

SD-CGAN: Conditional Sinkhorn Divergence GAN for DDoS Anomaly Detection in IoT Networks

Henry Onyeka¹, Emmanuel Samson¹, Liang Hong¹, Tariqul Islam², Imtiaz Ahmed³, Kamrul Hasan¹,

¹ Tennessee State University, Nashville, TN, USA

² University of Maryland Baltimore County, Baltimore, Maryland, USA

³ Howard University, Washington, DC, USA

Email: {honyeka,esamson,lhong}@tnstate.edu, {mtislam}@umbc.edu, {imtiaz.ahmed}@howard.edu, {mhasan1}@tnstate.edu

Abstract—The increasing complexity of IoT edge networks presents significant challenges for anomaly detection, particularly in identifying sophisticated Denial-of-Service (DoS) attacks and zero-day exploits under highly dynamic and imbalanced traffic conditions. This paper proposes SD-CGAN, a Conditional Generative Adversarial Network framework enhanced with Sinkhorn Divergence, tailored for robust anomaly detection in IoT edge environments. The framework incorporates CTGAN-based synthetic data augmentation to address class imbalance and leverages Sinkhorn Divergence as a geometry-aware loss function to improve training stability and reduce mode collapse. The model is evaluated on exploitative attack subsets from the CICDDoS2019 dataset and compared against baseline deep learning and GAN-based approaches. Results show that SD-CGAN achieves superior detection accuracy, precision, recall, and F1-score while maintaining computational efficiency suitable for deployment in edge-enabled IoT environments.

Index Terms—Generative Adversarial Network (GAN), Sinkhorn Divergence, Anomaly Detection, Machine Learning, Network Security

I. INTRODUCTION

The evolution of IoT edge networks has enabled ultra-low latency applications such as autonomous vehicles, industrial automation, and mission-critical connected systems. However, this proliferation has significantly increased the attack surface, making IoT infrastructures vulnerable to sophisticated cyber threats like Distributed Denial-of-Service (DDoS) attacks and zero-day exploits. The tight integration of IoT, SDN, and edge computing creates dynamic and heterogeneous traffic patterns that challenge existing anomaly detection systems [1]. Traditional Intrusion Detection System (IDS) solutions, such as rule-based methods and ML models like autoencoders and SVMs, struggle to adapt to these evolving traffic patterns and often fail to detect unseen attacks [2]. Even deep learning-based IDSs suffer from concept drift, mode collapse, and performance degradation under imbalanced data, which raises the need for more advanced methods [3]. In practical IoT deployments such as smart manufacturing lines, home automation hubs, and sensor-driven industrial monitoring systems, edge devices operate under strict computational budgets. These platforms often lack GPUs, rely on small memory footprints, and require millisecond inference for real-time responses. As a result, lightweight generative intrusion detection models that maintain accuracy under tight resource constraints are essential for feasible on-device deployment.

Generative Adversarial Networks (GANs) offer a promising direction by modeling high-dimensional network traffic distributions, enabling the detection of both known and zero-day anomalies. GANs consist of a Generator and a Discriminator trained adversarially to synthesize and evaluate data realism [4]. However, conventional GANs often face issues such as instability, mode collapse, and data imbalance when applied to network security [3].

To address these limitations, we propose SD-CGAN, a Conditional GAN framework enhanced with Sinkhorn Divergence for stable and geometry-aware training. Conditional GANs incorporate a class-based conditioning mechanism to improve sample specificity and diversity [5]. In our model, we leverage Sinkhorn Divergence [6], [7], an optimal transport-based loss that improves training stability and convergence. Additionally, we employ CTGAN (Conditional Tabular GAN) for synthetic data augmentation to mitigate class imbalance.

The key contributions of this work include:

- Design SD-CGAN for anomaly detection in high-dimensional IoT edge traffic, capable of detecting DDoS and zero-day attacks.
- Implement CTGAN to generate realistic minority-class samples and address data imbalance in the CICDDoS2019 dataset.
- Evaluate SD-CGAN against existing GAN-based anomaly detection models used for network intrusion detection, using the CIC-DDoS2019 dataset.
- Optimize SD-CGAN using Sinkhorn Divergence to enhance training stability, reduce mode collapse, and improve computational efficiency for real-time IoT edge network security applications.

The rest of this paper is organized as follows: Section II reviews related work; Section III presents the methodology; Section IV discusses the experimental results; and Section V concludes the study.

II. RELATED WORKS

A. Anomaly Detection in IoT edge Networks

To enhance anomaly detection in IoT edge networks, various machine learning and deep learning approaches have been proposed [8]. Maimo et al. [9] introduced a self-adaptive learning system that adjusts its parameters based on traffic

fluctuations, reducing false alarms and improving detection to over 70%. Hussain et al. [10] used CNNs to analyze CDRs for DDoS detection in cyberphysical IoT networks, achieving over 90% accuracy. However, these supervised models depend heavily on labeled data, which is often scarce and quickly outdated. Illiyasu et al. [11] trained a GAN solely on benign IoT traffic to define a normal profile, flagging deviations as anomalies and achieving 81.3% detection. While unsupervised learning improves adaptability, such methods still face high false positive rates if the normal boundary is poorly captured.

B. GAN-Based Anomaly Detection

GANs have gained traction in IDS research due to their ability to support one-class modeling and synthetic data augmentation. Schlegl et al. [12] developed AnoGAN for unsupervised image anomaly detection, which inspired similar applications in network domains. Yao et al. [13] applied BiGAN with Wasserstein distance for IoT-based IDS. Novaes et al. [14] trained a GAN-based IDS on CICDDoS2019, outperforming traditional DL methods like CNN and LSTM. These models show promise in zero-day detection but are often limited by training instability and poor convergence under adversarial loss settings. It is well known that traditional GANs suffer from training instability issues such as vanishing gradients, oscillatory convergence and mode collapse, especially in high-dimensional space and under distribution imbalance [3], [7]. These issues arise from the minimax optimization between the generator and discriminator where reaching equilibrium becomes difficult as distributions diverge. When applied in intrusion detection, the instability issue is amplified due to sparse anomalies and multi-modal benign traffic. As a result, GAN-based IDS models often fail to learn a stable representation of benign flows, motivating the need for more geometry-aware loss functions such as Sinkhorn Divergence [3].

C. Conditional GANs for Imbalanced Data

Conditional GANs (cGANs) have been applied to address class imbalance by generating targeted synthetic data. Ullah et al. [15] implemented one-class, binary, and multi-class cGANs to enhance minority-class detection across seven IoT datasets, achieving 98% accuracy. Ezeme et al. [16] proposed AD-CGAN to synthesize anomalies and adapt the decision boundary. Benaddi et al. [17] introduced a cGAN-based framework to generate adversarial traffic that improves CNN-LSTM IDS performance on zero-day threats by 40%.

D. Sinkhorn Divergence and Loss Stabilization

To address instability in GAN training, researchers have explored Optimal Transport (OT)-based metrics. Arjovsky et al. [18] introduced WGAN, using Earth Mover’s Distance to mitigate mode collapse. Cuturi [19] proposed Sinkhorn Divergence, which adds entropy regularization for improved computational efficiency. Genevay et al. [20] confirmed its utility in generating stable gradients and capturing distribution mismatches. Xu et al. [21] introduced COT-GAN with mixed

Sinkhorn divergence to address batch-level convergence. Tien et al. [22] applied debiased Sinkhorn divergence to improve image anomaly localization. BEGAN [3], [23] further stabilized training by using autoencoder-based discriminators with equilibrium constraints.

E. Research Gap

While prior work demonstrates the potential of GANs and cGANs in anomaly detection and data augmentation, several critical gaps remain. Existing GAN-based IDS frameworks often suffer from (i) unstable convergence under adversarial loss dynamics, (ii) high computational cost from discriminator-guided training, and (iii) limited generalization under real-world class imbalance. Moreover, Optimal-Transport-based divergences and conditional generation have been explored separately, but their integration in a Sinkhorn-based discriminator, a geometry-aware one-class anomaly detection framework has not been reported. Our work addresses these gaps by proposing SD-CGAN: a Sinkhorn-enhanced cGAN trained solely on benign traffic to learn stable, geometry-preserving latent representations for one-class anomaly detection, supported by CTGAN-based data augmentation, and designed for lightweight deployment in IoT networks.

III. PROPOSED METHODOLOGY

A. Overview of SD-CGAN Framework

To address the instability, inefficiency, and poor generalization of traditional GAN-based IDSs, we propose SD-CGAN—a Conditional Generative Adversarial Network trained on benign network flows and optimized using Sinkhorn Divergence. This model is designed to learn the underlying distribution of normal traffic and detect deviations as anomalies, including zero-day attacks. Unlike standard GANs, SD-CGAN relies on a direct statistical comparison between real and generated samples via Sinkhorn Divergence. This stabilizes training, mitigates mode collapse, and allows for lightweight deployment in edge environments.

The SD-CGAN framework involves four major steps: (i) data preprocessing and feature engineering, (ii) synthetic data generation using CTGAN to address class imbalance, (iii) training the conditional generator using Sinkhorn Divergence, and (iv) anomaly scoring based on distributional deviation from benign traffic. Figure 1 illustrates the full system architecture.

B. Conditional GAN Architecture and Training Objectives

In cGANs, the Generator maps a latent vector $z \sim \mathcal{N}(0, I)$ and condition c to a sample x , denoted as:

$$G : (z, c) \rightarrow x \in \mathbb{R}^n \quad (1)$$

where G learns to produce class-conditioned samples. Here, c encodes simple flow-type metadata to guide context-aware benign sample generation. Its loss objective:

$$\mathcal{L}_G = -\mathbb{E}_{z \sim p(z), c \sim p(c)} [\log D(G(z, c) | c)] \quad (2)$$

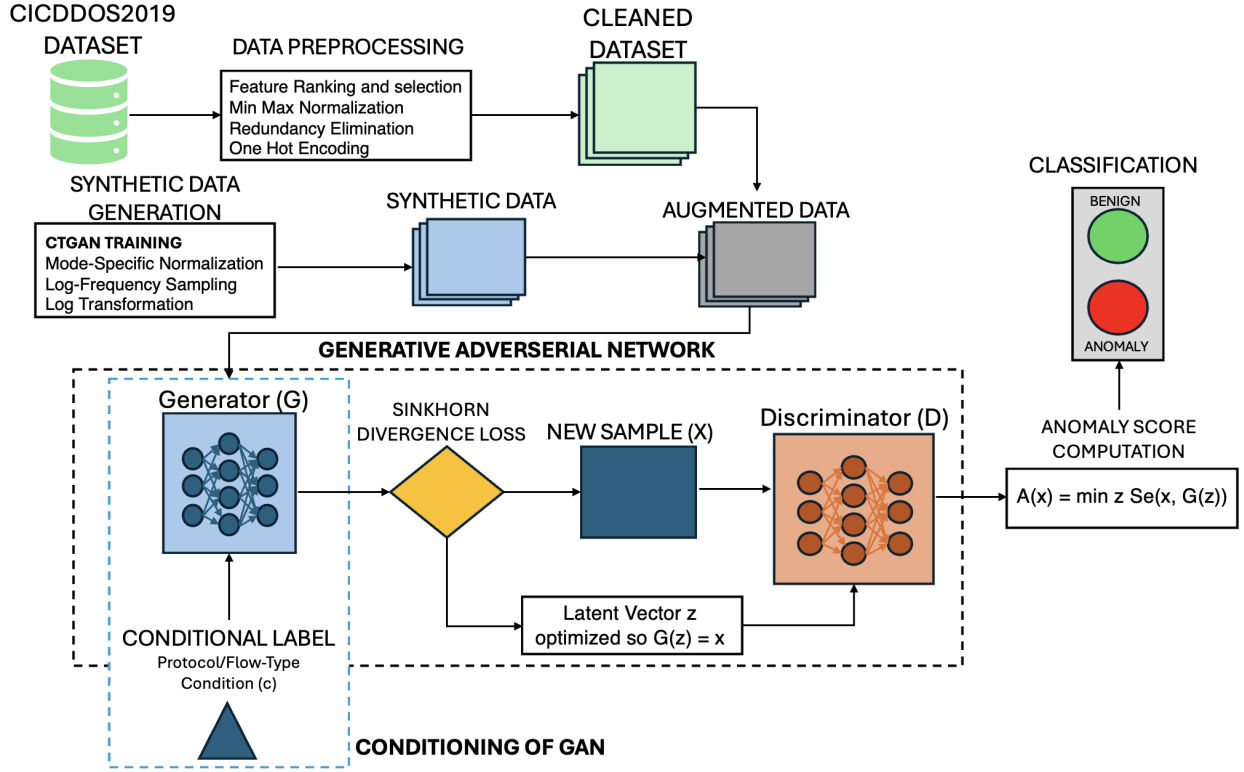


Fig. 1. SD-CGAN Model Architecture

encourages realistic synthesis while the Discriminator attempts to classify real vs. generated data:

$$\mathcal{L}_D = -\mathbb{E}_{x \sim p_{\text{data}}, c \sim p(c)}[\log D(x | c)] - \mathbb{E}_{z \sim p(z), c \sim p(c)}[\log(1 - D(G(z, c)))] \quad (3)$$

This adversarial game is optimized to approximate the true conditional distribution $p(x|c)$ [4], [5].

Conditional GANs are useful for anomaly intrusion detection environments because network traffic is highly imbalanced across classes, with minority attack behaviors exhibiting distinct statistical signatures. By conditioning generation on metadata such as flow type, the model avoids collapsing toward majority patterns and learns class-specific structure.

C. Sinkhorn Divergence

Traditional adversarial loss functions such as Jensen-Shannon and KL divergence are often unstable for training GANs, especially in high-dimensional spaces [24]. To address this, we adopt *Sinkhorn Divergence* [19], [20], an entropy-regularized optimal transport (OT) metric that stabilizes optimization and provides meaningful gradient signals even when data distributions have little overlap.

Given two empirical distributions μ and ν over a metric space, the entropic OT distance is defined as:

$$W_\epsilon(\mu, \nu) = \min_{\pi \in \Pi(\mu, \nu)} \sum_{x, y} \pi(x, y) c(x, y) + \epsilon \mathcal{H}(\pi) \quad (1)$$

where $\Pi(\mu, \nu)$ denotes the set of all couplings (joint distributions) between μ and ν , $c(x, y)$ is the ground cost function (typically Euclidean), and $\mathcal{H}(\pi) = -\sum_{x, y} \pi(x, y) \log \pi(x, y)$ is the entropy regularization term.

Since $W_\epsilon(\mu, \mu) \neq 0$, we use the debiased Sinkhorn divergence:

$$\mathcal{S}_\epsilon(\mu, \nu) = W_\epsilon(\mu, \nu) - \frac{1}{2} W_\epsilon(\mu, \mu) - \frac{1}{2} W_\epsilon(\nu, \nu) \quad (2)$$

In our SD-CGAN architecture, the Sinkhorn divergence replaces the adversarial loss and the Discriminator. The generator learns by directly minimizing the divergence between mini-batches of real and synthetic data, using the following training loss:

$$\mathcal{L}_{\text{Sink}}(\mu, \nu) = OT_\epsilon(\mu, \nu) - \frac{1}{2} [OT_\epsilon(\mu, \mu) + OT_\epsilon(\nu, \nu)] \quad (3)$$

where OT_ϵ is computed efficiently via the `geomloss` library, allowing scalable mini-batch optimization.

This loss function captures the geometry of the data distribution by measuring the "cost" of transporting mass from generated to real samples using a cost function such as $c(x, y) = \|f(x) - f(y)\|_2^2$, where $f(\cdot)$ maps samples into a learned feature space. As a result, Sinkhorn divergence offers several benefits over classical GAN losses: (a) stable and smooth gradient flow during training, (b) support for non-overlapping distributions, (c) enhanced mode coverage, and (d) suitability for unsupervised anomaly detection.

By combining conditional generation with Sinkhorn-based optimization, our model learns a robust representation of

benign traffic. During inference, any significant deviation from this distribution, indicated by high Sinkhorn divergence, is flagged as anomalous. This enables SD-CGAN to function effectively as a one-class anomaly detector [6], [7], while also ensuring training stability and reducing mode collapse.

D. Dataset

We utilize the CICDDoS2019 dataset [25], a flow-based dataset featuring benign and malicious traffic with over 80 statistical and time-based features. Our focus is narrowed to three high-impact, transport-layer exploitative attacks: UDP Flood, UDP-Lag, and SYN Flood, that emulate real-world threats in IoT edge networks due to their bursty, high-volume nature. The dataset’s IP-based traffic patterns and protocol-level behavior closely reflect conditions found in edge-deployed IoT systems, making it suitable for evaluating anomaly detection models in this domain.

E. Data Preprocessing

Preprocessing involved several stages: stratified sampling and schema verification, feature importance ranking via Random Forests [26], and Pearson correlation analysis for redundancy elimination. High-importance features with low multicollinearity were retained, followed by Min-Max normalization to [0,1] and one-hot encoding of categorical labels. This ensured consistent feature scaling and balance between model interpretability and training efficiency.

F. Synthetic Data Generation with CTGAN

To mitigate class imbalance, we used CTGAN [6], a generative model tailored for mixed-type tabular data. CTGAN applies mode-specific normalization for continuous features and log-frequency sampling for categorical conditioning. It was trained on the selected exploitative classes to generate realistic minority-class samples, which were merged with the original dataset to form a balanced training set. It is also particularly effective for IoT datasets because it models complex dependencies between mixed-type features such as port numbers, flag counts, durations and protocol fields. These attributes often exhibit long-tailed or multimodal behavior, and CTGAN’s mode-specific normalization prevents minority modes from being suppressed during training. Therefore, the augmented dataset more accurately reflects real traffic variability.

G. SD-CGAN Training and Anomaly Detection

SD-CGAN’s generator is a multi-layer perceptron trained solely on benign flows. At each iteration, a latent vector $z \sim \mathcal{N}(0, I)$ is mapped to a synthetic flow vector $G(z)$. The Sinkhorn loss is computed between batches of real and generated samples to update the generator. The generator is a 3-layer MLP with hidden dimensions [128, 256, 256] using ReLU activations and a linear output layer. Conditioning is applied by concatenating the latent vector z with the one-hot encoded flow-type vector c prior to the first layer. Training uses Adam 2×10^{-4} , default beta, and batch size 64. During

inference, an anomaly score is calculated by minimizing the Sinkhorn divergence between a test sample and its closest generated approximation. Scores above the 95th percentile (based on benign data) are flagged as anomalies. This enables one-class, unsupervised detection of zero-day behaviors without requiring a separate classifier. Algorithm 1 below gives an overview of the model.

SD-CGAN training alternates between generating synthetic flows and minimizing the Sinkhorn Divergence against real benign batches. In each iteration, a latent vector is sampled, concatenated with a conditional label, and passed through the generator to produce a synthetic sample. The Sinkhorn loss computes the transport cost between mini-batches of real and synthetic samples to provide smooth and informative gradients for updating generator parameters. Unlike adversarial GAN training, the OT-based optimization avoids saturation and reduces instability to enable robust representation of benign IoT traffic.

Algorithm 1: SD-CGAN Training with Sinkhorn Divergence

Input: Benign training set $X \in \mathbb{R}^{N \times d}$, latent dimension d_z , batch size B , learning rate η , Sinkhorn loss $\mathcal{L}_{\text{Sink}}$
Output: Trained generator G_θ
Initialize generator $G_\theta : \mathbb{R}^{d_z} \rightarrow \mathbb{R}^d$ with parameters θ
for $epoch = 1$ **to** E **do**
 for *each* mini-batch $X_i \subset X$ **do**
 Sample latent noise $Z_i \sim \mathcal{N}(0, I) \in \mathbb{R}^{B \times d_z}$
 Generate synthetic batch $\hat{X}_i = G_\theta(Z_i)$
 Compute Sinkhorn loss:
 $\mathcal{L}_{\text{Sink}} = \text{OT}_\varepsilon(X_i, \hat{X}_i) - \frac{1}{2}[\text{OT}_\varepsilon(X_i, X_i) + \text{OT}_\varepsilon(\hat{X}_i, \hat{X}_i)]$
 Update $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}_{\text{Sink}}$
return G_θ

IV. EVALUATION AND RESULTS

A. Experimental Setup and Computational Efficiency

All experiments were conducted on an Intel® Core™ i7-11700KF CPU (16 threads, 32GB RAM, Ubuntu 20.04) without GPU acceleration. The proposed SD-CGAN was trained solely on benign samples from the CICDDoS2019 exploitative attacks subset, with a latent dimension of 100, batch size of 64, Sinkhorn regularization = 0.05, 70/30 data split and learning rate $\eta = 0.0002$ over 10 epochs. The model has a training time of 4.6 seconds, Inference time of 10.9 seconds, and computes anomaly scores under 1 second per sample in batch mode, demonstrating edge-level feasibility. Although no embedded hardware benchmarks are included, the CPU-only training time (4.6s) and sub-second inference show that SD-CGAN is computationally light enough for edge-class devices. This aligns with recent findings showing that edge-class devices such as Raspberry Pi and Jetson Nano can only sustain real-time anomaly detection only when model

latency is kept within a tens-of-millisecond range and memory footprints remain under 50MB, as demonstrated in embedded AI evaluations. These studies further CPU-bound execution is typical in IoT deployment where devices lack GPUs and rely lightweight models, thereby supporting SD-CGAN’s practical relevance [27], [28]. Deep Learning classifier was also evaluated on the same environment, using the same dataset. This model had a training time of approximately 20 seconds and 19 seconds of inference time, further showcasing the efficiency of our SD-CGAN model. The baseline compared metrics are reported from their original publications and may involve different computational settings but are included for reference, evaluation and methodological context.

Prior GAN and hybrid models such as those by Novaes et al. [14] and Freitas et al. [29] required over 157 epochs of adversarial training without reporting total training time, and other hybrid approaches like DFNN-SAE-DCGAN involve deeper stacked modules that incur higher computational overhead. Similarly, the IDS by Freitas et al [29], which incorporates TCN and self-attention mechanisms, achieves strong accuracy but relies on GPU-accelerated infrastructure and LSTM-based layers with longer training cycles. Furthermore, training only on benign traffic enables robust zero-day detection, as deviations from the learned benign manifold are flagged as anomalies without requiring labeled attack data, hence supporting zero-day attack detection.

B. Evaluation Metrics

We evaluate SD-CGAN using four standard classification metrics: precision, recall, F1-score, and accuracy [30]. These metrics were computed using confusion matrix values obtained by comparing predicted labels against ground truth annotations.

C. Evaluation of Synthetic Data

To assess the quality of synthetic samples generated via CTGAN, we performed both qualitative distribution comparison and statistical testing.

1) Distribution Analysis: Kernel Density Estimation (KDE) plots for selected features confirmed strong overlap between real and synthetic samples, indicating CTGAN’s ability to capture multi-modal tabular distributions without mode collapse (Figure 2).

2) Kolmogorov–Smirnov Test: We applied the two-sample KS test to compare real vs. synthetic feature distributions. Table I shows that several features yielded low KS statistics, indicating high similarity. This supports the use of synthetic data for augmenting minority attack classes without introducing distributional bias [31]. Table 1 presents the KS test results across the top numerical features.

D. Experimental Results on CICDDoS2019

SD-CGAN achieved 98.6% precision, 98.5% recall, 98.5% F1-score, and 98.1% accuracy, outperforming baseline models including CNN, LSTM, RNN-GAN, FID-GAN, and LSTM-FUZZY [29], [32], [33]. Figure 3 is a graphical representation

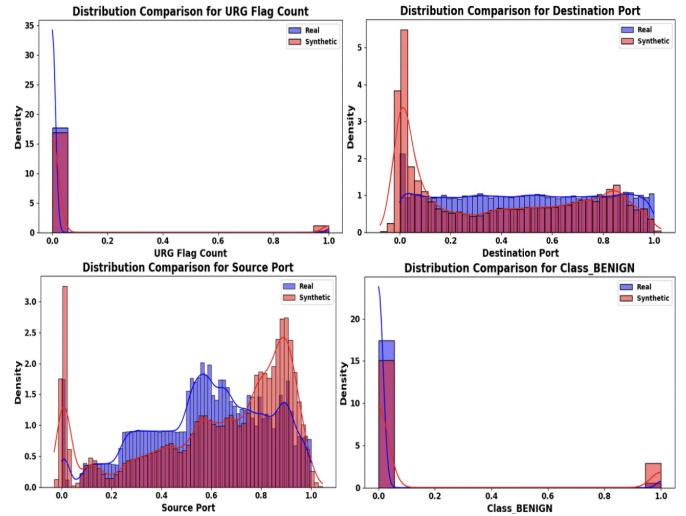


Fig. 2. Distribution Comparison Between Real and CTGAN-Generated Synthetic Samples for Selected Features

TABLE I
TOP 10 FEATURES BY KS STATISTIC (REAL VS SYNTHETIC)

Feature	KS Statistic	p-value
min_seg_size_forward	0.0082	0.0005
Class_WebDDoS	0.0072	0.0032
CWE Flag Count	0.0343	0.0000
URG Flag Count	0.0731	0.0000
Source Port	0.0985	0.0000
Class_BENIGN	0.1764	0.0000
Down/Up Ratio	0.1402	0.0000
Destination Port	0.2051	0.0000
Inbound	0.2077	0.0000
Class_UDP-lag	0.2435	0.0000

of the comparison. Figure 4 illustrates the ROC curve with an AUC of 0.9737, and Figure 5 shows the confusion matrix reflecting high classification accuracy.

TABLE II
PERFORMANCE COMPARISON WITH EXISTING MODELS ON CICDDoS2019

Model	Precision	Recall	F1 Score	Accuracy
LSTM-FUZZY [32]	97.89%	93.13%	95.45%	95.26%
GAN [14]	94.08%	97.89%	95.94%	94.38%
CNN [14]	96.32%	86.29%	91.02%	94.08%
LSTM [14]	90.12%	89.43%	89.77%	90.29%
RNN-GAN [33]	96.80%	97.20%	97.50%	97.90%
FID-GAN [29]	92.03%	92.03%	92.03%	92.03%
SD-CGAN (Proposed)	98.62%	99.36%	98.99%	98.03%

Also, we validated the original SD-CGAN model for zero-day attack detection by training on benign samples and testing against perturbed attack flows from the existing attack category in the dataset. The model achieved a 98.37% accuracy with near-perfect performance across all metrics, thereby confirming its ability to flag previously unseen threats as anomalies. An ablation comparison was conducted by training a GAN using the standard Jensen-Shannon divergence under the same environment, and SD-CGAN achieved a higher

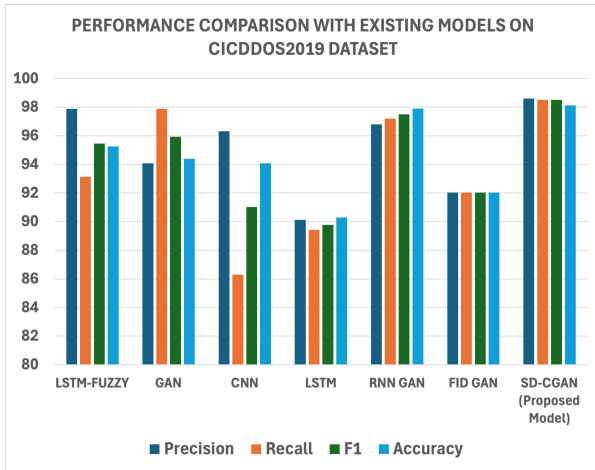


Fig. 3. Performance Comparison: SD-CGAN vs Existing Models




Fig. 4. ROC Curve of SD-CGAN (AUC = 0.9737)




Fig. 5. Confusion Matrix of SD-CGAN

accuracy (98.62%) compared to the GAN's (95.04%), which further confirms the contribution of Sinkhorn Divergence.

E. Training Stability and Convergence

To evaluate training dynamics, we monitored Sinkhorn loss across epochs. Unlike adversarial losses, the Sinkhorn objective exhibited a smooth decline from 0.59 to 0.21, with no oscillations or collapse. Traditional GANs are known to suffer from mode collapse, and this is modeled in Figure 6. A Deep Learning Classifier which we evaluated can be seen, in Figure 6, suffered from a steep and noisy loss trajectory. This validates geometry-aware divergence as a stabilizer and confirms convergence toward meaningful transport mappings, mitigating mode collapse.




Fig. 6. Comparison of Training loss of SD-CGAN against traditional GAN which suffers mode collapse

V. CONCLUSIONS AND FUTURE WORKS

This paper introduces SD-CGAN, a cGAN enhanced with Sinkhorn Divergence for anomaly detection in IoT edge networks. Trained exclusively on the CICDDoS2019 dataset with CTGAN-based augmentation, it effectively detects exploitative DDoS attacks under class imbalance. Experiments on the CICDDoS2019 dataset show that SD-CGAN surpasses existing models in accuracy, precision, recall, and F1-score, making it a strong candidate for real-time edge deployment. Mixed-attack evaluation was not included as this study focused on three high-impact transport-layer DDoS attacks to maintain a consistent and controlled experimental scope. Further works include extending SD-CGAN to multi-class scenarios in newer IoT-intensive datasets such as CICIoT2023. Generalizing SD-CGAN across additional IoT traffic datasets will also help evaluate robustness under diverse network scenarios. Explainable AI (XAI) will also be implemented in future works to interpret SD-CGAN's anomaly scoring and false negative behavior, provide feature-level explanations, as certain IoT traffic attributes may exhibit non-linear dependencies that influence optimal-transport-based modelling, and further justify zero-day generalization [34]. While SD-CGAN demonstrated strong performance under perturbed zero-day conditions, evaluating generalization to entirely new attack families is an important direction for future work.

ACKNOWLEDGMENT

This work is supported in part by the U.S. Department of Energy (DOE) under Award DE-NA0004189 and the National Science Foundation (NSF) under Award numbers 2409093 & 2219658.

REFERENCES

- [1] Y. Li, Z. Zhou, X. Xue, D. Zhao, and P. C. K. Hung, "Accurate anomaly detection with energy efficiency in iot-edge-cloud collaborative networks," *IEEE Internet of Things Journal*, vol. 10, no. 19, pp. 16959–16974, 2023.
- [2] M. A. Talukder, M. M. Islam, M. A. Uddin, K. F. Hasan, S. Sharmin, S. A. Alyami, and M. A. Moni, "Machine learning-based network intrusion detection for big and imbalanced data using oversampling, stacking feature embedding and feature extraction," *Journal of big data*, vol. 11, no. 1, p. 33, 2024.
- [3] C. Park, J. Lee, Y. Kim, J.-G. Park, H. Kim, and D. Hong, "An enhanced ai-based network intrusion detection system using generative adversarial networks," *IEEE Internet of Things Journal*, vol. 10, no. 3, pp. 2330–2345, 2022.
- [4] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [5] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [6] L. Xu, M. Skoularidou, A. Cuesta-Infante, and K. Veeramachaneni, "Modeling tabular data using conditional gan," *Advances in neural information processing systems*, vol. 32, 2019.
- [7] Z. Chen, J. Duan, L. Kang, and G. Qiu, "Supervised anomaly detection via conditional generative adversarial network and ensemble active learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 6, pp. 7781–7798, 2022.
- [8] E. J. Samson, K. Hasan, L. Hong, I. Ahmed, H. Onyeka, and S. Shetty, "Optimizing 5g network slices: Lstm and game theory synergy," in *2025 International Conference on Computing, Networking and Communications (ICNC)*, pp. 282–287, 2025.
- [9] L. F. Maimó, A. L. P. Gómez, F. J. G. Clemente, M. G. Pérez, and G. M. Pérez, "A self-adaptive deep learning-based system for anomaly detection in 5g networks," *Ieee Access*, vol. 6, pp. 7700–7712, 2018.
- [10] B. Hussain, Q. Du, B. Sun, and Z. Han, "Deep learning-based ddos-attack detection for cyber-physical system over 5g network," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 2, pp. 860–870, 2020.
- [11] A. S. Iliyasu and H. Deng, "N-gan: a novel anomaly-based network intrusion detection with generative adversarial networks," *International Journal of Information Technology*, vol. 14, no. 7, pp. 3365–3375, 2022.
- [12] T. Schlegl, P. Seeböck, S. M. Waldstein, G. Langs, and U. Schmidt-Erfurth, "f-anogan: Fast unsupervised anomaly detection with generative adversarial networks," *Medical image analysis*, vol. 54, pp. 30–44, 2019.
- [13] W. Yao, H. Shi, and H. Zhao, "Scalable anomaly-based intrusion detection for secure internet of things using generative adversarial networks in fog environment," *Journal of Network and Computer Applications*, vol. 214, p. 103622, 2023.
- [14] M. P. Novaes, L. F. Carvalho, J. Lloret, and M. L. Proença Jr, "Adversarial deep learning approach detection and defense against ddos attacks in sdn environments," *Future Generation Computer Systems*, vol. 125, pp. 156–167, 2021.
- [15] I. Ullah and Q. H. Mahmoud, "A framework for anomaly detection in iot networks using conditional generative adversarial networks," *IEEE Access*, vol. 9, pp. 165907–165931, 2021.
- [16] O. M. Ezeme, Q. H. Mahmoud, and A. Azim, "Design and development of ad-cgan: Conditional generative adversarial networks for anomaly detection," *IEEE Access*, vol. 8, pp. 177667–177681, 2020.
- [17] H. Benaddi, M. Jouhari, K. Ibrahim, A. Benslimane, and E. M. Amhoud, "Adversarial attacks against iot networks using conditional gan based learning," in *GLOBECOM 2022-2022 IEEE Global Communications Conference*, pp. 2788–2793, IEEE, 2022.
- [18] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *International conference on machine learning*, pp. 214–223, PMLR, 2017.
- [19] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," *Advances in neural information processing systems*, vol. 26, 2013.
- [20] A. Genevay, G. Peyré, and M. Cuturi, "Learning generative models with sinkhorn divergences," in *International Conference on Artificial Intelligence and Statistics*, pp. 1608–1617, PMLR, 2018.
- [21] T. Xu, L. K. Wenliang, M. Munn, and B. Acciaio, "Cot-gan: Generating sequential data via causal optimal transport," *Advances in neural information processing systems*, vol. 33, pp. 8798–8809, 2020.
- [22] T. D. Tien, A. T. Nguyen, N. H. Tran, T. D. Huy, S. Duong, C. D. T. Nguyen, and S. Q. Truong, "Revisiting reverse distillation for anomaly detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 24511–24520, 2023.
- [23] D. Berthelot, T. Schumm, and L. Metz, "Began: Boundary equilibrium generative adversarial networks," *arXiv preprint arXiv:1703.10717*, 2017.
- [24] Z. Pan, W. Yu, B. Wang, H. Xie, V. S. Sheng, J. Lei, and S. Kwong, "Loss functions of generative adversarial networks (gans): Opportunities and challenges," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 4, no. 4, pp. 500–522, 2020.
- [25] I. Sharafaldin, A. H. Lashkari, S. Hakak, and A. A. Ghorbani, "Developing realistic distributed denial of service (ddos) attack dataset and taxonomy," in *2019 international carmaham conference on security technology (ICCSST)*, pp. 1–8, IEEE, 2019.
- [26] Y. Zhou, G. Cheng, S. Jiang, and M. Dai, "Building an efficient intrusion detection system based on feature selection and ensemble classifier," *Computer networks*, vol. 174, p. 107247, 2020.
- [27] A. Imteaj, U. Thakker, S. Wang, J. Li, and M. H. Amini, "A survey on federated learning for resource-constrained iot devices," *IEEE Internet of Things Journal*, vol. 9, no. 1, pp. 1–24, 2021.
- [28] M. J. C. S. Reis and C. Seródio, "Edge ai for real-time anomaly detection in smart homes," *Future Internet*, vol. 17, no. 4, 2025.
- [29] P. F. de Araujo-Filho, M. Naili, G. Kaddoum, E. T. Fapi, and Z. Zhu, "Unsupervised gan-based intrusion detection system using temporal convolutional networks and self-attention," *IEEE Transactions on Network and Service Management*, vol. 20, no. 4, pp. 4951–4963, 2023.
- [30] C. Goutte and E. Gaussier, "A probabilistic interpretation of precision, recall and f-score, with implication for evaluation," in *European conference on information retrieval*, pp. 345–359, Springer, 2005.
- [31] G. Fasano and A. Franceschini, "A multidimensional version of the kolmogorov-smirnov test," *Monthly Notices of the Royal Astronomical Society*, vol. 225, no. 1, pp. 155–170, 1987.
- [32] M. P. Novaes, L. F. Carvalho, J. Lloret, and M. L. Proença, "Long short-term memory and fuzzy logic for anomaly detection and mitigation in software-defined network environment," *Ieee Access*, vol. 8, pp. 83765–83781, 2020.
- [33] C. M. V. S. Akana, A. Kumar, M. Tiwari, A. Z. Yunus, M. Singh, *et al.*, "An optimized ddos attack detection using deep convolutional generative adversarial networks," in *2023 5th International Conference on Inventive Research in Computing Applications (ICIRCA)*, pp. 668–673, IEEE, 2023.
- [34] A. Amin, K. Hasan, S. Zein-Sabatto, D. Chimba, I. Ahmed, and T. Islam, "An explainable ai framework for artificial intelligence of medical things," in *2023 IEEE Globecom Workshops (GC Wkshps)*, pp. 2097–2102, 2023.