

D²-DPM: Dual Denoising for Quantized Diffusion Probabilistic Models

Qian Zeng¹, Jie Song^{1*}, Han Zheng¹, Hao Jiang², Mingli Song¹

¹ Zhejiang University

² Alibaba Group

{qianz, sjie, h.zheng, haofeizhang, brooksong}@zju.edu.cn, aoshu.jh@alibaba-inc.com

Abstract

Diffusion models have achieved cutting-edge performance in image generation. However, their lengthy denoising process and computationally intensive score estimation network impede their scalability in low-latency and resource-constrained scenarios. Post-training quantization (PTQ) compresses and accelerates diffusion models without retraining, but it inevitably introduces additional quantization noise, resulting in mean and variance deviations. In this work, we propose D2-DPM, a dual denoising mechanism aimed at precisely mitigating the adverse effects of quantization noise on the noise estimation network. Specifically, we first unravel the impact of quantization noise on the sampling equation into two components: the mean deviation and the variance deviation. The mean deviation alters the drift coefficient of the sampling equation, influencing the trajectory trend, while the variance deviation magnifies the diffusion coefficient, impacting the convergence of the sampling trajectory. The proposed D2-DPM is thus devised to denoise the quantization noise at each time step, and then denoise the noisy sample through the inverse diffusion iterations. Experimental results demonstrate that D2-DPM achieves superior generation quality, yielding a 1.42 lower FID than the full-precision model while achieving 3.99x compression and 11.67x bit-operation acceleration.

Code — <https://github.com/TaylorJocelyn/D2-DPM>

Introduction

Diffusion models (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020; Song and Ermon 2019; Song et al. 2020) have rapidly emerged as predominant deep generative models. By leveraging intricate posterior probability modeling and stable training regimes, diffusion models effectively prevent mode collapse while achieving superior generation fidelity and diversity over GANs (Aggarwal, Mittal, and Battineni 2021) and VAEs (Kingma and Welling 2013). Recent multi-domain studies demonstrate that highly flexible diffusion models excel in various applications, including text-to-image (Zhu et al. 2023), image super-resolution (Wang et al. 2023), inpainting (Lugmayr et al. 2022), style transfer (Zhang et al. 2023), text-to-video (Singer et al. 2022) and interpretability modeling (Lee, Kim, and Kim 2022).

*Corresponding author.

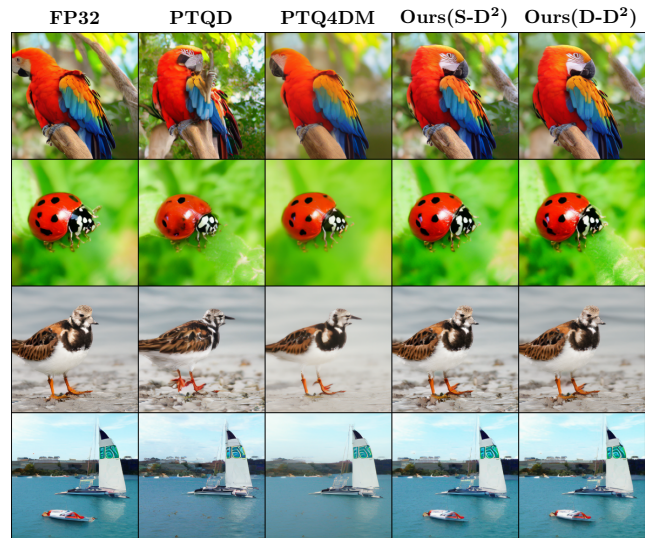


Figure 1: Comparison of generated samples on the ImageNet 256×256 between full-precision LDM-4 and its quantized versions using PTQ4DM, PTQD, and our proposed D²-DPM (comprising two variants, S-D² and D-D²).

However, the generation speed of diffusion models is constrained by two orthogonal factors: the lengthy denoising chain, involving up to 1000 steps (Ho, Jain, and Abbeel 2020), and expensive overhead at each iteration of the cumbersome noise estimation network, *i.e.*, the score estimation network (Song et al. 2020). The former challenge has been significantly alleviated through advanced learning-free samplers, which find more efficient sampling trajectories by providing high-precision numerical approximations for the stochastic differential equations (SDEs) (Dockhorn, Vahdat, and Kreis 2021) and ordinary differential equations (ODEs) (Lu et al. 2022; Liu et al. 2022) corresponding to the reverse diffusion process. However, the latter challenge remains formidable. While researchers have employed lightweight model paradigms such as pruning (Fang, Ma, and Wang 2024), knowledge distillation (Meng et al. 2023), and model quantization (Shang et al. 2023; Li et al. 2023) to further

accelerate each iteration and reduce runtime computational memory overhead for deployment on edge devices, the iterative nature of diffusion models inherently accumulates varying degrees of distortion.

Specifically, post-training quantization (PTQ) converts a pre-trained FP32 network directly into fixed-point networks with lower bit-width representations for weights and activations, bypassing the necessity for the original training pipeline. It has risen as a commonly embraced methodology owing to its practicality and ease of implementation. Nonetheless, PTQ unavoidably incurs quantization noise through the quantification of the noise estimation network, leading to deviations in the estimated mean during the inverse diffusion process and discrepancies with the predetermined variance schedule. The deviation in mean and variance significantly change the sampling trajectory, resulting in a decline in the fidelity of generated images.

In this work, we initially delve into the pattern of quantization noise based on its statistical properties, based on which we derive the denoising mechanism through the lens of the reverse-time SDE framework (Song et al. 2020). The reparameterization unveils that the mean and standard deviation induced by quantization noise impact distinct elements of the sampling equation at each time step. The mean deviation is integrated into the drift term of the reverse-time SDE, altering the direction of the sampling trajectory, while the variance deviation is superimposed on the stochastic term, resulting in increased volatility and divergence of the sampling trajectory. Building upon this insight, we propose **D²-DPM**, a dual denoising mechanism to precisely mitigate the adverse effects of quantization noise on the noise estimation network. Specifically, we propose to establish a joint Gaussian model for the quantized output and quantization noise at each timestep, enabling precise quantization noise modeling based on the quantized output during inference. With the quantization modeling, we propose two variants of D²-DPM, named *stochastic* dual denoising (S-D²) and *deterministic* dual denoising (D-D²), to eliminate the quantization noise during the inverse diffusion process. Experimental results demonstrate that the proposed method achieves significantly superior image generation quality (some examples shown in Fig. 1), in both conditional and unconditional image generation tasks.

In conclusion, we summarize our contributions as follow:

- We make the empirical observation that the quantization noise approximately follow a Gaussian distribution at each time step, which enables precise quantization noise modeling based on the quantized output at inference.
- We propose D²-DPM, a dual denoising mechanism to precisely mitigate the adverse effects of quantization noise on the noise estimation network. Innovatively, we customize various error correctors for sampling equations with different stochasticity capacities to fully utilize the additional standard deviation.
- Our extensive experiments demonstrate that, regardless of the stochasticity capacity of the sampling equations, D²-DPM achieves state-of-the-art post-training quantization performance for diffusion models.

Related Work

Diffusion Model Acceleration. To reduce inference costs while maintaining generation quality, two orthogonal diffusion model acceleration methods have emerged: i) optimizing model-agnostic sampling processes, and ii) developing more efficient score estimation models. The former includes advanced techniques such as diffusion scheme learning (Chung, Sim, and Ye 2022; Franzese et al. 2023), noise scale learning (Kingma et al. 2021; Kong and Ping 2021), and learning-free samplers based on SDE/ODE acceleration techniques. Meanwhile, the latter leverages model compress paradigms tailored to the intrinsic properties of diffusion models. (Salimans and Ho 2022; Song et al. 2023; Berthelot et al. 2023) employed knowledge distillation, utilizing ODE formulations parallels mapping the prior distribution to the target distribution through efficient paths within the distribution domain. (Fang, Ma, and Wang 2024) applied Diff-Pruning, achieving approximately a 50% reduction in FLOPs. (He et al. 2023) proposed a quantization-aware variant of the low-rank adapter (QALoRA) that can be merged with model weights and jointly quantized to low bit-width. Although these methods significantly lower per-iteration inference overhead, they require retraining the network, resulting in substantial additional time and computational costs.

Model Quantization. Quantization is a widely used technique for memory compression and computational acceleration (Liu et al. 2021; Fan et al. 2024; Lin et al. 2023). It includes quantization-aware training (QAT) (Nagel et al. 2022; Chu, Li, and Zhang 2024), and post-training quantization (PTQ) (Nagel et al. 2021; Yao et al. 2022; Xiao et al. 2023). QAT retrains the network to model quantization noise, incurring significant computational overhead. In contrast, PTQ uses a small calibration dataset to quantize network parameters to low-bit fixed-point values. PTQ commonly employs uniform asymmetric quantization, where the parameters include the scale factor s , zero-point z and quantization bit-width b . A floating-point value x is quantized to a fixed-point value x_{int} through the preceding parameters:

$$x_{\text{int}} = \text{clamp} \left(\left\lfloor \frac{x}{s} \right\rfloor + z, 0, 2^b \right), \quad (1)$$

where $\lfloor \cdot \rfloor$ is the round operation and clamp is a truncation function. In practice, quantization parameters are often derived by minimizing the MSE between the pre- and post-quantization weight or activation tensors.

PTQ on Diffusion Models. Until now, only a limited number of studies have delved into post-training quantization for diffusion models. PTQ4DM (Shang et al. 2023) introduces a calibration sampling strategy based on normal distribution but restricts its experimentation to low-resolution datasets. Q-diffusion (Li et al. 2023) proposes a time step-aware calibration strategy and shortcut-splitting quantization for Unet. PTQD (He et al. 2024) proposes a PTQ error correction method based on the assumption that the quantization noise is linear correlated with the quantized out, a premise that does not always hold true at various time step. In this work, we adopt a joint Gaussian distribution to modeling the quantization noise, which yield significantly superior performance in diverse experimental settings.

Preliminaries

Differential Equation Background in Diffusion

Diffusion models progressively add isotropic Gaussian noise with a variance schedule $\beta_1, \dots, \beta_T \in (0, 1)$ to real data \mathbf{x}_0 along the forward propagation chain $\{\mathbf{x}_0, \dots, \mathbf{x}_T\}$, then approximate the posterior probability $p(\mathbf{x}_{t-1}|\mathbf{x}_t)$ by learning the denoising process. During inference, they iteratively denoise the noisy sample along the learned posterior probability chain $q_\theta(\mathbf{x}_{i-1}|\mathbf{x}_i)$ to generate images.

The forward SDE. Song et al. (Song et al. 2020) extend the discrete-time propagation chain to continuous-time space by stochastic differential equations. In this theoretical framework, the diffusion process can be modeled as a solution to an Itô SDE:

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

where \mathbf{w} is the standard Wiener process (*a.k.a.*, Brownian motion), $\mathbf{f}(\cdot, t) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a vector-valued function called the drift coefficient of $\mathbf{x}(t)$, and $g(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ is a scalar function known as the diffusion coefficient of $\mathbf{x}(t)$, representing the stochasticity capacity.

Inverse Sampling Equation. The predominant sampling methodologies are categorized into *stochastic* and *deterministic* sampling. Stochastic sampling follows Anderson’s reverse-time SDE (Anderson 1982):

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

where $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ is the score function (Bao et al. 2022), $\bar{\mathbf{w}}$ is a standard Wiener process when time flows backwards from T to 0 , and dt is an infinitesimal negative timestep.

Deterministic sampling is typically formalized as the Probability Flow ODE sampling equation, derived by (Song et al. 2020) from the Fokker-Planck equation, ensuring that the corresponding probability densities $p_t(\mathbf{x})$ of the reverse-time SDE and ODE are equivalent at any given time t . The equation is as follows:

$$d\mathbf{x} = \left[\mathbf{f}(\mathbf{x}, t) - \frac{1}{2} g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) \right] dt, \quad (4)$$

To estimate the score $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ in Eqn. (3) and Eqn. (4), it is common to train a time-independent score-based model $\mathbf{s}_\theta(\mathbf{x}_t, t)$, which is linearly related to the denoising network:

$$\mathbf{s}_\theta(\mathbf{x}, t) \triangleq -\frac{\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)}{\sigma_t}, \quad (5)$$

where σ_t is the standard deviation of $p_{0t}(\mathbf{x}_t|\mathbf{x}_0)$, referred to as the noise schedule and $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)$ is the noise prediction network, from which quantization noise is introduced.

ODE-based samplers achieve faster sampling speeds due to the deterministic nature of their components. SDE-based samplers leverage the stochasticity provided by stochastic term $g(t)d\bar{\mathbf{w}}$ to achieve better generation quality. This stochasticity essentially functions as implicit Langevin diffusion, driving the sample towards the desired marginal distribution over time while correcting any errors made in earlier sampling steps. Inspired by this, we utilize the effective components of quantization errors to supplement stochasticity, thereby ensuring high generation quality.

Pre-analysis: Quantization Noise on Diffusion

We first make an empirical analysis of the quantization noise incurred by the quantized noise estimation model. At time step t , we use $\boldsymbol{\epsilon}_\theta^{(t)}$ to denote the full-precision output, $\hat{\boldsymbol{\epsilon}}_\theta^{(t)}$ the quantized output, and $\Delta\boldsymbol{\epsilon}_\theta^{(t)}$ the quantization noise, $\Delta\boldsymbol{\epsilon}_\theta^{(t)} = \hat{\boldsymbol{\epsilon}}_\theta^{(t)} - \boldsymbol{\epsilon}_\theta^{(t)}$. Fig. 2 depicts some example results on LDM-4 (Rombach et al. 2022) with the W4A8 (4-bit weights and 8-bit activations) quantization. Experiments on more diffusion models draw similar conclusions and are provided in the Appendix. From these results, we make following main findings.

Observation #1: *At each time step, the quantization noise $\Delta\boldsymbol{\epsilon}_\theta^{(t)}$ approximately follows a Gaussian distribution: $\Delta\boldsymbol{\epsilon}_\theta^{(t)} \sim \mathcal{N}(\boldsymbol{\mu}_\Delta(t), \boldsymbol{\Sigma}_\Delta(t))$.*

As a non-cherry-pick example, the distribution of the 3^{rd} element in $\Delta\boldsymbol{\epsilon}_\theta^{(0.5T)}$ is illustrated in the Fig. 2a. More results can be found in the Appendix. It is evident that the KDE-fitted probability density curve and the Gaussian density curve overlap almost exactly, which leads to the above observation.

Observation #2: *At each time step, the quantized output $\hat{\boldsymbol{\epsilon}}_\theta^{(t)}$ approximately follows a Gaussian distribution: $\hat{\boldsymbol{\epsilon}}_\theta^{(t)} \sim \mathcal{N}(\boldsymbol{\mu}_{\hat{\boldsymbol{\epsilon}}}(t), \boldsymbol{\Sigma}_{\hat{\boldsymbol{\epsilon}}}(t))$.*

This is a straightforward finding as the full-precision output $\boldsymbol{\epsilon}_\theta^{(t)}$ follows a Gaussian distribution by its nature. If the quantization noise $\Delta\boldsymbol{\epsilon}_\theta^{(t)}$ follows a Gaussian distribution (*Observation #1*), then the quantized output $\hat{\boldsymbol{\epsilon}}_\theta^{(t)}$ also follows a Gaussian distribution. Fig. 2b depicts distribution of the 5^{th} element of $\hat{\boldsymbol{\epsilon}}_\theta^{(0.5T)}$, which showcase that the quantized output $\hat{\boldsymbol{\epsilon}}_\theta^{(t)}$ approximately follows a Gaussian distribution.

Quantization Noise on SDE

With quantization noise $\Delta\boldsymbol{\epsilon}_\theta^{(t)} \sim \mathcal{N}(\boldsymbol{\mu}_\Delta(t), \boldsymbol{\Sigma}_\Delta(t))$ from *Observation #1*, the SDE-based sampling with quantization noise can be reformulated as follows:

$$d\mathbf{x} = \left[\mathbf{f}(\mathbf{x}, t) + g(t)^2 \frac{\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) + \boldsymbol{\mu}_\Delta(t)}{\sigma_t} \right] dt + \left[g(t) + \frac{g(t)^2 \sigma_\Delta(t) \sqrt{dt}}{\sigma_t} \right] d\bar{\mathbf{w}}, \quad (6)$$

where $\sigma_\Delta(t)$ represents the standard deviation of the isotropic standard Gaussian component of $\boldsymbol{\Sigma}_\Delta(t)$. From the above equation, it is evident that $\boldsymbol{\mu}_\Delta(t)$ and $\sigma_\Delta(t)$ independently affect the inverse sampling equation. Specifically, the mean alters the drift term, altering the sampling direction, while the variance increases the diffusion coefficient, impacting the fluctuation and convergence of the sampling trajectory. Therefore, we separately formulate equations for expectation and variance to conduct the analysis, followed by performing mean and variance corrections in a fully disentangled manner.

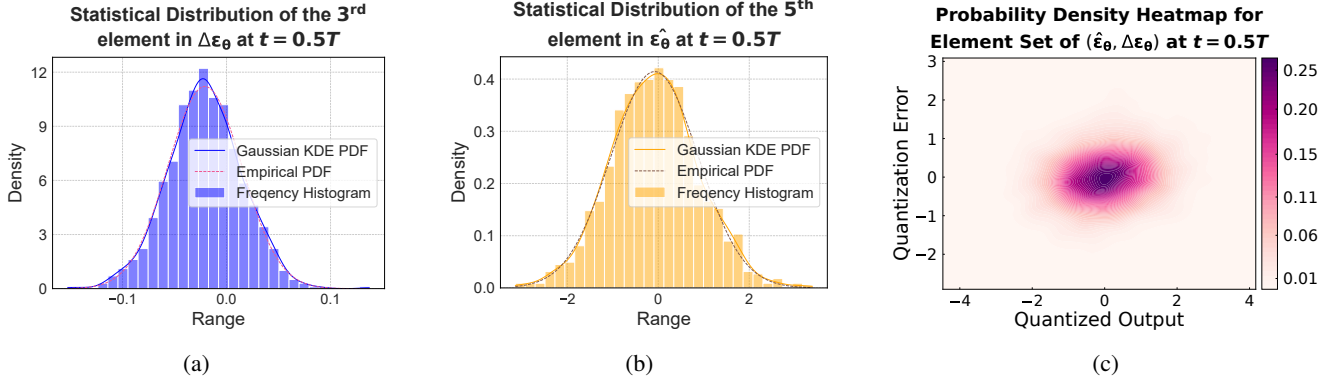


Figure 2: The statistical characteristics of $\Delta\epsilon_\theta$ and $\hat{\epsilon}_\theta$ on quantifying full-precision LDM-4 (Rombach et al. 2022) to W4A8 (4-bit for weights, 8-bit for activations) LDM-4. (a) The statistical distribution of the 3rd element of $\Delta\epsilon_\theta^{(0.5T)}$. (b) The statistical distribution of the 5th element of $\hat{\epsilon}_\theta^{(0.5T)}$. (c) The probability density heatmap for element set of $(\hat{\epsilon}_\theta^{(0.5T)}, \Delta\epsilon_\theta^{(0.5T)})$.

The Proposed Method

Time Step-aware Quantization Noise Modeling

For the sake of clarity, we use $\hat{\mathcal{E}}$ and $\Delta\mathcal{E}$ to denote the variables of $\hat{\epsilon}$ and $\Delta\epsilon$, respectively. With Observation #1 and #2, we can employ a Gaussian distribution to model the joint distribution of $\hat{\mathcal{E}}$ and $\Delta\mathcal{E}$ (as shown in Fig. 2c):

$$\begin{bmatrix} \hat{\mathcal{E}} \\ \Delta\mathcal{E} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \hat{\mu} \\ \Delta\mu \end{bmatrix}, \begin{bmatrix} \Sigma_{\hat{\mathcal{E}}, \hat{\mathcal{E}}} & \Sigma_{\hat{\mathcal{E}}, \Delta\mathcal{E}} \\ \Sigma_{\Delta\mathcal{E}, \hat{\mathcal{E}}} & \Sigma_{\Delta\mathcal{E}, \Delta\mathcal{E}} \end{bmatrix}\right) \quad (7)$$

Note that at different time step t , $\hat{\mathcal{E}}$ and $\Delta\mathcal{E}$ jointly follow a different Gaussian distribution (thus termed *time step-aware*). Here we omit the time step t for the symbol simplicity. $\hat{\mu}$ and $\Delta\mu$ denote the mean of the quantized output and the quantization noise, respectively. $\Sigma_{\hat{\mathcal{E}}, \Delta\mathcal{E}}$ denote the cross-covariance between the quantized output $\hat{\mathcal{E}}$ and the quantization noise $\Delta\mathcal{E}$. With Eqn. 7, we can derive the distribution of the quantization noise $\Delta\mathcal{E}$ conditioned on the quantized output $\hat{\mathcal{E}} = \hat{\epsilon}$ as follows:

$$\{\Delta\mathcal{E} | \hat{\mathcal{E}} = \hat{\epsilon}\} \sim \mathcal{N}\left(\mu_{\Delta\mathcal{E} | \hat{\mathcal{E}} = \hat{\epsilon}}, \Sigma_{\Delta\mathcal{E} | \hat{\mathcal{E}} = \hat{\epsilon}}\right), \quad (8)$$

$$\mu_{\Delta\mathcal{E} | \hat{\mathcal{E}} = \hat{\epsilon}} = \Sigma_{\Delta\mathcal{E}, \hat{\mathcal{E}}} \Sigma_{\hat{\mathcal{E}}, \hat{\mathcal{E}}}^{-1} (\hat{\epsilon} - \hat{\mu}) + \Delta\mu, \quad (9)$$

$$\Sigma_{\Delta\mathcal{E} | \hat{\mathcal{E}} = \hat{\epsilon}} = \Sigma_{\Delta\mathcal{E}, \Delta\mathcal{E}} - \Sigma_{\Delta\mathcal{E}, \hat{\mathcal{E}}} \Sigma_{\hat{\mathcal{E}}, \hat{\mathcal{E}}}^{-1} \Sigma_{\hat{\mathcal{E}}, \Delta\mathcal{E}}. \quad (10)$$

However, directly estimating the joint distribution in Eqn. 7 can be problematic due to the high dimensions of the joint space of $\hat{\mathcal{E}}$ and $\Delta\mathcal{E}$. In this work, we make the assumption that elements in $\hat{\mathcal{E}}$ ($\Delta\mathcal{E}$) are uncorrelated (but the i -th element in $\hat{\mathcal{E}}$ can be correlated to the i -th element in $\Delta\mathcal{E}$). With the assumption, the covariance matrices $\Sigma_{\hat{\mathcal{E}}, \hat{\mathcal{E}}}$, $\Sigma_{\hat{\mathcal{E}}, \Delta\mathcal{E}}$, $\Sigma_{\Delta\mathcal{E}, \hat{\mathcal{E}}}$ and $\Sigma_{\Delta\mathcal{E}, \Delta\mathcal{E}}$ become diagonal matrices. We further assume the distributions of $\hat{\mathcal{E}}$ and $\Delta\mathcal{E}$ to be isotropic (*i.e.*, $\Sigma = \sigma^2 I$), which significantly simplify the estimation of the joint distribution.

The Proposed D²-DPM

Now we provide the proposed dual denoising mechanism. We coin the proposed method ‘‘dual denoising’’ as it denoises two types of noise during the inverse diffusion process, including *quantization noise* and *diffusion noise*. Specifically, we propose two variants of dual denoising, named *stochastic* dual denoising (S-D²) and *deterministic* dual denoising (D-D²).

Stochastic Dual Denoising. In the stochastic variant of dual denoising, we recover the distribution of the diffusion noise by subtracting the estimated quantization noise $\Delta\mathcal{E}'$ from the quantized output $\hat{\mathcal{E}}$:

$$\mathcal{E}' = \hat{\mathcal{E}} - \Delta\mathcal{E}'. \quad (11)$$

Obviously \mathcal{E}' also follows a Gaussian distribution. If the estimated quantization noise $\Delta\mathcal{E}'$ accurately captures the real quantization noise $\Delta\mathcal{E}$, then the expectation and the covariance matrix of the recovered diffusion noise can be derived as follows:

$$\mathbf{E}[\mathcal{E}'] = \mathbf{E}[\mathcal{E}] + \mathbf{E}[\Delta\mathcal{E}] - \mathbf{E}[\Delta\mathcal{E}'] = \mathbf{E}[\mathcal{E}] \quad (12)$$

$$\begin{aligned} \mathbf{Var}[\mathcal{E}'] &= \mathbf{Var}[\mathcal{E}] + \mathbf{Var}[\Delta\mathcal{E}] + \mathbf{Var}[\Delta\mathcal{E}'] \\ &\quad + 2\mathbf{Cov}[\mathcal{E}, \Delta\mathcal{E}] - 2\mathbf{Cov}[\mathcal{E}, \Delta\mathcal{E}'] \\ &\quad - 2\mathbf{Cov}[\Delta\mathcal{E}, \Delta\mathcal{E}'] \\ &= \mathbf{Var}[\mathcal{E}] \end{aligned} \quad (13)$$

It can be seen that recovered diffusion noise follows the same distribution as that of the original diffusion noise. The sampling can be achieved by solving the following SDE:

$$\begin{aligned} dx &= \left[\mathbf{f}(x, t) + g(t)^2 \frac{\hat{\epsilon}_\theta(x_t, t) - \Delta\epsilon'(x_t, t)}{\sigma_t} \right] dt \\ &\quad + g(t) d\bar{w} \end{aligned} \quad (14)$$

Deterministic Dual Denoising In the deterministic variant of dual denoising, the distribution of the diffusion noise is recovered by subtracting the mean vector $\Delta\boldsymbol{\mu}$ of quantization noise from the quantized output $\hat{\mathcal{E}}$:

$$\mathcal{E}' = \hat{\mathcal{E}} - \Delta\boldsymbol{\mu}. \quad (15)$$

\mathcal{E}' again follows a Gaussian distribution. The expectation and the covariance matrix of the recovered diffusion noise can be derived as follows:

$$\mathbb{E}[\mathcal{E}'] = \mathbb{E}[\mathcal{E}] + \mathbb{E}[\Delta\mathcal{E}] - \mathbb{E}[\Delta\boldsymbol{\mu}] = \mathbb{E}[\mathcal{E}] \quad (16)$$

$$\begin{aligned} \text{Var}[\mathcal{E}'] &= \text{Var}[\mathcal{E}] + \text{Var}[\Delta\mathcal{E}] + \text{Var}[\Delta\boldsymbol{\mu}] \\ &\quad + 2\text{Cov}[\mathcal{E}, \Delta\mathcal{E}] - 2\text{Cov}[\mathcal{E}, \Delta\boldsymbol{\mu}] \\ &\quad - 2\text{Cov}[\Delta\mathcal{E}, \Delta\boldsymbol{\mu}] \\ &= \text{Var}[\mathcal{E}] + \sigma_{\Delta}^2 \mathbf{I} \end{aligned} \quad (17)$$

It can be seen that deterministic dual denoising introduces additional variance $\sigma_{\Delta}^2 \mathbf{I}$, which can be absorbed into diffusion term:

$$\begin{aligned} d\mathbf{x} &= \left[\mathbf{f}(\mathbf{x}, t) + g(t)^2 \frac{\hat{\epsilon}_{\theta}(\mathbf{x}_t, t) - \Delta\boldsymbol{\mu}}{\sigma_t} \right] dt \\ &\quad + \sqrt{g^2(t) - \frac{g^4(t)\sigma_{\Delta}^2(t)}{\sigma_t^2}} d\bar{\mathbf{w}} \end{aligned} \quad (18)$$

Algorithm 1 summarizes the procedure of the proposed dual denoising mechanism.

Experiments

Experiments Settings

Dataset and Metrics. We evaluated proposed D²-DPM using LDM (Rombach et al. 2022) across three standard datasets: ImageNet, LSUN-Bedrooms, and LSUN-Churches (Yu et al. 2015), each with a resolution of 256×256. To quantify generation performance, we employ metrics such as Fréchet Inception Distance (FID), Sliding Fréchet Inception Distance (sFID), Inception Score (IS), precision, and recall for comprehensive evaluation. For each evaluation, we generate 50,000 samples and calculate these metrics using the OpenAI’s evaluator (Dhariwal and Nichol 2021), with BOPs (Bit Operations) as the efficiency metric.

LDM settings. We primarily focus on the generative sampler parameters in LDM: classifier-free guidance *scale*, sampling *step* and variance schedule η . Since LDM employs the DDIM sampler, it degrades to an ODE-based sampler with zero stochasticity capacity when $\eta = 0$, becomes an SDE-based DDPM sampler with inherent stochasticity capacity when $\eta = 1$. Therefore, we simulate stochasticity capacity changes by adjusting the *scale*. In class-conditional generation, we set four parameter configurations: $\{scale = 3.0, \eta = 0.0|1.0, steps = 20\}$ and $\{scale = 1.5, \eta = 0.0|1.0, steps = 250\}$. For unconditional generation, we set two parameter configurations: $\{\eta = 0.0|1.0, steps = 200\}$.

Quantization Settings. We employ BRECQ (Li et al. 2021) as the PTQ baseline for extensive comparative experiments

Algorithm 1: The Methodological Framework of D²-DPM

Description: S : the number of sampling times; M : the number of sampling steps; TSQNM: Time Step-aware Quantization Noise Modeling. $\{\boldsymbol{\Sigma}_t = \sigma_t^2 \mathbf{I}\}_{t=1}^T$ and $\{\alpha_t\}_{t=1}^T$ are parameters within the DDIM sampler.

Input: Full precision model $model_{fp}$, quantization parameters q_{params} , sampling inputs $\{\mathbf{x}_T^i\}_{i=1}^N \sim \mathcal{N}(0, \mathbf{I})$

Output: Generated samples $\{\mathbf{x}_0^i\}_{i=1}^N$

```

1:  $\{(\mathbf{x}_t, t, c)\}_{i=1}^{M \times T} = \text{collect\_calibration}(model_{fp})$ 
2:  $model_q = \text{BRECQ}(model_{fp}, q_{params}, \{(\mathbf{x}_t, t, c)\}_{i=1}^{M \times T})$ 
3:  $\{(\hat{\mathcal{E}}, \Delta\mathcal{E})^i\}_{i=1}^{S \times T} = \text{collect\_quant\_error}(model_{fp}, model_q)$ 
4:  $\boldsymbol{\mu}_{\Delta\mathcal{E}}, \boldsymbol{\Sigma}_{\Delta\mathcal{E}} = \text{gaussian\_modeling}(\{(\hat{\mathcal{E}}, \Delta\mathcal{E})^i\}_{i=1}^{S \times T})$ 
5: for sample_num  $i$  in  $1, \dots, N$  do
6:    $\mathbf{x}_T = \mathbf{x}_T^i$ 
7:   for timestep  $t$  in  $T, \dots, 1$  do
8:      $\hat{\epsilon}_{\theta}^{(t)} = model_q(\mathbf{x}_t)$ 
9:      $\boldsymbol{\mu}_{\Delta\mathcal{E}|\hat{\epsilon}=\hat{\epsilon}_{\theta}^{(t)}}, \boldsymbol{\Sigma}_{\Delta\mathcal{E}|\hat{\epsilon}=\hat{\epsilon}_{\theta}^{(t)}} = \text{TSQNM}(\hat{\epsilon}_{\theta}^{(t)}, \boldsymbol{\mu}[t], \boldsymbol{\Sigma}[t])$ 
10:    if stochastic dual denoising then
11:       $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$ 
12:       $\Delta\mathcal{E}' = \boldsymbol{\mu}_{\Delta\mathcal{E}|\hat{\epsilon}=\hat{\epsilon}_{\theta}^{(t)}} + \boldsymbol{\Sigma}_{\Delta\mathcal{E}|\hat{\epsilon}=\hat{\epsilon}_{\theta}^{(t)}}^{1/2} \cdot \mathbf{z}$ 
13:       $\mathcal{E}' \leftarrow \hat{\epsilon}_{\theta}^{(t)} - \Delta\mathcal{E}'$ 
14:    else
15:       $\mathcal{E}' \leftarrow \hat{\epsilon}_{\theta}^{(t)} - \boldsymbol{\mu}_{\Delta\mathcal{E}|\hat{\epsilon}=\hat{\epsilon}_{\theta}^{(t)}}$ 
16:       $k = \sqrt{1 - \alpha_{t-1} - |\boldsymbol{\Sigma}_t|^{1/d}} - \sqrt{\frac{\alpha_{t-1}(1-\alpha_t)}{\alpha_t}}$ 
17:       $\boldsymbol{\Sigma}_t \leftarrow \boldsymbol{\Sigma}_t - k^2 \cdot \boldsymbol{\Sigma}_{\Delta\mathcal{E}|\hat{\epsilon}=\hat{\epsilon}_{\theta}^{(t)}}$ 
18:    end if
19:     $\boldsymbol{\epsilon}_t \sim \mathcal{N}(0, \mathbf{I})$ 
20:     $\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \left( \frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \mathcal{E}'}{\sqrt{\alpha_t}} \right)$ 
21:     $\quad + \sqrt{1 - \alpha_{t-1} - |\boldsymbol{\Sigma}_t|^{1/d}} \mathcal{E}' + \boldsymbol{\Sigma}_t^{1/2} \cdot \boldsymbol{\epsilon}_t$ 
22:   $\mathbf{x}_0^i = \mathbf{x}_0$ 
23: end for
24: return  $\{\mathbf{x}_0^i\}_{i=1}^N$ 

```

and implement an LDM-compatible version of Qrop (Wei et al. 2022). To ensure comparability, we keep all settings aligned with PTQD, specifically: 1) using Adaround (Nagel et al. 2020) as the weight quantizer; and 2) fixing the first and last layers to 8 bits, while quantizing other layers to the target bit-width. For calibration, we collect the diffusion model’s inputs at each sampling timestep as the calibration set. Notationally, WxAy indicates that weights and activations are quantized to x and y bits. In all experiments, we adopt two quantization configurations: W8A8 and W4A8.

Class-conditional Generation

We first compare the proposed D²-DPM with other works on class-conditional generation tasks. We conduct experiments using LDM-4 on ImageNet 256×256. The results for the configuration $\{scale = 3.0, \eta = 0.0|1.0, steps = 20\}$ are presented in Table 1. Regarding efficiency, W8A8 and W4A8 quantization achieve volume compression ratios of 3.99x and 7.95x, while reducing BOPs by 11.67x and 23.33x. In

Model	Method	Bits (W/A)	Size (MB)	BOPs (T)	IS \uparrow	FID \downarrow	sFID \downarrow	Precision \uparrow	Recall \uparrow
LDM-4 ($\eta = 0.0$)	FP	32/32	1742.72	102.20	366.03	11.13	7.834	93.93%	27.98%
	PTQ4DM	8/8	436.79	8.76	324.21	9.37	9.87	87.15%	31.77%
	Q-diffusion	8/8	436.79	8.76	327.16	8.72	10.46	86.91%	33.26%
	PTQD	8/8	436.79	8.76	324.64	8.46	10.12	87.68%	34.64%
	Ours _{S-D²}	8/8	436.79	8.76	332.55	8.11	8.02	87.35%	36.55%
	Ours _{D-D²}	8/8	436.79	8.76	333.89	8.12	7.92	88.56%	36.69%
	PTQ4DM	4/8	219.12	4.38	336.28	10.45	13.94	90.61%	28.63%
	Q-diffusion	4/8	219.12	4.38	347.52	11.13	9.07	90.89%	29.39%
	PTQD	4/8	219.12	4.38	355.10	10.41	8.45	92.13%	27.54%
	Ours _{S-D²}	4/8	219.12	4.38	358.14	9.75	6.60	92.25%	30.21%
Ours _{D-D²}	4/8	219.12	4.38	357.57	9.71	6.65	92.22%	30.21%	
LDM-4 ($\eta = 1.0$)	FP	32/32	1742.72	102.20	361.84	13.83	20.56	92.22%	19.58%
	PTQ4DM	8/8	436.79	8.76	332.18	12.24	18.63	87.21%	23.60%
	Q-diffusion	8/8	436.79	8.76	335.61	11.07	16.15	88.50%	24.93%
	PTQD	8/8	436.79	8.76	335.70	10.86	15.02	88.44%	25.24%
	Ours _{S-D²}	8/8	436.79	8.76	342.71	10.57	14.81	88.58%	26.02%
	Ours _{D-D²}	8/8	436.79	8.76	343.68	10.58	14.72	88.90%	26.03%
	PTQ4DM	4/8	219.12	4.38	340.10	13.68	22.05	89.50%	20.01%
	Q-diffusion	4/8	219.12	4.38	349.89	14.22	20.17	89.93%	20.57%
	PTQD	4/8	219.12	4.38	353.57	13.15	17.41	91.09%	20.42%
	Ours _{S-D²}	4/8	219.12	4.38	355.20	12.60	15.80	91.26%	20.94%
Ours _{D-D²}	4/8	219.12	4.38	356.39	12.65	15.47	91.44%	21.36%	

Table 1: Performance comparison of class-conditioned generation on ImageNet 256×256 using LDM-4 (scale=3.0, step=20).

terms of generation quality, S-D² and D-D² demonstrates superior performance across various sampling stochasticity capacities and quantization bit-width settings. Specifically, S-D² outperforms other works across all metrics, with FID scores in the best-case scenario up to 0.66 lower than PTQD, 1.26 lower than PTQ4DM, and 3.24 lower than the full-precision model. This initially demonstrates that the proposed quantization noise model precisely captures quantization noise during inference. Subsequently, S-D² effectively restores the distribution by implicitly correcting the standard deviation while correcting the mean. Additionally, its inherent stochasticity accumulates positive effects over prolonged iterations, steering the data toward a more optimal distribution. Moreover, the superior metrics of D-D² reaffirm the efficacy of noise modeling and distribution correction. When the stochastic capacity $g(t)$ is sufficiently large ($\eta = 1.0$), FID and sFID decrease by 1.06 and 5.47 on average compared to the full-precision model. This indicates that additional standard deviation is effectively utilized by the stochastic term without causing detrimental variance overflow. However, data shows that even when $g(t)$ is too minimal to absorb additional standard deviation, the performance remains superior. We attribute this to the limited additional variance from quantization noise effectively compensating for the stochastic term, creating a superior Langevin SDE over the original SDE (ODE), as discussed in preliminary work. This constructs a larger error buffer, smoothing out the sharp noise introduced at each step. Therefore, we suggest using the D-D² optimized low-stochasticity-capacity sam-

pler when the quantization bit-width is not too low, meaning the variance from quantization noise is limited, as it can effectively leverage the beneficial components of the noise.

To further validate the performance of our D²-DPM, we conduct two sets of high-density step generation experiments under the conditions $\{scale = 1.5, \eta = 0.0|1.0, steps=250\}$, with the results shown in Table 2. Evidently, our metrics consistently surpass PTQD, which demonstrated strong competitive advantages in earlier experiments.

Unconditional Generation

We evaluate D²-DPM on unconditional generation tasks, employing LDM-4 and LDM-8 models across the LSUN-Bedroom and LSUN-Church datasets, respectively. Table 3 and 4 show that our approach narrows the gap with full-precision model. Specifically, on LSUN-Bedroom dataset, S-D² reduces FID and sFID by 1.39 and 0.39 on average, compared to PTQD. Similarly, D-D² reduces FID and sFID by 1.73 and 0.60 on average. On LSUN-Church dataset, S-D² reduces the average FID and sFID by 1.14 and 0.37. In parallel, D-D² also achieves significant reductions, lowering FID and sFID by an average of 1.13 and 0.39. This demonstrates that our precise quantization noise modeling, along with the decoupled mean and standard deviation corrections in D²-DPM, more effectively restores the distribution. Finally, we observe a phenomenon consistent with previous findings: even when $g(t)$ is minimal, D-D² still shows superior performance, even partially surpassing S-D². This confirms the effectiveness of its deterministic mean correction

Model	Method	Bits (W/A)	Size (MB)	BOPs (T)	IS \uparrow	FID \downarrow	sFID \downarrow	Precision \uparrow	Recall \uparrow
LDM-4 ($\eta = 0.0$)	FP	32/32	1742.72	102.20	213.74	3.32	5.23	83.04%	53.31%
	PTQD	4/8	219.12	4.38	162.77	6.46	10.14	73.88%	58.10%
	Ours _{D-D²}	4/8	219.12	4.38	169.03	5.56	9.45	75.20%	58.14%
LDM-4 ($\eta = 1.0$)	FP	32/32	1742.72	102.20	250.97	3.54	5.07	87.10%	49.12%
	PTQD	4/8	219.12	4.38	153.01	7.90	7.87	71.75%	55.15%
	Ours _{S-D²}	4/8	219.12	4.38	171.49	6.91	7.49	72.82%	55.64%

Table 2: Performance comparisons of class-conditional generation on ImageNet256 \times 256 using LDM-4 (scale=1.5, step=250).

LDM-4 (steps=200, $\eta = 1.0$)					
Method	W/A	FID \downarrow	sFID \downarrow	Prec. \uparrow	Rec. \uparrow
FP	32/32	3.03	7.03	64.65%	47.60%
PTQD	8/8	9.16	12.94	51.99%	44.32%
Ours _{S-D²}	8/8	7.69	12.61	54.81%	45.03%
Ours _{D-D²}	8/8	7.55	12.56	55.60%	45.80%
PTQD	4/8	12.57	16.04	51.31%	42.40%
Ours _{S-D²}	4/8	11.26	15.60	51.45%	43.66%
Ours _{D-D²}	4/8	10.72	15.23	51.44%	43.90%

Table 3: Performance comparisons of unconditional image generation on LSUN-Bedroom 256 \times 256.

and further validates our earlier hypothesis: the additional variance from quantization noise effectively compensates for the stochastic term, thereby implicitly transforming the sampling equation into the original ODE with an enhanced Langevin diffusion term. The ODE aligns the marginal distribution $p_t(x)$, while the improved Langevin diffusion term better buffers against sharp noise from the quantized diffusion model during each iteration of noise estimation.

LDM-8 (steps = 200, $\eta = 0.0$)					
Method	W/A	FID \downarrow	sFID \downarrow	Prec. \uparrow	Rec. \uparrow
FP	32/32	4.17	12.91	66.00%	51.46%
PTQD	8/8	8.31	12.97	56.57%	54.15%
Ours _{S-D²}	8/8	7.82	12.52	56.75%	54.45%
Ours _{D-D²}	8/8	7.83	12.51	56.86%	54.54%
PTQD	4/8	12.96	15.42	50.23%	52.80%
Ours _{S-D²}	4/8	11.18	15.14	52.27%	53.78%
Ours _{D-D²}	4/8	11.18	15.11	52.15%	53.68%

Table 4: Performance comparisons of unconditional image generation on LSUN-Church 256 \times 256.

Ablation Study

As shown in Table 5, we perform ablation studies on the denoising components of dual denoising mechanisms, S-D² and D-D². Stochastic Joint Correction (SJC), which implicitly corrects the variance while correcting the mean using estimated noise, corresponds to S-D², while Deterministic

Method	W/A	IS \uparrow	FID \downarrow	sFID \downarrow
FP	32/32	250.97	3.54	5.07
PTQD	4/8	153.01	7.90	7.87
+ SJC (S-D ²)	4/8	171.49	6.91	7.49
+ DMC	4/8	159.30	7.14	7.67
+ DMC + DVC (D-D ²)	4/8	172.13	6.81	7.42

Table 5: Ablation study of denoising components using LDM-4 (scale=1.5, $\eta=1.0$, step=250) on ImageNet 256 \times 256.

Mean Correction (DMC) and Deterministic Variance Correction (DVC) are the key components of D-D². By applying SJC, we achieve FID and sFID reductions of 0.99 and 0.38 compared to PTQD, showing that S-D² successfully performs joint correction for more accurate distribution restoration. In D-D², the use of DMC alone reduces FID and sFID by 0.76 and 0.20, respectively, showing that we accurately estimate the conditional mean of quantization noise through joint distribution. Building on this, applying DVC further reduces FID and sFID by 0.33 and 0.25, indicating that the additional variance is also effectively absorbed by the stochastic term. The above experiments show that our D²-DPM more effectively mitigates the adverse effects of quantization noise and more precisely corrects distributions.

Conclusion and Future Work

In this paper, we propose a dual denoising paradigm to eliminate the residual quantization noise in quantized diffusion models. We first establish the joint distribution of quantized outputs and noise, allowing us to instantiate the conditional distribution of quantization noise during inference. We then design two variants, S-D² and D-D², to decouple and correct the mean and standard deviation shifts introduced by quantization noise. Extensive experiments demonstrate that our approach effectively corrects the distribution, achieving high-fidelity quantization of diffusion models.

In essence, this method provides technical support for high-fidelity, efficient compression of diffusion models aimed at reducing carbon emissions, and is therefore not limited by task type. It can be extended to various domains, including video generation, text modeling, and molecular design. Future work will focus on expanding this paradigm across multiple domains and pursuing a unified framework.

Acknowledgments

This work is supported by Zhejiang Province High-Level Talents Special Support Program "Leading Talent of Technological Innovation of Ten-Thousands Talents Program" (No. 2022R52046) and Alibaba-Zhejiang University Joint Research Institute of Frontier Technologies.

References

- Aggarwal, A.; Mittal, M.; and Battineni, G. 2021. Generative adversarial network: An overview of theory and applications. *International Journal of Information Management Data Insights*, 1(1): 100004.
- Anderson, B. D. 1982. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3): 313–326.
- Bao, F.; Li, C.; Zhu, J.; and Zhang, B. 2022. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. *arXiv preprint arXiv:2201.06503*.
- Berthelot, D.; Autef, A.; Lin, J.; Yap, D. A.; Zhai, S.; Hu, S.; Zheng, D.; Talbott, W.; and Gu, E. 2023. Tract: Denoising diffusion models with transitive closure time-distillation. *arXiv preprint arXiv:2303.04248*.
- Chu, X.; Li, L.; and Zhang, B. 2024. Make repvgg greater again: A quantization-aware approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 11624–11632.
- Chung, H.; Sim, B.; and Ye, J. C. 2022. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12413–12422.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Dockhorn, T.; Vahdat, A.; and Kreis, K. 2021. Score-based generative modeling with critically-damped langevin diffusion. *arXiv preprint arXiv:2112.07068*.
- Fan, Y.; Wei, X.; Gong, R.; Ma, Y.; Zhang, X.; Zhang, Q.; and Liu, X. 2024. Selective Focus: Investigating Semantics Sensitivity in Post-training Quantization for Lane Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 11936–11943.
- Fang, G.; Ma, X.; and Wang, X. 2024. Structural pruning for diffusion models. *Advances in neural information processing systems*, 36.
- Franzese, G.; Rossi, S.; Yang, L.; Finamore, A.; Rossi, D.; Filippone, M.; and Michiardi, P. 2023. How much is enough? a study on diffusion times in score-based generative models. *Entropy*, 25(4): 633.
- He, Y.; Liu, J.; Wu, W.; Zhou, H.; and Zhuang, B. 2023. Efficientdm: Efficient quantization-aware fine-tuning of low-bit diffusion models. *arXiv preprint arXiv:2310.03270*.
- He, Y.; Liu, L.; Liu, J.; Wu, W.; Zhou, H.; and Zhuang, B. 2024. Ptdq: Accurate post-training quantization for diffusion models. *Advances in Neural Information Processing Systems*, 36.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Kingma, D.; Salimans, T.; Poole, B.; and Ho, J. 2021. Variational diffusion models. *Advances in neural information processing systems*, 34: 21696–21707.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kong, Z.; and Ping, W. 2021. On fast sampling of diffusion probabilistic models. *arXiv preprint arXiv:2106.00132*.
- Lee, J. S.; Kim, J.; and Kim, P. M. 2022. ProteinSGM: Score-based generative modeling for de novo protein design. *bioRxiv*, 2022–07.
- Li, X.; Liu, Y.; Lian, L.; Yang, H.; Dong, Z.; Kang, D.; Zhang, S.; and Keutzer, K. 2023. Q-diffusion: Quantizing diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17535–17545.
- Li, Y.; Gong, R.; Tan, X.; Yang, Y.; Hu, P.; Zhang, Q.; Yu, F.; Wang, W.; and Gu, S. 2021. Brecq: Pushing the limit of post-training quantization by block reconstruction. *arXiv preprint arXiv:2102.05426*.
- Lin, J.; Tang, J.; Tang, H.; Yang, S.; Chen, W.-M.; Wang, W.-C.; Xiao, G.; Dang, X.; Gan, C.; and Han, S. 2023. Awq: Activation-aware weight quantization for llm compression and acceleration. *arXiv preprint arXiv:2306.00978*.
- Liu, L.; Ren, Y.; Lin, Z.; and Zhao, Z. 2022. Pseudo numerical methods for diffusion models on manifolds. *arXiv preprint arXiv:2202.09778*.
- Liu, Z.; Wang, Y.; Han, K.; Zhang, W.; Ma, S.; and Gao, W. 2021. Post-training quantization for vision transformer. *Advances in Neural Information Processing Systems*, 34: 28092–28103.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35: 5775–5787.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11461–11471.
- Meng, C.; Rombach, R.; Gao, R.; Kingma, D.; Ermon, S.; Ho, J.; and Salimans, T. 2023. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14297–14306.
- Nagel, M.; Amjad, R. A.; Van Baalen, M.; Louizos, C.; and Blankevoort, T. 2020. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, 7197–7206. PMLR.
- Nagel, M.; Fournarakis, M.; Amjad, R. A.; Bondarenko, Y.; Van Baalen, M.; and Blankevoort, T. 2021. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*.
- Nagel, M.; Fournarakis, M.; Bondarenko, Y.; and Blankevoort, T. 2022. Overcoming oscillations in

- quantization-aware training. In *International Conference on Machine Learning*, 16318–16330. PMLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Salimans, T.; and Ho, J. 2022. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*.
- Shang, Y.; Yuan, Z.; Xie, B.; Wu, B.; and Yan, Y. 2023. Post-training quantization on diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1972–1981.
- Singer, U.; Polyak, A.; Hayes, T.; Yin, X.; An, J.; Zhang, S.; Hu, Q.; Yang, H.; Ashual, O.; Gafni, O.; et al. 2022. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. PMLR.
- Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2023. Consistency models. *arXiv preprint arXiv:2303.01469*.
- Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Wang, Y.; Yang, W.; Chen, X.; Wang, Y.; Guo, L.; Chau, L.-P.; Liu, Z.; Qiao, Y.; Kot, A. C.; and Wen, B. 2023. SinSR: Diffusion-Based Image Super-Resolution in a Single Step. *arXiv preprint arXiv:2311.14760*.
- Wei, X.; Gong, R.; Li, Y.; Liu, X.; and Yu, F. 2022. Qdrop: Randomly dropping quantization for extremely low-bit post-training quantization. *arXiv preprint arXiv:2203.05740*.
- Xiao, G.; Lin, J.; Seznec, M.; Wu, H.; Demouth, J.; and Han, S. 2023. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, 38087–38099. PMLR.
- Yao, Z.; Yazdani Aminabadi, R.; Zhang, M.; Wu, X.; Li, C.; and He, Y. 2022. Zeroquant: Efficient and affordable post-training quantization for large-scale transformers. *Advances in Neural Information Processing Systems*, 35: 27168–27183.
- Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.
- Zhang, Y.; Huang, N.; Tang, F.; Huang, H.; Ma, C.; Dong, W.; and Xu, C. 2023. Inversion-based style transfer with diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10146–10156.
- Zhu, Y.; Li, Z.; Wang, T.; He, M.; and Yao, C. 2023. Conditional Text Image Generation With Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14235–14245.