

335 **A Experimental Details**

336 **A.1 Details of Diffusion Models**

337 In order to conduct a comprehensive analysis of the early-stage robustness in diffusion models,
 338 we employ five different diffusion models to generate various datasets, including CIFAR10 [24],
 339 FFHQ [27], CelebA-HQ [26], LSUN-bedrooms, and LSUN-churches [25]. To measure the Fréchet
 340 Inception Distance (FID), we utilize the torch-fidelity library ¹, following the methodology established
 341 in previous works [6, 8]. Subsequently, we apply the proposed RAQ method to perform high-
 342 resolution image generation tasks, encompassing both unconditional image generation for LSUN-
 343 bedrooms/LSUN-churches and conditional image generation using Stable Diffusion [8]. The details
 344 of the diffusion models utilized in these experiments are thoroughly presented in Table 2. Note that
 345 the calculation of σ_t in Eq. 2 is performed using η in Table 2 as specified in the following equation [6]:

$$\sigma_t = \eta \cdot \sqrt{(1 - \alpha_{t-1}) / (1 - \alpha_t)} \sqrt{1 - \alpha_t / \alpha_{t-1}} \quad (5)$$

Table 2: Implementation specifications of the diffusion models

| | CIFAR-10 | FFHQ/ CelebA-HQ | LSUN- Bedrooms | LSUN- Churches | Unconditional Generation |
|--------------|-------------------|--------------------|--------------------|--------------------|------------------------------------|
| Image Size | 32×32 | 256×256 | 256×256 | 256×256 | 512×512 |
| Architecture | DDIM ² | LDM-4 ³ | LDM-4 ³ | LDM-8 ³ | Stable Diffusion v1.4 ⁴ |
| Sampler | DDIM [6] | DDIM | DDIM | DDIM | PLMS [28] |
| Step Count | 100 | 200 | 200 | 200 | 50 |
| η | 0 | 0/1 | 1 | 0 | 0 |

346 **A.2 Details of Entropy Analysis**

347 In Section 3.1, we calculate the entropy of the latent variables x_t for each diffusion step. To facilitate
 348 this calculation, we transform the values of x_t into histogram bins. Specifically, we map x_t to a
 349 histogram bin using the following equation:

$$h(x_t) = \text{clamp}(\lfloor \frac{x_t}{256} \rfloor, -3, +3) \quad (6)$$

350 This equation ensures that the values of the histogram bins are constrained within the range of -3
 351 to $+3$, allowing us to effectively create a histogram with 256 bins. Once we have the histogram
 352 representation of x_t , we can calculate the entropy using the equation⁵:

$$H(X) = \sum_x -p(X) \log p(X) \quad (7)$$

353 **A.3 Intermediate Image Prediction during Reverse Diffusion Process**

354 In Fig. 1(a) and Fig. 2, we showcase the intermediate image prediction results during the reverse
 355 diffusion process, aiming to illustrate the characteristics of diffusion models. To visualize the
 356 prediction results, we utilize the following x_0 prediction of each diffusion step as stated in [6]:

$$x_0 = \frac{x_t - \sqrt{1 - \alpha_t} \epsilon_\theta(x_t, t)}{\sqrt{\alpha_t}} \quad (8)$$

¹<https://github.com/toshast/torch-fidelity>

²<https://github.com/ermongroup/ddim>

³<https://github.com/CompVis/latent-diffusion>

⁴<https://github.com/CompVis/stable-diffusion>

⁵Claude Elwood Shannon, "A Mathematical Theory of Communication", Bell system technical journal, 1948

357 **A.4 Details of Activation Quantization**

358 During the activation quantization process, we observed that the skip connections of ResBlocks,
359 the first convolutional layer responsible for transforming the latent variable into the input of the
360 denoising network, and the last convolutional layer responsible for transforming the output of the
361 denoising network (Fig. 1(b)) had a significant impact on the quality of the final image. However,
362 these components constitute a negligible fraction of the overall computation. Therefore, in order
363 to balance computational efficiency and image quality, the proposed RAQ adjust the activation bits
364 of the diffusion models to the desired bit precision while keeping the activation bits of these three
365 components fixed at 8 bits.

366 **B Additional Results**

367 **B.1 Activation Quantization across Diffusion Steps**

368 Fig. 9 complements the image generation results presented in Fig. 6. Fig. 9 showcases image genera-
369 tion with 4-bit activation quantization at different diffusion steps, alongside the image generation with
370 floating-point activations. The results of the activation quantization demonstrate a consistent trend
371 with the noise injection test (4). When 4-bit activation quantization is applied to the early stages, the
372 resulting images closely resemble those generated using floating-point activations, showcasing high
373 quality with minor shape variations. However, applying 4-bit activation quantization to the entire
374 diffusion process leads to a significant compromise in the generated image quality. This is primarily
375 due to the degradation in image quality caused by the quantization applied to the later diffusion steps.



Figure 8: Examples of 256×256 LSUN-Churches generation with FP32 activation or activation quantization across different diffusion steps. Example images with activation quantization are generated by applying 4-bit activation quantization in the target diffusion steps.

376 **B.2 Explanation on FID Improvement with Proposed RAQ on LSUN-Bedrooms**

377 In Table 1 of Section 5.1, we observe that the LSUN-Bedrooms images generated with the proposed
378 RAQ exhibit slightly better FID scores compared to Q-diffusion with W4A32 and W4A8. To
379 investigate the reason behind this FID improvement, we conduct a detailed comparison of the images
380 generated using full-precision activations and 4-bit activations in the early stage of the diffusion



Figure 9: Examples of 256×256 LSUN-Bedrooms generation with different activation precision.

381 process. We find that the models with full-precision activations sometimes generate images with
 382 complex structures that are not easily recognizable as bedrooms. However, when 4-bit activation
 383 quantization is applied to the early stage, it simplifies the complex structures and results in images that
 384 more closely resemble bedrooms. This observation suggests that the step-wise activation quantization
 385 strategy employed in the proposed RAQ method helps refine the generated images, leading to
 386 improved quality and better alignment with the target LSUN-Bedrooms dataset.

387 B.3 Activation Quantization of Stable Diffusion

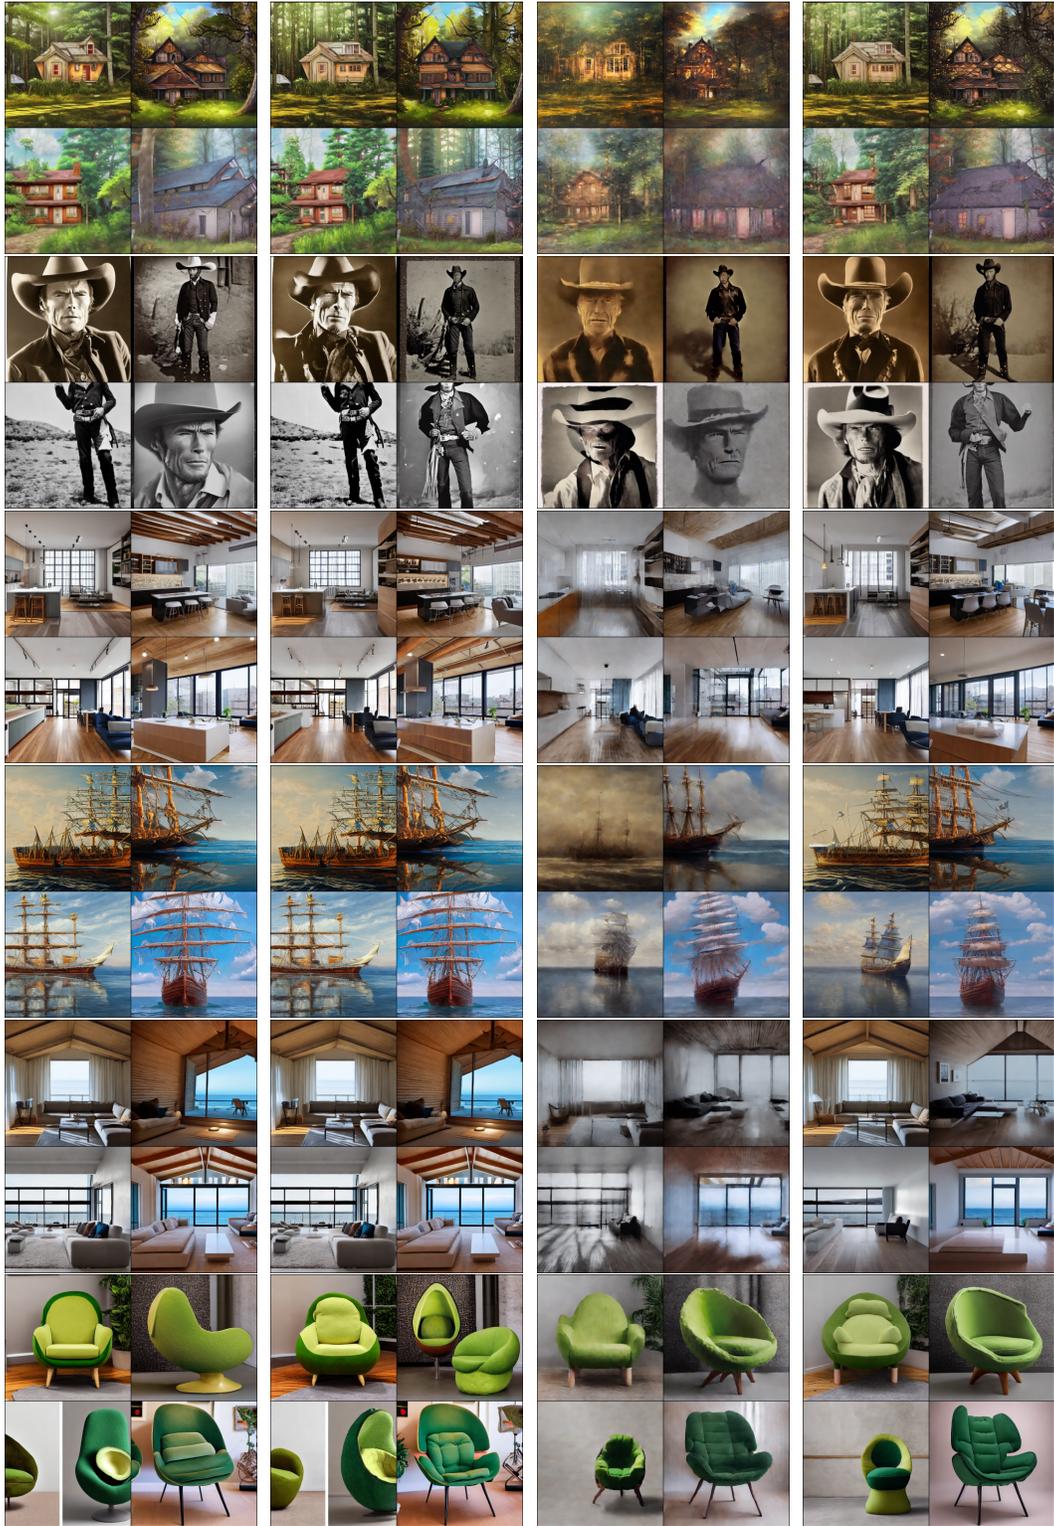
388 For the conditional image generation with Stable Diffusion, we utilize the prompt examples that are
 389 publicly available online^{6 7}. The prompts used in Section 5.2 are as follows:

- 390 1. a puppy wearing a hat
- 391 2. cluttered house in the woods in anime oil painting style*
- 392 3. Old photo of Clint Eastwood dressed as cowboy,1800s, centered, by professional photogra-
 393 pher, wide-angle lens, background saloon*
- 394 4. interior design, open plan, kitchen and living room, modular furniture with cotton textiles,
 395 wooden floor, high ceiling, large steel windows viewing a city Artstation and Antonio
 396 Jacobsen and Edward Moran, (long shot), clear blue sky, intricated details, 4k*
- 397 5. a tree on the hill, bright scene, highly detailed, realistic photo
- 398 6. a highly detailed, majestic royal tall ship on a calm sea,realistic painting, by Charles
 399 Gregory*
- 400 7. medium shot side profile portrait photo of the Takeshi Kaneshiro warrior chief, tribal panther
 401 make up, blue on red, looking away, serious eyes, 50mm portrait, photography, hard rim
 402 lighting photography -ar 2:3 -beta -upbeta
- 403 8. a picture of dimly lit living room, minimalist furniture, vaulted ceiling, huge room, floor to
 404 ceiling window with an ocean view, nighttim*
- 405 9. an armchair in the shape of an avocado, an armchair imitating an avocado*

406 In this section, we additionally present non-cherry-picked samples generated using Stable Diffusion
 407 with and without activation quantization. We use the prompts that are highlighted with asterisk (*).
 408 For each prompt, we generate four images to demonstrate the variety and quality of the generated
 409 results (Fig. 10 in the next page).

⁶<https://stablediffusion.fr/prompts>

⁷<https://mpost.io/best-100-stable-diffusion-prompts-the-most-beautiful-ai-text-to-image-prompts/>



(a) Full Precision (b) W4A8 (Q-Diffusion) (c) W4A6 (d) W4A6/8 (Proposed)

Figure 10: Text-guided 512×512 image generation results with Stable Diffusion.