Post-training Quantization with Progressive Calibration and Activation Relaxing for Text-to-Image Diffusion Models

Siao Tang^{1,3}, Xin Wang² \square \boxtimes , Hong Chen², Chaoyu Guan^{2,3}, Zewen Wu², Yansong Tang¹, and Wenwu Zhu² \square \boxtimes

¹ TBSI, Shenzhen International Graduate School, Tsinghua University
² Department of Computer Science and Technology, BNRIST, Tsinghua University ³ Tsingmao Intelligence {tsa22, h-chen20, guancy19, wuzw21}@mails.tsinghua.edu.cn tang.yansong@sz.tsinghua.edu.cn {xin_wang, wwzhu}@tsinghua.edu.cn

Abstract. High computational overhead is a troublesome problem for diffusion models. Recent studies have leveraged post-training quantization (PTQ) to compress diffusion models. However, most of them only focus on unconditional models, leaving the quantization of widely-used pretrained text-to-image models, e.g., Stable Diffusion, largely unexplored. In this paper, we propose a novel post-training quantization method PCR (Progressive Calibration and Relaxing) for text-to-image diffusion models, which consists of a progressive calibration strategy that considers the accumulated quantization error across timesteps, and an activation relaxing strategy that improves the performance with negligible cost. Additionally, we demonstrate the previous metrics for text-to-image diffusion model quantization are not accurate due to the distribution gap. To tackle the problem, we propose a novel QDiffBench benchmark, which utilizes data in the same domain for more accurate evaluation. Besides, QDiffBench also considers the generalization performance of the quantized model outside the calibration dataset. Extensive experiments on Stable Diffusion and Stable Diffusion XL demonstrate the superiority of our method and benchmark. Moreover, we are the first to achieve quantization for Stable Diffusion XL while maintaining the performance.

 ${\bf Keywords:} \ {\rm Diffusion} \ {\rm Models} \cdot {\rm Quantization} \cdot {\rm Text-to-image} \ {\rm Generation}$

1 Introduction

Diffusion models have demonstrated remarkable ability in various generative tasks [4, 11, 23, 24, 27, 28, 32], especially in text-to-image generation [4, 23, 24]. However, one critical limitation of diffusion models is the high computational cost, which mainly results from the following two aspects. On the one hand,

 $[\]bowtie$ denotes corresponding authors.

diffusion models typically require multiple denoising steps to generate images, which is time-consuming. On the other hand, the large network architecture of diffusion models usually requires a great deal of time and memory, especially for foundation models pretrained on large-scale datasets, e.g., Stable Diffusion [23] and Stable Diffusion XL [22]. In this paper, we focus on the latter problem, i.e., compressing the diffusion models to reduce the computational cost.

Recent studies have explored quantization methods to compress diffusion models [8,13,25,26,30], which mainly adopt Post-training Quantization (PTQ) [14, 18]. PTQ has no need for retraining or finetuning the network, making it more resource-friendly than Quantization-aware Training (QAT) [20,29]. Nevertheless, most of these previous works focus on the unconditional diffusion models, leaving the quantization of the widely used large pretrained text-to-image models, e.g., Stable Diffusion, largely unexplored. Only Q-diffusion [13] presents several quantitative quantization results for the text-to-image diffusion model. However, the results are limited because of the adopted inaccurate metrics without considering the data distribution-gap problem.

In this paper, we investigate the largely unexplored text-to-image diffusion model quantization problem, from both method and benchmark perspectives.

On the one hand, about the method, all previous works use the full-precision model to obtain the calibration data for each timestep, and thus ignore the accumulated quantization error of previous denoising steps, which may result in inaccurate quantization. Additionally, they neglect the sensitivity of image fidelity or text-image matching to different denoising steps, which may largely influence the quantization performance. On the other hand, about the metrics, previous works use a prompt subset of the COCO [15] dataset to perform the calibration and use the FID metric calculated with COCO images to evaluate the model. However, this metric leads to inaccurate evaluation because there exists a distribution gap between the generated images by the text-to-image diffusion model (pretrained on large-scale datasets) and the COCO images. Additionally, previous metrics only test on the COCO prompts, ignoring the prompt-generalization ability to unseen prompts outside the calibration dataset, which is important to the pretrained text-to-image diffusion models.

To tackle the problems, we propose a novel quantization method PCR (Progressive Calibration and Relaxing) and the QDiffBench benchmark for quantizing text-to-image diffusion models. Specifically, the proposed PCR method consists of a progressive calibration strategy and an activation relaxing strategy. The progressive calibration progressively quantizes each step with all the previous steps quantized, which can be aware of the accumulated quantization error across steps. The activation relaxing strategy relaxes a minor part of critical sensitive timesteps using higher bitwidth to improve image fidelity or text-image matching, with negligible additional cost.

As for benchmark, QDiffBench consists of an accurate FID calculation strategy and a prompt-generalization evaluation strategy. To deal with the data distribution-gap problem, QDiffBench calculates the FID [10] score between the images generated by the full-precision model and the quantized model. This strategy uses the data in the same domain for evaluation, making it more accurate in evaluating the degradation caused by quantization. Moreover, QDiffBench evaluates the generalization performance for the quantized text-to-image diffusion model on the prompts which have quite different styles from the calibration prompts. We emphasize the significance of QDiffBench, as the quantization for text-to-image diffusion models could be blocked without an effective benchmark.

To summarize, our contributions are listed as follows:

- 1. We propose a novel PTQ quantization method **PCR** for text-to-image diffusion models, which consists of a progressive calibration strategy and an activation relaxing strategy.
- 2. We propose a comprehensive and effective benchmark **QDiffBench**, and to our knowledge, it is the first effective benchmark to evaluate quantized text-to-image diffusion models.
- 3. Extensive experiments on foundation diffusion models (i.e., Stable Diffusion and Stable Diffusion XL) demonstrate the superiority of our proposed method and benchmark.
- 4. We are the first to achieve quantization on Stable Diffusion XL (3.5B parameters) which is one of the biggest diffusion models.

2 Related Work

2.1 Model Quantization

Model quantization [19] is one of the most significant techniques of model compression. Quantization methods compress neural networks by compressing the model weights and the activations to lower bits, e.g., converting 32/16-bit float weights/activations into formats with INT4 or INT8 bit-width. The quantization process can be formulated as:

$$w_q = \operatorname{clip}\left(\operatorname{round}\left(\frac{w}{s}\right) + z, q_{\min}, q_{\max}\right),$$

where s is the scaling factor and z is the zero-point.

Post-training Quantization. Quantization methods are usually divided into two classes, i.e., post-training quantization (PTQ) [14,18] and quantization-aware training (QAT) [20,29]. QAT methods involve training/finetuning, usually requiring massive training data and costing many computation resources. In contrast, PTQ methods only require a small subset of data (e.g., 128 images) for calibration, and have no need for training, making them more resource-friendly and faster than QAT methods.

Mixed-precision Quantization. Mixed-precision quantization [6,33] is a strategy that allocates different bitwidth to different layers/blocks, which is considered a trend [12,31] because it can further improve the model performance. Mixed-precision quantization can be applied in practical scenes, as more and more hardwares support it, such as Qualcomm Snapdragon 8 Gen 2 [1], NVIDIA A100 [5], and RTX 4090 [2]. Our activation relaxation strategy can also be regarded as a mixed-precision quantization in the dimension of timestep.

2.2 Quantization for Diffusion Model

Quantization methods have been used to compress and accelerate diffusion models [8, 13, 25, 26, 30]. For example, Shang et al. propose PTQ4DM [25] which collects calibration data in the denoising sampling process rather than simulating the training process. They also propose that the timesteps of calibration data should be generated from a skew-normal distribution. Q-diffusion [13] proposes to collect calibration data at certain timestep intervals and split quantization for shortcut layers whose input is a concatenated skip-connection feature. However, all the existing quantization methods ignore the accumulated error across sampling steps. Additionally, none of them tailor the design for the setting of text-to-image generation.

3 Method

In this section, we present our method **PCR**, which consists of a progressive calibration strategy that considers the accumulated quantization error across timesteps, and an activation relaxing strategy which can effectively improve performance with negligible costs.

3.1 Preliminaries for Diffusion Models

Diffusion models define a forward process that adds the noise to the real data \mathbf{x}_0 :

$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}, \ \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}),$$

where α_t and σ_t are functions of t and the signal-to-noise-ratio α_t/σ_t is strictly decreasing. Diffusion models train a neural network ϵ_{θ} to predict the noise of a noisy variable \mathbf{x}_t , whose training objective is:

$$\mathbf{E}_{\mathbf{x}_{0},\boldsymbol{\epsilon},t}[\|\boldsymbol{\epsilon}_{\theta}\left(\alpha_{t}\mathbf{x}_{0}+\sigma_{t}\boldsymbol{\epsilon},t\right)-\boldsymbol{\epsilon}\|_{2}^{2}],$$

In the reverse process, diffusion models generate samples by gradually denoising from a Gaussian noise $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ to \mathbf{x}_0 , along certain trajectories determined by sampling strategies [11, 16, 17].

3.2 Time-Accumulated Error Aware Progressive Calibration

Substantial studies [13, 25, 26, 30] point out that the distributions of activation obviously vary over different time steps, which is a challenge for quantizing diffusion models. To tackle this problem, some studies [26, 30] adopt time-aware quantization, i.e., separately quantizing the activations for each timestep, which has been shown to be effective. We also adopt the time-aware strategy.

5

Previous works use the full-precision model to obtain the calibration data for each timestep during the sampling process. However, they are not aware of the accumulated quantization error across timesteps in the iterative denoising process. The quantization errors in previous sampling steps will shift the distribution of the activations in subsequent steps, as shown in Fig. 2, which makes calibrating with the full-precision model sub-optimal. To tackle this problem, we propose a progressive calibration strategy for activations, which quantizes the activations at step t with all the previous steps t + 1, ..., T quantized, illustrated in Fig. 1.



Fig. 1: Progressive calibration for activations. The yellow bar means the activation has been quantized. We quantize the activations at step t with all the previous steps quantized, which considers the accumulated quantization error across timesteps.



Fig. 2: The distribution shift caused by the accumulated quantization error. We gather the statistics of sub feature maps at timestep 1, 10, and 20.

We further mathematically prove the necessity of considering the accumulated quantization errors, through Theorem 1 and its Corollary 1. The detailed proof of Theorem 1 is shown in Sec. A.1 (Appendix).

Notations. Let ϵ_{θ} be the denoising net and \mathbf{x}_t be the intermediate variable at timestep t. Let the mark $\hat{.}$ denote the quantized term, and $\hat{\epsilon}_{\theta}(\hat{\mathbf{x}}_t, t) = \epsilon_{\theta}(\hat{\mathbf{x}}_t, t) + \Delta_t$, where $\epsilon_{\theta}(\hat{\mathbf{x}}_t, t)$ is the predicted noise with quantized input but full-precision network, and $\hat{\epsilon}_{\theta}(\hat{\mathbf{x}}_t, t)$ is the predicted noise with both quantized input and network.

Theorem 1. Through the multi-step denoising process, the upper bound of the ultimate quantization error $\delta = \|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|$ can be approximated as the linear combination of $\{\|\Delta_t\|\}, t = 1, ..., T$.

Corollary 1. To minimize the upper bound of the final quantization error $\delta = \|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|$, we should minimize the quantization error $\|\Delta_t\|$ at each timestep t, t = 1, ..., T.

Note that, $\|\Delta_t\| = \|\hat{\boldsymbol{\epsilon}}_{\theta}(\hat{\mathbf{x}}_t, t) - \boldsymbol{\epsilon}_{\theta}(\hat{\mathbf{x}}_t, t)\|$ is related to the quantized input $\hat{\mathbf{x}}_t$ rather than the full-precision input \mathbf{x}_t . Hence, to minimize $\|\Delta_t\|$, we should use the data generated with all its previous steps quantized as the calibration data.

3.3 Time-wise Activation Relaxing



Fig. 3: Add perturbations to the predicted noise at timesteps in the interval t=[a,b]. (b) shows perturbations at steps near x_T make the image fail to match the prompt. (c) and (d) shows perturbations at the middle steps cause minor semantic changes such as color and computer pose. (e) shows the fidelity degradation brought by perturbations at steps near x_0 .

Firstly, we clarify the definitions of "text-image matching" and "image fidelity". Text-image matching means whether the semantics of the image matches its prompt, which is the high-level information. In contrast, image fidelity denotes whether the image has some low-level losses such as distortions and noises.

We observe the image fidelity and text-image matching of text-to-image diffusion models are sensitive to activation quantization. To identify this sensitivity with respect to timesteps, we manually add a small Gaussian perturbation to the predicted noise at certain timesteps, simulating quantization errors. Fig. 3 shows the visual results, from which we can draw the conclusion that: Image fidelity is sensitive to the timesteps near x_0 , while text-image matching is sensitive to the steps near x_T . We refer to this phenomenon as sensitivity discrepancy.

To further identify actual patterns of sensitivity discrepancy in practical quantization scenes, we first quantize a model to low-bit and then set certain timesteps to full-precision. According to the results in Fig. 4, we observe that Stable Diffusion is sensitive to image fidelity degradation (controlled by steps near x_0) rather than text-image matching degradation (controlled by steps near x_T), while Stable Diffusion XL is the opposite.



Fig. 4: Quantitative results of setting certain steps to full-precision. CLIP score measures the quality of text-image matching. FID to FP32 (presented in Sec. 4.1) measures image fidelity when text-image matching is good, but it can be quite poor when there is a bad text-image matching.

However, previous works neglect this sensitivity discrepancy, and allocate the same bitwidth to all timesteps, resulting in a sub-optimal solution. To tackle this issue, we propose to relax a minor part of timesteps using higher bitwidth to improve image fidelity or text-image matching. Let $b_t, t = 0, ..., T$ denote the bit widths for activation quantization at step t. For the quantized models which are sensitive to image fidelity degradation (e.g., Stable Diffusion), we relax the m steps near x_0 by setting $b_i, i = 0, ..., m - 1$ to a higher bitwidth such as 10-bit. For the models that are sensitive to text-image matching (e.g., Stable Diffusion XL), we set $b_i, i = T - m + 1, ..., T$ to a higher bit. The relaxation proportion $\tau = m/T$ conforms to the law of diminishing marginal utility (Sec. 5.3), so we propose to set τ no more than 0.20 for the trade-off between performance and extra computation. To identify the sensitivity of a given model to image fidelity or text-matching, users can perform simple experiments as Fig. 4, which only takes several minutes since we adopt PTQ.

Fig. 5 shows previous methods lose the ability to match the textual semantics in Stable Diffusion XL, while our method with relaxation maintains it well.

Although the relaxation remarkably improves the performance, it only requires negligible costs, since we only relax a small proportion of timesteps that are the most critical. For instance, in 8-bit quantization, relaxing 5% steps to 10 bits while keeping others the 8-bit actually results in the average 8.1 bit.

4 QDiffBench

Previous works rarely test their methods on text-to-image generation tasks. Only Q-diffusion [13] reports few quantitative results of Stable Diffusion. However, we observe that there exist some problems in the way of calculating FID used in Q-diffusion, which is caused by distribution-gap. Moreover, they ignore the prompt-generalization ability of the foundation text-to-image model. In this section, we discuss these issues of previous metrics and propose a novel benchmark



Fig. 5: Stable Diffusion XL 768x768 generation results. PTQ4DM and Q-diffusion fail to match the textual semantics, while our method with the activation relaxing can generate high-quality images matching the prompt.

QDiffBench, which consists of a more accurate FID strategy and a promptgeneralization evaluation strategy, to accurately and comprehensively evaluate the quantization of text-to-image diffusion models.

4.1 FID to Full Precision Model

FID ((Frechet Inception Distance) [10] is an important metric to evaluate the quality of generative models, which calculates the distance between generated images and real images, where a lower score indicates better performance. FID is also usually used to evaluate text-to-image diffusion models [23]. For example, Rombach et al. [23] evaluate latent text-to-image diffusion models on the COCO validation dataset. They generate images conditioned by COCO prompts, and calculate the FID score between generated images and real COCO images.

Previous studies also use the same way to evaluate quantized text-to-image diffusion models [13], i.e., generating images with a quantized model and calculating FID between the generated and the real ones. We name this strategy "FID to COCO". However, we demonstrate this strategy is inaccurate in the context of quantization. For instance, as shown in Tab. 1, the quantized models usually have similar or even lower "FID to COCO" scores than the full-precision model. However, the generated images by the quantized models are generally worse than the precision ones, taking Fig. 7 as an example. In addition, Tab. 1 also shows the W4A8 models have better "FID to COCO" scores than the W8A8 peers, which is also inconsistent with the visualized results, as shown in the Appendix (Fig. A.1). These phenomena prove the inaccuracy of "FID to COCO". The reason is that the distribution of the generated images by text-to-image diffusion models obviously differs from the COCO images. The manifestations of this distribution gap include style differences, fake details in generated images, etc. We detailedly analyze the reasons and manifestations of distribution-gap in the Appendix (Sec. A.5). This gap makes the distance to the COCO images inaccurate to measure the quality of the images generated by quantized models. Therefore, the worse quantized models can even have better "FID to COCO" scores.

To tackle this problem, we propose to directly compare the images generated by the quantized model with the images generated by the full precision model, i.e., calculating the FID score between them, as shown in Fig. 6. We name this strategy "FID to FP32". This strategy straightforwardly measures the degradation caused by the quantization algorithm. It compares the generated images to the data in the same domain, making it more accurate. FID to FP32 measures image fidelity when text-image matching is good, but it can be quite poor with a bad text-image matching,

Tab. 1 also shows our "FID to FP32' metric is basically consistent with the CLIP score. Fig. 7 and the samples in the Appendix also visually demonstrate our metric is consistent with human cognition.



Fig. 6: We calculate the FID to the images generated by the full-precision model rather than the COCO images. COCO images are photorealism, while Stable Diffusion can't achieve this perfectly. Additionally, there are obvious fake details in the generated plane images.

4.2 Prompt-Generalization Ability

An outstanding feature of text-to-image diffusion models is that they can generate images of various styles and themes, as long as the user enters the appropriate prompts. Although foundation diffusion models (e.g., Stable Diffusion and Stable Diffusion XL) are pretrained on billions of image-text pairs, researchers can usually only use a small set of data for quantization. For example, previous studies use a subset of prompts from the COCO dataset to perform calibration. In the evaluation, they only test the quantized models on the COCO prompts. However, users always input prompts with a quite different style from the calibration dataset, and generate images of various styles. Therefore, to evaluate the



Fig. 7: Comparison between the quantized models and the full-precision model. Use Stable diffusion 512 x 512. Relaxation proportion $\tau = 0.20$.

prompt-generalization of the quantized models, we propose to test them also on the unseen prompts that have a quite different style from the calibration dataset. In our experiments, we evaluate the quantized model on the dataset referred to as *Stable-Diffusion-Prompts* [7], which is a popular Stable Diffusion prompt dataset released on *hugging face* [3]. It consists of about 80,000 user-created high-quality prompts that can generate images of various styles. As shown in Fig. 8, the generated images by this dataset have quite different styles from the calibration dataset (COCO prompts in our settings), i.e.:

COCO prompts always generate realistic photographs, while Stable-Diffusion-Prompts tends to generate artistic images.

We use the first 5,000 prompts of *Stable-Diffusion-Prompts* as the promptgeneralization test set.



Fig. 8: Stable Diffusion 512 x 512 generation results on COCO prompts and *Stable-Diffusion-Prompts*. Relaxation proportion $\tau = 0.20$.

5 Experiments

5.1 Experimental Setup

Datasets. Our experiments involve with two datasets, i.e., COCO [15] and *Stable-Diffusion-Prompts* [7] (as discussed in Sec. 4.2). During the calibration, we

use a few prompts of the COCO [15] training dataset. During the evaluation, on the one hand, we use 5,000 prompts from the previously widely used COCO [15] validation dataset. On the other hand, we further use 5,000 prompts in the *Stable-Diffusion-Prompts* dataset to evaluate the prompt-generalization ability of the quantized model.

Metrics. We mainly adopt two metrics to evaluate generation quality. The first metric we adopt is the "FID to FP32", which is the FID score between the quantized-model generated images and full-precision-model generated images. Additionally, we also adopt the CLIP score [9] to evaluate whether the generated images match the text prompt. Here, we use the model ViT/L-14 to calculate CLIP scores. For the reference and the illustration of the problem in Sec. 4.1, we also present the metric "FID to COCO" adopted by previous works, but it is not proper to be used for comparison due to the distribution gap. Besides, We calculate the **BOPs** [34] (BOPs = FLOPs $\cdot b_w \cdot b_a$) metric to measure computation amount, where b_w and b_a denote the bitwidth of weights and activations respectively.

Baselines and implementation details. We compare our methods with two state-of-the-art methods Q-diffusion [13] and PTQ4DM [25]. For a fair comparison, we implement all the baselines and our method in a unified framework based on *diffusers* [21]. For quantization, we use block reconstruction [14] to adaptively round [18] the weight by minimizing the mean squared errors between the quantized and the full precision outputs, following previous works [8,13,25]. We use the Stable Diffusion v1-4 and the Stable Diffusion XL 1.0 checkpoints provided by *hugging face* [3]. We use the default configuration, i.e., 50-step PNDM sampling for Stable Diffusion and 50-step Euler sampling for Stable Diffusion XL. The calibration prompt number is 200 for Stable Diffusion and 30 for Stable Diffusion XL.

Relaxation settings. We relax the sensitive steps from 8-bit to 10-bit. By default, we set the relaxation proportion τ to 0.20 (i.e., 20%), but we also prove relaxing only 5% steps ($\tau = 0.05$) can also outperform the baselines by a large margin.

5.2 Main Results

Tab. 1 shows the results of Stable Diffusion on the COCO validation prompts, which prove the effectiveness of our proposed method and benchmark. Specifically, "FID to COCO" is the metric used in previous works, while "FID to FP32" is our proposed metric that calculates the distance between the quantized model and the full-precision 32-bit model. The quantized models have similar or even much lower "FID to COCO" than the full-precision model, which is inaccurate, as discussed in Sec. 4.1. On the other hand, the results of the CLIP score and our "FID to FP32" show our method reaches a new state-of-the-art performance.

The discussion about CLIP score. Tab. 1 also reveals that the CLIP score might not be able to precisely distinguish the quality of quantized models when their CLIP score is closer to the full-precision score. For example, the CLIP score of PCR 4/8 is a little higher than PCR 8/8 ($\tau = 0.20$), even though the

Table 1: Quantization results on COCO validation prompts for Stable Diffusion v1-4. \uparrow means higher metric is better, while \downarrow means lower is better. W/A means the bitwidth for weights and activations respectively. τ denotes the proportion of activation relaxing.

Method	Bits(W/A)	Size (GB)	BOPs (T)	FID to COCO	FID to FP32 \downarrow	CLIP score↑
FP32	32/32	3.44	693	27.50	0.00	26.46
Q-diffusion	8/8	0.87	43.31	29.98	18.64	26.15
PTQ4DM	8/8	0.87	43.31	28.00	14.60	26.33
$PCR(\tau = 0.05)$	8/8	0.87	43.85	25.16	9.92	26.46
$PCR(\tau = 0.20)$	8/8	0.87	45.47	25.83	8.35	26.47
Q-diffusion	4/8	0.44	21.66	27.87	20.42	26.15
PTQ4DM	4/8	0.44	21.66	25.64	17.73	26.25
$PCR(\tau = 0.05)$	4/8	0.44	21.93	23.86	14.39	26.35
$PCR(\tau = 0.20)$	4/8	0.44	22.74	22.04	14.25	26.48

Table 2: Quantization results on COCO validation prompts for Stable Diffusion XL.

Method	Bits(W/A)	Size (GB)	BOPs (T)	FID to COCO	FID to FP32 \downarrow	CLIP score↑
FP32	32/32	10.30	6933	27.28	0.00	26.44
Q-diffusion	8/8	2.61	433	42.85	38.19	15.97
PTQ4DM	8/8	2.61	433	42.33	38.65	15.98
$PCR(\tau = 0.05)$	8/8	2.61	438	30.87	22.08	19.73
$PCR(\tau = 0.20)$	8/8	2.61	455	26.16	12.00	24.05
Q-diffusion	4/8	1.32	217	45.43	44.07	15.91
PTQ4DM	4/8	1.32	217	43.37	46.45	15.92
$PCR(\tau = 0.05)$	4/8	1.32	219	34.99	30.87	19.51
$PCR(\tau = 0.20)$	4/8	1.32	227	24.27	18.27	23.85

quality of PCR 8/8 ($\tau = 0.20$), is obviously better, as shown in Fig. A.1 and Fig. A.2 in the Appendix. These phenomena support the significance of our "FID to FP32" metric to measure image fidelity. Therefore, it is significant to combine the CLIP score and our "FID to FP32" score to comprehensively measure the quality of quantized models.

Tab. 2 show the results of Stable Diffusion XL on COCO validation prompts. Our method outperforms previous methods by an obvious margin. Moreover, the visualized results in Fig. 5 show that previous methods fail to generate images that can properly match the prompt, while our method can generate high-quality images matching the text.

Tab. 1 and Tab. 2 also indicate relaxing only 5% steps ($\tau = 0.05$) can also outperform the baselines by a large margin.

To assess prompt generalization, we directly test quantized checkpoints (obtained on COCO) on the unseen dataset *Stable-Diffusion-Prompts*. Tab. 3 show the results of Stable Diffusion. Although our method has a slightly higher FID score in the setting of W4/A8, its CLIP score obviously exceeds the baselines. Tab. 4 show the results of Stable Diffusion XL. These results prove our method has better generalization on the unseen prompts outside the calibration dataset.

Comparing the results evaluated on COCO with the results on *Stable-Diffusion-Prompts*, we can observe some differences. For example, the gap between these quantized Stable Diffusion XL models is obviously smaller on the *Stable-Diffusion-Prompts* than the COCO. This may indicate slight overfitting to the calibration data, which also supports the necessity of the evaluation on the unseen prompts outside the calibration data.

Fig. 9 shows the generated samples by Stable Diffusion XL using *Stable-Diffusion-Prompts*. More samples can be found in the Appendix.

Bits(W/A) Size (GB) BOPs (T) FID to FP32↓ CLIP score↑ Method 28.79 27.19 FP32 32/32 3.44 693 0.00 Q-diffusio 8/8 0.87 43.31 16.22PTQ4DM 8/8 0.87 43.31 13.25 27.77 $PCR(\tau = 0.20)$ 8/8 0.87 45.479.5228.7417.43 17.28 27.31 27.39 Q-diffusion PTQ4DM 0.44 21.66 21.664/80.44 $PCR(\tau = 0.20)$ 22.74 28.050.44 17.95

Table 3: Results on Stable-Diffusion-Prompts dataset for Stable Diffusion v1-4.

Table 4: Results on Stable-Diffusion-Prompts dataset for Stable Diffusion XL.

Method	Bits(W/A)	Size (GB)	BOPs (T)	FID to FP32 \downarrow	CLIP score \uparrow
FP32	32/32	10.30	6933	0.00	29.74
Q-diffusion	8/8	2.61	433	22.81	20.01
PTQ4DM	8/8	2.61	433	22.85	20.09
$PCR(\tau = 0.20)$	8/8	2.61	455	11.71	25.74
Q-diffusion	4/8	1.32	217	24.98	19.93
PTQ4DM	4/8	1.32	217	28.28	20.13
$PCR(\tau = 0.20)$	4/8	1.32	227	18.25	25.71

The BOPs results in Tabs. 1 to 4 indicate that our PCR method only has a negligible extra computational cost.



Fig. 9: Stable Diffusion XL 768x768 generation using *Stable-Diffusion-Prompts*.

5.3 Ablation Study

The effects of our components. In this section, we analyze the effects of different components. Tab. 5 demonstrates that solely using the activation relaxing and solely using the progressive calibration can both improve the performance a lot. And combining these two strategies further improves the quality. Moreover, note that the progressive calibration is essential for maintaining the CLIP score. The quantized models suffer a remarkable degradation of CLIP score without the progressive calibration. For example, in the W4A8 setting, only adding the

Method	Bits(W/A)	FID to FP32 \downarrow	CLIP score \uparrow
FP32	32/32	0.00	26.46
Base	8/8	15.90	26.42
+ Relax	8/8	9.02	26.42
+ Prog	8/8	11.01	26.48
+ Prog & Relax	8/8	8.35	26.47
Base	4/8	25.29	26.21
+ Relax	4/8	15.63	26.35
+ Prog	4/8	17.54	26.49
+ Prog & Relax	4/8	14.25	26.48

Table 5: Ablation results on the COCO validation prompts for Stable Diffusion v1-4. "Relax" denotes the activation relaxing, and "Prog" denotes the progressive calibration.

activation relaxing to the base just has a 26.35 CLIP score (row 8), while further adding the progressive calibration can raise the score to 26.48 (row 10).

The effects of the relaxing proportion. We explore how the relaxation proportion τ affects the performance. Fig. 10 shows the score changes of Stable Diffusion (SD) and Stable Diffusion XL (SDXL) with respect to the proportion. The results show the gain from relaxing more steps decreases rapidly as the proportion grows, which means the steps nearest the x_0/x_T are the most important to the image-fidelity/text-image-matching.



Fig. 10: Score variations w.r.t. the relaxing proportion, on COCO validation prompts, W8/A8. The solid lines denote the scores of PCR. The dashed lines parallel to the x-axis denote the corresponding scores of the baselines.

6 Conclusion

This paper aims to advance the quantization of text-to-image diffusion models from two perspectives, method and benchmark. We propose a novel quantization method PCR and an effective benchmark QDiffBench. Extensive experiments prove the superiority of our method and benchmark.

Acknowledgement

This work is supported by the National Key Research and Development Program of China No.2023YFF1205001, National Natural Science Foundation of China (No. 62250008, 62222209, 62102222), Beijing National Research Center for Information Science and Technology under Grant No. BNR2023RC01003, BNR2023TD03006, and Beijing Key Lab of Networked Multimedia.

References

- https://www.qualcomm.com/products/mobile/snapdragon/smartphones/ snapdragon-8-series-mobile-platforms/snapdragon-8-gen-2-mobileplatform
- https://images.nvidia.com/aem-dam/Solutions/Data-Center/14/nvidiaada-gpu-architecture-whitepaper-v2.1.pdf
- 3. https://huggingface.co/
- Chen, H., Zhang, Y., Wang, X., Duan, X., Zhou, Y., Zhu, W.: Disenbooth: Disentangled parameter-efficient tuning for subject-driven text-to-image generation. arXiv preprint arXiv:2305.03374 (2023)
- Choquette, J., Gandhi, W., Giroux, O., Stam, N., Krashinsky, R.: Nvidia a100 tensor core gpu: Performance and innovation. IEEE Micro 41(2), 29–35 (2021)
- Dong, Z., Yao, Z., Gholami, A., Mahoney, M.W., Keutzer, K.: Hawq: Hessian aware quantization of neural networks with mixed-precision. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 293–302 (2019)
- 7. Gustavosta: https://huggingface.co/datasets/Gustavosta/Stable Diffusion-Prompts
- 8. He, Y., Liu, L., Liu, J., Wu, W., Zhou, H., Zhuang, B.: Ptqd: Accurate post-training quantization for diffusion models. arXiv preprint arXiv:2305.10657 (2023)
- Hessel, J., Holtzman, A., Forbes, M., Le Bras, R., Choi, Y.: Clipscore: A referencefree evaluation metric for image captioning. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing. pp. 7514–7528 (2021)
- 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)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems 33, 6840–6851 (2020)
- Koryakovskiy, I., Yakovleva, A., Buchnev, V., Isaev, T., Odinokikh, G.: One-shot model for mixed-precision quantization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7939–7949 (2023)
- 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)
- 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)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Computer Vision– ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13. pp. 740–755. Springer (2014)

- 16 S. Tang et al.
- 16. Liu, L., Ren, Y., Lin, Z., Zhao, Z.: Pseudo numerical methods for diffusion models on manifolds. In: International Conference on Learning Representations (2021)
- 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. Advances in Neural Information Processing Systems 35, 5775–5787 (2022)
- 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)
- 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)
- Nagel, M., Fournarakis, M., Bondarenko, Y., Blankevoort, T.: Overcoming oscillations in quantization-aware training. In: International Conference on Machine Learning. pp. 16318–16330. PMLR (2022)
- von Platen, P., Patil, S., Lozhkov, A., Cuenca, P., Lambert, N., Rasul, K., Davaadorj, M., Wolf, T.: Diffusers: State-of-the-art diffusion models. https:// github.com/huggingface/diffusers (2022)
- Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952 (2023)
- 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)
- 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)
- 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)
- So, J., Lee, J., Ahn, D., Kim, H., Park, E.: Temporal dynamic quantization for diffusion models. arXiv preprint arXiv:2306.02316 (2023)
- 27. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Scorebased generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2020)
- 29. Tailor, S.A., Fernandez-Marques, J., Lane, N.D.: Degree-quant: Quantizationaware training for graph neural networks. arXiv preprint arXiv:2008.05000 (2020)
- Wang, C., Wang, Z., Xu, X., Tang, Y., Zhou, J., Lu, J.: Towards accurate data-free quantization for diffusion models. arXiv preprint arXiv:2305.18723 (2023)
- Wang, K., Liu, Z., Lin, Y., Lin, J., Han, S.: Haq: Hardware-aware automated quantization with mixed precision. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8612–8620 (2019)
- Wang, X., Chen, H., Tang, S., Wu, Z., Zhu, W.: Disentangled representation learning. IEEE Transactions on Pattern Analysis and Machine Intelligence pp. 1–20 (2024). https://doi.org/10.1109/TPAMI.2024.3420937
- Wu, B., Wang, Y., Zhang, P., Tian, Y., Vajda, P., Keutzer, K.: Mixed precision quantization of convnets via differentiable neural architecture search. arXiv preprint arXiv:1812.00090 (2018)

PCR 17

34. Yu, H., Han, Q., Li, J., Shi, J., Cheng, G., Fan, B.: Search what you want: Barrier panelty nas for mixed precision quantization. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16. pp. 1–16. Springer (2020)