

MixDQ: Memory-Efficient Few-Step Text-to-Image Diffusion Models with Metric-Decoupled Mixed Precision Quantization

Tianchen Zhao^{12*}, Xuefei Ning^{1*†}, Tongcheng Fang^{12*}, Enshu Liu¹, Guyue Huang³, Zinan Lin⁴, Shengen Yan², Guohao Dai²⁵, and Yu Wang^{1†}

¹ Tsinghua University

² Infinigence AI

³ University of California Santa Barbara

⁴ Microsoft

⁵ Shanghai Jiaotong University

Abstract. Few-step diffusion models, which enable high-quality text-to-image generation with only a few denoising steps, have substantially reduced inference time. However, considerable memory consumption (5-10GB) still poses limitations for practical deployment on mobile devices. Post-Training Quantization (PTQ) proves to be an effective method for enhancing efficiency in both memory and operational complexity. However, when applied to few-step diffusion models, existing methods designed for multi-step diffusion face challenges in preserving both visual quality and text alignment. In this paper, we discover that the quantization is bottlenecked by highly sensitive layers. Consequently, we introduce a mixed-precision quantization method: **MixDQ**. Firstly, we identify some highly sensitive layers are caused by outliers in text embeddings, and design a specialized Begin-Of-Sentence (BOS)-aware quantization to address this issue. Subsequently, we investigate the drawback of existing sensitivity metrics, and introduce metric-decoupled sensitivity analysis to accurately estimate sensitivity for both image quality and content. Finally, we develop an integer-programming-based method to obtain the optimal mixed-precision configuration. In the challenging 1-step Stable Diffusion XL text-to-image task, current quantization methods fall short at W8A8. Remarkably, MixDQ achieves **W3.66A16** and **W4A8** quantization with negligible degradation in both visual quality and text alignment. Compared with FP16, it achieves **3-4**× reduction in model size and memory costs, along with a **1.5**× latency speedup.

Keywords: Diffusion Model · Quantization

* Equal contribution

† Corresponding Authors: Yu Wang (yu-wang@mail.tsinghua.edu.cn), Xuefei Ning (foxdoraame@gmail.com)




Fig. 1: The effectiveness of MixDQ. Left: MixDQ preserves both image quality and image-text alignment. Right: The efficiency improvements of MixDQ.

1 Introduction

Text-to-image diffusion models [30, 33, 34] attract substantial research attention due to their capability to generate high-quality images from textual prompts. However, their demanding computational and memory requirements present challenges for real-time application and deployment on mobile devices [2, 56]. Recent efforts on few-step diffusion models [18, 38, 42] have significantly alleviated the computational burden, which require only 1-8 steps of iteration to generate high-fidelity images, as opposed to the 10-100 steps needed in previous approaches [8, 41]. However, the memory cost of diffusion models remains excessive [35, 46]. For example, running a Stable Diffusion XL-turbo model [30] in the 16-bit floating-point (FP16) format with a batch size of 1 consumes a peak memory of 9.7GB, which exceeds the capacity of many mobile devices and even some desktop GPUs (e.g., an RTX 4070 has 8GB GPU memory).

Model quantization [10], compressing high bit-width floating-point parameters and activations into lower bit-width integers, proves to be an effective strategy for reducing both memory and computational cost. Many prior researches [4, 5] explore quantization for diffusion models. However, we observe that **few-step diffusion models are more sensitive to quantization than multi-step ones**. As depicted in Fig. 1, the quantized SDXL 30-step model and SDXL-turbo 2-step model have notably less degradation compared with the SDXL-turbo 1-step model. Prior research Q-diffusion [13] performs well on multi-step models but encounters challenges with W8A8 quantization for one-step SDXL-turbo [38]. The resulting image appears blurred and contains numerous artifacts. We conjecture this might be due to the absence of the iterative denoising process, which could compensate for the quantization error.

Additionally, prior research primarily focuses on preserving the image quality for quantized models. However, as illustrated in Fig. 1 and Fig. 2, quantization

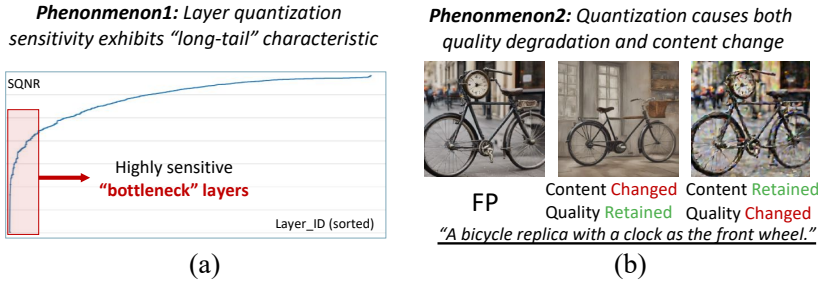


Fig. 2: Insightful findings for few-step text-to-image diffusion quantization. Left: The layer-sensitivity distribution has a “long-tail” characteristic. Right: Quantization affects both image quality and content.

impacts not only image quality but also content. The altered content may lead to **degradation in image-text alignment**, which refers to how well the generated image aligns with the given text instruction. For instance, in Fig. 1, the image generated by Q-diffusion depicts a polar bear, which contradicts the prompt mentioning “goats”. Similarly, the middle image in Fig. 2 loses the content “clock”.

To effectively quantize the few-step model while preserving the quality and alignment in the meantime, we investigate the reasons behind the failure of existing quantization methods. We discover two insightful phenomena, as presented in Fig. 2. Firstly, we measure each layer’s quantization sensitivity for the SDXL-turbo model with the signal-to-quantization-noise ratio (SQNR) following prior research [54]. As shown in Fig. 2 (a), the distribution exhibits a “long-tail” characteristic. Consequently, using uniform quantization settings for all layers, as in previous research, would result in the quantization being “bottlenecked” by some highly sensitive layers. We look into these sensitive layers and propose: (1) designing a specialized quantization method to “protect” the outliers in certain layers, significantly reducing their quantization error; and (2) adopting mixed-precision quantization to “protect” the sensitive layers with a higher bit-width, thus achieving higher quality. Secondly, as illustrated in Fig. 2 (b), quantization affects both the image content and quality. We identify that the entanglement of these two factors in evaluation will cause quantization methods to fail in preserving generation quality. Therefore, we propose “decoupling” the effects of quality and content in the sensitivity evaluation. This will be elaborated in Sec. 3.2.

In light of above, we introduce a mixed-precision quantization method: **MixDQ**. Firstly, we analyze the characteristics of highly sensitive layers and discover that many of them are associated with the quantization of text embeddings. We further investigate the feature distribution of text embeddings and design a specialized **Begin-of-sentence (BOS)-aware quantization** (Sec. 3.1) to address the outlier values. Secondly, we identify the shortcomings of existing quantization sensitivity evaluation and design an improved **Metric-decoupled sensitivity analysis** (Sec. 3.2) based on the idea of decoupling the impact of quantization on image content and quality. Finally, we design an **integer-programming-**

based mixed precision allocation (Sec. 3.3) to acquire optimal bit-width configuration based on the improved sensitivity analysis.

We summarize the contributions of this paper as follows:

1. We highlight the challenge of quantizing the few-step diffusion model, compared to quantizing the multi-step diffusion model. Additionally, we emphasize the often-overlooked necessity of preserving alignment when compressing the text-to-image generative models.
2. Based on the careful investigation of the data distribution and sensitivity of each layer, we design MixDQ, a mixed precision quantization method with improved sensitivity evaluation and quantization techniques.
3. We evaluate MixDQ in two settings: weight-only quantization and normal weight and activation quantization. MixDQ can achieve W4A16 and W5A8 quantization with a negligible 0.1 FID increase. Compared with the FP16 model, MixDQ can achieve W3.66A16 and W4A8 quantization within a 0.5 FID increase, resulting in a $3\times$ memory cost reduction and a $1.5\times$ latency speedup on Nvidia GPUs.

The techniques in MixDQ could benefit future research and applications of compression methods on other generative models and tasks. Firstly, BOS-aware quantization addresses the issue of outlier values in text embeddings, which is an inherent problem for transformer-based models (also recognized as the “attention sink” [52] in language models). Secondly, we believe the explicit and decoupled consideration of various metrics is important whenever compressing models on broader visual generation tasks. On one hand, from an application perspective, multiple metrics should be considered, especially for generative tasks. On the other hand, for better compression results, the explicit and decoupled consideration of multiple metrics can avoid the failure of compression methods. We will release all the code and models.

2 Related Work

Diffusion Models. [8, 40, 43] could generate high-quality images through an iterative denoising process. However, the excessive cost of repeated denoising iterations calls for improvements in the sampling efficiency, namely, reducing the timesteps. Some research [17, 41, 55] focuses on designing improved numerical solvers, and another line of research explores utilizing distillation [18, 22, 37, 38] to condense the sampling trajectory into fewer steps or even a single step. These few-step diffusion models significantly reduce the computational cost, however, as discussed in Sec. 1, **they face additional challenges for quantization.**

Network Quantization. Prior research, such as PTQD [5] and Q-DM [14] explores utilizing quantization techniques on diffusion models. Q-Diffusion [13] extends this application to large-scale stable diffusion models. Other research [9, 39, 44] continues to improve post-training quantization techniques on aspects like calibration and temporal adjustment. However, the majority of existing diffusion

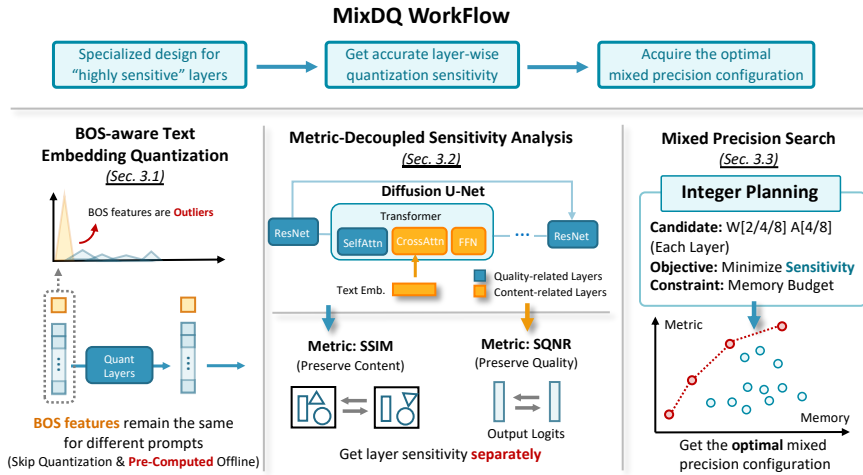


Fig. 3: Framework of the proposed mixed-precision quantization method: MixDQ. It consists of three key components, the BOS-aware quantization addresses the highly sensitive text embedding, the metric-decoupled scheme improves sensitivity analysis, and the integer programming acquires the optimal bit-width allocation.

quantization methods **employs uniform bit-width for layers with diverse sensitivity**. Yang et. al. [54] suggest preserving sensitive layers as high-precision based on SQNR metric. However, we observe that the SQNR metric tends to prioritize content change, potentially leading to quality degradation (as discussed in Sec. 3.2). Inspired by the prior success of mixed precision [3, 47], we design a mixed precision quantization method to address the imbalanced sensitivity.

Evaluation Metrics. The currently widely used metrics for evaluating text-to-image generation can be summarized into two major aspects: image fidelity (quality) and image-text alignment. To estimate image fidelity, metrics such as FID [7] and IS [36] are commonly used. These metrics measure the feature space distance between generated images and reference images. For image-text alignment, CLIP Score [6] is often utilized to calculate the similarity of the image and the text embeddings. To align the statistical scores and human preference, ImageReward [53] and HPS [51] collect user preference and train the model to predict generated image scoring, considering both fidelity and alignment. While previous methods solely discuss preserving the image quality, we investigate quantization’s effect on **both image fidelity and alignment**.

3 Methods

Fig. 3 presents our mixed precision quantization method **MixDQ** consisting of three consecutive steps. Firstly, we identify the highly sensitive layers and examine their distinctive properties. Based on these properties, we devise specialized

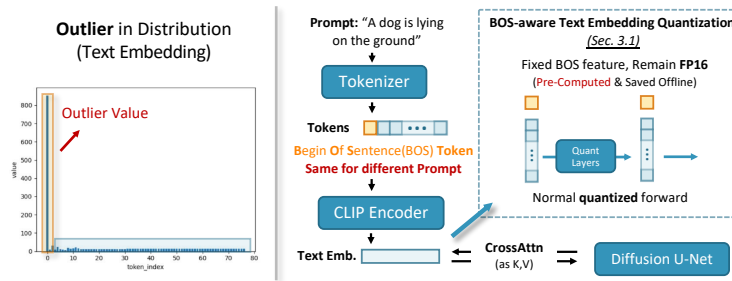


Fig. 4: Illustration of BOS-aware Quantization. Left: the first token has a significantly larger value than the others. Right: Since BOS token features remain the same for different prompts, we skip quantizing them and pre-compute them offline.

quantization techniques tailored for them (Sec. 3.1). Secondly, in Sec. 3.2, we introduce a metric-decoupled analysis that measures the quantization’s effects on image content and quality separately. Finally, based on the sensitivity, we employ integer programming to determine the optimal mixed precision configuration under a given budget (Sec. 3.3).

3.1 BOS-aware Text Embedding Quantization

Fig. 2 shows that the diffusion model quantization is “bottlenecked” by some highly sensitive layers. Upon going through the highly sensitive layers depicted in the “long-tail”, we find that a substantial portion of them corresponds to the “to_k” and “to_v” linear layers in the cross-attention. All of these layers take the text embedding, which is the output of the CLIP encoder [32], as their input. Fig. 4 (Left) shows the maximum magnitude of the text embedding of each token (average across 8 sentences). We observe that the 1st token has a significantly larger magnitude (823.5) compared to the rest of the tokens (10-15). Quantizing this tensor with such an outlier value to 8-bit would lead to the majority of values being close to 0, resulting in the loss of crucial textual information. This observation could potentially explain why current quantization methods struggle to maintain text-image alignment in Fig. 1 (Left).

As shown in Fig. 4 (Right), the first outlier token in the CLIP output corresponds to the “Begin Of Sentence (BOS)” in the tokenizer output. Actually, *the feature of this BOS token remains the same across different prompts*. Therefore, we can skip the quantization and computation for the “to_k/v” layers for this token. Concretely, we pre-compute the floating-point output feature of the “to_k/v” layers for the BOS token and then concatenate it with the dequantized features of other tokens. In this way, the quantization error of the CLIP embedding is significantly reduced as the outlier is removed. Additionally, only 640-1280 (the channel number of BOS features) elements need to be stored for each “to_k/v” layer, introducing only negligible overhead.




Fig. 5: Decoupling metrics and layers to separate the influence on image quality and content. Left: Existing SQNR-based sensitivity analysis needs improving. Right: SQNR’s problem of overemphasizing content change.

3.2 Metric-Decoupled Sensitivity Analysis

Preliminary Experiment and Analysis. We start by designing preliminary experiments. As discussed in Sec. 1, we measure the sensitivity with SQNR. It is computed as the ratio of the L2 norm of feature values to the quantization noise. The quantization noise is estimated by evaluating the impact on the network’s output logits with each layer quantized. We observe the “long-tailed” characteristics of sensitivity distribution in Fig. 2 (a). An intuitive solution is to allocate higher bit-widths for highly sensitive layers and lower bit-widths for less sensitive ones. We implement a straightforward bit-width allocation method based on this principle to achieve an average of W8A8. However, as seen in Fig. 5 (Left), the generated image faces severe quality degradation, suggesting that the current sensitivity evaluation relying solely on SQNR may require refinement.

We look into this issue and observe the following phenomena in Fig. 5 (Right): (1) Both changes in image content and quality impact the SQNR. In particular, compared with the 30.03dB image, the SQNR decline of an image with changed contents but high quality (-3.51dB) is much more significant than that of an image with similar contents but unacceptable quality (-0.26dB). This means that *when evaluating images with varying content and quality, SQNR tends to underestimate the performance degradation resulting from decreased image quality.* (2) *Image quality and content are mainly affected by different layer types.* Specifically, quantizing the cross-attention layers causes content changes, but the quality remains good. This aligns with the intuition that text supervision predominantly interacts with cross-attention layers, which control the image content. Conversely, when quantizing convolutions, the quality degrades and the content is preserved. Further details of the correlation between layer types and their impact on content or quality are provided in Appendix D.

As a result, during the bit-width allocation process, since SQNR tends to overemphasize content changes, the content-related layers are retained at very high bit-widths. Consequently, the bit-width for quality-related layers is decreased. This “*unfair competition*” between quality- and content-related layers explains the severe degradation of the image quality observed in Fig. 5 (Left) when the bit-width allocation is solely based on SQNR.

Method Design. Inspired by these findings, we improve the sensitivity analysis by decoupling metrics and layers to separate the quantization’s effect on quality and content. Therefore, we separate the layers into two groups: the content-related layers, i.e., the cross-attention layers and feed forward networks (FFNs), and the quality-related layers, i.e., self-attention layers and convolutions. Then, we conduct sensitivity analysis for each group separately with distinct metrics.

When performing the sensitivity analysis for the content-related layers, we adopt the Structural Similarity Index Measure (SSIM) [49] metric to assess image content change. Given a generated image x and reference image y , The SSIM metric combines the three components, the luminance (l), the contrast (c) and the structure (s):

$$l(x, y) = \frac{2\mu_x\mu_y}{\mu_x^2 + \mu_y^2}, c(x, y) = \frac{2\sigma_x\sigma_y}{\sigma_x^2 + \sigma_y^2}, s(x, y) = \frac{\sigma_{xy}}{\sigma_x\sigma_y}, \quad (1)$$

$$\text{SSIM}(x, y) = l(x, y)^\alpha \cdot c(x, y)^\beta \cdot s(x, y)^\gamma,$$

where μ, σ stand for the mean and variance of pixel values; and α, β, γ are weights that control the combination of three components. We set $\alpha = \beta = \gamma = 1$, and use the full image size as the window size for SSIM. Unlike Mean Squared Error (MSE) and Peak Signal-to-Noise Ratio (PSNR) that measure absolute errors, SSIM is adept at perceiving changes in structural information, making it well-suited for measuring content change.

When conducting sensitivity analysis for the quality-related layers, we utilize the SQNR metric. Given that quantizing only the quality-related layers does not significantly alter the image content, SQNR serves as a suitable metric for assessing the degree of image quality degradation in this scenario.

3.3 Integer Programming Bit-width Allocation

Having obtained the layer sensitivity, we can determine the mixed precision configuration by assigning higher bit-widths to more sensitive layers. Enlightened by prior research [12], we formulate the bit-width allocation into an integer programming problem: Given the resource budget \mathcal{B} , and candidate bit-widths $b \in \{2, 4, 8\}$, we aim to determine the bit-width choices \mathbf{c} (a one-hot indicator for each layer) that maximizes the sum of sensitivity $\mathcal{S} = \sum_{i=1}^N \sum_{b=2,4,8} c_{i,b} \cdot \mathcal{S}_{i,b}$:

$$\begin{aligned} & \underset{c_{i,b}}{\text{argmax}} \quad \sum_{i=1}^N \sum_{b=2,4,8} c_{i,b} \cdot \mathcal{S}_{i,b} \\ & \text{s.t.} \quad \sum_{b=2,4,8} c_{i,b} = 1, \quad \sum_{i=1}^N \sum_{b=2,4,8} c_{i,b} \cdot \mathcal{M}_{i,b} \leq \mathcal{B}, \\ & \quad c_{i,b} \in \{0, 1\}, \quad \forall i \in \{1, \dots, N\}, \forall b \in \{2, 4, 8\}, \end{aligned} \quad (2)$$

where N is the number of layers in the model; $c_{i,b} = 1$ indicates that the i -th layer will be quantized to b -bit, and $\mathcal{S}_{i,b}$ is the corresponding sensitivity score

(higher SQNR or SSIM is better). $\mathcal{M}_{i,b}$ denotes the resource cost of the i -th layer when it is quantized to b -bit.

We conduct the above integer programming separately for activation quantization and weight quantization, as well as for each of the two layer groups, i.e., the content-related cross-attention and FFN layers, and the quality-related self-attention and convolution layers.

4 Experiments

4.1 Experimental Settings

Quantization Scheme. We adopt the simplest and easily deployable asymmetric min-max quantization scheme similar to [25,29]. The quantization parameters (i.e., scaling factor, zero point) are shared within each tensor for the activation or within each output channel for the weight. The shortcut-splitting quantization technique [13] is applied to all methods in Tab. 1. We randomly sampled 1024 prompts from COCO [15] as our calibration dataset.

Mixed Precision Allocation. We do not quantize all nonlinear activation and normalization layers, and quantize all the linear and convolution layers. In our metric-decoupled sensitivity analysis, we iterate through all layers one by one. For each layer, our method quantizes only that layer and preserves the remaining layers as FP. Then, we measure two types of sensitivity metrics for this layer: SQNR and SSIM. The sensitivity score is averaged over 32 prompts. We quantize each layer to three bit-width choices: 2, 4, 8, and get their corresponding sensitivity scores. Then, we set a budget for the average bit-width of all elements (weights or activation) and use the layer-wise sensitivity scores to set up the integer programming. We use the OR-tools [45] library for integer programming. The efficiency of this implementation allows the bit-width allocation in seconds. For both layer groups (content-related and quality-related), we sweep 20 choices of the average bit-width budget, ranging uniformly from 3 to 8. Then, we keep the Pareto frontier of the $20 \times 20 = 400$ bit-width configurations (Fig. 8).

Hardware Profiling. We measure the latency and memory usage of MixDQ on the Nvidia RTX 4080 GPU using CUDA 12.1. All profiling is conducted with a batch size of 1. Specifically, we measure memory usage for all models using the PyTorch Memory Management APIs [31]. The inference latency of the FP16 models is profiled with NVIDIA Nsight tools [27]. For quantized models, we measure the latency of the quantized layers, the quantization operation, and the unquantized layers separately, and add them up to calculate the overall latency. We profile the FP and quantized models using stable-fast [19], a toolkit that provides state-of-the-art inference speed for diffuser models. We use the kernels from the Cutlass [28] library to implement the quantized layers and develop a quantization GPU kernel to reduce the quantization overhead.

4.2 Performance and Efficiency Comparison

We conduct experiments on widely-used few-step diffusion models, SDXL-turbo, and LCM-Lora, for text-to-image generation tasks using COCO2014 at the res-

Table 1: Performance and efficiency comparison of MixDQ and other quantization methods on full COCO annotations. The “CLIP” and “IR” denotes CLIP Score and ImageReward metric. The “Storage Opt.” and “Compute Opt.” denote equivalent savings of model size and computational complexity (measured in Bit Operations as in [29]). The bit-width “16” represents FP16 without quantization. The “weight only” setting represents the rows with activation bit-width of FP16, and the rest are the “normal” weight-activation quantization.

Model	Method	Bit-width (W/A)	Storage Opt.	Compute Opt.	FID(↓)	CLIP(↑)	IR(↑)
SDXL-turbo (1 step)	FP	16/16	-	-	17.15	0.2722	0.8631
	Naive PTQ	8/16	2×	1×	16.89	0.2740	0.8550
		4/16	4×	1×	301.49	0.1581	-2.2526
		8/8	2×	4×	103.96	0.1478	-1.7446
		4/8	4×	8×	358.894	0.1242	-2.2815
	Q-Diffusion	8/16	2×	1×	16.97	0.2735	0.8588
		4/16	4×	1×	22.58	0.2685	0.6847
		8/8	2×	4×	76.18	0.1772	-1.3112
		4/8	4×	8×	118.93	0.1662	-1.6353
	MixDQ(Ours)	4/16	4×	1×	17.23	0.2693	0.8254
		3.66/16	4.4×	1×	17.40	0.2682	0.7528
		8/8	2×	4×	17.03	0.2703	0.8415
		5/8	3.2×	8×	17.23	0.2697	0.8307
		4/8	4×	8×	17.68	0.2698	0.7822
LCM-lora (4 steps)	FP	16/16	-	-	25.56	0.2570	0.2122
	Naive PTQ	8/8	2×	4×	23.36	0.2548	0.0517
		4/8	4×	8×	87.36	0.2055	-1.6160
	Q-Diffusion	8/8	2×	4×	23.92	0.2561	0.1875
		4/8	4×	8×	57.73	0.2280	-1.1863
	MixDQ(Ours)	8/8	2×	4×	22.54	0.2552	0.1573
		4/8	4×	8×	33.48	0.2403	-0.6732

olution of 512×512. We adopt three different metrics: FID for fidelity, CLIP Score for image-text alignment, and ImageReward for human preference. The metrics are calculated on all 40,504 prompts. For the baseline methods, “naive PTQ” and “Q-diffusion,” we use a uniform bit-width for all layers. For MixDQ, we calculated the average bit-width weighted by each layer’s parameter size. Following [29], we measure the theoretical computational savings (“Compute Opt.”) in Bit Operations (BOPs). The “Storage Opt.” represents the model size reduction. The comparison of the performance and resource consumption of MixDQ quantization is presented in Tab. 1 and Fig. 6.

We experiment with two quantization settings: the “weight-only” scheme which seeks larger compression rate of model size, and the “normal” scheme focuses on both storage saving and latency speedup. During activation quantization, we observed that certain layers retain sensitivity and cannot be quantized to 8 bits without sacrificing performance. To address this, we choose to retain 1%

Table 2: Comparison with other quantization techniques for W8A8 on COCO annotations subset. The ‘‘CLIP’’ and ‘‘IR’’ denotes CLIP Score and ImageReward metric. The metrics are evaluated with the first 1024 prompts in COCO annotations. The ‘‘MP-only’’ represents adopting the mixed precision only.

Model	Method	FID(↓)	CLIP(↑)	IR(↑)
	FP16	84.51	0.26	0.84
SDXL-turbo (1 step)	Naive PTQ	165.92 (+81.4)	0.15 (-0.11)	-1.72 (-2.56)
	PTQD [5]	340.74 (+256.2)	0.12 (-0.14)	-2.28 (-3.12)
	Q-Diffusion [13]	149.15 (+64.6)	0.16 (-0.10)	-1.69 (-2.53)
	EDA-DM [16]	137.98 (+53.5)	0.16 (-0.10)	-1.71 (-2.55)
	MP-only [54]	114.70 (+30.2)	0.15 (-0.11)	-0.61 (-1.45)
	Non-Uniform (FP8) [23]	101.73 (+17.2)	0.24 (-0.01)	0.16 (-1.45)
	MixDQ	83.39 (-1.12)	0.27 (+0.01)	0.84 (+0.00)

of the most sensitive layers based on metric-decoupled sensitivity. Fig. 7 shows that it introduces minimal overhead but ensures the preservation of performance.

As evident from Fig. 6 and Tab.1, for the SDXL-turbo model, the baseline quantization methods can only maintain image quality with W8A16. The naive PTQ’s FID increases drastically from around 17.15 to 103.96 and 301.49 for W8A8 and W4A16, respectively. Q-diffusion, with the assistance of Adaround [24], manages to preserve performance for W4A16 but still fails at W8A8, generating images with ‘‘oil-painting’’-like quality degradation. For W4A8, both PTQ and Q-diffusion produce images that are hardly readable, exhibiting FID values exceeding 100, negative ImageReward, and CLIP Score below 0.2. In contrast, MixDQ generates images that are nearly identical to FP16 images, maintaining both the image content and quality. Even with W4A8 quantization, MixDQ incurs a 0.5 FID increase and a 2.5e-3 CLIP Score drop. As for LCM-Lora, the baseline methods encounter fewer failures since they involve 4 iteration steps. Nevertheless, MixDQ consistently outperforms them for all metrics.

We further compare MixDQ’s performance with other quantization techniques for W8A8 in Tab. 2. For PTQD [5], we conduct linear regression to determine K . However, we observe a relatively lower linear correlation (0.59) compared to the original paper. In the case of 1-step generation, the sampling process is deterministic, the ‘‘uncorrelated noise correction’’ inapplicable, and only the ‘‘correlated noise correction’’ is adopted. Despite being effective for multi-step quantization, we observe that PTQ causes notable quality degradation for the challenging 1-step quantization. EDA-DM [16] improves upon the quantization parameter tuning in Q-Diffusion [13] with layer-wise reconstruction. However, only a marginal improvement is witnessed. The term ‘‘MP-only’’ refers to ‘‘mixed precision only’’ with the SQNR-based method [54], where due to sub-optimal quantization sensitivity analysis, a significant performance degradation is still evident. The ‘‘Non-uniform’’ refers to the use of the FP8 (Floating-point 8-bit) format [23], which introduces exponential bits to handle a larger dynamic range of data distribution. While this notably improves performance, it still falls




Fig. 6: The FID with respect to memory cost of MixDQ and baseline quantization methods, with corresponding generated images. MixDQ achieves lossless quantization, whereas baseline methods fail to generate readable images.

short of the FP16 baseline. In contrast, MixDQ achieves W8A8 quantization without any performance loss.

4.3 Hardware Resource Savings

Memory footprint reduction. Fig. 7 (a) shows the GPU memory usage of MixDQ and the FP16 baseline on UNet inference. MixDQ can reduce the memory footprint from two aspects. Firstly, quantizing the model weights leads to smaller allocated memory to store the UNet model. Secondly, quantizing the activations saves allocated memory to store residual connection activations. Combining the two benefits, we can effectively reduce the peak memory footprint by $1.87\times$, $3.03\times$ and $3.03\times$ under the W8A8, W4A16 and W4A8, respectively.

Speedup of model inference. Fig. 7 (a) shows the latency of the UNet model inference on Nvidia RTX 4080. We use the W8A8 scenario to showcase the speedup. MixDQ W8A8 can accelerate the inference of UNet by $1.52\times$ over the FP16 baseline. 7 (b) gives a breakdown of the inference latency. There are three types of layers in the model: quantizable layers, including linear and convolution layers; non-quantizable layers, such as normalization and non-linear activation layers; and quantization layers, that perform the FP16-to-int8 conversion. The non-quantizable layers have the same latency in both the baseline and MixDQ. The quantizable layers are accelerated by $1.97\times$, approximately the same ratio between INT8 and FP16 hardware peak throughput on RTX4080 ($2\times$). MixDQ requires quantization layers for converting activations, but even with this overhead, the end-to-end acceleration still reaches $1.52\times$.




Fig. 7: The illustration of MixDQ’s hardware resource savings. Left: The comparison of efficiency and performance under different MixDQ mixed precision configurations. Right: (a) MixDQ’s optimization of peak memory, (b) MixDQ’s latency breakdown under W8A8.

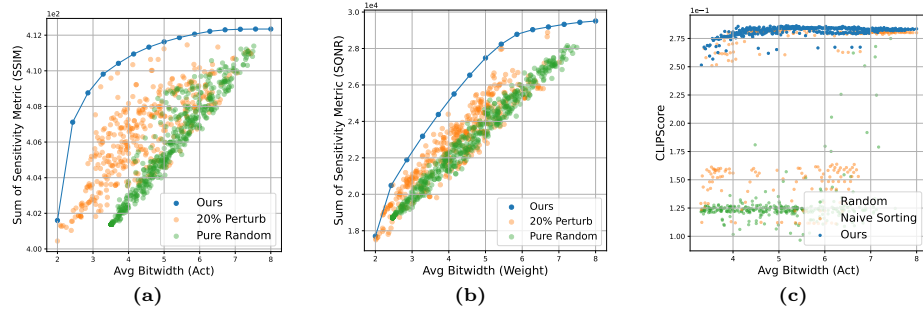


Fig. 8: The Pareto frontier of mixed precision configurations. The x-axis represents the averaged bit-width, the y-axis of Subfigures (a), (b), (c) presents the sensitivity metrics (SQNR & SSIM), and the final evaluation metric (CLIPScore).

5 Analysis

5.1 Analysis of Pareto Frontier

The Pareto frontier [48] is often applied to present the accuracy-efficiency trade-off. We present the scatter plot of three metrics with respect to averaged bit-width for SDXL-turbo in Fig. 8. In Fig. 8 (a),(b), we examine the effectiveness of integer planning by choosing the sensitivity metric (SQNR and SSIM) as the y-axis. As can be seen, compared with the “Pure random” baseline that randomly chooses bit-width, the “Ours” configuration achieves a significantly better trade-off. Furthermore, we randomly “perturb” 20% of layers’ bit-width in our configuration by replacing it with another bit-width and observing lower metric scores, denoting the superiority of our bit-width allocation. Furthermore, to demonstrate that our integer programming is non-trivial, we design a “Naive Sorting” baseline that repeatedly lowers the bit-width of the least sensitive layer to fit the average bit-width. We compare the final evaluation metric (CLIP Score) of “Ours” result, “Naive Sorting,” and “Random” in Fig. 8 (c). As can be seen,

our result achieves a superior trade-off, demonstrating not only the effectiveness of the integer programming but also the accuracy of our acquired sensitivity.

5.2 Analysis of BOS-aware Quantization

We further conduct experiments to demonstrate the effectiveness of BOS-aware Quantization. As discussed in Section 3.1, the first BOS token in CLIP embedding exhibits significantly larger values compared to others. While some specialized quantization techniques, such as non-uniform quantization, could address this outlier issue, introducing FP8 quantization for text embedding still incurs notable performance degradation (See Fig. 9). We also visualize the distribution of the BOS token and observe large variance along the channel dimension. Certain outlier channels have values $800 \times$ greater than other channels, as depicted in Fig. 10. Additionally, Fig. 10 illustrates that in the T5 [11] text encoder adopted by more recent papers [1, 50], although the first token is no longer an outlier, the channel imbalance problem persists. It's important to highlight that **BOS-aware quantization suggests that the quantization of the fixed BOS feature could be skipped**, resolving all its quantization issues beyond mere "outliers".



Fig. 9: The comparison of different quantization techniques for text embedding. The content preservation notably improves for FP8 quantization, however, is still inferior to BOS-aware

5.3 Ablation Studies

We conduct ablation studies by gradually incorporating MixDQ techniques into the W8A8 quantized SDXL-turbo model. As illustrated in Fig. 11 and Tab.3, the generated images improve from severe degradation to surpassing FP ones. The Pareto frontier in Fig. 8 could also assist in proving the effectiveness. We discuss the effectiveness of the proposed techniques individually as follows:

Effectiveness of BOS-aware Quantization. When adopting naive W8A8 PTQ for SDXL-turbo, as shown in Fig. 11 and Tab. 3, the generated quality

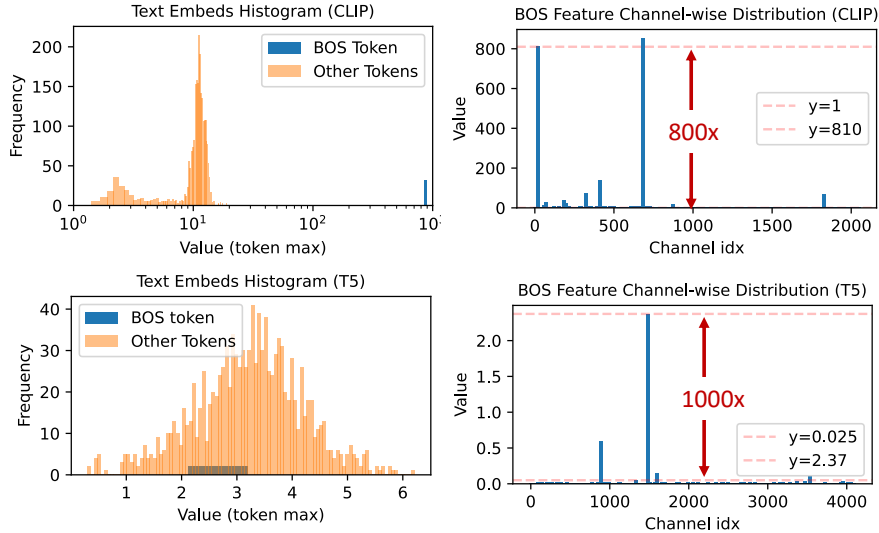


Fig. 10: The visualization of distribution for T5/CLIP text embeddings. Left: the histogram for BOS and other tokens. Right: the channel-wise distribution of the BOS token.

degrades, and the content changes significantly. After introducing BOS-aware quantization, the image content is recovered, and the metric values improve significantly (FID: 103.95 \rightarrow 31.65, CLIP Score: 0.1478 \rightarrow 0.2652).

Effectiveness of Metric-Decoupled Sensitivity. In Fig. 11 (c), we present the result of the bit-width allocation based on SQNR without our metric-decoupled scheme. It can be observed that compared with Fig. 11 (b), the quality worsened (FID: 31.65 \rightarrow 37.35, CLIP Score: 0.2652 \rightarrow 0.2624) after applying mixed precision. This reveals the insufficient accurate sensitivity and highlights the importance of our metric-decoupled technique. Also, Fig. 8 (c) helps prove the metric-decoupled sensitivity has a high correlation with final evaluation metrics.

Effectiveness of Mixed Precision. As shown in Tab.3 and Fig. 11 (d), when applying mixed precision with our metric-decoupled sensitivity analysis, MixDQ achieves lossless quantization with generated images nearly identical to FP and acquires similar metric values (FID: 17.03 vs. 17.15, CLIP Score: 0.2703 vs. 0.2722). Fig. 8 (a),(b) also illustrates that our integer programming strikes the optimal performance-efficiency trade-off on the Pareto frontier.

5.4 Analysis of Quantization Method

Takeaway Knowledge of Layer Sensitivity. We conclude key takeaways of layer sensitivity for text-to-image diffusion models as follows: (1) The CrossAttn & FFN affects image content while the SelfAttn & Conv affect quality. This is illustrated in Fig. 5 and Sec. 3.2. (2) The activation for “conv_in” and weight

Table 3: Ablation studies on MixDQ techniques. By gradually incorporating the proposed techniques on SDXL-turbo with W8A8 quantization, the generated images exhibit improvements from failure to surpassing the FP16 baseline.

<i>BOS-aware</i>	<i>Mixed-Precision</i>	<i>Metric-Decouple</i>	Bit-width(W/A)	FID (\downarrow)	CLIP (\uparrow)
-	-	-	FP16	17.15	0.2722
-	-	-	8/8	103.96	0.1478
✓	-	-	8/8	31.65	0.2652
✓	✓	-	8/8	37.35	0.2624
✓	✓	✓	8/8	17.03	0.2703




Fig. 11: The illustration of ablation studies on MixDQ techniques. From left to right, as techniques are progressively added, notable improvement in both generated image quality and alignment is witnessed.

for “conv_out” are highly sensitive compared with other convolution layers. (3) The sensitivity of activation for “to_k/v” layers notably decreases after BOS-aware quantization, but their weight’s sensitivity remain high. (4) In contrast to multi-step models, time embedding is not sensitive for few-step models. This fits intuition since distillation enables the network to denoise from any timestep.

Properties of MixDQ Quantization. (1) **Efficient Quantization:** Previous quantization methods require costly optimization of extensive AdaRound and scaling factor parameters, which takes tens of GPU hours. This process needs to be conducted repeatedly for different bit-width configurations. MixDQ adopts the naive minmax quantization without bells and whistles. After obtaining sensitivity, it only requires several minutes to acquire the entire Pareto frontier. (2) **Hardware Friendly.** When designing quantization schemes, we prioritize making them hardware-friendly to minimize the hardware overhead for supporting them. This includes employing tensor-wise scaling for activations and utilizing output-channel-wise scaling for weights. The granularity of mixed precision is set at the layer level, enabling the utilization of existing kernels for hardware resource savings and avoiding the cost of implementing customized kernels. (3) **Flexibility.** The framework of MixDQ is flexible and extensible, allowing it to be combined with other quantization techniques such as efficient QAT [46].

6 Conclusions

We propose a mixed-precision quantization method: MixDQ, which consists of three steps. Firstly, BOS-aware text embedding quantization addresses the highly sensitive layers. Secondly, metric-decoupled sensitivity analysis is introduced to consider preserving both the image quality and content. Finally, an integer programming is conducted for bit-width allocation. MixDQ achieves W4A8 with negligible performance loss, and practical hardware speedup.

Acknowledgements

This work was supported by National Natural Science Foundation of China (No. 62325405, 62104128, U19B2019, U21B2031, 61832007, 62204164), Tsinghua EE Xilinx AI Research Fund, and Beijing National Research Center for Information Science and Technology (BNRist). We thank for all the support from Infinigence-AI.

Appendix

A Application example of MixDQ

As discussed in the Sec. 3.3 and the Sec. 4 in the main paper, we introduce mixed-precision bit-width allocation methods based on integer programming. This approach proves effective in obtaining points on the Pareto frontier. In this section, we further provide a detailed description of an application case: **“How to determine the optimal bit-width configuration given a certain memory budget”?**

We provide a detailed description for MixDQ in this application in Algorithm 1. For simplicity, the algorithm only describes the bit-width allocation for weights, a similar process could be applied for activation bit-width allocation. Firstly, we perform sensitivity analysis after BOS-aware quantization (discussed in the main paper Sec. 3.1). Following the idea of “metric-decouple”, we categorize all layers of the model into content-related layers and quality-related layers. Specifically, the content-related layers include cross-attention layers and feed-forward networks (FFNs), while the quality-related layers encompass self-attention layers, convolutions, and other remaining layers (further discussed in Appendix D). Then we calculate the sensitivity separately for content/quality-related layers weight/activation, generating 4 groups of layer-wise relative sensitivities ($\mathcal{S}_{\text{content}}$ and $\mathcal{S}_{\text{quality}}$ in Algorithm 1).

Subsequently, we allocate the bit-width under budget constraints based on the acquired sensitivity. Specifically, we describe the budget with averaged bit-width (e.g., W5A8) multiplied by the number of parameters for weight and activation, respectively (the \mathcal{B}_{all}). We linearly scan through M nearby budgets \mathcal{B}_j within range $[\mathcal{B}_{\text{all}} - \Delta\mathcal{B}, \mathcal{B}_{\text{all}}]$ to ensure we get the optimal result. Then, we introduce a hyperparameter K (we choose linear space from 0.45 to 1.36 for weights, and 0.94 to 1.09 for activation) to describe the ratio between budgets assigned for the content-related and quality-related groups. It satisfies: $\mathcal{B}_{j,\text{all}} = \mathcal{B}_{\text{quality}}^{j,K} + \mathcal{B}_{\text{content}}^{j,K}$. After acquiring the budget for content or quality-related layers, integer programming could be used to get the mixed precision configuration: $\{b_{l,\text{content}}^{j,K}\}$, which is a layer-wise choice of candidate bit-width 2,4,8, the l indexes the layers. The aforementioned process is iterated for each $\mathcal{B}_{j,\text{model}}$ and K , and the optimal configurations are determined through a rapid evaluation of the generated images.

B Description of Challenges for Few-step Diffusion Quantization

As mentioned in the main paper Sec. 1 and Fig. 1, Fig. 2, the few-step diffusion model quantization faces the challenges of (1) The few-step diffusion models are more sensitive to quantization than multi-step ones. (2) The quantization’s effect on image content causes degradation in image-text alignment. In this section, we provide a detailed discussion of these two challenges.

Algorithm 1: MixDQ Bit-width Allocation Algorithm

Input: Sensitivity of content-related layers: S_{content} , Sensitivity of quality-related layers: S_{quality} , Parameter size of content-related layers : $\mathcal{P}_{\text{content}}$, Parameter size of quality-related layers : $\mathcal{P}_{\text{quality}}$, The target budget of the whole model, content and quality-related layers: $\mathcal{B}_{\text{all}}, \mathcal{B}_{\text{quality}}, \mathcal{B}_{\text{content}}$, The candidate budgets: $\{\mathcal{B}_j\} \in [\mathcal{B}_{\text{all}} - \Delta\mathcal{B}, \mathcal{B}_{\text{all}}]$ of length M , The candidate ratios: $\{k_i\}$ of length N , The FP model \mathcal{M} , the quantized model with bit-width configuration $\{b_l\}$: $\mathcal{M}_q^{\{b_l\}}$ **Output:** The mixed precision configuration $\{b_l\}$.

```

for  $j = 1$  to  $M$  do
  for  $k = k_0$  to  $k_{N-1}$  do
    // Split the budgets for content and quality-related layers.
     $\mathcal{B}_{\text{content}}^{j,k}, \mathcal{B}_{\text{quality}}^{j,k} = \text{SplitBudget}(\mathcal{B}_{j,\text{all}}, \mathcal{P}_{\text{content}}, \mathcal{P}_{\text{quality}}, k)$ ;
    // perform integer programming respectively
     $\{b_{l,\text{content}}^{u,k}\} = \text{IntegerProgramming}(\mathcal{B}_{\text{content}}^{u,k}, S_{\text{content}}, \mathcal{P}_{\text{content}})$ ;
     $\{b_{l,\text{quality}}^{u,k}\} = \text{IntegerProgramming}(\mathcal{B}_{\text{quality}}^{u,k}, S_{\text{quality}}, \mathcal{P}_{\text{quality}})$ ;
    // get the mixed precision configuration for the whole model
     $\{b_{l,\text{all}}^{u,k}\} = \{b_{l,\text{quality}}^{u,k}, b_{l,\text{content}}^{u,k}\}$ ;
    // infer with the quantized model
     $\text{imgs} = \mathcal{M}_Q^{\{b_{l,\text{all}}^{u,k}\}}(\text{prompts})$ ;
    // evaluate the generated images
     $\text{score}_{u,k} = \text{eval}(\text{imgs})$ ;

  // choose the optimal configuration with best score
   $\{b_l\} = \underset{\{b_{l,\text{all}}^{u,k}\}}{\text{argmax}}(\text{score}_{u,k})$ ;
return  $\{b_l\}$ 

```

B.1 Few-step Diffusion Models are More Sensitive to Quantization

As depicted in Fig. 1 of the main paper and discussed in Section 1, “the 2-step model exhibits significantly less degradation compared to the SDXL-turbo 1-step model”. We delve into the reasons behind this phenomenon. Examining the challenge in Fig. 12, we contrast the generation processes of the 2-step and 1-step models.

In the 2-step generation, the second step incorporates an image with both Gaussian noise and quantization noise as input. Recent literature on diffusion-based image editing, as highlighted in [21, 26], indicates that diffusion models can effectively recover content from partially disrupted image input (image inpainting). Consequently, **the additional denoising steps (2nd or more) facilitate the reconstruction of quantization errors**, resulting in improved image quality.

In contrast, the 1-step diffusion model lacks this recovery phase, making it more “vulnerable” to quantization noise.

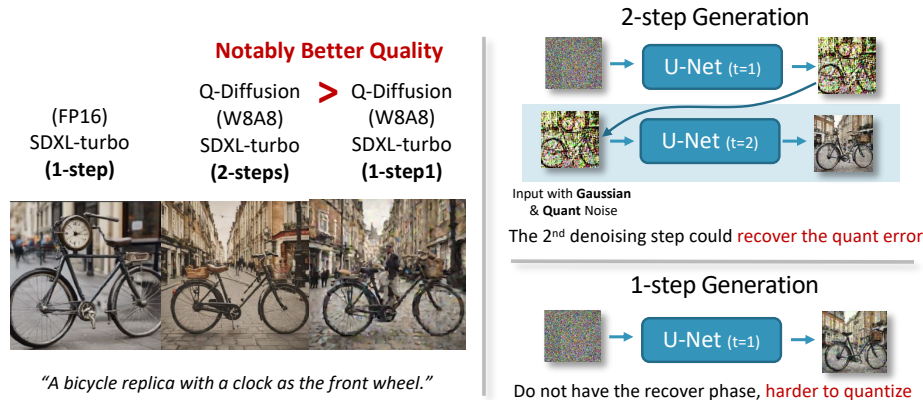


Fig. 12: The illustration of the few-step diffusion model is comparatively harder to quantize. Left: the 2-step SDXL-turbo quantized model exhibits notably better image quality than the 1-step one. Right: the explanation of why the 1-step model is harder to quantize.

B.2 Quantized Diffusion Models Lose Image-text Alignment

As outlined in Sec. 1 of the main paper, prior research predominantly concentrates on preserving image quality. Nevertheless, it’s crucial to recognize that quantization affects not just the image quality but also the content itself. The alteration in content can result in the degradation of image-text alignment. In Fig. 13, we provide additional examples of alignment degradation, illustrating

that such deterioration occurs not only in few-step models but also in less challenging multi-step diffusion models.

In the case of the prompt "A black Honda motorcycle parked in front of a garage," the W8A8 quantized SDXL-turbo 1-step model fails to include the "in front of garage" in the text instruction and instead generates an image of a man riding a bicycle. With the multi-step SDXL model, the output deviates even further, presenting a yellow truck instead of the specified black motorcycle.

Similarly, for the prompt "A room with blue walls and a white sink and door," both the W8A8 quantized models generate an image of a white room while omitting the "blue walls" component. In contrast, our MixDQ W8A8 produces images with **significantly improved image-text alignment**. They fulfill all components of the prompt, closely resembling the images generated by the FP16 model.



Fig. 13: The illustration of the "Image-text Alignment loss" challenge for diffusion quantization. The Q-diffusion witnesses alignment loss (lacking of components described in the text instruction) for both the multi-step and few-step models.

C Detailed Analysis of Hardware Experiments

In this section, we provide a detailed analysis of the hardware resource savings for different bit-width configurations and conclude our findings as follows:

- (1) Illustrated in Fig. 7b in the main paper Sec. 4.3, **the W8A8 bar displays an almost negligible red segment**, corresponding to the "1% most sensitive activation layers retained from FP16" as discussed in Sec. 4.2 of the main paper. This component introduces minimal overhead while significantly contributing to the preservation of image quality.

(2) Observing the table in Fig. 7 in the main paper Sec. 4.3, we note that **W4A16 quantization solely achieves memory savings, while the latency remains consistent with the FP16 baseline.** The reduction in memory primarily stems from the decreased size of the model parameters, which are transferred to the GPU and serve as the "static memory cost." However, the W4A16 model continues to utilize FP16 computation, thus maintaining similar latency.

(3) Observing the table in Fig. 7 of the main paper Sec. 4.3, **W4A16 and W4A8 exhibit similar memory optimization.** The key distinction lies in the fact that W4A8 quantizes the activations, thereby reducing the memory cost through saved activations, primarily affecting the shortcut feature for U-Net. However, we discover that the occupied size by these activations is relatively small, amounting to less than 100MB compared to the model size of 6GB. Consequently, the additional savings attributed to activation quantization are not evident.

(4) Observing the table in Fig. 7 of the main paper Sec. 4.3, **both W4A8 and W8A8 demonstrate a noteworthy $1.52\times$ latency speedup compared to the FP16 baseline.** This acceleration is attributed to the utilization of INT8 GPU kernels. It's worth noting that we haven't implemented the INT4 operators, making W4A8 and W8A8 quite similar in terms of latency improvement. However, W4A8 has the potential to achieve even higher latency speedup.

D Justification of Layer Grouping in Metric Decouple

As mentioned in the main paper Sec. 3.2, we split the layers into two groups based on their effect on image content and quality. Specifically, the cross-attention and feed-forward layers are regarded as the "content-related layers", and the convolution and self-attention layers as the "quality-related layers". In this section, we present experimental results to **demonstrate the rationale of the layer grouping.**

Firstly, we summarize 5 layer types for the text-to-image diffusion model: cross-attention layers, feed-forward neural networks (FFNs), convolutions, self-attention layers, and other miscellaneous layers. We independently quantize the weight and activation into 8-bit for each layer type, and compute their effects on the SSIM score. As seen in Fig. 14, the cross-attention and feed-forward layers have a significantly larger influence (0.8 and 0.05) on the image content compared with other layer types (less than 0.01). Therefore, we choose them as the "content-related layers", and the rest as "quality-related" layers.

E Qualitative Results

In this section, we provide more qualitative results to demonstrate the effectiveness of our MixDQ, especially in terms of **preserving the "image-text**




Fig. 14: SSIM Score when quantizing a certain group. Cross-attention (CrossAttn) and Feed-Forward Networks (FFNs) are the two types of layers that have the greatest impact on content.

alignment”, As can be seen from Fig. 15, even under W4A8, our method generates images with visual quality and image-text alignment similar to the FP16 model generated ones.

Fig. 16 highlights the superior text-image alignment achieved by MixDQ in comparison to baseline quantization methods. With W8A8 quantization, the Q-diffusion-generated images not only exhibit an "oil-painting" like blurring and tiny artifacts but also deviate from following the text instructions. For instance, it omits "in front of a garage" from the prompt and, instead, generates an image of a man riding a motorcycle. It also struggles to interpret “an older man skiing”; instead, it generates a cowboy riding a strange two-headed horse. Moreover, the quantized model loses the capability to adhere to explicit color instructions like "blue walls" and "gray and white kitten." In contrast, with the more challenging W4A8 setting, our MixDQ quantized model produces images with nearly identical content compared to the FP16 baseline.

F Limitations and Future Directions

We introduce MixDQ, a mixed-precision quantization method that handles both the imbalance sensitivity and alignment degradation problems for diffusion quantization. We summarize the current limitations and potential future improvements as follows:

- **Specialized Quantization Techniques for other Sensitive Layers.** In MixDQ, we pinpoint the bottleneck in text embedding quantization and craft BOS-aware quantization specifically tailored for it. Nevertheless, there are



"A cinematic shot of a baby racoon wearing an intricate italian priest robe."

Fig. 15: The comparison between images generated by full precision model and quantized one. Text-guided image generation from our W4A8 quantized SDXL-turbo with a fixed random seed.

other highly sensitive layers, such as "conv-in" and "conv-out," that could also gain advantages from the design of specialized quantization techniques.

- **Combine with Improved Quantization Techniques.** In MixDQ, we employ naive min-max quantization as the quantization method. Its performance can be enhanced by incorporating advanced quantization techniques, such as Adaround [24], or by introducing quantization-aware training [10].
- **Lower Bit-width.** In MixDQ, we opt for the 2, 4, and 8 as candidate bit-widths. When coupled with efficient lower-bit quantization methods, the mixed-precision bit-width allocation remains compatible with lower-bit-width combinations.
- **More Hardware Supports.** As discussed in Appendix C, we only utilize the INT8 GPU kernels for now, which restricts the performance for W4 quantization. The newest Nvidia TensorCore [20] supports INT4 computing, which could be further utilized for better efficiency.

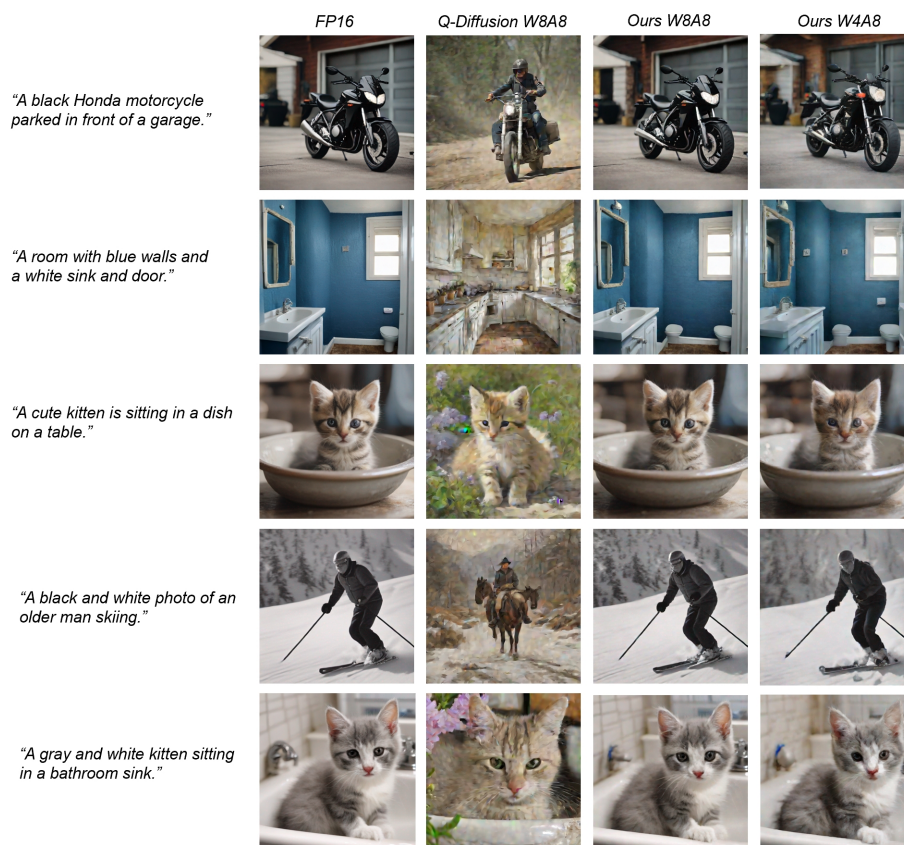


Fig. 16: The comparison between images generated by quantized model. Our quantization scheme better maintains the content and quality of the images.

References

1. Chen, J., Yu, J., Ge, C., Yao, L., Xie, E., Wu, Y., Wang, Z., Kwok, J., Luo, P., Lu, H., et al.: Pixart- α : Fast training of diffusion transformer for photorealistic text-to-image synthesis. arXiv preprint arXiv:2310.00426 (2023)
2. Chen, Y.H., Sarokin, R., Lee, J., Tang, J., Chang, C.L., Kulik, A., Grundmann, M.: Speed is all you need: On-device acceleration of large diffusion models via gpu-aware optimizations. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) pp. 4651–4655 (2023), <https://api.semanticscholar.org/CorpusID:258298971>
3. Dong, Z., Yao, Z., Gholami, A., Mahoney, M.W., Keutzer, K.: Hawq: Hessian aware quantization of neural networks with mixed-precision. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) pp. 293–302 (2019), <https://api.semanticscholar.org/CorpusID:148571720>
4. He, Y., Liu, J., Wu, W., Zhou, H., Zhuang, B.: Efficientdm: Efficient quantization-aware fine-tuning of low-bit diffusion models. ArXiv [abs/2310.03270](https://arxiv.org/abs/2310.03270) (2023), <https://api.semanticscholar.org/CorpusID:263672060>
5. He, Y., Liu, L., Liu, J., Wu, W., Zhou, H., Zhuang, B.: Ptdq: Accurate post-training quantization for diffusion models. ArXiv [abs/2305.10657](https://arxiv.org/abs/2305.10657) (2023), <https://api.semanticscholar.org/CorpusID:258762678>
6. Hessel, J., Holtzman, A., Forbes, M., Bras, R.L., Choi, Y.: Clipscore: A reference-free evaluation metric for image captioning. arXiv preprint arXiv:2104.08718 (2021)
7. 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 **30** (2017)
8. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
9. Huang, Y., Gong, R., Liu, J., Chen, T., Liu, X.: Tfmq-dm: Temporal feature maintenance quantization for diffusion models. ArXiv [abs/2311.16503](https://arxiv.org/abs/2311.16503) (2023), <https://api.semanticscholar.org/CorpusID:265466808>
10. Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A.G., Adam, H., Kalenichenko, D.: Quantization and training of neural networks for efficient integer-arithmetic-only inference. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition pp. 2704–2713 (2017), <https://api.semanticscholar.org/CorpusID:39867659>
11. Kale, M., Rastogi, A.: Text-to-text pre-training for data-to-text tasks. arXiv preprint arXiv:2005.10433 (2020)
12. Li, S., Ning, X., Hong, K., Liu, T., Wang, L., Li, X., Zhong, K., Dai, G., Yang, H., Wang, Y.: LLM-MQ: Mixed-precision Quantization for Efficient LLM Deployment. In: *NeurIPS 2023 Efficient Natural Language and Speech Processing Workshop* (2023)
13. Li, X., Lian, L., Liu, Y., Yang, H., Dong, Z., Kang, D., Zhang, S., Keutzer, K.: Q-diffusion: Quantizing diffusion models. *ICCV* (2023)
14. Li, Y., Xu, S., Cao, X., Sun, X., Zhang, B.: Q-dm: An efficient low-bit quantized diffusion model. In: *Neural Information Processing Systems* (2023), <https://api.semanticscholar.org/CorpusID:268096292>
15. Lin, T.Y., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *ECCV* (2014), <https://api.semanticscholar.org/CorpusID:14113767>
16. Liu, X., Li, Z., Xiao, J., Gu, Q.: Enhanced distribution alignment for post-training quantization of diffusion models. arXiv preprint arXiv:2401.04585 (2024)

17. 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)
18. Luo, S., Tan, Y., Patil, S., Gu, D., von Platen, P., Passos, A., Huang, L., Li, J., Zhao, H.: Lcm-lora: A universal stable-diffusion acceleration module. ArXiv **abs/2311.05556** (2023), <https://api.semanticscholar.org/CorpusID:265067414>
19. fast maintainers, S.: Stable-fast, <https://github.com/chengzeyi/stable-fast>
20. Markidis, S., Chien, S.W.D., Laure, E., Peng, I.B., Vetter, J.S.: Nvidia tensor core programmability, performance & precision. 2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW) pp. 522–531 (2018), <https://api.semanticscholar.org/CorpusID:3887305>
21. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: Sdedit: Guided image synthesis and editing with stochastic differential equations. In: International Conference on Learning Representations (2021), <https://api.semanticscholar.org/CorpusID:245704504>
22. Meng, C., Rombach, R., Gao, R., Kingma, D., Ermon, S., Ho, J., Salimans, T.: On distillation of guided diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14297–14306 (2023)
23. Micikevicius, P., Stosic, D., Burgess, N., Cornea, M., Dubey, P., Grisenthwaite, R., Ha, S., Heinecke, A., Judd, P., Kamalu, J., et al.: Fp8 formats for deep learning. arXiv preprint arXiv:2209.05433 (2022)
24. Nagel, M., Amjad, R.A., van Baalen, M., Louizos, C., Blankevoort, T.: Up or down? adaptive rounding for post-training quantization. ArXiv **abs/2004.10568** (2020), <https://api.semanticscholar.org/CorpusID:216056295>
25. Nagel, M., Fournarakis, M., Amjad, R.A., Bondarenko, Y., van Baalen, M., Blankevoort, T.: A white paper on neural network quantization. ArXiv **abs/2106.08295** (2021), <https://api.semanticscholar.org/CorpusID:235435934>
26. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)
27. NVIDIA: Nsight Systems. <https://docs.nvidia.com/nsight-systems/index.html>, <https://docs.nvidia.com/nsight-systems/index.html>
28. Nvidia: Nvidia cutlass release v3.4 (2024), <https://github.com/NVIDIA/cutlass>
29. Pandey, N.P., Nagel, M., van Baalen, M., Huang, Y.R., Patel, C., Blankevoort, T.: A practical mixed precision algorithm for post-training quantization. ArXiv **abs/2302.05397** (2023), <https://api.semanticscholar.org/CorpusID:256808627>
30. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Muller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. ArXiv **abs/2307.01952** (2023), <https://api.semanticscholar.org/CorpusID:259341735>
31. PyTorch: PyTorch Memory Management (2023), <https://pytorch.org/docs/stable/notes/cuda.html#memory-management>
32. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)

33. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: ICML. pp. 8821–8831. PMLR (2021)
34. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10684–10695 (2022)
35. Ryu, H., Lim, S., Shim, H.: Memory-efficient personalization using quantized diffusion model. ArXiv [abs/2401.04339](https://arxiv.org/abs/2401.04339) (2024), <https://api.semanticscholar.org/CorpusID:266899926>
36. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans **29** (2016)
37. Salimans, T., Ho, J.: Progressive distillation for fast sampling of diffusion models. arXiv preprint [arXiv:2202.00512](https://arxiv.org/abs/2202.00512) (2022)
38. Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial diffusion distillation. ArXiv [abs/2311.17042](https://arxiv.org/abs/2311.17042) (2023), <https://api.semanticscholar.org/CorpusID:265466173>
39. So, J., Lee, J., Ahn, D., Kim, H., Park, E.: Temporal dynamic quantization for diffusion models. ArXiv [abs/2306.02316](https://arxiv.org/abs/2306.02316) (2023), <https://api.semanticscholar.org/CorpusID:259075274>
40. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: International Conference on Machine Learning. pp. 2256–2265. PMLR (2015)
41. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint [arXiv:2010.02502](https://arxiv.org/abs/2010.02502) (2020)
42. Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. In: International Conference on Machine Learning (2023), <https://api.semanticscholar.org/CorpusID:257280191>
43. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. arXiv preprint [arXiv:2011.13456](https://arxiv.org/abs/2011.13456) (2020)
44. Tang, S., Wang, X., Chen, H., Guan, C., Wu, Z., Tang, Y., Zhu, W.: Post-training quantization with progressive calibration and activation relaxing for text-to-image diffusion models. ArXiv [abs/2311.06322](https://arxiv.org/abs/2311.06322) (2023), <https://api.semanticscholar.org/CorpusID:265149927>
45. Team, G.O.T.D.: Or-tools. <https://github.com/google/or-tools> (Year Accessed)
46. Wang, H., Shang, Y., Yuan, Z., Wu, J., Yan, Y.: Quest: Low-bit diffusion model quantization via efficient selective finetuning. ArXiv [abs/2402.03666](https://arxiv.org/abs/2402.03666) (2024), <https://api.semanticscholar.org/CorpusID:267500241>
47. Wang, K., Liu, Z., Lin, Y., Lin, J., Han, S.: Haq: Hardware-aware automated quantization with mixed precision. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 8604–8612 (2018), <https://api.semanticscholar.org/CorpusID:102350477>
48. Wikipedia: Pareto front (2024), https://en.wikipedia.org/wiki/Pareto_front, accessed: March 3, 2024
49. Wikipedia contributors: Structural similarity index measure. https://en.wikipedia.org/wiki/Structural_similarity_index_measure, accessed on: February 28, 2024
50. Wu, W., Li, Z., He, Y., Shou, M.Z., Shen, C., Cheng, L., Li, Y., Gao, T., Zhang, D., Wang, Z.: Paragraph-to-image generation with information-enriched diffusion model. arXiv preprint [arXiv:2311.14284](https://arxiv.org/abs/2311.14284) (2023)

51. Wu, X., Sun, K., Zhu, F., Zhao, R., Li, H.: Better aligning text-to-image models with human preference. arXiv preprint arXiv:2303.14420 (2023)
52. Xiao, G., Tian, Y., Chen, B., Han, S., Lewis, M.: Efficient streaming language models with attention sinks. ArXiv **abs/2309.17453** (2023), <https://api.semanticscholar.org/CorpusID:263310483>
53. Xu, J., Liu, X., Wu, Y., Tong, Y., Li, Q., Ding, M., Tang, J., Dong, Y.: Imagereward: Learning and evaluating human preferences for text-to-image generation. arXiv preprint arXiv:2304.05977 (2023)
54. Yang, Y., Dai, X., Wang, J., Zhang, P., Zhang, H.: Efficient quantization strategies for latent diffusion models. ArXiv **abs/2312.05431** (2023), <https://api.semanticscholar.org/CorpusID:266163100>
55. Zhao, W., Bai, L., Rao, Y., Zhou, J., Lu, J.: Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. arXiv preprint arXiv:2302.04867 (2023)
56. Zhao, Y., Xu, Y., Xiao, Z., Hou, T.: Mobicdiffusion: Subsecond text-to-image generation on mobile devices. ArXiv **abs/2311.16567** (2023), <https://api.semanticscholar.org/CorpusID:265466277>