

Understanding Flatness in Generative Models: Its Role and Benefits

Taehwan Lee^{1,*}Kyeongkook Seo^{1,*}Jaejun Yoo^{1,2,†}Sung Whan Yoon^{1,2,†}¹Graduate School of Artificial Intelligence, Ulsan National Institute of Science and Technology²Department of Electrical Engineering, Ulsan National Institute of Science and Technology

{taehwan, kyeongkookseo, jaejun.yoo, shyoon8}@unist.ac.kr

Abstract

Flat minima, known to enhance generalization and robustness in supervised learning, remain largely unexplored in generative models. In this work, we systematically investigate the role of loss surface flatness in generative models, both theoretically and empirically, with a particular focus on diffusion models. We establish a theoretical claim that flatter minima improve robustness against perturbations in target prior distributions, leading to benefits such as reduced exposure bias—where errors in noise estimation accumulate over iterations—and significantly improved resilience to model quantization, preserving generative performance even under strong quantization constraints. We further observe that Sharpness-Aware Minimization (SAM), which explicitly controls the degree of flatness, effectively enhances flatness in diffusion models even surpassing the indirectly promoting flatness methods—Input Perturbation (IP) which enforces the Lipschitz condition, ensembling-based approach like Stochastic Weight Averaging (SWA) and Exponential Moving Average (EMA)—are less effective. Through extensive experiments on CIFAR-10, LSUN Tower, and FFHQ, we demonstrate that flat minima in diffusion models indeed improve not only generative performance but also robustness.

1. Introduction

What does it mean for a generative model to have a flat loss landscape? While flat minima have been extensively studied in supervised learning, where they are known to enhance generalization and robustness to distribution shifts, e.g., promoting stable label predictions under input shifts such as domain changes [2, 5, 9, 14, 18, 24, 27], their role in generative models remains largely unexplored. At first glance, one might expect that flatness would be beneficial in

*Equal contribution

†Co-corresponding author




Figure 1. Loss surface of ADM (left) and ADM+SAM (right)

	FID ↓ (32 → 8)	$\ \epsilon_\theta\ ^2$ gap ↓	LPF ↓
ADM	34.47 $\xrightarrow{+13.65}$ 48.02	+11.39	0.097
+SAM	9.01 $\xrightarrow{-0.07}$ 8.94	+3.32	0.063

Table 1. Comparison of the 1) Fréchet Inception Distance (FID) scores under 32- and 8-bit precision with (+/-) changes, 2) the gap of $\|\epsilon\|_2$ with training, 3) flatness metric called LPF for ADM and +SAM in the CIFAR-10 experiments. ↓: a lower value is preferred.




Figure 2. Generated samples of 8-bit quantized models for CIFAR-10. Model parameters are directly converted from 32-bit to 8-bit after training, without further optimization or adaptation.

generative models as well. However, generative models—particularly those based on diffusion processes—operate under fundamentally different principles: rather than mapping structured inputs to labels, they take noise as input and iteratively refine it into coherent outputs through a learned denoising process. This key difference motivates a critical question: *Do flat minima compel the generative model to produce similar contents regardless of input variations?* If so, such behavior would contradict the goal of generative diversity and potentially impair generalization. If not, it remains unclear how flatness affects generative modeling, and if it is indeed beneficial, how can we induce it?

As a preview of our analysis, Figure 1, 2 and Table 1 hint

at the role and benefits of the generative models with flat minima. We find that SAM [9], an optimizer designed to promote flatter loss surfaces, significantly enhances flatness when applied to a baseline, i.e., Ablated Diffusion Model (ADM) [7], as shown in Fig. 1 and Tab. 1 (measured using Low-Pass-Filter (LPF) flatness metric [2]). Regarding the benefits from flatness, we observe improved FID, reduced exposure biases (as indicated by the $\|\epsilon\|^2$ gap [31]), and lower degradation from model quantization (as reflected in 8-bit precision FID and qualitative results in Fig. 2).

Returning to the key question, we set out to investigate the implication of flat loss landscapes in generative models, particularly diffusion models [13]. To make progress, we first seek to refine our core research question: “*Why might flatter minima be desirable in generative models, and what theoretical guarantees can we establish?*”

To address these questions, we conduct a theoretical and empirical analysis of flat minima in diffusion models. First, we establish a theoretical link between the model’s parameter space and data probability density (**Theorem 1**), showing that flatter minima improve robustness to perturbations in the target prior distribution. As a consequence, we prove that the divergence between the true distribution and the learned distribution is upper-bounded, indicating that flat minima enable diffusion models to generalize not only to the training distribution but also to variations in perturbed data densities (**Theorem 2**). This insight suggests that flat diffusion models exhibit superior robustness when out-of-distribution is given, demonstrating their ability to cope with the error accumulation problem during the iterative sampling process, known as *exposure bias problem* [21, 31, 32]. Moreover, the flat loss landscape in the parameter space directly implies the robustness against model parameter changes, which indicates that the flatness offers the robustness of generative performance when the model has been quantized.

To complement our theoretical findings, we empirically examine the effects of explicit flatness regularization. Initially, we expect that applying SAM to diffusion models would yield improvements similar to those observed in discriminative settings. However, an intriguing discovery emerges: diffusion models are inherently much flatter than anticipated. Unlike discriminative models, where typical SAM trade-off values significantly impact on flatness, applying standard regularization strengths to diffusion models produced almost no discernible change. Further analysis reveals that diffusion models naturally exhibit a flat loss landscape, likely due to the diverse range of noise levels they are trained to denoise. Building on this observation, we find that significantly stronger regularization is required in diffusion models for SAM to induce meaningful flatness, setting them apart from their discriminative counterparts.

To validate these findings, we conduct extensive exper-

iments on CIFAR-10, LSUN Tower, and FFHQ. Our results show that flat minima in diffusion models consistently improve generative performance with enhanced generalization. This improved generalization offers several practical advantages. For instance, we observe that flat diffusion models better mitigate exposure bias, where errors in noise estimation accumulate over iterations, leading to more stable sampling dynamics. We also show that flat diffusion models result in a low quantization error where the model precision is reduced from 32-bit to 8-bit.

We summarize our contributions as follows. Our work provides the first systematic investigation into the role of flat minima in generative models, offering both theoretical insights and practical implications:

- We explore diffusion models through the lens of flat loss landscapes, revealing their inherent flatness.
- We establish a theoretical link between flat minima and generalization, proving that flatness improves robustness via an upper-bound on the divergence between the target and learned distributions.
- We empirically analyze IP [32], SWA [14], and EMA, showing that while they improve FID, they do not significantly enhance flatness. In contrast, SAM reshapes the loss surface and effectively promotes flatness in diffusion models.
- We demonstrate how these insights address practical challenges in generative modeling through extensive evaluations on CIFAR-10, LSUN Tower, and FFHQ, analyzing FID performance, exposure bias, model quantization error, and loss landscape properties.

2. Preliminaries and Backgrounds

2.1. Optimization for flatness

EMA. EMA is a weight-averaging method that blends a fraction of newly updated parameters with a fraction of previously accumulated parameters at each step. Formally, for parameters θ , the EMA update is: $\theta_{\text{EMA}} \leftarrow (1 - \lambda)\theta_{\text{EMA}} + \lambda\theta$, $\lambda \in (0, 1)$. Heuristically ensembling different models from multiple iterations reaches flat local minima and improves generalization capacities [11, 17, 23].

SWA. SWA [14] stabilizes training by maintaining an ongoing average of model weights across training epochs. Rather than relying on the final parameters from a single run, SWA updates this average—typically during the latter part of training—leading to an overall smoothing effect in the loss landscape. Although SWA and EMA share the idea of parameter averaging, SWA is directly motivated by the search for flatter minima, which improves generalization.

SAM. SAM [9] is an optimization technique that promotes generalization by steering model parameters toward flatter regions of the loss landscape. It incorporates a sharpness term into the objective, defined as the maximal increase in

Algorithm 1 Pseudo algorithm for SWA, EMA, SAM

Input: w_0 : Initial weights, ρ : size of perturbation,
 c : Cycle length, λ : update momentum

Output: w, w_{SWA}, w_{EMA}

```
1:  $w, w_{SWA}, w_{EMA} \leftarrow w_0$  {Initialize weights with  $w_0$ }
2: for  $i = 0, \dots, n - 1$  do
3:   if SWA then
4:      $\alpha \leftarrow \alpha(i)$  {Calculate LR for the iteration}
5:   end if
6:   if SAM then
7:      $\hat{w} \leftarrow w + \rho \frac{\nabla \mathcal{L}_i(w)}{\|\nabla \mathcal{L}_i(w)\|}$  {Ascent step}
8:      $w \leftarrow w - \alpha \nabla \mathcal{L}_i(\hat{w})$  {Descent step}
9:   else
10:     $w \leftarrow w - \alpha \nabla \mathcal{L}_i(w)$ 
11:  end if
12:  if SWA &  $\text{mod}(i, c) = 0$  then
13:     $n_{models} \leftarrow i/c$  {Number of models}
14:     $w_{SWA} \leftarrow \frac{w_{SWA} \cdot n_{models} + w}{n_{models} + 1}$  {Update average}
15:  else if EMA then
16:     $w_{EMA} \leftarrow (1 - \lambda) \cdot w_{EMA} + \lambda \cdot w$  {Update average}
17:  end if
18: end for
```

loss within an ℓ_2 -ball of radius ρ around the current parameters θ : $\max_{\|\theta - p\|_2 \leq \rho} [L(\theta + p) - L(\theta)]$. By minimizing this sharpness term alongside the original loss, SAM encourages solutions that are robust to small perturbations in the parameters, thus improving generalization. The pseudocode for SWA, EMA, and SAM is provided in Algorithm 1.

IP. IP [32] was introduced to address the exposure bias problem, which inherently arises in the auto-regressive denoising process of diffusion models. They define exposure bias as a generalization failure to unseen inputs, caused by the accumulation of prediction errors over timesteps. To mitigate this issue, authors propose a smoothed diffusion model that satisfies a Lipschitz condition by training on randomly perturbed input data. This strategy aligns with the flat minima perspective, implicitly promoting generalization to nearby unseen inputs.

2.2. Score-based Generative Models

Score-based generative models (SGMs) [39–42] leverage the score function of a data distribution to iteratively generate samples by solving Stochastic Differential Equations (SDEs). Let $p_{\text{data}}(\mathbf{x})$ be the unknown data distribution. The *score function* is defined as the gradient of the log-density:

$$\nabla_{\mathbf{x}} \log p_{\text{data}}(\mathbf{x}). \quad (1)$$

Instead of modeling $p_{\text{data}}(\mathbf{x})$ explicitly, SGMs learn a neural network $s_{\theta}(\mathbf{x}, t)$ to approximate the score function under slight noise perturbations, t is timestep.

SGMs can be formulated as a stochastic process where noise is gradually added to the data in the forward direction and removed in the backward direction.

Forward SDE. The forward diffusion process is defined by the following SDE:

$$d\mathbf{x} = f(\mathbf{x}, t)dt + g(t)d\mathbf{w}, \quad (2)$$

where $f(\mathbf{x}, t)$ is the drift coefficient that controls deterministic changes, $g(t)$ is the diffusion coefficient that controls noise intensity, and $d\mathbf{w}$ denotes a standard Wiener process.

Reverse SDE. Using the learned score function, we can reverse the diffusion process to generate new samples. The backward SDE is given by:

$$d\mathbf{x} = [f(\mathbf{x}, t) - g^2(t)\nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt + g(t)d\bar{\mathbf{w}}, \quad (3)$$

where $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ is the time-dependent score function, approximated by $s_{\theta}(\mathbf{x}, t)$, and $d\bar{\mathbf{w}}$ is a standard Wiener process running in reverse time. To successfully estimate the score function, the score matching objective is defined as:

$$\mathcal{L}_{\text{SGM}} = \mathbb{E}_t \left[\lambda(t) \cdot \mathbb{E}_{p_t(\mathbf{x})} \left[\|s_{\theta}(\mathbf{x}, t) - \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x})\|_2^2 \right] \right]. \quad (4)$$

Diffusion models. Diffusion models (DPMs) [13, 38] can be viewed as a discrete-time approximation of SGMs, with the following forward and backward diffusion process:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}), \quad (5)$$

$$q(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\mu}(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t \mathbf{I}), \quad (6)$$

$$\tilde{\mu}(\mathbf{x}_t, \mathbf{x}_0) = \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_t}{1 - \bar{\alpha}_t} \mathbf{x}_0 + \frac{\sqrt{\bar{\alpha}_t} (1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \mathbf{x}_t, \quad (7)$$

$$\tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t \quad (8)$$

where β_t is a predefined parameter by variance scheduling, $\alpha_t := 1 - \beta_t$, and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$.

Exposure bias. Previous works [21, 31, 32] have raised an input discrepancy problem between the training and sampling phases, known as *exposure bias problem*. Specifically, during training, the DPM is provided with ground truth noisy images, whereas during sampling, it receives a denoised image that came from the previous timestep. The presence of errors in these inputs, amplified through the iterative denoising process, makes it challenging for the diffusion model to accurately predict noise. As these errors accumulate over iterative timesteps, they cause the generated samples to significantly deviate from the desired trajectory.

3. Theoretical Analysis on Flatness

We here theoretically analyze the effect of flat minima in SGMs, with an emphasis on diffusion models.

3.1. Notations

Our analysis is built upon the settings of SGM, which is a fundamental form of diffusion models. Beyond the brief preliminaries in Section 2.2, we introduce additional notations and formulations to be used in our analysis.

Following prior theoretical analysis [22], we formulate the simplified score model $s_\theta(\cdot, \cdot)$ as a random feature model consisting of an encoder and a decoder.

$$s_\theta(\mathbf{x}, t) := \frac{1}{m} \theta \sigma(\mathbf{W}^\top \mathbf{x} + \mathbf{U}^\top \mathbf{e}_t), \quad (9)$$

where $\mathbf{x} \in \mathbb{R}^{d \times 1}$ is the given input, \mathbf{e}_t is the timestep embedding for t , a group of parameters includes $\theta \in \mathbb{R}^{d \times m}$, $\mathbf{W} \in \mathbb{R}^{d \times m}$ and $\mathbf{U} \in \mathbb{R}^{d_e \times m}$, and a few positive integers include d , m , and d_e . By following the training procedures of diffusion models, θ is the learnable parameter, while others, i.e., \mathbf{W} and \mathbf{U} , which respectively embed \mathbf{x} and \mathbf{e}_t , are set to be frozen. For a given timestep t , the score matching loss objective defined in Eq. (4)

$$\mathcal{L}(\mathbf{x}; \theta, t, p_t) := \|s_\theta(\mathbf{x}, t) - \nabla_{\mathbf{x}} \log p_t(\mathbf{x})\|_2^2, \quad (10)$$

where $p_t(\mathbf{x})$ is the probability density function for the target distribution at time t .

3.2. Mathematical Claims

For the main claims, we follow the formulation of Equation (10), while omitting timestep t without loss of generality. Our mathematical claims are valid for all timesteps. We start from defining flat minima in SGMs (Definition 1), and the robustness against the gap of ground truth prior $p(\mathbf{x})$ and the estimated prior $\hat{p}(\mathbf{x})$ (Definition 2). For the flat minima definition, we rehearse the definition of flat minima for the supervised learning case [37] and extend it to the Δ -flat minima of SGMs.

Definition 1. (Δ -flat minima) *Let us consider a SGM with loss function $\mathcal{L}(\mathbf{x}; \theta, p)$. A minimum θ^* is Δ -flat minima when the following constraints are hold:*

$$\begin{aligned} \forall \delta \in \mathbb{R}^{d \times m} \text{ s.t. } \|\delta\|_2 \leq \Delta, \quad \mathcal{L}(\mathbf{x}; \theta^* + \delta, p) &= l^* \\ \exists \delta \in \mathbb{R}^{d \times m} \text{ s.t. } \|\delta\|_2 > \Delta, \quad \mathcal{L}(\mathbf{x}; \theta^* + \delta, p) &> l^*, \end{aligned}$$

where $l^* := \mathcal{L}(\mathbf{x}; \theta^*, p)$ and $\Delta \in \mathbb{R}^+$.¹

To extend to robustness against the erroneous estimation of prior probability, we further define the \mathcal{E} -distribution gap robustness of SGMs, where \mathcal{E} indicates the divergence between the ground truth p and the estimated \hat{p} :

Definition 2. (\mathcal{E} -distribution gap robustness) *A minimum θ^* is \mathcal{E} -distribution gap robust when the following constraints are hold:*

$$\begin{aligned} \forall \hat{p}(\mathbf{x}) \text{ s.t. } D(p||\hat{p}) \leq \mathcal{E}, \quad \mathcal{L}(\mathbf{x}; \theta^*, \hat{p}) &= l^* \\ \exists \hat{p}(\mathbf{x}) \text{ s.t. } D(p||\hat{p}) > \mathcal{E}, \quad \mathcal{L}(\mathbf{x}; \theta^*, \hat{p}) &> l^*, \end{aligned}$$

where $D(\cdot||\cdot)$ is the divergence between two probability density functions, \hat{p} is the perturbed prior distribution of \mathbf{x} , and \mathcal{E} is a positive real number.

Let us then bridge the flatness and the robustness. First, based on the definitions of flat minima (Definition 1) and distributional gap robustness (Definition 2), we try to formulate how the perturbation of the parameter, i.e., $\theta + \delta$, links to the perturbed prior distribution \hat{p} , with the given equality as follows:

$$\mathcal{L}(\mathbf{x}; \theta + \delta, p) = \mathcal{L}(\mathbf{x}; \theta, \hat{p}). \quad (11)$$

This means how the parameter perturbation is translated into the shift of the prior, while keeping the loss unchanged. Let us provide a mathematical claim for it:

Theorem 1. (*A perturbed distribution*) *For a given prior distribution of $p(\mathbf{x})$ and the δ -perturbed minimum, i.e., $\theta + \delta$, the following $\hat{p}(\mathbf{x})$ satisfies the equality (11):*

$$\hat{p}(\mathbf{x}) = e^{-I(\mathbf{x}, \delta)} p(\mathbf{x}), \quad (12)$$

where $I(\mathbf{x}, \delta) := \frac{1}{2} \mathbf{x}^\top (\delta \mathbf{W}^\top) \mathbf{x} + \mathbf{x}^\top \delta (\mathbf{U}^\top \mathbf{e}) + C$ with $\delta \mathbf{W}^\top$ being symmetric, and $C \in \mathbb{R}$ is set to satisfy $\int_{\mathbb{R}^d} e^{-I(\mathbf{x}, \delta)} p(\mathbf{x}) d\mathbf{x} = 1$.

For analytical convenience, Theorem 1 derives a perturbed density under symmetric $\delta \mathbf{W}^\top$, showing that parameter perturbation translates into a ‘‘scaling’’ of probability density. Note that C absorbs the normalizing constant. However, the scaling depends on various parameters. Thus, it is non-trivial to imagine the behavior of \hat{p} . Let us narrow our view to the diffusion model, where it predicts a noise signal sampled from a Gaussian distribution, i.e., $\mathcal{N}(0, \mathbf{I})$. It changes our focus on handling the noise distribution, i.e., ϵ , rather than the distribution of data side, i.e., p . The following corollary describes how the perturbed $\hat{\epsilon}$ looks like:

Corollary 1. (*Diffusion version of Theorem 1*) *For a given prior Gaussian distribution of noise $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ and the δ -perturbed minimum, i.e., $\theta + \delta$, the following $\hat{\epsilon}$ satisfies the equality (11):*

$$\hat{\epsilon} = e^{-I(\mathbf{x}, \delta)} \epsilon = \mathcal{N}(\boldsymbol{\mu}_\delta, \Sigma_\delta), \quad (13)$$

where $\Sigma_\delta := \left(\mathbf{I} + \frac{\delta \mathbf{W}^\top}{m} \right)^{-1}$, $\boldsymbol{\mu}_\delta := \frac{1}{m} \Sigma_\delta \delta \mathbf{U}^\top \mathbf{e}$.

¹ \forall means ‘‘for all,’’ \exists means ‘‘there exists,’’ and \mathbb{R}^+ indicates the set of positive real numbers.

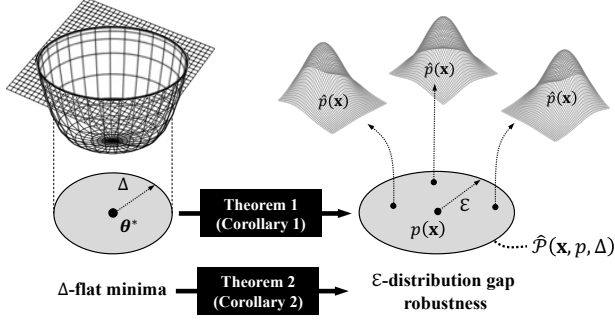


Figure 3. A conceptual illustration of theoretical analysis. **Theorem 1 (Corollary 1)** for diffusion model) translates the perturbation in the parameter space into the set of perturbed distributions. **Theorem 2 (Corollary 2)** for diffusion model) shows that flat minima lead to robustness against the distribution gap.

Remark 1. ($\hat{\epsilon}$ becomes perturbed Gaussian) For the diffusion models, we emphasize that the perturbation in the parameter space, δ , leads to the perturbation of distribution, which follows the Gaussian distribution.

Next, let us consider all possible perturbations within the ball of norm, i.e., $\|\delta\|_2 \leq \Delta$ and $\delta \mathbf{W}^\top = \mathbf{W} \delta^\top$. The resulting set of perturbed distributions induced by the parameter perturbation is defined as follows:

Definition 3. (A set of perturbed distributions) For a given distribution of $p(\mathbf{x})$, a set of distributions $\hat{\mathcal{P}}(\mathbf{x}; p, \Delta)$ is defined as the set of perturbed distributions $\hat{p}(\mathbf{x})$:

$$\hat{\mathcal{P}}(\mathbf{x}; p, \Delta) := \left\{ e^{-I(\mathbf{x}, \delta)} p(\mathbf{x}) \mid \|\delta\|_2 \leq \Delta, \delta \mathbf{W}^\top = \mathbf{W} \delta^\top \right\}. \quad (14)$$

From the definition, Δ -flat minimum flattens all loss values of the perturbed parameters. Consequently, the flat minimum achieves flat loss values for all possible distributions within $\hat{\mathcal{P}}$, as formally shown in the following proposition:

Proposition 1. (A link from Δ -flatness to $\hat{\mathcal{P}}$) A Δ -flat minimum θ^* achieves the flat loss values for all distributions sampled from the set of perturbed distributions:

$$\forall p \sim \hat{\mathcal{P}}(\mathbf{x}; p, \Delta), \mathcal{L}(\mathbf{x}; \theta^*, p) = l^* \quad (15)$$

$$\exists p \sim \hat{\mathcal{P}}(\mathbf{x}; p, \Delta), \mathcal{L}(\mathbf{x}; \theta^*, p) > l^*. \quad (16)$$

We finally link flatness to the distributional gap, i.e., connecting **Definition 1** and **2**. Let us anticipate the distribution in $\hat{\mathcal{P}}$ with the maximal divergence from p . The distribution gap \mathcal{E} is then upper bounded by the divergence of the most outreached distribution:

Theorem 2. (A link from Δ -flatness to \mathcal{E} -gap robustness) A Δ -flat minimum achieves \mathcal{E} -distribution gap robustness, such that \mathcal{E} is upper-bounded as follows:

$$\mathcal{E} \leq \max_{\hat{p} \sim \hat{\mathcal{P}}(\mathbf{x}; p, \Delta)} D(p \parallel \hat{p}). \quad (17)$$

Let us further manipulate the upper bound for the diffusion model case. From **Corollary 1**, \hat{p} is a form of Gaussian distribution; thus, it is possible to achieve the closed form of the upper bound as follows:

Corollary 2. (Diffusion version of Theorem 2) For a diffusion model, a Δ -flat minimum achieves \mathcal{E} -distribution gap robustness, such that \mathcal{E} is upper-bounded as follows:

$$\mathcal{E} \leq \max_{\|\delta\|_2 \leq \Delta} \frac{1}{2} \left[\log |\Sigma_\delta| - d + \text{tr}(\Sigma_\delta^{-1}) + \boldsymbol{\mu}_\delta^\top \Sigma_\delta^{-1} \boldsymbol{\mu}_\delta \right] \leq \frac{1}{2} \left[\sum_{i=1}^d (\sigma_i - \log \sigma_i) - d + \frac{\sigma_d}{m^2} \|\mathbf{U}^\top \mathbf{e}\|_2^2 \Delta^2 \right], \quad (18)$$

where σ_i is an eigenvalue of Σ_δ^{-1} with the increasing order of $\sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_d$.

Remark 2. (Flatter minima enhance distribution gap robustness) It is noteworthy that flatter minima with a large Δ flatten the loss values for the far-pushed away perturbed distribution \hat{p} with large divergence, leading to the larger bound of \mathcal{E} . It indicates that flattening the loss surface makes the generative model zero-force the loss values of diversified prior distribution, thus enhancing the robustness against the prior distribution gap.

Flatness reduces exposure bias (a gap of $\hat{\epsilon}$ and ϵ): A clear advantage of flatness is robustness to exposure bias, where the errors in noise estimation accumulate over iterations, thus severely deteriorating the generative performance. In formula, $\hat{\epsilon}$ deviates from ϵ . We argue that a generative model with a flat minimum suppresses the loss values of perturbed estimation, so that the error accumulation is sufficiently relieved. To be shown in experiments, we empirically confirm that flat minima tend to show a smaller exposure bias, leading to robust generative performance.

Flatness becomes robust to model quantization (a compression from θ to $\hat{\theta}$): For another benefit, flat minima make the model robust to the performance degradation caused by the model quantization. Recently, the quantization of generative models has become crucial, when enabling real-time applications of generative tasks [25, 26, 45]. We claim that flat minima inherently bolster the robustness against the degradation due to the quantization, because the quantized model, $\hat{\theta}$, can be viewed as a perturbed version of the full-precision parameter, i.e., $\hat{\theta} = \theta + \Delta$; thus the loss remains flat for the perturbation [26]. In experiments, we clearly demonstrate that the diffusion model with flatness is remarkably resilient to degradation from quantization compared to other methods.

A conceptual overview of our theoretical analysis is illustrated in Fig. 3. The proof of **Theorem 1**, and **Corollary 1, 2** are given in Appendix 1 in Supplementary.

FID Score	Dataset	CIFAR-10 (32x32)		LSUN Tower (64x64)		FFHQ (64x64)	
	T'	20 steps	100 steps	20 steps	100 steps	20 steps	100 steps
Algorithms	ADM	34.47	8.80	36.65	8.57	30.81	7.53
	+EMA	10.63	4.06	7.87	2.49	19.03	6.19
	+SWA	11.00	3.78	8.72	2.31	17.93	5.49
	+IP	20.11	7.23	25.77	7.00	15.03	13.55
	+IP+EMA	9.10	3.46	7.66	2.43	11.72	4.00
	+IP+SWA	9.04	3.07	8.55	2.34	12.99	3.54
	+SAM	9.01	3.83	16.02	4.79	11.59	5.29
	+SAM+EMA	7.00	3.18	6.66	2.30	11.41	5.04
	+SAM+SWA	7.27	2.96	6.50	2.27	12.15	4.17

Table 2. FID Scores for ADM baselines with different algorithms on CIFAR-10, LSUN Tower, and FFHQ datasets. We use DDPM sampling with shorter resampling timesteps, $T' = 20, 100$.

4. Experiments

4.1. Experiments Settings

Datasets. Similar to the setup in [32], we use CIFAR-10 (32x32), LSUN-tower (64x64), and FFHQ (64x64).

Baselines. We use ADM baselines² [7] for training unconditional DDPM. We incorporate flatness-enhancing approaches such as SWA [9] and SAM [9] to deliberately enhance flatness and reveal the effects of that. Because EMA may influence flatness [17], we distinguish between baselines with and without EMA. Additionally, we compare DDPM-IP [32] that introduces input perturbation during training to encourage smoothness in the diffusion model.

Experimental details. We trained the model for 200K steps on the CIFAR-10, LSUN Tower, and FFHQ datasets. We employed the Adam optimizer for all experiments with a learning rate of $1e^{-4}$, following [7, 30]. For +IP, we used an input perturbation strength of 0.1 following [32], and we set the cycle length c to 100 and start averaging from 180K steps for SWA. Finally, we tuned the $\rho \in [1e^{-1}, 1e^{-2}, 1e^{-3}]$ parameters for +SAM.

Evaluation metrics. We evaluate generative performance by comparing Fréchet Inception Distance (FID) scores [12] using subsequence of timesteps ($T' < T$). We assess Low-Pass Filter (LPF) [2] values and loss plots under perturbation [5] to determine whether the baselines successfully identify a flat loss surface. To further investigate generalization ability of diffusion models, we investigate exposure bias and quantization error. For exposure bias [31], we compare the square norm of the predicted noise, $\|\epsilon_\theta\|^2$ when models are conditioned on ground truth noisy images dur-

²<https://github.com/openai/guided-diffusion>

LPF ↓	w/o	+EMA	+SWA
ADM	0.097	0.099	0.099
+IP	0.103	0.101	0.102
+SAM	0.063	0.063	0.063

↓: a lower value is preferred.

Table 3. Flatness measure on CIFAR-10. We calculate the loss with the perturbed model with Gaussian noise. Lower values indicate a flatter loss landscape.

ing training versus when they receive error-contained inputs at different sampling timesteps [30]. For quantization error, we compare the FID before and after quantization.

4.2. Generative Performance

The quantitative and qualitative results presented in Tab. 2 and Fig. 4 demonstrate the generative performance of baselines. From Tab. 2, we observe that applying EMA consistently improves FID scores across all datasets and respaced timesteps, demonstrating its effectiveness in stabilizing model updates and enhancing sample quality. SWA also provides notable improvements but does not outperform EMA in most cases.³ While EMA and SWA can always be applied as auxiliary options, both standalone SAM and **+SAM+SWA** achieve comparable or better FID scores, especially under 20 timesteps, where exposure bias worsens. This suggests that explicit flatness control improves generalization, in terms of enhancing sample quality and robustness. For IP, the FID score is comparable to that of SAM at 100 steps, since the perturbation in the data distribution translates into parameter perturbations (**Theorem 1**). However, IP performs poorly at 20 steps, as it lacks a principled

³Because of limited space, we include the discussion and experimental results of the post-hoc EMA (pEMA) [8] in Appendix 2.3.

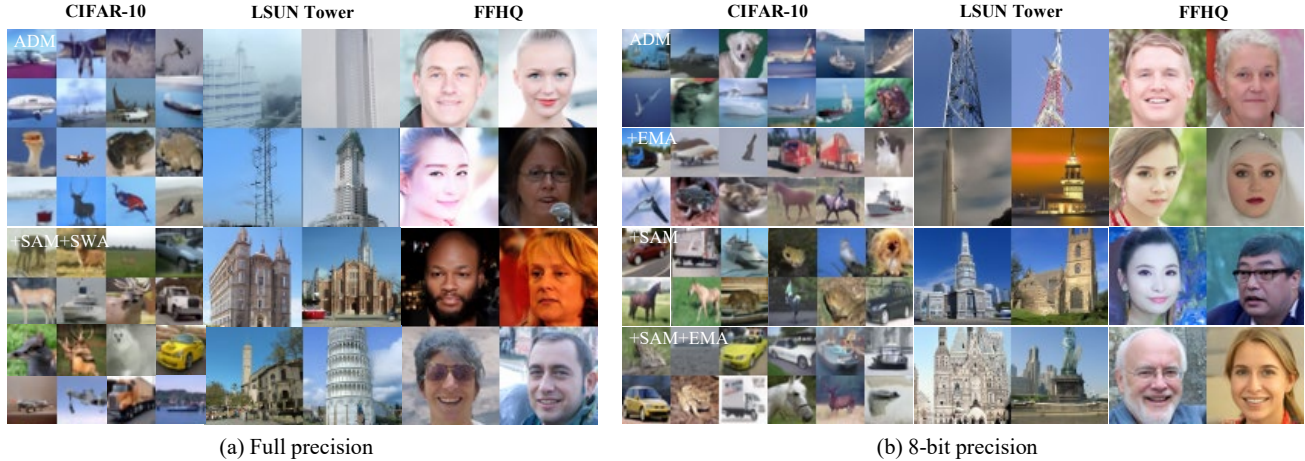


Figure 4. Generation results of (a) full precision (32-bit) and (b) 8-bit quantized models. We use resampled timesteps $T' = 20$ for sampling. For the 8-bit case, we selectively compare +EMA, +SAM, and their combinations, where the relatively flatter minima have been found.




Figure 5. Loss plots under perturbation for CIFAR-10. As the norm of model perturbation increases, we plot the corresponding loss values. A smaller slope indicates a flatter loss landscape.

noise direction and tends to converge to suboptimal minima. Figure 4 further supports these findings by providing visual comparisons of generated images ⁴.

4.3. Flatness Measurement

In Tab. 3, we compare LPF flatness measures [2], where lower values indicate a flatter loss landscape. We observe that ADM already possesses a certain level of flatness. Interestingly, we find that ensemble-like averaging, such as +SWA and +EMA, fail to induce additional flatness. In contrast, +SAM finds a flatter loss surface by explicitly perturbing model weights, as evidenced by SAM-applied models consistently achieving the lowest LPF values and significantly superior performance to other baselines. Also, we plot another measurement to support these findings further. Figure 5 shows how the loss value increases as the model perturbation is imposed for CIFAR-10. +IP, +EMA, and +SWA show no significant differences, coincide with the result of LPF, where these fail to find additional flatness. In this setting, +SAM shows a flatter behavior than other algorithms. It shows that diffusion models already show flatter loss landscapes, and ensemble-like averaging, such as

⁴We include additional sampling images in Appendix 2.2




Figure 6. L2 norm of the predicted noise for CIFAR-10. As the gap with “Training” decreases, the predicted noise norm during sampling approaches the ground truth.

+SWA and EMA, shows less impact. Moreover, finding flat minima explicitly leads to the flatness in diffusion models. We provide plots for other algorithms, including +IP+EMA, +IP+SWA, +SAM+EMA, +SAM+SWA, in the Appendix. 2.1.

4.4. Flat Diffusion Models are Robust

We show that flat diffusion models perform well under input shifts. As an instantiation, we set two cases: 1) exposure bias and 2) quantization error. Both scenarios share the common characteristic that diffusion models receive error-accumulated inputs, leading to performance degradation.

Exposure bias. In Fig. 6, we compare the L2 norm of the predicted noise in diffusion models. A model that closely aligns with the training trajectory is expected to exhibit lower exposure bias, thereby mitigating distribution shifts over iterative sampling steps. We observe that SAM-applied models consistently maintain L2 norms closer to the training curve, indicating a lower degree of exposure bias. This result aligns well with their superior FID performance reported in Tab. 2, suggesting that flat diffusion models generalize better across sampling timesteps and effectively combat error accumulation.

Model quantization error. Tab. 4 presents the FID scores

32-bit → 8-bit	20 steps	100 steps
ADM	34.47 $\xrightarrow{+13.65}$ 48.02	8.80 $\xrightarrow{+3.98}$ 12.78
+EMA	10.63 $\xrightarrow{+10.02}$ 20.65	4.06 $\xrightarrow{+3.3}$ 7.36
+SAM	9.01 $\xrightarrow{-0.07}$ 8.94	3.83 $\xrightarrow{+0.19}$ 4.02
+SAM+EMA	7.00 $\xrightarrow{+0.2}$ 7.20	3.18 $\xrightarrow{-0.06}$ 3.12

Table 4. FID performance on CIFAR-10 when model parameters are quantized to 8-bit. We apply direct quantization by clipping and rounding 32-bit parameters without additional fine-tuning. The 32-bit performance results are identical to those in Table 2.

of quantization across baselines, when reducing precision from 32-bit to 8-bit. Here, we quantize model parameters to 8-bit values using scaling and rounding, without any additional fine-tuning or retraining. Thanks to +SAM that aims to the robustness of model perturbation, flatter diffusion models results in surprisingly better robustness under quantization. These findings also support that +SAM is robust to model quantization.

5. Related works

Diffusion models. Diffusion probabilistic models (DPMs) [7, 13, 39, 40, 42] have recently emerged as powerful generative frameworks, achieving state-of-the-art performance in image [33, 36] and video [3, 16] synthesis. Despite extensive research on improving diffusion models through architecture modifications [15, 29, 33], enhanced training strategies [44], and novel noise scheduling techniques [42, 43], relatively little attention has been given to understanding their loss landscape properties.

Flat minima in various tasks. Flat minima have garnered significant attention in classification and domain generalization tasks with the view of loss landscape [10, 14, 20] and objective function [9, 18]. Other tasks also adopt the flatness for their performance. In federated learning, for example, where client devices collect and train on heterogeneous datasets, ensuring robust generalization across diverse data distributions is a core challenge. Recent works have employed flat minima strategies to mitigate this data heterogeneity issue by obtaining the model robust to distribution shift [4, 19, 34] and enhance the robustness of compressed neural networks [28, 48].

Flat minima on diffusion models.

Recent findings suggest that diffusion models inherently exhibit a flatter loss landscape [46]. However, they do not elucidate why the landscape is flat or how it relates to generalization and robustness, leaving the role of flatness in diffusion models unexplored. While prior works [1, 6, 32] leverage smoothness via Lipschitz continuity or gradient regularization, they do not directly examine flatness. To the best of our knowledge, we are the first to systematically investigate how flat minima influence diffusion models, their genera-

tive performance, and robustness—motivating our theoretical and empirical analysis.

6. Discussion and Limitation

The generative modeling community has long sought to understand generalization in terms of preventing models from overfitting to the training set while ensuring the generation of diverse and high-quality samples, particularly for underrepresented classes [35, 47]. While our work primarily investigates the role of flat minima in addressing exposure bias and quantization error, We conjecture that flat diffusion models may offer broader implications, such as cross-domain generation capability, few-shot personalization, outlier sample generation, mitigating memorization of generative models, and adversarial attack.

FID ↓	
ADM	15.97 (+5.34)
+SAM	11.72 (+4.72)
+SAM (7× noise)	15.72 (+8.72)

Table 5. FID degradations under adversarial attack.

In Tab. 5, we simulate an adversarial attack by maximizing the diffusion loss w.r.t. \tilde{z}_T . +SAM consistently achieves better FID than ADM and maintains its performance even under 7× amplified adversarial noise. Given the link between flatness and generalization, this insight suggests promising directions for future research into their role in generative modeling. As an initial attempt, we validate the benefits of flatness on relatively simple, low-resolution tasks. Notably, our theoretical claim is not confined to data dimensionality; extending the analysis to more complex settings would further support our theoretical claims.

7. Conclusion

We presented an in-depth investigation of the role and benefits of flat loss surfaces in generative models, which have remained unknown despite the great success of deep generative models. Our study explored this unknown by analyzing how flatness influences generative modeling and demonstrating its possible impact on robustness. First, our theoretical analysis reveals that the loss flatness of the generative models enhances the robustness against the disruptions of the target data distributions or the perturbations of model parameters. These insights naturally link to the reduced exposure bias problem and the model quantization error. Based on the evaluations on CIFAR-10, LSUN Tower, and FFHQ datasets, we demonstrated that 1) SAM is suitable to seek flatter minima of diffusion models rather than SWA and EMA, 2) flat minima reduce exposure bias with a minimal gap between testing and training, and 3) a flat model keeps its generative performance even with a strong model quantization; coinciding with our theory.

Acknowledgments

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.2022R1C1C100849612, No.RS-2024-00459023), Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT):(No.RS-2020-II201336, Artificial Intelligence Graduate School Program at UNIST), (RS-2025-25442824, AI Star Fellowship Program at UNIST), (No.2022-0-00959, No.RS-2022-II220959 (Part 2) Few-Shot Learning of Causal Inference in Vision and Language for Decision Making), (RS-2025-25441996, Development of a Virtual Tactile Signal Generation Platform Technology Based on Multimodal Vision-Tactile Integrated AGI), (RS-2025-25442149), Korea Health Technology R&D Project through the Korea Health Industry Development Institute (KHIDI), funded by the Ministry of Health & Welfare, Republic of Korea (No.RS-2025-02223382) and 2025 Research Fund (1.250006.01) of UNIST(Ulsan National Institute of Science & Technology).

References

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017. 8
- [2] Devansh Bisla, Jing Wang, and Anna Choromanska. Low-pass filtering sgd for recovering flat optima in the deep learning optimization landscape. In *International Conference on Artificial Intelligence and Statistics*, pages 8299–8339. PMLR, 2022. 1, 2, 6, 7
- [3] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023. 8
- [4] Debora Caldarola, Barbara Caputo, and Marco Ciccone. Improving generalization in federated learning by seeking flat minima. In *European Conference on Computer Vision (ECCV)*, pages 654–672. Springer, 2022. 8
- [5] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems*, 34: 22405–22418, 2021. 1, 6
- [6] Dingfan Chen, Tribhuvanesh Orekondy, and Mario Fritz. Gswgan: A gradient-sanitized approach for learning differentially private generators. *Advances in Neural Information Processing Systems*, 33:12673–12684, 2020. 8
- [7] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 2, 6, 8
- [8] Karras et al. Analyzing and improving the training dynamics of diffusion models. In *CVPR*, 2024. 6
- [9] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021. 1, 2, 6, 8
- [10] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31, 2018. 8
- [11] Diego Granzio, Xingchen Wan, Samuel Albanie, and Stephen Roberts. Iterative averaging in the quest for best test error. *arXiv preprint arXiv:2003.01247*, 2020. 2
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 6
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 3, 8
- [14] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018. 1, 2, 8
- [15] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022. 8
- [16] Kihong Kim, Haneol Lee, Jihye Park, Seyeon Kim, Kwanghee Lee, Seungryong Kim, and Jaejun Yoo. Hybrid video diffusion models with 2d triplane and 3d wavelet representation. In *European Conference on Computer Vision*, pages 148–165. Springer, 2024. 8
- [17] Frank Klinker. Exponential moving average versus moving exponential average. *Mathematische Semesterberichte*, 58: 97–107, 2011. 2, 6
- [18] Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. Asam: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *International Conference on Machine Learning*, pages 5905–5914. PMLR, 2021. 1, 8
- [19] Taehwan Lee and Sung Whan Yoon. Rethinking the flat minima searching in federated learning. In *Forty-first International Conference on Machine Learning*, 2024. 8
- [20] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems*, 31, 2018. 8
- [21] Mingxiao Li, Tingyu Qu, Ruicong Yao, Wei Sun, and Marie-Francine Moens. Alleviating exposure bias in diffusion models through sampling with shifted time steps. *arXiv preprint arXiv:2305.15583*, 2023. 2, 3
- [22] Puheng Li, Zhong Li, Huishuai Zhang, and Jiang Bian. On the generalization properties of diffusion models. *Advances in Neural Information Processing Systems*, 36:2097–2127, 2023. 4
- [23] Siyuan Li, Zicheng Liu, Juanxi Tian, Ge Wang, Zedong Wang, Weiyang Jin, Di Wu, Cheng Tan, Tao Lin, Yang Liu,

- et al. Switch ema: A free lunch for better flatness and sharpness. [arXiv preprint arXiv:2402.09240](#), 2024. 2
- [24] Tao Li, Pan Zhou, Zhengbao He, Xinwen Cheng, and Xiaolin Huang. Friendly sharpness-aware minimization. In [Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition \(CVPR\)](#), pages 5631–5640, 2024. 1
- [25] Xiuyu Li, Yijiang Liu, Long Lian, Huanrui Yang, Zhen Dong, Daniel Kang, Shanghang Zhang, and Kurt Keutzer. Q-diffusion: Quantizing diffusion models. In [Proceedings of the IEEE/CVF International Conference on Computer Vision](#), pages 17535–17545, 2023. 5
- [26] Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. Brecq: Pushing the limit of post-training quantization by block reconstruction. [arXiv preprint arXiv:2102.05426](#), 2021. 5
- [27] Yong Liu, Siqi Mai, Minhao Cheng, Xiangning Chen, Chou-Jui Hsieh, and Yang You. Random sharpness-aware minimization. In [Advances in Neural Information Processing Systems](#), pages 24543–24556. Curran Associates, Inc., 2022. 1
- [28] Clara Na, Sanket Vaibhav Mehta, and Emma Strubell. Train flat, then compress: Sharpness-aware minimization learns more compressible models. [arXiv preprint arXiv:2205.12694](#), 2022. 8
- [29] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In [International conference on machine learning](#), pages 8162–8171. PMLR, 2021. 8
- [30] Mang Ning, Mingxiao Li, Jianlin Su, Albert Ali Salah, and Itir Onal Ertugrul. Elucidating the exposure bias in diffusion models. [arXiv preprint arXiv:2308.15321](#), 2023. 6
- [31] Mang Ning, Mingxiao Li, Jianlin Su, Albert Ali Salah, and Itir Onal Ertugrul. Elucidating the exposure bias in diffusion models. [arXiv preprint arXiv:2308.15321](#), 2023. 2, 3, 6
- [32] Mang Ning, Enver Sangineto, Angelo Porrello, Simone Calderara, and Rita Cucchiara. Input perturbation reduces exposure bias in diffusion models. [arXiv preprint arXiv:2301.11706](#), 2023. 2, 3, 6, 8
- [33] William Peebles and Saining Xie. Scalable diffusion models with transformers. In [Proceedings of the IEEE/CVF international conference on computer vision](#), pages 4195–4205, 2023. 8
- [34] Zhe Qu, Xingyu Li, Rui Duan, Yao Liu, Bo Tang, and Zhuo Lu. Generalized federated learning via sharpness aware minimization. In [International Conference on Machine Learning \(ICML\)](#), pages 18250–18280. PMLR, 2022. 8
- [35] Sai Niranjan Ramachandran, Rudrabha Mukhopadhyay, Madhav Agarwal, CV Jawahar, and Vinay Namboodiri. Understanding the generalization of pretrained diffusion models on out-of-distribution data. In [Proceedings of the AAAI Conference on Artificial Intelligence](#), pages 14767–14775, 2024. 8
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In [Proceedings of the IEEE/CVF conference on computer vision and pattern recognition](#), pages 10684–10695, 2022. 8
- [37] Guangyuan SHI, JIAXIN CHEN, Wenlong Zhang, Li-Ming Zhan, and Xiao-Ming Wu. Overcoming catastrophic forgetting in incremental few-shot learning by finding flat minima. In [Advances in Neural Information Processing Systems](#), pages 6747–6761. Curran Associates, Inc., 2021. 4
- [38] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In [International conference on machine learning](#), pages 2256–2265. pmlr, 2015. 3
- [39] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. [Advances in neural information processing systems](#), 32, 2019. 3, 8
- [40] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. [Advances in neural information processing systems](#), 33:12438–12448, 2020. 8
- [41] Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. In [Uncertainty in artificial intelligence](#), pages 574–584. PMLR, 2020.
- [42] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. [arXiv preprint arXiv:2011.13456](#), 2020. 3, 8
- [43] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. 2023. 8
- [44] Zhendong Wang, Yifan Jiang, Huangjie Zheng, Peihao Wang, Pengcheng He, Zhangyang Wang, Weizhu Chen, Mingyuan Zhou, et al. Patch diffusion: Faster and more data-efficient training of diffusion models. [Advances in neural information processing systems](#), 36:72137–72154, 2023. 8
- [45] Junyi Wu, Haoxuan Wang, Yuzhang Shang, Mubarak Shah, and Yan Yan. Ptq4dit: Post-training quantization for diffusion transformers. [Advances in Neural Information Processing Systems](#), 37:62732–62755, 2025. 5
- [46] Tianshuo Xu, Peng Mi, Ruilin Wang, and Yingcong Chen. Why diffusion models are stable and how to make them faster: An empirical investigation and optimization. 8
- [47] TaeHo Yoon, Joo Young Choi, Sehyun Kwon, and Ernest K Ryu. Diffusion probabilistic models generalize when they fail to memorize. In [ICML 2023 workshop on structured probabilistic inference](#) $\{\&\}$ generative modeling, 2023. 8
- [48] Yefan Zhou, Yaoqing Yang, Arin Chang, and Michael W Mahoney. A three-regime model of network pruning. In [International Conference on Machine Learning](#), pages 42790–42809. PMLR, 2023. 8