Differentially Private Synthetic Adversarial Network Traffic Generation Based on Tabular Diffusion Processes

Minjae Kang[1], Gunhee Cho[2], Sungju Yun[1], Yeonjoon Lee[3]
[1]*Major in Bio-Artificial Intelligence, Department of Computer Science and Engineering*
[2]*Department of Intelligent Information Convergence Engineering*
[3]*College of Computing*
Hanyang University
Ansan 15588, Republic of Korea
minjae0110@hanyang.ac.kr, no1gun2@hanyang.ac.kr, tjdwn77777@hanyang.ac.kr, yeonjoonlee@hanyang.ac.kr

***Abstract*—IDS is becoming increasingly important as networks become more vulnerable to a variety of attacks. However, traditional ML and DL-based IDS have limited accuracy and false positive rates due to the lack of open adversarial network datasets and vendors' concerns regarding potential data leakage during the training process. To address these challenges, VAE and GAN-based network packet generation models have emerged, but they still face issues with low data fidelity. In this context, TabDDPM is a suitable option for generating network packets, as it surpasses VAE and GAN in accurately capturing data distribution characteristics and effectively learning the inherent structure of network packet formats. Despite its potential, research applying TabDDPM to network packet generation has not yet been explored.**

In this paper, we propose and evaluate DP-NetDDPM, a novel framework that enables adversarial network traffic generation while preserving both fidelity and privacy guarantees. Specifically, we examine the use of diffusion to generate adversarial network packet datasets and compare its performance with baseline models. Our framework focuses on three key criteria: data fidelity, adaptability for machine learning applications, and data privacy. The results show that DP-NetDDPM surpasses traditional models in both fidelity and adaptability, achieving a notable 72% improvement in fidelity and a 45% enhancement in adaptability over baselines.

***Index Terms*—Adversarial Network Traffic Generation, Data Synthesis, Tabular Diffusion, Differential Privacy.**

I. Introduction

With the increasing complexity and scale of networks, the frequency of unauthorized access and malicious activities has risen, emphasizing the critical need for robust Intrusion Detection Systems (IDS). Traditional IDS techniques—such as signature-based, rule-based, and heuristic-based approaches that depend on predefined detection conditions—are widely implemented to monitor and classify various types of malicious traffic and attacks [1]. However, these conventional methods face significant challenges in adapting to the growing sophistication of contemporary network threats, particularly those that exploit new vulnerabilities or evade traditional detection mechanisms. While machine learning (ML) and deep learning (DL) models have been introduced to improve IDS performance, their efficacy has been limited by concerns over data leakage from vendors and the lack of openly available adversarial network datasets [2].

To address these limitations, network packet synthesis methods have been developed, leveraging generative models such as Tabular Variational Autoencoders (TVAE) [3] and Tabular Generative Adversarial Networks (TGAN) [4], [5]. The adoption of tabular data formats for network traffic synthesis offers several key advantages: it enables efficient reconstruction of network flows, preserves complete flow header information, and provides computational efficiency compared to image-based approaches. These methods allow for the creation of new datasets that preserve the statistical properties of the original data. However, both Variational Autoencoders (VAE) and Generative Adversarial Networks (GAN) exhibit inherent limitations, including mode collapse, poor locality, and challenges in handling sparse data, such as missing values. These issues hinder their effectiveness in learning complex data distributions, resulting in synthesized network traffic that fails to achieve classification accuracy comparable to real-world datasets.

To overcome the limitations of traditional generative models, diffusion models have recently gained attention, with the Tabular Diffusion Probabilistic Model (TabDDPM) emerging as a leading example [6]. Diffusion models have demonstrated superior performance in learning data distributions, effectively preserving the fidelity of raw data. Despite their potential, research into the use of TabDDPM for network packet generation remains scarce, presenting a gap in our understanding of its effectiveness in network packet generation applications.

In this paper, we suggest end-to-end framework,

DP-NetDDPM. Our framework demonstrates several key capabilities in network data generation: (1) In DP-NetDDPM, we provide both DP-enabled and standard models, allowing users to choose based on their specific requirements. By offering this flexibility, our framework accommodates diverse use cases in network security research and applications, from privacy-sensitive scenarios requiring robust data protection to cases where maximum fidelity is paramount for accurate analysis and testing. (2) DP-NetDDPM effectively synthesizes both packet and flow-level network traffic data, capturing the intricate characteristics of header traces and payload information through entropy values while preserving the complex relationships between numerical and categorical features via its diffusion process architecture. (3) DP-NetDDPM demonstrated a 45% superior accuracy compared to baseline models in downstream tasks, showcasing its potential for expansion into additional downstream applications. (4) By incorporating of DP-SGD, we ensure privacy preservation while maintaining high-quality synthetic data generation. Our experimental results demonstrate that this privacy-preserving approach maintains competitive performance in both packet and flow-level analysis. Through this comprehensive approach, we address the current shortage of IDS datasets and propose a method to enhance the reliability of research in the network security domain while maintaining privacy guarantees.

II. Related works

A. Diffusion Models for Tabular Data Generation

Diffusion models, which gained initial recognition through success in image generation, have emerged as a promising approach for tabular data synthesis. These models operate by adding and reversing noise systematically during training, demonstrating remarkable versatility across various domains [7]. The introduction of DDPM by Ho et al [8]. marked a significant advancement, establishing superior stability and performance compared to traditional generative approaches.

In the specific domain of tabular data synthesis, multiple approaches have been developed to address the challenges of limited dataset sizes. Traditional methods like TVAE [3] have shown success through feature embeddings and conditional dependencies, effectively preserving statistical properties and inter-feature relationships [9]. GAN-based approaches, particularly CTGAN and TGAN, have contributed by addressing challenges related to mixed data types and imbalanced distributions [3], [10].

The latest advancement in this field is represented by specialized diffusion-based models. TabDDPM [6] pioneered the adaptation of diffusion processes for tabular data, introducing carefully designed noise schedules and denoising steps to preserve feature relationships. This approach has shown particular promise in maintaining statistical integrity while offering improved stability. Recent developments, including multinomial diffusion models [11], have further enhanced the handling of categorical variables within tabular structures, while ongoing research focuses on optimizing these models for better adaptability and efficiency [12].

B. Adversarial Traffic Generation

The challenge of collecting real malicious traffic data, constrained by ethical and legal considerations, has led to increased research focus on synthetic malicious traffic data generation. In this domain, various approaches have been developed to address these limitations.

Gulrajani et al. introduced Wasserstein GAN with Gradient Penalty (WGAN-GP), addressing the persistent challenges of instability and mode collapse in traditional GANs [13]. Building on this foundation, Zolbayar et al. proposed NIDSGAN [14], demonstrating enhanced capabilities in generating adversarial network traffic specifically designed to challenge ML-based IDS systems. Zilong et al. utilized low-dimensional latent representations to enhance data diversity and quality [15].

Recent research has explored network traffic image synthesis using diffusion models [16]. However, these approaches are limited by their focus on generating image representations rather than tabular data. While image synthesis can capture certain visual patterns, it presents challenges in reconstructing realistic network datasets since the original tabular structure and relationships between features may be lost in the image-to-data conversion process.

III. Methdology

In this section, we present DP-NetDDPM based on TabDDPM [6], a diffusion-based adversarial network traffic synthetic generation model. To overcome the limitations of existing models like GANs, such as mode collapse and difficulty in capturing local distributions, this model employs a diffusion process. Additionally, it uses feature vectors of network packet header traces as tabular input in table form. In tabular data, it's important to thoroughly understand the distributions of heterogeneous features to uncover patterns inherent to each column. At this point, the ability to process categorical data through TabDDPM's multinomial diffusion process is essential.

The proposed framework in Fig. 1 consists of three main phases: (1) data pre-processing, (2) training and synthesis phases and (3) post-processing.

In the pre-processing phase (1), we begin with the adversarial dataset in raw PCAP format. To handle raw packet data without predefined labels, we implement an automated labeling mechanism based on the dataset's metadata, including labeled flow data provided by the authors [17], timestamps, and attack

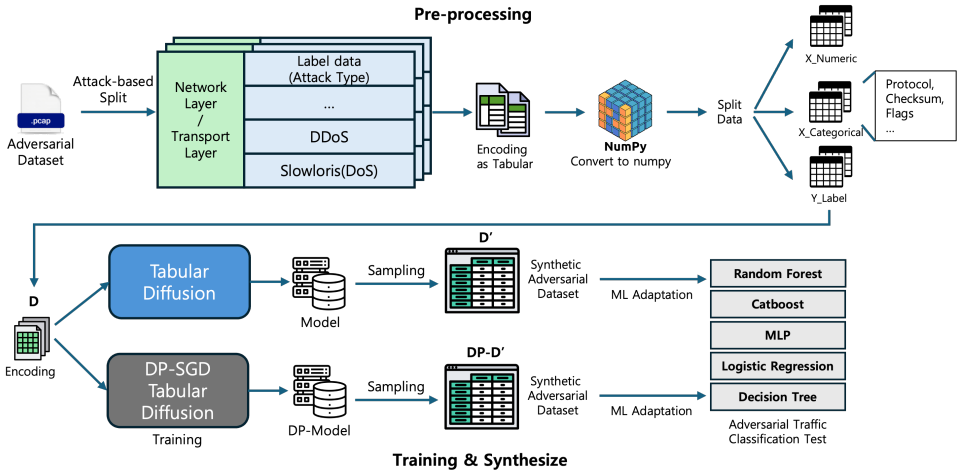

Fig. 1: Overview of DP-NetDDPM framework.

IP addresses, to classify distinct network attacks. Additionally, we extract both network and transport layers header and payload features from the raw packets focusing on header information that characterizes various types of network behaviors and attack patterns. The data is then converted into a tabular format, preserving essential packet characteristics.

After the converting process, the data is converted to *NumPy* arrays to facilitate efficient processing and manipulation. The final pre-processing step involves splitting the data into three distinct components: numeric features *(x_{num})*, categorical features *(x_{cat})* containing protocol information, checksum, and flags, and labels *(y_{label})* representing attack classifications.

Numerical Features (x_{num}) are processed through a Quantile Transformer, which ensures that the numerical value N_{num} are normalized into a uniform distribution $x_{num} \in R^{N_{num}}$. Categorical Features (x_{cat}) are handled through a One-hot Encoder, which converts K_i categorical features with C numbers of x_{cat_i} and $x_{cat_i}^{ohe} \in \{0,1\}^{K_i}$ into a numeric vector, where each bit represents a category.

The processed data x_{in}'s dimension is $(N_{num} + \sum K_i)$. The reverse diffusion step is modeled by a multi-layer neural network that has an output of the same dimensionality as x_0, where the first N_{num} coordinates are the predictions of ϵ for the Gaussian diffusion and the rest are the predictions of $x_{cat_i}^{ohe}$ for the multinomial diffusions.

After processing, the system computes a residual term e for the categorical component, to refine the categorical output. The Softmax layer ensures that the final output remains a valid probability distribution over the categories. The entire system in Fig. 2 thus processes both continuous and categorical data in an integrated pipeline.

The training and synthesis phase (2) employs two parallel approaches: standard Tabular Diffusion and DP-SGD (Differentially Private Stochastic Gradient Descent) Tabular Diffusion. The standard tabular diffusion model prioritizes maintaining high-fidelity data distributions and complex feature relationships, making it suitable for scenarios where maximum accuracy in synthetic data generation is required.

In contrast, the DP-SGD variant incorporates privacy-preserving mechanisms during the training process, offering formal privacy guarantees at the cost of potentially reduced distribution accuracy. In DP-NetDDPM, the parameter update as: $\theta_{t+1} = \theta_t - \eta \nabla L(\theta_t)$ where θ_t represents the model parameters at step t, η is the learning rate, and $\nabla L(\theta_t)$ is the gradient of the loss function. DP-SGD modifies this process with three key steps:

- **Per-example gradient computation:** For each example i in the batch, compute individual gradients: $g_i = \nabla L(\theta_t, x_i)$
- **Gradient clipping:** Clip the gradient norm to a threshold C: $\bar{g}_i = g_i \cdot \min(1, \frac{C}{||g_i||_2})$
- **Noise addition:** Add Gaussian noise calibrated to the clipping threshold: $\tilde{g} = \frac{1}{B}(\sum_{i=1}^{B} \bar{g}_i + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I}))$

The final update rule becomes: $\theta_{t+1} = \theta_t - \eta \tilde{g}$ where B is the batch size, C is the clipping threshold, and δ is the noise multiplier. The privacy guarantee (ϵ) is calculated based on these parameters and the number of training steps, using the moments accountant method.

Although DP-SGD presents challenges in terms of computational costs and reduced accuracy and reproducibility, these limitations can be effectively addressed by leveraging the data reproduction capabilities in the diffusion process. This advancement enables us to mitigate the traditional drawbacks of privacy-preserving approaches while maintaining high-quality synthetic data generation.

Our framework allows users to choose between DP and standard models based on their specific requirements. When privacy is main consideration, users can opt for the DP model. Conversely, when high-fidelity data reproduction is the primary concern, users can

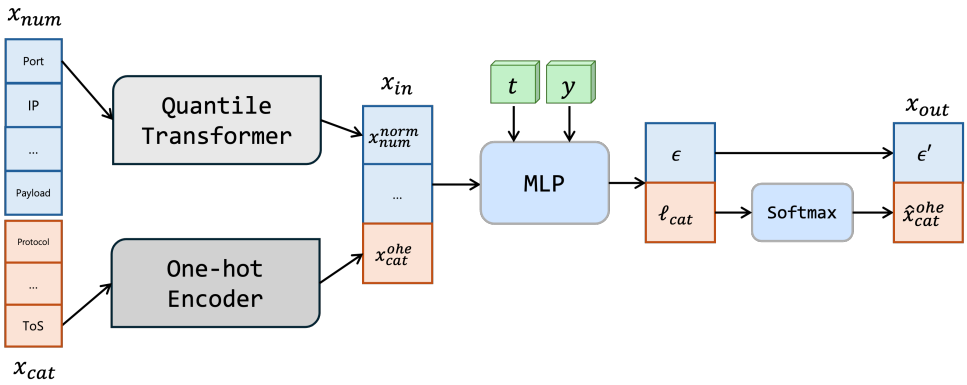

Fig. 2: Encoding process on DP-NetDDPM.

select the standard model. This dual-track approach ensures that the framework can accommodate varying priorities in synthetic network traffic generation while maintaining appropriate balance between privacy preservation and data utility.

After generating all synthetic data, we restore the label-encoded data its original representation. The data is converted into CSV format and then reconstructed into its original PCAP or flow data format. This process enables the creation of new datasets that maintain the structural integrity of network traffic data. These reconstructed synthetic datasets are evaluated using a comprehensive methodology to assess their quality and utility for network security applications.

IV. Evaluation

In this section, we provide a overview of our evaluation methods.

TABLE I: Overview of datasets features

Packet Header Fields (16)	**Flow Header Fields (13)**
1. Source port number	1. Destination port
2. Destination port number	2. Protocol
3. Time To Live (TTL)	3. Flow duration
4. Packet size	4. Total forward packets
5. Protocol	5. Total backward packets
6. Payload size	6. Total length of forward packets
7. Payload Entropy	7. Total length of backward packets
8. Type of service	8. Forward packet length maximum
9. Total length	9. Forward packet length mean
10. Identification	10. Flow bytes per second
11. Flags	11. Flow packets per second
12. Fragment offset	12. Packet length standard deviation
13. Header checksum	13. Forward IAT mean
14. Label (benign/attack)	
15. Source IP address (numeric)	
16. Destination IP address (numeric)	
Dataset Summary	
Dataset	**Labels**
CICIDS-2018(Flow)	benign, bot, brute force, DDoS, DoS, infiltration
CICIDS-2017(Packet)	benign, Bot, DDoS, DoS (GoldenEye, Hulk, Slowhttptest, slowloris), FTP-Patator, PortScan

A. Evaluation Setup

Datasets. To systematically evaluate the performance of tabular data generation models, we select two public datasets (one flow-based and one packet-based dataset). For packet header datasets, we consider 16 fields in the packet records. The field extraction process was implemented according to the protocol hierarchy of the Network layer (IP) and Transport layer (TCP, UDP). For flow header datasets, we analyze 13 key fields in the flow records, each labeled as either benign traffic or a specific type of attack. For flow-based header extraction, we selected a subset of universally applicable features to prevent model overfitting on dataset-specific characteristics.

We evaluate DP-NetDDPM and baseline models on datasets with 300,000 samples each, split into training, validation, and test sets in a ratio of 8:1:1. The dataset labels include four types of benign, bot, brute force, DDoS, DoS, Infilteration for *CICIDS-2018* and benign, Bot, DDoS, DoS GoldenEye, DoS Hulk, DoS Slowhttptest, DoS slowloris, FTP-Patator, PortScan for *CICIDS-2017*. The information about the datasets we used is in the Table I.

- **CICIDS-2017(Packet traces) [17]:** The CICIDS2017 dataset captures full packet payloads from a complete network infrastructure with diverse attack scenarios, making it suitable for intrusion detection research.
- **CICIDS-2018(Flow traces) [17]:** The CSE-CIC-IDS2018 dataset was created to provide comprehensive network traffic data, featuring seven attack scenarios executed across an infrastructure containing 50 attacking machines and 420 victim machines.

Baselines. Given the extensive variety of generative models currently proposed for tabular data generation, this study focuses on comparing state-of-the-art model such as GAN and VAE models.

- **TVAE** [3]: VAE model designed specifically for tabular data, employs Gumbel-Softmax techniques to handle both categorical and numerical features effectively. It maintains its position as the leading VAE-based approach due to its robust performance and public availability [3].
- **CTABGAN** [10]: CTABGAN addresses fundamental challenges in tabular data generation through its GAN-based architecture, particularly focusing on class imbalance and mixed data

types using Conditional GAN structures. While the model demonstrates superior performance in classification tasks, it shows limitations when handling regression applications that require continuous outputs [5].

- **CTABGAN+** [10]: CTABGAN+ introduces advanced features including dynamic reweighting and differentiable augmentation to better handle rare categories and high-dimensional data [10]. CTABGAN+ is widely used in industries such as finance and healthcare for data generation and privacy preservation.

B. Experiment results

Data Distribution. Through t-SNE visualizations, we analyze high-dimensional data in lower-dimensional space (2D and 3D) to evaluate the quality of synthetic data generated by various models. Fig. 3 presents t-SNE comparisons between real data and synthetic data generated by CTABGAN, CTABGAN-PLUS, DDPM, and TVAE.

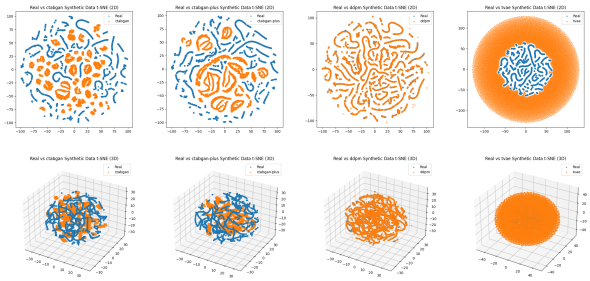

Fig. 3: 2D and 3D data distribution based on t-SNE.

The visualizations reveal that CTABGAN and CTABGAN-PLUS exhibit noticeable distribution differences from the real data, suggesting potential mode collapse issues. In contrast, NetDDPM demonstrates superior performance with closer distribution alignment in both 2D and 3D representations, suggesting higher fidelity in synthetic data generation. TVAE exhibits intermediate performance but still maintains distinct separations between real and synthetic distributions. These results highlight DDPM's superior capability in capturing and synthesizing the true data distribution.

Fidelity. In fidelity evaluation metrics, we employ two complementary metrics to assess the fidelity of synthetic data generation. *Jensen-Shannon Divergence (JSD)* is used for categorical data, measuring distributional similarities with values approaching 0 indicating higher similarity. For continuous data, we utilize *Earth Mover's Distance (EMD)*, which quantifies the minimal cost of transforming one distribution into another.

This dual-metric approach addresses the inherent challenges in evaluating network traffic traces, where continuous fields exhibit varying ranges that complicate traditional metric applications. By utilizing JSD

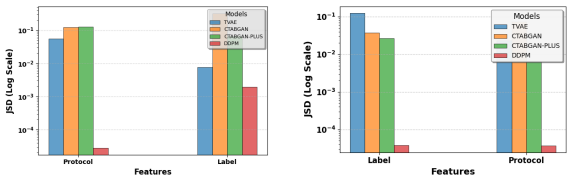

(a) CICIDS 2017(Packet) JSD (b) CICIDS 2018(Flow) JSD

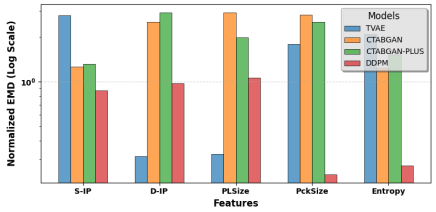

(c) CICIDS 2017(Packet) EMD

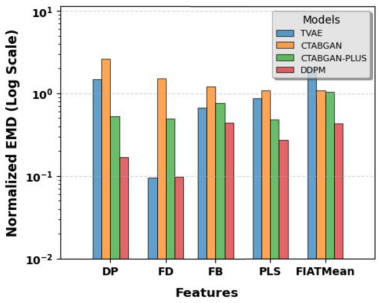

(d) CICIDS 2018(Flow) EMD

Fig. 4: JSD(↓) and normalized EMD(↓) between real and synthetic distributions on CICIDS 2017-2018

for categorical features and EMD for numerical data, we ensure appropriate evaluation across different data types while maintaining measurement accuracy.

Overall in Fig. 4, our analysis demonstrates that NetDDPM achieves 77% better performance in maintaining distribution fidelity across both packet and flow-level metrics. Fig. 4 provides a detailed quantitative comparison of different models across CICIDS 2017 and 2018 datasets, showcasing NetNetDDPM's consistent advantages over baseline approaches. We observe that while baseline models occasionally perform well on specific metrics, NetDDPM maintains consistent performance across all distribution metrics.

For categorical features, NetDDPM demonstrates significantly lower JSD values for both label and protocol distributions, particularly notable in label distributions, where NetDDPM achieves approximately an order of magnitude better performance compared to baselines.

In terms of continuous features, NetDDPM consistently maintains lower normalized EMD values across various metrics. For packet-level analysis, NetDDPM shows superior performance in preserving critical features such as packet size (PckSize) and payload size (PLSize), with normalized EMD values consistently below 1.0. Similarly, in flow-level analysis, NetDDPM

TABLE II: Performance Comparison of ML-Adaptation on CICIDS Datasets

ML-Adapation	Model	Packet (CICIDS 2017)		Flow (CICIDS 2018)	
		Accuracy	F1-Score	Accuracy	F1-Score
TVAE	Catboost	0.8564	0.306	0.6000	0.1071
	MLP	0.6001	0.1111	0.5991	0.1091
	Decision Tree	0.6106	0.1203	0.5970	0.1089
	LR	0.6000	0.1111	0.5881	0.1060
	Random Forest	0.6008	0.1103	0.6000	0.1090
	Average	0.6536	0.1518	0.5968	0.1080
CTABGAN	Catboost	0.9262	0.5304	0.6079	0.2058
	MLP	0.9199	0.5307	0.6492	0.3177
	Decision Tree	0.7933	0.4053	0.5217	0.1684
	LR	0.8607	0.3724	0.6017	0.1369
	Random Forest	0.8934	0.5335	0.6033	0.2045
	Average	0.8787	0.4745	0.5968	0.2067
CTABGAN+	Catboost	0.9058	0.4432	0.7877	0.4780
	MLP	0.9017	0.4468	**0.7916**	**0.5102**
	Decision Tree	0.5924	0.2158	0.7464	0.5526
	LR	0.7680	0.3008	**0.6780**	**0.3382**
	Random Forest	0.7101	0.2542	0.8065	0.5462
	Average	0.7756	0.3322	0.7620	**0.4850**
NetDDPM	Catboost	**0.9641**	**0.7463**	**0.8408**	**0.5857**
	MLP	**0.9694**	**0.8171**	0.7678	0.4556
	Decision Tree	**0.9958**	**0.9828**	**0.7869**	**0.5878**
	LR	**0.9066**	**0.4022**	0.6203	0.1788
	Random Forest	**0.9857**	**0.9271**	**0.8395**	**0.5981**
	Average	**0.9643**	**0.7751**	**0.7711**	0.4812
Real Data	Catboost	0.8755	0.6531	0.8753	0.6534
	MLP	0.9892	0.9532	0.8278	0.5792
	Decision Tree	1.0000	1.0000	0.8663	0.7437
	LR	0.9210	0.4982	0.6505	0.2660
	Random Forest	1.0000	1.0000	0.8786	0.7462
	Average	0.9571	0.8209	0.8197	0.5977

maintains better fidelity across destination port (DP), flow duration (FD), and flow byte (FB) distributions, as evidenced by lower EMD values in logarithmic scale.

ML adaptability. Here, we measure how effectively synthetic data can be utilized for downstream attack classification tasks by comparing the performance of various ML models trained on synthetic versus real data. Our analysis demonstrates the practical utility of synthetic data generation approaches through comprehensive accuracy and F1-score measurements.

The experimental results show that NetDDPM significantly outperforms other generative models across both packet and flow-level analysis. For packet-level data (CICIDS 2017), NetDDPM achieved an average accuracy of 0.9643 and F1-score of 0.7751, substantially exceeding baseline models. Particularly notable is NetDDPM's performance with Decision Tree and Random Forest classifiers, achieving accuracy rates of 0.9958 and 0.9857 respectively, closely approaching the performance metrics of real data.

For flow-level analysis (CICIDS 2018), NetDDPM maintained its superior performance with an average accuracy of 0.7711 and F1-score of 0.4812, demonstrating consistent effectiveness across different data types.

These results indicate that NetDDPM generates high-quality synthetic data that closely resembles real network traffic patterns, enabling more effective training of machine learning models for attack detection. The consistent performance across different classifier types and datasets suggests that NetDDPM-generated data maintains essential data characteristics and relationships.

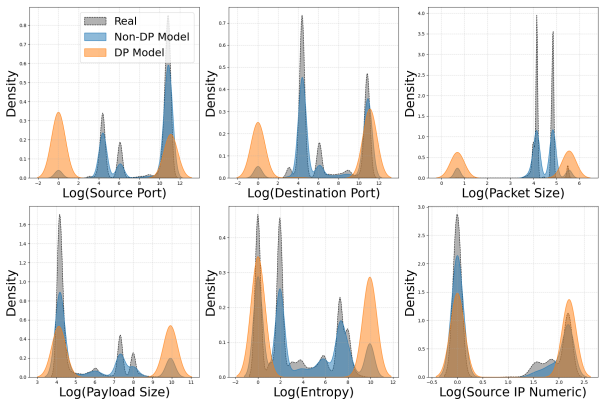

Fig. 5: Density distribution on real and synthetic data.

TABLE III: Privacy results on DP-NetDDPM

Privacy Metric	DP-NetDDPM		NetDDPM	
	DCR	Epsilon	DCR	Epsilon
CICIDS-2017(Packet)	**1.8586**	**2.0∼13.2**	0.16	∞
CICIDS-2018(Flow)	**1.3270**	**2.85∼15.23**	0.13	∞

DP-ML Adaptation Results				
Model	CICIDS-2017(Packet)		CICIDS-2018(Flow)	
	Acc	F1	Acc	F1
Catboost	0.7359	0.1977	0.6676	0.4594
MLP	0.5984	0.0917	0.5914	0.1240
Decision Tree	0.4617	0.1081	0.6046	0.4170
LR	0.5847	0.0820	0.5925	0.1241
Random Forest	0.6081	0.0911	0.6299	0.2745
Avg	0.5978	0.1141	0.6172	0.2798

Privacy. In our DP-NetDDPM framework, we utilize DP-SGD for data synthesis and apply it during the pre-training phase with public datasets, addressing potential privacy concerns from the outset. We fixed the δ value at 10^{-5} during pre-training, while optimizing other hyperparameters through Optuna.

Table III shows the effectiveness of DP-NetDDPM in synthetic data generation. Using DCR metrics and epsilon values to measure privacy protection levels, our framework achieves a DCR value of 1.8, indicating a balanced trade-off between fidelity and privacy. The epsilon values of 2.0 at 1000 epochs and 13.2 at 20000 epochs demonstrate strong privacy protection.

As shown in Fig. 5, the data fidelity distribution difference between DP-SGD and standard models remains minimal, indicating that the Diffusion model maintains strong learning capabilities while satisfying privacy requirements. Furthermore, in terms of ML Adaptability, our DP-enabled model outperforms baseline models like TVAE and CTABGAN in accuracy metrics, demonstrating superior fidelity despite privacy constraints.

V. Conclusion

In this study, we provide DP-NetDDPM framework for synthetic network traffic generation. Our empirical analysis shows that TabDDPM significantly outperforms conventional GAN and VAE models in data fidelity, adaptability, and privacy preservation. The results confirm that diffusion-based models offer promising solutions for generating high-quality synthetic data in network security applications, though future research opportunities remain in handling time-series data and packet payloads.

Acknowledgment

This work was supported by Korea Research Institute for defense Technology planning advancement(KRIT) grant funded by the Korea government(DAPA(Defense Acquisition Program Administration)) (No. KRIT-CT-22-021, Space Signal Intelligence Research Laboratory, 2022).

References

[1] H.-J. Liao, C.-H. R. Lin, Y.-C. Lin, and K.-Y. Tung, "Intrusion detection system: A comprehensive review," *Journal of Network and Computer Applications*, vol. 36, no. 1, pp. 16–24, 2013.

[2] X. Jiang, S. Liu, A. Gember-Jacobson, A. N. Bhagoji, P. Schmitt, F. Bronzino, and N. Feamster, "Netdiffusion: Network data augmentation through protocol-constrained traffic generation," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 8, no. 1, pp. 1–32, 2024.

[3] 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.

[4] Z. Ding, X.-Y. Liu, M. Yin, and L. Kong, "Tgan: Deep tensor generative adversarial nets for large image generation," *arXiv preprint arXiv:1901.09953*, 2019.

[5] Z. Zhao, A. Kunar, R. Birke, and L. Y. Chen, "Ctab-gan: Effective table data synthesizing," in *Asian Conference on Machine Learning*. PMLR, 2021, pp. 97–112.

[6] A. Kotelnikov, D. Baranchuk, I. Rubachev, and A. Babenko, "Tabddpm: Modelling tabular data with diffusion models," in *International Conference on Machine Learning*. PMLR, 2023, pp. 17 564–17 579.

[7] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," *arXiv preprint arXiv:2011.13456*, 2020.

[8] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.

[9] D. P. Kingma, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.

[10] Z. Zhao, A. Kunar, R. Birke, H. Van der Scheer, and L. Y. Chen, "Ctab-gan+: Enhancing tabular data synthesis," *Frontiers in big Data*, vol. 6, p. 1296508, 2024.

[11] E. Hoogeboom, D. Nielsen, P. Jaini, P. Forré, and M. Welling, "Argmax flows and multinomial diffusion: Learning categorical distributions," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12 454–12 465, 2021.

[12] F. Mazé and F. Ahmed, "Diffusion models beat gans on topology optimization," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 37, no. 8, 2023, pp. 9108–9116.

[13] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," *Advances in neural information processing systems*, vol. 30, 2017.

[14] B.-E. Zolbayar, R. Sheatsley, P. McDaniel, M. J. Weisman, S. Zhu, S. Zhu, and S. Krishnamurthy, "Generating practical adversarial network traffic flows using nidsgan," *arXiv preprint arXiv:2203.06694*, 2022.

[15] Z. Lin, Y. Shi, and Z. Xue, "Idsgan: Generative adversarial networks for attack generation against intrusion detection," in *Advances in Knowledge Discovery and Data Mining*, J. Gama, T. Li, Y. Yu, E. Chen, Y. Zheng, and F. Teng, Eds. Cham: Springer International Publishing, 2022, pp. 79–91.

[16] N. Sivaroopan, D. Bandara, C. Madarasingha, G. Jourjon, A. P. Jayasumana, and K. Thilakarathna, "Netdiffus: Network traffic generation by diffusion models through time-series imaging," *Computer Networks*, vol. 251, p. 110616, 2024.

[17] I. Sharafaldin, A. H. Lashkari, A. A. Ghorbani *et al.*, "Toward generating a new intrusion detection dataset and intrusion traffic characterization." *ICISSp*, vol. 1, pp. 108–116, 2018.