively groups anomalies into well-separated and internally cohesive clusters.

Table 8. GMM Clustering Quality Metrics

Metric	Value	Interpretation
Silhouette Score (↑)	99.01%	High cohesion within clusters
Adjusted Rand Index (ARI ↑)	98.97%	Excellent agreement with ground truth
Normalized Mutual Information (NMI ↑)	98.91%	Well-separated clusters; high mutual information
Davies–Bouldin Index (DBI ↓)	0.0091	Low value indicates better cluster separation

The multiclass confusion matrix (Figure 9) illustrates the predicted cluster assignments versus the true attack types. High values along the diagonal indicate accurate clustering, while off-diagonal entries correspond to misclassified anomalies. This visualization provides an intuitive assessment of the interpretability and accuracy of the clustering stage.

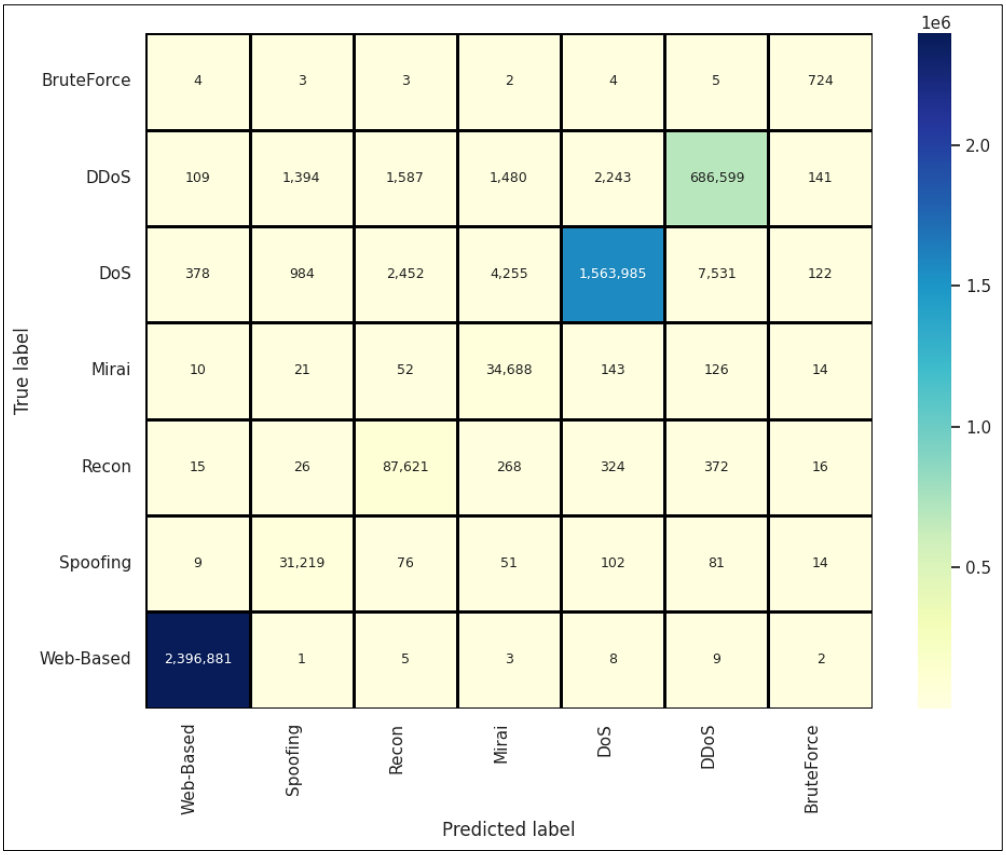

Figure 9. GMM Clustering Evaluation via Multiclass Confusion Matrix

The results confirm that the GMM stage effectively organizes detected anomalies into interpretable attack families with high accuracy, low FNR, and excellent clustering quality. By selecting representative attack types, the framework avoids sparsity issues while demonstrating generalization across heterogeneous IoT datasets. These findings support actionable threat intelligence mapping, such as leveraging frameworks like MITRE ATT&CK, and reinforce the framework's suitability for distributed smart farming IoT environments.

The interpretability of the clustering results is particularly valuable for automated threat response systems. Distinct clusters for attacks such as DDoS, Mirai, and BruteForce allow the system to map anomalies to known threat categories and support the integration of threat-intelligence frameworks like MITRE ATT&CK. Low false-negative rates across all attack types indicate that critical malicious behaviors are rarely missed, while high ARI and NMI values confirm strong agreement with ground truth labels. These interpretations show that the GMM-based clustering not only enhances transparency but also enables actionable prioritization for farm-network security teams.

4.8. Comparative Performance of the Proposed Framework

Table 9 presents a comparative analysis of the proposed multi-stage generative–probabilistic intrusion detection framework against 17 representative IDS approaches reported in the literature. The comparison highlights key numerical metrics such as Accuracy, F1-score, Precision, Recall, and AUC where available. Across diverse datasets, including Smart-Farm-IDS, X-IIOTID, BoT-IoT, CIC IoT 2023, and CIC DIAD IoT 2024, the proposed framework consistently demonstrates superior detection performance. Notably, it achieves an accuracy range of 98.47–99.15%, F1-scores between 0.992–0.996, and AUC values of 0.981–0.987, surpassing both shallow and deep learning-based IDS models as well as federated and hybrid approaches.

The proposed framework's key advantages stem from its distributed architecture, which balances computational load across IoT, fog, and cloud layers, ensuring both real-time responsiveness and energy efficiency. The integration of IDSGAN for generative adversarial modeling enables robust detection of adaptive and unseen attacks, while the Gaussian Mixture Model (GMM) probabilistic clustering allows interpretable attack-type classification and facilitates actionable mitigation. In contrast, existing literature methods largely rely on single-stage detection, shallow ML/DL models, or isolated federated learning techniques, which often compromise on scalability, interpretability, or robustness against adversarial and probabilistic threats.

Despite its superior performance, the proposed framework introduces additional system complexity due to its multi-stage design and requires careful resource management at the cloud layer. Nonetheless, the framework addresses critical gaps identified in smart agriculture IDS literature, including distributed adaptability, generative robustness, and probabilistic characterization. Collectively, these findings validate the effectiveness, scalability, and interpretability of the proposed approach, establishing it as a comprehensive and practical solution for IoT-enabled smart farming environments.

Table 9. Comparative Performance of the Proposed Framework with Existing Literature

Reference	Dataset	Methodology	Accuracy (%)	F1-score (%)	Precision (%)	Recall (%)	Key Insights
Ileri (2025) [17]	Smart-Farm-IDS	Hybrid FS + ANN	95.00- 97.00	51.57	100	34.74	Poor recall; dataset-specific ANN
Zidi et al. (2024) [18]	X-IIOTID	DKPLS + KELM	99.92	99.92	100	99.84	High accuracy; lacks real-time adaptability
Aburasain (2024) [19]	BoT-IoT	EBWO + Hybrid DL	98.35	82.79	84.85	80.95	Lower recall; single dataset
Praharaj et al. (2025) [23]	Custom Smart Farms	FedTL + Grad. Compression	96.00	97.00	–	–	Limited Non-IID testing; testbed-based
Aldossary (2024) [28]	Dry Beans, Soil Type	ML/DL hybrids	97.00	95.00	95.00	97.00	Limited cross-dataset generalization
Mohamed & Ismael (2023) [29]	UNSW-NB15, ToN_IoT	GA-BPNN + Fog	96.47	96.47	96.53	96.47	Diverse attack generalization not shown
Tariq et al. (2024) [30]	NSL-KDD, CICIDS2017	Fog-edge + FL + SVM	98.00	97.00	98.00	97.00	Device constraints; adversarial resilience
Tawfik (2024) [31]	NSL-KDD, UNSW-NB15, AWID	SAE + CatBoost + Ensemble	99.70–99.90	99.80	99.00	99.00	High accuracy; lacks interpretability
Aldhaheri & Alhuzali (2023) [32]	CICIDS2017	Self-attention GAN	97.87	98.98	98.41	95.36	Needs real-world validation
Poongodi & Hamdi (2023) [25]	KDD, SWAT	Multilevel GAN + FL	98.92	97.00	99.00	98.92	Limited metrics; generalizability unclear
Yang et al. (2025) [26]	NSL-KDD, UNSW-NB15	CE-GAN + Ensemble	99.71	99.83	99.81	99.85	Complex; imbalance handling focus
Md Shakil Siddique et al. (2025) [27]	SWAT, WADI	GAN-LSTM hybrid	98.12–99.99	91.00	97.00	99.00	High resource demands
Althunayyan et al. (2024) [24]	Car-Hacking	ANN + LSTM-AE + H-FL	98.59	95.00-99.00	91.00	99.99	Single dataset; limited scope
Alsagri (2025) [22]	Credit Card Fraud	iForest + XGBoost	99.98	93.60	–	95.8	Traditional method limitations
Polat et al. (2024) [21]	IoTID20	1D-CNN + Decision Tree	99.99	99.99	99.98	99.99	High accuracy; single dataset
Sajid et al. (2024) [33]	CIC IDS 2017, UNSW NB15	CNN-LSTM	98.55	97.43	98.38	–	Difficulty detecting unseen attacks
Ling et al. (2025) [20]	UNSW-NB15	ISFT + Two-stage classifier	96.99	96.93	97.11	96.99	Low-sample attack limitation
Our Proposed: Multi-Stage Generative–Probabilistic Framework	CIC IoT 2023, CIC DIAD IoT 2024	(PCA + iForest + IDSGAN + GMM)	98.47–99.15	0.992–0.996	99.94–99.97	98.48–99.16	Balanced high performance; interpretable; scalable distributed design

This comparative study highlights that the proposed framework fills several important research gaps that existing IDS solutions do not fully address, including adversarial resilience, distributed resource allocation, and probabilistic attack interpretation. Despite the high performance of certain single-dataset models in the literature, they often fail to generalize across heterogeneous or evolving IoT environments. In contrast, the proposed approach consistently maintains high accuracy and balanced precision-recall metrics across multiple datasets, underscoring its robustness. These insights demonstrate that integrating PCA, iForest, IDSGAN, and GMM in a distributed architecture results in a more comprehensive and practically deployable IDS for smart agriculture.

5. Discussion

5.1. Key Findings and Implications

The proposed framework consistently demonstrates high detection performance across both benchmark datasets, achieving excellent Accuracy, Precision, Recall, F1-score, Specificity, and AUC-ROC values. Binary confusion matrices (Figures 6a–b) show minimal misclassification between benign and attack flows, confirming the pipeline's ability to reliably separate normal and anomalous traffic. Comparative evaluation with baseline models (Random Forest, LSTM-based IDS, Autoencoder-only IDS) highlights the superiority of the proposed multi-stage approach, with notable improvements across all metrics. ROC curves further reinforce its strong discriminative capability, reflecting high true positive rates and low false positive rates.

The GMM-based clustering stage effectively organizes detected anomalies into seven representative attack types: BruteForce, DDoS, DoS, Mirai, Recon, Spoofing, and Web-Based attacks. These types represent diverse attack behaviors, including volumetric, integrity-based, reconnaissance, and command-level intrusions. Clustering performance metrics—including Precision, Recall, Specificity, F1-score, Accuracy, and False Negative Rate (FNR)—demonstrate high detection fidelity, with low FNR values highlighting minimal missed detections across all attack categories (Table 7).

Clustering quality metrics—Silhouette Score (99.01%), Adjusted Rand Index (ARI 98.97%), Normalized Mutual Information (NMI 98.91%), and Davies–Bouldin Index (DBI 0.0091)—indicate that anomalies are grouped into well-separated, internally cohesive clusters. The multiclass confusion matrix (Figure 9) visually confirms accurate cluster assignments, with high diagonal values and minimal off-diagonal misclassifications. These results enhance interpretability and enable actionable threat intelligence mapping, supporting frameworks such as MITRE ATT&CK.

The distributed deployment evaluation further demonstrates the framework's scalability and efficiency. IoT devices perform lightweight feature extraction with low latency (4.2 ms) and minimal energy consumption (0.8 J), fog nodes handle PCA and preliminary anomaly scoring with moderate overhead (latency 12.6 ms, energy 2.1 J), and cloud servers manage resource-intensive generative modeling and clustering (latency 45.4 ms, energy 6.5 J) while maintaining high throughput. This layered design ensures real-time detection and optimizes resource utilization across IoT, fog, and cloud layers.

Importantly, the framework exhibits robustness to moderate hyperparameter variations. Preliminary experiments varying the number of GMM components (K) and iForest estimators (n) by ±10–15% showed less than 1% change in accuracy, indicating stable performance under typical parameter adjustments. Nonetheless, as farm conditions and IoT traffic patterns evolve, periodic retraining or online adaptation (e.g., incremental learning, federated updates) may be required to maintain optimal detection performance.

The CIC IoT 2023 and CIC DIAD IoT 2024 datasets encompass a wide range of attack types and heterogeneous IoT traffic patterns, providing a realistic testbed for smart farming networks. While certain conditions—such as sensor failures, packet loss, seasonal variations, or entirely novel attacks—are not fully represented, the framework demonstrates strong robustness. IDSGAN-based adversarial augmentation combined with GMM clustering allows reliable detection of deviations from learned patterns, maintaining high precision (≥99.94%) and recall (≥98.48%). Future work will evaluate online adaptation and live deployment to ensure resilience against evolving attacks and dynamic farm network conditions.

Overall, the multi-stage generative–probabilistic pipeline effectively handles non-IID traffic distributions from heterogeneous farms, ensuring reliable anomaly detection, high interpretability, and practical applicability in distributed smart farming IoT environments.

6. Conclusion

This study introduces a distributed, multi-stage generative–probabilistic intrusion detection framework tailored for IoT-based smart farming environments. The design integrates PCA and Isolation Forest at the fog layer for efficient anomaly isolation, while leveraging IDSGAN and a Gaussian Mixture Model (GMM) at the cloud layer for advanced generative modeling and probabilistic clustering. Experiments conducted on the CIC IoT 2023 and CIC DIAD IoT 2024 datasets show consistently high performance, achieving Accuracy between 98.47–99.15%, F1-scores of 0.992–0.996, and AUC values of 0.981–0.987. The framework outperforms baseline approaches—including Random Forest, LSTM-based IDS, and Autoencoder-only IDS—demonstrating robust detection capability across diverse attack patterns. The distributed architecture minimizes computational burden on IoT endpoints (latency 4.2 ms, energy 0.8 J) while allocating complex generative and clustering tasks to cloud servers (latency 45.4 ms, throughput 3500 flows/s), confirming scalability and real-time operational feasibility. The GMM-based clustering mechanism provides high-quality attack grouping (Silhouette Score 99%, ARI 98.97%) with low false-negative rates (<3%), enabling detailed characterization of attacks such as BruteForce, DDoS, Mirai, Reconnaissance, and other attack vectors. Importantly, the framework effectively handles non-IID traffic distributions from heterogeneous farms, ensuring robust anomaly detection across diverse IoT environments. These outcomes highlight the system's interpretability and its value for proactive cyber threat intelligence. While the framework delivers strong results, limitations remain regarding system complexity, offline training dependencies, and restricted evaluation of low-frequency attack types. Future work will focus on adaptive online learning, hierarchical or multi-label attack clustering, and large-scale real-world deployment in operational smart farming settings. Overall, the proposed framework provides a scalable, high-performing, and resource-efficient intrusion detection solution, contributing to secure and resilient IoT ecosystems in distributed smart agriculture environments.

6.1. Limitations and Future Directions

Despite the strong performance of the proposed framework, several limitations should be acknowledged, alongside opportunities for future enhancements.

Attack Coverage and Generalizability: The GMM-based clustering stage primarily targets seven dominant attack categories observed in the selected datasets. Consequently, rare, emerging, or low-frequency attack patterns are underrepresented, which may limit generalizability to diverse real-world threat landscapes. Future work could extend the model to support multi-label, fine-grained, or hierarchical attack taxonomies, improving adaptability across a broader spectrum of cyber threats.

Resource-Constrained Deployment and Hardware Considerations: Although the distributed IoT–fog–cloud design demonstrates efficiency, certain ultra-resource-constrained IoT devices may still struggle with even minimal feature extraction under strict energy and memory limits. The framework has been optimized for practical deployment in developing agricultural regions: lightweight PCA operations can run on low-cost microcontrollers or single-board computers (e.g., Raspberry Pi, Arduino, ESP32, $20–$50 per unit) with low energy consumption (<1 J) and minimal latency (~4 ms). Mid-tier fog nodes (e.g., Intel NUC, NVIDIA Jetson Nano, $200–$500) handle preliminary anomaly scoring and PCA aggregation with moderate latency (~12–15 ms), while resource-intensive generative modeling, GMM clustering, and threat intelligence mapping are selectively offloaded to cloud servers, minimizing bandwidth and cost. Additional improvements could include model compression, lightweight neural architectures, or edge pruning strategies to further support highly constrained nodes.

Adaptive and Privacy-Preserving Learning: The current framework relies on offline-trained models, limiting responsiveness to rapidly evolving threats and zero-day attacks. Integrating adaptive online learning, continual retraining, or reinforcement-driven self-updating mechanisms could enhance real-time adaptability. In multi-farm collaborative scenarios, privacy-preserving approaches—such as federated learning or secure multi-party computation—can allow local PCA and iForest models to be trained on-site, with only aggregated updates shared centrally. IDSGAN can incorporate differential privacy or encrypted gradient updates, and GMM clustering can use secure aggregation, enabling collaborative IDS deployment while preserving sensitive IoT traffic and maintaining high detection performance and interpretability.

Evaluation Scope and Benchmarking: The current evaluation relies on widely adopted baseline models (Random Forest, LSTM-based IDS, Autoencoder-only IDS) and benchmark datasets, which may not fully capture operational variability in heterogeneous IoT environments. Future studies should explore real-world deployments under noisy, cross-domain, or multi-tenant conditions and extend comparative evaluation to advanced graph-based or transformer-based IDS techniques to comprehensively assess detection performance, scalability, and resource efficiency.

Overall, addressing these limitations—through improved attack coverage, cost-effective edge deployment, adaptive and privacy-preserving learning, and broader evaluation—will strengthen the robustness, adaptability, and practical applicability of the proposed intrusion detection framework for next-generation smart farming IoT networks.

7. Declarations

7.1. Author Contributions

Conceptualization, M.T. and S.Y.; methodology, M.T.; software, M.T.; validation, M.T., S.Y., S.F.A.R., M.S.S., and R.S.; formal analysis, M.T. and R.S.; investigation, M.T.; resources, S.Y. and M.S.S.; data curation, M.T.; writing—original draft preparation, M.T.; writing—review and editing, S.Y., S.F.A.R., M.S.S., and R.S.; visualization, M.T.; supervision, S.Y., S.F.A.R., and M.S.S.; project administration, S.Y.; funding acquisition, S.Y. All authors have read and agreed to the published version of the manuscript.

7.2. Data Availability Statement

The data presented in this study are available on request from the corresponding author.

7.3. Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

7.4. Institutional Review Board Statement

Not applicable.

7.5. Informed Consent Statement

Not applicable.

7.6. Declaration of Competing Interest

The authors declare that there are no conflicts of interest concerning the publication of this manuscript. Furthermore, all ethical considerations, including plagiarism, informed consent, misconduct, data fabrication and/or falsification, double publication and/or submission, and redundancies have been completely observed by the authors.

8. References

[1] Alsaleh, S. S., El Bachir Menai, M., & Al-Ahmadi, S. (2024). Federated Learning-Based Model to Lightweight IDSs for Heterogeneous IoT Networks: State-of-the-Art, Challenges, and Future Directions. IEEE Access, 12, 134256–134272. doi:10.1109/ACCESS.2024.3460468.

[2] Zolghadri, M., Asghari, P., Dashti, S. E., & Hedayati, A. (2025). Dynamic task offloading for IoT-Fog-Cloud systems: a network traffic-aware decision tree approach. Computing, 107(4), 94. doi:10.1007/s00607-025-01449-4.

[3] Mohy-eddine, M., Guezzaz, A., Benkirane, S., & Azrour, M. (2024). Malicious detection model with artificial neural network in IoT-based smart farming security. Cluster Computing, 27(6), 7307–7322. doi:10.1007/s10586-024-04334-5.

[4] Pal, B., Islam, M. S., & Liew, A. W. C. (2025). Calibrated Uncertainty Estimation for Trustworthy Deep IoT Attack Detection. IEEE Transactions on Dependable and Secure Computing, 22(6), 7571–7584. doi:10.1109/TDSC.2025.3598350.

[5] GabAllah, N., Farrag, I., Khalil, R., Sharara, H., & ElBatt, T. (2023). IoT systems with multi-tier, distributed intelligence: From architecture to prototype. Pervasive and Mobile Computing, 93. doi:10.1016/j.pmcj.2023.101818.

[6] Sagar, A. S. M. S., Islam, M. Z., Haider, A., & Kim, H. S. (2024). Uncertainty-Aware Federated Reinforcement Learning for Optimizing Accuracy and Energy in Heterogeneous Industrial IoT. Applied Sciences (Switzerland), 14(18), 8299. doi:10.3390/app14188299.

[7] Ferreira, R., Bispo, I., Rabadão, C., Santos, L., & Costa, R. L. de C. (2025). Farm-flow dataset: Intrusion detection in smart agriculture based on network flows. Computers and Electrical Engineering, 121, 109892. doi:10.1016/j.compeleceng.2024.109892.

[8] Talpini, J., Sartori, F., & Savi, M. (2024). Enhancing trustworthiness in ML-based network intrusion detection with uncertainty quantification. Journal of Reliable Intelligent Environments, 10(4), 501–520. doi:10.1007/s40860-024-00238-8.

[9] Chakravarthy, V., Bell, D., & Bhaskaran, S. (2025). Emergent Intrusion Detection System for Fog Enabled Smart Agriculture Using Federated Learning and Blockchain Technology: A Review. 2nd International Conference on IT Innovations and Knowledge Discovery, ITIKD 2024, 1–7. doi:10.1109/ITIKD63574.2025.11005327.

[10] Srichandan, S. K., Majhi, S. K., Jena, S., Mishra, K., & Bhat, R. (2024). A Secure and Distributed Placement for Quality of Service-Aware IoT Requests in Fog-Cloud of Things: A Novel Joint Algorithmic Approach. IEEE Access, 12, 56730–56748. doi:10.1109/ACCESS.2024.3390723.

[11] Khan, J., Elfakharany, R., Saleem, H., Pathan, M., Shahzad, E., Dhou, S., & Aloul, F. (2025). Can Machine Learning Enhance Intrusion Detection to Safeguard Smart City Networks from Multi-Step Cyberattacks? Smart Cities, 8(1), 13. doi:10.3390/smartcities8010013.

[12] Logeswari, G., Deepika Roselind, J., Tamilarasi, K., & Nivethitha, V. (2025). A Comprehensive Approach to Intrusion Detection in IoT Environments Using Hybrid Feature Selection and Multi-Stage Classification Techniques. IEEE Access, 13, 24970–24987. doi:10.1109/ACCESS.2025.3532895.

[13] Benameur, R., & Dahane, A. (2025). Sfedrl-ids: secure federated deep reinforcement learning-based intrusion detection system for agricultural internet of things. Cluster Computing, 28(6), 403. doi:10.1007/s10586-024-05091-1.

[14] Yang, Y. M., Chang, K. C., & Luo, J. N. (2025). Hybrid Neural Network-Based Intrusion Detection System: Leveraging LightGBM and MobileNetV2 for IoT Security. Symmetry, 17(3), 314. doi:10.3390/sym17030314.

[15] Rahman, S., Pal, S., Mittal, S., Chawla, T., & Karmakar, C. (2024). SYN-GAN: A robust intrusion detection system using GAN-based synthetic data for IoT security. Internet of Things (Netherlands), 26. doi:10.1016/j.iot.2024.101212.

[16] Chatterjee, S., Shaw, V., & Das, R. (2024). Multi-stage intrusion detection system aided by grey wolf optimization algorithm. Cluster Computing, 27(3), 3819–3836. doi:10.1007/s10586-023-04179-4.

[17] Ileri, K. (2025). A hybrid feature selection method for anomaly detection using shallow and deep ANN classifiers in smart farming. Journal of Ambient Intelligence and Smart Environments, 18761364251359885. doi:10.1177/18761364251359885.

[18] Zidi, K., Ben Abdellafou, K., Aljuhani, A., Taouali, O., & Harkat, M. F. (2024). Novel intrusion detection system based on a downsized kernel method for cybersecurity in smart agriculture. Engineering Applications of Artificial Intelligence, 133. doi:10.1016/j.engappai.2024.108579.

[19] Aburasain, R. Y. (2024). Enhanced Black Widow Optimization With Hybrid Deep Learning Enabled Intrusion Detection in Internet of Things-Based Smart Farming. IEEE Access, 12, 16621–16631. doi:10.1109/ACCESS.2024.3359043.

[20] Ling, J., Zhang, L., Liu, C., Xia, G., & Zhang, Z. (2025). Machine Learning-Based Multilevel Intrusion Detection Approach. Electronics (Switzerland), 14(2), 323. doi:10.3390/electronics14020323.

[21] Polat, O., Türkoğlu, M., Polat, H., Oyucu, S., Üzen, H., Yardımcı, F., & Aksöz, A. (2024). Multi-Stage Learning Framework Using Convolutional Neural Network and Decision Tree-Based Classification for Detection of DDoS Pandemic Attacks in SDN-Based SCADA Systems. Sensors, 24(3), 1040. doi:10.3390/s24031040.

[22] Alsagri, H. S. (2025). Hybrid Machine Learning-Based Multi-Stage Framework for Detection of Credit Card Anomalies and Fraud. IEEE Access, 13, 77039–77048. doi:10.1109/ACCESS.2025.3565612.

[23] Praharaj, L., Gupta, D., & Gupta, M. (2025). Efficient federated transfer learning-based network anomaly detection for cooperative smart farming infrastructure. Smart Agricultural Technology, 10. doi:10.1016/j.atech.2024.100727.

[24] Althunayyan, M., Javed, A., & Rana, O. (2024). A robust multi-stage intrusion detection system for in-vehicle network security using hierarchical federated learning. Vehicular Communications, 49. doi:10.1016/j.vehcom.2024.100837.

[25] Poongodi, M., & Hamdi, M. (2023). Intrusion detection system using distributed multilevel discriminator in GAN for IoT system. Transactions on Emerging Telecommunications Technologies, 34(11), e4815. doi:10.1002/ett.4815.

[26] Yang, Y., Liu, X., Wang, D., Sui, Q., Yang, C., Li, H., Li, Y., & Luan, T. (2025). A CE-GAN based approach to address data imbalance in network intrusion detection systems. Scientific Reports, 15(1), 7916. doi:10.1038/s41598-025-90815-5.

[27] Siddique, M. S., Khan, M. A. R., Ahammad, I., Nath, N., Das, J. R., & Rahman, F. (2025). An intelligent intrusion detection system for cyber-physical systems using GAN-LSTM networks. Franklin Open, 11. doi:10.1016/j.fraope.2025.100281.

[28] Aldossary, M., Alharbi, H. A., & Anwar Ul Hassan, C. (2024). Internet of Things (IoT)-Enabled Machine Learning Models for Efficient Monitoring of Smart Agriculture. IEEE Access, 12, 75718–75734. doi:10.1109/ACCESS.2024.3404651.

[29] Mohamed, D., & Ismael, O. (2023). Enhancement of an IoT hybrid intrusion detection system based on fog-to-cloud computing. Journal of Cloud Computing, 12(1), 41. doi:10.1186/s13677-023-00420-y.

[30] Tariq, N., Alsirhani, A., Humayun, M., Alserhani, F., & Shaheen, M. (2024). A fog-edge-enabled intrusion detection system for smart grids. Journal of Cloud Computing, 13(1), 43. doi:10.1186/s13677-024-00609-9.

[31] Tawfik, M. (2024). Optimized intrusion detection in IoT and fog computing using ensemble learning and advanced feature selection. PLoS ONE, 19(8 August), 304082. doi:10.1371/journal.pone.0304082.

[32] Aldhaheri, S., & Alhuzali, A. (2023). SGAN-IDS: Self-Attention-Based Generative Adversarial Network against Intrusion Detection Systems. Sensors, 23(18), 7796. doi:10.3390/s23187796.

[33] Sajid, M., Malik, K. R., Almogren, A., Malik, T. S., Khan, A. H., Tanveer, J., & Rehman, A. U. (2024). Enhancing intrusion detection: a hybrid machine and deep learning approach. Journal of Cloud Computing, 13(1), 685. doi:10.1186/s13677-024-00685-x.

[34] Alotaibi, S. D., Yadav, K., Aledaily, A. N., Alkwai, L. M., Yousef Dafhalla, A. K., Almansour, S., & Lingamuthu, V. (2022). Deep Neural Network-Based Intrusion Detection System through PCA. Mathematical Problems in Engineering, 2022(1), 6488571. doi:10.1155/2022/6488571.

[35] Elsaid, S. A., & Binbusayyis, A. (2024). An optimized isolation forest based intrusion detection system for heterogeneous and streaming data in the industrial Internet of Things (IIoT) networks. Discover Applied Sciences, 6(9), 6165. doi:10.1007/s42452-024-06165-w.

[36] De Araujo-Filho, P. F., Naili, M., Kaddoum, G., Fapi, E. T., & Zhu, Z. (2023). Unsupervised GAN-Based Intrusion Detection System Using Temporal Convolutional Networks and Self-Attention. IEEE Transactions on Network and Service Management, 20(4), 4951–4963. doi:10.1109/TNSM.2023.3260039.

[37] Liu, L., & Xu, M. (2025). A network intrusion detection method based on contrastive learning and Bayesian Gaussian Mixture Model. Cybersecurity, 8(1), 3647. doi:10.1186/s42400-025-00364-7.

[38] Wang, Q., Wang, W., Wang, Y., Ren, J., & Zhang, B. (2025). Multi-Stage Network Attack Detection Algorithm Based on Gaussian Mixture Hidden Markov Model and Transfer Learning. IEEE Transactions on Automation Science and Engineering, 22, 3470–3484. doi:10.1109/TASE.2024.3395355.

[39] Neto, E. C. P., Dadkhah, S., Ferreira, R., Zohourian, A., Lu, R., & Ghorbani, A. A. (2023). CICIoT2023: A Real-Time Dataset and Benchmark for Large-Scale Attacks in IoT Environment. Sensors, 23(13), 5941. doi:10.3390/s23135941.

[40] Rabbani, M., Gui, J., Nejati, F., Zhou, Z., Kaniyamattam, A., Mirani, M., Piya, G., Opushnyev, I., Lu, R., & Ghorbani, A. A. (2025). Device Identification and Anomaly Detection in IoT Environments. IEEE Internet of Things Journal, 12(10), 13625–13643. doi:10.1109/JIOT.2024.3522863.