

Enhanced Privacy Leakage from Noise-Perturbed Gradients via Gradient-Guided Conditional Diffusion Models

Jiayang Meng^{1*}, Tao Huang^{2*}, Hong Chen^{1†}, Chen Hou^{2†}, Guolong Zheng²

¹School of Information, Renmin University of China, Beijing 100872, China

²School of Computer Science and Big Data, Minjiang University, Fuzhou, Fujian 350108, China
{jiayangmeng, chong}@ruc.edu.cn, {huang-tao, houchen, gzheng}@mju.edu.cn

Abstract

Federated learning synchronizes models through gradient transmission and aggregation. However, these gradients pose significant privacy risks, as sensitive training data is embedded within them. Existing gradient inversion attacks suffer from significantly degraded reconstruction performance when gradients are perturbed by noise—a common defense mechanism. In this paper, we introduce gradient-guided conditional diffusion models for reconstructing private images from leaked gradients, without prior knowledge of the target data distribution. Our approach leverages the inherent denoising capability of diffusion models to circumvent the partial protection offered by noise perturbation, thereby improving attack performance under such defenses. We further provide a theoretical analysis of the reconstruction error bounds and the convergence properties of the attack loss, characterizing the impact of key factors—such as noise magnitude and attacked model architecture—on reconstruction quality. Extensive experiments demonstrate our attack’s superior reconstruction performance with Gaussian noise-perturbed gradients, and confirm our theoretical findings.

Code — <https://github.com/mjyyhxx/GGSS-R>

Extended version — <https://arxiv.org/abs/2511.10423>

1 Introduction

Federated Learning (FL) is a prominent paradigm for collaborative model training, which aggregates locally computed gradients from multiple clients without sharing raw private data (Kairouz et al. 2021; Li et al. 2020). However, gradients inherently encode sensitive training data, making them vulnerable to Gradient Inversion Attacks (GIAs) (Zhu, Liu, and Han 2019; Jeon et al. 2021a; Fang et al. 2023a)—where adversaries aim to recover private training data from leaked gradients. This privacy risk is especially critical in sensitive domains such as medical imaging and biometric authentication, where private training data is highly valuable and costly to acquire, making it an attractive target for adversaries.

Existing GIAs based on pixel-space differentiation (Zhu, Liu, and Han 2019; Zhao, Mopuri, and Bilal 2020; Geiping et al. 2020a) struggle with high-resolution images, as

*These authors contributed equally.

†Corresponding authors: Hong Chen and Chen Hou.

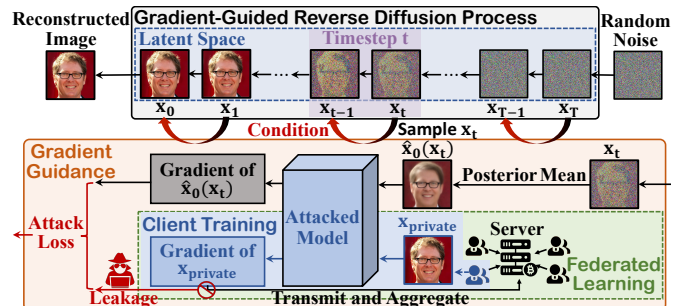


Figure 1: An overview of our proposed gradient-guided conditional diffusion model, with a detailed view of the gradient guidance procedure applied at timestep t .

iterative differentiation of a randomly initialized image becomes computationally prohibitive at higher resolutions. GIAs based on generative models, such as Generative Adversarial Networks (GANs) and diffusion models, require an additional assumption that the pre-training data distribution of the generative model aligns with that of the target images. Furthermore, GAN-based GIAs (Yin et al. 2021a; Jeon et al. 2021b; Li et al. 2022a; Fang et al. 2023b) often suffer from model collapse and training instability, resulting in blurry or artifact-ridden reconstructions. Diffusion model-based GIAs (Ho, Jain, and Abbeel 2020a; Song and Ermon 2019; Song et al. 2020; Meng et al. 2024), which typically rely on model fine-tuning, inherently struggle to precisely recover image details due to the lack of pixel-level guidance.

To mitigate privacy risks associated with gradient leakage, a common strategy is to perturb gradients with noise (e.g., Gaussian or Laplacian) before transmission (Geiping et al. 2020b; Yin et al. 2021b; Li et al. 2022b; Jeon et al. 2021a; Fang et al. 2023a). Empirical studies (Meng et al. 2024; Wang et al. 2022; El Mestari, Lenzini, and Demirci 2024) confirm that this noise injection effectively mitigates GIAs, significantly degrading GIA performance when reconstructing training data from these noise-perturbed gradients.

This paper explores the potential of diffusion models’ denoising capability for partially circumventing protections based on noise perturbation, thereby improving attack performance. Our contributions are summarized as follows.

- We introduce a pixel-level, prior-free GIA method to

reconstruct high-fidelity private images from noise-perturbed gradients. Our proposed attack is implemented through Gradient-Guided Conditional Diffusion Models (GG-CDMs), as illustrated in Figure 1.

- Theoretically, we derive the reconstruction error bounds and characterize the convergence properties of the attack loss, quantifying how key factors—such as noise scale and model architecture—affect the reconstruction fidelity. Furthermore, we explicitly propose the RV metric to quantify a model’s intrinsic vulnerability to GIAs.
- We conducted extensive experiments to validate the effectiveness of our proposed GIA method, with results consistently supporting our theoretical findings.

2 Related Work

2.1 Pixel-Space Differentiation-Based GIAs

Gradient inversion attacks (GIAs) based on pixel-space differentiation (Zhu, Liu, and Han 2019; Zhao, Mopuri, and Bilen 2020; Geiping et al. 2020a) typically initialize a differentiable dummy image and iteratively update it by minimizing an attack loss function. Specifically, (Zhu, Liu, and Han 2019; Zhao, Mopuri, and Bilen 2020) minimize the Mean Squared Error (MSE) for pixel-wise accuracy. (Geiping et al. 2020a) improves the reconstruction fidelity by introducing a carefully designed attack loss function. However, these approaches face significant scalability challenges when reconstructing high-resolution images, due to the inherent computational complexity of differentiation. Furthermore, their performance is severely degraded when the leaked gradients are perturbed by noise, as the target image, which corresponds to these noise-perturbed gradients, is distorted.

2.2 Generative Model-Based GIAs

Recent advances leverage generative models for high-resolution image reconstruction, avoiding the computational burden of differentiation-based methods. Specifically, (Yin et al. 2021a; Jeon et al. 2021b; Li et al. 2022a; Fang et al. 2023b) have demonstrated that Generative Adversarial Networks (GANs) can serve as effective priors by approximating the natural image manifold, thus improving GIA performance. However, GAN-based methods require distributional alignment between the target image and the GAN’s training set; moreover, they often suffer from training instability. Diffusion models provide a promising alternative, with (Meng et al. 2024) demonstrating that model fine-tuning enables stable and noise-free reconstructions. Nevertheless, this fine-tuning paradigm suffers from several limitations: (i) high computational cost due to its iterative nature, (ii) reliance on the distributional prior of target data, (iii) model overfitting that degrades generation quality, and (iv) the absence of pixel-level guidance for fine-grained detail recovery.

3 Preliminaries

3.1 Diffusion Models

Diffusion models (Nichol and Dhariwal 2021; Rombach et al. 2022) are a class of generative models capable of producing high-quality and diverse samples. In the forward

diffusion process, diffusion models gradually perturb clean data $\mathbf{x}_0 \sim p_{\text{data}}$ by adding Gaussian noise until it becomes pure noise. For both DDPM (Ho, Jain, and Abbeel 2020b) and DDIM (Song, Meng, and Ermon 2020), the posterior distribution of any $\mathbf{x}_t, t \in [0, T]$ given \mathbf{x}_0 is defined as:

$$\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \epsilon_t, \epsilon_t \sim \mathcal{N}(0, I), \quad (1)$$

where $\{\alpha_t\}_{t=1}^T$ is the pre-defined noise schedule. The objective is to train a Denoiser, $\epsilon_\theta(\mathbf{x}_t, t)$, to predict the noise ϵ_t added at each timestep t . This is achieved by minimizing the expected discrepancy between the true and predicted noise:

$$L_{\text{Denoiser}} = \mathbb{E}_{t, \mathbf{x}_0, \epsilon_t} [\|\epsilon_t - \epsilon_\theta(\mathbf{x}_t, t)\|^2]. \quad (2)$$

During the reverse diffusion process, the joint distribution $p(\mathbf{x}_{0:T})$ is given by:

$$p(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t), \quad (3)$$

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t)). \quad (4)$$

Specifically, our attack follows DDIM’s reverse sampling:

$$\begin{aligned} \mathbf{x}_{t-1} = & \sqrt{\alpha_{t-1}} \underbrace{\left(\frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\alpha_t}} \right)}_{\text{predicted } \mathbf{x}_0} \\ & + \underbrace{\sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \epsilon_\theta(\mathbf{x}_t, t)}_{\text{direction pointing to } \mathbf{x}_t} + \underbrace{\sigma_t \epsilon_t}_{\text{random noise}}, \end{aligned} \quad (5)$$

where $\sigma_t = \eta \sqrt{\frac{1 - \alpha_{t-1}}{1 - \alpha_t}} \sqrt{1 - \frac{\alpha_t}{\alpha_{t-1}}}$, and $\epsilon_t \sim \mathcal{N}(0, I)$.

3.2 Conditional Diffusion Models

Conditional diffusion models (Dhariwal and Nichol 2021; Ho and Salimans 2022; Chung et al. 2022b; Yu et al. 2023; He et al. 2023; Yang et al. 2024) incorporate a given condition or measurement \mathbf{y} by introducing an additional likelihood term $p(\mathbf{x}_t | \mathbf{y})$. According to the Bayes’ rule,

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t | \mathbf{y}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \mathbf{x}_t). \quad (6)$$

With DDIM as the prior, the term $\nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \mathbf{x}_t)$ can be utilized to perform an additional correction step:

$$\mathbf{x}_{t-1} = \underbrace{\text{DDIM}(\mathbf{x}_t, \epsilon_\theta(\mathbf{x}_t, t))}_{\text{reverse sampling step}} - \underbrace{\gamma \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \mathbf{x}_t)}_{\text{conditional correction step}}. \quad (7)$$

4 Methodology

4.1 Problem Statement

In federated learning, consider a client-side model $F(\mathbf{x}; W)$ parameterized by W under attack. For a private local image $\mathbf{x}_{\text{private}} \in \mathbb{X}$, the corresponding gradient of $F(\mathbf{x}; W)$ is:

$$\mathbf{g}(\mathbf{x}_{\text{private}}) = \nabla_W F(\mathbf{x}_{\text{private}}; W). \quad (8)$$

Under a defense using Gaussian noise perturbation, the leaked gradient $\mathbf{g}_{\text{leaked}} = \mathbf{g}(\mathbf{x}_{\text{private}}) + \mathcal{N}(0, \sigma^2 I) = \mathcal{N}(\mathbf{g}(\mathbf{x}_{\text{private}}), \sigma^2 I)$ follows a Gaussian distribution. Given leaked gradient $\mathbf{g}_{\text{leaked}}$, access to the attacked model F (allowing gradient computation for arbitrary inputs), and a diffusion model $\epsilon_\theta(\mathbf{x})$ pre-trained on an arbitrary dataset without requirement of distributional alignment with $\mathbf{x}_{\text{private}}$, the objective is to generate an image \mathbf{x} such that $\mathbf{g}(\mathbf{x})$ approximates $\mathbf{g}(\mathbf{x}_{\text{private}})$, thereby ensuring \mathbf{x} approximates $\mathbf{x}_{\text{private}}$.

4.2 Posterior Distribution Approximation

To achieve computationally efficient and fine-grained pixel-level reconstruction, we guide the DDIM reverse sampling process using leaked gradients as a conditioning signal. This approach enables GIAs within a single reverse diffusion process. A key challenge in this formulation is to derive $p(\mathbf{y} = \mathbf{g}_{\text{leaked}}|\mathbf{x}_t)$ in Equation (7), as \mathbf{x}_t is time-dependent. A tractable approximation can be factorized as:

$$p(\mathbf{y} = \mathbf{g}_{\text{leaked}}|\mathbf{x}_t) = \int p(\mathbf{y} = \mathbf{g}_{\text{leaked}}|\mathbf{x}_0)p(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0. \quad (9)$$

From DDIM’s forward diffusion sampling (Equation (1)), the posterior mean $\hat{\mathbf{x}}_0$ conditioned on \mathbf{x}_t can be obtained as:

$$\begin{aligned} \hat{\mathbf{x}}_0(\mathbf{x}_t) &:= \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t] = \mathbb{E}_{\mathbf{x}_0 \sim p(\mathbf{x}_0|\mathbf{x}_t)}[\mathbf{x}_0] \\ &= \frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - (1 - \alpha_t)\epsilon_\theta(\mathbf{x}_t, t)). \end{aligned} \quad (10)$$

We then approximate:

$$p(\mathbf{g}_{\text{leaked}}|\mathbf{x}_t) \simeq p(\mathbf{g}_{\text{leaked}}|\hat{\mathbf{x}}_0(\mathbf{x}_t)). \quad (11)$$

The error induced by the posterior mean approximation is naturally quantified by the Jensen gap.

Definition 4.1 (Jensen Gap (Chung et al. 2022a)). Let \mathbf{x} be a random variable with distribution $p(\mathbf{x})$. For some function f that may or may not be convex, Jensen gap is defined as:

$$\mathcal{J}(f, \mathbf{x} \sim p(\mathbf{x})) = \mathbb{E}[f(\mathbf{x})] - f(\mathbb{E}[\mathbf{x}]), \quad (12)$$

where the expectation is taken over $p(\mathbf{x})$.

Specifically, the reconstruction error introduced by the posterior mean approximation is given by $\mathbb{E}_{\mathbf{x}_0|\mathbf{x}_t}[f(\mathbf{x}_0)] - f(\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t])$, where f is a function such as the gradient computation. A smaller Jensen gap implies a lower approximation error and thus higher reconstruction fidelity.

4.3 Attack Methodology

In this section, we propose the *Gradient-guided Gaussian Spherical Sampling-Reconstruction (GGSS-R)*, a novel GIA method that leverages Gradient-Guided Conditional Diffusion Models (GG-CDMs) to reconstruct private images from leaked gradients via a single reverse diffusion process.

We first introduce a symmetric and differentiable attack loss function $\mathcal{L}(\cdot, \cdot) = \|\cdot - \cdot\|$. Given the leaked gradient $\mathbf{g}_{\text{leaked}}$, the reverse sampling at timestep t is formulated as:

$$\arg \min_{\mathbf{x}_{t-1}} [\nabla_{\mathbf{x}_t} \mathcal{L}(\mathbf{g}(\hat{\mathbf{x}}_0(\mathbf{x}_t)), \mathbf{g}_{\text{leaked}})]^T (\mathbf{x}_{t-1} - \mathbf{x}_t), \quad (13)$$

where $\mathbf{x}_{t-1} \in CI_{1-\delta}$, with $CI_{1-\delta}$ denoting the $1 - \delta$ confidence intervals for the Gaussian distribution defined in Equation (4). This constrained optimization aims to steer the reverse sampling toward directions that most rapidly minimize the attack loss while enforcing sampling within high-confidence intervals under the DDIM Gaussian distribution.

Due to the high-dimensional nature of gradients, direct optimization is challenging. Previous theoretical work (Ledoux 2001) has established that Gaussian-distributed variables exhibit concentration around their mean. The Laurent-Massart bound (Laurent and Massart 2000) further provides bounds on the deviation of a Gaussian variable from its mean, as formalized in Lemma 4.1.




Figure 2: An overview of GGSS at timestep t , mapping from data manifold M_t to M_{t-1} . The green annulus represents the concentration region of samples in unconditional DDIM reverse sampling, and the blue annulus denotes the concentration region under accurate conditional guidance.

Algorithm 1: GGSS-R Process.

Input: Pre-trained diffusion model ϵ_θ with maximum timestep T ; Pre-defined noise schedule $\{\alpha_t\}_{t=1}^T$ and denoising scale $\{\sigma_t\}_{t=1}^T$; Attacked model $F(\mathbf{x}; W)$ where $\mathbf{x} \in \mathbb{R}^n$; Leaked gradient $\mathbf{g}_{\text{leaked}}$; Attack loss function \mathcal{L} .

Output: Reconstructed image \mathbf{x}_0 .

- 1: Randomly initialize a dummy image: $\mathbf{x}_T \sim \mathcal{N}(0, I)$.
- 2: **for** $t = T$ to 1 **do**
- 3: Predict added noise at forward step t : $\hat{\epsilon}_t \leftarrow \epsilon_\theta(\mathbf{x}_t, t)$.
- 4: Approximate the posterior mean:
 $\hat{\mathbf{x}}_0(\mathbf{x}_t) \leftarrow \frac{1}{\sqrt{\alpha_t}} (\mathbf{x}_t - (1 - \alpha_t)\hat{\epsilon}_t)$.
- 5: Compute the mean of \mathbf{x}_{t-1} :
 $\mu_\theta(\mathbf{x}_t, t) = \sqrt{\alpha_{t-1}}\hat{\mathbf{x}}_0(\mathbf{x}_t) + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \hat{\epsilon}_t$.
- 6: Input $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ to F to calculate $\mathbf{g}(\hat{\mathbf{x}}_0(\mathbf{x}_t))$.
- 7: Perform gradient-guided conditional correction:
 $\mathbf{x}_{t-1} \leftarrow \mu_\theta(\mathbf{x}_t, t) - \sqrt{n}\sigma_t \cdot \frac{\nabla_{\mathbf{x}_t} \mathcal{L}(\mathbf{g}(\hat{\mathbf{x}}_0(\mathbf{x}_t)), \mathbf{g}_{\text{leaked}})}{\|\nabla_{\mathbf{x}_t} \mathcal{L}(\mathbf{g}(\hat{\mathbf{x}}_0(\mathbf{x}_t)), \mathbf{g}_{\text{leaked}})\|}$.
- 8: **end for**
- 9: **return** Reconstructed image \mathbf{x}_0 .

Lemma 4.1 (Laurent-Massart Bound (Laurent and Massart 2000)). For an n -dimensional isotropy Gaussian distribution $\mathbf{x} \sim \mathcal{N}(\mu, \sigma^2 I)$, it satisfies:

$$P(\|\mathbf{x} - \mu\|^2 \geq n\sigma^2 + 2n\sigma^2(\sqrt{\epsilon} + \epsilon)) \leq e^{-n\epsilon}, \quad (14)$$

$$P(\|\mathbf{x} - \mu\|^2 \leq n\sigma^2 + 2n\sigma^2\sqrt{\epsilon}) \leq e^{-n\epsilon}, \quad (15)$$

where n is \mathbf{x} 's dimension.

When n is large, $e^{-n\epsilon}$ approaches 0. In this case, \mathbf{x} is distributed on a sphere with radius $\sqrt{n}\sigma$ centered at μ . The distribution of \mathbf{x} follows: $\mathbf{x} \sim S_{\mu_\theta(\mathbf{x}_t, t), \sqrt{n}\sigma_t}^n = \{\mathbf{x} \in \mathbb{R}^n, \|\mathbf{x} - \mu_\theta(\mathbf{x}_t, t)\| = \sqrt{n}\sigma_t\}$. With this spherical constraint, the optimization problem has a closed-form solution:

$$\mathbf{x}_{t-1}^* = \mu_\theta(\mathbf{x}_t, t) - \sqrt{n}\sigma_t \frac{\nabla_{\mathbf{x}_t} \mathcal{L}(\mathbf{g}(\hat{\mathbf{x}}_0(\mathbf{x}_t)), \mathbf{g}_{\text{leaked}})}{\|\nabla_{\mathbf{x}_t} \mathcal{L}(\mathbf{g}(\hat{\mathbf{x}}_0(\mathbf{x}_t)), \mathbf{g}_{\text{leaked}})\|}. \quad (16)$$

We term this sampling as *Gradient-guided Gaussian Spherical Sampling (GGSS)*. Figure 2 provides an illustration of GGSS at timestep t , and Algorithm 1 outlines the GGSS-R.

In practice, deterministic sampling under a Gaussian spherical constraint minimizes the attack loss most rapidly.

However, noise-perturbed gradients make the corresponding target image non-smooth and deviate from the original target. To address this, we combine DDIM’s unconditional sampling direction with the optimal conditional sampling direction, thereby enhancing both generation quality and diversity. Specifically, line 7 in Algorithm 1 is refined as:

$$\begin{aligned} \mathbf{x}_{t-1} &= \mu_\theta(\mathbf{x}_t, t) + r \frac{d_m}{\|d_m\|}, \\ d_m &= d^{\text{sample}} + m_r(d^* - d^{\text{sample}}), \end{aligned} \quad (17)$$

where $d^* = -\sqrt{n}\sigma_t \cdot \frac{\nabla_{\mathbf{x}_t} \mathcal{L}(\mathbf{g}(\hat{\mathbf{x}}_0), \mathbf{g}^{\text{leaked}})}{\|\nabla_{\mathbf{x}_t} \mathcal{L}(\mathbf{g}(\hat{\mathbf{x}}_0), \mathbf{g}^{\text{leaked}})\|}$ represents the optimal adjustment step derived from the leaked gradient $\mathbf{g}^{\text{leaked}}$, while $d^{\text{sample}} = \sigma_t \epsilon_t$ (with $\epsilon_t \sim \mathcal{N}(0, I)$) denotes DDIM’s stochastic unconditional sampling step. Their weighted combination, d_m , is controlled by the guidance rate $m_r \in [0, 1]$, which balances two objectives, and r dictates the overall step size of the conditional guidance.

This refined guidance mechanism strategically balances conditional guidance with sample quality, not only ensuring adherence to gradient-based conditions but also promoting smoother, higher-fidelity image generation. By exploiting DDIM’s inherent generative capability, we mitigate the adverse effects of noise-induced perturbations, potentially generating reconstructions closer to the original target.

5 Theoretical Analysis

This section establishes theoretical upper and lower bounds on the reconstruction error arising from the posterior mean approximation, which is quantified by the Jensen gap in Section 4.2. In addition, we analyze the convergence behavior of the attack loss. Furthermore, our analysis reveals the significant variation in vulnerability to GIAs across model architectures and proposes the RV metric to quantify a model’s intrinsic vulnerability to such attacks. Detailed proofs of Theorems are provided in Appendix A of the extended version.

5.1 Bounds of the Reconstruction Error

For Algorithm 1, Theorem 5.1 establishes the upper bound of the reconstruction error.

Theorem 5.1 (Upper Bound of the Reconstruction Error). *Under the attacked model $F(\mathbf{x}; W)$, we assume that the adversary obtains a Gaussian noise-perturbed gradient $\mathbf{g}^{\text{leaked}} = \nabla_W F(\mathbf{x}_{\text{private}}; W) + \mathcal{N}(0, \sigma^2 I)$. The gradient inversion attack follows Algorithm 1, using the approximation $p(\mathbf{g}^{\text{leaked}} | \mathbf{x}_t) \simeq p(\mathbf{g}^{\text{leaked}} | \hat{\mathbf{x}}_0(\mathbf{x}_t))$, where $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ denotes the posterior mean of \mathbf{x}_0 given \mathbf{x}_t . The reconstruction error—measured by the Jensen gap—is upper bounded by:*

$$\begin{aligned} \mathcal{J}(\mathbf{g}(\mathbf{x}), p(\mathbf{x}_0 | \mathbf{x}_t)) &= |\mathbb{E}[\mathbf{g}(\mathbf{x}_0)] - \mathbf{g}(\mathbb{E}[\mathbf{x}_0])| \\ &\leq \frac{n}{\sqrt{2\pi}\sigma^2} \|\nabla_{\mathbf{x}} \mathbf{g}(\mathbf{x})\| \int \|\mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t)\| p(\mathbf{x}_0 | \mathbf{x}_t) d\mathbf{x}_0, \end{aligned}$$

where n is \mathbf{x} ’s dimension, and $\mathbf{g}(\mathbf{x}) = \nabla_{\mathbf{x}} \nabla_W F(\mathbf{x}; W)$.

The upper bound of the reconstruction error characterizes the worst-case reconstruction quality. The term $\int \|\mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t)\| p(\mathbf{x}_0 | \mathbf{x}_t) d\mathbf{x}_0$, which depends on the pre-trained diffusion model, is sensitive to the noise magnitude σ^2 . Our

analysis further identifies an architecture-dependent factor, $\|\nabla_{\mathbf{x}} \mathbf{g}(\mathbf{x})\| = \|\nabla_{\mathbf{x}} \nabla_W F(\mathbf{x}; W)\|$. To quantify the vulnerability of a model F to GIAs, we introduce the Reconstruction Vulnerability (RV) metric, defined as the Frobenius norm of the upper bound on this quantity as follows.

Definition 5.1 (Reconstruction Vulnerability (RV)). *Let $F(\mathbf{x}; W)$ be a model parameterized by W , where \mathbf{x} denotes the input. The Reconstruction Vulnerability (RV) of F with respect to a private dataset \mathbb{X} is formally defined as:*

$$RV = \max_{\mathbf{x} \in \mathbb{X}} \|\nabla_{\mathbf{x}} \nabla_W F(\mathbf{x}; W)\|_F, \quad (18)$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

A larger RV value indicates a greater sensitivity of the parameter gradients to input variations, which amplifies input-specific leakage and facilitates higher-fidelity GIAs. As directly computing the RV is computationally intensive, we project $\nabla_W F(\mathbf{z}; W)$ onto M random orthogonal directions $\{\mathbf{v}_j\}_{j=1}^M$, where $\mathbf{v}_j \sim \mathbf{p}_{\mathbf{v}}$. Given a model F trained on a dataset $\mathbb{X} = \{\mathbf{x}_i\}_{i=1}^{|\mathbb{X}|}$, $\mathbf{x}_i \sim \mathbf{p}_{\text{data}}$, RV can be estimated as:

$$\begin{aligned} RV &= \mathbb{E}_{\mathbf{v} \sim \mathbf{p}_{\mathbf{v}}} \mathbb{E}_{\mathbf{x} \sim \mathbf{p}_{\text{data}}} \|\nabla_{\mathbf{x}} (\mathbf{v}^T \nabla_W F(\mathbf{x}; W))\| \\ &\simeq \frac{1}{N} \frac{1}{M} \sum_{i=1}^N \sum_{j=1}^M \|\nabla_{\mathbf{x}} (\mathbf{v}_j^T \nabla_W F(\mathbf{x}_i; W))\| \end{aligned} \quad (19)$$

We then examine the lower bound of the reconstruction error. As shown in Equation (16), the properties of the attack loss \mathcal{L} influence the lower bound of the Jensen gap. Let $\mathcal{L}(\mathbf{x}_t) = \mathcal{L}(\mathbf{g}(\mathbf{x}_t)) = \mathcal{L}(\mathbf{g}(\hat{\mathbf{x}}_0(\mathbf{x}_t)), \mathbf{g}^{\text{leaked}})$ for simplicity. We rely on Assumptions 5.1-5.3 to derive the lower bound.

Assumption 5.1. $\mathcal{L}(\mathbf{g}(\mathbf{x}))$ is α -smooth and β -convex with respect to $\mathbf{g}(\mathbf{x})$;

Assumption 5.2. $\mathbf{g}(\mathbf{x})$ is L_g -smooth with respect to \mathbf{x} ;

Assumption 5.3. Jacobian $\mathbf{J}_{\mathbf{g}}(\mathbf{x}) = \nabla_{\mathbf{x}} \mathbf{g}(\mathbf{x})$ is full rank.

Theorem 5.2 (Lower Bound of the Reconstruction Error). *Under conditions of Theorem 5.1 and Assumptions 5.1-5.3, for $\mathbf{x} \in \mathbb{R}^n$, the reconstruction error is lower bounded by:*

$$\begin{aligned} \mathcal{J}(\mathcal{L}(\mathbf{g}(\mathbf{x}), \mathbf{g}^{\text{leaked}}), p(\mathbf{x}_0 | \mathbf{x}_t)) &= |\mathbb{E}[\mathcal{L}(\mathbf{g}(\mathbf{x}_0), \mathbf{g}^{\text{leaked}})] - \mathcal{L}(\mathbf{g}(\mathbb{E}[\mathbf{x}_0]), \mathbf{g}^{\text{leaked}})| \\ &\geq \frac{1}{2} (\beta \lambda_{\min}(\mathbf{J}_{\mathbf{g}}(\mathbf{x})^T \mathbf{J}_{\mathbf{g}}(\mathbf{x})) - \alpha L_g) \sum_{i=1}^n \sigma_i^2, \end{aligned}$$

where $\lambda_{\min}(\cdot)$ represents the smallest eigenvalue of \cdot , L_g can be refined to a specific value $\frac{n}{\sqrt{2\pi}\sigma^2} \exp(-\frac{1}{2\sigma^2})$ under Gaussian noise-perturbed isotropic gradients, and $\{\sigma_i^2\}_{i=1}^n$ are the eigenvalues of the covariance matrix of $p(\mathbf{x}_0 | \mathbf{x}_t)$.

In Theorem 5.2, $\lambda_{\min}(\mathbf{J}_{\mathbf{g}}(\mathbf{x})^T \mathbf{J}_{\mathbf{g}}(\mathbf{x}))$ can be interpreted as the information loss associated with the mapping $\mathbf{x} \mapsto \mathbf{g}$. This suggests that reconstructing an exact replica of the original private image from leaked gradients—even without explicit perturbation noise—is impossible. The degree of information loss is scaled by the randomness inherent in the DDIM’s denoising steps, specifically by $\sum_{i=1}^n \sigma_i^2$. Moreover, Theorem 5.2 demonstrates that a larger Gaussian noise variance σ^2 monotonically decreases L_g , thus increasing the lower bound on the Jensen gap. This indicates a positive correlation between the reconstruction error and the noise scale.

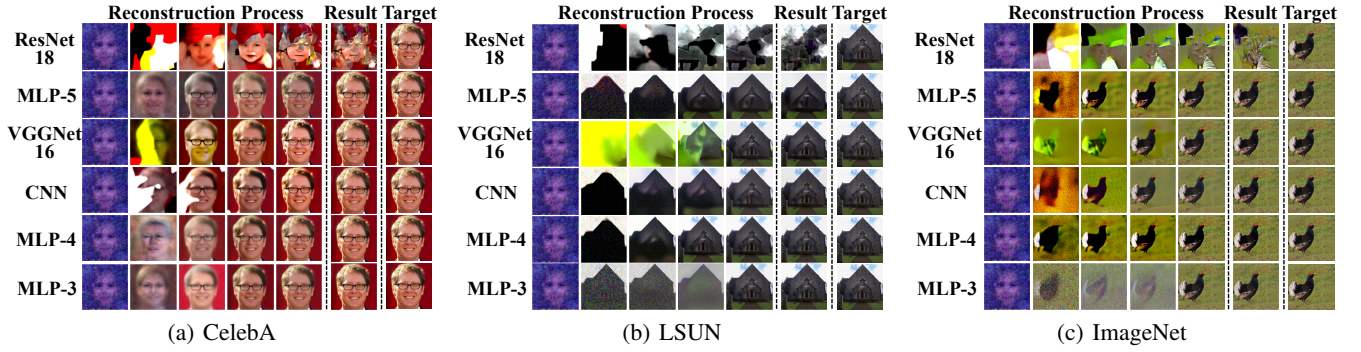


Figure 3: Reconstruction process of GGSS-R across diverse datasets and attacked models. The *Reconstruction Process* columns display the intermediate results $\hat{\mathbf{x}}_0(\mathbf{x}_t)$, with final reconstructions and target image presented in the *Result* and *Target* columns.

5.2 Convergence of the Attack Loss

In the context of GIA, a critical issue is whether the attack loss decreases as the reconstructed samples are generated progressively. Formally, we examine whether the sequence $\{\mathcal{L}(\mathbf{x}_t)\}_{t=1}^T$ is monotonically decreasing. If this holds, it is essential to characterize the rate of this decrease. Specifically, Theorem 5.3 states the monotonic decrease of the attack loss, providing a lower bound on its decrease rate, and Theorem 5.4 derives an upper bound on this rate.

Theorem 5.3 (Convergence of the Attack Loss). *Assume that the attack loss $\mathcal{L}(\mathbf{x})$ is α' -smooth, β' -strongly convex with respect to \mathbf{x} . When GIA is executed following Algorithm 1, the sequence $\{\mathcal{L}(\mathbf{x}_t)\}_{t=1}^T$ monotonically decreases as the denoising steps progress from timestep $t = T$ to 1.*

Theorem 5.4 (Upper Bound on the Convergence Rate of the Attack Loss). *Under Assumptions 5.1-5.3, and assume that the attack loss $\mathcal{L}(\mathbf{x})$ is α' -smooth, β' -strongly convex with respect to \mathbf{x} . The optimization using Algorithm 1 satisfies:*

$$\mathcal{L}(\mathbf{x}_{t-1}) - \mathcal{L}(\mathbf{x}_t) \geq \frac{\beta' n \sigma_t^2}{2} + \beta' k \sqrt{n} \sigma_t + \frac{\beta' k^2}{2} - (k + \sqrt{n} \sigma_t) (\sqrt{\lambda_{\max}(\mathbf{J}_{\mathbf{g}}(\mathbf{x}) \mathbf{J}_{\mathbf{g}}(\mathbf{x})^T)} \cdot L_g),$$

where n is \mathbf{x} 's dimension, k is a positive constant, $\lambda_{\max}(\cdot)$ is the largest eigenvalue of \cdot , and $L_g = \frac{n}{\sqrt{2\pi\sigma^2}} \exp(-\frac{1}{2\sigma^2})$.

As discussed in Theorem 5.2, the private information loss associated with the mapping $\mathbf{x} \mapsto \mathbf{g}$ is quantified by the term $\sqrt{\lambda_{\max}(\mathbf{J}_{\mathbf{g}}(\mathbf{x}) \mathbf{J}_{\mathbf{g}}(\mathbf{x})^T)}$. The convergence rate of the attack loss is further influenced by L_g , with $L_g \leq \frac{n}{\sqrt{2\pi\sigma^2}}$. Consequently, increasing the noise variance σ^2 tightens the lower bound on $\mathcal{L}(\mathbf{x}_{t-1}) - \mathcal{L}(\mathbf{x}_t)$, which, in turn, slows the convergence of attack loss and deteriorates reconstruction quality.

6 Experiments

6.1 Experimental Setup

Datasets and Models. Experiments were conducted on three datasets: CelebA (Karras et al. 2017a), LSUN (Yu et al. 2015), and ImageNet (Deng et al. 2009), with images of 256×256 resolution. The attacked models include MLPs

with 3/4/5 layers (MLP-3/4/5), CNN, ResNet18 (He et al. 2016), and VGGNet16 (Simonyan and Zisserman 2014). The pre-trained diffusion model, sourced from (Chung et al. 2022a), was pre-trained on the FFHQ face dataset (Karras et al. 2017b), processing 256×256 resolution images.

Hyperparameters. We employ conditional reverse sampling in Equation (17), a common practice. We set a gradient acquisition batch size of 1 and a guidance rate m_r of 0.20. The attack loss is measured using the Euclidean distance. All RV values are computed with $M = 1000$ and $N = 310$.

Evaluation Metrics. Reconstruction quality is evaluated using three widely adopted computer vision metrics: MSE, PSNR, and LPIPS, which quantify the similarity between the reconstructed samples and the original target images.

Baselines. We compared our GGSS-R with state-of-the-art baselines, including DLG (Zhu, Liu, and Han 2019), IG (Geiping et al. 2020b), GI (Yin et al. 2021b), GIAS (Jeon et al. 2021a), GGL (Li et al. 2022b), GIFD (Fang et al. 2023a), and FinetunedDiff (Meng et al. 2024).

More detailed experimental settings are provided in Appendix B of the extended version.

6.2 Experimental Results

GIA from Original Gradients and RV Metric Validation.

This section presents reconstruction results obtained with original, unperturbed gradients. For each attacked model, we first compute its RV value and then perform our GGSS-R attack. Figure 3 visually illustrates the reconstruction results, with attacked models ordered by increasing RV values. Table 1 further quantifies the reconstruction performance.

Based on these results, we highlight two key findings: (i) Robustness across diverse distributions. Despite diffusion models being pre-trained on the FFHQ face dataset, GGSS-R effectively reconstructs high-fidelity images across diverse distributions. This demonstrates our attack's robustness, enabling efficient GIAs without requiring prior knowledge of the target data distribution—a common limitation in existing generative model-based approaches; (ii) RV as an effective indicator of a model's vulnerability to GIAs. Models with higher RV values consistently exhibit higher recon-

Dataset	Model	RV	MSE↓	PSNR↑	LPIPS↓
CelebA	ResNet18	0.17	9.46e-2	10.79	4.14e-4
	MLP-5	8.71	3.13e-5	45.05	9.84e-8
	VGGNet16	10.74	1.26e-4	40.76	2.43e-7
	CNN	18.51	7.72e-5	41.12	2.05e-7
	MLP-4	20.85	3.08e-5	45.12	8.92e-8
	MLP-3	52.73	4.55e-5	43.42	1.60e-7
LSUN	ResNet18	0.17	1.00e-1	9.99	3.10e-4
	MLP-5	8.71	2.27e-3	26.44	3.09e-5
	VGGNet16	10.74	6.63e-3	22.79	4.88e-5
	CNN	18.51	2.54e-3	25.95	2.44e-5
	MLP-4	20.85	3.61e-4	34.43	5.60e-7
ImageNet	MLP-3	52.73	1.48e-4	38.29	4.03e-7
	ResNet18	0.17	4.45e-2	13.52	2.54e-4
	MLP-5	8.71	3.34e-4	25.35	3.58e-5
	VGGNet16	10.74	1.07e-3	29.71	1.48e-6
	CNN	18.51	1.15e-4	27.39	2.55e-5
	MLP-4	20.85	3.09e-4	25.10	4.05e-5
	MLP-3	52.73	1.05e-4	28.00	1.63e-5

Table 1: Quantitative analysis of the reconstruction quality across various datasets and attacked models.

struction quality, indicating an elevated privacy risk. Specifically, we observe a convex increase in reconstruction performance with RV. This finding suggests designing models with lower RV values, as they inherently present lower gradient privacy risks. Furthermore, architectural modifications to low-RV models warrant particular caution, as even minor increases in RV can significantly increase privacy risks.

GIA from Gaussian Noise-Perturbed Gradients. When the leaked gradient is perturbed by Gaussian noise, the target corresponding to $\mathbf{g}_{\text{leaked}}$ becomes non-smooth and deviates from the original. We define this distorted target as $\mathbf{x}_{\text{noisy}}$. This section investigates the effectiveness of the added noise in impeding our attack and validates our theoretical findings.

Figure 4 illustrates the reconstruction process on the CelebA dataset under varying Gaussian noise magnitudes σ^2 added to CNN gradients. It reveals that the PSNR between the intermediate reconstructions $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ and the original target $\mathbf{x}_{\text{private}}$ exhibits a non-monotonic behavior, peaking before subsequently declining. This phenomenon is attributed to the inherent denoising capability of diffusion models. Rather than directly converging from a randomly initialized dummy image to the distorted target $\mathbf{x}_{\text{noisy}}$, the reconstruction process produces smoother intermediate results that are more accurate approximations of $\mathbf{x}_{\text{private}}$. However, as reconstruction iterations progress, diffusion model’s denoising capability becomes insufficient to counteract the noise inherent in $\mathbf{x}_{\text{noisy}}$. Consequently, in later stages, the reconstruction converges toward $\mathbf{x}_{\text{noisy}}$ and deviates from $\mathbf{x}_{\text{private}}$, leading to the decline in GIA performance. The quantitative results in Table 2 corroborate these observations.

Experimental results further corroborate our theoretical findings: (i) Effect of noise scale on reconstruction error. Reconstruction quality degrades noticeably with increasing noise levels. This observation aligns with our theoretical analysis (Theorems 5.1-5.2), which establishes that larger noise levels make accurate reconstruction provably more

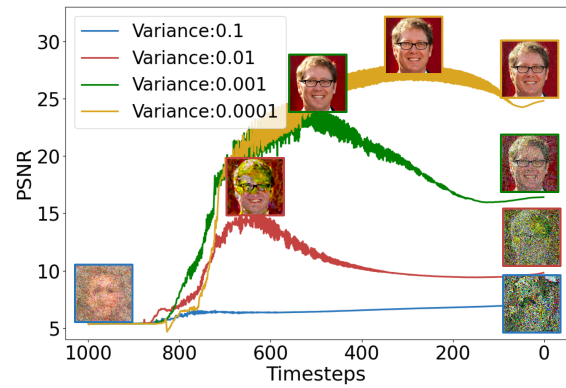


Figure 4: Reconstruction processes on CelebA with Gaussian noise-perturbed CNN gradients. Reconstructions at peak performance and upon GIA completion are displayed.

Gaussian Variance	MSE↓	PSNR↑	LPIPS↓
10^{-4}	0.0016	28.0343	7.9935e-6
10^{-3}	0.0039	24.0567	4.5263e-5
10^{-2}	0.0366	14.9784	2.4762e-4
10^{-1}	0.1252	9.0236	7.1237e-4

Table 2: Quantitative analysis on reconstruction quality on CelebA with Gaussian noise-perturbed CNN gradients. Metrics are reported between the intermediate reconstructions $\hat{\mathbf{x}}_0(\mathbf{x}_t)$ at peak performance and the original target $\mathbf{x}_{\text{private}}$.

challenging; (ii) Reduction of attack loss during the reconstruction process. According to Theorem 5.3, the attack loss in Algorithm 1 decreases monotonically as the process evolves from random noise towards the original target $\mathbf{x}_{\text{private}}$, achieving optimal reconstruction upon convergence. In our implementation, line 7 of Algorithm 1 is modified to Equation (17), leveraging the DDIM generative prior to improve reconstruction smoothness and quality. This modification yields earlier PSNR peaks, though reconstruction quality eventually declines due to the noise-induced gradient perturbations; (iii) Convergence rate of attack loss. Theorem 5.4 states that the attack loss reduction rate is governed by the noise scale: higher noise levels yield smaller decreases $\mathcal{L}(\mathbf{x}_{t-1}) - \mathcal{L}(\mathbf{x}_t)$, leading to slower convergence and poorer final reconstructions. This trend is empirically validated in Figure 4, which exhibits a clear correlation between increased noise variance and both slower PSNR improvement and reduced reconstruction fidelity.

Comparison with Baselines. To demonstrate the superior performance of GGSS-R, we conducted a comprehensive comparison with state-of-the-art baselines on CelebA, using CNN as the attacked model. Table 3 reports the quantitative results of GIAs from the original gradients. As shown in Table 3, our approach successfully guides the diffusion model to reconstruct high-resolution images with markedly improved fidelity, underscoring the heightened privacy risks when adversaries gain access to the original gradients.

We further extended the comparison to assess reconstruc-

Method	MSE↓	PSNR↑	LPIPS↓
DLG (Zhu, Liu, and Han 2019)	0.0480	13.1835	2.3453e-4
IG (Geiping et al. 2020b)	0.0196	17.0756	1.1612e-4
GI (Yin et al. 2021b)	0.0223	16.5109	1.4274e-4
GIAS (Jeon et al. 2021a)	0.0458	13.3885	3.9351e-4
GGL (Li et al. 2022b)	0.0179	17.4923	1.1937e-4
GIFD (Fang et al. 2023a)	0.0098	20.0534	6.6672e-5
FinetunedDiff (Meng et al. 2024)	0.0011	29.4490	5.4441e-6
GGSS-R (Ours)	0.0001	41.1229	2.0519e-7

Table 3: Quantitative analysis of baselines and our method with original, unperturbed gradients.

Method	MSE↓	PSNR↑	LPIPS↓
DLG (Zhu, Liu, and Han 2019)	0.2194	6.5871	1.8540e-3
IG (Geiping et al. 2020b)	0.0747	11.2677	3.2103e-4
GI (Yin et al. 2021b)	0.0892	10.4968	5.6760e-4
GIAS (Jeon et al. 2021a)	0.0613	12.1276	3.4494e-4
GGL (Li et al. 2022b)	0.0394	13.8982	2.5247e-4
GIFD (Fang et al. 2023a)	0.0425	13.7118	2.5901e-4
FinetunedDiff (Meng et al. 2024)	0.0683	11.6675	3.0877e-4
GGSS-R (Ours)	0.0366	14.9784	2.4762e-4

Table 4: Quantitative analysis of baselines and our method with Gaussian noise ($\sigma^2 = 0.01$)-perturbed gradients.

tion performance using CNN gradients perturbed by Gaussian noise ($\sigma^2 = 0.01$). As detailed in Table 4, our methods demonstrate improved robustness against noisy gradients. Specifically, GGSS-R achieves the highest PSNR, surpassing the optimal baseline, GGL, by 1.0802 dB. These results underscore that simple noise addition is insufficient to ensure gradient privacy, highlighting the need for more sophisticated mitigation strategies in practical applications.

6.3 Ablation Studies

Impact of Batch Size. We analyze the effect of gradient acquisition batch size \mathcal{B} on GIA performance using CelebA. As illustrated in Table 5, across all methods, reconstruction quality degrades as \mathcal{B} increases. This deterioration stems from the averaging effect of batching, which obscures image-specific gradient information and consequently reduces the reconstruction fidelity. Nevertheless, GGSS-R consistently outperforms all baselines at every batch size, benefiting from the strong generative capability of diffusion models coupled with pixel-level conditional guidance.

Impact of Guidance Rate. Table 6 reports the effect of the guidance rate m_r on reconstruction quality. The guidance rate, explored over the range 0.01–0.80, governs the balance between conditional guidance and generated image quality. The optimal GIA performance is achieved at $m_r = 0.20$, yielding the lowest MSE ($7.7215e-5$), the highest PSNR (41.1230) and the lowest LPIPS ($2.0519e-7$).

Impact of Noise Type. To evaluate how noise type influences GIA performance, we varied the magnitude of Laplacian noise applied to gradients. Table 7 quantitatively evaluates the corresponding peak performance. Comparing these

Method	$\mathcal{B} = 1$	$\mathcal{B} = 2$	$\mathcal{B} = 4$	$\mathcal{B} = 8$
DLG (Zhu, Liu, and Han 2019)	13.1835	11.0356	9.9452	8.3562
IG (Geiping et al. 2020b)	17.0756	15.2562	11.3003	9.5382
GI (Yin et al. 2021b)	16.5109	14.8351	11.0352	9.2350
GIAS (Jeon et al. 2021a)	13.3885	11.4684	10.2465	8.9462
GGL (Li et al. 2022b)	17.4923	15.5526	-	-
GIFD (Fang et al. 2023a)	20.0534	18.3562	13.9636	10.3425
FinetunedDiff (Meng et al. 2024)	29.4490	21.4583	15.6387	11.4570
GGSS-R (Ours)	41.1229	24.4602	17.4791	12.1425

Table 5: Impact of batch size for gradient acquisition on reconstruction quality. Note that StyleGAN2’s latent vector contains a large number of parameters and the CMA-ES optimizer used in GGL does not support large-scale optimization. As a result, GGL is limited to batch sizes $\mathcal{B} \leq 2$.

Guidance Rate	MSE↓	PSNR↑	LPIPS↓
0.01	2.0406e-3	26.9023	6.9703e-6
0.10	2.1778e-4	36.6198	4.1730e-7
0.20	7.7215e-5	41.1229	2.0519e-7
0.30	1.4769e-4	38.3065	2.8713e-7
0.40	1.0778e-4	39.6746	1.9574e-7
0.60	2.0407e-4	36.9023	9.0661e-7
0.80	2.7478e-4	35.6101	1.1251e-6

Table 6: Impact of guidance rate on reconstruction quality.

Laplacian Variance	MSE↓	PSNR↑	LPIPS↓
10^{-4}	0.0016	27.9395	9.3180e-6
10^{-3}	0.0038	24.2073	3.7156e-5
10^{-2}	0.0304	15.1657	2.3746e-4
10^{-1}	0.1390	8.5697	5.6100e-4

Table 7: Quantitative analysis on reconstruction quality with Laplacian noise-perturbed CNN gradients.

results with those obtained under Gaussian noise perturbations (Table 2), we observe that the reconstruction quality is determined primarily by the noise scale, while the specific noise distribution has only a negligible effect.

More details of running time and GPU usage are provided in Appendix C of the extended version.

7 Conclusion

This paper proposes a novel GIA method based on conditional diffusion models to recover high-quality images from noise-perturbed gradients. Our approach requires minimal modifications to the reverse diffusion sampling and operates without prior knowledge. By exploiting the inherent generative capability of diffusion models, our attack effectively mitigates the impact of noise perturbation. Moreover, we derive theoretical bounds on the reconstruction error and analyze the convergence properties of the attack loss. Our findings further reveal an intrinsic vulnerability of diverse model architectures to GIAs, and we propose the RV metric to quantify this vulnerability. Extensive experiments confirm the superior performance of our method in recovering high-quality reconstructions from noise-perturbed gradients.

Acknowledgments

This work is supported by National Natural Science Foundation of China (U24B20144, U23A20299, 62172424, 62276270, 62322214, 62436010, 62441230), The Natural Science Foundation of Fujian Province (2024J08276, 2024J08278), The Key Research Project for Young and Middle-aged Researchers by the Fujian Provincial Department of Education (JZ230044).

References

- Chung, H.; Kim, J.; Mccann, M. T.; Klasky, M. L.; and Ye, J. C. 2022a. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*.
- Chung, H.; Sim, B.; Ryu, D.; and Ye, J. C. 2022b. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35: 25683–25696.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- El Mestari, S. Z.; Lenzini, G.; and Demirci, H. 2024. Preserving data privacy in machine learning systems. *Computers & Security*, 137: 103605.
- Fang, H.; Chen, B.; Wang, X.; Wang, Z.; and Xia, S.-T. 2023a. GIFD: A Generative Gradient Inversion Method with Feature Domain Optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4967–4976.
- Fang, H.; Chen, B.; Wang, X.; Wang, Z.; and Xia, S.-T. 2023b. GIFD: A Generative Gradient Inversion Method with Feature Domain Optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4967–4976.
- Geiping, J.; Bauermeister, H.; Dröge, H.; and Moeller, M. 2020a. Inverting gradients-how easy is it to break privacy in federated learning? *Advances in Neural Information Processing Systems*, 33: 16937–16947.
- Geiping, J.; Bauermeister, H.; Dröge, H.; and Moeller, M. 2020b. Inverting gradients-how easy is it to break privacy in federated learning? *Advances in Neural Information Processing Systems*, 33: 16937–16947.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- He, Y.; Murata, N.; Lai, C.-H.; Takida, Y.; Uesaka, T.; Kim, D.; Liao, W.-H.; Mitsufuji, Y.; Kolter, J. Z.; Salakhutdinov, R.; et al. 2023. Manifold preserving guided diffusion. *arXiv preprint arXiv:2311.16424*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020a. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; Jain, A.; and Abbeel, P. 2020b. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Jeon, J.; Lee, K.; Oh, S.; Ok, J.; et al. 2021a. Gradient inversion with generative image prior. *Advances in neural information processing systems*, 34: 29898–29908.
- Jeon, J.; Lee, K.; Oh, S.; Ok, J.; et al. 2021b. Gradient inversion with generative image prior. *Advances in neural information processing systems*, 34: 29898–29908.
- Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A. N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. 2021. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2): 1–210.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017a. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017b. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Laurent, B.; and Massart, P. 2000. Adaptive estimation of a quadratic functional by model selection. *Annals of statistics*, 1302–1338.
- Ledoux, M. 2001. *The concentration of measure phenomenon*. 89. American Mathematical Soc.
- Li, T.; Sahu, A. K.; Talwalkar, A.; and Smith, V. 2020. Federated learning: Challenges, methods, and future directions. *IEEE signal processing magazine*, 37(3): 50–60.
- Li, Z.; Zhang, J.; Liu, L.; and Liu, J. 2022a. Auditing privacy defenses in federated learning via generative gradient leakage. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10132–10142.
- Li, Z.; Zhang, J.; Liu, L.; and Liu, J. 2022b. Auditing privacy defenses in federated learning via generative gradient leakage. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10132–10142.
- Meng, J.; Huang, T.; Chen, H.; and Li, C. 2024. Is Diffusion Model Safe? Severe Data Leakage via Gradient-Guided Diffusion Model. *arXiv preprint arXiv:2406.09484*.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, 8162–8171. 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.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.

Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.

Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.

Wang, J.; Guo, S.; Xie, X.; and Qi, H. 2022. Protect privacy from gradient leakage attack in federated learning. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*, 580–589. IEEE.

Yang, L.; Ding, S.; Cai, Y.; Yu, J.; Wang, J.; and Shi, Y. 2024. Guidance with spherical gaussian constraint for conditional diffusion. *arXiv preprint arXiv:2402.03201*.

Yin, H.; Mallya, A.; Vahdat, A.; Alvarez, J. M.; Kautz, J.; and Molchanov, P. 2021a. See through gradients: Image batch recovery via gradinversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16337–16346.

Yin, H.; Mallya, A.; Vahdat, A.; Alvarez, J. M.; Kautz, J.; and Molchanov, P. 2021b. See through gradients: Image batch recovery via gradinversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16337–16346.

Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.

Yu, J.; Wang, Y.; Zhao, C.; Ghanem, B.; and Zhang, J. 2023. Freedom: Training-free energy-guided conditional diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23174–23184.

Zhao, B.; Mopuri, K. R.; and Bilen, H. 2020. idlg: Improved deep leakage from gradients. *arXiv preprint arXiv:2001.02610*.

Zhu, L.; Liu, Z.; and Han, S. 2019. Deep leakage from gradients. *Advances in neural information processing systems*, 32.