

# Timestep-Aware Correction for Quantized Diffusion Models

Yuzhe Yao<sup>\*1,2</sup>, Feng Tian<sup>2,3</sup>, Jun Chen<sup>\*\*4,5</sup>, Haonan Lin<sup>2,3</sup>, Guang Dai<sup>5</sup>,  
Yong Liu<sup>4,5</sup>, and Jingdong Wang<sup>7</sup>

<sup>1</sup> School of Automation Science and Engineering, Xi'an Jiaotong University

<sup>2</sup> MOEKLINNS Lab, Xi'an Jiaotong University

<sup>3</sup> School of Computer Science and Technology, Shaanxi Province Key Laboratory of  
Big Data Knowledge Engineering, Xi'an Jiaotong University

<sup>4</sup> Zhejiang Normal University

<sup>5</sup> SGIT AI Lab, State Grid Corporation of China

<sup>6</sup> Zhejiang University

<sup>7</sup> Baidu Inc

{yuzheyao.22,junc.change,guang.gdai}@gmail.com,  
fengtian@mail.xjtu.edu.cn, linhaonan@stu.xjtu.edu.cn,  
yongliu@iipc.zju.edu.cn, wangjingdong@outlook.com

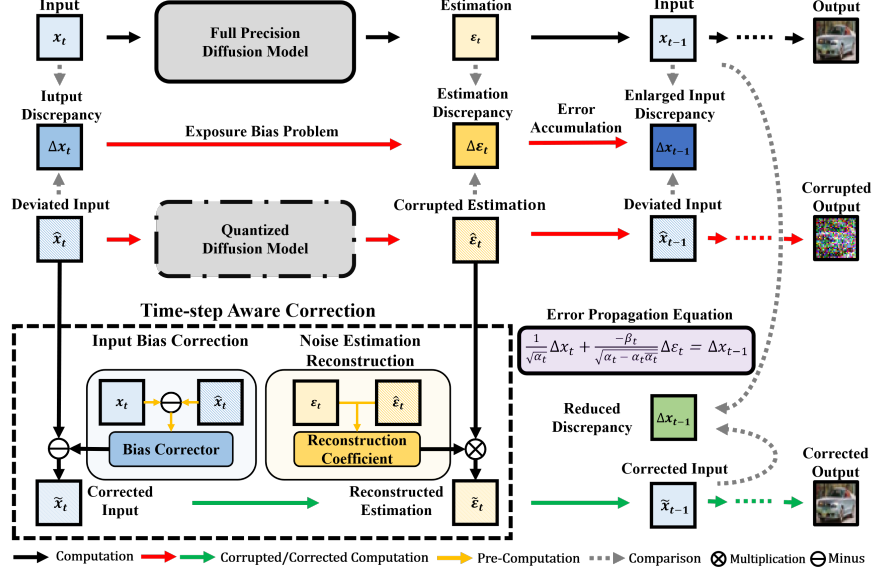
**Abstract.** Diffusion models have marked a significant breakthrough in the synthesis of semantically coherent images. However, their extensive noise estimation networks and the iterative generation process limit their wider application, particularly on resource-constrained platforms like mobile devices. Existing post-training quantization (PTQ) methods have managed to compress diffusion models to low precision. Nevertheless, due to the iterative nature of diffusion models, quantization errors tend to accumulate throughout the generation process. This accumulation of error becomes particularly problematic in low-precision scenarios, leading to significant distortions in the generated images. We attribute this accumulation issue to two main causes: error propagation and exposure bias. To address these problems, we propose a timestep-aware correction method for quantized diffusion model, which dynamically corrects the quantization error. By leveraging the proposed method in low-precision diffusion models, substantial enhancement of output quality could be achieved with only negligible computation overhead. Extensive experiments underscore our method’s effectiveness and generalizability. By employing the proposed correction strategy, we achieve state-of-the-art (SOTA) results on low-precision models.

**Keywords:** Diffusion Models · Post-training Quantization

---

\* This work was completed during an internship at SGIT AI Lab, State Grid Corporation of China.

\*\* Corresponding Author



**Fig. 1: (Upper)** Illustration of error accumulation in diffusion models. Inherent to the design of DMs, discrepancy in the input not only propagates to the next timestep but also leads to significant discrepancy in noise estimation, due to exposure bias. This cascading effect amplifies errors in subsequent stages, cumulatively impairing the quality of the final output. **(Lower-left)** To address this challenge, our method focuses on reducing the accumulated error  $\Delta \mathbf{x}_0$  through two key strategies: 1) minimizing the discrepancy  $\Delta \mathbf{x}_{t-1}$  at each timestep  $t < T$ , and 2) decomposing  $\Delta \mathbf{x}_{t-1}$  into two distinct components—the input discrepancy  $\Delta \mathbf{x}_t$  and the noise estimation discrepancy  $\Delta \epsilon_t$ —and rectifying them separately to enhance error correction.

## 1 Introduction

Diffusion models (DMs) [17, 49, 50] have emerged as powerful deep generative models for various applications, including image synthesis, Text-to-Image generation, video generation, and medical image reconstruction [6, 11, 39, 42–44, 52, 53, 55]. Despite their ability to generate images with high fidelity and diversity, diffusion models are impeded by a time-consuming and computationally intensive synthesis process. This complexity arises primarily from two factors. First, diffusion models employ a complex deep neural network, such as U-Net [41], for noise estimation. Second, to maintain the quality of the synthesized images, diffusion models necessitate an iterative process that progressively denoises the input image. This procedure can require up to 1,000 iterations, substantially contributing to the overall computational burden.

In response to the significant computational demands of the diffusion model’s noise estimation process, researchers have been pursuing various methods regarding the two factors that slow down the inference. One approach focuses on the refinement of the sampling trajectory, which can effectively shorten the noise es-

timization phase, as explored in several studies [3, 30, 31, 35, 50]. Another approach aims to reduce the latency of the noise estimation network in diffusion models by employing model quantization [38]. While Quantization-Aware Training (QAT) [12, 21, 48] requires retraining the neural network with simulated quantization and hyper-parameter search, it is more computation-intensive and requires more engineering effort for deployment [10, 34]. In contrast, Post-Training Quantization (PTQ) emerges as a preferred choice for lightweight diffusion models due to its straightforward approach [15, 24, 47]. However, even though PTQ for models like MobileNet [18] and ResNet [13] is a well-studied technique [2, 4, 29, 33, 34], its application in PTQ in diffusion models warrants further investigation, particularly due to the dynamic nature of DMs.

Recent advancements in PTQ for DMs have enabled the quantization of model weights to 4-bit precision [15, 19, 24, 47]. However, these methods often face challenges in preserving image quality at lower bit-widths (e.g., 3-bit and 2-bit weight parameters), primarily due to error accumulation in DMs [15, 24, 26].

We attribute error accumulation in diffusion models to two main factors: error propagation and exposure bias. Mathematical analysis by recent studies [25] reveals that errors in the inputs of DMs will propagate to subsequent denoising steps. This is particularly problematic in low-precision DMs, where each inference step inherently introduces additional errors, leading to significant cumulative effects. Moreover, the input discrepancy causes exposure bias problem, exacerbating the issue by causing a significant discrepancy in the noise estimation from the ground truth [23, 36, 37, 46], resulting in compounded errors in the model’s output. The error accumulation challenge is depicted in the upper section of Fig. 1.

To counteract error accumulation, we introduce a timestep-aware correction strategy. This approach, integrating Noise Estimation Reconstruction (NER) and Input Bias Correction (IBC), dynamically mitigates quantization-induced errors at every generation step, significantly improving the fidelity of images produced by low-precision DMs. Importantly, our method does not require additional training or tuning, allowing for seamless integration into existing PTQ frameworks. The operational details of our method are illustrated in the lower-left segment of Fig. 1.

Our study’s core contributions are as follows:

1. We identify error accumulation as a key challenge in PTQ for diffusion models, attributing it primarily to error propagation and exposure bias. In response, we present Timestep-Aware Correction for Quantized Diffusion Models (TAC-Diffusion), a novel approach for dynamically mitigating errors across the generative process.
2. We decompose the error in the model’s input into two main components: corrupted noise estimation and deviated input from previous timesteps. We propose Noise Estimation Reconstruction (NER) and Input Bias Correction (IBC) to correct each aspect.
3. Extensive experiments across various diffusion models and samplers for both conditional and unconditional image generation tasks demonstrate our method’s

effectiveness. Notably, TAC-Diffusion significantly outperforms the state-of-the-art Q-Diffusion method [24], achieving a FID improvement of 7.76 on CIFAR-10 with 3-bit weight and 8-bit activation quantization.

## 2 Related work

### 2.1 Diffusion Models.

Diffusion models are latent variable models designed to generate images that match the training set’s distribution  $q(\mathbf{x}_0)$  from random Gaussian noise [17, 39, 49, 50]. The generating process can be modeled as:

$$p_\theta(\mathbf{x}_0) := \int p_\theta(\mathbf{x}_{0:T}) d\mathbf{x}_{1:T}, \quad (1)$$

where  $\mathbf{x}_{0:T}$  represents a sequence of latent variables from the initial state  $\mathbf{x}_0$  to the final state  $\mathbf{x}_T$ ,  $\theta$  denotes the parameters,  $p_\theta(\mathbf{x}_0)$  denotes the model distribution and  $p_\theta(\mathbf{x}_{0:T})$  denotes the reverse process [50].

The training of diffusion models involves two key processes: the forward process and the reverse process. In the forward process, Gaussian noise, scheduled by the parameter  $\beta_t$ , is incrementally added to the latent variable  $\mathbf{x}_{t-1}$  at each timestep. This process can be mathematically described as follows:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t \mathbf{I}). \quad (2)$$

With an appropriately set noise schedule  $\beta_t$ , the variable  $\mathbf{x}_t$  gradually converges to a standard Gaussian distribution as  $t$  increases, for all initial values  $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ . This convergence is crucial for the reverse process, where the model generates images by sampling noise from Gaussian distribution and iteratively refining the sample through a series of steps to produce the final output image. In the reverse process, each iteration of DMs can be formulated as:

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)), \quad (3)$$

where  $p(\mathbf{x}_T) \sim \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$ , with  $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$  denotes the model estimation for noise, and  $\boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)$  denotes the variance for sampling which can be fixed to constants [17]. When the training process ends, Eq. (3) is used for inference. While in DDPM [17], the noise estimation process is a Markov process, the inference of diffusion models is inevitably time-consuming and computationally resource-intensive.

### 2.2 Post-Training Quantization.

Post-Training Quantization (PTQ) is a prominent technique for compressing neural network models to reduce inference latency [1, 4, 5, 29, 34]. PTQ converts full-precision (typically FP32) models to a fixed-point format by reducing the

bit-width of weights and activations. The method comprises two principal operations: quantization, which utilizes a lower bit width to enhance computation efficiency, and de-quantization, which reverses this process for operations that necessitate floating-point accuracy. The processes are defined as follows:

$$X_q = \text{Clip} \left( \text{Round} \left( \frac{X_{fp}}{s} \right) - z, q_{\min}, q_{\max} \right), \quad X_{deq} = s(X_q + z), \quad (4)$$

where  $s$  is the scaling factor determined by the range of full-precision data,  $z$  represents the zero point, and  $q_{\max}$  and  $q_{\min}$  denote the clamping thresholds.

Calibration is then introduced to determine the optimal quantization thresholds ( $q_{\max}$  and  $q_{\min}$ ), ensures that the quantized output retains the most significant data [20]. During these processes, PTQ introduces two types of errors: rounding error and clamping error. Rounding error occurs due to the quantization of continuous values to discrete counterparts, while clamping error results from constraining values within the specified  $q_{\min}$  and  $q_{\max}$  thresholds to prevent overflow or underflow. Efforts to mitigate rounding errors include task-specific rounding strategy [32] and block reconstruction [27]. To address clamping error, Esseret *al.* [7] carefully selects the quantization threshold during calibration.

### 2.3 Challenges in PTQ for Diffusion Models

One primary challenge in PTQ for diffusion models is the variability of activation ranges, where differences across timesteps lead to clamping errors due to the use of a fixed quantization step size  $s$ . In response, Shang et al. [47] devised a calibration set that specifically incorporates samples near  $t = 0$ , using a skew-normal distribution. Li et al. [24] introduced uniform timestep sampling and shortcut-splitting quantization as strategies to alleviate quantization errors in UNet structures. He et al. [15] recommended varying the bit-widths for activation quantization across timesteps, coupled with variance re-scheduling, to effectively manage quantization errors. Similarly, Ho et al. [48] introduced a novel method that dynamically adjusts the quantization interval based on timestep information. These innovative approaches have successfully preserved the performance of quantized DMs with 4-bit weights and 8-bit activations.

Despite the advancements in PTQ for DMs, another challenge, the error accumulation, remains underexplored. Error accumulation refers to a phenomenon where initial errors propagate and amplify in later timesteps due to exposure bias [23, 25, 36, 37]. In our work, by addressing this key issue during models' inference, we aim to close the gap between quantized diffusion models and their full-precision counterparts.

## 3 Problem Assumption and Statement

**Notation** Given a full-precision noise estimation network in diffusion models,  $\epsilon_\theta$ , we denote its quantized version as  $\hat{\epsilon}_\theta$ , their input latent variable at timestep  $t$  as  $\mathbf{x}_t$  and  $\hat{\mathbf{x}}_t$ , their noise estimation as  $\epsilon_\theta(\mathbf{x}_t, t)$  and  $\hat{\epsilon}_\theta(\hat{\mathbf{x}}_t, t)$ , respectively.

In DDPM [17], the input of a full-precision noise estimation network at time step  $t - 1$  can be formulated as:

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}, \quad (5)$$

where  $\alpha_t$ ,  $\beta_t$ ,  $\bar{\alpha}_t$  and  $\sigma_t$  are fixed full-precision constants, and  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . We define the error on the input of a quantized model at timestep  $t - 1$  as:

$$\begin{aligned} \Delta \mathbf{x}_{t-1} &= \hat{\mathbf{x}}_{t-1} - \mathbf{x}_{t-1} \\ &= \frac{1}{\sqrt{\alpha_t}} (\hat{\mathbf{x}}_t - \mathbf{x}_t) - \frac{\beta_t}{\sqrt{\alpha_t - \alpha_t \bar{\alpha}_t}} (\hat{\boldsymbol{\epsilon}}_\theta(\hat{\mathbf{x}}_t, t) - \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)) \\ &= \frac{1}{\sqrt{\alpha_t}} \Delta \mathbf{x}_t - \frac{\beta_t}{\sqrt{\alpha_t - \alpha_t \bar{\alpha}_t}} \Delta \boldsymbol{\epsilon}_t. \end{aligned} \quad (6)$$

Based on the above decomposition, we can deduce that the input error at  $t - 1$  stems from two sources: the input discrepancy and the noise estimation discrepancy at the previous timestep, represented by  $\Delta \mathbf{x}_t$  and  $\Delta \boldsymbol{\epsilon}_t$ , respectively. Therefore any input discrepancy will propagate to next timestep. Since the noise estimation network  $\boldsymbol{\epsilon}_\theta$  is non-linear, and the quantized model  $\hat{\boldsymbol{\epsilon}}_\theta$  has a different input-output mapping compared to its full-precision version  $\boldsymbol{\epsilon}_\theta$ , it is difficult to predict the input error's influence on the noise estimation. As a consequence, we propose to minimize these two terms separately. We have:

$$\begin{aligned} \|\Delta \mathbf{x}_{t-1}\| &= \left\| \frac{1}{\sqrt{\alpha_t}} \Delta \mathbf{x}_t - \frac{\beta_t}{\sqrt{\alpha_t - \alpha_t \bar{\alpha}_t}} \Delta \boldsymbol{\epsilon}_t \right\| \\ &\leq \frac{1}{\sqrt{\alpha_t}} \|\Delta \mathbf{x}_t\| + \frac{\beta_t}{\sqrt{\alpha_t - \alpha_t \bar{\alpha}_t}} \|\Delta \boldsymbol{\epsilon}_t\|. \end{aligned} \quad (7)$$

Above formulation indicates that, if we can effectively rectify  $\Delta \mathbf{x}_t$  and  $\Delta \boldsymbol{\epsilon}_t$ , we might success in reducing  $\Delta \mathbf{x}_{t-1}$ . This further leads us to one solution for the error accumulation challenge in quantized DMs: by rectifying the error in the input and the noise estimation for all timestep, the error accumulated at the end  $\Delta \mathbf{x}_0$  shall be reduced, which indicates that high-quality synthesis images would be obtainable with quantized diffusion models.

**Assumption 1** [25] *Given a quantized diffusion model and its full-precision version, at timestep  $t$ , reducing their input discrepancy  $\Delta \mathbf{x}_t$ , and their noise estimation discrepancy,  $\Delta \boldsymbol{\epsilon}_t$ , is to reduce the discrepancy between their final outputs,  $\Delta \mathbf{x}_0$ .*

Consequently, we propose to reduce the input discrepancy and the noise estimation discrepancy separately for each timestep.

## 4 Timestep-Aware Correction

In this section, we present the timestep-aware correction method, which dynamically reduces the quantization error accumulated on the latent image. At each timestep, the correction consists of two processes: **noise estimation reconstruction** (NER) and **input bias correction** (IBC). With the proposed method, we can effectively correct the quantization error and alleviate the error accumulation during the inference of low-precision DMs.

### 4.1 Noise Estimation Reconstruction

To reduce the noise estimation discrepancy between the full-precision model and its quantized version, an intuitive method is to introduce an additional neural network to predict error in the corrupted estimation. However, the complexity of this approach is not consistent with the simplicity of PTQ. Therefore, we aim to find a simple yet effective way to reconstruct the noise estimation.

Previous work in PTQ [1, 5] succeeded in reconstructing the activation map in intermediate layers using damaged activation. Inspired by this approach, we assume that the noise estimation from a full-precision model  $\epsilon_t$  can be partially reconstructed from its corrupted version  $\hat{\epsilon}_t$ , with a linear function.

**Assumption 2** *Given a quantized diffusion model, if its input distortion to the full-precision one  $\hat{\mathbf{x}}_t - \mathbf{x}_t$  is small enough, e.g.  $\hat{\mathbf{x}}_t \approx \mathbf{x}_t$ , the noise estimation from the full-precision diffusion model  $\epsilon_\theta(\mathbf{x}_t, t)$  can be partially reconstructed by scaling the noise estimation from the quantized model  $\hat{\epsilon}_\theta(\hat{\mathbf{x}}_t, t)$  with a channel-wise scaling factor.*

We then introduce the timestep-dependant reconstruction coefficient  $\mathbf{K} \in \mathbb{R}^{T \times C}$ , where  $T$  denotes the number of denoising steps, and  $C$  denotes the number of channels of estimated noise. Then the reconstructed noise in timestep  $t$  can be represented as follows:

$$\tilde{\epsilon}_\theta(\hat{\mathbf{x}}_t, t) = \mathbf{K}_t \cdot \hat{\epsilon}_\theta(\hat{\mathbf{x}}_t, t), \quad (8)$$

where  $(\cdot)$  denote the channel-wise multiplication. We now focus on finding a loss function  $\mathcal{L}$ , by minimizing which, we can efficiently reconstruct  $\epsilon_\theta(\mathbf{x}_t, t)$ ,  $\forall t \in [0, T]$ :

$$\min \mathcal{L}(\epsilon_\theta(\mathbf{x}_t, t), \tilde{\epsilon}_\theta(\hat{\mathbf{x}}_t, t)). \quad (9)$$

For brevity, in Sec. 4.1, we omit the input of the noise estimation network, using  $\epsilon_{t,i}$ ,  $\hat{\epsilon}_{t,i}$  and  $\tilde{\epsilon}_{t,i}$  to denote the estimation  $\epsilon_{t,i}(\mathbf{x}_t, t)$ ,  $\hat{\epsilon}_{t,i}(\hat{\mathbf{x}}_t, t)$  and  $\tilde{\epsilon}_{t,i}(\hat{\mathbf{x}}_t, t)$ ,  $i \in [0, C-1]$ , respectively.

While simply minimizing MSE between two images might cause blurriness, inspired by previous work in PTQ [9], we measure both the absolute error and the relative error in our loss function  $\mathcal{L}$ . For the convenience of computation and

optimization, we use the inverse root quantization to noise ratio (rQNSR) [9] to measure the signal's relative distortion caused by the quantization noise:

$$\text{rQNSR}(\hat{\epsilon}_{t,i}, \epsilon_{t,i}) = \sqrt{\frac{\sum_j^H \sum_k^W (\hat{\epsilon}_{t,i,j,k} - \epsilon_{t,i,j,k})^2}{\sum_j^H \sum_k^W \epsilon_{t,i,j,k}^2}}, \quad (10)$$

where  $H$  and  $W$  denote the height and weight of the noise estimation, respectively. By taking the  $\lambda_1$  as the weight of rQNSR penalty in  $\mathcal{L}$  and  $\lambda_2$  the regularization coefficient, we define the reconstruction loss for channel  $i$  at timestep  $t$  as:

$$\begin{aligned} \mathcal{L}(\mathbf{K}_{t,i}, \hat{\epsilon}_{t,i}, \epsilon_{t,i}) &= (1 - \lambda_1) \cdot \text{MSE}(\tilde{\epsilon}_{t,i}, \epsilon_{t,i})^2 + \lambda_1 \cdot \text{rQNSR}(\tilde{\epsilon}_{t,i}, \epsilon_{t,i})^2 \\ &\quad + \lambda_2 \cdot (\mathbf{K}_{t,i} - 1)^2. \end{aligned} \quad (11)$$

Our next object is to find the optimal  $\mathbf{K}_{t,i}$  that reduces the MSE and the rQNSR simultaneously. By expanding Eq. (11), we have:

$$\begin{aligned} \mathcal{L}(\mathbf{K}_{t,i}, \hat{\epsilon}_{t,i}, \epsilon_{t,i}) &= (1 - \lambda_1) \cdot \frac{1}{N} \sum_j^H \sum_k^W (\mathbf{K}_{t,i} \hat{\epsilon}_{t,i,j,k} - \epsilon_{t,i,j,k})^2 \\ &\quad + \lambda_1 \cdot \frac{1}{N} \sum_j^H \sum_k^W \left( \frac{\mathbf{K}_{t,i} \hat{\epsilon}_{t,i,j,k} - \epsilon_{t,i,j,k}}{\epsilon_{t,i,j,k}} \right)^2 + \lambda_2 (\mathbf{K}_{t,i} - 1)^2 \\ &= A_{t,i} \cdot \mathbf{K}_{t,i}^2 + B_{t,i} \cdot \mathbf{K}_{t,i} + D_{t,i}, \end{aligned} \quad (12)$$

where

$$\begin{cases} A_{t,i} = \frac{1-\lambda_1}{N} \sum_j^H \sum_k^W \hat{\epsilon}_{t,i,j,k}^2 + \lambda_1 \sum_j^H \sum_k^W \frac{\hat{\epsilon}_{t,i,j,k}^2}{\epsilon_{t,i,j,k}^2} + \lambda_2, \\ B_{t,i} = \frac{2\lambda_1-2}{N} \sum_j^H \sum_k^W \hat{\epsilon}_{t,i,j,k} \epsilon_{t,i,j,k} - 2\lambda_1 \sum_j^H \sum_k^W \frac{\hat{\epsilon}_{t,i,j,k}}{\epsilon_{t,i,j,k}} - 2\lambda_2, \\ D_{t,i} = \frac{1-\lambda_1}{N} \sum_j^H \sum_k^W \epsilon_{t,i,j,k}^2 + \lambda_1 N + \lambda_2, \\ N = C \cdot H \cdot W. \end{cases} \quad (13)$$

The second derivative of the loss with respect to  $\mathbf{K}_{t,i}$  is:

$$\frac{\partial^2 \mathcal{L}_{t,i}}{\partial \mathbf{K}_{t,i}^2} = 2A_{t,i}. \quad (14)$$

By assigning both  $\lambda_1$  and  $\lambda_2$  as positive coefficients, we have  $A_{t,i} > 0$ . As a consequence, the reconstruction loss is convex, permitting us to minimize it with the closed-form solution which satisfies:

$$\frac{\partial \mathcal{L}(\mathbf{K}_{t,i})}{\partial \mathbf{K}_{t,i}} = 0. \quad (15)$$

Plugging Eq. (13) into Eq. (15), we have:

$$\mathbf{K}_{t,i} = -\frac{B_{t,i}}{2A_{t,i}} = \frac{(1 - \lambda_1) \sum_j^H \sum_k^W \hat{\epsilon}_{t,i,j,k} \epsilon_{t,i,j,k} + \lambda_1 N \sum_j^H \sum_k^W \frac{\hat{\epsilon}_{t,i,j,k}}{\epsilon_{t,i,j,k}} + \lambda_2 N}{(1 - \lambda_1) \sum_j^H \sum_k^W \hat{\epsilon}_{t,i,j,k}^2 + \lambda_1 N \sum_j^H \sum_k^W \frac{\hat{\epsilon}_{t,i,j,k}^2}{\epsilon_{t,i,j,k}^2} + \lambda_2 N}. \quad (16)$$



In practice, we pre-calculate  $\mathbf{K}_{t,i}$  using a small batch of samples. Considering that larger values in  $\epsilon_{t,i}$  might indicate more change to  $\mathbf{x}_{t,i}$  [50], correcting the error in larger value other than that in smaller values might lead to more effective reconstruction. Therefore, when calculating  $\mathbf{K}_{t,i}$ , we only consider pixels whose values are larger than a threshold:

$$\tau_t = k_{threshold} * \frac{1}{N} \sum_i^C \sum_j^H \sum_k^W |\epsilon_{t,i,j,k}|, \tau \in \mathbb{R}^T, \quad (17)$$

with  $k_{threshold}$  denoting the coefficient for threshold calculation. We use the element-wise mask  $M \in \mathbb{R}^{T \times C \times H \times W}$  to filter pixels with small activation:

$$M_{t,i,j,k} = \begin{cases} 1 & \text{if } \epsilon_{t,i,j,k} > \tau_t, \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

By replacing  $\hat{\epsilon}_{t,i}$  and  $\epsilon_{t,i}$  in Eq. (16) with  $\hat{\epsilon}_{t,i} \cdot M$  and  $\epsilon_{t,i} \cdot M$ , we calculate the reconstruction coefficient  $K_{t,i}$  with following formula:

$$\mathbf{K}_{t,i} = \frac{(1 - \lambda_1) \sum_j^H \sum_k^W \hat{\epsilon}_{t,i,j,k} \epsilon_{t,i,j,k} M_{t,i} + \lambda_1 N \sum_j^H \sum_k^W \frac{\hat{\epsilon}_{t,i,j,k}}{\epsilon_{t,i,j,k}} M_{t,i} + \lambda_2 N}{(1 - \lambda_1) \sum_j^H \sum_k^W \hat{\epsilon}_{t,i,j,k}^2 M_{t,i} + \lambda_1 N \sum_j^H \sum_k^W \frac{\epsilon_{t,i,j,k}^2}{\epsilon_{t,i,j,k}^2} M_{t,i} + \lambda_2 N}. \quad (19)$$

Both multiplication and division in Eq. (19) are element-wise. After pre-calculating  $\mathbf{K}_{t,i}$  for a small batch, this coefficient is directly used for estimation reconstruction during inference, as shown in Fig. 1.

## 4.2 Input Bias Correction

Diffusion models are trained on the objective [17]:

$$L_{\text{simple}}(\theta) := \mathbb{E}_{t, \mathbf{x}_0, \epsilon} \left[ \left\| \epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t) \right\|^2 \right], \quad (20)$$

where  $\theta$  denotes the model parameters,  $\epsilon$  denotes the target noise,  $\epsilon_\theta$  denotes the noise estimation network,  $\bar{\alpha}_t$  denotes the hyper-parameter and  $\mathbf{x}_0$  denotes the input image. While in the training, the DMs are given the ground-truth input  $\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ , they never receive it as input during the inference, resulting in noise estimation error. Such training-inference discrepancy in auto-regressive generative models is named as exposure bias [46]. Recent works [23, 36, 37] systematically analyze the exposure bias problem in full-precision diffusion models and succeed in improving the models' performance by addressing this problem.

We argue that in diffusion models quantized with PTQ methods, the exposure bias problem is more severe, considering that the error introduced by quantization is far larger than the training-inference discrepancy studied in [23, 25, 36, 37]. Thus it necessitates a solution to it.

According to [17], the training objective Eq. (20) is simplified from:

$$L_{t-1} - C = \mathbb{E}_{\mathbf{x}_0, \epsilon} \left[ \frac{1}{2\sigma_t^2} \left\| \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t(\mathbf{x}_0, \epsilon) - \frac{\beta_t}{\sqrt{1-\alpha_t}} \epsilon \right) - \boldsymbol{\mu}_\theta(\mathbf{x}_t(\mathbf{x}_0, \epsilon), t) \right\|^2 \right], \quad (21)$$

where  $C$  is a constant independent of  $\theta$ , and  $\boldsymbol{\mu}_\theta(\mathbf{x}_t(\mathbf{x}_0, \epsilon), t)$  is the model’s prediction for the mean of  $\mathbf{x}_{t-1}$ . Eq. (21) indicates that the  $\boldsymbol{\mu}_\theta$  must predict  $\frac{\beta_t}{\sqrt{1-\alpha_t}} \epsilon$  given  $\mathbf{x}_t$  [17]. Therefore, if the mean of  $\hat{\mathbf{x}}_t$  is close to  $\mathbf{x}_t$ , low-precision DMs might be able to predict  $\epsilon_t$  better.

We propose Input Bias Correction (IBC) to correct the input discrepancy. We first calculate the average element-level bias in the deviated model input using a mini-batch of  $S$  sample:

$$\mathbf{B}_{t,i,j,k} = \frac{1}{S} \sum_{s=0}^S (\hat{\mathbf{x}}_{t,s,i,j,k} - \mathbf{x}_{t,s,i,j,k}), \mathbf{B} \in \mathbb{R}^{T \times C \times H \times W}, \quad (22)$$

during the inference, the input discrepancy can be rectified with this correction:

$$\tilde{\mathbf{x}}_{t-1} = \hat{\mathbf{x}}_{t-1} - \mathbf{B}_{t-1}. \quad (23)$$

We further compare the IBC with the estimation bias correction strategy (solely correct the mean of  $\hat{\epsilon}_t$ ), and find that IBC performs better, which indicates that directly addressing the exposure bias problem in low-precision DMs is an effective method (see Appendix B).

The overall timestep-aware correction pipeline, including the pre-process and the inference process, is described in Appendix C. During inference, the corrected sampling process for low-precision DMs can be formulated as:

$$\tilde{\mathbf{x}}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( \hat{\mathbf{x}}_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \mathbf{K}_t \hat{\epsilon}_\theta(\hat{\mathbf{x}}_t - \mathbf{B}_t, t) - \mathbf{B}_t \right) + \mathbf{z}, \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (24)$$

## 5 Experiments

### 5.1 Experiments Setup

Following previous work for diffusion quantization [14, 15, 24, 47], we evaluate our method on both unconditional and conditional generation. For unconditional generation, we conduct experiments with quantized DDIM on CIFAR-10 [22], LDM-8 on LSUN-Church [54] and LDM-4 on LSUN-Bedroom. For conditional generation, we evaluate our method with Stable-Diffusion v1.4 [40] on text-to-image task. Notably, even with unseen prompts, TAC-Diffusion still achieves considerable improvement in generated images (see Appendix M). In experiments of 3-bit and 2-bit weight quantization, we compare our model with the state-of-the-art PTQ methods, including PTQ4DM [47], Q-Diffusion [24], and

PTQD [15], implementing their result by rerunning their official codes. The primary metric for diffusion models is Fréchet Inception Distance (FID, the lower the better) [16], we report it for all experiments. We also evaluate the Inception Score (IS, the higher the better) [45] for experiments on CIFAR-10. To ensure consistency with previous works, all results are obtained by sampling 50,000 images. To further investigate the portability of the proposed method, we extend it to another fast sampling method, DPM-solver++ [31] (see Appendix D).

## 5.2 Main Results

We conduct extensive experiments on CIFAR-10 [22] using a 100-step DDIM, across four quantization precisions: W8A8, W4A8, W3A8, and W3A6. Quantitative results are shown in Tab. 1 (see the Appendix J for corresponding qualitative results). We observe that, across all precision levels, TAC-Diffusion achieves better image fidelity and diversity, even outperforming the QAT method EfficientDM [14] under W8A8 quantization. At lower precisions, such as W3A8 and W3A6, TAC-Diffusion demonstrates robust performance, maintaining a relatively low FID of 9.55 and 31.77, respectively, thus narrowing the performance gap between a quantized model and its full-precision counterpart.

**Table 1:** Quantization results for unconditional image generation with 100 steps DDIM on CIFAR-10 ( $32 \times 32$ ).

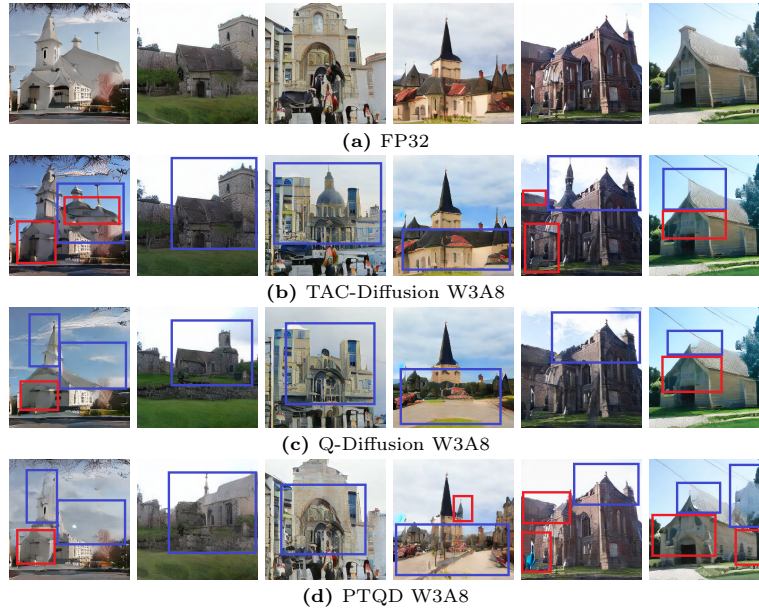
Model	Method	Bits(W/A)	Size(Mb)	IS $\uparrow$	FID $\downarrow$
DDIM (steps = 100 eta = 0.0 )	FP	32/32	143.2	9.12	4.22
	PTQ4DM [47]	8/8	34.26	9.31	10.55
	Q-diffusion [24]	8/8	35.8	9.48	3.75
	TDQ [48]	8/8	34.30	8.85	5.99
	EDA-DM [28]	8/8	–	9.40	3.72
	TFMQ-DM [19]	8/8	–	9.07	4.24
	EfficientDM (QAT) [14]	8/8	34.30	9.38	3.75
	Ours	8/8	35.8	<b>9.49</b>	<b>3.68</b>
	PTQ4DM [47]	4/8	17.22	7.92	37.79
	Q-diffusion [24]	4/8	17.31	9.12	4.93
	EDA-DM [28]	4/8	–	9.29	4.16
	TFMQ-DM [19]	4/8	–	9.13	4.78
	EfficientDM (QAT) [14]	4/8	17.26	<b>9.41</b>	<b>3.80</b>
	Ours	4/8	17.31	9.15	4.89
	PTQ4DM [47]	3/8	13.4	4.21	115.4
	Q-diffusion [24]	3/8	13.4	8.53	17.31
	Ours	3/8	13.4	<b>8.86</b>	<b>9.55</b>
	PTQ4DM [47]	3/6	13.4	4.07	124.97
	Q-diffusion [24]	3/6	13.4	<b>8.54</b>	33.21
	Ours	3/6	13.4	8.27	<b>31.77</b>

Evaluations using LDM-4 and LDM-8 on the LSUN-Church and LSUN-Bedroom datasets are detailed in Tab. 2. The quantitative outcomes for these high-resolution datasets mirror the performance improvements previously observed in experiments with lower resolutions (see Tab. 1). Remarkably, TAC-Diffusion significantly boosts the performance of W3A8/W2A8 quantized LDM-4 models, achieving FID reductions of 2.03 and 1.69, respectively, compared to the baseline reported in [24], thereby pushing the limits of post-training quantization (PTQ) for diffusion models to 2-bit precision for the first time.

**Table 2:** Quantization results for unconditional image generation with 500 steps LDM-8 on LSUN-Churches ( $256 \times 256$ ) and 200 steps LDM-4 on LSUN-Bedrooms ( $256 \times 256$ ).

Model	Method	Bits (W/A)	Size (Mb)	TBops	FID↓
LDM-8 (steps = 500 eta = 0.0 )	Full-precision	32/32	1179.9	22.17	4.06
	PTQD [15]	8/8	295.0	2.68	10.76
	QuEST [51]	8/8	330.6	–	6.55
	TFMQ-DM [19]	8/8	295.0	–	4.01
	EDA-DM [28]	8/8	–	–	3.83
	Q-diffusion [24]	8/8	295.0	2.68	3.65
	Ours	8/8	295.0	2.68	<b>3.37</b>
	PTQD [15]	4/8	147.5	1.34	8.41
	QuEST [51]	4/8	189.9	–	7.33
	TFMQ-DM [19]	4/8	147.5	–	4.14
	EDA-DM [28]	4/8	–	–	4.01
	Q-diffusion [24]	4/8	147.5	1.34	4.12
	Ours	4/8	147.5	1.34	<b>3.81</b>
	PTQD [15]	3/8	110.6	1.01	12.68
	Q-diffusion [24]	3/8	110.6	1.01	9.80
	Ours	3/8	110.6	1.01	<b>7.78</b>
LDM-4 (steps = 200 eta = 1.0 )	Full-precision	32/32	1096.2	107.17	2.98
	PTQD [15]	4/8	137.0	6.48	5.94
	Q-diffusion [24]	4/8	137.0	6.48	5.32
	Ours	4/8	137.0	6.48	<b>4.94</b>
	PTQD [15]	3/8	102.75	4.86	6.46
	Q-diffusion [24]	3/8	102.75	4.86	7.17
	Ours	3/8	102.75	4.86	<b>5.14</b>
	Q-Diffusion	2/8	68.5	2.43	9.48
	Ours	2/8	68.5	2.43	<b>7.79</b>

The corresponding qualitative results on LSUN-Church and LSUN-Bedroom are presented in Fig. 2 and the Fig. 3. We compare TAC-Diffusion with the full-precision model, Q-Diffusion [24]. These visualization demonstrates TAC-Diffusion’s superiority in preserving detailed textures and structural integrity, which are highlighted by red and blue boxes, respectively. We argue that this



**Fig. 2:** Results on LSUN-Church, with 500 steps W3A8 LDM-8 [39]. Red boxes highlight areas where our model preserves intricate details more effectively. Blue boxes show regions where our model maintains structural accuracy.



**Fig. 3:** Results on LSUN-Bedroom ( $256 \times 256$ ) with 200 steps LDM-4.

success is due to TAC-Diffusion’s ability to reduce errors during both the early and later stages, when structure and details begin to emerge [8].

### 5.3 Ablation Study

**Table 3:** Ablation study of each component in our method, with W3A8 100 steps DDIM on CIFAR-10 ( $32 \times 32$ ).

Method	Bits(W/A)	IS $\uparrow$	FID $\downarrow$
Baseline [24]	3/8	8.53	17.31
+IBC+ Timestep-Aware	3/8	8.60 (+0.07)	16.16 (-1.15)
+NER + IBC	3/8	8.67 (+0.14)	13.42 (-3.89)
+NER + IBC+ Timestep-Aware (Ours)	3/8	<b>8.86</b> (+0.33)	<b>9.55</b> (-7.76)

To evaluate the individual contributions of each component in our proposed methods, we conduct a thorough ablation study using the CIFAR-10 dataset with a W3A8 100-step DDIM sampler. The results are documented in Tab. 3, where 'IBC' represents Input Bias Correction, 'NER' stands for Noise Estimation Reconstruction, and 'Timestep-Aware' indicates that corrections are applied across all timesteps, as opposed to methods where corrections are only applied at the first timestep. We initiated our evaluation with the strong Q-Diffusion baseline [24], to evaluate the adaptability of our methods. The findings from our ablation study are summarized as follows:

(a) IBC effectively improves model's noise estimation accuracy, as evidenced by a 1.15 decrease in FID, underscoring our previous discussion on exposure bias in Sec. 4.2. (b) NER successfully reconstructs the noise estimation from the corrupted noise estimation, leading to a notable 7.76 decrease in FID. This also verifies our assumption in Sec. 4.1 that we can partially reconstruct the original noise estimation from the corrupted one with a appropriate scale factor. (c) Applying a singular correction at the initial timestep yields a 3.89 FID reduction. This is attributed to the reduction of the error at the beginning greatly helps alleviate the error introduced in subsequent timesteps. thereby mitigating the error at minimal cost. (d) Timestep-Aware Correction, by actively minimizing errors at each timestep, substantially lowers cumulative errors in the final output, leading to a 3.87 decrease in FID compared to the singular correction strategy.

## 6 Conclusion

In this paper, we present TAC-Diffusion, a Timestep-Aware Correction method that dynamically rectifies the quantization error in the denoising process of diffusion models. First, we assume that reducing the input discrepancy for each timestep will also minimize the error accumulated on the final output. Next, We propose to decompose the input discrepancy into two parts and rectify each one separately. Extensive experiments demonstrate that when employing our proposed method, diffusion models quantized to low precision can generate images that are on par with or even better than the full-precision model.

## Acknowledgements

This work was supported by National Key Research and Development Program of China (2022ZD0117102), National Natural Science Foundation of China (62293551, 62177038, 62277042, 62137002, 61721002, 61937001, 62377038). Project of China Knowledge Centre for Engineering Science and Technology.

## References

1. Bai, S., Chen, J., Shen, X., Qian, Y., Liu, Y.: Unified data-free compression: Pruning and quantization without fine-tuning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5876–5885 (October 2023)
2. Banner, R., Nahshan, Y., Soudry, D.: Post training 4-bit quantization of convolutional networks for rapid-deployment. *Advances in Neural Information Processing Systems* **32** (2019)
3. Bao, F., Li, C., Zhu, J., Zhang, B.: Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models (2022)
4. Bondarenko, Y., Nagel, M., Blankevoort, T.: Understanding and overcoming the challenges of efficient transformer quantization. *arXiv preprint arXiv:2109.12948* (2021)
5. Chen, J., Bai, S., Huang, T., Wang, M., Tian, G., Liu, Y.: Data-free quantization via mixed-precision compensation without fine-tuning. *Pattern Recognition* p. 109780 (2023)
6. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
7. Esser, S.K., McKinstry, J.L., Bablani, D., Appuswamy, R., Modha, D.S.: Learned step size quantization (2020)
8. Fang, G., Ma, X., Wang, X.: Structural pruning for diffusion models. *Advances in neural information processing systems* **36** (2024)
9. Finkelstein, A., Almog, U., Grobman, M.: Fighting quantization bias with bias. *arXiv preprint arXiv:1906.03193* (2019)
10. Gholami, A., Kim, S., Dong, Z., Yao, Z., Mahoney, M.W., Keutzer, K.: A survey of quantization methods for efficient neural network inference. In: *Low-Power Computer Vision*, pp. 291–326. Chapman and Hall/CRC (2022)
11. Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B.: Vector quantized diffusion model for text-to-image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10696–10706 (2022)
12. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149* (2015)
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
14. He, Y., Liu, J., Wu, W., Zhou, H., Zhuang, B.: EfficientDM: Efficient quantization-aware fine-tuning of low-bit diffusion models. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=UmMa3UNDAz>

15. He, Y., Liu, L., Liu, J., Wu, W., Zhou, H., Zhuang, B.: PTQD: Accurate post-training quantization for diffusion models. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=Y3g1PV5R91>
16. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
17. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
18. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017)
19. Huang, Y., Gong, R., Liu, J., Chen, T., Liu, X.: Tfmq-dm: Temporal feature maintenance quantization for diffusion models. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2024)
20. Hubara, I., Nahshan, Y., Hanani, Y., Banner, R., Soudry, D.: Improving post training neural quantization: Layer-wise calibration and integer programming. *arXiv preprint arXiv:2006.10518* (2020)
21. Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H., Kalenichenko, D.: Quantization and training of neural networks for efficient integer-arithmetic-only inference. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2704–2713 (2018)
22. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. Rep. 0, University of Toronto, Toronto, Ontario (2009), <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>
23. Li, M., Qu, T., Yao, R., Sun, W., Moens, M.F.: Alleviating exposure bias in diffusion models through sampling with shifted time steps. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=ZSD3Ml0Ke6>
24. Li, X., Liu, Y., Lian, L., Yang, H., Dong, Z., Kang, D., Zhang, S., Keutzer, K.: Q-diffusion: Quantizing diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17535–17545 (2023)
25. Li, Y., van der Schaar, M.: On error propagation of diffusion models. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=RtAct1E2zS>
26. Li, Y., Xu, S., Cao, X., Sun, X., Zhang, B.: Q-DM: An efficient low-bit quantized diffusion model. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=sFGkL5BsPi>
27. Li, Y., Gong, R., Tan, X., Yang, Y., Hu, P., Zhang, Q., Yu, F., Wang, W., Gu, S.: Brecq: Pushing the limit of post-training quantization by block reconstruction. *arXiv preprint arXiv:2102.05426* (2021)
28. Liu, X., Li, Z., Xiao, J., Gu, Q.: Enhanced distribution alignment for post-training quantization of diffusion models. *arXiv preprint arXiv:2401.04585* (2024)
29. Liu, Z., Wang, Y., Han, K., Zhang, W., Ma, S., Gao, W.: Post-training quantization for vision transformer. *Advances in Neural Information Processing Systems* **34**, 28092–28103 (2021)
30. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *arXiv preprint arXiv:2206.00927* (2022)
31. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models (2023)



32. Nagel, M., Amjad, R.A., Van Baalen, M., Louizos, C., Blankevoort, T.: Up or down? adaptive rounding for post-training quantization. In: International Conference on Machine Learning. pp. 7197–7206. PMLR (2020)
33. Nagel, M., Baalen, M.v., Blankevoort, T., Welling, M.: Data-free quantization through weight equalization and bias correction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1325–1334 (2019)
34. Nagel, M., Fournarakis, M., Amjad, R.A., Bondarenko, Y., Van Baalen, M., Blankevoort, T.: A white paper on neural network quantization. arXiv preprint arXiv:2106.08295 (2021)
35. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: International Conference on Machine Learning. pp. 8162–8171. PMLR (2021)
36. Ning, M., Li, M., Su, J., Salah, A.A., Ertugrul, I.O.: Elucidating the exposure bias in diffusion models. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=xEJMo1SpX>
37. Ning, M., Sangineto, E., Porrello, A., Calderara, S., Cucchiara, R.: Input perturbation reduces exposure bias in diffusion models. arXiv preprint arXiv:2301.11706 (2023)
38. Polino, A., Pascanu, R., Alistarh, D.: Model compression via distillation and quantization. arXiv preprint arXiv:1802.05668 (2018)
39. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
40. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022), <https://github.com/CompVis/latent-diffusion><https://arxiv.org/abs/2112.10752>
41. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
42. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22500–22510 (2023)
43. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH 2022 Conference Proceedings. pp. 1–10 (2022)
44. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems* **35**, 36479–36494 (2022)
45. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. *Advances in neural information processing systems* **29** (2016)
46. Schmidt, F.: Generalization in generation: A closer look at exposure bias. In: Birch, A., Finch, A., Hayashi, H., Konstas, I., Luong, T., Neubig, G., Oda, Y., Sudoh, K. (eds.) *Proceedings of the 3rd Workshop on Neural Generation and Translation*. pp. 157–167. Association for Computational Linguistics, Hong Kong (Nov 2019). <https://doi.org/10.18653/v1/D19-5616>, <https://aclanthology.org/D19-5616>

47. Shang, Y., Yuan, Z., Xie, B., Wu, B., Yan, Y.: Post-training quantization on diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1972–1981 (2023)
48. So, J., Lee, J., Ahn, D., Kim, H., Park, E.: Temporal dynamic quantization for diffusion models. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=D1sECc9fiG>
49. Sohl-Dickstein, J., Weiss, E.A., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics (2015)
50. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. International Conference on Learning Representations (2020), [dblp.org/rec/journals/corr/abs-2010-02502](https://dblp.org/rec/journals/corr/abs-2010-02502)
51. Wang, H., Shang, Y., Yuan, Z., Wu, J., Yan, Y.: Quest: Low-bit diffusion model quantization via efficient selective finetuning (2024)
52. Xie, Y., Li, Q.: Measurement-conditioned denoising diffusion probabilistic model for under-sampled medical image reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 655–664. Springer (2022)
53. Yang, L., Zhang, Z., Song, Y., Hong, S., Xu, R., Zhao, Y., Zhang, W., Cui, B., Yang, M.H.: Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys* **56**(4), 1–39 (2023)
54. Yu, F., Zhang, Y., Song, S., Seff, A., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. arXiv preprint arXiv:1506.03365 (2015)
55. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023)