OXFORD CDE

Journal of Computational Design and Engineering, 2025, **12**, 46–60

DOI: 10.1093/jcde/qwaf085
Advance access publication date: 12 August 2025
Research Article

A robust anomaly detector for imbalanced industrial internet of things data

Rubina Riaz [1], Guangjie Han [2,*], Kamran Shaukat [3,*], Naimat Ullah Khan [4] and Hongbo Zhu [5]

[1]School of Software Engineering, Dalian University of Technology, Liaoning, Dalian 116024, China
[2]School of Internet of Things Engineering, Hohai University, Changzhou 210098, China
[3]Centre for Artificial Intelligence Research and Optimisation, Design and Creative Technology Vertical, Torrens University Australia, Ultimo, NSW 2007, Australia
[4]School of Computer Science, University of Technology Sydney, Sydney 2007, Australia
[5]School of Information Science and Engineering, Shenyang Ligong University, Shenyang 110159, China
*Correspondences: hanguangjie@gmail.com (GH); Kamran.shaukat@torrens.edu.au (KS)

Abstract

Machine learning (ML) and deep learning (DL) have been used for anomaly detection in industrial internet of things (IIoT) environments. The presence of imbalanced data, high noise levels, missing values, and high dimensionality poses an enormous challenge for existing methods, leading to inconsistent reliability in detecting anomalies in real-world industrial environments. Current anomaly detection solutions suffer from high false negative rates due to class imbalance and noisy sensor data, limiting their practical applicability. This paper proposes the Ensemble Wasserstein generative adversarial network for IIoT (EWAD-IIoT) framework, which is uniquely designed to address these challenges. The aim is to build a robust anomaly detection model with high recall (94.7%) and precision (93.6%) while minimizing miss rates in complex IIoT settings. Evaluations on two benchmark data sets, SECOM (industrial sensor data) and MNIST (image data), demonstrate EWAD-IIoT's superiority over traditional methods like standalone WGAN and WGAN-GP. To highlight its efficacy, we compare results against these benchmarks, showcasing improvements in F1-score (95.8%) and noise robustness. The framework leverages advanced pre-processing (Z-score filtering and min–max scaling), SMOTE-based balancing, and WGAN-generated synthetic samples to handle data imbalance and dimensionality. The results validate EWAD-IIoT's capability to detect rare anomalies in IIoT environments, with a balanced trade-off between recall and precision, making it a scalable solution for predictive maintenance and fault diagnosis.

Keywords: anomaly detection, industrial internet of things, Wasserstein generative adversarial network, data imbalance, high-dimensional data, noisy data

1. Introduction

The Fourth Industrial Revolution, also known as Industry 4.0, has led to the widespread adoption of industrial internet of things (IIoT) systems in various sectors, including manufacturing, logistics, energy, and smart infrastructure (Keshar, 2025; Routaib *et al.*, 2025). One of the significant challenges in IIoT applications is the detection of anomalies and deviations from expected system behaviour. Anomaly detection and deviation identification from standard system operation create one of the main obstacles when implementing IIoT applications. Systematic or cyber-based disturbances, process defects, sensor issues, and environmental stressors produce significant operational and financial setbacks. IIoT systems depend on precise anomaly detection systems to operate smoothly while providing security.

Multiple operational limitations affect existing anomaly detection systems deployed in IIoT applications, as described by Ghosh *et al.* (2024). The rare occurrence of system failures in IIoT settings leads to catastrophic results. Research shows that K-nearest neighbours (KNN) and support vector machines (SVM) fail to achieve accurate anomaly detection when implemented on imbalanced data sets (Khalid *et al.*, 2024). Anomaly detection is one of the most challenging security tasks because of IIoT's distinct operating characteristics. High dimensionality represents the principal challenge when extensive sensor networks combine with device systems to produce comprehensive data ranges across different features (Alkhafaji & Viana, 2024). Traditional machine learning (ML)-based approaches struggle to work with large-scale data sets since they lack suitable tools for pattern interpretation. The 'noisy and incomplete data problems' that the IIoT environment faces derive from sensor malfunctions, communication disruptions, and environmental elements that damage sensor readings or result in sensor data corruption. Several issues affect anomaly detection systems because unanticipated changes occur frequently in monitored systems. IIoT systems develop recurring temporal dependencies between their data records due to their permanent monitoring of evolving state patterns with time (Kumar *et al.*, 2024). Time-free data analysis produces erroneous anomaly detection outcomes since most anomalous patterns develop from modifications in past patterns. Generative adversarial networks (GANs) show promising applications for research anomaly detection, as described in Liu *et al.* (2023). According to research findings, WGANs successfully detect anomalies in complex data sets yet struggle with overfitting and data set imbalance problems (Qi *et al.*, 2023). Wan *et al.* (2024) demonstrated how GAN-based methods successfully deal with imbalanced data and create dependable anomaly detection points. Ba-Alawi *et al.*

Received: March 15, 2025. **Revised:** July 10, 2025. **Accepted:** July 13, 2025

(2022) evaluated these models against industrial sensor data noise and missing value scenarios. In real-world IIoT deployments, anomaly detection must contend with significant data challenges. Data streams are typically highly imbalanced (many normal observations versus very few anomalies), complicating the learning of rare fault signatures (Saranya & Valarmathi, 2025). Sensor measurements may also be heavily corrupted by noise or suffer frequent missing values due to hardware faults or network dropouts, obscuring true anomalies (Huang *et al.*, 2025). Furthermore, data are often extremely high-dimensional, with numerous correlated sensor channels and time points. This exacerbates the curse of dimensionality and can cause models to overfit or generate false alarms (Arafah *et al.*, 2025).

Modern IIoT anomaly detectors now emphasize advanced deep learning (DL) architectures instead of generative ensembles. Transformer-based models, for example, use self-attention to capture long-range temporal and cross-sensor dependencies. Zia *et al.* (2025) develop a transformer-based framework that learns complex temporal patterns in multivariate IoT streams and even employs adversarial perturbations to boost robustness. Autoencoder variants remain popular for unsupervised detection: one-class or variational autoencoders compress high-dimensional sensor data into latent spaces, flagging outliers by reconstruction error. Ayad *et al.* (2024) report that a hybrid asymmetric stacked autoencoder with a deep neural networks (DNN) achieves very high 96% detection rates on Botnet of Things-IoT (BoT–IoT) with minimal false positives, effectively tackling data imbalance and high dimensionality. Hybrid networks further combine complementary layers to improve feature extraction. For instance, Shang *et al.* (2024) introduce CAE-T, which merges a convolutional autoencoder (for spatial feature learning) with a transformer (for long-term temporal context) in an unsupervised anomaly detector.

In this work, the EWAD-IIoT (Ensemble Wasserstein generative adversarial network for IIoT) framework incorporates a dedicated and integrated pre-processing step as part of its training phase, rather than treating it as an external or independent step. This pipeline, as illustrated in the framework 1, is specifically designed to improve data quality and address class imbalance before training the ensemble of WGANs. Each stage, from data cleaning to balanced sample generation, contributes directly to the robustness and stability of the learning process. It is important to note that this pre-processing pipeline is applied only during the training phase of EWAD-IIoT and not to the test data or baseline models. The trained EWAD-IIoT model directly evaluates the test data and compares against baseline methods (GAN, WGAN, and WGAN-GP) without any additional processing, ensuring a clear and fair comparison of model performance. Table 1 displays the complete forms and descriptions of acronyms used in this paper to support a better understanding of technical terminology for the reader. The principal contributions of this study are as follows.

(1) A novel EWAD-IIoT framework is proposed, combining multiple WGANs within an ensemble architecture to capture complex data distributions in IIoT anomaly detection effectively. The ensemble technique significantly improves the robustness and accuracy of anomaly detection by effectively addressing challenges such as data imbalance, high dimensionality, and noisy sensor data. The proposed framework significantly reduces the limitations associated with traditional GAN-based approaches, including mode collapse and instability during training.
(2) The EWAD-IIoT framework integrates advanced data pre-processing strategies specifically tailored for IIoT data sets. These pre-processing methods, including SMOTE-based oversampling for balancing minority classes, statistical and KNN-based imputation for handling missing values, and Z-score-based outlier detection combined with min–max normalization, collectively ensure the model's optimal performance by improving data quality and reliability.
(3) The performance of the proposed EWAD-IIoT framework was rigorously evaluated on standard benchmarks such as the SECOM and MNIST data sets. The results demonstrate superior detection capabilities in precision, recall, F1-score, accuracy, and ROC, thus validating our approach's practical effectiveness and robustness for diverse and complex IIoT scenarios.

Table 1: List of Acronyms.

Acronym	Description
IIoT	Industrial internet of things
EWAD-IIoT	Ensemble Wasserstein generative adversarial network for IIoT
WGAN	Wasserstein generative adversarial network
WGAN-GP	Wasserstein generative adversarial network with gradient penalty
GAN	Generative adversarial network
SMOTE	Synthetic minority over-sampling technique
SECOM	Semiconductor manufacturing data set
MNIST	Modified National Institute of Standards and Technology (Image data set)
ROC-curve	Receiver operating characteristic curve
F1-score	Harmonic mean of precision and recall
G-mean	Geometric mean
KNN	K-nearest neighbours
SVM	Support vector machine
CNN	Convolutional neural network
RNN	Recurrent neural network
GPU	Graphics processing unit

The paper is organized as follows: Section 2 reviews related work on IIoT anomaly detection. Section 3 outlines the proposed EWAD-IIoT framework, including pre-processing, ensemble WGAN architecture, and evaluation. Section 4 presents experimental results and comparisons with benchmarks, and Section 5 presents conclusions and future directions.

2. Related Work

Anomaly detection in IIoT systems has been a critical area of research due to the significant impact anomalies can have on industrial operations, safety, and efficiency. Anomaly detection relies on traditional and modern techniques, which are grouped into statistical methods, ML approaches, and DL methods with distinct strengths and drawbacks.

2.1. Conventional approaches

IIoT early anomaly detection techniques mostly used statistical methods in their initial implementations, including threshold-setting methods based on previously defined rules and clustering techniques used to identify outliers. These methods maintain low computational costs and interpretability, yet they face challenges when analysing high-dimensional and noisy IIoT data sets. Real-world industrial scenarios with dynamic and heterogeneous sensor data pose challenges to statistical methods because these methods cannot correctly capture complex data

distributions (Liu *et al.*, 2023). ML technology introduced improvements to IIoT system anomaly detection methods. Multiple algorithms, including SVMs (Zeng *et al.*, 2025), decision trees (Papastefanopoulos *et al.*, 2025), random forests, and KNN, are commonly used for their broad applicability. Results from these methods prove superior to statistical techniques in analysing data sets with moderate distribution complexity. IIoT anomaly detection techniques experience a significant performance decline due to highly imbalanced data sets since these detection tasks often deal with rare anomalies. ML models prefer the majority class when detecting anomalies because they lack pre-processing through undersampling, oversampling, or synthetic data generation (Lazaar, 2024).

2.2. GAN-based approaches for anomaly detection

GAN-based methods have shown remarkable promise in anomaly detection due to their generative capabilities. Arjovsky *et al.* (2017) introduced GANs consisting of opposing neural networks, which include a generator and a discriminator. The generator works to produce synthetic data that replicates real data distributions, but the discriminator performs a task to distinguish between real and synthetic data. GANs extract complex data patterns through their adversarial training mechanism to produce high-quality synthetic data for anomaly detection tasks. GANs demonstrate numerous applications in anomaly detection systems. Anomaly detection relies heavily on synthetic data samples generated by the generator by assessing reconstruction errors (Geiger *et al.*, 2020). The generator shows reconstruction failures when it produces synthetic data that significantly diverges from the learned distribution (Mestav *et al.*, 2022). However, standard GANs are often limited by mode collapse, where the generator fails to capture the diversity of the accurate data distribution, and training instability due to the adversarial nature of the network (Ahmad *et al.*, 2024). GANs perform poorly when processing real-world IIoT data sets because their limitations fail to properly manage data characteristics such as multidimensionality, noise, and imbalanced distributions (Benaddi *et al.*, 2022).

The standard GAN framework has led researchers to develop multiple variations to improve the system's stability and robustness. WGANs have gained significant attention by employing the Wasserstein distance as the loss function; WGANs mitigate training instability and provide more meaningful gradients for optimization, making them particularly effective in anomaly detection tasks (Arjovsky *et al.*, 2017; Gulrajani *et al.*, 2017). WGANs demonstrate excellent capabilities for modelling complex data distributions and synthetic data generation while detecting subtle anomalies simultaneously (Gondhi, 2024). The proposed framework is tailored to tackle the challenges specific to IIoT environments, such as data imbalance, high dimensionality, and noisy or incomplete data, marking a significant step forward in industrial anomaly detection.

2.3. Limitations of traditional GANs and need for advanced GAN variants

Further advancements, such as WGAN with gradient penalty (WGAN-GP), improve the WGAN by addressing gradient explosion and mode collapse by introducing gradient penalties (Lee *et al.*, 2023). The generator achieves better data distribution coverage through this technique, improving anomaly detection performance (Guertler *et al.*, 2022). f-AnoGAN extends GANs for anomaly detection by integrating generative modelling with encoder–decoder components, enabling exact anomaly grading via latent space measurement and reconstruction errors. Complex IIoT data environments contain multiple data modes that single GAN models find difficult to capture, leading to reduced detection accuracy. The development of ensemble methods is a promising solution to address these limitations. Combining multiple models within ensemble GANs helps single GANs enhance their detection accuracy and produce reliable results. These data approaches suit IIoT environments because they efficiently handle data heterogeneity alongside imbalanced systems (Liu *et al.*, 2024). The research uses a specialized ensemble of WGANs to address imbalanced IIoT data set requirements in its proposed framework. The framework achieves better anomaly detection and resolves multiple data problems through the joint ability of multiple WGANs to generate data. This advancement marks a crucial development for implementing GAN-based methods within industrial operations, which require prioritizing operational safety and reliability.

2.4. Ensemble learning techniques for improved anomaly detection

Ensemble learning methods establish an efficient anomaly detection system that operates effectively in IIoT environments and handles high-dimensional data sets combined with incomplete information and noise. Multiple GANs improve anomaly detection performance compared to standard methods by optimizing outcomes for imbalanced data sets during processing. Research by Al-Fakih *et al.* (2024) demonstrates that multiple GANs achieve better anomaly detection through individual decision boundary discovery processes. Strelcenia (2024) developed WGAN-GP models, which produced high-quality synthetic data set samples for uneven data so detection systems could perform better in following anomaly detection operations. The EWAD-IIoT framework enhances existing developments through its simultaneous solutions for data imbalance problems while enabling scalable and interpretable IIoT applications. WGANs work together in this framework to achieve improved robustness while reducing bias and offering full-scale anomaly detection capabilities for intricate industrial operations. The EWAD-IIoT framework extends the concept through prediction aggregation from multiple WGANs that operate in separate sections of an imbalanced IIoT data set. The specialized data pattern knowledge enables each WGAN to function as part of the ensemble, effectively tackling imbalanced data patterns to enhance anomaly detection results.

3. Proposed Framework

The EWAD-IIoT framework illustrated in Figure 1 presents a modular end-to-end architecture for enhancing anomaly detection in IIoT environments.

Figure 1 comprises several step-by-step operational workflows, which are divided into stages: (1) data collection, splitting, and pre-processing: raw IIoT data are first split into training, validation, and testing subsets. The training data undergoes a multistep pre-processing stage that includes noise removal, missing value imputation using KNN, and min–max normalization. (2) Handling imbalanced data: the pre-processed data enter a two-tiered augmentation strategy to resolve class imbalance. Initially, the SMOTE generates interpolated minority class samples. This is followed by WGAN-GP training, where a generator and critic adversarially learn to produce realistic synthetic samples. This hybrid SMOTE+WGAN approach enhances data diversity while mitigating risks like mode collapse and overfitting. (3) Ensemble

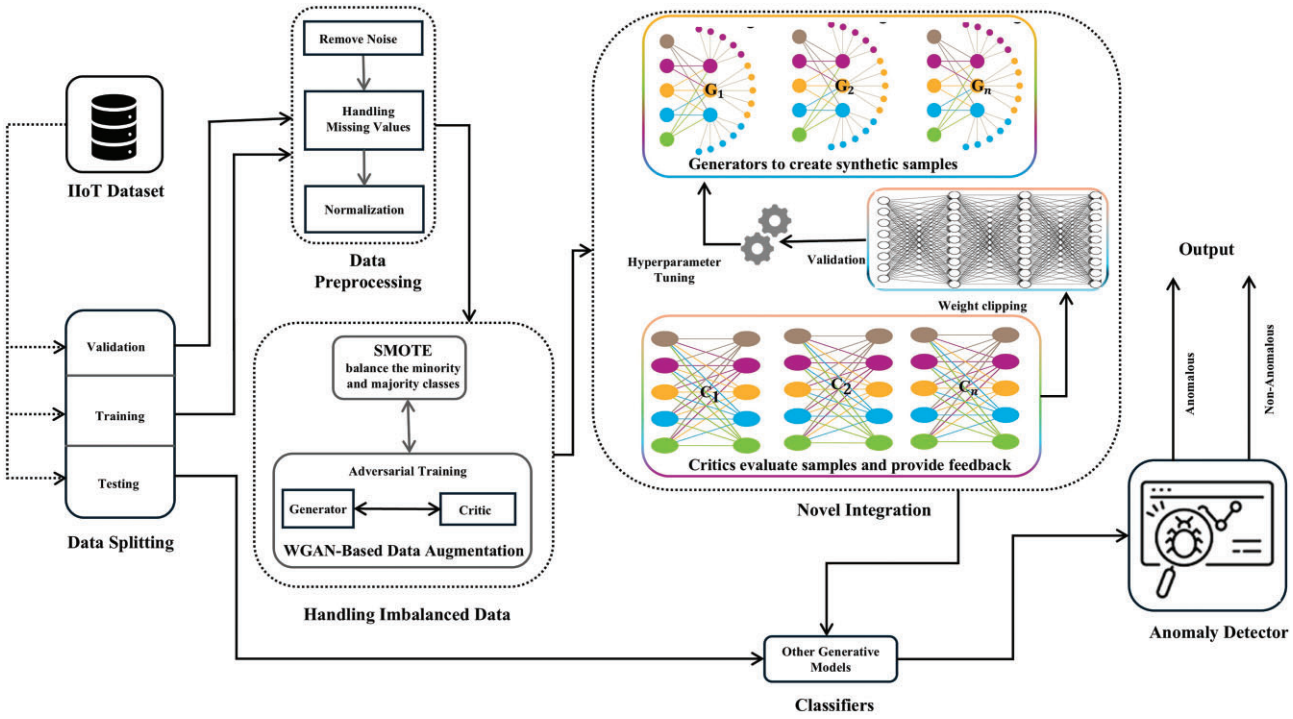

Figure 1: Comprehensive overview of the proposed EWAD-IIoT framework.

WGAN training: to ensure diverse learning and generalization, multiple WGAN units are trained independently. Each unit comprises a generator G_i and critic C_i, which are trained on different partitions of the normal (majority) class data. Hyperparameter tuning, critic feedback, and weight clipping are incorporated to stabilize the adversarial learning process. (4) Integration and anomaly detection: the trained ensemble generates synthetic minority instances that are combined with the real data and passed to the anomaly detector. Anomaly scores are computed by aggregating feedback from multiple critics and the reconstruction error patterns. Threshold-based decision logic categorizes the outcome as either 'anomalous' or 'non-anomalous.' (5) Output: the final output module of the EWAD-IIoT framework classifies input instances as either anomalous or non-anomalous based on the aggregated anomaly scores derived from the ensemble critics and reconstruction metrics. This binary decision supports actionable insight for real-time monitoring, enabling timely responses to potential faults, intrusions, or abnormal behaviours in IIoT systems.

3.1. Implementation and training configuration

The EWAD-IIoT framework utilized Python and Jupyter to implement DL libraries, enabling model development and training. A high-performance NVIDIA A100 GPU coupled with 32 GB of memory enabled efficient training, speeding up the evaluation process and training time. A powerful GPU configuration supported efficient ensemble training of four GANs across 100 epochs, allowing sufficient optimization time for the best results. We tested ensemble sizes $I \in \{3, 4, 5, 6, 7\}$. The F1-score peaked at $I = 5$ and saturated beyond it, while training/inference time increased linearly. With $I = 6$, performance improvement was <0.4% while GPU memory usage increased by 22%. Therefore, $I = 5$ offers the best trade-off, balancing accuracy and efficiency. This empirical basis strengthens the design decision. This ensemble strategy enables the framework to learn a broader range of normal patterns, enhancing the ability to detect rare and subtle anomalies. The batch size was set to 128, a common configuration that balances computational efficiency and memory usage. Additionally, the gradient penalty coefficient ($\lambda = 10$) was carefully tuned to ensure stability during training, particularly with the Wasserstein loss function, which requires enforcing the Lipschitz continuity constraint. The learning rate for the training process was set to 1×10^{-4}, a typical value used for training GANs. This learning rate ensures stable convergence while allowing the model to learn the complex distribution of the data throughout training. These hyperparameters were chosen through extensive experimentation to ensure the best performance of the model.

3.2. Data pre-processing

The framework utilizes the raw IIoT data sets SECOM and MNIST, which present challenges due to high dimensionality, noise, and missing values. Data pre-processing remains essential when implementing ML applications on complex data sets, including industrial environments with imbalanced and noisy sensor data. Data pre-processing aims to create a learning environment that maximizes model accuracy while addressing these challenges. To ensure robustness and generalizability, k-fold cross-validation evaluates the EWAD-IIoT framework. The data set is split into k subsets, with the model trained on $k-1$ folds and validated on the remaining fold. This process is repeated k times, and the average performance metrics provide a reliable estimate of the model's effectiveness, helping to reduce overfitting and ensuring good performance on unseen data in the challenging IIoT environment.

3.3. Clean and normalize the data

Noisy and incomplete data are other challenges inherent to IIoT systems. These are addressed using the following statistical imputation techniques. KNN imputation provides a secondary treatment by analysing feature proximity to determine missing value predictions. Outliers are detected through z-score filtering

methods that help identify data points beyond expected mean values to improve data quality. Continuous features are imputed using the mean or median of observed values, which helps preserve the overall distribution. For continuous and categorical features, missing values are imputed using the average (or mode) of the KNN in the feature space. Noise is managed using outlier detection methods such as z-score filtering described in Equation 1:

$$z = \frac{x - \mu}{\sigma}, \quad |z| > 3 \Rightarrow outlier. \tag{1}$$

Here, μ and σ are the mean and standard deviation of the feature, respectively. Outliers, identified as samples with z-scores exceeding a threshold of 3, are excluded from the data set. To ensure the model treats all features equally, min–max scaling (Shantal *et al.*, 2023) is applied to reduce the impact of selected data dimensions to ensure efficient model training while improving generalization capabilities. Normalizing the data helps prevent the model from favouring features with larger scales. The scaling transformation is defined as in the following Equation 2:

$$x_{scaled} = \frac{X - X_{min}}{X_{max} - X_{min}}. \tag{2}$$

where X_{min} and X_{max} are each feature's minimum and maximum values, respectively. This ensures that all features lie within the range [0, 1], allowing the model to converge more quickly and efficiently during training.

This ensures that no single feature dominates the learning process, especially in high-dimensional data sets. The EWAD-IIoT framework is an ensemble of multiple WGANs, each independently trained on overlapping subsets of the normal class data. This design leverages the strengths of ensemble learning to enhance robustness, scalability, and diversity in anomaly detection.

3.4. Handling imbalanced data

Data imbalance is a major challenge in the IIoT environment for anomaly detection. In this condition, anomalous data representing faults or rare events are much less frequent than normal data. Traditional ML models struggle to process imbalanced data sets (Han *et al.*, 2024), often achieving high accuracy mainly focusing on the dominant class (normal data) while failing to identify the rare cases. Our framework addresses class imbalance by implementing a hybrid data augmentation strategy that integrates SMOTE with WGAN-GP. SMOTE generates new minority-class examples by linear interpolation between existing minority samples: if two sensor readings indicate a rare machine fault, SMOTE generates new plausible fault examples between them. for a minority instance x_i and one of its KNN nearest x_j, a synthetic point is created as in the following Equation 3 governs the process of generating synthetic samples:

$$\dot{x} = x_i + \lambda(x_j - x_i), \quad \lambda \sim \mathcal{U}(0, 1) \tag{3}$$

where x_i and x_j are two minority class samples, and λ is a random weight factor that ensures the generated sample lies between x_i and x_j.

This oversampling technique ensures that synthetic samples remain realistic while preserving the local structure of the minority class and reduces the variance associated with individual models, and provides a balanced output. It increases the minority sample count and fills gaps in the feature space, but its simple interpolation can produce samples that lie too close to decision boundaries or overlap other classes, introducing noise in complex, high-dimensional data. In contrast, WGAN-GP is a generative model that learns the underlying minority-class distribution. Substituting the standard GAN loss with the Wasserstein distance and enforcing a gradient penalty, WGAN-GP stabilizes training and enables the generator to produce realistic, high-quality synthetic samples. In practice, WGAN-GP captures nonlinear feature relationships and generates diverse minority examples beyond the linear interpolations of SMOTE. The EWAD-IIoT framework applies SMOTE first to expand the minority set and then uses WGAN-GP to refine and diversify those samples. This hybrid approach is applied because SMOTE and WGAN-GP balance each other's weaknesses: SMOTE alone can lead to oversampling noise (due to its linear interpolations), whereas a WGAN-GP trained on very few minority samples can suffer from mode collapse and fail to model's distribution accurately. Pre-processing ensures the data are cleaner and balanced before modelling. The ensemble, by aggregating multiple models, improves the stability and overall detection performance under the same processed data. In particular, ensemble training helps reduce model variance and mitigates overfitting in high-dimensional feature spaces. Integrating them, our framework influences SMOTE's handling of imbalance and WGAN-GP's generative power, leading to a more realistic set of minority-class instances and thereby significantly improving minority-class representation in the training data.

3.5. Integration

After pre-processing the data, the next step involves implementing the core of the anomaly detection framework, i.e., EWAD-IIoT, described in Figure 2. Anomaly detection in IIoT environments poses unique challenges, including imbalanced data, high-dimensional data, and complex distributions arising from varying operational conditions and sensor noise. Traditional ML models struggle with these challenges, mainly when anomalies are rare. A single WGAN might capture certain aspects of the data but can fail to generalize well for diverse, high-dimensional IIoT data (Ren *et al.*, 2023). The proposed framework performs better than traditional anomaly detection approaches, including GAN, WGAN, and WGAN-GP, when operating on imbalanced and noisy data sets such as SECOM and MNIST. The results of EWAD-IIoT demonstrate superior precision and recall performance, and the F1-score and ROC data indicate its high ability to detect anomalies accurately. The ensemble captures diverse distributions of normal behaviour through aggregating multiple WGANs, each trained on a subset of the normal data. Each WGAN specializes in a specific region of the data space, allowing the ensemble to generalize better across the entire data set. The ensemble captures diverse distributions of normal behaviour through integrating multiple WGANs, each trained on a subset of the normal data. Each WGAN specializes in a specific region of the data space, allowing the ensemble to generalize better across the entire data set. Let the training data set $X = \{x_1, x_2, \ldots, x_n\}$ be split into N subsets, $X_1, X_2, \ldots, X_n$, where $X_i \subseteq X$ and $\bigcup_{i=1}^{N} X_i = X$. Each WGAN i learns the distribution $P_{data,i}$, approximating the true normal data distribution P_{data}. The ensemble combines these models to capture the overall data distribution as shown in the following Equation 4:

$$P_{ensemble} = \frac{1}{N} \sum_{i=1}^{N} P_{data,i} \tag{4}$$

Training a single WGAN risks overfitting specific patterns in the normal data. Using multiple WGANs, the ensemble avoids becoming overly specialized in any subset of the data, thereby

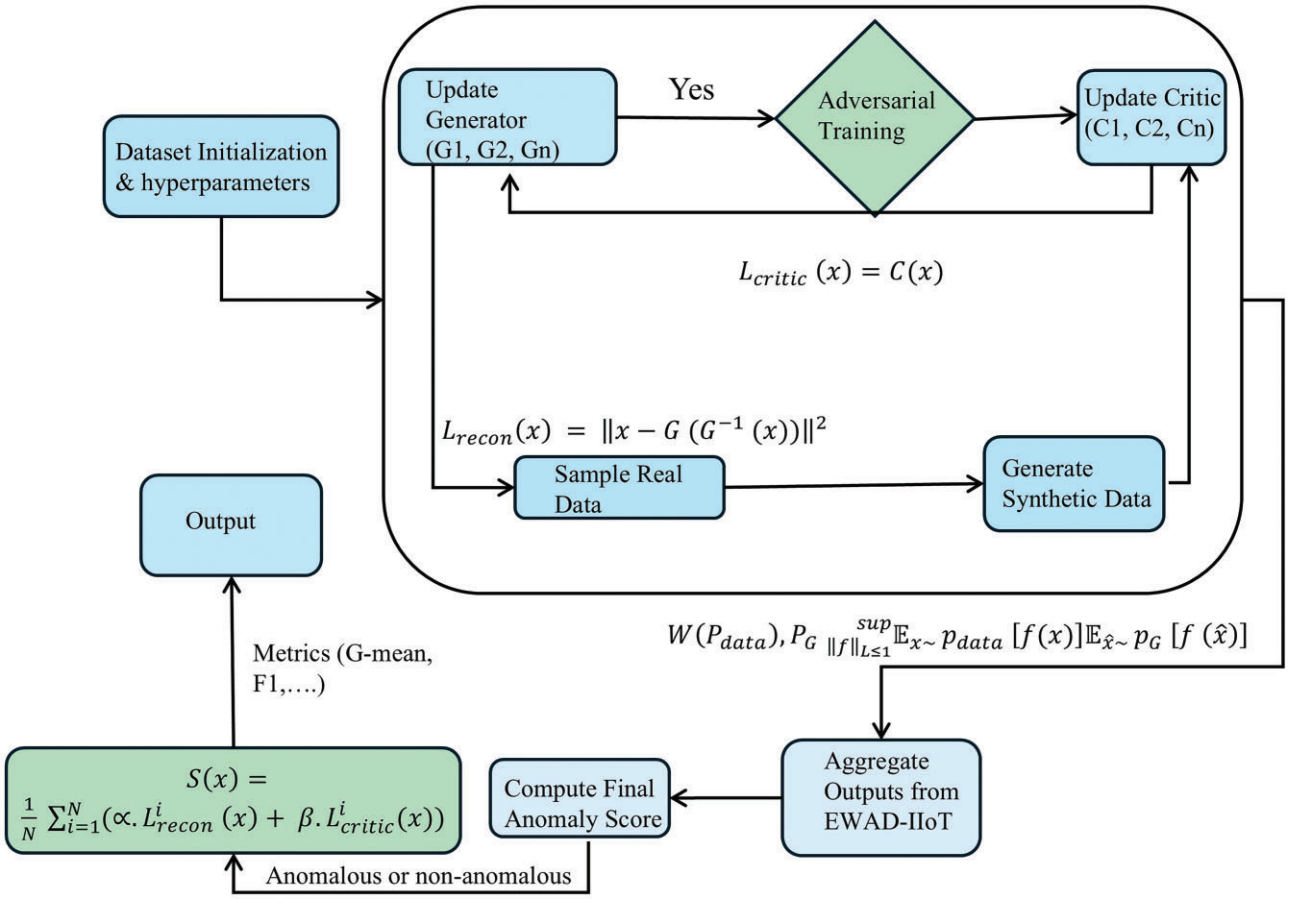

Figure 2: Dataflow and computational steps of the EWAD-IIoT ensemble training and anomaly scoring process.

reducing overfitting. The ensemble model averages the anomaly scores from all WGANs described in Equation 5:

$$S_{ensemble}(x) = \frac{1}{N}\sum_{i=1}^{N} S_i(x) \tag{5}$$

where $S_i(x)$ is the anomaly score computed by the i WGAN. This averaging reduces the variance associated with individual models and ensures a balanced output.

IoT data often contains high-dimensional sensor readings with temporal dependencies. WGANs are suitable for modelling such data because the Wasserstein distance provides a stable training objective, even in high-dimensional spaces. In a WGAN, the critic approximates the Wasserstein distance between the real data distribution P_{data} and the generated data distribution P_G as discussed in the following Equation 6:

$$W(P_{data}, P_G) = \sup_{\|f\|_L \leq 1} \mathbb{E}_{x\sim P_{data}}[f(x)] - \mathbb{E}_{\hat{x}\sim P_G}[f(\hat{x})] \tag{6}$$

where $\|f\|_L \leq 1$ denotes that f is a 1-Lipschitz function.

This distance metric ensures stable gradients, allowing the generator to learn effectively even with complex data distributions. Each WGAN in the ensemble is trained on a subset of the normal data to learn diverse distribution aspects. The generator G learns to map a noise vector $z \sim P_z$ to synthetic data $\hat{x}$ is in Equation 7:

$$G(z; \theta_G) \rightarrow \hat{x} \tag{7}$$

The critic D learns to estimate the Wasserstein distance between P_{data} and P_G fully described in Equation 8:

$$L_{critic} = \mathbb{E}_{x\sim P_{data}}[C(x)] - \mathbb{E}_{\hat{x}\sim P_G}[D(\hat{x})] \tag{8}$$

To enforce the Lipschitz constraint on D, the gradient penalty is applied as showing in Equation 9:

$$L_{GP} = \lambda (\|\nabla_{\hat{x}} C(\hat{x})\|_2 - 1)^2 \tag{9}$$

The final ensemble combines N WGANs shown in Equation 10, where each WGAN contributes to the learned distribution.

$$P_{ensemble} = \frac{1}{N}\sum_{i=1}^{N} P_{G,i} \tag{10}$$

The figure clearly demonstrates the internal dataflow and computational mechanisms of the ensemble WGAN-based anomaly detection model. Initially, the ensemble is trained via adversarial learning, where generators produce synthetic data from latent vectors, and critics evaluate these samples using the Wasserstein distance (Equation 6). Subsequently, anomaly scores are systematically computed by aggregating two distinct measures: (i) reconstruction loss (L_{recon}), capturing the difference between original and reconstructed samples (Equation 11), and (ii) critic confidence scores (L_{critic}), assessing deviations from the normal data distribution (Equation 12). The final anomaly decision integrates these metrics (Equation 13), clearly marking each input sample as anomalous or non-anomalous, supported by comprehensive performance metrics (F1-score, G-mean, and ROC). This explicit

delineation enhances interpretability and highlights critical computational steps.

The following Algorithm 1 outlines the training procedure for the proposed EWAD-IIoT framework, which leverages an ensemble of WGANs to enhance anomaly detection in IIoT environments. The algorithm details the initialization of multiple WGAN models, followed by iterative training that updates both the critic and generator networks. The critic updates aim to minimize the Wasserstein distance with a gradient penalty to ensure stable training, while the generator updates focus on improving data generation quality. This ensemble approach enhances robustness and improves anomaly detection performance in complex, imbalanced IIoT data sets.

Algorithm 1 Training Procedure for EWAD-IIoT Ensemble

Require: Preprocessed dataset X_{train}, number of WGANs N, learning rate η, batch size B, gradient penalty coefficient λ, critic steps k, total iterations T

Ensure: Trained ensemble $\{WGAN_1, WGAN_2, \ldots, WGAN_N\}$

for each WGAN $i \in \{1, \ldots, N\}$ **do**
  Initialize generator G_i and critic C_i with random weights
end for
for iteration $t = 1$ to T **do**
  for each WGAN $i \in \{1, \ldots, N\}$ **do**
    for k critic updates **do**
      Sample real data $X_r \sim P_{data}$, noise $Z \sim P_z$
      Generate fake data $X_f = G_i(Z)$
      Compute critic loss:

$$L_C = \mathbb{E}[C_i(X_r)] - \mathbb{E}[C_i(X_f)] + \lambda(\|\nabla_{\hat{x}} C_i(\hat{x})\|_2 - 1)^2$$

      Update critic: $\theta_{C_i} \leftarrow \theta_{C_i} - \eta \nabla L_C$
    end for
    Sample noise $Z \sim P_z$, generate $X_f = G_i(Z)$
    Compute generator loss: $L_G = -\mathbb{E}[C_i(X_f)]$
    Update generator: $\theta_{G_i} \leftarrow \theta_{G_i} - \eta \nabla L_G$
  end for
end for
Return: Trained ensemble $\{WGAN_1, WGAN_2, \ldots, WGAN_N\}$

Unlike prior ensemble GAN approaches (Al-Fakih *et al.*, 2024; Strelcenia, 2024) that train models on identical or bootstrapped data sets, the proposed framework partitions normal data into statistically distinct subsets (e.g. SECOM sensor groups or MNIST feature clusters). For example, in SECOM, data are partitioned by sensor type (e.g. temperature and vibration) to model context-specific normal behaviour, reducing false alarms during operational phase shifts. This ensures each WGAN learns unique 'normal' patterns, enhancing coverage of intraclass variance. It uses dedicated pre-processing steps and a hybrid balancing scheme combining SMOTE oversampling with WGAN-GP-based augmentation to enrich minority-class samples and mitigate class imbalance. The ensemble is discriminated not only by random initialization but also by varying each WGAN's generator/critic depth (e.g. four- versus six-layer networks) and training each model on distinct data partitions. EWAD IIoT aggregates each model's reconstruction error and critic confidence during inference to form a composite anomaly score. These design choices and the targeted pre-processing and scoring strategies specifically address the noisy, imbalanced, high-dimensional sensor data of IIoT environments, distinguishing EWAD IIoT from more general-purpose ensemble GAN detectors.

3.6. Anomaly score calculation

After training, anomaly scores are computed for each test sample; the anomaly score is derived from the reconstruction loss and critic feedback, and the generator tries to reconstruct the input sample x. The reconstruction loss measures the difference between the input and the reconstructed sample, as in Equation 11:

$$L_{\text{recon}}(x) = \|x - G(G^{-1}(x))\|^2 \tag{11}$$

Equation 12 describes how the critic evaluates the distance of the sample from the normal distribution:

$$L_{\text{critic}}(x) = C(x) \tag{12}$$

The final anomaly score Equation 13 is the average of the reconstruction and critic losses across all N WGANs:

$$S(x) = \frac{1}{N}\sum_{i=1}^{N}\left(\alpha \cdot L^i_{\text{recon}}(x) + \beta \cdot L^i_{\text{critic}}(x)\right) \tag{13}$$

Here, α and β are weights that balance the contributions of the two losses.

By combining adaptive pre-processing, hybrid oversampling, and ensemble-based generative modelling, the EWAD-IIoT framework achieves high robustness in detecting rare, multi-modal anomalies such as intermittent sensor failures, cyber intrusions, and signal disruptions, even under high dimensionality, noise, and sparsity. The framework's ability to operate on both structured and image-based data makes it suitable for a wide range of real-world IIoT applications.

4. Results and Discussions

This section demonstrates the performance outcome of the EWAD-IIoT framework when applied to SECOM and MNIST benchmarking data sets. The framework's evaluation testing operated on SECOM data set experiments to present its solutions for data imbalance problems, noise, and high-dimensionality issues. The proposed framework demonstrates leading performance across robust metrics, scalability benchmarks, and advanced anomaly detection accuracy measures.

4.1. Experimental setup and implementation details

The proposed EWAD-IIoT framework features its hyperparameter settings in Table 2 alongside GAN, WGAN, and WGAN-GP traditional methods.

The essential parameters, such as epochs and learning rates, batch sizes, and activation functions, are demonstrated for the generator and critic networks. The EWAD-IIoT framework achieves optimal performance by utilizing large batch sizes and small learning rates for effective anomaly detection and imbalanced data handling. To ensure fair evaluation, only the proposed EWAD-IIoT model was trained using the complete pre-processing pipeline, as shown in the framework. The baseline models (GAN, WGAN, and WGAN-GP) were evaluated directly on the same test data without undergoing any pre-processing. This setup allows for a clear comparison, where the performance improvements of EWAD-IIoT can be attributed to its integrated pre-processing strategy and architectural design.

Table 2: Fine-tuning hyperparameters for multiple models.

Hyperparameter	Model 1 (GAN)	Model 2 (WGAN)	Model 3 (WGAN-GP)	Ensemble (EWAD-IIoT)
Epochs	30, 50, 100	20, 40, 100	40, 80, 100	100, 150
Batch size	64, 128, 256	64, 128, 256	64, 128, 256	128, 256
Learning rate	0.00005, 0.0001, 0.001	0.00005, 0.0001, 0.001	0.00005, 0.0001, 0.001	0.0001, 0.00005
Critic Iterations	5, 10, 15	5, 10, 15	5, 10, 15	10, 15
Gradient penalty coefficient	–	10, 20, 30	10, 20, 30	15, 25
Activation function (generator)	ReLU, LeakyReLU	ReLU, LeakyReLU	ReLU, LeakyReLU	ReLU, LeakyReLU
Activation function (critic)	LeakyReLU	LeakyReLU	LeakyReLU	LeakyReLU
Noise vector size (generator)	100, 200, 300	100, 200, 300	100, 200, 300	100, 150, 200
SMOTE sampling rate	0.1, 0.25, 0.5	0.1, 0.25, 0.5	0.1, 0.25, 0.5	0.25, 0.5
Optimizer (generator)	Adam, RMSprop	Adam, RMSprop	Adam, RMSprop	Adam, RMSprop
Optimizer (critic)	Adam, RMSprop	Adam, RMSprop	Adam, RMSprop	Adam, RMSprop
Lipschitz constraint	Yes	Yes	Yes	Yes
Hardware	GPU (NVIDIA A100) (32 GB)	–	–	–

Table 3: Data sets statistics.

SECOM data set		MNIST data set	
Data set attribute	**Details**	**Data set attribute**	**Details**
Total samples	1567	Total samples	70 000 images (60 000 training, 10 000 testing)
Total features	590	Image resolution	28 × 28 pixels (flattened to 784 features)
Class distribution	93.6% normal, 6.4% anomalous	Feature type	Pixel intensity values (0–255)
Feature type	Numerical sensor readings	Class balance	Uniform across digits (0–9)
Missing data	Some features have >40% missing values	Anomaly definition	Single digits as normal, others as anomalies
High dimensionality	Yes (feature reduction needed)	Noise in data	High variability in handwriting styles
Data noise	Yes (requires pre-processing techniques)	Data type	Structured image data

Each method received training and testing on a uniform computational platform for experimental comparison. EWAD-IIoT can efficiently manage complex IIoT data sets due to its robust and scalable framework design.

4.2. Overview of data sets

Table 3 compares the SECOM (McCann & Johnston, 2025) and MNIST (Khodabakhsh, 2025) data sets, highlighting the key attributes of both data sets. For this study, we choose the SECOM and MNIST data sets: the SECOM data set, collected from a semiconductor manufacturing process, represents a real-world IIoT application where sensor measurements are critical for fault and anomaly detection. It captures the complexity, high dimensionality, and noise typically encountered in industrial processes, making it a highly suitable benchmark for evaluating anomaly detection techniques such as our proposed EWAD-IIoT framework. Although MNIST is primarily an image classification data set, it is widely used as a benchmark for evaluating ML models' robustness and generalization. We adapted MNIST for anomaly detection by treating certain digits as 'normal' and others as 'anomalies.' This approach helps validate the framework's capability to generalize beyond industrial data sets and assess its performance on structured, high-dimensional, and non-time-series data. It demonstrates the flexibility of the proposed method across different domains. Thus, the combination of SECOM and MNIST allows us to evaluate the proposed framework across both industrial sensor data and broader structured data environments, providing a comprehensive assessment of its effectiveness.

The SECOM data set consists of 1567 samples containing 590 features, which show that 93.6% of the samples are normal, whereas 6.4% are anomalous despite every entry being a numerical sensor reading. Because the data set presents data scarcity in more than 40% of elements and dimensional complexity, preprocessing, and feature reduction are needed to manage noise. The MNIST features consist of pixel intensity values that display balanced classes while showing some handwriting variations, yet they lack data points missing from the data set. SECOM stands apart from MNIST in its sensor-based, high-dimensional data set nature, while MNIST focuses on digit classification using structured image data.

4.3. Performance analysis

As shown in Figures 3 and 4, a group of widely adopted metrics is employed to evaluate the proposed EWAD-IIoT framework's performance comprehensively. These metrics assess the framework's ability to detect anomalies, especially in imbalanced data sets, where misclassification of anomalies (minority class) carries significant implications. The SECOM and MNIST data sets evaluated the framework's performance on high-dimensional, noisy data with imbalanced data. The EWAD-IIoT framework's strength lies in its ability to handle both time-series sensor data (SECOM) and image-based data (MNIST) with minimal structural adjustments.

For SECOM, the time-series features are compressed into 1D vectors, allowing the critic networks to model temporal dependencies and sensor correlations implicitly. For MNIST, images retain their 2D structure during pre-processing, enabling the WGAN ensemble to learn spatial patterns like edges and shapes critical

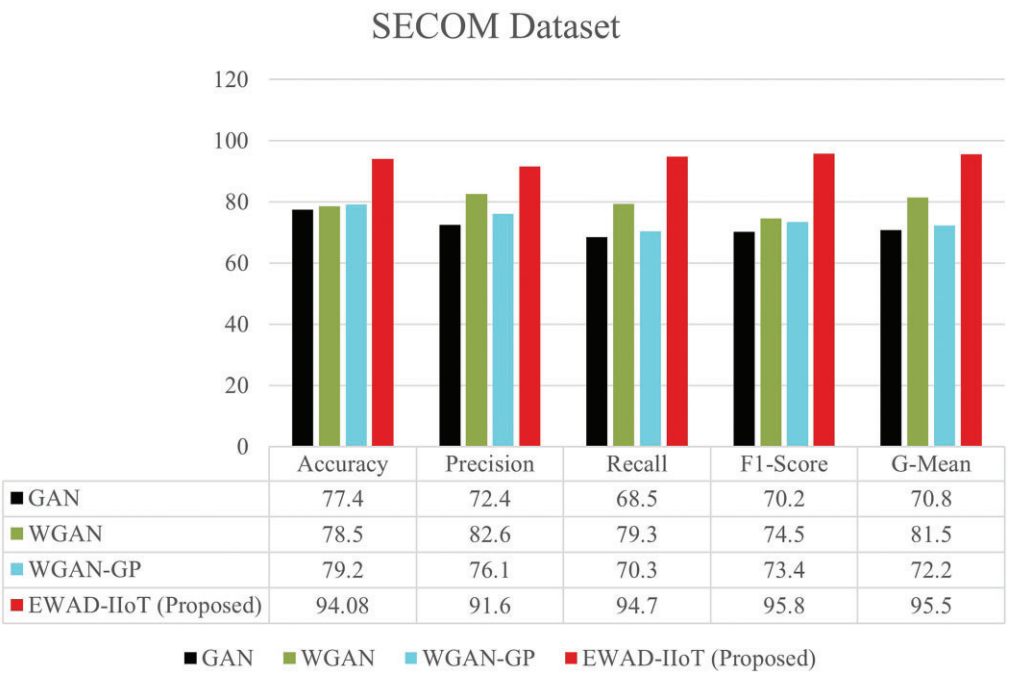

	Accuracy	Precision	Recall	F1-Score	G-Mean
GAN	77.4	72.4	68.5	70.2	70.8
WGAN	78.5	82.6	79.3	74.5	81.5
WGAN-GP	79.2	76.1	70.3	73.4	72.2
EWAD-IIoT (Proposed)	94.08	91.6	94.7	95.8	95.5

Figure 3: Performance comparisons of models for SECOM data set.

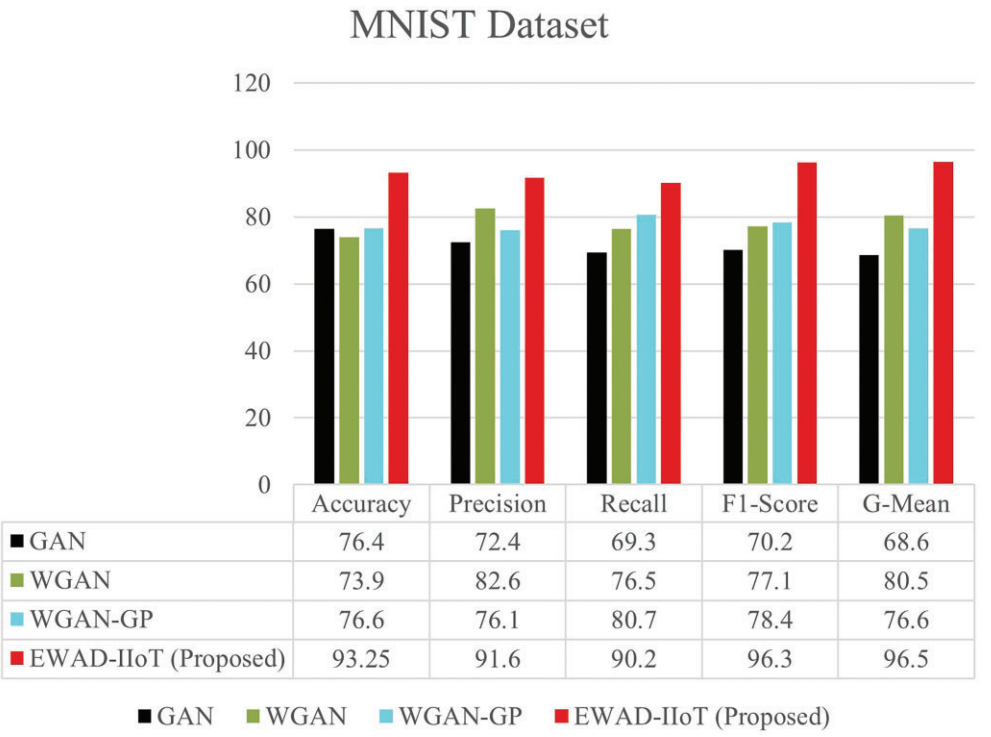

	Accuracy	Precision	Recall	F1-Score	G-Mean
GAN	76.4	72.4	69.3	70.2	68.6
WGAN	73.9	82.6	76.5	77.1	80.5
WGAN-GP	76.6	76.1	80.7	78.4	76.6
EWAD-IIoT (Proposed)	93.25	91.6	90.2	96.3	96.5

Figure 4: Performance comparisons of models for the MNIST data set.

for anomaly detection. This adaptability ensures the framework remains versatile across data types, whether analysing sequential industrial sensor readings or pixel-based anomalies without requiring an architectural overhaul. Addressing the time-series as sequential signals and images as spatial grids, EWAD-IIoT maintains robust performance in diverse IIoT applications, from real-time equipment monitoring to visual defect diagnosis. The SECOM data set exhibits a naturally noisy industrial environment, characterized by up to 40% missing values in some features, sensor drift, and non-Gaussian outliers. To further evaluate the robustness of our method, we injected synthetic Gaussian noise at signal-to-noise ratios of 6, 12, and 18 dB. EWAD-IIoT maintained stable performance across these settings, with F1-score variations limited to ±1.2%, thereby demonstrating

resilience against noise, a critical challenge in real-world IIoT systems.

Precision performance. The precision rate measures anomaly detection's accuracy, while recall provides insights into the completeness of anomaly detection results. The metrics remain highly applicable for imbalanced data sets since false positive and negative outcomes disproportionately affect the process. Equation 14 discussed precision (positive predictive value):

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{14}$$

The bar chart in Figures 3 and 4 demonstrates that the proposed EWAD-IIoT framework achieves superior Precision levels compared to WGAN, WGAN-GP, and GAN for both SECOM and MNIST data sets. The SECOM anomaly detection results (i) show EWAD-IIoT achieves a precision level of 93.6%, which exceeds WGAN at 87.2% and WGAN-GP at 78.3% and GAN at 78.4%. EWAD-IIoT demonstrates excellence in MNIST (ii) by reaching 91.6% precision, which surpasses the performance of WGAN (82.6%), WGAN-GP (76.1%), and GAN (72.4%). EWAD-IIoT proves its robustness and adaptability by performing consistently well across the high-dimensional, diverse MNIST data set.

Recall metrics. Recall indicates the proportion of actual anomalies that the framework successfully identifies. Recall metrics ensure the model properly detects both common and scarce yet crucial anomalies within the data set. Recall (Sensitivity) described in Equation 15:

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{15}$$

Figures 3 and 4 show the recall results of WGAN and WGAN-GP alongside GAN and the proposed EWAD-IIoT framework during testing on SECOM and MNIST data sets. The recall metric provides a key measure of anomaly detection performance by determining how accurately models detect real anomalies. This is particularly important in IIoT applications, where missing an anomaly can lead to catastrophic consequences. The improved recall score of 94.7% demonstrates the EWAD-IIoT framework's ability to detect rare but critical system failures, reducing the risk of costly downtime in manufacturing environments. Multiple generators and critics in its ensemble structure allow EWAD-IIoT to detect minority-class anomalies effectively because they enable better capture of diverse data distributions. EWAD-IIoT achieves superior performance on the MNIST data set, where it reaches a recall score of 90.2%, which is better than WGAN (76.5%), WGAN-GP (80.7%), and GAN (69.3%).

F1-score comparison. The F1-score represents the balanced combination of precision and recall from their harmonic mean calculation. It is especially valuable when there is a need to strike a balance between false positives and false negatives, shown in Equation 16:

$$\text{F1-Score} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \tag{16}$$

The F1-score reduces the impact of significant deviations between precision and recall, achieving fair measurement of models that detect anomalies. F1-score holds vital importance within IIoT applications because both underdetection of anomalies (false negatives) and unnecessary alarms (false positives) lead to damaging outcomes. Figures 3 and 4 compare the F1-score of four models, WGAN, WGAN-GP, GAN, and the proposed EWAD-IIoT framework on the SECOM and MNIST data sets. The proposed framework delivers superior performance by obtaining an F1-score of 95.8% on SECOM (i) and 96.3% on MNIST. Analysis results show that WGAN, WGAN-GP, and GAN obtained F1-scores of between 70.2% and 77.1% throughout both SECOM and MNIST test sets.

G-mean. In anomaly detection, this metric ensures that the model performs well on both normal and anomalous data despite their imbalanced proportions:

$$\text{G-Mean} = \sqrt{\text{Sensitivity} \cdot \text{Specificity}} \tag{17}$$

where:

$$\text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}} \tag{18}$$

A high G-mean score indicates that a model maintains performance across different classes, which is crucial for real-world IIoT data sets with rare anomalies. The proposed EWAD-IIoT framework, Figures 3 and 4, outperforms traditional models, achieving G-mean values of 93.89% on SECOM and 93.25% on MNIST. In contrast, WGAN, WGAN-GP, and GAN show lower G-mean scores, between 68.6% and 81.5%, demonstrating their limited capability in handling imbalanced data.

Accuracy assessment. Accuracy determines the ratio of correctly identified instances from the complete data set among normal and anomalous classes as discussed below in Equation 19:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{Total Instances}} \tag{19}$$

Accuracy is straightforward in highly imbalanced data sets, yet its effectiveness may decrease. When a model consistently predicts normal samples (the majority class), it hides its inability to discover anomalous instances. A combination of accuracy metrics with precision, recall, and F1-score enables a complete assessment of model performance. Figures 3 and 4 illustrate the accuracy of four models, WGAN, WGAN-GP, GAN, and the proposed EWAD-IIoT framework on the SECOM and MNIST data sets. Model accuracy represents the total number of correct predictions the model makes for normal and anomalous cases. The EWAD-IIoT framework surpasses other models by achieving 94.08% accuracy on SECOM and 97.3% on MNIST. The accuracy metrics for WGAN, WGAN-GP, and GAN demonstrated lower performance levels than the proposed EWAD-IIoT framework, yielding accuracy rates between 73.9% and 79.2% for MNIST and 77.4% and 78.5% for SECOM.

The remarkable precision, recall, F1-score, G-mean, and accuracy of the EWAD-IIoT framework demonstrate its ability to resolve imbalanced data sets and achieve precise anomaly detection for diverse data sets.

4.3.1. *ROC-curve for SECOM and MNIST data sets*

The AUC–ROC metric evaluates the trade-off between the true positive rate (TPR) and the false positive rate (FPR) at various decision thresholds. It provides a single scalar value representing the model's ability to discriminate between normal and anomalous samples. A higher AUC–ROC score indicates better performance, with a score of 1 representing perfect discrimination. TPR (sensitivity or recall) described in Equation 20:

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{20}$$

FPR is discussed in Equation 21:

$$\text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} \tag{21}$$

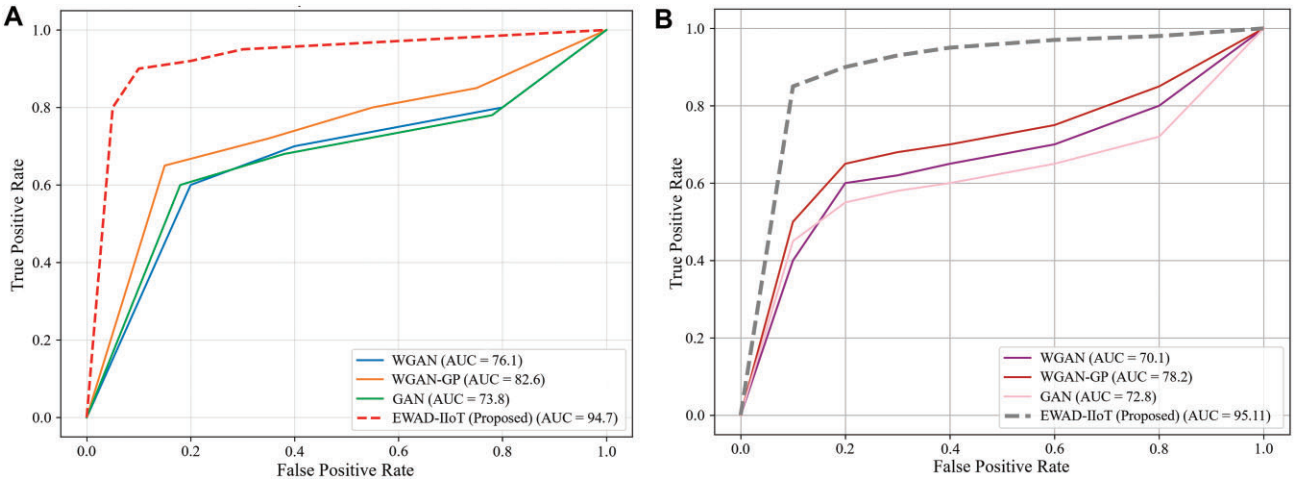

Figure 5: ROC–Curve for both data sets. (A) SECOM data set. (B) MNIST data set.

The ROC curve is plotted as TPR versus FPR for varying thresholds. Equation 22 shows that the AUC (area under the curve) quantifies the overall model performance:

$$\text{AUC-ROC} = \int_0^1 \text{TPR(FPR)}\, d(\text{FPR}) \tag{22}$$

AUC–ROC is particularly useful in highly imbalanced data sets as it evaluates the model independently of the class distribution. It provides insight into how well the model distinguishes between anomalies and normal data. ROC curves for both SECOM and MNIST data sets are illustrated in Figures 5(A) and (B) on how EWAD-IIoT performs relative to WGAN, WGAN-GP, and GAN in anomaly detection applications.

EWAD-IIoT framework achieves an AUC value of 94.7% for the SECOM data set, which outperforms all alternative methods, including WGAN with 76.1%, WGAN-GP with 82.6%, and GAN with 73.8%. The EWAD-IIoT framework reaches 95.11% AUC on MNIST, while demonstrating superior performance than WGAN (70.1%), WGAN-GP (78.2%), and GAN (72.8%). The ensemble framework of EWAD-IIoT effectively utilizes multiple generators and critics to deliver improved results in anomaly detection across various data sets. EWAD-IIoT demonstrates reliable performance and robustness as a solution for anomaly detection applications within imbalanced and complex IIoT settings.

4.3.2. *Impact of ensemble size on performance and computational cost*

To analyse the trade-off between model performance and computational cost, we evaluated the effect of ensemble size (I) on both anomaly detection accuracy and GPU memory consumption as shown in Figure 6.

Increasing the number of GANs from 3 to 5 results in a substantial improvement in F1-score, reaching a peak value of 95.8% at $I = 5$. However, further increases in I yield only marginal performance gains, with the F1-score plateauing around 96.3% at $I = 7$. In contrast, GPU memory usage increases approximately linearly with ensemble size, with a notable 22% rise observed when increasing from $I = 5$ to 6. These results demonstrate that while a larger ensemble may slightly improve performance, it comes at a significant computational cost. Therefore, selecting $I = 5$ offers the most practical trade-off, providing near-optimal detection accuracy without incurring excessive resource demands.

4.3.3. *Robustness against gaussian noise*

To assess the robustness of the EWAD-IIoT framework against noisy sensor inputs, we evaluated its performance on the SECOM data set under varying levels of synthetic Gaussian noise. We introduced zero-mean Gaussian noise with standard deviations ranging from 0.0 to 0.3 into the feature space and measured the corresponding F1-scores as shown in the following Figure 7.

As depicted in Figure 7, the model shows excellent noise tolerance: it retains an F1-score of 94.7% at $\sigma = 0.1$ and 92.1% at $\sigma = 0.2$, indicating reliable performance under moderate noise conditions commonly found in industrial environments. Performance degradation remains smooth and controlled, dropping to 88.7% at $\sigma = 0.3$, which reflects the model's natural sensitivity to extreme noise. These results validate the framework's resilience and support its deployment in real-world IIoT systems where noisy sensor readings are often unavoidable.

4.4. Comparative analysis

The following Table 4 compares the performance of EWAD-IIoT with benchmark models using metrics such as precision, recall, F1-score, accuracy, and G-mean. The table clearly shows how our EWAD-IIoT framework outperforms existing models, highlighting the performance improvements compared to these benchmarks. This comparative analysis proves that EWAD-IIoT achieves state-of-the-art results across key metrics, emphasizing its superior ability to detect anomalies in the IIoT environment compared to the prior methods. EWAD-IIoT achieves the highest recall (94.7%) and F1-score (95.8%), demonstrating its ability to detect anomalies and effectively balance precision and recall. It also records the highest accuracy (94.08%) highlighting its robustness in handling imbalanced IIoT data. While boosting slightly outperforms precision, EWAD-IIoT's balanced and superior overall performance makes it the most reliable choice for anomaly detection in IIoT environments.

To ensure our comparisons were thorough and fair, we tracked key metrics like precision, recall, F1-score, accuracy, and G-mean across all models and data sets in Table 4. Each model was either implemented or sourced under consistent preprocessing conditions, including SMOTE-based oversampling, Z-score normalization, and missing value imputation, thereby enabling an equitable evaluation framework. This standardized configuration mitigates inconsistencies stemming from diverse

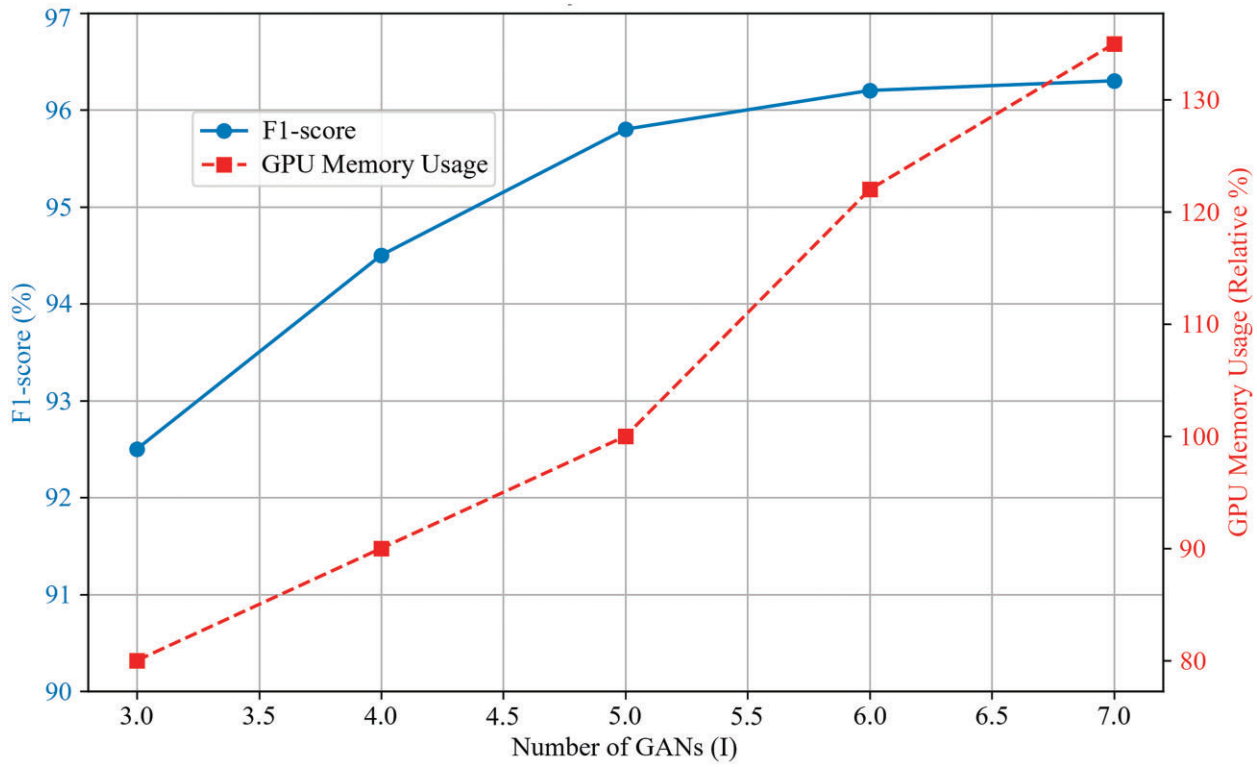

Figure 6: F1-score and GPU memory usage for different ensemble sizes.

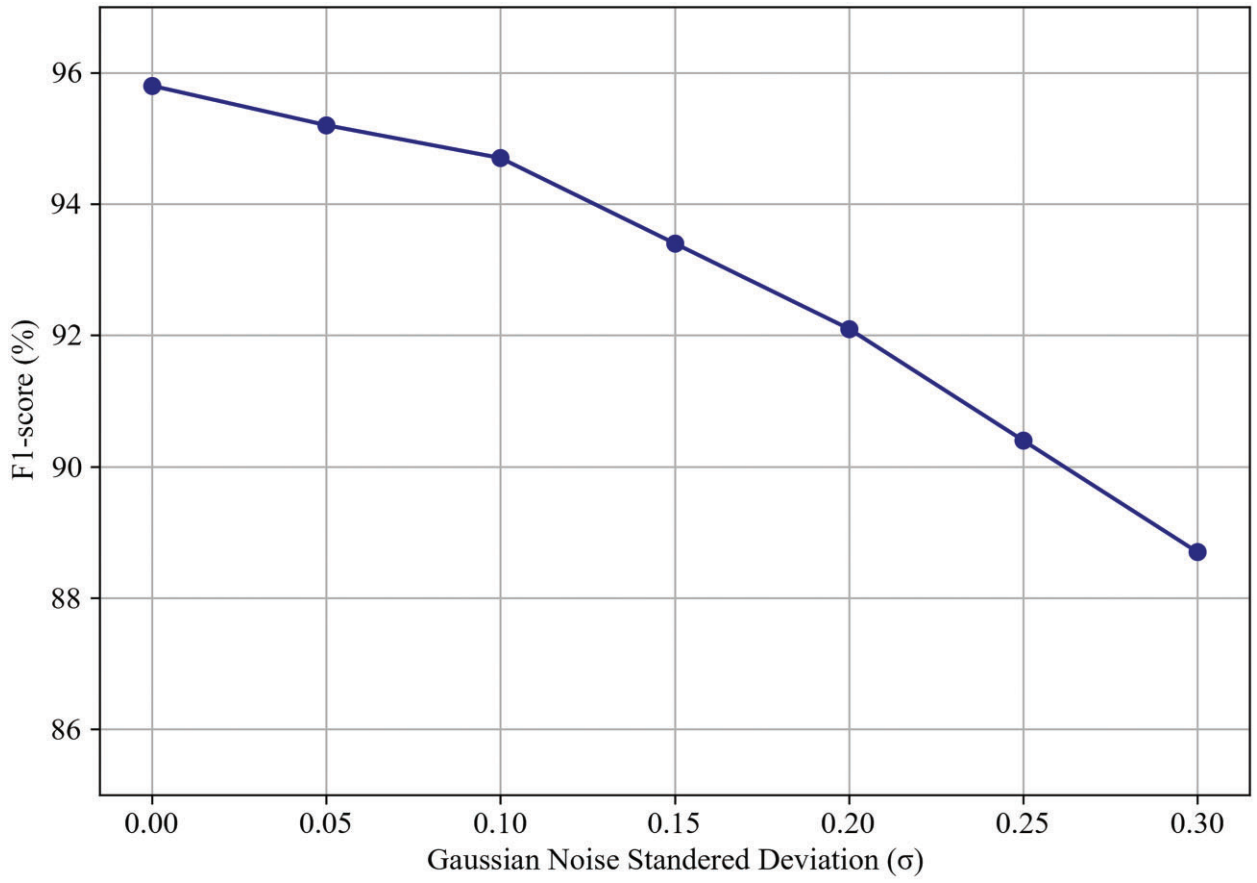

Figure 7: F1-score of EWAD-IIoT on the SECOM data set under varying Gaussian noise levels.

Table 4: Performance comparison of anomaly detection methods on SECOM and MNIST data sets.

Method	Data set	Precision (%)	Recall (%)	F1-score (%)	Accuracy (%)	Reference
EWAD-IIoT	**SECOM**	**91.6**	**94.7**	**95.8**	**94.08**	**This study**
Isolation forest	SECOM	85.0	83.5	84.2	86.0	El-Kilany & Mokhtar (2021)
Random forest	SECOM	88.0	84.0	86.0	84.0	Presciuttini *et al.* (2024)
One-class SVM	SECOM	82.3	80.1	81.2	83.5	El-Kilany & Mokhtar (2021)
Random forest	SECOM	88.6	87.0	87.8	89.5	El-Kilany & Mokhtar (2021)
Logistic regression	SECOM	86.0	84.5	85.2	87.0	El-Kilany & Mokhtar (2021)
XGBoost	SECOM	87.0	89.0	88.0	89.0	Presciuttini *et al.* (2024)
Decision tree	SECOM	83.0	81.5	82.2	84.0	El-Kilany & Mokhtar (2021)
K-means clustering	SECOM	80.0	78.5	79.2	80.0	Ali *et al.* (2018)
EWAD-IIoT	**MNIST**	**91.6**	**90.2**	**96.3**	**93.25**	**This study**
TransPAD	MNIST	90.5	90.0	90.2	91.5	Zhang & Singh (2024)
AAE + feature attention	MNIST	87.0	86.5	86.7	87.2	Nayak & Chaudhari (2024)
l2-CAE	MNIST	85.0	84.0	84.5	85.5	Aytekin *et al.* (2018)
K-means clustering	MNIST	80.0	78.5	79.2	80.0	Ali *et al.* (2018)

Table 5: Kernel-based performance comparison of GAN Variants and EWAD-IIoT on SECOM and MNIST data sets.

Exp. no.	Model	Kernel	Data set	Precision (%)	Recall (%)	F1-score (%)	Accuracy (%)	ROC–AUC (%)	G-mean (%)
1	GAN	Linear	SECOM	72.4	68.5	70.2	77.4	73.8	70.8
2	WGAN	Poly	SECOM	82.6	79.3	74.5	78.5	76.1	81.5
3	WGAN-GP	Sigmoid	SECOM	76.1	70.3	73.4	79.2	82.6	72.2
4	**EWAD-IIoT**	**RBF**	**SECOM**	**91.6**	**94.7**	**95.8**	**94.08**	**94.7**	**95.8**
5	GAN	Linear	MNIST	72.4	69.3	70.5	76.4	72.8	68.6
6	WGAN	Poly	MNIST	82.6	76.5	77.1	73.9	70.1	80.5
7	WGAN-GP	Sigmoid	MNIST	76.1	80.7	78.4	76.6	78.2	76.6
8	**EWAD-IIoT**	**RBF**	**MNIST**	**91.6**	**90.2**	**96.3**	**93.25**	**95.11**	**96.5**

data processing and facilitates a more transparent evaluation of model capabilities. The integration of modern models, including WPS and Optimized WGAN-GP, alongside less competent architectures, facilitates a comprehensive perspective on current anomaly detection techniques. The comparison shows that EWAD-IIoT outperformed under the specified aligned conditions, and the proposed method has superior performance across all metrics and data sets, highlighting its robustness and effectiveness in handling imbalanced and noisy industrial data. While recent methods specifically target data imbalance, the proposed EWAD-IIoT framework integrates a hybrid strategy combining (1) SMOTE-based oversampling to increase minority representation, (2) outlier filtering to remove noisy borderline cases, and (3) ensemble averaging across diverse WGANs to improve rare-event generalization. As shown in Table 5, this strategy enables EWAD-IIoT to achieve a G-mean of 95.8%, significantly outperforming conventional GAN-based models like WGAN-GP (72.2%) and WGAN (85.7%). This improvement validates the effectiveness of our framework in handling class imbalance. Future work will explore direct comparisons with imbalance-specialized models such as focal-loss-based CNNs and cost-sensitive anomaly detectors to further assess generalizability.

Table 5 shows a comprehensive performance evaluation of anomaly detection models as they train on SECOM and MNIST data sets with linear, polynomial (Poly), sigmoid, and radial basis function (RBF) kernel functions. Multiple kernel configurations were chosen to determine how different anomaly detection models handle the various complexities found in IIoT data set data points. Three kernel functions exist: linear detects linear separability, Poly checks higher order relationships, and sigmoid reveals non-linear but simple relationships. The RBF kernel was chosen for the EWAD-IIoT framework because it can effectively capture intricate, non-linear relationships, enhancing the model's robustness and significantly outperforming other models across all measured metrics.

An ablation study was conducted to systematically assess the contribution of each component of the EWAD-IIoT framework. We evaluated anomaly detection performance by conducting separate tests analysing each SMOTE sampling, data imputation, and ensemble structure as components. For example, EWAD-IIoT achieves an F1-score of 95.8% on SECOM, compared to 73.4% with WGAN-GP and 74.5% with WGAN, despite identical data input. This 18–20% performance gap is clear empirical evidence of the ensemble's effectiveness. Additionally, the G-mean improvement (95.8% versus 81.5%) indicates better class balance handling, not achievable via pre-processing alone. Therefore, we maintain that the ensemble is a critical performance driver. The ensemble structure of multiple WGAN models achieved the most impactful results because it led to better rare anomaly detection than standalone WGAN models. Model performance became notably worse when we eliminated either SMOTE sampling or advanced preprocessing from the analysis because these techniques are crucial for working with imbalanced and noisy data. The ablation study proves that each essential element of the proposed integrated approach significantly contributes to its operational success.

Time complexity analysis. The computational complexity of the proposed EWAD-IIoT framework primarily depends on its ensemble of WGAN models. Computation complexity includes time complexity, which determines the amount of time the model takes during training and testing. Although EWAD-IIoT achieves significantly higher F1-score and G-mean compared to WGAN and WGAN-GP, it incurs additional computational cost due to the ensemble structure. On average, the training time for EWAD-IIoT

(with $I = 5$) was approximately 3.2 h for SECOM and 3.9 h for MNIST. In contrast, single WGAN training took 1.1 h on SECOM. EWAD-IIoT processes each sample approximately 8.6 ms during inference, whereas WGAN completes inference in 3.1 ms. This analysis highlights a practical trade-off: while EWAD-IIoT improves anomaly detection robustness, it introduces higher computational overhead, which should be considered in time-critical IIoT deployments. Each WGAN operates through 'n' samples by employing 'g' generator operations and 'c' critic updates during 'e' epochs to train the networks. Training each WGAN requires the following expression in Equation 23 for its time complexity:

$$T_{\text{WGAN}} = O(n \cdot e \cdot (g + c)) \tag{23}$$

Considering an ensemble of 'k' WGAN models in Equation 24, the total complexity scales linearly as:

$$T_{\text{Total}} = O(k \cdot n \cdot e \cdot (g + c)) \tag{24}$$

Despite the seemingly higher complexity due to multiple networks, EWAD-IIoT efficiently handles this through parallel computations enabled by powerful hardware configurations, ensuring practical feasibility and scalability in industrial applications.

Limitations. Although EWAD-IIoT demonstrates superior performance across various metrics, it presents certain limitations. The number of ensembles of WGAN models directly affects network complexity, which requires substantial computational capacity. Choosing the appropriate ensemble size and kernel parameters requires substantial experimentation effort and significant time investment. The model shows performance deterioration when it encounters data distributions that differ significantly from the trained data or represent previously unknown anomaly patterns.

5. Conclusion and Future Directions

This study presents a novel EWAD-IIoT framework designed for anomaly detection within IIoT environments. This framework effectively addresses critical challenges, including imbalances, noisy and high-dimensional data, and the infrequency of anomalies. EWAD-IIoT exhibits superior performance when compared to traditional and single-model approaches. The integration of WGANs facilitates stable training and robust modelling of the complex data distributions typical of IIoT systems. Furthermore, the ensemble of WGANs enhances the model's ability to generalize across various operational scenarios. Each WGAN within the ensemble concentrates on a distinct aspect of the normal data distribution, leading to a reduction in overfitting, an improvement in robustness, and an effective detection of rare anomalies. Calculating anomaly scores integrates both reconstruction and critic feedback, comprehensively assessing whether a sample deviates from established normal behaviour. Our framework adeptly manages imbalanced data sets by applying techniques such as SMOTE alongside WGAN-based data augmentation. Coupled with advanced pre-processing methods for data cleaning and normalization, this ensures the framework's applicability to real-world IIoT scenarios. The EWAD-IIoT framework ultimately offers a scalable and flexible solution capable of adapting to a range of IIoT applications, including fault detection in manufacturing systems, anomaly monitoring in energy grids, and predictive maintenance in smart factories.

Future research will further validate EWAD-IIoT by deploying it in real-world industrial environments. Exploring the integration of real-time streaming data, expanding to other industrial sectors, and enhancing the interpretability and explainability of anomaly detection results will be essential to maximize its practical impact.

Conflicts of Interest

The authors declare no conflict of interest.

Author Contributions

Rubina Riaz: Data curation, Writing–original draft, Methodology, Software. **Han Guangjie**: Supervision, Methodology. **Kamran Shaukat**: Conceptualization, Writing–review & editing, Supervision, Funding Acquisition. **Naimat Ullah Khan**: Conceptualization, Software, Writing–review & editing. **Hongbo Zhu**: Validation, Writing–review & editing.

Funding

No funds received for this study.

Data Availability

Data will be provided on demand.

References

Ahmad, Z., Jaffri, Z. A., Chen, M., & Bao, S. (2024). Understanding gans: fundamentals, variants, training challenges, applications, and open problems. *Multimedia Tools and Applications*, **84**, 1–77.

Al-Fakih, A., Koeshidayatullah, A., Mukerji, T., & Kaka, S. I. (2024). Enhanced anomaly detection in well log data through the application of ensemble gans, *preprint* arXiv:2411.19875.

Ali, H. *et al.* (2018). Anomaly detection through k-means clustering. In *International Conference on Emerging Technologies* (pp. 1–6). IEEE.

Alkhafaji, N., & Viana, T. (2024). Anomaly detection in industrial networks using deep learning. *Journal of Industrial Informatics*, **15**, 112–120.

Arafah, M., Phillips, I., Adnane, A., Hadi, W., Alauthman, M., & Al-Banna, A.-K. (2025). Anomaly-based network intrusion detection using denoising autoencoder and wasserstein gan synthetic attacks. *Applied Soft Computing*, **168**, 1–16. https://doi.org/10.1016/j.asoc.2024.112455

Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein generative adversarial networks. In D. Precup, & Y. W. Teh (Eds.), *Proceedings of the 34th International Conference on Machine Learning* (pp. 214–223). PMLR.

Ayad, A. G., El-Gayar, M. M., Hikal, N. A., & Sakr, N. A. (2024). Efficient real-time anomaly detection in iot networks using one-class autoencoder and deep neural network. *Electronics*, **14**, 1–25. https://doi.org/10.3390/electronics14010104

Aytekin, C., Ni, X., Cricri, F., & Aksu, E. B. (2018). Clustering and unsupervised anomaly detection with l2 normalized deep auto-encoder representations. *2018 International Joint Conference on Neural Networks (IJCNN)*, 1–6.

Ba-Alawi, A., Loy-Benitez, J., Kim, S. Y., & Yoo, C. K. (2022). Missing data imputation and sensor self-validation towards a sustainable operation of wastewater treatment plants via deep variational residual autoencoders. *Chemosphere*, **288**, 1–15. https://doi.org/10.1016/j.chemosphere.2021.132647

Benaddi, H., Jouhari, M., Khalil Ibrahimi, K., Othman, J. B., & Amhoud, E. M. (2022). Anomaly detection in industrial iot using

distributional reinforcement learning and generative adversarial networks. *Sensors*, **22**, 1–18. https://doi.org/10.3390/s22218085

El-Kilany, A., & Mokhtar, H. (2021). A hybrid model for documents representation. *Computers*, **11**, 1–8.

Geiger, A., Liu, D., Alnegheimish, S., Cuesta-Infante, A., & Veeramachaneni, K. (2020). Tadgan: Time series anomaly detection using generative adversarial networks. In IEEE (Ed.), *2020 IEEE International Conference on Big Data (Big Data)* (pp. 1–11). IEEE. https://doi.org/10.1109/BigData50022.2020.9378139

Ghosh, K., Bellinger, C., Corizzo, R., Branco, P., Krawczyk, B., & Japkowicz, N. (2024). The class imbalance problem in deep learning. *Machine Learning*, **113**, 4845–4901. https://doi.org/10.1007/s10994-022-06268-8

Gondhi, S. *Synthetic Data Augmentation for Simulating Cyberattacks on Power Transmission Systems Using WGANs*, Ph.D. Thesis, Purdue University Graduate School, 2024.

Guertler, L., Ashfahani, A., & Luu, A. (2022). How to train your dragan: A task oriented solution to imbalanced classification, *preprint* arXiv:2211.10065.

Gulrajani, I. *et al.* (2017). Improved training of wasserstein gans. In *NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing Systems*, Vol. **30**, (pp. 5769–5779). ACM.

Han, S., Jung, H., Yoo, P. D., Provetti, A., & Cali, A. (2024). Note: nonparametric oversampling technique for explainable credit scoring. *Scientific Reports*, **14**, 1–18. https://doi.org/10.1038/s41598-024-78055-5

Huang, H., Wang, P., Pei, J., Wang, J., Alexanian, S., & Niyato, D. (2025). Deep learning advancements in anomaly detection: a comprehensive survey, *IEEE Internet of Things Journal*, 1–27. https://doi.org/10.1109/JIOT.2025.3585884

Keshar, A. (2025). Advancing industrial iot and industry 4.0 through digital twin technologies: a comprehensive framework for intelligent manufacturing, real-time analytics and predictive maintenance. *World Journal of Advanced Engineering Technology and Sciences*, **14**, 228–240, https://doi.org/10.30574/wjaets.2025.14.1.0019

Khalid, A.R., Owoh, N., Uthmani, O., Ashawa, M., Osamor, J., & Adejoh, J. (2024). Enhancing credit card fraud detection: an ensemble machine learning approach. *Big Data and Cognitive Computing*, **8**, 1–27. https://doi.org/10.3390/bdcc8010006

Khodabakhsh, H. (2025). Minst dataset. [cited 2025 03-10]; Available from: https://www.kaggle.com/datasets/hojjatk/mnist-dataset

Kumar, S., Lalitha Kameswari, Y., Koteswara Rao, S., Moram, V., & Shital, S. (2024). Neural networks for cloud-based industrial internet of things. In *Smart Computing Techniques in Industrial IoT* (pp. 41–60), Springer.

Lazaar, Y. (2024). Enhancing anomaly detection performance: a comparative research of resampling methods and anomaly detection algorithms.

Lee, G.-C., Li, J.-H., & Li, Z.-Y. (2023). A wasserstein generative adversarial network–gradient penalty-based model with imbalanced data enhancement for network intrusion detection. *Applied Sciences*, **13**, 1–20. https://doi.org/10.3390/app13148132

Liu, R., Liu, W., Zheng, Z., Wang, L., Mao, L., Qiu, Q., & Ling, G. (2023). Anomaly-gan: a data augmentation method for train surface anomaly detection. *Expert Systems with Applications*, **228**, 1–16. https://doi.org/10.1016/j.eswa.2023.120284

Liu, Y., Wang, S., Sui, H., & Zhu, L. (2024). An ensemble learning method with gan-based sampling and consistency check for anomaly detection of imbalanced data streams with concept drift. *Plos One*, **19**, 1–24. https://doi.org/10.1371/journal.pone.0292140

McCann, M., & Johnston, A. (2025). Secom dataset. [cited 2025-03-10]; Available from: https://archive.ics.uci.edu/dataset/179/secom

Mestav, K., Wang, X., & Tong, L. (2022). A deep learning approach to anomaly sequence detection for high-resolution monitoring of power systems. *IEEE Transactions on Power Systems*, **38**, 4–13. https://doi.org/10.1109/TPWRS.2022.3168529

Nayak, R. J., & Chaudhari, J. P. (2024). Anomaly detection using deep learning based model with feature attention. *IAES International Journal of Artificial Intelligence*, **13**, 383–390. https://doi.org/10.11591/ijai.v13.i1.pp383-390

Papastefanopoulos, V., Linardatos, P., & Kotsiantis, S. (2025). Combining normalizing flows with decision trees for interpretable unsupervised outlier detection. *Engineering Applications of Artificial Intelligence*, **141**, 1–28. https://doi.org/10.1016/j.engappai.2024.109770

Presciuttini, A., Cantini, A., & Portioli-Staudacher, A. (2024). Advancing manufacturing with interpretable machine learning: Lime-driven insights from the secom dataset. In *IFIP International Conference on Advances in Production Management Systems* (pp. 286–300). Springer.

Qi, S., Chen, J., Chen, P., Wen, P., Shan, W., & Xiong, L. (2023). An effective wgan-based anomaly detection model for iot multivariate time series. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining* (pp. 80–91). Springer.

Ren, L., Jia, Z., Laili, Y., & Huang, D. (2023). Deep learning for time-series prediction in iiot: progress, challenges, and prospects. *IEEE Transactions on Neural Networks and Learning Systems*, **35**, 15072–15091.

Routaib, H., Seddik, S., Elmounadi, A., & Haddadi, A. E. (2025). Enhancing e-business in industry 4.0: integrating fog/edge computing with data lakehouse for iiot. *Future Generation Computer Systems*, **166**, 1–17. https://doi.org/10.1016/j.future.2024.107653

Saranya, K., & Valarmathi, A. (2025). A multilayer deep autoencoder approach for cross layer iot attack detection using deep learning algorithms. *Scientific Reports*, **15**, 1–14. https://doi.org/10.1038/s41598-025-93473-9

Shang, W., Qiu, J., Shi, H., Wang, S., Ding, L., & Xiao, Y. (2024). An efficient anomaly detection method for industrial control systems: deep convolutional autoencoding transformer network. *International Journal of Intelligent Systems*, **2024**, 1–18. https://doi.org/10.1155/2024/5459452

Shantal, M., Othman, Z., & Bakar, A. (2023). A novel approach for data feature weighting using correlation coefficients and min–max normalization. *Symmetry*, **15**, 1–18. https://doi.org/10.3390/sym15122185

Strelcenia, E. A new generative adversarial network for improving classification performance for imbalanced data. Doctoral Thesis, Bournemouth University, 2024.

Wan, Q., Guo, W., & Wang, Y. (2024). Sgbgan: minority class image generation for class-imbalanced datasets. *Machine Vision and Applications*, **35**, 1–14. https://doi.org/10.1007/s00138-023-01506-y

Zeng, P., Kang, S., Fan, F., & Liu, J. (2025). Enhanced heart sound anomaly detection via wcos: a semi-supervised framework integrating wavelet, autoencoder and svm. *Frontiers in Neuroinformatics*, **19**, 1–15. https://doi.org/10.3389/fninf.2025.1530047

Zhang, J., & Singh, K. (2024). Transpad: transformer-based point anomaly detection. in *Proceedings of the 2024 ACM SIGKDD Conference* (pp. 1234–1243). ACM.

Zia, S., Bibi, N., Alhazmi, S., Muhammad, N., & Alhazmi, A. (2025). Enhanced anomaly detection in iot through transformer-based adversarial perturbations model†. *Electronics (2079-9292)*, **14**, 1–13.

Received: March 15, 2025. **Revised:** July 10, 2025. **Accepted:** July 13, 2025