

A Method for Privacy-Safe Synthetic Health Data

Xiaohui Luo

School of Software Engineering, Chengdu University of Information Technology, Chengdu 610225, China

Abstract: Private health records are important for medical research but hard to get because of legal rules. This shortage of data can be solved by using generative models like GANs, which make new, similar data. But GANs might leak private information. To fix this, we made a new kind of GAN with a privacy protection part called DP-ACTGAN. It uses differential privacy to keep the original data safe. We also put a classifier in the GAN to make sure the new data is very close to the real data. Experiments show that DP-ACTGAN can make good quality data without giving away private information. This means we can use data well without breaking privacy, which is good for ethical research and making new things while keeping privacy.

Keywords: Generative Adversarial Networks, Differential privacy, Health records, Synthetic data.

1. Introduction

Private health information records are useful for medical studies [1]. However, because of the legal rules, this kind of personal health information is often not readily available to researchers, leading to a shortage of data. When faced with data scarcity issues, using generative models can be a smart solution. Generative Adversarial Networks (GANs) [2] has been very successful in generation field. GANs essentially learn from real data and then generate new, high-quality synthetic data based on what they've learned. This generated data can address data scarcity issues without compromising the privacy of real data [3-5].

The central privacy concern with GANs is the potential for inadvertent disclosure of sensitive information from the training dataset. When GANs are trained to produce new data points, they may inadvertently capture and replicate specific and potentially identifying patterns from the original dataset. If an adversary were to extensively analyze the output data, it becomes possible to detect these patterns, which could reveal private attributes or even reconstruct original data points. This is particularly true if the GAN overfits to the training data, essentially memorizing specific details rather than learning the broader distribution [6]. Numerous inference attack techniques have been developed to exploit this vulnerability by examining the synthetic data with the intention of unearthing or deducing the source data, posing a significant threat to the confidentiality of the data used in training the GANs [7-9]. In experiments involving attackers, GAN models that incorporate differential privacy can effectively resist their attacks.

This paper tackles the problem of privacy breaches in data generated by GANs by integrating a differential privacy module into the architecture of the GAN framework, naming it DP-ACTGAN, which stands for Differentially Private - Auxiliary Classifier Conditional Tabular Generative Adversarial Network. Differential privacy is a mathematical framework designed to provide strong privacy guarantees when analyzing and sharing information from original data. By adding Gaussian noise during the training process to achieve differential privacy, we aim to ensure that the generated data cannot be used to infer sensitive information about individuals in the original dataset. However, integrating a differential privacy module may impact the quality of the

generated data. To ensure the quality of the generated data, this paper adds a classifier to the original GAN model architecture to supervise the generator in producing data that more closely resembles the real data.

In the experimental results, the data generated by DP-ACTGAN not only had the best ability to withstand MIA but also maintained data utility similar to that of the real data. This implies that the model can create synthetic datasets that not only closely mimic the statistical properties of the original data but also incorporate mechanisms to preserve the privacy of the individuals represented in the dataset.

By achieving the balance between data utility and privacy, the use of DP-ACTGAN facilitates a more ethical approach to data science, where the value of data can be harnessed without compromising the privacy rights of individuals. This creates new opportunities for research and innovation with privacy in mind.

2. Related Work

2.1. Generative Adversarial Networks for Data Privacy

Generative Adversarial Networks (GANs) and their variants have been increasingly recognized for their potential in facilitating data sharing while preserving privacy. By generating new data instances that mimic the statistical properties of original datasets, GANs provide a mechanism for sharing information without disclosing actual records, thus protecting individual privacy.

In the field of medical imaging, GANs are used to create realistic images such as X-rays or MRI scans for data sharing and model training, without disclosing the information of real patients. These synthetic images can enhance the size and diversity of datasets, aiding in the training of machine learning models, while simultaneously protecting patient privacy [10, 11]. Researchers [13, 14] use GANs to generate health records. This synthetic data can be used for medical research and model training without compromising patient privacy, demonstrating potential for advancing healthcare analytics while adhering to privacy regulations.

Under Membership Inference Attacks (MIA), GANs may leak information about the real data as attackers might be able to identify that certain synthetic data points were generated based on specific real training data. This type of attack is

particularly likely to occur when GANs capture and replicate unique or identifiable patterns in the training data, allowing attackers to infer that certain individuals or information were used in the training of the model, leading to potential privacy breaches [7-9].

2.2. Generative Adversarial Networks with Differential Privacy

Differential privacy enhances the defense of datasets against MIA by introducing noise into the data processing procedure, thereby ensuring individual privacy.

Dwork et al. [15] introduced an innovative technique for protecting privacy of data exposure risks, known as Differential Privacy (DP) which defined as follows:

$$\mathbb{P}(M(d) \in S) \leq e^\epsilon \cdot \mathbb{P}(M(d') \in S) + \delta \quad (1)$$

In the formula, d and d' represent two adjacent datasets that differ by the data of just one individual. ϵ is the privacy budget which control the level of the privacy. An algorithm M that meets the above formula satisfies (ϵ, δ) -DP. This approach aims to minimize the potential impact on individual privacy during data processing and analysis. In this privacy protection method, the released data is moderately disturbed, making it difficult to reconstruct or identify an individual's information, thereby effectively safeguarding privacy. Mironov et al. [16] proposed a new definition of privacy called Rényi Differential Privacy (Rényi DP). Rényi DP shares many properties with traditional differential privacy, making it a natural extension of it. Moreover, Rényi DP can capture privacy loss more accurately through the use of composition theorems.

Acs et al. [17] were the first to attempt to build a GAN model incorporating DP. Their method divides the entire

training dataset into k subsets using a differentially private k -means algorithm, and then trains a local generative model on each subset independently. Xie et al. [18] proposed DPGAN, which achieves differential privacy by clipping weights and adding Gaussian noise. It employs the loss function from WGAN [19] to produce better results while also combating mode collapse, which is critical for balancing privacy and utility. During the training phase, adding Gaussian noise to the gradient calculations and clipping the updated weights helps to achieve smaller privacy losses.

Incorporating differential privacy techniques into GANs makes it more challenging to infer any individual information from the generated data, as differential privacy ensures that even minor changes, such as the addition or removal of individual data, do not significantly alter the output results.

3. Methods

3.1. Problem Definition

This article defines real health records as $\mathcal{D}_{\text{data}} = \{(\mathbf{X}, \mathbf{Y})\}$, where $\mathbf{X} = \{\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^m\}$ is table-type data with $\mathbf{X}^m \in \mathcal{R}^n$, meaning there are m samples of tabular health data, each with n features. For the labels, $\mathbf{Y} = \{\mathbf{Y}^1, \mathbf{Y}^2, \dots, \mathbf{Y}^m\}$, where $\mathbf{Y}^m \in \mathcal{R}^2$, forming the label set. These variables follow an unknown joint distribution, each row is a sample from this joint distribution, and each row is independently sampled, meaning the order of the rows doesn't need to be considered. The goal of this article is to train a generative model to produce a table \mathbf{T} that meets the following criteria: First, the generated data \mathbf{T} should be used to train a classification model and achieve similar results on a real test set as the real training set. Second, the generated table \mathbf{T} should have a similar statistical distribution to the original dataset. Lastly, the generated table \mathbf{T} should not expose any private information from the real data.

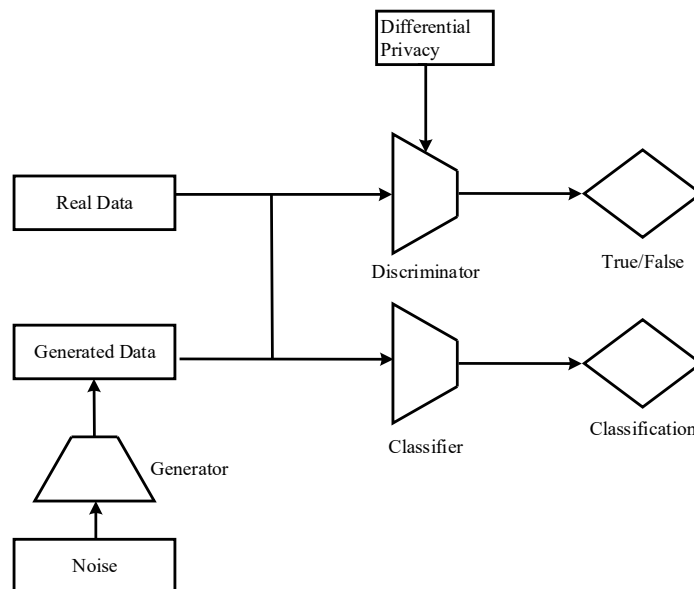


Figure 1. Structure of DP-ACTGAN

3.2. Model Structure

The DP-ACTGAN is an improvement on past studies like [12, 18], adding new features that improve the quality of generated data and provide a guarantee of privacy protection.

Figure 1 displays the overall structure of DP-ACTGAN. For training the generator and discriminator, we employ a loss function that is consistent with the one described in [12].

To improve the quality of generated data, we added a classifier to the original GAN structure. This classifier

employs Binary Cross-Entropy (BCE) as its loss function and the formula for BCE is shown as Equation 2. This classifier is trained using the real labels from the original tabular data, allowing it to learn the relationship between labels and features in the real data. When a generated record is presented, the classifier can assess whether the relationship between features and labels is correct. If the classifier detects an error, it can provide feedback to the generator, helping it create more accurate and realistic data. This can enhance the quality and credibility of the generated data. Figure 2 displays the network structure of DP-ACTGAN.

$$L_{class} = -\frac{1}{k} \sum_i^k y_i \log(p_i) + (1 - y_i) \log(1 - p_i) \quad (2)$$

To address privacy concerns with data, we've incorporated a differential privacy module into our model to protect the privacy of the original data. By introducing random noise to the data, we can mask individual contributions and make it hard to distinguish any single individual's data from the dataset.

Laplace and Gaussian mechanisms are two popular methods used to add noise for differential privacy [16, 20]. The reason for using the Gaussian mechanism is that it strikes a balance between data utility and privacy. To get similar

privacy protection, the amount of Gaussian noise added will be less than the amount of Laplace noise [16]. It allows us to maintain the statistical properties of the dataset while still providing strong privacy guarantees.

During training step, we add Gaussian noise to the gradient of discriminator. The model keeps an eye on the privacy budget used by checking it each time we put noise into the gradient. We use the Rényi Differential Privacy Accountant [16], which gives a more accurate privacy budget estimate compare with moment accountant.

The formula of Rényi Differential Privacy Accountant to calculate the privacy budget is shown as Equation 3.

$$\epsilon(\alpha) = \frac{\alpha \cdot \Delta_2(f)^2}{2\sigma^2} \quad (3)$$

$\Delta_2(f)^2$ is the L_2 sensitivity (i.e., the maximum Euclidean norm difference on any two adjacent datasets). σ is the standard deviation of the added Gaussian noise.

With the same privacy budget, we can add less noise to keeps important information better during training while protecting the privacy. In this way, we can improve the quality of the generated data while ensuring privacy.

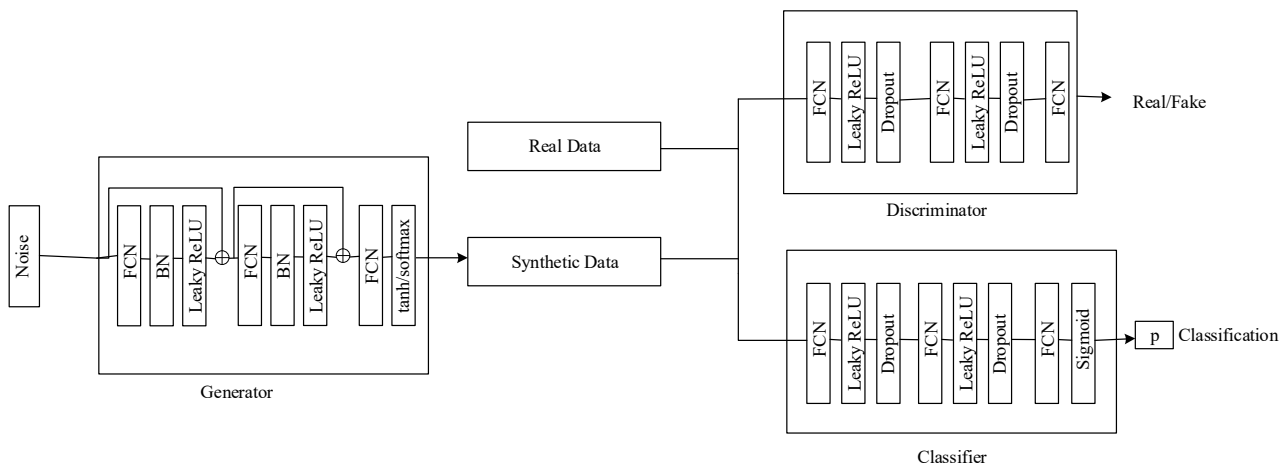


Figure 2. Network Structure of DP-ACTGAN

4. Experiments

Table 1. Datasets description

Datasets	Samples
Pulmonary complications	17356
ICU admission	12240
Cardiovascular adverse events	17356

The experiments in this paper validate the utility of the generated data, statistical similarity between real data and generated data, and the privacy protection of the real data. The experiments use data from a hospital database.

4.1. Evaluation Metrics

4.1.1. Machine Learning Utility

Machine learning utility means dividing the original data into a training set and a test set in a 7:3 ratio. First, we train a GAN model and a prediction classifier using the training set. Then, we use the GAN to generate data and train another

classifier with this synthetic data. Finally, we evaluate both sets of classifiers using the real test set data to compare their performance, which assesses the utility of the generated data. In the classification model, we use F1 score to measure the model's performance. This metric measures the utility of the generated data.

4.1.2. Statistical Similarity

This study measures the statistical similarity between real and generated data using two metrics: Jensen-Shannon divergence [21] and Wasserstein distance [22]. By using both metrics together, we can more comprehensively assess the statistical similarity between the generated and real data, providing deeper insights into the study of the generated data.

Jensen-Shannon divergence (JSD) is a measure of the difference between two probability distributions, with values ranging from 0 to 1. A JSD of 0 indicates identical probability distributions between two datasets, while a JSD of 1 indicates completely different distributions. In machine learning, JSD is often used to measure differences in probability distributions between datasets. In this paper, the smaller the JSD value, the smaller the difference in probability distribution between the generated dataset and the real data,

indicating better generation quality. The formula for JSD is shown as Equation 4, where P and Q represent the probability distributions of the real and generated datasets, respectively.

$$JSD(P \parallel Q) = \frac{1}{2} \sum p(x) \log \left(\frac{p(x)}{\frac{p(x)+q(x)}{2}} \right) + \frac{1}{2} \sum q(x) \log \left(\frac{q(x)}{\frac{p(x)+q(x)}{2}} \right) \quad (4)$$

The Wasserstein distance (WD) is a metric used to measure the similarity of two probability distributions, where a smaller value indicates more similarity. Unlike other distance measures, the Wasserstein distance can handle probability distributions with different supports, not just measuring distance over a common set. In this paper, the smaller the WD value, the less difference there is between the generated dataset and the real data, indicating better generation quality. The formula for WD is shown as Equation 5. Here, P and Q are the probability distributions of the real and generated datasets, respectively, and $\Pi(P, Q)$ represents the set of all possible joint distributions combining P and Q. The application of these two metrics helps to quantify the similarity between the generated data and real data, providing a powerful tool for evaluating generative models.

$$WD(P, Q) = \inf_{\gamma \in \Pi(P, Q)} E_{x, y \sim \gamma} [\|x - y\|] \quad (5)$$

4.1.3. Privacy Metrics

This study used two methods to verify the privacy protection of the generated data for the real data: Distance to Closest Record [23] and Membership Inference Attack [24].

Distance to Closest Record (DCR) measures how susceptible the generated data is to re-identification attacks. DCR refers to the shortest distance from data points in dataset A to data points in dataset B, calculated using the Euclidean distance. A smaller DCR value indicates poorer privacy

protection of the generated data, potentially exposing real information, while a larger DCR value suggests better privacy protection. In experiments, for each piece of generated data, we select the closest 10 real data points, average their distances, and then calculate the average DCR for the generated dataset to get an average DCR for a dataset. The paper conducted three DCR experiments on three datasets to analyze the risk and possibility of privacy exposure.

Membership Inference Attack (MIA) aims to find out if specific data was used to train a model by looking at the model's outputs. Using an MIA model that targets generative models to assess the privacy of generated data can further understand how the data holds up against privacy attacks.

This study used a generative model-specific MIA model, DOMIAS [25], to simulate potential privacy attack scenarios. This kind of evaluation helps to fully grasp the privacy challenges that generated data may encounter in real-world use, and gives focused advice to improve the privacy features of generative models. The MIA model uses Accuracy as an evaluation metric, which is a measure of how well a classification model performs, ranging from 0 to 1. A value closer to 1 means better model performance, while a value closer to 0 indicates poorer classification performance. In the MIA model, a lower Accuracy value suggests that it's harder for an attacker to correctly identify the training data of the generative model from the generated data, meaning better privacy protection against membership inference attacks. If the Accuracy is higher, it means the MIA model can more accurately identify the source of the generated data, which could pose a privacy risk to the real data.

4.2. Experiment Setup

We compared DP-ACTGAN with the CTGAN model and the DPGAN model on three datasets, with the differential privacy budget ϵ set to 10 and 100, δ set to 10^{-5} . We used the above metrics to measure the overall performance of the GANs.

4.3. Results Analysis

4.3.1. Machine Learning Utility

Table 2. Result of ML utility

Model	ϵ	Pulmonary complications	ICU admission	Cardiovascular adverse events
Origin	-	0.502	0.509	0.303
CTGAN	-	0.503	0.501	0.306
DPGAN	10	0.177	0.204	0.099
	100	0.262	0.353	0.159
DP-ACTGAN	10	0.375	0.269	0.244
	100	0.485	0.496	0.301

In Table 2, the model list represents data generated using different models, while Origin indicates the use of original data. In theory, a smaller privacy budget ϵ provides stronger privacy protection but adds more noise to the generative model, affecting the utility of the generated data. To explore the balance between privacy and utility, we set ϵ to 10 and 100 for our experiments. The table 2 shows that when ϵ is 10, the utility of the data generated by the two GAN models with differential privacy is poor. This is because smaller ϵ values result in more noise, which degrades the utility of the generated data. When ϵ is 100, the privacy protection decreases, but the utility of the data generated by the GAN models improves. The prediction model trained with data

generated by DP-ACTGAN has similar performance to models trained with real data.

The table 2 also shows that under the same privacy budget constraints, the utility of data generated by DP-ACTGAN is better than that of DPGAN. This is because using the Rényi accountant method allows for a stricter calculation of ϵ , meaning less noise can be added under the same privacy budget. This ensures privacy while more effectively preserving key information during training, thus enhancing the utility of the generated data.

4.3.2. Statistical Similarity

The statistical similarity experiments in this chapter use the JSD and WD metrics to measure the statistical similarity

between the generated data and the real training set. The experiments show the statistical similarity of the data generated by the three models on three datasets when the privacy budget ϵ is set to 10 and 100, respectively. The experimental results are shown in Table 3.

Both JSD and WD are metrics where smaller values indicate greater similarity between the generated data and the

real data. Comparing the results when ϵ is set to 100 with when it is set to 10, we find that the data generated by the two GAN models is more similar to the real data when ϵ is set to 100. This suggests that larger amounts of noise not only affect the utility of the generated data but also reduce its statistical similarity to the real data.

Table 3. Result of Statistical similarity

Datasets	Model	ϵ	WD	JSD
Pulmonary complications	CTGAN	-	0.032	0.097
		10	0.046	0.108
		100	0.036	0.092
	DP-ACTGAN	10	0.042	0.109
		100	0.033	0.09
		100	0.032	0.091
ICU admission	CTGAN	-	0.032	0.091
		10	0.084	0.132
		100	0.051	0.096
	DP-ACTGAN	10	0.074	0.13
		100	0.048	0.094
		100	0.032	0.097
Cardiovascular adverse events	CTGAN	-	0.032	0.097
		10	0.041	0.107
		100	0.033	0.091
	DP-ACTGAN	10	0.043	0.1
		100	0.034	0.085
		100	0.034	0.085

4.3.3. Privacy Metrics

This paper uses DCR and MIA to verify the privacy protection of the generated data for the real data, with the DCR experimental results shown in Table 4. Firstly, it can be observed that the DCR values for the data generated by the three GAN models all exceed those of the original dataset. This indicates that the GAN models do not simply memorize and generate samples similar to the original data, but have successfully learned the patterns in the data, creating new

samples that differ from the original dataset but have similar characteristics as the generated dataset. From the table, it can be also seen that in the three different prediction tasks, the privacy protection provided by setting the generative model's privacy budget ϵ to 10 is higher than that provided when setting it to 100. This is consistent with the definition of differential privacy, which states that setting a smaller ϵ for the model introduces more noise, thus providing stronger privacy protection.

Table 4. Result of DCR

Model	ϵ	Pulmonary complications	ICU admission	Cardiovascular adverse events
Origin	-	1.86	2.06	1.86
CTGAN	-	2.37	2.48	2.37
DPGAN	10	2.42	2.69	2.42
	100	2.34	2.56	2.39
DP-ACTGAN	10	2.58	2.85	2.81
	100	2.4	2.59	2.5

Upon further analysis of the experimental results in the table, it can be observed that even when ϵ is set to 100, the DCR metrics for data generated by DP-ACTGAN across all three prediction tasks are still better than the GAN model which not employing differential privacy techniques. This result indicates that even with a more relaxed privacy budget setting, the use of differential privacy technology can effectively enhance the privacy protection level of the generated data, reducing the potential risk of privacy

exposure.

The MIA experiment results are shown in Table 5. The table indicates that compared to the GAN model without the differential privacy module, the data generated by DP-ACTGAN and DPGAN can more effectively resist membership inference attacks, enhancing the privacy protection of the generated data for the real data. Furthermore, as the privacy budget decreases, the protection for the original data improves.

Table 5. Result of MIA

Model	ϵ	Pulmonary complications	ICU admission	Cardiovascular adverse events
CTGAN	-	0.525	0.493	0.518
DPGAN	10	0.488	0.477	0.469
	100	0.493	0.48	0.487
DP-ACTGAN	10	0.472	0.463	0.466
	100	0.487	0.467	0.484

Synthesizing the DCR and MIA experiment results, it is evident that the privacy protection for data generated by DP-

ACTGAN is consistently better than that for DPGAN. Moreover, across the three datasets, compared to the CTGAN

model which lacks a differential privacy module, the use of GAN models with differential privacy protection has enhanced the privacy protection of the generated data relative to the real data. Even in the face of membership inference attacks, the data generated by GANs with differential privacy maintains a higher degree of privacy protection, reducing the risk of individual privacy information being identified and leaked.

And combining experiments on machine learning utility, statistical similarity, and privacy exposure risk, it is demonstrated that introducing differential privacy techniques into GANs does not sacrifice all data utility. Although using differential privacy impacts the quality of the generated data to some extent, by adjusting the privacy budget parameters and other settings, a balance between privacy protection and data quality can be found. This allows the generated data to retain practical utility similar to the real data while providing privacy protection, supporting subsequent prediction tasks or academic sharing. This type of generated data, which offers both privacy protection and utility, is important for postoperative risk prediction because it ensures data privacy and security, and can be used for academic sharing, providing researchers with a wider range of data resources and promoting collaborative academic research.

5. Conclusion

DP-ACTGAN is capable of addressing the challenge of accessing real personal health data which is often restricted due to privacy policies. Doctors can utilize the data generated by this model for academic sharing, offering a viable solution for medical research and data analysis that adheres to privacy regulations, thereby overcoming the issue of data scarcity caused by privacy concerns.

References

- [1] Kalkman S, van Delden J, Banerjee A, et al. Patients' and public views and attitudes towards the sharing of health data for research: a narrative review of the empirical evidence[J]. *Journal of medical ethics*, 2022, 48(1): 3-13.
- [2] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems* (pp. 2672-2680).
- [3] Antoniou, A., Storkey, A., & Edwards, H. (2017). Data Augmentation Generative Adversarial Networks. *arXiv preprint arXiv:1711.04340*.
- [4] Skandarani Y, Painchaud N, Jodoin P M, et al. On the effectiveness of GAN generated cardiac MRIs for segmentation[J]. *arXiv preprint arXiv:2005.09026*, 2020.
- [5] Chen R J, Lu M Y, Chen T Y, et al. Synthetic data in machine learning for medicine and healthcare[J]. *Nature Biomedical Engineering*, 2021, 5(6): 493-497.
- [6] Ma C, Li J, Ding M, et al. RDP-GAN: AR\enyi-Differential Privacy based Generative Adversarial Network[J]. *arXiv preprint arXiv:2007.02056*, 2020.
- [7] Hitaj B, Ateniese G, Perez-Cruz F. Deep models under the GAN: information leakage from collaborative deep learning[C]//*Proceedings of the 2017 ACM SIGSAC conference on computer and communications security*. 2017: 603-618.
- [8] Hu H, Pang J. Membership inference attacks against gans by leveraging over-representation regions[C]//*Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. 2021: 2387-2389.
- [9] Chen D, Yu N, Zhang Y, et al. Gan-leaks: A taxonomy of membership inference attacks against generative models[C]//*Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*. 2020: 343-362.
- [10] Frid-Adar, M., Diamant, I., Klang, E., Amitai, M., Goldberger, J., & Greenspan, H. (2018). GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification. *Neurocomputing*, 321, 321-331.
- [11] Mao X, Li Q, Xie H, et al. Least squares generative adversarial networks[C]//*Proceedings of the IEEE international conference on computer vision*. 2017: 2794-2802.
- [12] Xu L, Skoularidou M, Cuesta-Infante A, et al. Modeling tabular data using conditional gan[J]. *Advances in neural information processing systems*, 2019, 32.
- [13] Yahi A, Vanguri R, Elhadad N, et al. Generative adversarial networks for electronic health records: A framework for exploring and evaluating methods for predicting drug-induced laboratory test trajectories[J]. *arXiv preprint arXiv:1712.00164*, 2017.
- [14] Choi E, Biswal S, Malin B, et al. Generating multi-label discrete patient records using generative adversarial networks[C]//*Machine learning for healthcare conference*. PMLR, 2017: 286-305.
- [15] Dwork C, Roth A. The algorithmic foundations of differential privacy[J]. *Foundations and Trends® in Theoretical Computer Science*, 2014, 9(3-4): 211-407.
- [16] Mironov I. Rényi differential privacy[C]//*2017 IEEE 30th computer security foundations symposium (CSF)*. IEEE, 2017: 263-275.
- [17] Acs G, Melis L, Castelluccia C, et al. Differentially private mixture of generative neural networks[J]. *IEEE Transactions on Knowledge and Data Engineering*, 2018, 31(6): 1109-1121.
- [18] Xie L, Lin K, Wang S, et al. Differentially private generative adversarial network[J]. *arXiv preprint arXiv:1802.06739*, 2018.
- [19] Gulrajani I, Ahmed F, Arjovsky M, et al. Improved training of wasserstein gans[J]. *Advances in neural information processing systems*, 2017, 30.
- [20] Huang Z, Mitra S, Dullerud G. Differentially private iterative synchronous consensus[C]//*Proceedings of the 2012 ACM workshop on Privacy in the electronic society*. 2012: 81-90.
- [21] Lin J. Divergence measures based on the Shannon entropy[J]. *IEEE Transactions on Information theory*, 1991, 37(1): 145-151.
- [22] Ramdas A, García Trillos N, Cuturi M. On wasserstein two-sample testing and related families of nonparametric tests[J]. *Entropy*, 2017, 19(2): 47.
- [23] Lu P H, Wang P C, Yu C M. Empirical evaluation on synthetic data generation with generative adversarial network[C]//*Proceedings of the 9th International Conference on Web Intelligence, Mining and Semantics*. 2019: 1-6.
- [24] Shokri R, Stronati M, Song C, et al. Membership inference attacks against machine learning models[C]//*2017 IEEE symposium on security and privacy (SP)*. IEEE, 2017: 3-18.
- [25] van Breugel B, Sun H, Qian Z, et al. Membership inference attacks against synthetic data through overfitting detection[J]. *arXiv preprint arXiv:2302.12580*, 2023.