PAC PRIVACY PRESERVING DIFFUSION MODELS

Qipan Xu **Rutgers University** qx67@scarletmail.rutgers.edu **Youlong Ding** Hebrew University of Jerusalem dingyoulon@gmail.com

Rutgers University xinxi.zhang@rutgers.edu

Hao Wang* **Rutgers University** hw488@cs.rutgers.edu

Xinxi Zhang

Jie Gao

Rutgers University jg1555@cs.rutgers.edu

ABSTRACT

Data privacy protection is garnering increased attention among researchers. Diffusion models (DMs), particularly with strict differential privacy, can potentially produce images with both high privacy and visual quality. However, challenges arise such as in ensuring robust protection in privatizing specific data attributes, areas where current models often fall short. To address these challenges, we introduce the PAC Privacy Preserving Diffusion Model, a model leverages diffusion principles and ensure Probably Approximately Correct (PAC) privacy. We enhance privacy protection by integrating a private classifier guidance into the Langevin Sampling Process. Additionally, recognizing the gap in measuring the privacy of models, we have developed a novel metric to gauge privacy levels. Our model, assessed with this new metric and supported by Gaussian matrix computations for the PAC bound, has shown superior performance in privacy protection over existing leading private generative models according to benchmark tests. Our code is available at https://github.com/XUQIPAN/P3DM.

1 INTRODUCTION

Modern deep learning models, fortified with differential privacy as defined by Dwork et al. (2006), have been instrumental in significantly preserving the privacy of sensitive data (Dwork et al., 2014). DP-SGD Abadi et al. (2016), a pioneering method for training deep neural networks within the differential privacy framework, applies gradient clipping at each step of the SGD (stochastic gradient descent) to enhance privacy protection effectively.

Diffusion models (DMs) (Song & Ermon, 2019; 2020; Dhariwal & Nichol, 2021) have emerged as state-of-the-art generative models, setting new benchmarks in various applications, particularly in generating high-quality images. When trained under strict differential privacy protocols, these DMs can produce images that safeguard privacy while maintaining high visual fidelity. For instance, DPGEN (Chen et al., 2022) leverages a randomized response technique to privatize the recovery direction in the Langevin MCMC process for image generation. The images produced by DPGEN are not only visually appealing but also compliant with differential privacy standards, although its privacy mechanism has been shown to be data-dependent later on. Moreover, the Differentially Private Diffusion Models (DPDM) (Dockhorn et al., 2022) adapt DP-SGD and introduce noise multiplicity in the training process of diffusion models, demonstrating that DPDM can indeed produce high-utility images while strictly adhering to the principles of differential privacy.

While diffusion models integrated with differential privacy (DP) mark a significant advance in privacy-preserving generative modeling, several challenges and limitations remain.

^{*}Corresponding Author.

- Training modern deep learning models with differential privacy is notoriously difficult, as evidenced by Tramer & Boneh (2021); He et al. (2023); Ding et al. (2024). Recent progress relies heavily on the use of large pre-trained models (Li et al., 2022; Golatkar et al., 2022; Yu et al., 2022; Li et al., 2022) or extensive public training data (De et al., 2022). We believe that such a paradigm is unlikely to be a holy grail solution, as it does not provide a rigorous privacy guarantee but is more heuristic in nature. We refer the reader to the recent criticism against this paradigm (Tramèr et al., 2024).
- Most research on diffusion models with DP has concentrated on the privatization of overall image features. The need to privatize specific attributes, such as facial expressions in human portraits, has not been adequately addressed. This oversight suggests a gap in the nuanced application of DP in generative modeling.
- The absence of a robust privacy measurement for models poses a critical challenge. Without a clear metric, it becomes problematic to assess and compare the data privacy protection performance across different models. This lack of standardized evaluation complicates the advancement and adoption of privacy-preserving techniques in the field.

These issues highlight the need for continued research and development to overcome the current limitations of diffusion models with DP and to push the boundaries of privacy protection in generative modeling.

Recently, Xiao & Devadas (2023) introduces a novel definition of privacy known as Probably Approximately Correct (PAC) Privacy, representing a significant evolution in privacy-preserving methodologies. PAC Privacy characterizes the information-theoretic hardness to recover sensitive data given arbitrary information disclosure/leakage during/after any processing. Compared with differential privacy, it has the following advantages.

- Differential privacy requires bounded sensitivity, which cannot be tightly computed for modern machine learning models. Artificial modifications are typically required to decompose algorithms into multiple simpler and analyzable components, e.g., gradient clipping in DP-SGD. The exact privacy analysis of DP can be, in general, NP-hard. Especially for for deep learning, tight DP analysis are challenging. However, PAC privacy can be applied to any data processing algorithm (as a black-box procedure), where security parameters can be produced with arbitrarily high confidence via Monte-Carlo simulation.
- On the utility side, the magnitude of (necessary) perturbation required in PAC Privacy is not lower bounded by $\Theta(\sqrt{d})$ for a *d*-dimensional release, but could be O(1) (depending on data distribution) for many practical data processing tasks, which is in contrast to the input-independent worst-case information-theoretic lower bound like DP. Therefore, PAC Privacy analysis in many applications can produce much sharpened utility-privacy tradeoffs.

To tackle the aforementioned challenges, we have introduced PAC Privacy Preserving Diffusion Models (P3DM). Drawing from the foundations of DPGEN and harnessing insights from conditional classifier guidance (Dhariwal & Nichol, 2021; Batzolis et al., 2021; Ho & Salimans, 2022), our P3DM incorporates a conditional private attribute guidance module during the Langevin sampling process. This addition empowers the model to specifically target and privatize certain image attributes with greater precision.

Furthermore, we have crafted a set of privacy evaluation metrics. These metrics operate by measuring the output class labels of the two nearest neighbor images in the feature space of the Inception V3 model (Szegedy et al., 2016), using a pretrained classifier. Additionally, we quantify the noise addition B necessary to assure PAC privacy in our model and conduct comparative analyses against the mean L2-norm of B from various other models.

Through meticulous evaluations that utilize our privacy metrics and benchmarks for noise addition, our model has proven to offer a superior degree of privacy. It exceeds the capabilities of state-of-the-art (SOTA) models in this critical aspect, while simultaneously preserving the high quality of synthetic image samples. These samples remain on par with those produced by the state-of-the-art models, illustrating that our model does not sacrifice quality for privacy. This achievement underscores our model's potential to set new precedents in the domain of privacy-preserving image generation and data protection at large.

Our contributions are summarized as follows:

- We propose the first diffusion model with analysis on its PAC privacy guarantees.
- We incorporate conditional private classifier guidance into the Langevin Sampling Process, enhancing the protection of privacy for specific attributes in images.
- We introduce a new metric that we developed for assessing the extent of privacy provided by models.
- We compute the noise addition matrix to establish the Probably Approximately Correct (PAC) upper bound and have conducted a comparative analysis of the norm of this matrix across various models.
- Through extensive evaluations, we demonstrate that our model sets a new standard in privacy protection of specific attributes, achieving state-of-the-art (SOTA) results, while maintaining image quality at a level that is comparable to other SOTA models.

2 RELATED WORKS

2.1 EARLY WORKS ON DIFFERENTIALLY PRIVATE IMAGE GENERATION

Image Synthesis with differential privacy has been studied extensively during the research. Recently, there has been an increased emphasis on utilizing sophisticated generative models to improve the quality of differentially private synthetic data (Hu et al., 2023). Some approaches employ Generative Adversarial Networks (GANs) (Goodfellow et al., 2014), or GANs that have been trained using the Private Aggregation of Teacher Ensembles (PATE) framework (Xie et al., 2018; Chen et al., 2020; Harder et al., 2021; Torkzadehmahani et al., 2019). Other contributions leverage variational autoencoders(VAEs) (Pfitzner & Arnrich, 2022; Jiang et al., 2022; Takagi et al., 2021), or take advantage of customized architectures (Cao et al., 2021; Harder et al., 2023). However, there are several limitations for those DP synthesizers: (1) Failure when applying to high-dimensional data, primarily due to the constraints imposed by discretization. (2) Limited image quality and lack of expressive generator networks (Cao et al., 2021).

2.2 DIFFERENTIALLY PRIVATE DIFFUSION MODELS

Diffusion models (DMs) (Song & Ermon, 2019; 2020; Dhariwal & Nichol, 2021), recognized for setting new standards in image generation, can produce high-quality, privacy-compliant images when trained with differential privacy protocols. For example, DPGEN (Chen et al., 2022) employs a data-dependant randomized response method to privatize the recovery direction in the Langevin MCMC process for image generation. Furthermore, Differentially Private Diffusion Models (DPDM) (Dockhorn et al., 2022), which adapt DP-SGD and introduce noise multiplicity, both demonstrate the feasibility of generating visually appealing, privacy-protective images. Subsequent advancements, including fine-tuning existing models and employing novel diffusion model architectures, have been made to boost the effectiveness of differentially private image generation, as detailed in Ghalebikesabi et al. (2023); Lyu et al. (2023). Nevertheless, as previously noted in the introduction, there remains three key challenges to be addressed. In the following sections, we propose solutions and methods to tackle these issues.

3 Methods

We introduce the PAC Privacy Preserving Diffusion Model (P3DM), which aims to safeguard privacy for specific attributes. Our method is inspired by DPGEN (Chen et al., 2022). While DPGEN asserts compliance with stringent ϵ -differential privacy, recent findings from Dockhorn et al. (2022) indicate that DPGEN is, in fact, data-dependent. The RR mechanism M(d) in DPGEN holds validity only within the perturbed dataset. Specifically, if an element z belongs to the output set O but not to the perturbed dataset d, then Pr[M(d) = O] = 0, which contravenes the differential privacy definition. Different from DPGEN, which turns out to have no formal privacy guarantees, our method satisfies PAC privacy.

Algorithm 1 PAC-Private Adapted Randomized Response Algorithm

Require: a sensitive dataset $\{x_i : i = 1, 2, ..., m\}_{i=1}^m$; sample number k

- Sampling k image candidates from {x_i}^m_{i=1}, with the sampling probability for each image Eqn. 2, to construct set X ← {x_j : max(x_i − x_j)/σ_j ≤ β, j ∈ [m]}
- 2: Privatize images in X with RR from Eqn. 1 and obtain $H(x_i)$
- 3: $\tilde{x}_i = H(x_i) + B; B \sim \mathcal{N}(0, \Sigma_B)$
- 4: **Return** \tilde{x}_i

Given a sensitive dataset $\{x_i : i = 1, 2, ..., m\}_{i=1}^m$, we first sample an image by $\tilde{x}_i \leftarrow \mathcal{N}(x_i, \sigma^2 I)$. Next, we utilizes the random response method as follows:

$$Pr[H(\tilde{x}_i) = w] = \begin{cases} \frac{e^{\epsilon}}{e^{\epsilon} + k - 1}, & w = x_i \\ \frac{1}{e^{\epsilon} + k - 1}, & w = x'_i \in X \setminus x_i \end{cases}$$
(1)

In the equation, $X = \{x_j : max(\tilde{x}_i - x_j)/\sigma_j \le \beta, j \in [m]\}$ (where "max" is over the dimensions of $\tilde{x}_i - x_j$), $|X| = k \ge 2$, where the hyperparameter k denotes the number of selected candidates from m training images, and the sampling probability for each image is given by

$$p(x_i) = \exp(-d_{\infty}(x_i, \tilde{x}_i, \sigma_i)) / \sum_{a=1}^{m} \exp(-d_{\infty}(x_i, \tilde{x}_i, \sigma_a))$$
(2)

In other words, the mechanism $H(\cdot)$ consists of 2 steps:

- 1. Sampling k image candidates from $\{x_i\}_{i=1}^m$, with the sampling probability for each image from Eqn. 2, to construct set X.
- 2. Privatize images in X with RR from Eqn. 1.

After applying $H(\cdot)$, we then add Gaussian noise B to the image to get $\tilde{x}_i \leftarrow H(x_i) + B$ for achieving PAC privacy. With these methods, \tilde{x}_i outputs one of its k nearest neighbors with certain probability $Pr[H(\tilde{x}_i) = x_i^r]$, as in Chen et al. (2022), giving us the privatized the recovery direction $d_i^r = (\tilde{x}_i - x_i^r)/\sigma_i^2$.

Following perturbation and privatization, we learn an energy function $q_{\theta}(\tilde{x})$ by optimizing the following loss function:

$$l(\theta, \sigma) = \frac{1}{2} \mathbb{E}_{p(x)} \mathbb{E}_{\tilde{x} \sim \mathcal{N}(x, \sigma^2)} \left[\left\| d - \nabla_x \log q_\theta(\tilde{x}) \right\|^2 \right]$$
(3)

After getting the optimal $\nabla_x \log q_\theta(\tilde{x}) = (\tilde{x} - x^r)/\sigma^2$ from Eqn. 3, we can then synthesize images using the Langevin MCMC sampler with the optimal output $\nabla_x \log q_\theta(\tilde{x})$ as:

$$x_t \leftarrow x_{t-1} + \frac{\alpha_i}{2} q_\theta(x_{t-1}) + \sqrt{\alpha_i} z_t, t = 0, 1, 2, \dots T.$$
(4)

where α_i denotes the step size.

3.1 CONDITIONAL PRIVATE LANGEVIN SAMPLING

We introduce a method of conditional private guidance within the Langevin sampling algorithm, detailed in Algorithm 2. This method is designed to protect specific attributes within the original datasets against adversarial attacks.

Prior to commencing the sampling iterations, it is essential to obtain class labels from a balanced attribute of the dataset. The necessity for a dataset attribute that possesses an approximately equal number of negative and positive samples is crucial for training classifiers to achieve exceptional performance. Subsequently, we select a random label y_i from y and fix it to initialize x_0 with a predetermined distribution, such as the standard normal distribution.

Algorithm 2 PAC-Private	Conditional	Guidance I	Langevin D	ynamics Sampling
-------------------------	-------------	------------	------------	------------------

Require: class labels $y = \{y_1, y_2, ..., y_n\}$ from one of balanced dataset attributes; gradient scale k; $\{\delta_i\}_{i=1}^L, \epsilon, T;$ optimal output q_{θ} from Eqn. 3; pretrained classifier model on noisy images c_{θ} 1: Randomly sample y_n from y2: Initialize x_0 3: **for** *i* from 1 to *L* **do** 4: stepsize $\alpha_i \leftarrow \delta_i^2 / \delta_I^2$ 5: for t from 1 to T do Sample noise $z_t \sim \mathcal{N}(0, I)$ 6: $x_{t} \leftarrow x_{t-1} + \frac{\alpha_{i}}{2} q_{\theta}(x_{t-1}, \delta_{i}) + \sqrt{\alpha_{i}} z_{t} + k \Sigma_{\theta}(x_{t-1}) \nabla_{x_{t-1}} \log c_{\theta}(y_{n} | x_{t-1})$ 7: 8: end for 9: $x_0 \leftarrow x_T$ 10: **end for** 11: Return x_T

Drawing on conditional image generation (Dhariwal & Nichol, 2021; Batzolis et al., 2021; Ho & Salimans, 2022), we adapt the model from the vanilla Langevin dynamic samplings with the selected attributes with conditional guidance

$$k\Sigma_{\theta}(x_{t-1})\nabla_{x_{t-1}}\log c_{\theta}(y_n|x_{t-1}) \tag{5}$$

where k is the gradient scale that can be tuned according to the performance of the model, $\Sigma_{\theta}(x_{t-1})$ is the covariance matrix from reverse process Eqn. 8. In Eqn. 5, the term $\nabla_{x_{t-1}} \log c_{\theta}(y_n | x_{t-1})$ directs the Langevin sampling process toward a specific class label y_n , which is sampled prior to the inference.

This private guidance during the inference phase ensures that the synthesized images are protected from privacy breaches related to designated attributes. At its core, this method involves intentionally modifying certain generated trajectories by randomly perturbing the image label y_n , which is then used by a pretrained classifier to guide the image generation process. This strategy is aimed at diverting certain image attributes that we wish to protect. For instance, consider a scenario where an original dataset image depicting Celebrity A wearing eyewear, a known trait of the celebrity. If our goal is to privatize the attribute of "wearing glasses," private classifier guidance can effectively achieve this. When sampling occurs under the influence of this guidance, the attributes are likely to vary, creating a chance that the resulting synthesized image based on Celebrity A might be rendered without glasses. This modification effectively protects the attribute of 'glasses' from being a consistent element in the generated depictions of Celebrity A. Consequently, the images generated with this approach offer a higher degree of privacy compared to those produced by the original DPGEN method.

3.2 PRIVACY METRICS

Most privacy-preserving generative models prioritize assessing the utility of images for downstream tasks, yet often overlook the crucial metric of the models' own privacy. To bridge this gap in evaluating privacy extent, we have developed a novel algorithm that computes a privacy score for the models obtained. As per Algorithm 3, we commence by preparing images x_i generated from Algorithm 2, the original dataset images x_k^G , and alongside the classifier model which trained separately with clean images (different from the classifier from Algorithm 2). Subsequently, we process all synthesized images through InceptionV3 (Szegedy et al., 2016).

After obtaining the feature vector output from InceptionV3, we locate the feature vector of a ground truth image that has the smallest L2 distance to that of the synthesized image, effectively finding the nearest ground truth neighbor. We then test whether a pretrained classifier model can differentiate these two images. Inability of the classifier to distinguish between the two indicates that the specific attributes, even in images most similar to the original, are well-protected. Thus, our method successfully preserves privacy for specific attributes. Finally, we compute the average probability of incorrect classification by the pretrained model to establish our privacy score. The greater the privacy score a model achieves, the more robust its privacy protection is deemed to be.

Algorithm 3 Privacy Score

Require: images x_i generated by algorithm 2; ground truth images x_{k}^{G} ; sample number n; pretrained InceptionV3 model I_{θ} ; pretrained classifier model c_{θ} . 1: private score $s \leftarrow 0$ 2: **for** x_i from x_0 to x_n **do** find $argmin_k ||I_{\theta}(x_i), I_{\theta}(x_k^G)||_2$ 3: 4: if $c_{\theta}(x_i) \neq c_{\theta}(x_k^G)$ then $s \leftarrow s + 1$ 5: end if 6: 7: end for 8: **Return** s/n



Figure 1: CelebA images generated from DPDM, DPGEN and our model from left to right with image resolution 64×64 .

4 EXPERIMENTS

In this section, we compare P3DM with existing methods.



Figure 2: **Privacy score and FID curve of all datapoints from different models.** Each datapoint in the figure consists of mean and standard deviation from 3 experimental results with the same epsilon and different random seeds. Top-right corner is preferred. The curve from our method, pushes the frontier to the upper-right over DPGEN (Chen et al., 2022), DPDM (Dockhorn et al., 2022) and DP-MEPF (Harder et al., 2023). For more details related to data points in this figure, please refer to Appendix D.

4.1 DATASETS

Our experiments were carried out using the CelebA (Liu et al., 2015) datasets. We specifically targeted attributes that have a balanced distribution of positive and negative samples, as noted in Rudd et al. (2016), to facilitate the training of classifiers. Consequently, we selected attributes like gender and smile from the CelebA dataset (referred to as CelebA-gender and CelebA-smile, respectively) for sampling in Algorithm 2. All the images used in these experiments were of the resolution 64×64 .

4.2 **BASELINES**

In our study, we consider DPGEN (Chen et al., 2022), DPDM (Dockhorn et al., 2022) and DP-MEPF (Harder et al., 2023) as baseline methods. These approaches, except for DPGEN, which is not a rigorous DP, excel in synthesizing images under differential privacy (DP) constraints, and they stand out for their exceptional sample quality in comparison to other DP generative models. These models serve as important benchmarks against which we evaluate the performance and efficacy of our proposed method.

4.3 EVALUATION METRICS

In our evaluation process, we validate the capability of our PAC Privacy Preserving Diffusion Model to generate high-resolution images using the key metric: the Frechet Inception Distance (Heusel et al., 2017). This metric is widely recognized and utilized in the field of generative models to assess the visual fidelity of the images they produce.

Additionally, to demonstrate our model's effectiveness in preventing privacy leakage, we employ our unique privacy metrics as outlined in Algorithm 3. We compare the mean norm of the Gaussian noise $\mathbb{E}||B||_2$ as detailed in Theorem A.2. For our experiments, we have chosen the hyperparameters $\nu = \beta = 0.5$ and $\gamma = 0.01$. This approach allows us to comprehensively assess not just the quality of the images generated, but also the strength of privacy protection our model offers. It is important to note that hyperparameter tuning does incur an additional privacy cost, albeit modest (Liu & Talwar, 2019; Papernot & Steinke, 2022; Ding & Wu, 2023; Xiang et al., 2024). For simplicity we ignore the privacy overhead due to hyperparameter tuning for all methods.

4.4 Empirical Results And Analysis

It is important to note that the ξ and τ presented within the tables below is merely a hyperparameter derived from the Randomized Response (RR) as indicated in Eqn. 1. This ξ and τ should not be confused with the ε from ε -Differential Privacy (DP), where ξ indicates the model is PAC private, and τ indicates the model is neither PAC private nor DP. We will later illustrate how our model assures privacy through the automatic control of mutual information for a PAC privacy guarantee in this section.

Fig. 2 details the evaluation results of image visual quality and privacy score on the CelebA dataset with a resolution of 64×64 , where each datapoint in the figure, or each entry in the table, consists of mean and standard deviation from 3 experimental results with the same epsilon and different random seeds. By examining the table, we can see that our model achieves image quality comparable to the state-of-the-art model DPGEN (Chen et al., 2022), a conclusion also supported by Fig. 1.

Additionally, from Fig. 2, our model registers the highest performance in privacy score when having similar FIDs, signaling an enhancement in our model's ability to preserve privacy without substantially impacting image utility. The assertion is further supported by Fig. 3, wherein the CelebA dataset is filtered based on the 'smile' attribute. Even upon querying the nearest images of P3DM samples, we can distinctly observe that while all P3DM samples exhibit the absence of a smile, the nearest neighbors predominantly display smiling faces. This highlights that, in contrast to other models, the closest images between the generated dataset and the ground truth datasets are notably similar in terms of the 'smile' expression, while the images generated by our model demonstrate distinctive features. Hence, the P3DM model effectively conceals the 'smile' attribute.



(a) DPGEN for $\varepsilon = \infty$



(b) DPDM for $\varepsilon = 10$



(c) DPGEN for $\varepsilon = 10$



(d) P3DM-Smile (Ours) for $\xi = 10$

Figure 3: Generated images (the second row) and their nearest neighbors measured by the l_2 distance between images from CelebA-smile dataset (the first row), with image resolution 64×64 .

IC I. Estimated I		ise D to apply	UNIMALCITY C	lisuic TAC pil
Methods	DP	PAC Privacy	Heuristic	$\mathbb{E} B _2\downarrow$
P3DM-Gender		$\xi = 5$		283.03±2.25
P3DM-Smile		$\dot{\xi} = 5$		$280.60{\pm}2.57$
DPGEN			$\tau = 5$	281.3 ± 2.78
DP-MEPF	$\varepsilon = 5$			$325.6 {\pm} 3.62$
P3DM-Gender		$\xi = 10$		328.80±1.24
P3DM-Smile		$\xi = 10$		327.98±1.45
DPGEN			$\tau = 10$	$329.48{\pm}1.6$
DPDM	$\varepsilon = 10$			$335.83 {\pm} 1.21$
DP-MEPF	$\varepsilon = 10$			$330.75 {\pm} 1.33$
P3DM-Gender		$\xi = \infty$		332.87±1.15
P3DM-Smile		$\tilde{\xi} = \infty$		331.67±1.3
DPGEN			$\tau = \infty$	332.02 ± 1.32
DPDM	$\varepsilon = \infty$			$335.83{\pm}1.56$
DP-MEPF	$\varepsilon = \infty$			333.48±1.79

Table 1: Estimated level of noise B to approximately ensure PAC privacy

Furthermore, in Table 1, we compute the multivariate Gaussian matrix B with a $1 - \gamma$ noise determination for the deterministic mechanism M followed the work by Xiao et.al Xiao & Devadas (2023). From this, we show that under the same confidence level of $1 - \gamma = 0.99$ to ensure $MI(X; M(X) + B) \leq 1$, the PAC Privacy Preserving Diffusion Model exhibits the smallest norm on $\mathbb{E}||B||_2$. This demonstrates that data processed by our model retains the least mutual information with the original dataset, thus affirming our model's superior performance in privacy protection.

5 CONCLUSION

We introduce the PAC Privacy Preserving Diffusion Model (P3DM), which incorporates conditional private classifier guidance into the Langevin Sampling process to selectively privatize image features. In addition, we have developed and implemented a unique metric for evaluating privacy. This metric involves comparing a generated image with its nearest counterpart in the dataset to assess

whether a pretrained classifier can differentiate between the two. Furthermore, we calculate the necessary additional noise B to ensure PAC privacy and benchmark the noise magnitude against other models. Our thorough empirical and theoretical testing confirms that our model surpasses current state-of-the-art private generative models in terms of privacy protection while maintaining comparable image quality.

REFERENCES

- Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016.
- Georgios Batzolis, Jan Stanczuk, Carola-Bibiane Schönlieb, and Christian Etmann. Conditional image generation with score-based diffusion models. *arXiv preprint arXiv:2111.13606*, 2021.
- Tianshi Cao, Alex Bie, Arash Vahdat, Sanja Fidler, and Karsten Kreis. Don't generate me: Training differentially private generative models with sinkhorn divergence. Advances in Neural Information Processing Systems, 34:12480–12492, 2021.
- Dingfan Chen, Tribhuvanesh Orekondy, and Mario Fritz. Gs-wgan: A gradient-sanitized approach for learning differentially private generators. Advances in Neural Information Processing Systems, 33:12673–12684, 2020.
- Jia-Wei Chen, Chia-Mu Yu, Ching-Chia Kao, Tzai-Wei Pang, and Chun-Shien Lu. Dpgen: Differentially private generative energy-guided network for natural image synthesis. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8387–8396, 2022.
- Soham De, Leonard Berrada, Jamie Hayes, Samuel L Smith, and Borja Balle. Unlocking high-accuracy differentially private image classification through scale. *arXiv preprint arXiv:2204.13650*, 2022.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. Advances in neural information processing systems, 34:8780–8794, 2021.
- Youlong Ding and Xueyang Wu. Revisiting hyperparameter tuning with differential privacy, 2023. URL https://arxiv.org/abs/2211.01852.
- Youlong Ding, Xueyang Wu, Yining Meng, Yonggang Luo, Hao Wang, and Weike Pan. Delving into differentially private Transformer. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pp. 11049–11071, 2024.
- Tim Dockhorn, Tianshi Cao, Arash Vahdat, and Karsten Kreis. Differentially Private Diffusion Models. *arXiv:2210.09929*, 2022.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pp. 265–284. Springer, 2006.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations* and *Trends*® in *Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Sahra Ghalebikesabi, Leonard Berrada, Sven Gowal, Ira Ktena, Robert Stanforth, Jamie Hayes, Soham De, Samuel L Smith, Olivia Wiles, and Borja Balle. Differentially private diffusion models generate useful synthetic images. arXiv preprint arXiv:2302.13861, 2023.
- Aditya Golatkar, Alessandro Achille, Yu-Xiang Wang, Aaron Roth, Michael Kearns, and Stefano Soatto. Mixed differential privacy in computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8376–8386, 2022.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. Advances in neural information processing systems, 27, 2014.

- Frederik Harder, Kamil Adamczewski, and Mijung Park. Dp-merf: Differentially private mean embeddings with randomfeatures for practical privacy-preserving data generation. In *International conference on artificial intelligence and statistics*, pp. 1819–1827. PMLR, 2021.
- Frederik Harder, Milad Jalali, Danica J Sutherland, and Mijung Park. Pre-trained perceptual features improve differentially private image generation. *Transactions on Machine Learning Research*, 2023.
- Jiyan He, Xuechen Li, Da Yu, Huishuai Zhang, Janardhan Kulkarni, Yin Tat Lee, Arturs Backurs, Nenghai Yu, and Jiang Bian. Exploring the limits of differentially private deep learning with group-wise clipping. In *International Conference on Learning Representations*, 2023.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598, 2022.
- Yuzheng Hu, Fan Wu, Qinbin Li, Yunhui Long, Gonzalo Munilla Garrido, Chang Ge, Bolin Ding, David Forsyth, Bo Li, and Dawn Song. Sok: Privacy-preserving data synthesis. arXiv preprint arXiv:2307.02106, 2023.
- Dihong Jiang, Guojun Zhang, Mahdi Karami, Xi Chen, Yunfeng Shao, and Yaoliang Yu Dp2vae. Differentially private pre-trained variational autoencoders. *arXiv preprint arXiv:2208.03409*, 2022.
- Xuechen Li, Florian Tramer, Percy Liang, and Tatsunori Hashimoto. Large language models can be strong differentially private learners. In *International Conference on Learning Representations*, 2022.
- Jingcheng Liu and Kunal Talwar. Private selection from private candidates. In *Proceedings of the* 51st ACM Symposium on Theory of Computing, pp. 298–309, 2019.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Saiyue Lyu, Margarita Vinaroz, Michael F Liu, and Mijung Park. Differentially private latent diffusion models. arXiv preprint arXiv:2305.15759, 2023.
- Nicolas Papernot and Thomas Steinke. Hyperparameter tuning with renyi differential privacy. In *International Conference on Learning Representations*, 2022.
- Bjarne Pfitzner and Bert Arnrich. Dpd-fvae: Synthetic data generation using federated variational autoencoders with differentially-private decoder. *arXiv preprint arXiv:2211.11591*, 2022.
- Ethan M Rudd, Manuel Günther, and Terrance E Boult. Moon: A mixed objective optimization network for the recognition of facial attributes. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V 14*, pp. 19–35. Springer, 2016.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. Advances in neural information processing systems, 32, 2019.
- Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. Advances in neural information processing systems, 33:12438–12448, 2020.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Shun Takagi, Tsubasa Takahashi, Yang Cao, and Masatoshi Yoshikawa. P3gm: Private highdimensional data release via privacy preserving phased generative model. In 2021 IEEE 37th International Conference on Data Engineering (ICDE), pp. 169–180. IEEE, 2021.

- Reihaneh Torkzadehmahani, Peter Kairouz, and Benedict Paten. Dp-cgan: Differentially private synthetic data and label generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 0–0, 2019.
- Florian Tramer and Dan Boneh. Differentially private learning needs better features (or much more data). In *International Conference on Learning Representations*, 2021.
- Florian Tramèr, Gautam Kamath, and Nicholas Carlini. Position: Considerations for differentially private learning with large-scale public pretraining. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pp. 48453–48467, 2024.
- Zihang Xiang, Tianhao Wang, Chenglong Wang, and Di Wang. Revisiting differentially private hyper-parameter tuning, 2024. URL https://arxiv.org/abs/2402.13087.
- Hanshen Xiao and Srinivas Devadas. Pac privacy: Automatic privacy measurement and control of data processing. In *Annual International Cryptology Conference*, pp. 611–644. Springer, 2023.
- Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. Differentially private generative adversarial network. *arXiv preprint arXiv:1802.06739*, 2018.
- Da Yu, Saurabh Naik, Arturs Backurs, Sivakanth Gopi, Huseyin A Inan, Gautam Kamath, Janardhan Kulkarni, Yin Tat Lee, Andre Manoel, Lukas Wutschitz, et al. Differentially private fine-tuning of language models. In *International Conference on Learning Representations*, 2022.

A APPENDIX

A.1 DIFFERENTIAL PRIVACY

A randomized mechanism \mathcal{M} is said (ε, δ) - differentially private if for any two adjacent datasets D and D' differing in a single datapoint for any subset of outputs S as follows (Dwork et al., 2014):

$$\mathbf{Pr}[\mathcal{M}(D) \in S] \le e^{\varepsilon} \cdot \mathbf{Pr}[\mathcal{M}(D') \in S] + \delta \tag{6}$$

Here, ε is the upper bound on the privacy loss corresponding to \mathcal{M} , and δ is the probability of violating the DP constraint.

Differential privacy is a mathematical approach designed to protect individual privacy within datasets. It offers a robust privacy assurance by enabling data analysis without disclosing sensitive details about any specific person in the dataset.

A.2 CONDITIONAL DIFFUSION MODELS

Dhariwal & Nichol (2021) proposed a diffusion model that is enhanced by classifier guidance; it has been shown to outperform existing generative models. By using true labels of datasets, it is possible to train a classifier on noisy images x_t at various timesteps $p_{\phi}(y|x_t, t)$, and then use this classifier to guide the reverse sampling process $\nabla_{x_t} \log p_{\phi}(y|x_t, t)$. What begins as an unconditional reverse noising process is thus transformed into a conditional one, where the generation is directed to produce specific outcomes based on the given labels

$$p_{\theta,\phi}(x_t|x_{t+1}, y) = Zp_{\theta}(x_t|x_{t+1})p_{\phi}(y|x_t)$$
(7)

Where Z is a normalizing constant. According to unconditional reverse process, which predicts timestep x_t from x_{t-1} leveraging Gaussian distribution, we have

$$p_{\theta}(x_t | x_{t+1}) \sim \mathcal{N}(\mu, \Sigma) \tag{8}$$

$$\log p_{\theta}(x_t | x_{t+1}) = -\frac{1}{2} (x_t - \mu)^T \Sigma^{-1} (x_t - \mu) + C$$
(9)

If we assume that $\log_{\phi} p(y|x_t)$ has low curvature compared to Σ^{-1} , then we can approximate $\log p_{\phi}(y|x_t)$ using a Taylor expansion around $x_t = \mu$ as

$$\log p_{\phi}(y|x_{t}) \approx \log p_{\phi}(y|x_{t})|_{x_{t}=\mu} + (x_{t}-\mu)\nabla_{x_{t}} \log p_{\phi}(y|x_{t})|_{x_{t}=\mu}$$

$$= (x_{t}-\mu)g + C_{1}$$
(10)

where g is the gradient of classifier $g = \nabla_{x_t} \log p_{\phi}(y|x_t)|_{x=\mu}$ and C_1 is the constant.

Therefore, combing Eqn. 9 and 10 gives us

$$\log(p_{\theta}(x_t|x_{t+1})p_{\phi}(y|x_t)) \sim \mathcal{N}(\mu + \Sigma g, \Sigma)$$
(11)

The investigation has led to the conclusion that the conditional transition operator can be closely estimated using a Gaussian. And the Gaussian resembles the unconditional transition operator, with the distinction that its mean is adjusted by the product of the covariance matrix, Σ , and the vector g. This methodology allows for the generation of high-quality, targeted synthetic data.

A.3 PAC PRIVACY

A.3.1 BASIC DEFINITION

Definition A.1. $((\delta, \rho, D) \text{ PAC Privacy})$. As discussed in Xiao & Devadas (2023), for a data processing mechanism \mathcal{M} , given some data distribution D, a measure function $\rho(.,.)$, and a finite set X^* , we say \mathcal{M} satisfies (δ, ρ, D) -PAC Privacy if the following experiment is impossible: A user generates data X from distribution D and sends $\mathcal{M}(\mathcal{X})$ to an adversary. The adversary who knows D and \mathcal{M} is asked to return an estimation $\tilde{X} \in X^*$ on X such that with posterior success probability at least $(1 - \delta)$, $\rho(\tilde{X}, X) = 1$.

Definition A.2. $((\Delta_f \delta, \rho, D) \text{ PAC Advantage Privacy})$ Equivalantly, as discussed in Xiao & Devadas (2023), \mathcal{M} could be defined as $(\Delta_f \delta, \rho, D)$ PAC Advantage Privacy if the posterior advantage measured in f-divergence \mathcal{D}_f satisfies

$$\Delta_f \delta = \mathcal{D}_f(\mathbf{1}_\delta \| \mathbf{1}_{\delta_o^{\rho}}) = \delta_o^{\rho} f(\frac{\delta}{\delta_o^{\rho}}) + (1 - \delta_o^{\rho}) f(\frac{1 - \delta}{1 - \delta_o^{\rho}}), \tag{12}$$

where $(1 - \delta_{\alpha}^{\rho})$ represents the optimal prior success rate of recovering X.

$$1 - \delta_o^{\rho} = \sup_{\tilde{X} \in \mathcal{X}^*} \Pr_{X \sim \mathsf{D}} \left(\rho(\tilde{X}, X) = 1 \right), \tag{13}$$

and $\mathbf{1}_{\delta}$ and $\mathbf{1}_{\delta_{0}^{\rho}}$ represent two Bernoulli distributions of parameters δ and δ_{0}^{ρ} , respectively.

The definition above depicts the reconstruction hardness for the attackers to recover the private data distribution $\mathcal{M}(\mathcal{X})$. With a lower bound probability $(1 - \delta)$, the measure function $\rho(.,.)$ cannot distinguish the recovery data from the original data. However, the limitation of the naive definition is that the prior distribution of the public dataset is unknown, resulting in the failure of adversarial inferences.

Definition A.3. (Mutual Information). For two random variables x and w in some joint distribution, the mutual information MI(x; w) is defined as

$$MI(x,w) = \mathcal{H}(x) - \mathcal{H}(x|w) = \mathcal{D}_{KL}(P_{x,w}||P_x \otimes P_w), \tag{14}$$

where \mathcal{D}_{KL} denotes the KL divergence.

Theorem A.1. As discussed in Xiao & Devadas (2023), for any selected f-divergence \mathcal{D}_f , a mechanism $\mathcal{M} : \mathcal{X}^* \to \mathcal{Y}$ satisfies $(\Delta_f \delta, \rho, \mathsf{D})$ PAC Advantage Privacy if

$$\Delta_f \delta = \mathcal{D}_f \left(\mathbf{1}_{\delta} \| \mathbf{1}_{\delta_o^{\rho}} \right) \le \inf_{\mathsf{P}_W} \mathcal{D}_f \left(\mathsf{P}_{X, \mathcal{M}(X)} \| \mathsf{P}_X \otimes \mathsf{P}_W \right).$$
(15)

In particular, when we select \mathcal{D}_f to be the KL-divergence and $\mathcal{P}_W = \mathcal{P}_{\mathcal{M}(X)}$, \mathcal{M} satisfies $(\Delta_{KL}\delta, \rho, \mathsf{D})$ PAC Advantage Privacy where

$$\Delta_{KL}\delta = \mathcal{D}_{KL}(\mathbf{1}_{\delta} \| \mathbf{1}_{\delta_{o}^{\rho}}) \le MI(X; \mathcal{M}(X)).$$
(16)

In summary, $\mathcal{D}_f(\mathbf{1}_{\delta} \| \mathbf{1}_{\delta_o^{\rho}})$ quantifies the divergence between optimal a priori and posterior reconstruction, effectively measuring the difficulty of inference. A higher value of $\mathcal{D}_f(\mathbf{1}_{\delta} \| \mathbf{1}_{\delta_o^{\rho}})$ signifies greater privacy leakage. Moreover, since $MI(X; \mathcal{M}(X))$ provides an upper bound for \mathcal{D}_f , a lower value of $MI(X; \mathcal{M}(X))$ indicates stronger privacy protection. Thus, theorem A.1 establishes a general method for linking the difficulty of arbitrary inference to the well-known concept of mutual information. With theorem A.1, the goal of PAC privacy is explicit: determining the bound $MI(X; \mathcal{M}(X))$ with high confidence.

A.3.2 NOISE DETERMINATION AND SIMULATED PRIVACY GUARANTEE WITH HIGH CONFIDENCE

The natural idea to achieve information leakage control is perturbation: when $MI(X; \mathcal{M}(X))$ is not small enough to produce satisfied PAC Privacy, we may add additional Gaussian noise B, to produce smaller $MI(X; \mathcal{M}(X) + B)$.

Theorem A.2. (Xiao & Devadas, 2023) When the mutual information MI(X; M(X)) is insufficient to ensure PAC privacy, additional Gaussian noise $B \sim \mathcal{N}(0, \Sigma_B)$ can be introduced to yield a reduced mutual information MI(X; M(X) + B), such that MI(X; M(X) + B) satisfies

$$\mathsf{MI}(X; \mathcal{M}(X) + \mathbf{B}) \le \frac{1}{2} \cdot \log \det \left(\mathbf{I}_d + \Sigma_{\mathcal{M}(X)} \cdot \Sigma_{\mathbf{B}}^{-1} \right).$$
(17)

In particular, let the eigenvalues of $\Sigma_{\mathcal{M}(\mathbf{X})}$ be $(\lambda_1, \ldots, \lambda_d)$, then there exists some Σ_B such that $\mathbb{E}[\|\mathbf{B}\|_{2}^2] = (\sum_{j=1}^d \sqrt{\lambda_j})^2$, and $MI(\mathbf{X}; \mathcal{M}(\mathbf{X}) + \mathbf{B}) \leq 1/2$.

Therefore, we have a simple upper bound on the mutual information after perturbation which only requires the knowledge of the covariance of $\mathcal{M}(X)$. Another important and appealing property is that the noise calibrated to ensure the target mutual information bound is *not* explicitly dependent on the output dimensionality d but instead on the square root sum of eigenvalues of $\Sigma_{\mathcal{M}(X)}$. When $\mathcal{M}(X)$ is largely distributed in a p-rank subspace of \mathbb{R}^d , the above theorem suggests that a noise of scale $O(\sqrt{p})$ is needed. Depending on data distribution, if p is a constant, then the noise can be as low as O(1). This is different from DP where the expected l_2 norm of noise is in a scale of $\Theta(\sqrt{d})$ given constant L_2 -norm sensitivity.

Based on Theorem A.2, we can use an automatic protocol Algorithm Xiao & Devadas (2023) to determine Σ_B and produce an upper bound such that $MI(\mathbf{X}; \mathcal{M}(\mathbf{X}) + \mathbf{B}) \leq (\nu + \beta)$ with high confidence, where ν and β are positive parameters selected as explained below. After sufficiently many simulations, the following theorem ensures that we can obtain accurate enough estimation with arbitrarily high probability.

Theorem A.3. Xiao & Devadas (2023) Assume that $\mathcal{M}(\mathbf{X}) \in \mathbb{R}^d$ and $\|\mathcal{M}(\mathbf{X})\|_2 \leq r$ for some constant r uniformly for any \mathbf{X} , and apply Algorithm 1 to obtain the Gaussian noise covariance Σ_B for a specified mutual information bound $v + \beta$. v and β can be chosen independently, and c is a security parameter. Then, there exists a fixed and universal constant κ such that one can ensure $MI(\mathbf{X}; \mathcal{M}(\mathbf{X}) + \mathbf{B}) \leq \nu + \beta$ with confidence at least $(1 - \gamma)$ once the selections of c, m and γ satisfy,

$$c \ge \kappa r \left(\max\left\{ \sqrt{\frac{d + \log(4/\gamma)}{m}}, \frac{d + \log(4/\gamma)}{m} \right\} + \sqrt{\frac{d \log(4/\gamma)}{m}} \right).$$
(6)

Another way to interpret this noise determination is that a smaller $\mathbb{E}||B||_2$ implies a smaller covariance matrix Σ_B and lesser noise addition on M(X), which indicates X and M(X) have less mutual information. Therefore, less noise addition signifies a higher privacy protection of the model M.

B AUTOMATIC PROTOCOL ALGORITHM XIAO & DEVADAS (2023)

Algorithm 4 (1 - γ)-Confidence Noise Determination of Deterministic Mechanism

Require: Diffusion-Privacy model M, data distribution D, sampling complexity m, security parameter c, and mutual information quantities ν and β .

1: for $k = 1, 2, \ldots, m$ do

2: Independently generate data $X^{(k)}$ from distribution D.

3: Record
$$y^{(k)} = \mathcal{M}(X^{(k)})$$
.

- 5: Calculate empirical mean $\hat{\mu} = \frac{1}{m} \sum_{k=1}^{m} y^{(k)}$ and the empirical covariance estimation $\hat{\Sigma} = \frac{1}{m} \sum_{k=1}^{m} (y^{(k)} \hat{\mu}) (y^{(k)} \hat{\mu})^T$.
- 6: Apply singular value decomposition (SVD) on $\hat{\Sigma}$ and obtain the decomposition as $\hat{\Sigma} = \hat{U}\hat{\Lambda}\hat{U}^T$, where $\hat{\Lambda}$ is the diagonal matrix of eigenvalues $\lambda_1 \ge \lambda_2 \ge \ldots \ge \lambda_d$.
- 7: Determine the maximal index $j_0 = \arg \max_j \lambda_j$ for those $\lambda_j > c$.
- 8: if $\min_{1 \le j \le j_0, 1 \le d} (\lambda_j \lambda_j) > r \sqrt{d/c + 2c}$ then then

9: **for**
$$j = 1, 2, ..., d$$
 do

10: Determine the *j*-th element of a diagonal matrix A_B as

$$\lambda_{B,j} = \frac{2\nu}{\sqrt{\lambda_j + 10c\nu/\beta} \cdot \left(\sum_{j=1}^d \sqrt{\lambda_j + 10c\nu/\beta}\right)}$$

- 11: end for
- 12: Determine the Gaussian noise covariance as $\Sigma_B = \hat{U} A_B \hat{U}^T$.

13: else

- 14: Determine the Gaussian noise covariance as $\Sigma_B = \left(\sum_{j=1}^d \lambda_j + dc/(2\nu)\right) \cdot I_d$.
- 15: end if
- 16: **Return** Gaussian covariance matrix Σ_B .

C PAC PRIVACY PROOF OF OUR MODEL

Since the randomized response method adds Gaussian noise Z in the random response mechanism, it can be proven to be PAC private using Theorem A.1 and Theorem A.2 in the paper:

- Algorithm 1 can be combined and written as H(X) + B, where H(X) represents the RR mechanism and B denotes Gaussian noise with B ~ N(0, Σ_B).
- Applying Theorem A.2 and Theorem A.3, we can yield a satisfied upper bound (with arbitrary high confidence) for the whole process $MI(X; \mathcal{M}(X)) \leq \frac{1}{2} \cdot \log \det(\mathbf{I}_d + \Sigma_{\mathcal{M}(X)} \cdot \Sigma_{\mathbf{B}}^{-1})$, where $\mathcal{M}(X) = H(X) + \mathbf{B}$.
- Hence, $MI(X; \mathcal{M}(X))$ can be used to produce an upper bound for the attack success probability via Theorem A.1.

What's more, Algorithm 2 is a heuristic method that further enhances privacy. We calculate noise B to measure its privacy strength, achieving the best performance as shown in Table 1. Of independent interest to ensure the PAC privacy of Algorithm 2 itself, we simply need to add noise B to the final x_T in Algorithm 2.

D DETAILED DATA IN FIGURE 2

Table 2: Perceptual and privacy score comparisons on CelebA with image resolution 64×64 . In our model, we train with data on a specific label respectively, and Model-Gender means our model is trained with CelebA-Gender dataset. ξ indicates a hyperparameter derived from the Randomized Response (RR).

	DPGEN		P3DM-Gender		P3DM-Smile	
ξ	$FID\downarrow$	Privacy Score ↑	$FID\downarrow$	Privacy Score ↑	$FID\downarrow$	Privacy Score ↑
1			175±2.5	$0.6{\pm}0.05$	170±1.85	0.63±0.08
5	155 ± 1.5	$0.46 {\pm} 0.076$	125.8 ± 3.6	$0.55 {\pm} 0.082$	$107.85 {\pm} 4.7$	0.6±0.075
10	$39.2 {\pm} 0.7$	$0.4{\pm}0.03$	$44.82 {\pm} 4.48$	$0.5 {\pm} 0.05$	37.96±1.85	$0.56 {\pm} 0.045$
15	$38.5 {\pm} 0.7$	$0.4{\pm}0.06$	$40 {\pm} 0.5$	$0.45 {\pm} 0.05$	37.9±0.6	0.5±0.025
∞	$\textbf{34.4}{\pm}\textbf{1.8}$	$0.376 {\pm} 0.03$	$34.64{\pm}2.18$	$0.4{\pm}0.04$	36.03 ± 1.32	

Table 3: Perceptual and privacy score comparisons on CelebA with image resolution 64×64 . In our model, we train with data on a specific label respectively, and Model-Gender means our model is trained with CelebA-Gender dataset. ε represents the DP parameter of baselines.

		DPDM	DP-MEPF		
ε	$FID\downarrow$	Privacy Score ↑	FID ↓	Privacy Score ↑	
1			67.5±1.45	$0.48 {\pm} 0.02$	
5	170 ± 2	$0.57 {\pm} 0.06$	$61.2{\pm}1.8$	$0.45 {\pm} 0.03$	
10	117 ± 1.5	$0.47 {\pm} 0.04$	57.5 ± 1.2	$0.4{\pm}0.04$	
15	113 ± 1.2	$0.35 {\pm} 0.05$	55.8 ± 1.4	$0.38 {\pm} 0.02$	
∞	$110{\pm}1.2$	$0.3 {\pm} 0.06$	52 ± 2.3	$0.34{\pm}0.025$	