

Privacy-Preserving Low-Rank Adaptation Against Membership Inference Attacks for Latent Diffusion Models

Zihao Luo^{1*}, Xilie Xu^{2*}, Feng Liu³, Yun Sing Koh¹,
Di Wang⁴, Jingfeng Zhang^{1 4†}

¹The University of Auckland

²The National University of Singapore

³The University of Melbourne

⁴King Abdullah University of Science and Technology

Abstract

Low-rank adaptation (LoRA) is an efficient strategy for adapting latent diffusion models (LDMs) on a private dataset to generate specific images by minimizing the adaptation loss. However, the LoRA-adapted LDMs are vulnerable to membership inference (MI) attacks that can judge whether a particular data point belongs to the private dataset, thus leading to the privacy leakage. To defend against MI attacks, we first propose a straightforward solution: Membership-Privacy-preserving LoRA (MP-LoRA). MP-LoRA is formulated as a min-max optimization problem where a proxy attack model is trained by maximizing its MI gain while the LDM is adapted by minimizing the sum of the adaptation loss and the MI gain of the proxy attack model. However, we empirically find that MP-LoRA has the issue of unstable optimization, and theoretically analyze that the potential reason is the unconstrained local smoothness, which impedes the privacy-preserving adaptation. To mitigate this issue, we further propose a Stable Membership-Privacy-preserving LoRA (SMP-LoRA) that adapts the LDM by minimizing the ratio of the adaptation loss to the MI gain. Besides, we theoretically prove that the local smoothness of SMP-LoRA can be constrained by the gradient norm, leading to improved convergence. Our experimental results corroborate that SMP-LoRA can indeed defend against MI attacks and generate high-quality images.

Code —

<https://github.com/WilliamLUO0/StablePrivateLoRA>

1 Introduction

Generative diffusion models (Ho, Jain, and Abbeel 2020; Song et al. 2021) are leading a revolution in AI-generated content, renowned for their unique generation process and fine-grained image synthesis capabilities. Notably, the Latent Diffusion Model (LDM) (Rombach et al. 2022; Podell et al. 2024) stands out by executing the diffusion process

in latent space, enhancing computational efficiency without compromising image quality. Thus, LDMs can be efficiently adapted to generate previously unseen contents or styles (Meng et al. 2022; Gal et al. 2023; Ruiz et al. 2023; Zhang, Rao, and Agrawala 2023), thereby catalyzing a surge across multiple fields, such as facial generation (Huang et al. 2023; Xu et al. 2024) and medicine (Kazerouni et al. 2022; Shavlokhova et al. 2023).

Among various adaptation methods, Low-Rank Adaptation (LoRA) (Hu et al. 2022) is the superior strategy for adapting LDMs by significantly reducing computational resources while ensuring commendable performance with great flexibility. Compared to the full fine-tuning method which fine-tunes all parameters, LoRA optimizes the much smaller low-rank matrices, making the training more efficient and lowering the hardware requirements for adapting LDMs (Hu et al. 2022). By performing the low-rank decomposition of the transformer structure within the LDM, LoRA offers performance comparable to fine-tuning all LDM parameters (Cuenca and Paul 2023). Moreover, LoRA allows flexible sharing of a pre-trained LDM to build numerous small LoRA modules for various tasks.

However, recent studies (Pang and Wang 2023; Pang et al. 2023; Dubiński et al. 2024) have pointed out that adapted LDMs are facing the severe risk of privacy leakage. The leakage primarily manifests in the vulnerability to Membership Inference (MI) attacks (Shokri et al. 2017), which utilize the model’s loss of a data point to differentiate whether it is a member of the training dataset or not. As shown in Figure 1d, the LoRA-adapted LDM (red circle marker) exhibits an incredibly high Attack Success Rate (ASR) of 82.27%.

To mitigate the issue of privacy leakage, we make the first effort to propose a Membership-Privacy-preserving LoRA (MP-LoRA) method, which is formulated as a min-max optimization problem. Specifically, in the inner maximization step, a proxy attack model is trained to maximize its effectiveness in inferring membership privacy which is quantitatively referred to as MI gain. In the outer minimization step, the LDM is adapted by minimizing the sum of the adaptation loss and the MI gain of the proxy attack model to enhance the preservation of membership privacy.

However, the vanilla MP-LoRA encounters an issue of

*These authors contributed equally.

†Correspondence to: Jingfeng Zhang <jingfeng.zhang@auckland.ac.nz>.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

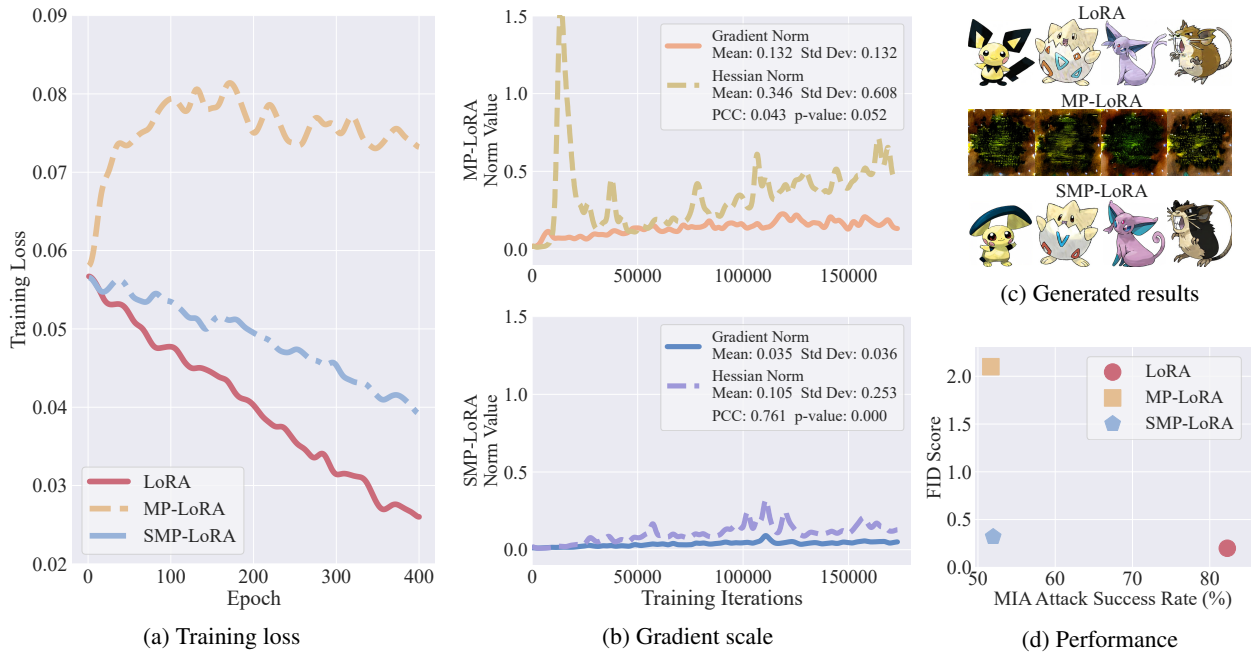


Figure 1: Figure 1a shows the trajectory of the training loss during the adaptation process via LoRA, MP-LoRA, and SMP-LoRA on the Pokemon dataset. Figure 1b displays the mean and standard deviation of the gradient norms and Hessian norms for MP-LoRA and SMP-LoRA throughout the training iterations. It also presents the Pearson correlation coefficients (PCC) and p-values assessing their correlation. Note that each epoch contains 433 training iterations. Figures 1c and 1d demonstrate the generated images and a comparison of evaluation metrics including FID Score and MI attack success rate (ASR). MP-LoRA preserves membership privacy but compromises image generation capability. In contrast, SMP-LoRA effectively preserves membership privacy while maintaining the quality of the generated image, demonstrating its effectiveness in defending against MI attacks without significant loss of functionality. Extensive generated images are visualized in Appendix H.

effective optimization of the training loss, as evidenced in the orange dashed line of Figure 1a. We theoretically find that during MP-LoRA, the local smoothness, quantified by the Hessian norm (the norm of the Hessian matrix) (Bubeck et al. 2015), is independent of and not bounded by the gradient norm (see Proposition 1 for details). This independence hinders the privacy-preserving adaptation of the MP-LoRA (Zhang et al. 2019), thus impeding optimizing the training loss. Besides, we empirically show that the correlation between the Hessian norm and the gradient norm during MP-LoRA is insignificant. This is manifested by the Pearson Correlation Coefficient (PCC) of 0.043 and the p-value above 0.05, as shown in the upper panel of Figure 1b, which corroborates our theoretical analyses.

To stabilize the optimization procedure of MP-LoRA, we further propose a Stable Membership-Privacy-preserving LoRA (SMP-LoRA) method, which incorporates the MI gain into the denominator of the adaptation loss instead of directly summing it. We theoretically demonstrate that this modification ensures a positive correlation (see Proposition 2 for details). Specifically, the local smoothness (that is quantified by the Hessian norm) is positively correlated with and upper bounded by the gradient norm during adaptation, which can improve convergence. Furthermore, we empirically corroborate that during SMP-LoRA, the Hessian norm is positively correlated with the gradient norm, as ev-

idenced by the higher PCC (0.761) and the p-value of less than 0.001 in the lower panel of Figure 1b. The constrained local smoothness allows the SMP-LoRA to achieve better optimization, as shown in the blue dash-dot line of Figure 1a.

To evaluate the performance of the SMP-LoRA, we conducted adapting experiments using the Stable Diffusion v1.5 (CompVis 2022) on the Pokemon (Pinkney 2022) and CelebA (Liu et al. 2015) datasets, respectively. Figure 1d shows that, although MP-LoRA (orange square marker) lowers the ASR to near-random levels, it significantly degrades the image generation capability of LoRA, as evidenced by a high FID score of 2.10 and the poor visual quality in Figure 1c. In contrast, the SMP-LoRA (blue pentagon marker) effectively preserves membership privacy without sacrificing generated image quality significantly, as evidenced by its FID score of 0.32 and ASR of 51.97%.

2 Background and Preliminary

This section outlines the related work and preliminary concepts in diffusion models, low-rank adaptation, and membership inference attacks.

Diffusion Model

Diffusion models (DMs) (Ho, Jain, and Abbeel 2020; Song et al. 2021) have shown remarkable performance in image synthesis. Compared with other generative models such

as GAN (Goodfellow et al. 2014), DMs can mitigate the problems of the training instability and the model collapse (Rombach et al. 2022), while achieving state-of-the-art results in numerous benchmarks (Dhariwal and Nichol 2021). Among various implementations of DMs, the Latent Diffusion Model (LDM) is renowned for generating high-quality images with limited computational resources, thereby widely utilized across many applications. Fundamentally, LDM is a Denoising Diffusion Probabilistic Model (DDPM) (Ho, Jain, and Abbeel 2020) built in the latent space, effectively reducing computational demand while enabling high-quality and flexible image generation (Rombach et al. 2022). Therefore, LDMs have been widely utilized for adaptation (Gal et al. 2023; Hu et al. 2022; Ruiz et al. 2023), capable of delivering high-performance models even when adapted on small-scale datasets.

Adapting LDMs involves a training process that progressively adds noise to the data and then learns to reverse the noise, finely tailoring the model’s output. To be specific, an image $x \in \mathcal{X}$ is initially mapped to a latent representation by a pre-trained encoder $\mathcal{E} : \mathcal{X} \rightarrow \mathcal{Z}$. In the diffusion process, Gaussian noise $\epsilon \sim \mathcal{N}(0, 1)$ is progressively added at each time step $t = 1, 2, \dots, T$, evolving $\mathcal{E}(x)$ into $z_t = \sqrt{\alpha_t}\mathcal{E}(x) + \sqrt{1 - \alpha_t}\epsilon$, where $\alpha_t \in [0, 1]$ is a decaying parameter. Subsequently, the model f_θ is trained to predict and remove noise ϵ , therefore recovering $\mathcal{E}(x)$. Building on this, a pre-trained decoder reconstructs the image from the denoised latent representation. Furthermore, to incorporate conditional information y from various modalities, such as language prompts, a domain-specific encoder τ_ϕ is introduced, projecting y into an intermediate representation. Given a pair (x, y) consisting of an image x and the corresponding text y , the adaptation loss for LDM is defined as follows:

$$\ell_{\text{ada}}(x, y; t, \epsilon, f_\theta) = \|\epsilon - f_\theta(z_t, t, \tau_\phi(y))\|_2^2, \quad (1)$$

where τ_ϕ refers to the pre-trained text encoder from CLIP (Radford et al. 2021).

Low-Rank Adaptation (LoRA)

To unleash the power of large pretrained models, many adaptation methods have emerged. In particular, Low-Rank Adaptation (LoRA) (Hu et al. 2022) provides an efficient and effective solution by freezing the pre-trained model weights and introducing the trainable low-rank counterparts, significantly reducing the number of trainable parameters and memory usage during the LoRA adaptation process. Therefore, LoRA not only lessens the demand for computational resources but also allows for the construction of multiple lightweight portable low-rank matrices on the same pre-trained LDM, addressing various downstream tasks (Cuenca and Paul 2023; Huggingface 2023).

Specifically, when adapting LDMs via LoRA, a low-rank decomposition is performed on each attention layer in the LDM backbone f_θ . During LoRA, assuming that the original pre-trained weight $\theta \in \mathbb{R}^{d \times k}$, a trainable LoRA module \mathbf{BA} is randomly initialized and added to the pre-trained weights θ , where $\mathbf{B} \in \mathbb{R}^{d \times r}$, $\mathbf{A} \in \mathbb{R}^{r \times k}$. Note that the rank

r is significantly less than d or k , which ensures the computational efficiency of LoRA. During the adaptation process via LoRA, for the augmented LDM backbone $f_{\bar{\theta}+\mathbf{BA}}$, all trainable LoRA module parameters \mathbf{B} and \mathbf{A} are updated, while the original parameters $\bar{\theta}$ are frozen. Given an image-text pair (x, y) , the adaptation loss during LoRA is formulated as follows:

$$\ell_{\text{ada}}(x, y; f_{\bar{\theta}+\mathbf{BA}}) = \|\epsilon - f_{\bar{\theta}+\mathbf{BA}}(z_t, t, \tau_\phi(y))\|_2^2. \quad (2)$$

For notational simplicity, we omit the variables t and ϵ in the adaptation loss ℓ_{ada} . Given the training dataset $\mathcal{D}_{\text{tr}} = \{(x_i, y_i)\}_{i=1}^n$ composed of $n \in \mathbb{N}^+$ image-text pairs, the training loss can be calculated as follows:

$$\mathcal{L}_{\text{ada}}(f_{\bar{\theta}+\mathbf{BA}}, \mathcal{D}_{\text{tr}}) = \frac{1}{n} \sum_{i=1}^n \ell_{\text{ada}}(x_i, y_i; f_{\bar{\theta}+\mathbf{BA}}). \quad (3)$$

Note that during the adaptation process via LoRA, the objective function that optimizes the parameters \mathbf{B} and \mathbf{A} is formulated as $\min_{\{\mathbf{B}, \mathbf{A}\}} \mathcal{L}_{\text{ada}}(f_{\bar{\theta}+\mathbf{BA}}, \mathcal{D}_{\text{tr}})$.

Membership Inference Attack

Membership Inference (MI) attack (Shokri et al. 2017) aims to determine whether a particular data point is part of a model’s training set. Recent studies (Carlini et al. 2023; Duan et al. 2023) have shown that DMs are particularly vulnerable to MI attacks, thus undergoing high risks of privacy leakage. MI attacks on diffusion models can be categorized based on the adversary’s capabilities into white-box (Hu and Pang 2023; Matsumoto, Miura, and Yanai 2023; Pang et al. 2023), gray-box (Duan et al. 2023; Kong et al. 2024; Fu et al. 2023), and black-box (Wu et al. 2022; Matsumoto, Miura, and Yanai 2023; Pang and Wang 2023; Zhang et al. 2024) attacks.

In the black-box setting, Wu et al. (2022) were the first to explore MI attacks on DMs and achieved the highest success rate among the black-box attacks mentioned above. They noted that DMs, when replicating training images, consistently produce outputs with higher fidelity and greater alignment with textual captions, indicating significant behavioural differences. Therefore, they utilized the L2 distance between the embeddings of a given image and its corresponding caption-generated image to infer membership.

In the white-box setting, the gradient-based MI attack developed by Pang et al. (2023) stands as the most effective method for DMs in terms of attack performance. They leveraged the model’s gradient to train an attack model for inference. Additionally, Hu and Pang (2023) and Matsumoto, Miura, and Yanai (2023) employed a threshold-based MI attack by analyzing model loss at specific diffusion steps, which we refer to as the loss-based MI attack, also yielding significant attack performance. In contrast, the black-box and gray-box MI attacks previously mentioned, which lack access to internal model information, achieve lower success rates compared to the white-box loss-based and gradient-based MI attacks (Pang et al. 2023).

Conventional techniques to defend against MI attacks include differential privacy (Dwork 2008; Abadi et al. 2016),

min-max membership privacy game (Nasr, Shokri, and Houmansadr 2018), data augmentation (DeVries and Taylor 2017; Cubuk et al. 2020), early stopping, etc. To the best of our knowledge, these techniques have not been systematically applied or comprehensively evaluated in DMs. Notably, some studies (Dockhorn et al. 2023; Ghalebikesabi et al. 2023; Lyu et al. 2023) have applied differential privacy to DMs to achieve a better privacy-accuracy trade-off, yet they have not evaluated its effectiveness against MI attacks in DMs. Therefore, our paper takes the first step by developing a defensive adaptation based on the min-max membership privacy game to defend against MI attacks in DMs.

3 Membership-Privacy-Preserving LoRA

In this section, we first use the min-max optimization to formulate the learning objective of MP-LoRA. Then, we disclose the issue of unstable optimization of MP-LoRA. Finally, we propose the stable SMP-LoRA and its implementation.

A Vanilla Solution: MP-LoRA

Objective function. In MI attack, the conflicting objectives of defenders and adversaries can be modelled as a privacy game (Shokri et al. 2012; Manshaei et al. 2013; Alvim et al. 2017). Adversaries can adjust their attack models to maximize MI gain against the target model, which requires that the defense can anticipate and withstand the strongest inference attacks. Consequently, the defender’s goal is to enhance the preservation of membership privacy in worst-case scenarios where the adversary achieves the maximum MI gain while maintaining the model performance. Inspired by Nasr, Shokri, and Houmansadr (2018), we propose MP-LoRA to defend against MI attacks which is formulated as a min-max optimization problem as follows:

$$\min_{\{\mathbf{B}, \mathbf{A}\}} \left(\underbrace{\mathcal{L}_{\text{ada}}(f_{\bar{\theta}+\mathbf{B}\mathbf{A}}, \mathcal{D}_{\text{tr}})}_{\text{Adaptation loss}} + \lambda \underbrace{\max_{\omega} G(h_{\omega}, \mathcal{D}_{\text{aux}}, f_{\bar{\theta}+\mathbf{B}\mathbf{A}})}_{\text{Membership inference gain}} \right), \quad (4)$$

where $\mathcal{L}_{\text{ada}}(f_{\bar{\theta}+\mathbf{B}\mathbf{A}}, \mathcal{D}_{\text{tr}})$ refers to the adaptation loss for the LDM with LoRA module $f_{\bar{\theta}+\mathbf{B}\mathbf{A}}$ on the training dataset \mathcal{D}_{tr} , h_{ω} is the proxy attack model parameterized by ω , $G(h_{\omega}, \mathcal{D}_{\text{aux}}, f_{\bar{\theta}+\mathbf{B}\mathbf{A}})$ represents the MI gain of the proxy attack model h_{ω} on the auxiliary dataset \mathcal{D}_{aux} .

Therein, the inner maximization aims to search for the most effective proxy attack model h_{ω} for a given adapted LDM $f_{\bar{\theta}+\mathbf{B}\mathbf{A}}$ via maximizing the MI gain. The outer minimization, conversely, searches for the LDM $f_{\bar{\theta}+\mathbf{B}\mathbf{A}}$ that can best preserve membership privacy under the strong proxy attack model h_{ω} while being able to adapt on the training dataset.

Updating the proxy attack model in inner maximization.

The proxy attack model h_{ω} equipped with white-box access to the target LDM $f_{\bar{\theta}+\mathbf{B}\mathbf{A}}$, aims to infer whether a specific image-text pair (x, y) is from the training dataset \mathcal{D}_{tr} for adapting the target LDM $f_{\bar{\theta}+\mathbf{B}\mathbf{A}}$. The model achieves this by constructing an auxiliary dataset \mathcal{D}_{aux} , which consists of half of the member data from \mathcal{D}_{tr} , denoted as $\mathcal{D}_{\text{aux}}^{\text{m}}$, and

an equal amount of local non-member data $\mathcal{D}_{\text{aux}}^{\text{nm}}$. Using the auxiliary dataset \mathcal{D}_{aux} , h_{ω} trains a binary classifier based on the adaptation loss of the target LDM $f_{\bar{\theta}+\mathbf{B}\mathbf{A}}$ to predict the probability of (x, y) for being a member of the \mathcal{D}_{tr} . Consequently, the MI gain of h_{ω} can be quantified based on its performance on the \mathcal{D}_{aux} as follows:

$$G(h_{\omega}, \mathcal{D}_{\text{aux}}, f_{\bar{\theta}+\mathbf{B}\mathbf{A}}) = \frac{1}{2|\mathcal{D}_{\text{aux}}^{\text{m}}|} \sum_{(x,y) \in \mathcal{D}_{\text{aux}}^{\text{m}}} \log(h_{\omega}(\ell_{\text{ada}}(x, y; f_{\bar{\theta}+\mathbf{B}\mathbf{A}}))) + \frac{1}{2|\mathcal{D}_{\text{aux}}^{\text{nm}}|} \sum_{(x,y) \in \mathcal{D}_{\text{aux}}^{\text{nm}}} \log(1 - h_{\omega}(\ell_{\text{ada}}(x, y; f_{\bar{\theta}+\mathbf{B}\mathbf{A}}))). \quad (5)$$

In the inner maximization, the proxy attack model optimizes the parameters ω by maximizing the MI gain, i.e., $\max_{\omega} G(h_{\omega}, \mathcal{D}_{\text{aux}}, f_{\bar{\theta}+\mathbf{B}\mathbf{A}})$.

Adapting the LDM in outer minimization. MP-LoRA optimizes the LDM by directly minimizing a weighted sum of the MI gain for the h_{ω} and the adaptation loss, which enables it to adapt to the training data and protect the private information of the training dataset simultaneously. To be specific, the training loss of MP-LoRA is formulated as

$$\mathcal{L}_{\text{PL}} = \mathcal{L}_{\text{ada}}(f_{\bar{\theta}+\mathbf{B}\mathbf{A}}, \mathcal{D}_{\text{tr}}) + \lambda \cdot G(h_{\omega}, \mathcal{D}_{\text{tr}}, f_{\bar{\theta}+\mathbf{B}\mathbf{A}}), \quad (6)$$

where $\lambda \in \mathbb{R}$ controls the importance of optimizing the adaptation loss versus protecting membership privacy. In the outer minimization of MP-LoRA, the parameters \mathbf{B} and \mathbf{A} is updated by minimizing the \mathcal{L}_{PL} , i.e., $\min_{\{\mathbf{B}, \mathbf{A}\}} \mathcal{L}_{\text{PL}}$.

MP-LoRA is realized by one step of inner maximization to obtain a power proxy attack model by maximizing the MI gain in Equation (5) and one step of outer minimization to update \mathbf{A} and \mathbf{B} by minimizing the training loss in Equation (6). The algorithm of MP-LoRA is shown in Algorithm 2 (Appendix A of the extended version (Luo et al. 2024)).

Unstable Issue of MP-LoRA

In this subsection, we theoretically demonstrate that the convergence for MP-LoRA cannot be guaranteed due to unconstrained local smoothness. Then we validate the theoretical analyses with empirical evidence.

Definition 1 (Relaxed Smoothness Condition from Zhang et al. (2019)). *A second order differentiable function f is (L_0, L_1) -smooth if*

$$\|\nabla^2 f(x)\| \leq L_0 + L_1 \|\nabla f(x)\|. \quad (7)$$

Lemma 1 (Zhang et al. (2019)). *Let f be a second-order differentiable function and (L_0, L_1) -smooth. If the local smoothness, quantified by the Hessian norm (the norm of the Hessian matrix), is positively correlated with the gradient norm (i.e., $L_1 > 0$), then the gradient norm upper bounds the local smoothness, facilitating faster convergence and increasing the likelihood of converging to an optimal solution.*

Proposition 1. *MP-LoRA does not satisfy the positive correlation as described in Lemma 1, therefore the convergence cannot be guaranteed and the model may settle at a suboptimal solution.*

Proof. We establish the Relaxed Smoothness Condition for MP-LoRA as follows:

$$\left\| \frac{\partial^2 \mathcal{L}_{\text{PL}}}{\partial \mathbf{BA}^2} \right\| \leq L_0 + L_1 \left\| \frac{\partial \mathcal{L}_{\text{PL}}}{\partial \mathbf{BA}} \right\|,$$

where $L_0 = \left\| \frac{\partial^2 \mathcal{L}_{\text{ada}}}{\partial \mathbf{BA}^2} \right\| + \lambda \left\| \frac{\partial^2 G}{\partial \mathbf{BA}^2} \right\|$, $L_1 = 0$, (8)

in which \mathcal{L}_{ada} represents the adaptation loss and G represents the MI gain. The detailed derivation is presented in Appendix B of the extended version (Luo et al. 2024). The value of L_1 being zero indicates that the Hessian norm is independent of and not bounded by the gradient norm, suggesting that the local smoothness is unconstrained. \square

Next, we provide empirical evidence to support our theoretical analyses. We tracked the gradient norm and the Hessian norm of the training loss at each training iteration, and calculated their Pearson Correlation coefficient (PCC) and p-value as shown in Figure 1b. The details for calculating the gradient norm and the Hessian norm can be found in Appendix C. In Figure 1b, the low PCC of 0.043 for MP-LoRA suggests a very weak correlation between the Hessian norm and the gradient norm. Additionally, with the p-value of 0.052, there is insufficient evidence to reject the hypothesis of no correlation. This indicates that the Hessian norm is unbounded, implying that the local smoothness, quantified by the Hessian norm (Bubeck et al. 2015), is unconstrained. Such unconstrained local smoothness leads to the unstable optimization issue in MP-LoRA, and even to the failure of adaptation, as evidenced in the orange dashed line of Figure 1a and the poor visual quality of the generated images in Figure 1c.

Stabilizing MP-LoRA

To mitigate the aforementioned optimization issue of MP-LoRA, we propose SMP-LoRA by incorporating the MI gain into the denominator of the adaptation loss. The objective function of SMP-LoRA is formulated as follows:

$$\min_{\{\mathbf{B}, \mathbf{A}\}} \left(\frac{\mathcal{L}_{\text{ada}}(f_{\bar{\theta}+\mathbf{BA}}, \mathcal{D}_{\text{tr}})}{1 - \lambda \max_{\omega} G(h_{\omega}, \mathcal{D}_{\text{aux}}, f_{\bar{\theta}+\mathbf{BA}})} \right). \quad (9)$$

To optimize Equation (9), SMP-LoRA targets to minimize the following training loss function, i.e.,

$$\mathcal{L}_{\text{SPL}} = \frac{\mathcal{L}_{\text{ada}}(f_{\bar{\theta}+\mathbf{BA}}, \mathcal{D}_{\text{tr}})}{1 - \lambda \cdot G(h_{\omega}, \mathcal{D}_{\text{tr}}, f_{\bar{\theta}+\mathbf{BA}}) + \delta}, \quad (10)$$

where δ is a stabilizer with a small value such as $1e-5$. This prevents the denominator from approaching zero and ensures stable calculation.

The implementation of SMP-LoRA is detailed in Algorithm 1. At each training step, SMP-LoRA will first update the proxy attack model by maximizing the MI gain and then update the LDM by minimizing the training loss \mathcal{L}_{SPL} .

Algorithm 1: Stable Membership-Privacy-preserving LoRA

Input: Training dataset \mathcal{D}_{tr} for adaptation process, Auxiliary dataset $\mathcal{D}_{\text{aux}} = \mathcal{D}_{\text{aux}}^{\text{m}} \cup \mathcal{D}_{\text{aux}}^{\text{nm}}$, a pre-trained LDM f_{θ} , a proxy attack model h_{ω} parameterized by ω , learning rate η_1 and η_2

Output: a SMP-LoRA for LDMs

- 1: Perform low-rank decomposition on f_{θ} to obtain $f_{\bar{\theta}+\mathbf{BA}}$ (\mathbf{B} and \mathbf{A} are trainable LoRA modules)
 - 2: **for** each epoch **do**
 - 3: **for** each training iteration **do**
 - 4: Sample batches S^{m} and S^{nm} from $\mathcal{D}_{\text{aux}}^{\text{m}}$ and $\mathcal{D}_{\text{aux}}^{\text{nm}}$
 - 5: Calculate the MI gain G^* on $S^{\text{m}} \cup S^{\text{nm}}$
 - 6: Update the parameters $\omega \leftarrow \omega + \eta_1 \cdot \nabla_{\omega} G^*$.
 - 7: Sample a fresh batch from \mathcal{D}_{tr}
 - 8: Calculate the training loss $\mathcal{L}^* = \mathcal{L}_{\text{SPL}}$
 - 9: Update parameters $\mathbf{A} \leftarrow \mathbf{A} - \eta_2 \cdot \nabla_{\mathbf{A}} \mathcal{L}^*$ and $\mathbf{B} \leftarrow \mathbf{B} - \eta_2 \cdot \nabla_{\mathbf{B}} \mathcal{L}^*$, respectively
 - 10: **end for**
 - 11: **end for**
-

Proposition 2. *SMP-LoRA satisfies the positive correlation as described in Lemma 1, thus promoting faster convergence, and the model is more likely to converge to an optimal solution.*

Proof. We establish the Relaxed Smoothness Condition for SMP-LoRA as follows:

$$\left\| \frac{\partial^2 \mathcal{L}_{\text{SPL}}}{\partial \mathbf{BA}^2} \right\| \leq L'_0 + L'_1 \left\| \frac{\partial \mathcal{L}_{\text{SPL}}}{\partial \mathbf{BA}} \right\|,$$

where $\mu = \frac{\partial \mathcal{L}_{\text{ada}}}{\partial \mathbf{BA}}$, $\nu = \lambda \frac{\partial G}{\partial \mathbf{BA}}$,

$$L'_0 = \frac{1}{1 - \lambda G + \delta} \cdot \left\| \frac{\partial^2 \mathcal{L}_{\text{ada}}}{\partial \mathbf{BA}^2} \right\| + \frac{\lambda \mathcal{L}_{\text{ada}}}{(1 - \lambda G + \delta)^2} \cdot \left\| \frac{\partial G^2}{\partial \mathbf{BA}^2} \right\|,$$

$$L'_1 = \frac{2\|\nu\|}{1 - \lambda G + \delta}. \quad (11)$$

Please refer to Appendix B of the extended version (Luo et al. 2024) for detailed derivation. The value of L'_1 being greater than zero indicates that the Hessian norm is positively correlated with and upper bounded by the gradient norm, suggesting that the gradient norm constrains the local smoothness during adaptation. \square

Subsequently, we further corroborate our theoretical analyses with the following empirical evidence. Compared to MP-LoRA's insignificant correlation, SMP-LoRA demonstrates a strong positive correlation between the Hessian norm and the gradient norm, evidenced by the PCC of 0.761 and the p-value less than 0.001 in the lower panel of Figure 1b. This indicates that the Hessian norm, which represents the local smoothness, is upper bounded by the gradient norm, resulting in lower mean (0.105) and standard deviation (0.253) than MP-LoRA. Consequently, the constrained local smoothness mitigates the issue of unstable optimization and enables the SMP-LoRA to converge to a more optimal solution, as demonstrated by the progressively decreasing training loss shown in the blue dash-dot line of Figure 1a

Dataset	Method	FID ↓	KID ↓	ASR (%) ↓	$\frac{ AUC-0.5 }{0.5}$ ↓	TPR @5%FPR (%) ↓
Pokemon	LoRA	0.20±0.04	0.003±0.00	82.27±4.38	0.73±0.09	4.44±1.45
	MP-LoRA	2.10±0.51	0.121±0.00	51.67±2.73	0.07±0.04	2.23±1.27
	SMP-LoRA	0.32±0.07	0.004±0.00	51.97±1.20	0.14±0.02	4.45±2.12
CelebA_Small	LoRA	N/A	0.05±0.00	91.53±2.27	0.94±0.02	81.33±9.68
	MP-LoRA	N/A	0.30±0.04	55.67±2.91	0.12±0.05	8.00±3.46
	SMP-LoRA	N/A	0.03±0.01	56.00±2.52	0.24±0.08	17.33±5.21
CelebA_Large	LoRA	0.52±0.01	0.06±0.00	87.83±0.17	0.87±0.02	66.83±6.00
	MP-LoRA	2.34±0.71	0.30±0.05	53.58±1.52	0.07±0.04	2.50±0.87
	SMP-LoRA	0.60±0.04	0.05±0.00	48.83±1.17	0.19±0.05	1.67±1.01
CelebA_Gender	LoRA	N/A	0.06±0.00	84.63±1.67	0.79±0.06	36.67±5.83
	MP-LoRA	N/A	0.32±0.04	55.43±2.07	0.14±0.06	5.83±1.67
	SMP-LoRA	N/A	0.06±0.00	54.20±0.40	0.15±0.05	4.17±3.00
CelebA_Varying	LoRA	N/A	0.06±0.00	87.30±0.92	0.83±0.04	47.10±23.37
	MP-LoRA	N/A	0.26±0.06	55.73±0.75	0.07±0.04	3.73±1.07
	SMP-LoRA	N/A	0.04±0.01	53.97±1.82	0.15±0.01	4.57±3.27

Table 1: Performance of LoRA, MP-LoRA, and SMP-LoRA across five datasets, as measured by FID, KID, ASR, AUC, and TPR at 5% FPR. Results are presented as mean \pm standard error, based on three independent runs with different seeds. Due to the limited size of the CelebA_Small, CelebA_Gender, and CelebA_Varying datasets, the FID scores are not available. It is important to note that an AUC value closer to 0.5 (random level) indicates a stronger defensive capability against MI attacks.

and the superior performance on both FID and ASR metrics illustrated by the blue pentagon marker in Figure 1d.

Notably, SMP-LoRA also exhibits lower mean and standard deviation of the gradient norm compared to MP-LoRA. We provide further empirical analysis in Appendix D of the extended version (Luo et al. 2024).

4 Experiments

In this section, we first evaluate the performance of LoRA, MP-LoRA, and SMP-LoRA in terms of image generation capability and effectiveness in defending against MI attacks. Then, we conduct ablation studies on the important hyperparameters and further extend our membership-privacy-preserving method to full fine-tuning and Dream-Booth (Ruiz et al. 2023) methods. Subsequently, we compare SMP-LoRA with traditional techniques for stabilizing the gradient and evaluate the effectiveness of SMP-LoRA in defending against MI attacks in different settings.

Dataset. In our experiment, we utilized four datasets: Pokemon (Pinkney 2022), CelebA (Liu et al. 2015), AFHQ (Choi et al. 2020), and MS-COCO (Lin et al. 2014). We created several subsets from CelebA, including CelebA_Small and CelebA_Large, both balanced with equal image contribution per individual, as well as CelebA_Gender and CelebA_Varying, which are imbalanced with a 7 : 3 gender ratio and varied image contributions per individual, respectively. Additionally, we constructed CelebA_Large_5X, which is five times larger than the CelebA_Large along with comparably sized AFHQ_Large_5X and MS-COCO_Large_5X. For more details, please refer to Table 4 in Appendix E of the extended version (Luo et al. 2024).

Model hyperparameters. We utilized the official pre-trained Stable Diffusion v1.5 (CompVis 2022) for LDMs to build the LoRA module, with specific model hyperparameters detailed in Table 5 (Appendix E). We employed a 3-layer MLP as the proxy attack model h_ω and a structurally similar new attack model h' to evaluate the effectiveness of adapted LDMs in defending against MI attacks.

Evaluation metrics. We employed Attack Success Rate (ASR) (Choquette-Choo et al. 2021), Area Under the ROC Curve (AUC), and True Positive Rate (TPR) to evaluate the effectiveness of MP-LoRA and SMP-LoRA in defending against MI attacks. Lower values for ASR and TPR indicate a more effective defense against MI attacks. Consistent with many prior studies (Chen, Yu, and Fritz 2022; Ye et al. 2022; Duan et al. 2023; Dubiński et al. 2024), we default that the TPR measures the capability of the attack model to identify samples as members correctly. Consequently, an AUC value closer to 0.5, i.e., a smaller value of $\frac{|AUC-0.5|}{0.5}$, represents a stronger defense capability, as MI attacks involve determining both membership and non-membership. For assessing the image generation capability of the adapted LDMs, we utilized the Fréchet Inception Distance (FID) (Heusel et al. 2017) and the Kernel Inception Distance (KID) (Biniowski et al. 2018), with lower values denoting better image quality. Please refer to Appendix E for detailed explanations.

Effectiveness of SMP-LoRA in Defending Against MI Attacks

In Table 1, we report the performance of LoRA, MP-LoRA, and SMP-LoRA, across the Pokemon, CelebA_Small, CelebA_Large, CelebA_Gender, CelebA_Varying datasets.

Dataset	Method	FID ↓	KID ↓	ASR (%) ↓	$\frac{ AUC-0.5 }{0.5}$ ↓	TPR @5%FPR (%) ↓
CelebA_Large_5X	LoRA	0.53	0.051	92.80	0.94	85.80
	SMP-LoRA	0.59	0.062	51.85	0.15	1.70
AFHQ_Large_5X	LoRA	0.39	0.025	88.20	0.86	85.00
	SMP-LoRA	0.51	0.041	56.00	0.26	12.00
MS-COCO_Large_5X	LoRA	0.37	0.014	80.40	0.78	68.50
	SMP-LoRA	0.72	0.023	46.10	0.14	4.1

Table 2: Performance of LoRA and SMP-LoRA across three larger datasets, as measured by FID, KID, AUC, and TPR.

Compared to LoRA, MP-LoRA displays considerably higher FID and KID scores, indicating lower quality of generated images. Meanwhile, MP-LoRA exhibits near-random levels of ASR and AUC, along with low TPR values at 5% TPR, showcasing its effectiveness in defending against MI attacks. Notably, the FID and KID scores for SMP-LoRA closely align with those of LoRA, suggesting that SMP-LoRA only makes a minor sacrifice to the quality of generated images. Also, SMP-LoRA achieves near-random levels of ASR and AUC, and low TPR values at 5% FPR. These results demonstrate that compared to MP-LoRA, SMP-LoRA effectively preserves membership privacy against MI attacks without significantly compromising image generation capability.

In Table 2, we present the performance of LoRA and SMP-LoRA on the CelebA_Large_5X, AFHQ_Large_5X, and MS-COCO_Large_5X datasets. SMP-LoRA remains effective on larger datasets, consistently defending against MI attacks and generating high-quality images.

Ablation Study

This subsection presents ablation studies on the important hyperparameters: the coefficient λ , the learning rate η_2 , and the LoRA’s rank r . We also extended the application of SMP-LoRA to the full fine-tuning and DreamBooth (Ruiz et al. 2023) method to assess the generalizability of our membership-privacy-preserving method. Additionally, we compared the performance of SMP-LoRA with gradient clipping and normalization techniques. Furthermore, we evaluated the effectiveness of SMP-LoRA in preserving membership privacy under the black-box MI attacks (Wu et al. 2022) and the white-box gradient-based MI attacks (Pang et al. 2023). Due to space constraints, only the key conclusions are presented in the main paper, with all tables and detailed analyses located in Appendix F of the extended version (Luo et al. 2024).

Coefficient λ . Table 6 (Appendix F) presents the performance of SMP-LoRA with different coefficient $\lambda \in \{1.00, 0.50, 0.10, 0.05, 0.01\}$ across the Pokemon, CelebA_Small, and CelebA_Large datasets. As λ decreases from 1.00 to 0.01, the FID and KID scores gradually decrease, while ASR increases and AUC deviates further from 0.5, suggesting that a lower λ shifts the focus more towards minimizing adaptation loss rather than protecting membership privacy.

Extending to the full fine-tuning and DreamBooth Methods. In Table 9 (Appendix F), we present the performance of SMP-LoRA and its extension to the full fine-tuning and DreamBooth (Ruiz et al. 2023) methods on the Pokemon dataset. Our membership-privacy-preserving method continues to effectively protect membership privacy when applied to these methods, highlighting its potential applicability across different adaptation methods.

Comparing with gradient clipping and normalization techniques. In Table 10 (Appendix F), we report the performance of SMP-LoRA and MP-LoRA enhanced with gradient clipping and normalization techniques on the Pokemon dataset. These traditional techniques for stabilizing the gradient cannot address the unstable optimization issue in MP-LoRA.

Defending against MI attacks in different settings. In Table 11 (Appendix F), we display the attack performance on LoRA and SMP-LoRA using the black-box MI attack (Wu et al. 2022) and the white-box gradient-based MI attack (Pang et al. 2023), which is currently the most potent MI attack targeting DMs. This renders further comparisons with weaker attacks unnecessary, such as gray-box MI attacks (Duan et al. 2023; Kong et al. 2024; Fu et al. 2023). Compared to LoRA, SMP-LoRA, specifically designed to defend against white-box loss-based MI attacks, consistently provides enhanced membership privacy protection when facing MI attacks in different settings.

5 Conclusion

In this paper, we proposed membership-privacy-preserving LoRA (MP-LoRA), a method based on low-rank adaptation (LoRA) for adapting latent diffusion models (LDMs), while mitigating the risk of privacy leakage. We first highlighted the unstable issue in MP-LoRA. Directly minimizing the sum of the adaptation loss and MI gain can lead to unconstrained local smoothness, which results in unstable optimization. To mitigate this issue, we further proposed a stable membership-privacy-preserving LoRA (SMP-LoRA) method, which constrains the local smoothness through the gradient norm to improve convergence. Detailed theoretical analyses and comprehensive empirical results demonstrate that the SMP-LoRA can effectively preserve membership privacy against MI attacks and generate high-quality images.

Acknowledgments

Feng Liu is supported by the Australian Research Council (ARC) with grant numbers DP230101540 and DE240101089, and the NSF&CSIRO Responsible AI program with grant number 2303037.

Di Wang is supported in part by the funding BAS/1/1689-01-01, URF/1/4663-01-01, REI/1/5232-01-01, REI/1/5332-01-01, and URF/1/5508-01-01 from KAUST, and funding from KAUST - Center of Excellence for Generative AI, under award number 5940.

References

- Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 308–318.
- Alvim, M. S.; Chatzikokolakis, K.; Kawamoto, Y.; and Palamidessi, C. 2017. Information leakage games. In *Decision and Game Theory for Security: 8th International Conference, GameSec 2017, Vienna, Austria, October 23-25, 2017, Proceedings*, 437–457. Springer.
- Bińkowski, M.; Sutherland, D. J.; Arbel, M.; and Gretton, A. 2018. Demystifying MMD GANs. In *International Conference on Learning Representations*.
- Bubeck, S.; et al. 2015. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4): 231–357.
- Carlini, N.; Hayes, J.; Nasr, M.; Jagielski, M.; Sehwag, V.; Tramèr, F.; Balle, B.; Ippolito, D.; and Wallace, E. 2023. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23)*, 5253–5270.
- Chen, D.; Yu, N.; and Fritz, M. 2022. RelaxLoss: Defending Membership Inference Attacks without Losing Utility. In *International Conference on Learning Representations*.
- Choi, Y.; Uh, Y.; Yoo, J.; and Ha, J.-W. 2020. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8188–8197.
- Choquette-Choo, C. A.; Tramèr, F.; Carlini, N.; and Papernot, N. 2021. Label-only membership inference attacks. In *International conference on machine learning*, 1964–1974. PMLR.
- CompVis. 2022. Stable Diffusion. <https://github.com/CompVis/stable-diffusion>. Accessed on January 16, 2024.
- Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2020. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 702–703.
- Cuenca, P.; and Paul, S. 2023. Using LoRA for Efficient Stable Diffusion Fine-Tuning. Accessed on January 16, 2024.
- DeVries, T.; and Taylor, G. W. 2017. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Dockhorn, T.; Cao, T.; Vahdat, A.; and Kreis, K. 2023. Differentially Private Diffusion Models. *Transactions on Machine Learning Research*.
- Duan, J.; Kong, F.; Wang, S.; Shi, X.; and Xu, K. 2023. Are diffusion models vulnerable to membership inference attacks? In *International Conference on Machine Learning*, 8717–8730. PMLR.
- Dubiński, J.; Kowalczyk, A.; Pawlak, S.; Rokita, P.; Trzciński, T.; and Morawiecki, P. 2024. Towards More Realistic Membership Inference Attacks on Large Diffusion Models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 4860–4869.
- Dwork, C. 2008. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, 1–19. Springer.
- Fu, W.; Wang, H.; Gao, C.; Liu, G.; Li, Y.; and Jiang, T. 2023. A Probabilistic Fluctuation based Membership Inference Attack for Generative Models. *arXiv preprint arXiv:2308.12143*.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-or, D. 2023. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. In *The Eleventh International Conference on Learning Representations*.
- Ghalebikesabi, S.; Berrada, L.; Goyal, S.; Ktena, I.; Stanforth, R.; Hayes, J.; De, S.; Smith, S. L.; Wiles, O.; and Balle, B. 2023. Differentially private diffusion models generate useful synthetic images. *arXiv preprint arXiv:2302.13861*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hu, E. J.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Hu, H.; and Pang, J. 2023. Membership inference of diffusion models. *arXiv preprint arXiv:2301.09956*.
- Huang, Z.; Chan, K. C.; Jiang, Y.; and Liu, Z. 2023. Collaborative diffusion for multi-modal face generation and editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6080–6090.
- Huggingface. 2023. Conceptual Guides LoRA. https://huggingface.co/docs/peft/conceptual_guides/lora. Accessed on January 16, 2024.

- Kazerouni, A.; Aghdam, E. K.; Heidari, M.; Azad, R.; Fayyaz, M.; Hacihaliloglu, I.; and Merhof, D. 2022. Diffusion models for medical image analysis: A comprehensive survey. *arXiv preprint arXiv:2211.07804*.
- Kong, F.; Duan, J.; Ma, R.; Shen, H. T.; Shi, X.; Zhu, X.; and Xu, K. 2024. An Efficient Membership Inference Attack for the Diffusion Model by Proximal Initialization. In *The Twelfth International Conference on Learning Representations*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Luo, Z.; Xu, X.; Liu, F.; Koh, Y. S.; Wang, D.; and Zhang, J. 2024. Privacy-Preserving Low-Rank Adaptation for Latent Diffusion Models. *arXiv preprint arXiv:2402.11989*.
- Lyu, S.; Liu, M. F.; Vinaroz, M.; and Park, M. 2023. Differentially private latent diffusion models. *arXiv preprint arXiv:2305.15759*.
- Manshaei, M. H.; Zhu, Q.; Alpcan, T.; Başçar, T.; and Hubaux, J.-P. 2013. Game theory meets network security and privacy. *ACM Computing Surveys (CSUR)*, 45(3): 1–39.
- Matsumoto, T.; Miura, T.; and Yanai, N. 2023. Membership inference attacks against diffusion models. In *2023 IEEE Security and Privacy Workshops (SPW)*, 77–83. IEEE.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2022. SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Nasr, M.; Shokri, R.; and Houmansadr, A. 2018. Machine learning with membership privacy using adversarial regularization. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, 634–646.
- Pang, Y.; and Wang, T. 2023. Black-box Membership Inference Attacks against Fine-tuned Diffusion Models. *arXiv preprint arXiv:2312.08207*.
- Pang, Y.; Wang, T.; Kang, X.; Huai, M.; and Zhang, Y. 2023. White-box membership inference attacks against diffusion models. *arXiv preprint arXiv:2308.06405*.
- Pinkney, J. N. M. 2022. Pokemon BLIP captions. <https://huggingface.co/datasets/lambdalabs/pokemon-blip-captions/>. Accessed on January 16, 2024.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2024. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In *The Twelfth International Conference on Learning Representations*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22500–22510.
- Shavlokhova, V.; Vollmer, A.; Zouboulis, C. C.; Vollmer, M.; Wollborn, J.; Lang, G.; Kübler, A.; Hartmann, S.; Stoll, C.; Roeder, E.; et al. 2023. Finetuning of GLIDE stable diffusion model for AI-based text-conditional image synthesis of dermoscopic images. *Frontiers in Medicine*, 10.
- Shokri, R.; Stronati, M.; Song, C.; and Shmatikov, V. 2017. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, 3–18. IEEE.
- Shokri, R.; Theodorakopoulos, G.; Troncoso, C.; Hubaux, J.-P.; and Le Boudec, J.-Y. 2012. Protecting location privacy: optimal strategy against localization attacks. In *Proceedings of the 2012 ACM conference on Computer and communications security*, 617–627.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Wu, Y.; Yu, N.; Li, Z.; Backes, M.; and Zhang, Y. 2022. Membership inference attacks against text-to-image generation models. *arXiv preprint arXiv:2210.00968*.
- Xu, J.; Motamed, S.; Vaddamanu, P.; Wu, C. H.; Haene, C.; Bazin, J.-C.; and De la Torre, F. 2024. Personalized Face Inpainting with Diffusion Models by Parallel Visual Attention. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5432–5442.
- Ye, J.; Maddi, A.; Murakonda, S. K.; Bindschaedler, V.; and Shokri, R. 2022. Enhanced membership inference attacks against machine learning models. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 3093–3106.
- Zhang, J.; He, T.; Sra, S.; and Jadbabaie, A. 2019. Why Gradient Clipping Accelerates Training: A Theoretical Justification for Adaptivity. In *International Conference on Learning Representations*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Zhang, M.; Yu, N.; Wen, R.; Backes, M.; and Zhang, Y. 2024. Generated Distributions Are All You Need for Membership Inference Attacks Against Generative Models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 4839–4849.