

# PAC Privacy Preserving Diffusion Models

Anonymous authors

Paper under double-blind review

## 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.

## 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.

- 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 core objective of DP is to evaluate the distinguishability between processed and original data. Yet, a significant hurdle is the difficulty of providing adversarial worst-case proofs, especially when comparing two distinct datasets. This difficulty presents a substantial barrier to validating the efficacy of DP methods in practical scenarios.

- 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. This definition diverges from traditional Differential Privacy (DP) by focusing on the difficulty of reconstructing data using any measure function, thereby offering broader applicability beyond the adversarial worst-case scenarios considered by DP. PAC privacy’s strength lies in its model-agnostic nature: it does not depend on an in-depth understanding of a model’s internal workings or its training process. This attribute makes it exceptionally versatile, enabling its application across a wide array of models and contexts with relative ease. Furthermore, leveraging PAC privacy comes with a framework that autonomously determines the minimal noise addition necessary for effective data protection (Xiao & Devadas, 2023). This framework simplifies the process of achieving privacy-preserving data handling, making it more accessible for widespread use in the field of data science and beyond.

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 Work

### 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 Background

### 3.1 Differential Privacy

A randomized mechanism  $\mathcal{M}$  is said  $(\epsilon, \delta)$ -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):

$$\Pr[\mathcal{M}(D) \in S] \leq e^\epsilon \cdot \Pr[\mathcal{M}(D') \in S] + \delta \quad (1)$$

Here,  $\epsilon$  is the upper bound on the privacy loss corresponding to  $\mathcal{M}$ , and  $\delta$  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.

### 3.2 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

$$p_{\theta, \phi}(x_t|x_{t+1}, y) = Z p_\theta(x_t|x_{t+1}) p_\phi(y|x_t) \quad (2)$$

Where  $Z$  is a normalizing constant. According to unconditional reverse process,

$$p_\theta(x_t|x_{t+1}) \sim \mathcal{N}(\mu, \Sigma) \quad (3)$$

And then we take the logarithm, it could be written as

$$\log(p_\theta(x_t|x_{t+1})p_\phi(y|x_t)) \sim \mathcal{N}(\mu + \Sigma g, \Sigma) \quad (4)$$

where  $g$  is the gradient of classifier

$g = \nabla_{x_t} \log p_\phi(y|x_t)|_{x=\mu}$ . 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,  $\Sigma$ , and the vector  $g$ . This methodology allows for the generation of high-quality, targeted synthetic data.

### 3.3 PAC Privacy

**Definition 3.3.1.**  $((\delta, \rho, D)$  PAC Privacy (Xiao & Devadas, 2023)). For a data processing mechanism  $\mathcal{M}$ , given some data distribution  $D$ , a measure function  $\rho(\cdot, \cdot)$ , and a finite set  $X^*$ , we say  $\mathcal{M}$  satisfies  $(\delta, \rho, D)$ -PAC Privacy if the following experiment is impossible: A user generates data  $X$  from distribution  $D$  and sends  $\mathcal{M}(X)$  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 probability at least  $(1 - \delta)$ ,  $\rho(\tilde{X}, X) = 1$ .

Definition above depicts the reconstruction hardness for the attackers to recover the private data distribution  $\mathcal{M}(X)$ . With a lower bound probability  $(1 - \delta)$ , such that the measure function  $\rho(\cdot, \cdot)$  is not able to 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 to the failure of adversarial inferences.

Hence, the author introduces mutual information to measure the mutual dependence of two variables,

$$MI(x, w) = \mathcal{H}(x) - \mathcal{H}(x|w) = \mathcal{D}_{KL}(P_{x,w} || P_x \otimes P_w), \quad (5)$$

where  $\mathcal{D}_{KL}$  denotes the KL divergence. Utilizing the mutual information, the Definition can be rewritten as  $\mathcal{M}$  satisfies  $(\Delta_{KL}\delta, \rho, D)$  PAC Advantage Privacy where

$$\Delta_{KL}\delta = \mathcal{D}_{KL}(1_\delta || 1_{\delta^e}) \leq MI(X; M(X)). \quad (6)$$

Therefore, the goal for the rest of this paper, following Theorem above, is to ascertain or establish limits for the mutual information  $MI(X; M(X))$  with a high level of certainty.

## 4 Methods

We introduce the PAC Privacy Preserving Diffusion Model (P3DM), which aims to safeguard privacy for specific attributes, building upon the foundation set by DPGEN (Chen et al., 2022). DPGEN introduces a Randomized Response (RR) mechanism, outlined as follows:

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

With the method,  $\tilde{x}_i$  is designed to point to one of its  $k$  nearest neighbors with certain probability  $Pr[H(\tilde{x}_i) = x_i^r]$ , which is designed to enforce  $\epsilon$ -differential privacy by privatizing the recovery direction  $d = (\tilde{x}_i - x_i^r)/\sigma_i^2$ . Following perturbation, DPGEN learns an energy score function:

$$l(\theta, \sigma) = \frac{1}{2} \mathbb{E}_{p(x)} \mathbb{E}_{\tilde{x} \sim \mathcal{N}(x, \sigma^2)} \left[ \|d - \nabla_x \log q_\theta(\tilde{x})\|^2 \right] \quad (7)$$

Subsequently, images are synthesized via the Langevin MCMC process.

**Algorithm 1** Private Conditional Guidance Langevin Dynamics Sampling

---

**Require:** class labels  $y = \{y_1, y_2, \dots, y_n\}$  from one of balanced dataset attributes;  
 gradient scale  $k$ ;  $\{\delta_i\}_{i=1}^L$ ,  $\epsilon$ ,  $T$ ;  
 pretrained scorenet model  $s_\theta$ ;  
 pretrained classifier model on noisy images  $c_\theta$

- 1: Randomly sample  $y_n$  from  $y$
- 2: Initialize  $x_0$
- 3: **for**  $i$  from 1 to  $L$  **do**
- 4:   stepsize  $\alpha_i \leftarrow \delta_i^2 / \delta_L^2$
- 5:   **for**  $t$  from 1 to  $T$  **do**
- 6:     Sample noise  $z_t \sim \mathcal{N}(0, I)$
- 7:      $x_t \leftarrow x_{t-1} + \frac{\alpha_i}{2} s_\theta(x_{t-1}, \delta_i) + \sqrt{\alpha_i} z_t + k \Sigma_\theta(x_{t-1}) \nabla_{x_{t-1}} \log c_\theta(y_n | x_{t-1})$
- 8:   **end for**
- 9:    $x_0 \leftarrow x_T$
- 10: **end for**
- 11: **Return**  $x_T$

---

While DPGEN asserts compliance with stringent  $\epsilon$ -differential privacy, recent findings 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.

#### 4.1 Conditional Private Langevin Sampling

We introduce a method of conditional private guidance within the Langevin sampling algorithm, detailed in Algorithm 1. 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.

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

$$k \Sigma_\theta(x_{t-1}) \nabla_{x_{t-1}} \log c_\theta(y_n | x_{t-1}) \quad (8)$$

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. 3. In Eqn. 8, 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.

In essence, private classifier guidance can be interpreted as a method of deliberately altering certain generated trajectories. To illustrate, consider an image that is autonomously generated to depict Celebrity A donning eyewear, thereby reflecting the known fact that Celebrity A habitually wears glasses. However, when sampling is conducted with the application of private classifier guidance, the attributes are subject to variation. Consequently, there is a possibility that Celebrity A may be depicted without eyewear in the resulting image. This alteration effectively safeguards the attribute of 'glasses' from being a consistent feature in the generated representation of Celebrity A. As a result, the images generated under this guidance have a higher privacy extent when compared to those produced by the original DPGEN.

**Algorithm 2** Privacy Score

---

**Require:** images  $x_i$  generated by algorithm 1;  
ground truth images  $x_k^G$ ;  
sample number  $n$ ;  
pretrained InceptionV3 model  $I_\theta$ ;  
pretrained classifier model  $c_\theta$ .

- 1: private score  $s \leftarrow 0$
- 2: **for**  $x_i$  from  $x_0$  to  $x_n$  **do**
- 3:   find  $\text{argmin}_k \|I_\theta(x_i), I_\theta(x_k^G)\|_2$
- 4:   **if**  $c_\theta(x_i) \neq c_\theta(x_k^G)$  **then**
- 5:      $s \leftarrow s + 1$
- 6:   **end if**
- 7: **end for**
- 8: **Return**  $s/n$

---

## 4.2 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 2, we commence by preparing images  $x_i$  generated from Algorithm 1, the original dataset images  $x_k^G$ , and alongside the classifier model which trained separately with clean images (different from the classifier from Algorithm 1). 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.

## 4.3 PAC Gaussian Noise Addition Calculation

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  $MI(X; M(X) + B)$  (Xiao & Devadas, 2023). PAC privacy provides a  $1 - \gamma$  confidence in the noise determination for the deterministic mechanism  $M$  shown in Algorithm 3, to compute the Gaussian noise covariance  $\Sigma_B$ , ensuring that  $MI(X; M(X) + B) \leq \nu + \beta$ . In practical terms, this means we can evaluate the privacy extent of the model by comparing the mean norm  $\mathbb{E}\|B\|_2$  under the same confidence level. A smaller  $\mathbb{E}\|B\|_2$  implies a smaller covariance matrix and lesser noise addition, signifying a higher privacy extent of the model.

# 5 Experiments

## 5.1 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 1. All the images used in these experiments were of the resolution  $64 \times 64$ .

**Algorithm 3** (1 -  $\gamma$ )-Confidence Noise Determination of Deterministic Mechanism

**Require:** Diffusion-Privacy model  $M$ , data distribution  $D$ , sampling complexity  $m$ , security parameter  $c$ , and mutual information quantities  $\nu$  and  $\beta$ .

- 1: Train  $M$  model with data  $X$ .
- 2: **for**  $k = 1, 2, \dots, m$  **do**
- 3:   sample images  $y^{(k)} = M(X^{(k)})$ .
- 4: **end for**
- 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  $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$ .
- 7: Determine the maximal index  $j_0 = \arg \max_j \lambda_j$  for those  $\lambda_j > c$ .
- 8: **if**  $\min_{1 \leq j \leq j_0, 1 \leq d} (\lambda_j - \hat{\lambda}_j) > r\sqrt{d/c + 2c}$  **then**
- 9:   **for**  $j = 1, 2, \dots, d$  **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  $\Sigma_B$ .

## 5.2 Baselines

In our study, we consider DPGEN (Chen et al., 2022) and DPDM (Dockhorn et al., 2022) as baseline methods. Both of these approaches 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.

## 5.3 Evaluation Metrics

In our evaluation process, we validate the capability of our PAC Privacy Preserving Diffusion Model to generate high-resolution images using two key metrics: (1) the Frechet Inception Distance (Heusel et al., 2017) and (2) the Inception Score (Salimans et al., 2016). These metrics are 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 2. We compare the mean norm of the Gaussian noise  $\mathbb{E}||B||_2$  as detailed in Sec. 4.3. 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.

## 5.4 Empirical Results And Analysis

It is important to note that the  $\varepsilon$  presented in our model within the tables below is merely a hyperparameter derived from the Randomized Response (RR) as indicated in Eqn. 4. This  $\varepsilon$  should not be confused with the  $\varepsilon$  from  $\varepsilon$ -Differential Privacy (DP), as our model uses PAC privacy rather than  $\varepsilon$ -DP. We will later

Table 1: Perceptual and privacy score comparisons on CelebA with image resolution  $64 \times 64$ . In our model, we train with data on a specific label respectively, and Model-Gender means our model is trained with CelebA-Gender dataset. The same applies to the following.

$\epsilon$	P3DM-Gender		P3DM-Smile		DPGEN		DPDM	
	FID $\downarrow$	Privacy Score $\uparrow$	FID $\downarrow$	Privacy Score $\uparrow$	FID $\downarrow$	Privacy Score $\uparrow$	FID $\downarrow$	Privacy Score $\uparrow$
1	175 $\pm$ 2.5	0.6 $\pm$ 0.05	<b>170<math>\pm</math>1.85</b>	<b>0.63<math>\pm</math>0.08</b>				
5	125.8 $\pm$ 3.6	0.55 $\pm$ 0.082	<b>107.85<math>\pm</math>4.7</b>	<b>0.6<math>\pm</math>0.075</b>	155 $\pm$ 1.5	0.46 $\pm$ 0.076	170 $\pm$ 2	0.57 $\pm$ 0.06
10	44.82 $\pm$ 4.48	0.5 $\pm$ 0.05	<b>37.96<math>\pm</math>1.85</b>	<b>0.56<math>\pm</math>0.045</b>	39.16 $\pm$ 0.68	0.4 $\pm$ 0.03	117 $\pm$ 1.5	0.47 $\pm$ 0.04
15	40 $\pm$ 0.5	0.45 $\pm$ 0.05	<b>37.9<math>\pm</math>0.6</b>	<b>0.5<math>\pm</math>0.025</b>	38.5 $\pm$ 0.73	0.4 $\pm$ 0.06	113 $\pm$ 1.2	0.35 $\pm$ 0.05
$\infty$	34.64 $\pm$ 2.18	0.4 $\pm$ 0.04	36.03 $\pm$ 1.32	<b>0.48<math>\pm</math>0.055</b>	<b>34.4<math>\pm</math>1.8</b>	0.376 $\pm$ 0.03	110 $\pm$ 1.2	0.3 $\pm$ 0.06



(a) DPDM FID = 117, Privacy Score = 0.47 (b) DPGEN FID = 39.16, Privacy Score = 0.4 (c) P3DM(Ours) FID = 37.96, Privacy Score = 0.56

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

illustrate how our model assures privacy through the automatic control of mutual information for a PAC privacy guarantee in this section.

Table 1 and Fig.3 details the evaluation results of image visual quality and privacy score on the CelebA dataset with a resolution of  $64 \times 64$ , 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.3, 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. 2, 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.

Table 2:  $\mathbb{E}\|B\|_2$  to ensure  $MI(X; M(X) + B) \leq 1$  on CelebA dataset

Methods	$\epsilon$	$\mathbb{E}\ B\ _2 \downarrow$
<b>P3DM-Gender</b>	5	283.03 $\pm$ 2.25
<b>P3DM-Smile</b>	5	<b>280.60<math>\pm</math>2.57</b>
DPGEN	5	281.3 $\pm$ 2.78
<b>P3DM-Gender</b>	10	328.80 $\pm$ 1.24
<b>P3DM-Smile</b>	10	<b>327.98<math>\pm</math>1.45</b>
DPGEN	10	329.48 $\pm$ 1.6
DPDM	10	335.83 $\pm$ 1.21
<b>P3DM-Gender</b>	$\infty$	332.87 $\pm$ 1.15
<b>P3DM-Smile</b>	$\infty$	<b>331.67<math>\pm</math>1.3</b>
DPGEN	$\infty$	332.02 $\pm$ 1.32
DPDM	$\infty$	335.83 $\pm$ 1.56





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



(b) DPDM for  $\varepsilon = 10$



(c) DPGEN for  $\varepsilon = 10$



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

Figure 2: 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  $64 \times 64$ .




Figure 3: **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) and DPDM (Dockhorn et al., 2022).

Furthermore, in Table 2, we compute the multivariate Gaussian matrix  $B$  with a  $1 - \gamma$  noise determination for the deterministic mechanism  $M$  followed the work by 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.

## 6 Limitations

The dataset attributes analyzed in our study predominantly consist of features with balanced positive and negative samples. Attributes characterized by imbalanced distributions remain underexplored. Furthermore, it should be noted that not all datasets include annotated attributes. Therefore, utilizing image-to-text models such as CLIP (Radford et al., 2021) may prove beneficial for the generation of new attribute labels. Lastly, our model faces challenges in synthesizing higher-resolution privacy-preserving images (e.g., 256x256 pixels) and handling complex datasets such as the CUB dataset (Wah et al., 2011)).

## 7 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.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- 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.
- 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 random features 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.
- 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 Dp2-vae. Differentially private pre-trained variational autoencoders. *arXiv preprint arXiv:2208.03409*, 2022.

- 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.
- Bjarne Pfitzner and Bert Arnrich. Dpd-fvae: Synthetic data generation using federated variational autoencoders with differentially-private decoder. *arXiv preprint arXiv:2211.11591*, 2022.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- 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.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 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 high-dimensional 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.
- C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. Caltech-ucsd birds-200-2011. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- 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.