

# PPGAN: Privacy-preserving Generative Adversarial Network

Yi Liu<sup>1</sup>, Jialiang Peng<sup>1,\*</sup>, James J.Q. Yu<sup>2</sup>, Yi Wu<sup>1</sup>

<sup>1</sup>*School of Data Science and Technology, Heilongjiang University*

<sup>2</sup>*Department of Computer Science and Engineering, Southern University of Science and Technology*  
97liuyi@ieee.org, pengjialiang@hlju.edu.cn, yujq3@sustech.edu.cn, 1995050@hlju.edu.cn

**Abstract**—Generative Adversarial Network (GAN) and its variants serve as a perfect representation of the data generation model, providing researchers with a large amount of high-quality generated data. They illustrate a promising direction for research with limited data availability. When GAN learns the semantic-rich data distribution from a dataset, the density of the generated distribution tends to concentrate on the training data. Due to the gradient parameters of the deep neural network contain the data distribution of the training samples, they can easily *remember* the training samples. When GAN is applied to private or sensitive data, for instance, patient medical records, as private information may be leakage. To address this issue, we propose a *Privacy-preserving Generative Adversarial Network (PPGAN)* model, in which we achieve differential privacy in GANs by adding well-designed noise to the gradient during the model learning procedure. Besides, we introduced the *Moments Accountant* strategy in the PPGAN training process to improve the stability and compatibility of the model by controlling privacy loss. We also give a mathematical proof of the differential privacy discriminator. Through extensive case studies of the benchmark datasets, we demonstrate that PPGAN can generate high-quality synthetic data while retaining the required data available under a reasonable privacy budget.

**Index Terms**—Privacy leakage, GAN, deep learning, differential privacy, moments accountant.

## I. INTRODUCTION

In recent years, researchers have used a large number of training data to perform data mining tasks, in the field of medical and health informatics, such as disease prediction and auxiliary diagnosis. Deep learning models are employed to remember the characteristics of a large number of training samples for classification or prediction purposes. However, organizations such as hospitals and research institutes are paying more and more attention to the protection of data. Additionally, the *General Data Protection Regulation (GDPR)*[1] issued by the European Union prohibits organizations from sharing private data. It is increasingly difficult for researchers to obtain training data unlimited legally.

Fortunately, the generative model provides us with a solution to the issue of data scarcity [2], yet data privacy leakage issues may arise. StyleGAN [3] shown impressive performance in generating fake face images. In principle, it can

memorize data distribution from the small amount of training data, rendering indistinguishable high-quality “fake” samples. However, for most people, they expect their face data not to be used as a training sample. GAN can implicitly disclose the privacy information of training samples. GAN model produces high-quality “fake” samples through continuous training and resampling. This training method grants hackers the opportunity to restore the original samples. Therefore, we not only need high-quality sample generation approaches but also need to achieve a reasonable level of data privacy [4].

Based on the above findings, we propose a **Privacy-preserving GAN (PPGAN)**. PPGAN combines with differential privacy [5] to ensure that the exact training samples can not be revealed by adversaries from the trained model, resulting in well-protected data privacy. In particular, we added well-designed noise to the gradients in the training process in PPGAN and used the framework of the WGAN [6] model as the main skeleton of PPGAN. The proposed model does not suffer from a privacy leakage issue whose proportional to the volume of data thanks to the introduced average aggregator that offsets the privacy overhead of large datasets.

We would like to point out our main contributions as follows:

- We propose the PPGAN framework that can generate high-quality data points while protecting data privacy. PPGAN combines noise well-designed in the differential privacy with training gradients to disturb the distribution of the original data. Finally, we give a rigorous proof of the differential privacy discriminator in mathematics.
- We introduced the *Moments Accountant* strategy that maintains the boundedness of the function, controls the privacy level and significantly improves the stability of the model training.
- We evaluated PPGAN with benchmark datasets. The results show that PPGAN can generate high-quality data with adequately protected privacy under a reasonable privacy budget.

The overall structure of this paper is as follows. First, we introduce the proposed PPGAN framework and its theoretical proof in Section II. We assess the performance of our framework in Section III. Finally, this paper is concluded in Section IV.

\* Corresponding Author: pengjialiang@hlju.edu.cn. This work is supported by the Ministry of Education of China and the School of Entrepreneurship Education of Heilongjiang University (Grant NO.201910212133) and Heilongjiang Provincial Natural Science Foundation of China (Grant NO.QC2016091).

## II. METHODOLOGY

### A. Differential Privacy

Differential privacy (DP) [7], [5], [8] constitutes a solid standard for privacy guarantee for algorithms on the database. Then we give the definition of the distance between two databases  $x$  and  $y$  as follows:

**Definition 1: (Distance Between Databases)**

The  $\ell_1$  norm of a database  $x$  is denoted  $\|x\|_1$  and is defined to be:

$$\|x\|_1 = \sum_{i=1}^{|\mathbb{N}|} |x_i| \quad (1)$$

The  $\ell_1$  distance between two databases  $x$  and  $y$  is  $\|x - y\|_1$ . In particular, when  $\|x - y\|_1 = 1$ ,  $x$  and  $y$  are mutually referred to as neighboring datasets.

**Definition 2:  $((\epsilon, \delta)$ -DP)**

A randomized algorithm  $\phi(\cdot)$  with domain  $\Phi^{|\mathbb{N}|}$  is  $(\epsilon, \delta)$ -DP if for all  $O \subseteq \text{Range}(\phi)$  and for all  $d, d' \in \Phi^{|\mathbb{N}|}$  (for any neighbouring datasets) such that  $\|d - d'\| \leq 1$ :

$$\Pr[\phi(d) \in O] \leq e^\epsilon \Pr[\phi(d') \in O] + \delta \quad (2)$$

Noted that  $\epsilon$  stands for privacy budget, which controls the level of privacy guarantee achieved by mechanism  $\phi$ . And when  $\epsilon = \infty$ , this case is non-private.

Among the mechanisms for achieving differential privacy, the two most widely used are the *Laplace mechanism* and the *Gaussian noise mechanism (GNM)* [9]. Due to the combined properties of the GNM, it is prevalent in many DP protection models. In PPGAN, we use the GNM because the moments accountant (detailed in Section II-D) provides an improved privacy boundary analysis and is well-matched to the combined properties of the GNM. The GNM is defined as follows:

$$\phi(x) \triangleq f(x) + N(0, \sigma^2 s_f^2) \quad (3)$$

where  $s_f$  is defined as *sensitivity*, which is only related to query type  $f$ . The sensitivity is defined as follows:

**Definition 3: ( $L_2$  norm-Sensitivity)**

We given the neighboring datasets  $x$  and  $x'$  and given a query  $f : x \rightarrow \Omega$ , the *sensitivity* of  $f$  as follows:

$$\Delta f = \max_{x, x'} \|f(x) - f(x')\|_2 \quad (4)$$

Noted that it records the largest difference between query results on datasets  $x$  and  $x'$ .

According to the algorithm  $\phi(\cdot)$  in Definition 2 is stochastic and is not related to the distribution of the output data. Moreover, the Gaussian noise mechanism adds a well-design noise to a single gradient without affecting the entire gradient aggregation. Therefore, we can use this attribute with GAN so that GAN can generate high-quality data while satisfying differential privacy.

### B. GAN and WGAN

Generative adversarial network (GAN) [10], [11] is a class of deep neural network architectures comprised of two networks, pitting one against the other (thus the “*adversarial*”). The objective function of the discriminative model is as follows:

$$\max_D E_{x \sim P_r}[\log(D(x))] + E_{x \sim P_g}[\log(1 - D(x))] \quad (5)$$

The goal of a similar from distinguishing is to prevent them from real records and the generated ones. The entire optimization objective function is as follows:

$$\min_G \max_D V(G, D) = E_{x \sim P_{data}(x)}[\log(D(x))] + E_{z \sim P_z(z)}[\log(1 - D(G(z)))] \quad (6)$$

WGAN [6] uses the *Wasserstein* distance instead of the *Jensen-Shannon* distance. Compared with the original GAN, WGAN’s parameters are less sensitive and the training process is smoother. It solves a minimax two-player game that finds the balance point of each other:

$$\min_G \max_{w \in W} E_{x \sim P_{data}(x)}[f_w(x)] - E_{z \sim P_z(z)}[f_w(G(z))] \quad (7)$$

### C. PPGAN framework

In this section, we present the proposed *Privacy-preserving Generative Adversarial Network (PPGAN) model*, which is detailed in Algorithm 1 and illustrated in Fig. 1.

---

**Algorithm 1** Privacy-preserving Generative Adversarial Network (PPGAN)

---

**Require:**

The learning rate:  $\alpha$ . The clipping parameter:  $c$ . The mini-batch size:  $m$ . The number of discriminator iterations per generator iteration:  $n_d$ . Generator iteration:  $n_g$ . Noise scale:  $\sigma_n$ .

**Ensure:**

DP generator  $\theta$ ;

- 1: Initialize generator parameters and discriminator parameters  $\omega_0, \theta_0$ , respectively.
  - 2: **for**  $t_1 = 1, \dots, n_g$  **do**
  - 3:   **for**  $t_2 = 1, \dots, n_d$  **do**
  - 4:      $\{x^{(i)}\}_{i=1}^m \sim P_r$  a mini-batch from the real data.
  - 5:      $\{z^{(i)}\}_{i=1}^m \sim p(z)$  a mini-batch of prior samples.
  - 6:      $g_\omega \leftarrow g_\omega \min(1, C/\|g_\omega\|) + N(0, \sigma_n^2 c_g^2 I)$  (**adding noise**)
  - 7:      $\omega \leftarrow \text{clip}(\omega + \alpha \cdot \text{SGD}(\omega, g_\omega), -c, c)$
  - 8:   **end for**
  - 9:    $g_\delta \leftarrow g_\delta \min(1, C/\|g_\delta\|)$
  - 10:    $\theta \leftarrow \theta - \alpha \cdot \text{SGD}(\theta, g_\delta)$
  - 11: **end for**
  - 12: **return**  $\theta$ ;
-

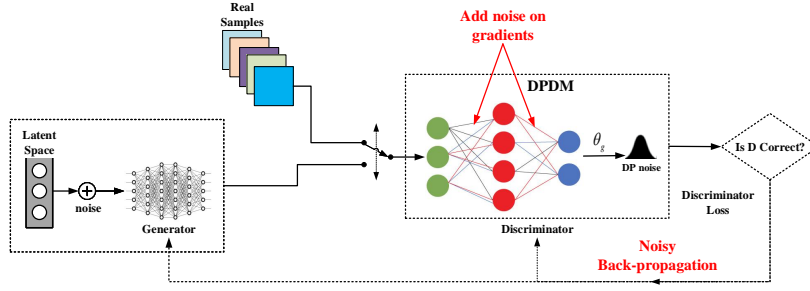


Fig. 1. Overview of our Privacy-preserving Generative Adversarial Network (PPGAN) framework.

#### D. Privacy Guarantees of PPGAN

To show that PPGAN in Algorithm 1 does satisfy the differential privacy, we prove that the parameters of the generator guarantee the differential privacy relative to the sample training point under the condition that the discriminator parameters satisfy the differential privacy. Therefore, the generated data from  $G$  satisfies the differential privacy, which means that  $G$  does not leakage the privacy of the datasets. Through moment accountant strategy, we can control the boundary of  $g_w(x^{(i)}, z^{(i)})$  and calculate the final privacy loss. Along with Definition 2, intuitively, we have the definition of privacy loss at  $\tau$ :

**Definition 4: (Privacy Loss)**

$$c(\tau; \phi, aux, d, d') \triangleq \log \frac{\mathbb{P}[\phi(aux, d) = \tau]}{\mathbb{P}[\phi(aux, d') = \tau]} \quad (8)$$

We introduce privacy loss to measure the distribution difference between two changing data. The privacy loss random variable is derived from the Definition 2, which is used to describe the privacy budget of  $\phi(d)$ . For a given mechanism  $\phi$ , we define the  $v^{th}$  moment  $\beta_\phi(v; aux, d, d')$  as the log of the moment generating function evaluated at the value:

**Definition 5: (Log moment generating function)**

$$\beta_\phi(v; aux, d, d') \triangleq \log E_{o \sim \phi}[e^{vC(\phi, aux, d, d')}] \quad (9)$$

**Definition 6: (Moments Accountant)**

$$\beta_\phi(v) \triangleq \max_{aux, d, d'} \beta_\phi(v; aux, d, d') \quad (10)$$

The basic idea behind the moments accountant is to accumulate the privacy expenditure by framing the privacy loss as a random variable and using its moment-generating functions to understand that variables distribution better. This property makes the PPGAN model training more stable. The tail bound can also be applied to privacy guarantee (In [8]). Since the moments accountant saves a factor of  $\sqrt{\log(n_g/\delta)}$ , according to Definition 2, this is a significant improvement for the large iteration  $n_g$ .

The following theorem, a proof of which can be found in [2], [8], [12], allows us to move the burden of differential privacy to the discriminator; the differential privacy of the generator will follow by the theorem.

#### Theorem 1: (Post-processing)

Let  $\phi$  be an  $(\epsilon, \delta)$ -differentially private algorithm and let  $f : \xi \rightarrow \xi'$  where  $\xi'$  is any arbitrary space. Then  $f \circ \phi$  meets  $(\epsilon, \delta)$ -differentially private.

Next, we present the mathematical reasoning proof that the discriminator satisfies the differential privacy. First, we propose a lemma that PPGAN satisfies the definition of DP.

**Lemma 1:** Under the definition of GNM and  $L_2$ -sensitivity (in Definition 3), for any  $\delta \in (0, 1)$ ,  $\sigma > \frac{\sqrt{2 \ln(1.25/\delta)} \Delta f}{\epsilon}$ , we have noise  $Y \sim N(0, \sigma^2)$  satisfies  $(\epsilon, \delta)$ -DP.

**Proof 1:** We assume that  $\Delta f$  is the  $L_2$ -sensitivity, and according to the Definition. 2, then we have:

$$\left| \ln \frac{e^{-\frac{1}{2\sigma^2}x^2}}{e^{-\frac{1}{2\sigma^2}(x+\Delta f)^2}} \right| = \left| \frac{1}{2\sigma^2} (2x\Delta f + (\Delta f)^2) \right| \leq \epsilon \therefore |x| \leq \frac{\sigma^2 \epsilon}{\Delta f} - \frac{\Delta f}{2} \Delta f. \quad (11)$$

Let  $t = \frac{\sigma^2 \epsilon}{\Delta f} - \frac{\Delta f}{2}$ , if and only if  $\|x\| \leq t$ , the distribution satisfies DP, and when  $\|x\| > t$ , we want the probability of privacy leakage to be less than  $\delta$ , so we have:  $P(x > t) < \frac{\delta}{2}$ , where  $P(\cdot)$  denotes the probability of revealing privacy. Next, we prove that the Gaussian distribution function is bounded above:

$$\begin{aligned} P(x > t) &= \frac{1}{\sqrt{2\pi}\sigma} \int_t^\infty e^{-\frac{x^2}{2\sigma^2}} dx < \frac{1}{\sqrt{2\pi}\sigma} \int_t^\infty \frac{x}{t} e^{-\frac{x^2}{2\sigma^2}} dx \\ &= \frac{\sigma}{\sqrt{2\pi}t} e^{-\frac{t^2}{2\sigma^2}} (\because x > t) \end{aligned} \quad (12)$$

Then the problem is converted to:

$$\frac{\sigma}{\sqrt{2\pi}t} e^{-\frac{t^2}{2\sigma^2}} < \frac{\delta}{2}, \frac{t}{\sigma} e^{\frac{t^2}{2\sigma^2}} > \frac{2}{\sqrt{2\pi}\sigma}, \ln \frac{t}{\sigma} + \frac{t^2}{2\sigma^2} > \ln \frac{2}{\sqrt{2\pi}\sigma} \quad (13)$$

$$\therefore \begin{cases} \ln \frac{t}{\sigma} \geq 0 \\ \frac{t^2}{2\sigma^2} > \ln \frac{2}{\sqrt{2\pi}\sigma} \end{cases} \quad (14)$$

For the left two terms of Equation 14, because  $t = \frac{\sigma^2 \epsilon}{\Delta f} - \frac{\Delta f}{2}$ , let  $\sigma = c \frac{\Delta f}{\epsilon}$ , then  $t = c\sigma - \frac{\Delta f}{2}$ , thus we have:

$$\frac{t}{\sigma} = c - \frac{\Delta f}{2\sigma} = c - \frac{\epsilon}{2c}. \quad (15)$$

Here  $\epsilon < 1, c \geq 1$ , then

$$\ln(c - \frac{\epsilon}{2c}) > \ln(c - \frac{1}{2}) \geq 0. \quad (16)$$

By Equation 16, we have  $c \geq \frac{3}{2}$ . By Equation 15, we have:

$$\frac{t^2}{2\sigma^2} = \frac{1}{2}(c^2 - \varepsilon + \frac{\varepsilon^2}{4c^2}). \quad (17)$$

Because  $\varepsilon < 1, c \geq \frac{3}{2}$ , we have:

$$c^2 - \varepsilon + \frac{\varepsilon^2}{4c^2} > c^2 - \frac{8}{9} > 2 \ln \frac{1}{\sqrt{2\pi\delta}} \quad (18)$$

$$c^2 > \ln \frac{2}{\pi} e^{\frac{8}{9}} + 2 \ln \frac{1}{\delta}, \therefore \ln \frac{2}{\pi} e^{\frac{8}{9}} > 1.25^2, \therefore c^2 > 2 \ln \frac{1.25}{\delta} \quad (19)$$

In the above equations, let  $\sigma = c \frac{\Delta f}{\varepsilon}$ , so we have  $\sigma > \frac{\sqrt{2 \ln(1.25/\delta)} \Delta f}{\varepsilon}$ . In particular, in the SGD algorithm, Gaussian noise meets the definition of satisfying differential privacy as long as it satisfies  $\sigma \geq c \frac{q \sqrt{T \ln(\frac{1}{\delta})}}{\varepsilon}$ , where  $q$  is the sampling probability and  $T$  is the iteration round. ■

According to [2], the conditions for the discriminator to guarantee differential privacy are given as follows:

$$\sigma_n = 2q \sqrt{n_d \log(\frac{1}{\delta})} / \varepsilon \quad (20)$$

where  $q$  is the sampling probability and  $n_d$  is the number of iterations of the discriminator in each loop.

*Theorem 2:* Equation 20 represents the relationship between the noise level  $\sigma_n$  and the privacy level  $\varepsilon$ . When we give a fixed perturbation  $\sigma_n$  on the gradient, according to Equation 20, we know that the larger the  $q$  is, the  $D$  gets the fewer privacy guarantee. Because the  $D$  calculates more data, the privacy that can be allocated on each data point is limited. In addition, due to the data provides more information, more iterations ( $n_d$ ) will result in fewer privacy guarantees. The facts described above require us to be cautious when choosing parameters to achieve a reasonable level of privacy.

PPGAN modifies the GAN framework to keep differentially private while relying on Theorem 1,2 and Lemma 1 to change the differentially private  $G$  to train the differentially private  $D$ .

### III. EXPERIMENTS

#### A. Data preprocessing

First, we only use the extracted ICD9 code (The ICD9 code represents the type of disease, and the range of coding is  $C \in [1, 1071]$ .) and use the first three digits for encoding. We then record the patient's admission to the disease and turn it into a vector  $x$ . For example, patient  $P$  was diagnosed with three diseases at admission, and the disease codes are indicated by 9, 42, 146, respectively. (So the ICD9 code consists of 9, 42 and 146.) We use the vector  $x$  to indicate the patient's access record, where the vector is at position 9, the 42nd and 146th bits are set to 1, and the rest are set to 0. Then we aggregate the patient's longitudinal record into a single fixed-size vector  $x \in \mathbb{Z}^+$ , where  $|C| = 1071$  for dataset.

#### B. Relationship between Privacy budget and Generation Performance

In this section, we mainly explore the relationship between privacy budget and generation performance. Considering the combined properties data of *Gaussian noise*, we add Gaussian noise in the process of stochastic gradient descent. Different Gaussian noises can produce different levels of privacy. We input the same set of MNIST image datasets and observe the output generated samples. In the experiments,  $\alpha_d = 5.0 \times 10^{-5}$  learning rate of discriminator;  $\alpha_g = 5.0 \times 10^{-5}$ , learning rate of generator; moments accountant parameter  $C = 1.0 \times 10^{-2}$ ; noise scale  $\delta = 1.0 \times 10^{-5}$ , and the number of iterations on discriminator  $t_d$  and generator  $t_g$  are 5 and  $5.0 \times 10^5$ , respectively. The experimental results are shown in Fig. 2. As shown in Fig. 2, as the privacy budget increases, the quality of the generated images is getting worse. We add well-designed noise that disturbs the data point distribution of the image. Since the noise is randomly added, the distribution of disturbing data points is not fixed, thus ensuring differential privacy.

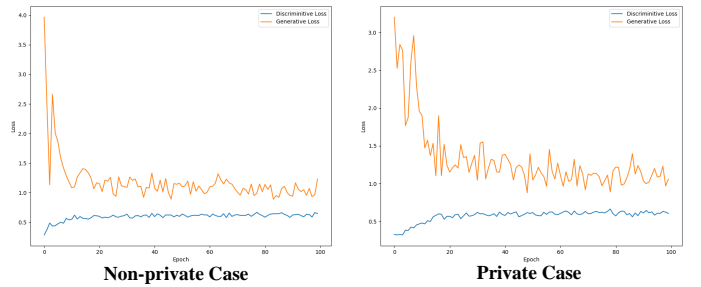


Fig. 3. Loss of Non-private Case ( $\varepsilon = \infty$ ) and Private Case ( $\varepsilon \neq \infty$ ).

Next, we will focus on the impact of noise on PPGAN's loss function during training. The results are shown in Fig. 3. In the non-private case, we observe the training loss of the first 100 epoch in training. The result indicates that the loss of GAN is smooth and stable, and no large fluctuations exist in this round of training. When the loss of the PPGAN with noise starts to fluctuate at the tail of the curve, PPGAN can still converge. As can be inspected from Fig. 3, the convergence rate of PPGAN is acceptable as the compromise of the introduced privacy preservation capability.

#### C. Relationship between Privacy budget and High-quality Datasets

In this section, we quantitatively evaluate the performance of PPGAN. Specifically, we first compare generated data with real data based on statistical characteristics. We propose a *Generate score* to measure the quality of data generated by GAN. We proposed *Generate score* ( $GS(P_g)$ ) to measure the quality of data generated by PPGAN, which can be formally defined as follows for  $P_g$ :




Fig. 2. Four different  $\epsilon$  values are generated for four different quality pictures on MNIST dataset. ( $\epsilon = \infty, \epsilon = 20, \epsilon = 10, \epsilon = 5; \delta = 1.0 \times 10^{-5}$ )

**Definition 7: (Generate scores):**

$$IS(P_g) = e^{E_{x \sim P_g}[KL(PM(y|x)||PM(y))]} \quad (21)$$

$$GS(P_g) = \left| \frac{IS(P_g) - \text{mean}(IS(P_g))}{\max(IS(P_g)) - \min(IS(P_g))} \right|$$

where  $IS(P_g)$  is Inception score which is measure of the performance of the GAN.

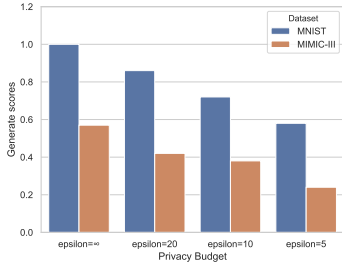


Fig. 4. Generate scores of generative data on MNIST.

The experimental result is shown in Fig. 4. The generated data's (generated by PPGAN) generate score is compared to the real data of the MNIST dataset with different privacy budgets. The larger the score value, the better the quality of the data generated by the generator. The figure shows the distribution of the generate scores of PPGAN in the case of  $\epsilon = 20, 10, 5$ . It can be seen from the figure that the score is very close to the real data generated by the WGAN (non-private case,  $\epsilon = \infty$ ). When  $\epsilon = 20$ , the PPGAN generate score is only 0.14 different from the WGAN generate score, which indicates that the PPGAN generation quality is close to the WGAN.

To evaluate the performance of PPGAN, we compare three solutions, namely dp-GAN [6], DPGAN [2] and WGAN [13] (Non-private Case) in terms of the quality of the generated data. As can be seen from Fig. 5, the data quality generated by PPGAN is better than dp-GAN and DPGAN.

#### IV. CONCLUSION

In this paper, we propose the PPGAN model that preserves the privacy of training data in a differentially private case. PPGAN mitigates information leakage by adding well-designed noise to the gradient during the learning process. We conducted two experiments to show that the proposed algorithm can converge under the noise and constraints of the training data and generate high-quality data. Also, our experimental results verify that PPGAN does not suffer from mode collapse

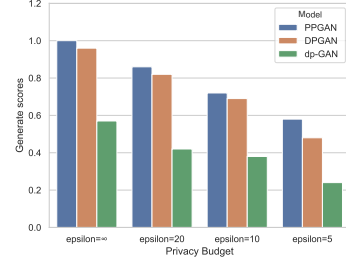


Fig. 5. Generate scores of generative data on model PPGAN, DPGAN and dp-GAN. ( $\delta = 1.0 \times 10^{-5}$ )

or gradient disappearance during training, thus maintaining excellent stability and scalability of model training.

#### ACKNOWLEDGMENTS

This work is supported by the Ministry of Education of China and the School of Entrepreneurship Education of Heilongjiang University (Grant NO.201910212133) and Heilongjiang Provincial Natural Science Foundation of China (Grant NO.QC2016091).

#### REFERENCES

- [1] P. Voigt and A. V. D. Bussche, *The EU General Data Protection Regulation (GDPR)*, 2017.
- [2] L. Xie, K. Lin, S. Wang, F. Wang, and J. Zhou, "Differentially private generative adversarial network," *arXiv preprint arXiv:1802.06739*, 2018.
- [3] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," 2018.
- [4] J. Kang, Z. Xiong, D. Niyato, Y. Zou, Y. Zhang, and M. Guizani, "Reliable federated learning for mobile networks," *IEEE Wireless Communications Magazine*, to be published, Oct. 2019.
- [5] C. Dwork, "Differential privacy," in *International Colloquium on Automata, Languages, & Programming*, 2006.
- [6] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein gan," 2017.
- [7] B. Hitaj, G. Ateniese, and F. Perez-Cruz, "Deep models under the gan: Information leakage from collaborative deep learning," 2017.
- [8] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," 2016.
- [9] J. Dong, A. Roth, and W. J. Su, "Gaussian differential privacy," 2019.
- [10] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, X. Bing, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *International Conference on Neural Information Processing Systems*, 2014.
- [11] X. Yuan, X. Wang, C. Wang, J. Weng, and K. Ren, "Enabling secure and fast indexing for privacy-assured healthcare monitoring via compressive sensing," *IEEE Transactions on Multimedia (TMM)*, vol. 18, no. 10, pp. 1–13, 2016.
- [12] J. Jordan, J. Yoon, and M. van der Schaar, "Pate-gan: generating synthetic data with differential privacy guarantees," 2018.
- [13] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, "Improved training of wasserstein gans," 2017.