

CADC: Content Adaptive Diffusion-Based Generative Image Compression

Xihua Sheng¹, Lingyu Zhu¹, Tianyu Zhang², Dong Liu², Shiqi Wang¹, Jing Wang^{3*}

¹City University of Hong Kong ²University of Science and Technology of China

³Central Media Technology Institute, Huawei

{xihsheng, shiqwang}@cityu.edu.hk, lingyuzhu-c@my.cityu.edu.hk,

zhangtianyu@mail.ustc.edu.cn, dongeliu@ustc.edu.cn, wangjing215@huawei.com

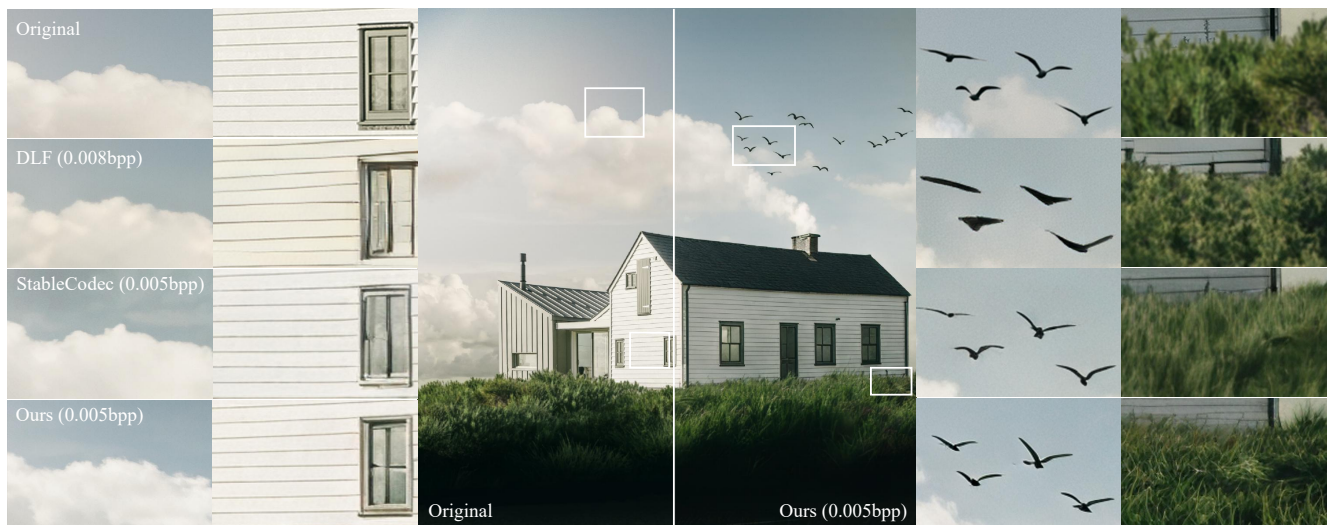


Figure 1. A qualitative comparison between our codec, StableCodec [66], and DLF [59] when compressing a 2K-resolution image of the test set of CLIC 2020 Professional [57] under ultra-low bitrate conditions. Our codec produces images with high visual quality, especially in regions with complex texture. In contrast, DLF and StableCodec exhibit noticeable artifacts, such as blurring and color shifting.

Abstract

Diffusion-based generative image compression has demonstrated remarkable potential for achieving realistic reconstruction at ultra-low bitrates. The key to unlocking this potential lies in making the entire compression process content-adaptive, ensuring that the encoder’s representation and the decoder’s generative prior are dynamically aligned with the semantic and structural characteristics of the input image. However, existing methods suffer from three critical limitations that prevent effective content adaptation. First, isotropic quantization applies a uniform quantization step, failing to adapt to the spatially varying complexity of image content and creating a misalignment with the diffusion model’s noise-dependent prior. Second, the information concentration bottleneck—arising from the di-

mensional mismatch between the high-dimensional noisy latent and the diffusion decoder’s fixed input—prevents the model from adaptively preserving essential semantic information in the primary channels. Third, existing textual conditioning strategies either need significant textual bitrate overhead or rely on generic, content-agnostic textual prompts, thereby failing to provide adaptive semantic guidance efficiently. To overcome these limitations, we propose a content-adaptive diffusion-based image codec (CADC) with three technical innovations: 1) an Uncertainty-Guided Adaptive Quantization (UGAQ) method that learns spatial uncertainty maps to adaptively align quantization distortion with content characteristics; 2) an Auxiliary Decoder-Guided Information Concentration (ADGIC) method that uses a lightweight auxiliary decoder to enforce content-aware information preservation in the primary latent channels; and 3) a Bitrate-Free Adaptive Textual Condition-

* Corresponding author

ing (BFATC) method that derives content-aware textual descriptions from the auxiliary reconstructed image, enabling semantic guidance without bitrate cost. Comprehensive experimental results show that our codec achieves state-of-the-art perceptual quality at ultra-low bitrates.

1. Introduction

Digital images account for a substantial portion of internet traffic, driving continuous demand for efficient compression technologies that reduce storage and transmission costs while maintaining visual quality. Traditional image codecs [9, 54, 56, 58] have been widely adopted for decades. To pursue higher compression performance, learned image compression [1, 3, 6, 7, 13, 16, 17, 19, 22–24, 27, 32, 41, 42, 44, 45, 48, 67] has emerged in recent years. However, at ultra-low bitrates, even state-of-the-art learned image codecs tend to produce reconstructions plagued by blurring and loss of fine details, as they primarily optimize for pixel-level signal fidelity rather than perceptual quality [8].

Generative image compression addresses this limitation by leveraging strong generative models to produce visually realistic reconstructions. Existing generative image codecs can be broadly categorized into three classes: Generative Adversarial Network (GAN)-based codecs [4, 5, 11, 37, 43], which employ adversarial training [20] to enhance visual quality; vector quantization (VQ)-based codecs [26, 47, 60], which learn discrete representations for efficient compression; and diffusion-based codecs [10, 18, 28, 30, 38, 39, 50, 51, 60, 66], which utilize diffusion denoising to reconstruct images. While GAN-based and vector quantization codecs have shown promising results, their generative capabilities remain inherently constrained. In contrast, diffusion-based codecs—benefiting from the exceptional generative capacity of diffusion models—have recently achieved remarkable performance, particularly under ultra-low bitrate conditions.

Despite these advances, we identify three fundamental limitations in current diffusion-based image codecs that prevent them from achieving effective content adaptation, which is crucial for optimally leveraging generative priors across diverse image structures and semantics. First, the prevalent use of *isotropic quantization* applies a uniform quantization step across the compact latent representation, ignoring the spatial heterogeneity of image content. This content-agnostic quantization creates a mismatch with the diffusion model’s noise-dependent prior: textured regions receive insufficient generative intervention while smooth regions are over-regularized, limiting the overall rate-perception performance. Second, these codecs suffer from an *information concentration bottleneck* caused by the architectural mismatch between the high-dimensional

noisy latent and the fixed 4-channel input of the pre-trained diffusion decoder [52]. Without explicit supervision, the model may fail to concentrate essential semantic information into the primary channels, leading to a non-adaptive latent representation that does not prioritize critical content. Third, existing methods struggle with *ineffective textual conditioning*: they either incur substantial bitrate overhead by transmitting textual side information or rely on generic prompts that lack content relevance, failing to provide content-adaptive textual guidance without bitrate costs.

To tackle the aforementioned limitations and establish a content-adaptive diffusion-based image codec, we introduce three technical innovations that address the identified limitations. First, we propose an Uncertainty-Guided Adaptive Quantization (UGAQ) method that fundamentally redesigns the quantization process. Our key innovation lies in learning a spatially-varying uncertainty map from the residual between the main latent representation and the upsampled hyperprior latent. This uncertainty map modulates the quantization noise level across different spatial locations, ensuring that textured regions receive stronger generative intervention while smooth regions maintain structural fidelity. This method effectively aligns the quantization-induced distortion with the diffusion model’s noise-dependent denoising strategy, enabling content-aware noise shaping. Second, we develop an Auxiliary Decoder-Guided Information Concentration (ADGIC) method that explicitly solves the information concentration bottleneck. We introduce a lightweight auxiliary decoder that operates exclusively on the first four channels of the noisy latent, producing an auxiliary reconstruction and computing its distortion against the original image. This design forces the compression model to concentrate semantically critical information into the primary channels used by the diffusion decoder, ensuring that the essential visual content is properly preserved and accessible to the generative prior, thereby enforcing a content-driven information allocation. Third, we design a Bitrate-Free Adaptive Textual Conditioning (BFATC) method that enables effective textual conditioning without textual bitrate overhead. We generate content-adaptive captions using a pre-trained BLIP model [34] from the auxiliary reconstruction image, providing semantically meaningful guidance to the diffusion process while requiring zero additional textual bitrate, effectively bridging the gap between conditioning quality and bitrate efficiency and achieving dynamic, content-specific conditioning.

In summary, our contributions are as follows:

- We identify three key limitations in current diffusion-based image codecs that hinder content adaptation: i) the mismatch from content-agnostic isotropic quantization, ii) the non-adaptive latent representation due to the information concentration bottleneck, and iii) the inability

- ity to provide content-aware textual guidance efficiently.
- We propose three novel methods that collectively establish a content-adaptive diffusion-based image codec: Uncertainty-Guided Adaptive Quantization for content-aware noise shaping, Auxiliary Decoder-Guided Information Concentration for content-driven information allocation, and Bitrate-Free Adaptive Textual Conditioning for content-specific semantic guidance.
 - We validate our codec across various datasets and evaluation criteria, showing improved quantitative and qualitative results, particularly at ultra-low bitrates.

2. Related Work

2.1. Learned Image Compression

Learned image compression has gained significant research interest in recent years. Most methods typically build upon three core components: non-linear transforms, quantization, and entropy models. Early work by Ballé *et al.* [6] established a foundational framework using convolutional networks with generalized divisive normalization (GDN) for non-linear analysis and synthesis transforms. To enable gradient-based training, they approximated quantization with additive uniform noise during optimization, combined with a factorized entropy model. Subsequent efforts have focused on enhancing each of these components. For transform design, attention mechanisms [13, 16] and transformer architectures [41, 42, 67] have been incorporated to better capture both global structures and local textures. For quantization, several alternatives have been proposed to improve quantizer differentiability and efficiency while maintaining inference-time discretization, including soft-to-hard quantization [3], soft-then-hard quantization [22], universal quantization [1], and non-uniform quantization [19]. Entropy modeling has also seen considerable advances. Early hyperprior models [7] introduced side information to capture spatial dependencies, while autoregressive model [45] further improved accuracy through spatial conditioning. More recent methods combine these ideas, such as channel-wise autoregressive models [44], checkerboard context models [23], and hybrid designs [17, 24, 27]. Transformer-based entropy models [32, 48] have also been explored to leverage long-range dependencies for more accurate probability estimation.

2.2. Diffusion-Based Image Compression

The success of diffusion models in image generation has motivated their recent adoption in extreme image compression, enabling highly realistic reconstruction under ultra-low bitrates. These methods typically leverage generative priors from pre-trained diffusion models by conditioning the decoding process on the noisy latent representation extracted from the input image. For example, PerCo [10] fine-

tunes a diffusion model using vector-quantized spatial features and global text descriptions generated by BLIP-2 [36]. ResULIC [28] further analyzes semantic residuals between the original and reconstructed images, transmitting text-guided residuals to steer the diffusion process. DiffEIC [38] and its extension RDEIC [39] demonstrate that conditioning solely on VAE-compressed latents can yield competitive performance without textual input. Alternatively, Relic *et al.* [50] and later work [51] model quantization errors in the latent space as noise and recover the image via a diffusion denoising process. Despite gains in perceptual realism, such methods often incur high decoding latency due to multi-step sampling strategies like DDIM [55]. To mitigate this, recent methods [21, 60, 65, 66] have incorporated one-step diffusion models [53], substantially improving inference efficiency while preserving high perceptual quality. Notwithstanding these advances, several key challenges persist that hinder content adaptation, motivating our work toward further improving compression performance.

3. Limitations of Diffusion-Based Compression

3.1. Isotropic Quantization

A critical yet underexplored challenge in diffusion-based generative compression lies in the mismatch between the content-agnostic quantization strategy and the operational principle of diffusion models. For example, most methods [21, 38, 60, 66] directly quantized the latent representation with a fixed step size, treating the quantization errors as uniform noise to be removed by diffusion process. While Relic *et al.* [51] introduced universal quantization to better align the quantization errors with Gaussian diffusion noise schedule through a derived signal-to-noise ratio (SNR)-matching scheme, it still relies on a global isotropic quantization parameter that is applied uniformly across the entire compact latent representation. This common practice of isotropic quantization implicitly assumes that the distortion introduced by quantization—which the diffusion model is tasked with denoising—affects all image regions equally. However, this assumption contradicts both the nature of image content and the behavior of diffusion models. Natural images exhibit significant spatial heterogeneity, comprising textured regions rich in high-frequency details and smooth regions with minimal information. Diffusion models, in turn, possess noise-level-dependent priors: at high noise levels, they excel at hallucinating realistic textures from a strong generative prior, while at lower noise levels, they act more as conservative denoisers, preserving transmitted structural information. Global isotropic quantization forces a single, compromise noise level upon the entire compact latent representation, failing to adapt to the varying content complexity. Consequently, textured regions suffer from insufficient generative intervention, leading to blurred de-

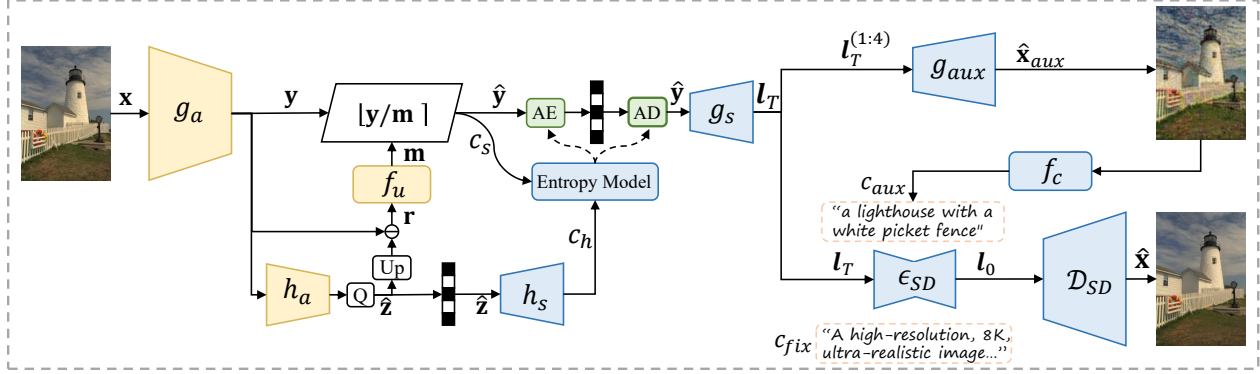


Figure 2. On the encoder side, an analysis transform g_a encodes the input image \mathbf{x} into a compact latent representation \mathbf{y} . An uncertainty map \mathbf{m} is estimated by f_u to guide the quantization of \mathbf{y} . The quantized latent $\hat{\mathbf{y}}$ is encoded into a bitstream via an arithmetic encoder (AE) and transmitted. On the decoder side, a synthesis transform g_s upsamples $\hat{\mathbf{y}}$ to produce a noisy latent \mathbf{l}_T at the spatial resolution required by the pre-trained Stable Diffusion VAE decoder \mathcal{D}_{SD} [52]. In learned codecs, \mathbf{l}_T typically has a high channel count (e.g., 320), while \mathcal{D}_{SD} is fixed to accept only 4-channel input. To resolve this, the entire \mathbf{l}_T is commonly input to the Unet ϵ_{SD} (a new input channel number is set to the first convolutional layer of the Unet) to utilize all available context for estimating more accurate 4-channel noise (the output channel number of the Unet is still 4) [66]. The denoising process is applied exclusively to the first four noisy channels $\mathbf{l}_T^{(1:4)}$, yielding the standard 4-channel clean latent \mathbf{l}_0 for \mathcal{D}_{SD} . To concentrate essential semantic information into $\mathbf{l}_T^{(1:4)}$, a lightweight auxiliary decoder g_{aux} takes $\mathbf{l}_T^{(1:4)}$ as inputs to reconstruct an auxiliary image $\hat{\mathbf{x}}_{aux}$. To produce a content-adaptive textual description c_{aux} , $\hat{\mathbf{x}}_{aux}$ is captioned by f_c (a frozen BLIP [34]). c_{aux} is then combined with a fixed description c_{fix} to condition a one-step diffusion denoising process [53].

tails, while smooth regions are subjected to unnecessary and potentially disruptive artificiality. This isotropic quantization strategy fails to harness the full potential of the diffusion model’s adaptive denoising capabilities due to its inability to perform content-aware noise shaping, creating a suboptimal alignment between the encoded signal and the decoder’s generative prior.

3.2. Information Concentration Bottleneck

Diffusion-based image codecs face a fundamental information concentration bottleneck that limits its efficiency and prevents content-aware representation learning. This bottleneck arises from the architectural mismatch between the high-dimensional noisy latents produced by learned codecs and the fixed-dimensional input expected by the pre-trained Stable Diffusion VAE decoder [52]. As shown in Fig. 2, learned codecs commonly generate a noisy latent \mathbf{l}_T with channel dimensions significantly exceeding the standard 4-channel latent space of the VAE decoder. However, only the first four channels $\mathbf{l}_T^{(1:4)}$ are denoised and fed into the subsequent VAE decoder \mathcal{D}_{SD} . This creates a critical constraint: the most semantically meaningful information must be concentrated within $\mathbf{l}_T^{(1:4)}$ for effective reconstruction. Without explicit guidance, the learning process may fail to adaptively concentrate essential semantic information into the critical first four channels based on image content. The diffusion model’s generative prior cannot be fully leveraged if the noisy latent lacks the necessary information density in its primary channels, resulting in a non-adaptive latent

representation that fails to preserve content-specific details.

3.3. Ineffective Textual Conditioning

Despite the remarkable capability of diffusion models to leverage textual guidance for high-quality image generation, existing diffusion-based compression methods struggle to effectively integrate content-adaptive text conditioning in a rate-efficient manner. Several methods [10, 18, 28] leverage multimodal large models to automatically generate textual descriptions, which are then losslessly compressed and transmitted as side information. While this provides content-aware guidance for the diffusion denoising process, the associated textual bitrate constitutes some overheads in ultra-low bitrate scenarios. For example, for the mobile satellite image communication application, the transmission packet size is limited (e.g., 450 bytes). Even low-bitrate texts consume a portion of the precious bit budget, leaving fewer bits for image transmission. Therefore, Zhang *et al.* [66] proposed to avoid transmitting textual descriptions and instead condition the diffusion model on a fixed, generic prompt such as “A high-resolution, 8K, ultra-realistic image with sharp focus, vibrant colors, and natural lighting”. Although this eliminates textual bitrate cost, the prompt is inherently agnostic to image content, failing to provide semantically meaningful guidance and limiting the model’s ability to reconstruct content-specific details. The fundamental limitation of these methods is their inability to provide content-adaptive textual conditioning without incurring additional textual bitrate costs, thereby hindering the potential for semantically-aware reconstruction.

4. Method

4.1. Uncertainty-Guided Adaptive Quantization

To address the limitations of content-agnostic global isotropic quantization in diffusion-based compression, we propose an Uncertainty-Guided Adaptive Quantization method that aligns the quantization process with the diffusion model’s noise-dependent generative prior, thereby achieving content-adaptive quantization.

Specifically, as illustrated in Fig. 2, our method begins by upsampling the hyperprior latent $\hat{\mathbf{z}}$ to match the spatial resolution of the main latent representation \mathbf{y} :

$$\bar{\mathbf{z}} = \text{UP}(\hat{\mathbf{z}}), \quad (1)$$

where $\text{UP}(\cdot)$ denotes bilinear upsampling. We then compute the residual between \mathbf{y} and $\bar{\mathbf{z}}$:

$$\mathbf{r} = \mathbf{y} - \bar{\mathbf{z}}. \quad (2)$$

This residual reflects the uncertainty between \mathbf{y} and $\bar{\mathbf{z}}$ —the larger the residual in a region, the less information $\bar{\mathbf{z}}$ conveys about \mathbf{y} , indicating more complex image textures and serving as a basis for content-aware adaptation. The residual is then processed through a lightweight uncertainty estimation network f_u to predict an uncertainty map:

$$\mathbf{m} = f_u(\mathbf{r}), \quad (3)$$

where each element of \mathbf{m} is restricted to be $m_{i,j} \geq 1$.

We then modulate the latent representation \mathbf{y} using the uncertainty map \mathbf{m} to achieve content-adaptive scaling:

$$\bar{\mathbf{y}} = \mathbf{y}/\mathbf{m}, \quad (4)$$

where “/” denotes element-wise division. The modulated latent $\bar{\mathbf{y}}$ is then quantized:

$$\hat{\mathbf{y}} = Q(\bar{\mathbf{y}}) = \left\lfloor \frac{\bar{\mathbf{y}}}{\Delta} \right\rfloor \cdot \Delta, \quad (5)$$

where $Q(\cdot)$ represents the quantization function, $\lfloor \cdot \rfloor$ denotes rounding to the nearest integer, and Δ is the quantization bin width. The quantization process can be approximated as the addition of independent uniform noise:

$$\hat{\mathbf{y}} - \bar{\mathbf{y}} \approx \boldsymbol{\epsilon}, \quad \epsilon_{i,j} \sim \mathcal{U}(-\Delta/2, \Delta/2), \quad (6)$$

where $\boldsymbol{\epsilon}$ represents the quantization error approximated as i.i.d. uniform noise, and $\mathcal{U}(-\Delta/2, \Delta/2)$ denotes a uniform distribution over the interval $[-\Delta/2, \Delta/2]$. The quantized latent representation is therefore:

$$\hat{\mathbf{y}} \approx \bar{\mathbf{y}} + \boldsymbol{\epsilon} = \mathbf{y}/\mathbf{m} + \boldsymbol{\epsilon}. \quad (7)$$

At the decoder, $\hat{\mathbf{y}}$ is directly feed into the diffusion model without applying the inverse scaling. This formulation reveals that although the quantization noise $\boldsymbol{\epsilon}$ has a constant

variance $\sigma_\epsilon^2 = \Delta^2/12$ [51], the pre-quantization modulation by \mathbf{m} creates a locally varying signal-to-noise ratio (SNR) at the decoder input, enabling content-aware noise shaping. The effective local SNR for a latent element $y_{i,j}$ can be characterized as:

$$\text{SNR}_{i,j} \propto \frac{\mathbb{E}[\bar{y}_{i,j}^2]}{\sigma_\epsilon^2} = \frac{\mathbb{E}[y_{i,j}^2]}{m_{i,j}^2 \cdot \sigma_\epsilon^2}, \quad (8)$$

where $\mathbb{E}[\bar{y}_{i,j}^2]$ is the power of the modulated latent. This quantitative relationship leads to our central content-adaptive mechanism:

- High uncertainty regions ($m_{i,j}$ is large): The signal power is reduced by a factor of $m_{i,j}^2$, resulting in low local SNR. The diffusion model therefore relies more heavily on its generative prior to synthesize details, adaptively enhancing texture generation where needed.
- Low uncertainty regions ($m_{i,j}$ is small): The signal power remains largely unchanged, maintaining high local SNR. The diffusion model thus prioritizes faithful preservation of the transmitted structural information, adaptively conserving fidelity in smooth areas.

Our UGAQ method thereby reduces the misalignment between the quantization distortion and the diffusion model’s noise-dependent denoising strategy through content-aware noise shaping, adaptively leveraging its generative capabilities based on local image content.

4.2. Auxiliary Decoder-Guided Information Concentration

To address the information concentration bottleneck and enable content-aware latent representation learning, we propose an Auxiliary Decoder-Guided Information Concentration method. Our key insight is that without explicit supervision, it may be difficult for the first four channels of the noisy latent contain essential semantic information for high-quality reconstruction, resulting in a non-adaptive latent distribution. Therefore, we introduce a lightweight auxiliary decoder g_{aux} that operates exclusively on these primary channels, providing direct supervision to optimize their information-carrying capacity and enforce content-driven information allocation.

Formally, given the full noisy latent \mathbf{l}_T with C channels where $C \gg 4$, as shown in Fig. 2, we feed its first four channels $\mathbf{l}_T^{(1:4)}$ into an auxiliary decoder g_{aux} to produce an auxiliary reconstructed image $\hat{\mathbf{x}}_{aux}$:

$$\hat{\mathbf{x}}_{aux} = g_{aux}(\mathbf{l}_T^{(1:4)}). \quad (9)$$

We then compute an auxiliary reconstruction loss between this output and the original image \mathbf{x} :

$$\mathcal{L}_{aux} = \|\mathbf{x} - \hat{\mathbf{x}}_{aux}\|_2^2. \quad (10)$$

This auxiliary reconstruction loss will be incorporated into the overall loss function [66] to optimize the model.

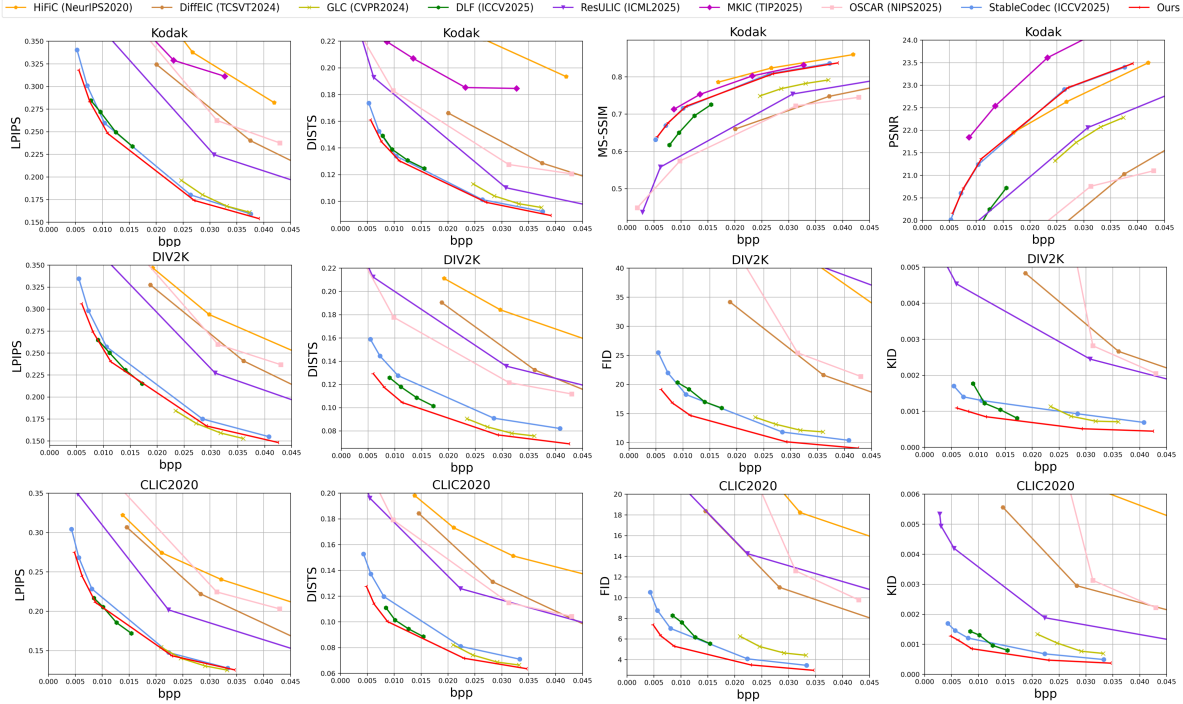


Figure 3. Quantitative comparisons of different generative image codecs on Kodak, DIV2K Val, and CLIC 2020 Test.

4.3. Bitrate-Free Adaptive Textual Conditioning

To overcome the limitation of being unable to obtain content-adaptive textual conditioning without additional textual bitrate costs, we propose a Bitrate-Free Adaptive Textual Conditioning method. Our key insight is to leverage the auxiliary reconstruction $\hat{\mathbf{x}}_{aux}$ —generated by the lightweight auxiliary decoder g_{aux} introduced in Section 4.2—as a proxy to infer a content-aware textual description. Since $\hat{\mathbf{x}}_{aux}$ is derived entirely from the noisy latent $l_T^{(1:4)}$, our method provides semantically meaningful textual guidance with zero textual bitrate cost.

As shown in Fig. 2, the auxiliary reconstructed image $\hat{\mathbf{x}}_{aux}$ is fed into an image captioning model f_c (a frozen BLIP [34]) to produce a content-adaptive textual description c_{aux} :

$$c_{aux} = f_c(\hat{\mathbf{x}}_{aux}). \quad (11)$$

To enhance robustness and stability, we then combine c_{aux} with a fixed, generic textual description c_{fix} (“A high-resolution, 8K, ultra-realistic image with sharp focus, vibrant colors, and natural lighting”) [63, 66] using the simple string concatenation:

$$c = c_{aux} + c_{fix}. \quad (12)$$

This combined textual description c is used as the conditions for the diffusion denoising process.

4.4. Implementation

Our codec follows the prevalent auto-encoder architecture commonly adopted in learned image compression. To balance generative capability against decoding complexity, we employ a distilled version [53] of Stable Diffusion 2.1 [52] to achieve **one-step diffusion**. A lightweight *blip-image-captioning-base* model is used for auxiliary reconstructed image captioning. For entropy modeling, we integrate both a hyperprior c_h and a spatial prior c_s , where the latter is generated by a 4-step autoregressive entropy model with quadtree partitioning [35].

We train our codec on the training sets of DF2K [40] and CLIC 2020 Professional [57] with the rate-distortion objective:

$$\mathcal{L} = \lambda \mathcal{R} + \mathcal{D}, \quad (13)$$

where λ is the Lagrange multiplier, \mathcal{R} is the bitrate. In addition to our proposed auxiliary reconstruction loss, the distortion term \mathcal{D} incorporates several components from [66], including MSE, LPIPS (with VGG features), a CLIP distance [49], and a GAN loss. More implementation details can be found in the supplementary materials.

5. Experiments

5.1. Experimental Settings

Test Datasets. We use Kodak [29], the validation set of DIV2K (DIV2K Val) [2], and the test set of CLIC 2020 Pro-




Figure 4. Qualitative comparison of different generative image codecs on the Kodak dataset under ultra-low bitrate conditions.

fessional (CLIC 2020 Test) [57] for evaluation. The Kodak dataset contains 24 images with a resolution of 768×512 . The DIV2K Val and the CLIC 2020 Test contain 100 and 428 high-quality 2K-resolution images, respectively. All images are evaluated with the original resolution.

Evaluation Metrics. We measure bitrate cost in bits per pixel (bpp). Following prior work [51, 59, 66], we evaluate perceptual quality using DISTS [15], LPIPS [64], FID [25], and KID [63]. Note that FID and KID are omitted on the Kodak dataset due to its limited size. Results of PSNR and MS-SSIM are also provided in the supplementary material.

Comparison Methods. We compare our codec against several generative image codecs, including: HiFiC [43], a representative GAN-based codec [20]; DLF [59] and GLC [26], which are leading vector quantization-based codecs; and a range of diffusion-based codecs, namely DiffeIC [38], ResULIC [28], MKIC [18], OSCAR [21], and StableCodec (our redesigned variant; see the supplementary material for details) [66].

5.2. Quantitative and Qualitative Comparisons

As illustrated in Fig. 3, we evaluate the quantitative compression performance of different generative image codecs under ultra-low bitrate conditions. Our codec consistently outperforms other diffusion-based codecs—including DiffeIC, ResULIC, MKIC, OSCAR, and StableCodec—across all evaluated datasets in terms of LPIPS, DISTS, FID, and KID, validating the effectiveness of our design. We further provide a qualitative comparison on the Kodak dataset in Fig. 4 under comparable bitrates. Visual results show

Table 1. Ablation studies of our proposed methods on Kodak. Negative BD-rate (%) values indicate better compression performance. The distortion is measured by LPIPS and DISTS.

Models	UGAQ	ADGIC	BFATC	LPIPS	DISTS
M_0	✗	✗	✗	0.00	0.00
M_1	✓	✗	✗	-3.7	-2.7
M_2	✓	✓	✗	-5.3	-3.5
M_3	✓	✓	✓	-6.8	-5.5

that competing codecs such as StableCodec, DLF, MKIC, and OSCAR produce overly smooth regions or lose fine-grained textures. In contrast, our reconstructions preserve more high-frequency details and exhibit more natural visual characteristics. More results can be found in the supplementary material.

5.3. Ablation Study

In this section, we conduct ablation studies to validate the effectiveness of our proposed methods.

Uncertainty-Guided Adaptive Quantization. We first evaluate the contribution of our UGAQ method by integrating it into the baseline Model M_0 to construct Model M_1 . As shown in Tab. 1, UGAQ yields BD-rate reductions of 3.7% and 2.7% on LPIPS and DISTS, respectively. This improvement is attributed to the ability of UGAQ to provide content-aware quantization, in contrast to isotropic quantization which applies uniform quantization strength regardless of local content complexity. As visualized in Fig. 5, the




Figure 5. Analysis of Uncertainty-Guided Adaptive Quantization (UGAQ) and the isotropic quantization on the DIV2K dataset.

residual $y - \bar{z}$ serves as an indicator of content uncertainty, with larger values corresponding to highly textured regions. Our method translates this signal into a spatially-varying uncertainty map m , where regions with more complex texture are assigned larger values—opposing existing spatial scaling-based quantization method [33]. This map guides the quantization process in a content-adaptive manner, producing quantization residuals $y - \hat{y}$ that exhibit clear semantic alignment and strongly correlate with content complexity. This demonstrates the ability of UGAQ to actively shape quantization errors, thereby aligning the quantization distortion with the diffusion model’s noise-dependent denoising behavior through content-aware noise allocation.

Auxiliary Decoder-Guided Information Concentration. We proceed to validate our ADGIC method by progressively integrating it into Model M_1 . The results in Tab. 1 show that ADGIC yields a further BD-rate reduction of 1.6% (from -3.7% to -5.3%) on LPIPS and 0.8% (from -2.7% to -3.5%) on DISTS. To better understand its working mechanism, we analyze the energy distribution across the first four channels of the noisy latent $l_T^{(1:4)}$ with and without ADGIC under varying bitrate conditions on the Kodak dataset. Following [12], we employ variance as a measure of channel energy. As shown in Fig. 6, ADGIC enhances energy concentration in the primary channels, demonstrating its role in enforcing content-aware information allocation by prioritizing essential semantics adaptively.

Bitrate-Free Adaptive Textual Conditioning. We further validate our BFATC method by integrating it into Model M_2 . As reported in Tab. 1, BFATC brings an additional 1.5% BD-rate improvement (from -5.3% to -6.8%) on LPIPS and 2.0% (from -3.5% to -5.5%) on DISTS. This performance gain demonstrates that providing content-adaptive textual guidance is crucial for enhancing compression efficiency. Our BFATC method remains robust across bitrates. As visualized in Fig. 7, even under severe reconstruction noise at low bitrates, the auxiliary reconstructed image retains sufficient semantic content to produce textual descriptions that remain semantically consistent.

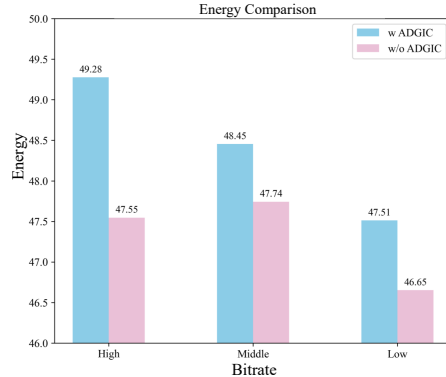


Figure 6. Energy comparison of the first four channels of the noisy latents of the codecs with and without Auxiliary Decoder-Guided Information Concentration (ADGIC) on the Kodak dataset.

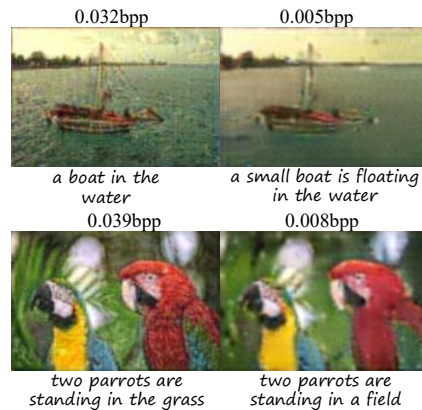


Figure 7. Illustration of the textual descriptions extracted from auxiliary reconstructed images under different bitrate conditions.

6. Conclusion

In this work, we address three key limitations hindering content adaptation in diffusion-based image compression: the misalignment from isotropic quantization, the information concentration bottleneck in latent representation, and the inefficiency of textual conditioning. To this end, we propose a content-adaptive diffusion-based image codec (CADC) built upon three technical innovations: an Uncertainty-Guided Adaptive Quantization method that aligns effective quantization distortion with content characteristics through content-aware noise shaping; an Auxiliary Decoder-Guided Information Concentration method that ensures essential semantic information is adaptively preserved in the primary latent channels via content-driven allocation; and a Bitrate-Free Adaptive Textual Conditioning method that derives content-specific semantic guidance without textual bitrate overhead. Extensive experiments validate that our codec achieves the state-of-the-art perceptual quality, particularly at ultra-low bitrates.

7. Acknowledgement

This research was supported in part by the HongKong Innovation and Technology Commission (ITC) grant GHP/044/21SZ and PRP/036/24FX, in part by the General Research Fund(GRF) of the RGC of Hong Kong under Grants 11200323, in part by National Natural Science Foundation of China (NSFC)/ResearchGrants Council (RGC) Joint Research Scheme N_CityU198/24.

References

- [1] Eirikur Agustsson and Lucas Theis. Universally quantized neural compression. *NeurIPS*, 33:12367–12376, 2020. [2](#), [3](#)
- [2] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *CVPRW*, pages 126–135, 2017. [6](#), [1](#)
- [3] Eirikur Agustsson, Fabian Mentzer, Michael Tschannen, Lukas Cavigelli, Radu Timofte, Luca Benini, and Luc V Gool. Soft-to-hard vector quantization for end-to-end learning compressible representations. *Advances in neural information processing systems*, 30, 2017. [2](#), [3](#)
- [4] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative adversarial networks for extreme learned image compression. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 221–231, 2019. [2](#)
- [5] Eirikur Agustsson, David Minnen, George Toderici, and Fabian Mentzer. Multi-realism image compression with a conditional generator. In *CVPR*, pages 22324–22333, 2023. [2](#)
- [6] Johannes Ballé, Valero Laparra, and Eero P. Simoncelli. End-to-end optimized image compression. In *ICLR*, 2017. [2](#), [3](#)
- [7] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *ICLR*, 2018. [2](#), [3](#)
- [8] Yochai Blau and Tomer Michaeli. Rethinking lossy compression: The rate-distortion-perception tradeoff. In *International Conference on Machine Learning*, pages 675–685. PMLR, 2019. [2](#)
- [9] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer Ohm. Overview of the versatile video coding (vvc) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10):3736–3764, 2021. [2](#)
- [10] Marlene Careil, Matthew J Muckley, Jakob Verbeek, and Stéphane Lathuilière. Towards image compression with perfect realism at ultra-low bitrates. In *ICLR*, 2023. [2](#), [3](#), [4](#)
- [11] Bolin Chen, Zhao Wang, Binzhe Li, Shiqi Wang, and Yan Ye. Compact temporal trajectory representation for talking face video compression. *CSVT*, 33(11):7009–7023, 2023. [2](#)
- [12] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learning image and video compression through spatial-temporal energy compaction. In *CVPR*, pages 10071–10080, 2019. [8](#)
- [13] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *CVPR*, 2020. [2](#), [3](#)
- [14] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*, 2023. [1](#)
- [15] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. *IEEE TPAMI*, 44(5):2567–2581, 2020. [7](#)
- [16] Donghui Feng, Zhengxue Cheng, Shen Wang, Ronghua Wu, Hongwei Hu, Guo Lu, and Li Song. Linear attention modeling for learned image compression. In *CVPR*, pages 7623–7632, 2025. [2](#), [3](#)
- [17] Haisheng Fu, Feng Liang, Jie Liang, Yongqiang Wang, Zhenman Fang, Guohe Zhang, and Jingning Han. Fast and high-performance learned image compression with improved checkerboard context model, deformable residual module, and knowledge distillation. *IEEE TIP*, 2024. [2](#), [3](#)
- [18] Junlong Gao, Zhimeng Huang, Qi Mao, Siwei Ma, and Chuanmin Jia. Exploring multimodal knowledge for image compression via large foundation models. *IEEE TIP*, 2025. [2](#), [4](#), [7](#), [1](#)
- [19] Ziqing Ge, Siwei Ma, Wen Gao, Jingshan Pan, and Chuanmin Jia. Nlic: Non-uniform quantization-based learned image compression. *IEEE TCSVT*, 34(10):9647–9663, 2024. [2](#), [3](#)
- [20] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [2](#), [7](#)
- [21] Jinpei Guo, Yifei Ji, Zheng Chen, Kai Liu, Min Liu, Wang Rao, Wenbo Li, Yong Guo, and Yulun Zhang. Oscar: One-step diffusion codec across multiple bit-rates. *arXiv preprint arXiv:2505.16091*, 2025. [3](#), [7](#), [1](#)
- [22] Zongyu Guo, Zhizheng Zhang, Runsen Feng, and Zhibo Chen. Soft then hard: Rethinking the quantization in neural image compression. pages 3920–3929. PMLR, 2021. [2](#), [3](#)
- [23] Dailan He, Yaoyan Zheng, Baocheng Sun, Yan Wang, and Hongwei Qin. Checkerboard context model for efficient learned image compression. In *CVPR*, pages 14771–14780, 2021. [3](#)
- [24] Dailan He, Ziming Yang, Weikun Peng, Rui Ma, Hongwei Qin, and Yan Wang. Elic: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In *CVPR*, pages 5718–5727, 2022. [2](#), [3](#)
- [25] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30, 2017. [7](#)
- [26] Zhaoyang Jia, Jiahao Li, Bin Li, Houqiang Li, and Yan Lu. Generative latent coding for ultra-low bitrate image compression. In *CVPR*, pages 26088–26098, 2024. [2](#), [7](#), [1](#)
- [27] Wei Jiang, Jiayu Yang, Yongqi Zhai, Peirong Ning, Feng Gao, and Ronggang Wang. Mlic: Multi-reference entropy

- model for learned image compression. In *ACM MM*, pages 7618–7627, 2023. 2, 3
- [28] Anle Ke, Xu Zhang, Tong Chen, Ming Lu, Chao Zhou, Jiawen Gu, and Zhan Ma. Ultra lowrate image compression with semantic residual coding and compression-aware diffusion. *ICML*, 2025. 2, 3, 4, 7, 1
- [29] Eastman Kodak. Kodak lossless true color image suite (photo cd pcd0992). URL <http://r0k.us/graphics/kodak>, 6(2):5, 1993. 6
- [30] Haowei Kuang, Yiyang Ma, Wenhan Yang, Zongming Guo, and Jiaying Liu. Consistency guided diffusion model with neural syntax for perceptual image compression. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 1622–1631, 2024. 2
- [31] Nupur Kumari, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Ensembling off-the-shelf models for gan training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10651–10662, 2022. 1
- [32] Daxin Li, Yuanchao Bai, Kai Wang, Junjun Jiang, Xianming Liu, and Wen Gao. Groupedmixer: An entropy model with group-wise token-mixers for learned image compression. *IEEE TCSVT*, 34(10):9606–9619, 2024. 2, 3
- [33] Jiahao Li, Bin Li, and Yan Lu. Hybrid spatial-temporal entropy modelling for neural video compression. In *ACMMM*, pages 1503–1511, 2022. 8
- [34] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. pages 12888–12900. PMLR, 2022. 2, 4, 6
- [35] Jiahao Li, Bin Li, and Yan Lu. Neural video compression with diverse contexts. In *CVPR*, pages 22616–22626, 2023. 6, 1
- [36] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. 3
- [37] Meng Li, Shangyin Gao, Yihui Feng, Yibo Shi, and Jing Wang. Content-oriented learned image compression. In *European Conference on Computer Vision*, pages 632–647. Springer, 2022. 2
- [38] Zhiyuan Li, Yanhui Zhou, Hao Wei, Chenyang Ge, and Jingwen Jiang. Towards extreme image compression with latent feature guidance and diffusion prior. *IEEE TCSVT*, 2024. 2, 3, 7, 1
- [39] Zhiyuan Li, Yanhui Zhou, Hao Wei, Chenyang Ge, and Ajmal Saeed Mian. Diffusion-based extreme image compression with compressed feature initialization. 2024. 2, 3
- [40] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017. 6
- [41] Jinming Liu, Heming Sun, and Jiro Katto. Learned image compression with mixed transformer-cnn architectures. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14388–14397, 2023. 2, 3
- [42] Ming Lu, Peiyao Guo, Huiqing Shi, Chuntong Cao, and Zhan Ma. Transformer-based image compression. *arXiv preprint arXiv:2111.06707*, 2021. 2, 3
- [43] Fabian Mentzer, George D Toderici, Michael Tschannen, and Eirikur Agustsson. High-fidelity generative image compression. *NeurIPS*, 33:11913–11924, 2020. 2, 7, 1
- [44] David Minnen and Saurabh Singh. Channel-wise autoregressive entropy models for learned image compression. In *ICIP*, pages 3339–3343. IEEE, 2020. 2, 3
- [45] David Minnen, Johannes Ballé, and George Toderici. Joint autoregressive and hierarchical priors for learned image compression. In *NeurIPS*, pages 10794–10803, 2018. 2, 3
- [46] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 1
- [47] Linfeng Qi, Zhaoyang Jia, Jiahao Li, Bin Li, Houqiang Li, and Yan Lu. Generative latent coding for ultra-low bitrate image and video compression. *IEEE Transactions on Circuits and Systems for Video Technology*, 2025. 2
- [48] Yichen Qian, Ming Lin, Xiuyu Sun, Zhiyu Tan, and Rong Jin. Entroformer: A transformer-based entropy model for learned image compression. *arXiv preprint arXiv:2202.05492*, 2022. 2, 3
- [49] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 6
- [50] Lucas Relic, Roberto Azevedo, Markus Gross, and Christopher Schroers. Lossy image compression with foundation diffusion models. In *ECCV*, pages 303–319. Springer, 2024. 2, 3
- [51] Lucas Relic, Roberto Azevedo, Yang Zhang, Markus Gross, and Christopher Schroers. Bridging the gap between gaussian diffusion models and universal quantization for image compression. In *CVPR*, pages 2449–2458, 2025. 2, 3, 5, 7
- [52] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2, 4, 6
- [53] Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. In *ECCV*, pages 87–103. Springer, 2024. 3, 4, 6
- [54] Athanassios Skodras, Charilaos Christopoulos, and Touradj Ebrahimi. The jpeg 2000 still image compression standard. *IEEE Signal processing magazine*, 18(5):36–58, 2002. 2
- [55] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3
- [56] Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (hevc) standard. *IEEE Transactions on circuits and systems for video technology*, 22(12):1649–1668, 2012. 2

- [57] George Toderici, Lucas Theis, Nick Johnston, Eirikur Agustsson, Fabian Mentzer, Johannes Ballé, Wenzhe Shi, and Radu Timofte. Clic 2020: Challenge on learned image compression. *Retrieved March, 29:2021*, 2020. [1](#), [6](#), [7](#)
- [58] Gregory K Wallace. The jpeg still picture compression standard. *Communications of the ACM*, 34(4):30–44, 1991. [2](#)
- [59] Naifu Xue, Zhaoyang Jia, Jiahao Li, Bin Li, Yuan Zhang, and Yan Lu. Dlf: Extreme image compression with dual-generative latent fusion. *ICCV*, 2025. [1](#), [7](#), [3](#)
- [60] Naifu Xue, Zhaoyang Jia, Jiahao Li, Bin Li, Yuan Zhang, and Yan Lu. One-step diffusion-based image compression with semantic distillation. *arXiv preprint arXiv:2505.16687*, 2025. [2](#), [3](#)
- [61] Weihao Yu and Xinchao Wang. Mambaout: Do we really need mamba for vision? In *CVPR*, pages 4484–4496, 2025. [1](#), [2](#)
- [62] Weihao Yu, Pan Zhou, Shuicheng Yan, and Xinchao Wang. Inceptionnext: When inception meets convnext. In *CVPR*, pages 5672–5683, 2024. [1](#), [2](#)
- [63] Aiping Zhang, Zongsheng Yue, Renjing Pei, Wenqi Ren, and Xiaochun Cao. Degradation-guided one-step image super-resolution with diffusion priors. *arXiv preprint arXiv:2409.17058*, 2024. [6](#), [7](#)
- [64] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. [7](#)
- [65] Tianyu Zhang, Dong Liu, and Chang Wen Chen. Ultra-low bitrate perceptual image compression with shallow encoder. *arXiv preprint arXiv:2512.12229*, 2025. [3](#)
- [66] Tianyu Zhang, Xin Luo, Li Li, and Dong Liu. Stablecodec: Taming one-step diffusion for extreme image compression. *ICCV*, 2025. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [67] Renjie Zou, Chunfeng Song, and Zhaoxiang Zhang. The devil is in the details: Window-based attention for image compression. In *CVPR*, pages 17492–17501, 2022. [2](#), [3](#)