

Towards Efficient Low-rate Image Compression with Frequency-aware Diffusion Prior Refinement

Yichong Xia^{*1,2}, Yimin Zhou^{*1}, Jinpeng Wang³, Bin Chen^{†3}

¹Tsinghua Shenzhen International Graduate School, Tsinghua University

²Peng Cheng Laboratory

³Harbin Institute of Technology, Shenzhen

xiayc23,zhouym24,wjp20@mails.tsinghua.edu.cn, chenbin2021@hit.edu.cn

Abstract

Recent advancements in diffusion-based generative priors have enabled visually plausible image compression at extremely low bit rates. However, existing approaches suffer from slow sampling processes and suboptimal bit allocation due to fragmented training paradigms. In this work, we propose Accelerate **Diffusion-based Image Compression via Consistency Prior Refinement (DiffCR)**, a novel compression framework for efficient and high-fidelity image reconstruction. At the heart of DiffCR is a Frequency-aware Skip Estimation (FaSE) module that refines the ϵ -prediction prior from a pre-trained latent diffusion model and aligns it with compressed latents at different timesteps via Frequency Decoupling Attention (FDA). Furthermore, a lightweight consistency estimator enables fast **two-step decoding** by preserving the semantic trajectory of diffusion sampling. Without updating the backbone diffusion model, DiffCR achieves substantial bitrate savings (27.2% BD-rate (LPIPS) and 65.1% BD-rate (PSNR)) and over 10 \times speed-up compared to SOTA diffusion-based compression baselines.

Introduction

In recent years, remarkable progress has been made in image compression technology. However, compressing high-definition images at low bit-rate still poses significant challenges. Traditional image compression standards (Bellard 2014; Wallace 1992) rely on manually designed transform methods. Nevertheless, in low bit-rate (≤ 0.1 bpp) scenarios, prominent block artifacts and color distortion issues occur. Although learning-based end-to-end neural compression models (Cheng et al. 2020; Ballé, Laparra, and Simoncelli 2016; Ballé et al. 2018; He et al. 2022) outperform traditional methods in terms of performance, their distortion-driven optimization objectives still lead to substantial loss of texture and details in the reconstructed images when bandwidth is limited. (Blau and Michaeli 2019; Agustsson et al. 2023) summarizes this phenomenon as a triple trade-off between bit rate, distortion, and realism, revealing the inherent contradiction that extreme compression inevitably entails semantic information loss and a decline in perceptual quality.

*These authors contributed equally.

†Corresponding author(s)

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

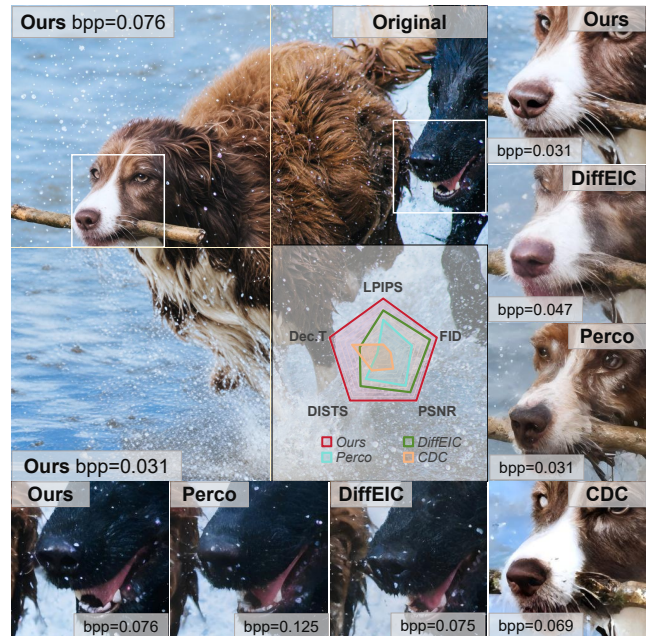


Figure 1: Qualitative and quantitative comparison between DiffEIC, Perco, HiFiC, MS-ILLM, CDC, and our proposed approach. DiffCR (Ours) not only possesses realism but also faithfully restores the details of the original image. The radar chart in the bottom right corner illustrates the performance comparison among various methods on the CLIC20 dataset. *Best viewed when zoomed in.*

To address this issue, recent studies (Mentzer et al. 2020; Muckley et al. 2023; Lu et al. 2024; Qin et al. 2023) have introduced generative adversarial networks (GAN) to optimize perceptual quality. Based on this, (Lee et al. 2024) introduces semantic information at both the encoding and decoding ends to enhance the image restoration effect. Constrained by the GAN structure and the single perceptual loss, these methods still have considerable room for improvement in statistical fidelity. Meanwhile, some 3D generative methods based on neural representation and Gaussian splatting (Liu et al. 2025a,b) have also been applied to lossy compression tasks for videos and images.

In contrast, diffusion models (Dhariwal and Nichol 2021; Rombach et al. 2022) demonstrate superior generative quality through a progressive denoising process. The statistical characteristics of their reconstructed images are closer to the distribution of natural images. (Yang and Mandt 2024; Hoogeboom et al. 2023) first applied the DDPM (Ho, Jain, and Abbeel 2020) framework to the field of image compression and achieved breakthrough progress. However, this method requires the complete training of diffusion model components and has problems such as high computational costs and generalization being limited by the training data.

More recently, several studies based on Latent Diffusion Models (LDMs) (Rombach et al. 2022) have introduced novel concepts to overcome the issues present in the GAN framework. For instance, (Lei et al. 2023) endeavors to manipulate the pre-trained LDM by encoding sketch and text semantics. Utilize the sampling process of the diffusion model for image decoding. Given the difficulties in jointly optimizing semantic embeddings and the transfer representation for compression objectives, they have dedicated significant resources to iterative semantic embedding and alignment. (Careil et al. 2024) makes use of a fully-trained LDM and encodes conditional images with a trainable codebook. However, this approach has two drawbacks. Firstly, it demands careful optimization of the model on a dataset consisting of millions of images. Secondly, the independent training of the compressor and the diffusion model diminishes the learning capacity of the compressor. (Li et al. 2024a) utilizes a pre-trained foundation LDM. Nevertheless, the joint training method it employs makes it arduous for the compressor to learn compact representations by capitalizing on the diffusion prior. (Ke et al. 2025) harnesses Multimodal Large Language Models (MLLM) to acquire high-quality text semantics and attains remarkable performance improvements at ultra low bit-rate. Nevertheless, this comes at the cost of extremely high encoding latency.

In general, existing image compression schemes based on diffusion suffer from two key deficiencies. (1) Suboptimal bit allocation due to training paradigms and latent misalignment. In current research, the training of the compressor and diffusion model is either conducted separately (Yang and Mandt 2024; Careil et al. 2024) or the diffusion loss and compressor loss are misaligned (Li et al. 2024a; Ke et al. 2025). (Xia et al. 2025) attempts have been made to employ proxy losses to enable the compressor to capture diffusion priors during training, yet the most compact compression representation remains elusive. (2) Methods based on diffusion rely on standard diffusion models and encounter challenges in sampling as they require multiple iterations to generate realistic image outputs. This not only leads to significant decoding delays but also hinders their seamless integration with perceptual losses applied to the final image output, consequently diminishing decoding fidelity.

To address the aforementioned critical issues, we introduce the Accelerate **D**iffusion-based Image Compression via **C**onsistency Prior **R**efinement (DiffCR). Reflecting upon the training paradigm of ϵ -Prediction in the fundamental diffusion model, we devise a Consistency Refinement Estimator (CRE) to actively correct and enhance diffusion pri-

ors, aligning the model’s focus on optimizing both the realism of generated images and the quest for compact compression representations. Specifically, we craft the Frequency-aware Skip Estimation (FaSE) to align diffusion priors with compressed latents at different time steps and employ Frequency Decoupling Attention (FDA) to capture contextual relationships between the two at various frequencies. Furthermore, we maintain the consistency of FaSE sampling trajectories with original ODE trajectories through consistency losses. Unlike conventional consistency distillation methods, our approach achieves high-fidelity decoding in just two steps with only 1/100 of the parameters of the original denoising network (8M vs.800M). Furthermore, DiffCR leverages mixed semantic flows to further fortify and stabilize the generated results. Benefitting from these characteristics, DiffCR decodes highly realistic images at extremely low bitrates (≤ 0.05 bpp) in a mere 0.4 seconds, as depicted in fig. 1.

Our contributions can be summarized as follows:

- We introduce the Frequency-aware Skip Estimation to refine (FaSE) and align pre-trained diffusion priors with compressed latents. This not only facilitates the compressor in acquiring more compact compression representations but also allows our model to decode high-quality results in just two steps through consistency constraints.
- We devise the Frequency Decoupling Attention (FDA) to synchronize diffusion priors generated at different time steps with compressed latents. FDA decouples the two in the Fourier domain and modulates them across different time steps using a temporal mask.
- Building upon these key points, we propose DiffCR. DiffCR not only enhances encoding efficiency but also utilizes the semantic branch to stabilize and enhance decoding performance. We validate DiffCR on multiple high-quality datasets, demonstrating its superior perceptual performance. Compared to a state-of-the-art compression approach (Li et al. 2024a), which is also based on SD 2.1, DiffCR achieves savings in BD-rate of 27.2% (LPIPS), 32.8% (FID), and 65.1% (PSNR).

Background

Learned Latent Compression

Learnable image compression (Ballé, Laparra, and Simoncelli 2016; Ballé et al. 2018; Minnen, Ballé, and Toderici 2018; He et al. 2022; Cheng et al. 2020; Qin et al. 2024) involves extracting image representations through an encoder, estimating the probability distribution of these representations using entropy models, and subsequently performing entropy coding. Following Shannon’s information theory (Shannon 1948), the training paradigm can be unified into the following form:

$$\mathcal{L} = R_x + \lambda D(x, \mathcal{M}(x)) \quad (1)$$

Here, $\mathcal{M}(\cdot)$ represents the compressor, and R_x denotes the bitrate of the image x . The function $D(\cdot)$ typically selects pixel-level distortions such as Mean Squared Error (MSE) or

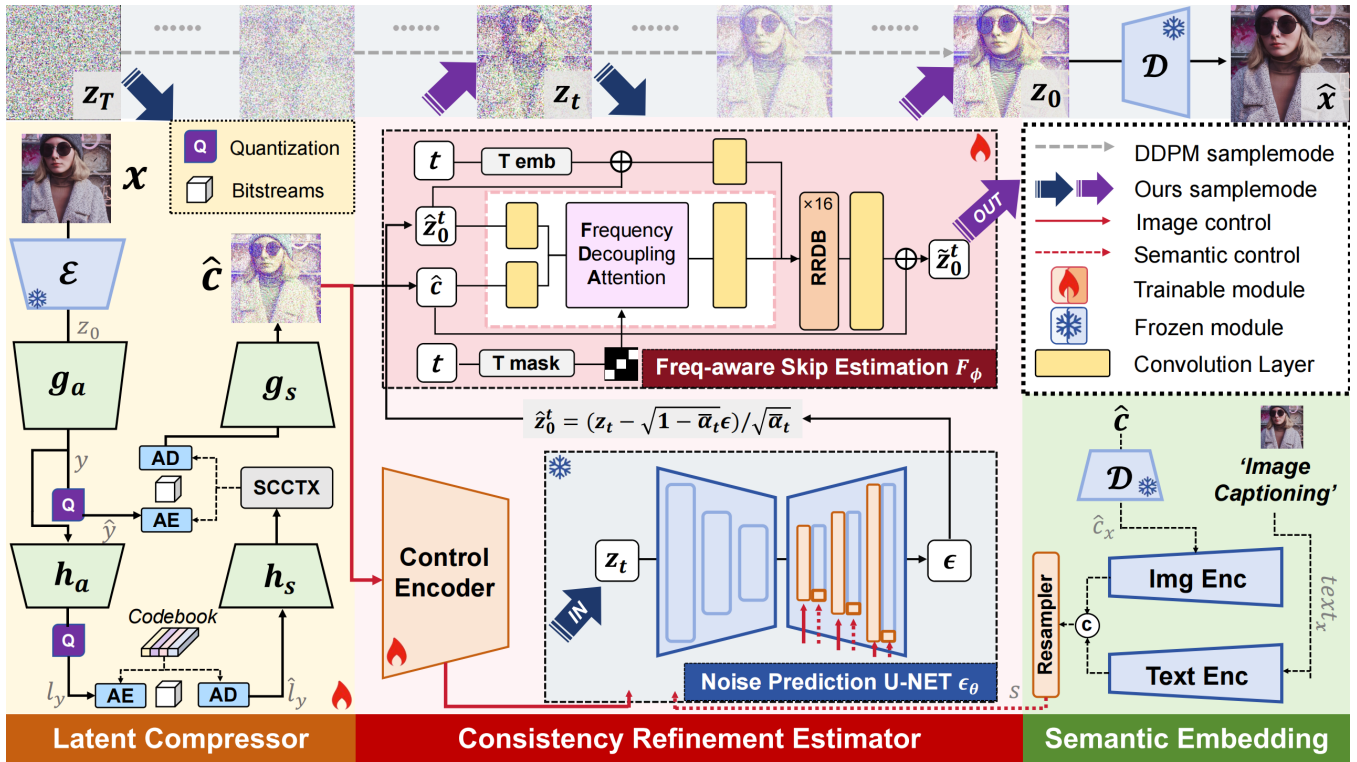


Figure 2: Illustration of the proposed DiffCR. RRDB signifies Residual in Residual Dense Block (Wang et al. 2018). DiffCR utilizes two branches to guide pre-trained denoising networks in generation and reinforces prior correction and refinement through the FaSE module, aligning it with the reconstruction objective of the compressor. Unlike the ordinary sampling process (depicted by the grey dashed line), DiffCR permits a two-step sampling approach (illustrated by the blue-purple arrows).

MS-SSIM (Wang, Simoncelli, and Bovik 2003). In estimating some generative compression tasks (Careil et al. 2024; Mao et al. 2024; Li et al. 2024a; Yang and Mandt 2024), the compressor replaces encoding the image with a latent extracted by a pretrained analysis encoder. In this scenario, the training objective is replaced by:

$$\mathcal{L} = R_{\mathcal{E}(x)} + \lambda D(\mathcal{E}(x), \mathcal{M}(\mathcal{E}(x))) \quad (2)$$

Here, $\mathcal{E}(\cdot)$ represents the pretrained analysis encoder. This approach eliminates the need for designing intricate feature extraction modules for the compressor, while ensuring that the decoded features reside in the encoding space of the pre-trained generative model, thereby enhancing the quality of the final reconstructed image. Despite lightening the load on the compressor, this method introduces a potential inconsistency between distortions in the feature domain and distortions in the image domain, which could lead the compressor astray from the optimal rate-distortion trade-off solution.

Latent Diffusion Models

Diffusion models (DMs), or score-based generative models (Ho, Jain, and Abbeel 2020; Song et al. 2021), are a type of generative model that gradually injects Gaussian noise into the data and then generates samples from the noise through a reverse denoising process. To enable diffusion models to train on limited computing resources while preserving generation quality, Latent Diffusion Models (LDMs) (Rombach

et al. 2022) encode an image x into a latent representation z_0 via an encoder \mathcal{E} and reconstruct it using a decoder \mathcal{D} .

In LDMs' forward process, Gaussian noise is gradually added to the clean latent feature z_0 , with noise intensity at each step controlled by schedule β_t . This process is formulated as:

$$z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad t \in \{1, 2, \dots, T\}, \quad (3)$$

where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ is standard Gaussian noise. Here, $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$. As t increases, the corrupted z_t progressively approximates a Gaussian distribution. The reverse process of modeling LDMs is equivalent to training a noise prediction network ϵ_θ with the diffusion loss:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{z_0, t, \epsilon \sim \mathcal{N}(0, \mathbf{I})} \|\epsilon - \epsilon_\theta(z_t, c, t)\|_2^2. \quad (4)$$

During inference, LDMs predict noise using the pre-trained denoising network $\epsilon_\theta(z_t, c, t)$ with text condition c , generating latent z_{t-1} to sequentially obtain the final latent z_0 (As indicated by the grey dashed line in fig. 2). Stable Diffusion, as pre-trained LDMs, have been shown to have priors that can be widely applied in low-level image tasks (Lin et al. 2023; Wang et al. 2024; Li et al. 2024b).

Consistency Models

Consistency Models (CMs) (Song et al. 2023) use consistency mapping to directly map any point in an ODE trajectory back to its origin, enabling fast few-step generation.

Formally, this mapping is defined as $\mathbf{f} : (\mathbf{x}_t, t) \mapsto \mathbf{x}_\epsilon$ where ϵ is a fixed small positive number. A key property is self-consistency:

$$\mathbf{f}(\mathbf{x}_t, t) = \mathbf{f}(\mathbf{x}_{t'}, t'), \forall t, t' \in [\epsilon, T]. \quad (5)$$

To ensure $\mathbf{f}_\theta(\mathbf{x}, \epsilon) = \mathbf{x}$, the model is parameterized as:

$$\mathbf{f}_\theta(\mathbf{x}, t) = c_{\text{skip}}(t)\mathbf{x} + c_{\text{out}}(t)\mathbf{F}_\theta(\mathbf{x}, t), \quad (6)$$

where $c_{\text{skip}}(0) = 1, c_{\text{out}}(0) = 0$. During training, self-consistency is enforced using a target model θ^- —an exponential moving average (EMA) of θ : $\theta^- \leftarrow \mu\theta^- + (1-\mu)\theta$. The consistency loss is:

$$\begin{aligned} \mathcal{L}(\theta, \theta^-; \Phi) = \\ \mathbb{E}_{\mathbf{x}, t} \left[d \left(\mathbf{f}_\theta(\mathbf{x}_{t_{n+1}}, t_{n+1}), \mathbf{f}_{\theta^-}(\hat{\mathbf{x}}_{t_n}^\Psi, t_n) \right) \right]. \quad (7) \end{aligned}$$

Here, $d(\cdot, \cdot)$ is a distance metric. Ψ denotes a one-step ODE solver estimating \mathbf{x}_{t_n} from $\mathbf{x}_{t_{n+1}}$. (Luo et al. 2023) applies this to pre-trained LDMs with skip-step scheduling, ensuring consistency between non-adjacent steps $t_{n+k} \rightarrow t_n$ instead of $t_{n+1} \rightarrow t_n$.

Method

Overview Framework

The fig. 2 illustrates the overall framework of DiffCR. For a detailed algorithmic flow, please refer to the appendix. The image \mathbf{x} is first passed through the encoder $\mathcal{E}(\cdot)$ to obtain the latent representation \mathbf{z}_0 of the LDMs, which is then fed into the latent compressor for lossy compression, resulting in a bitstream and its decoded output $\hat{\mathbf{c}}$. $\hat{\mathbf{c}}$ guides the denoising network through the image control branch. Simultaneously, we mix the semantic information provided by the distorted image $\mathcal{D}(\hat{\mathbf{c}})$ and the textual semantics from \mathbf{x} , embedding them into the denoising network.

Next, DiffCR corrects and refines the output of the denoising network, ϵ , and image-level control $\hat{\mathbf{c}}$, through Frequency-aware Skip Estimation, mapping any input at any time ($\hat{\mathbf{z}}_0^t$) back to the origin status \mathbf{z}_0 ($\hat{\mathbf{z}}_0^t$). Unlike the sampling path of DDPM, DiffCR allows for two-step fast sampling. The final sampling result is obtained through the decoder $\mathcal{D}(\cdot)$ to yield the ultimate decoded output $\hat{\mathbf{x}}$.

Consistency Refinement for Latent Compression

Frequency-aware Skip Estimation The foundational diffusion model optimizes eq. (4) by infusing conditions into the denoising model, gradually enhancing the realism of generated images. However, this objective cannot serve as the reconstruction target for the compressor, manifesting as pattern collapse (Li et al. 2024a). Instead, existing methods align latent eq. (2) to train the compressor, yet the information in the latent stems from the priors of the encoder-decoder \mathcal{E} and \mathcal{D} , not the denoising network ϵ_θ . FaSE aims to establish a consistency estimator $\mathbf{f}_{\phi, \theta}(\mathbf{z}_t, \hat{\mathbf{c}}, t)$ that transforms the diffusion prior and compressed latent conditions into the same predictive target \mathbf{z}_0 :

$$\mathbf{f}_{\phi, \theta}(\mathbf{z}_t, \hat{\mathbf{c}}, t) = c_{\text{skip}}(t)\mathbf{z}_t + c_{\text{out}}(t)(\mathbf{F}_\phi(\epsilon_\theta(\mathbf{z}_t), \hat{\mathbf{c}}, t)). \quad (8)$$

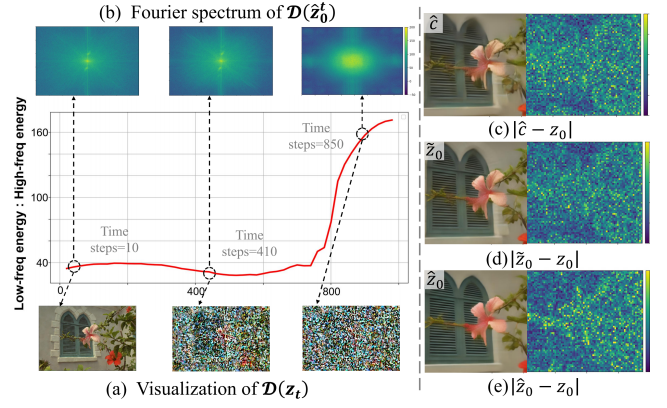


Figure 3: (Left) The trend of frequency energy variation in reconstruction by diffusion models at different time steps. (Right) Quantitative performance comparison of the compressor, denoising network, and FaSE in proceeding \mathbf{z}_0 -prediction.

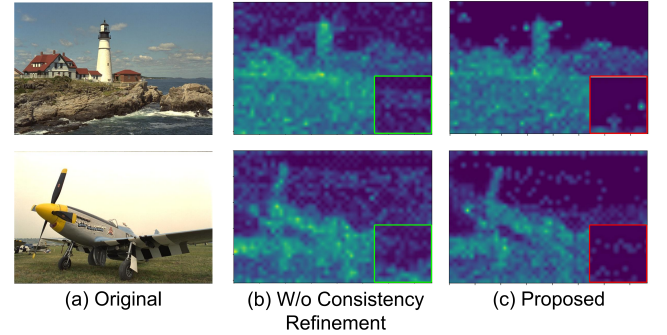


Figure 4: Visualization of bit-rate allocation on the Kodak.

Where the objective of $\mathbf{F}_\phi(\cdot, \cdot, \cdot)$ is to align the ϵ -Prediction prior with the image condition $\hat{\mathbf{c}}$, predicting the solution of the PF-ODE through consistency loss (Luo et al. 2023):

$$\begin{aligned} \mathcal{L}_C(\phi, \phi^-; \Psi) = \\ \mathbb{E}_{\mathbf{z}, \hat{\mathbf{c}}, n} \left[d \left(\mathbf{F}_\phi(\epsilon_\theta(\mathbf{z}_{t_{n+k}}), \hat{\mathbf{c}}, t_{n+1}), \mathbf{F}_{\phi^-}(\epsilon_\theta(\hat{\mathbf{z}}_{t_n}^\Psi), \hat{\mathbf{c}}, t_n) \right) \right]. \quad (9) \end{aligned}$$

Where Ψ represents the DDIM-solver (Song, Meng, and Ermon 2020), $d(\cdot, \cdot)$ denotes the mean squared error, and $k = 20$. In contrast to consistency distillation, $\mathbf{F}_\phi(\cdot, \cdot, \cdot)$ begins with zero initialization, necessitating the \mathbf{z}_0 -Prediction loss to stabilize the prediction process:

$$\mathcal{L}_F = w(t) \cdot \|\mathbf{F}_\phi(\epsilon_\theta(\mathbf{z}_t), \hat{\mathbf{c}}, t) - \mathbf{z}_0\|_2^2. \quad (10)$$

Here, $w(t) = \frac{1}{2} \left(\frac{\bar{\alpha}t-1}{1-\bar{\alpha}t-1} - \frac{\bar{\alpha}t}{1-\bar{\alpha}t} \right)$.

We refer to the predicted \mathbf{z}_0 obtained through ϵ -Prediction (eq. (3)) as $\hat{\mathbf{z}}_0$, and the prediction derived from FaSE as $\tilde{\mathbf{z}}_0$. The right side of fig. 3 illustrates three distinct patterns of \mathbf{z}_0 prediction. $\hat{\mathbf{c}}$ provides the roughest prediction (fig. 3(c)), with the predicted differences evenly distributed throughout the image. Although $\hat{\mathbf{z}}_0$ refines many textures using priors (fig. 3(e)), it exhibits significant inconsistencies

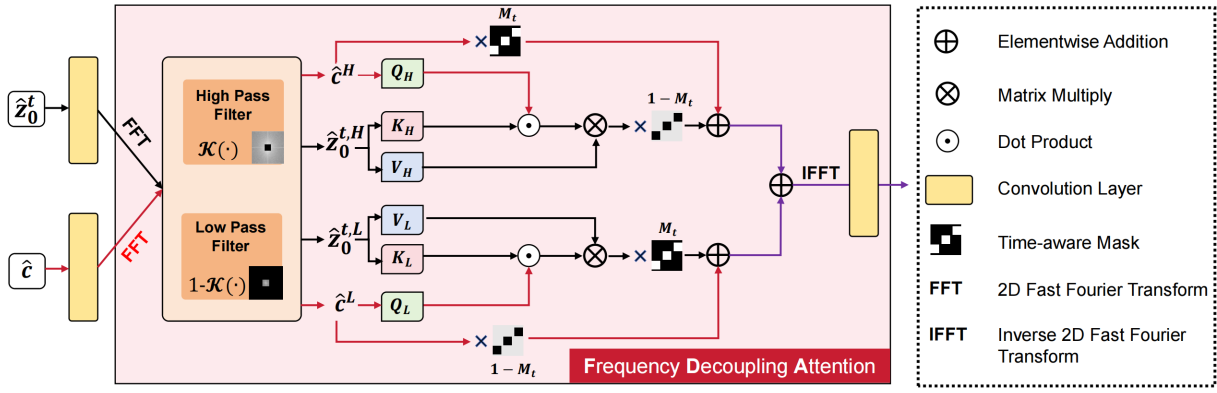


Figure 5: Illustration of the FDA module, where $Q.K.V.$ corresponds to Query, Key, and Value in the attention mechanism.

and distortions in specific details. In contrast, FaSE effectively aligns \hat{c} and \hat{z}_0 , combining detail consistency with texture richness (fig. 3(d)). Crucially, the joint training of the compressor with FaSE enables it to actively capture the attention of the diffusion prior in different regions to extract more compact compression representations. As shown in fig. 4, compared to not utilizing consistency refinement, the compressor tends to waste bits on flat, highly repetitive patterns in natural images (such as sky, grass, etc.), as shown in fig. 4(b). However, these common textures can be easily synthesized by the diffusion prior. In contrast, FaSE perfectly captures these prior signals, significantly optimizing the bit rate allocation of the compressor.

Frequency Decoupling Attention The signals recovered by the diffusion model at different time steps exhibit non-uniformity in the frequency domain. Specifically, in the early time steps, the denoising network tends to focus on recovering low-frequency signals, while high-frequency signals are synthesized more in later time steps (Qian et al. 2024; Yuan et al. 2025). Our experiments also confirm this observation, as shown on the left side of fig. 3. This implies that when designing $F_\phi(\cdot, \cdot, \cdot)$, we need to consider the alignment of \hat{z}_0^t and \hat{c} in different frequency domains for different time steps t . To address this, we introduce Frequency-Decoupling Attention (FDA), as illustrated in fig. 5. Specifically, FDA first transforms \hat{z}_0^t and \hat{c} to the Fourier domain using fast Fourier transform $FFT(\cdot)$, and then obtains the high and low-frequency components of both through high-pass and low-pass filters:

$$\begin{aligned} \hat{c}^H &= FFT(\hat{c}) \odot \mathcal{K}, \hat{c}^L = FFT(\hat{c}) \odot (1 - \mathcal{K}), \\ \hat{z}_0^{t,H} &= FFT(\hat{z}_0^t) \odot \mathcal{K}, \hat{z}_0^{t,L} = FFT(\hat{z}_0^t) \odot (1 - \mathcal{K}). \end{aligned} \quad (11)$$

We then employ two separate cross-attention mechanisms to individually match the high and low-frequency components:

$$\begin{aligned} Crossattn^H(\hat{z}_0^{t,H}, \hat{c}^H) &= \text{Softmax}\left(\frac{Q^H K^H}{\sqrt{d}}\right) \cdot V^H, \\ Crossattn^L(\hat{z}_0^{t,L}, \hat{c}^L) &= \text{Softmax}\left(\frac{Q^L K^L}{\sqrt{d}}\right) \cdot V^L. \end{aligned} \quad (12)$$

Subsequently, we modulate the contribution of the prior \hat{z}_0^t in the prediction using a linear temporal mask, following

best practices and previous research experience, where the proportion of the prior’s high-frequency content in the mixture increases gradually as the time step approaches zero. Taking the high-frequency component as an example:

$$\hat{z}_0^{t,H,o} = C attn^H(\hat{z}_0^{t,H}, \hat{c}^H) * M_t + \hat{c}^H * (1 - M_t). \quad (13)$$

Where $M_t = 1 - t/T$, with T being the total number of timesteps in the diffusion model. We present ablation experiments on this module in table 2, demonstrating that compared to directly using cross-attention, FDA achieves significant performance gains through the temporal attention mechanism and frequency decoupling.

Control Foundation Diffusion with Hybrid Control

Inspired by (Xia et al. 2025), we integrate image control branch reconstructed by the compressor and semantic control branch extracted by CLIP (Radford et al. 2021). These are injected into the denoising network via control encoder (Zhang, Rao, and Agrawala 2023) and cross-attention (fig. 2). The image control enhances edge and texture generation for better fidelity, while the semantic control stabilizes generation and corrects color shifts caused by denoising.

For image control, an end-to-end compressor performs latent-space compression and reconstruction. Feature z_0 is downsampled to y by g_a . A VQ-based categorical hyper-prior model (Jia et al. 2024) extracts and compresses hyper-priors, which are then used with the SCCTX entropy model (He et al. 2022) for \hat{y} ’s entropy estimation. Finally, g_s decodes \hat{y} to obtain image-level control \hat{c} .

For the semantic control, an image captioning model (Wang et al. 2022) extracts image textual descriptions (transmission cost ≤ 0.0001 bpp). Two Clip-Encoders extract image and textual semantics separately; these embeddings are modulated via a resampling layer (Ye et al. 2023) and injected into the denoising network. Compared to sole textual control, such multimodal semantics enhance decoding performance without transmission redundancy.

Training Strategy

Our model is trained in two stages. In the first stage of training, the compressor, control module, and skip estimation

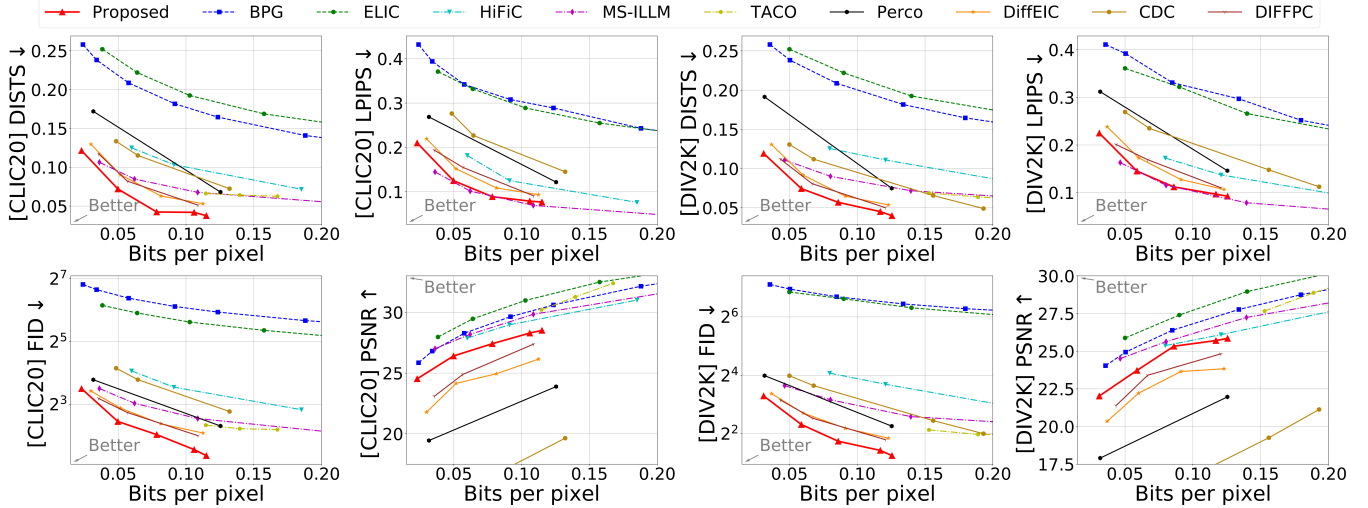


Figure 6: Qualitative comparison on the CLIC20 and DIV2K datasets. Grey arrows indicate the direction of better performance.

module are jointly trained. While optimizing the compressor’s ability to extract representations, the guiding capabilities of the control module and skip estimation are enhanced. For the compressor, we need the rate-distortion loss to balance distortion and compression rate:

$$\mathcal{L}_{LC} = \lambda_1 \|z_0 - \hat{c}\|_2^2 + \lambda_2 R(\hat{y}) + \mathcal{L}_{\text{codebook}}, \quad (14)$$

where λ_1 and λ_2 are used to balance distortion and bit-rate, and the definition of $\mathcal{L}_{\text{codebook}}$ is:

$$\mathcal{L}_{\text{codebook}} = \left\| \text{sg}(\mathbf{l}_y) - \hat{\mathbf{l}}_y \right\|_2^2 + \beta \left\| \text{sg}(\hat{\mathbf{l}}_y) - \mathbf{l}_y \right\|_2^2, \quad (15)$$

where $\text{sg}(\cdot)$ denotes the stop-gradient operator and $\beta = 0.25$.

For the skip estimation module $F_\phi(\cdot, \cdot, \cdot)$, we need to maintain its self-consistency while preserving its functionality for z_0 -Prediction:

$$\mathcal{L}_{\text{FaSE}} = \lambda_1 \mathcal{L}_F + \mathcal{L}_C(\phi, \phi^-; \Psi). \quad (16)$$

Since \mathcal{L}_F effectively facilitates the optimization of the compressor’s distortion, we select hyperparameters that are in line with the distortion settings of the compressor.

For the control module, we need to optimize the original diffusion loss eq. (4) to integrate it into the denoising net:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{z_0, \hat{c}, s, t, \epsilon \sim \mathcal{N}(0, \mathbf{I})} \|\epsilon - \epsilon_\theta(z_t, \hat{c}, s, t)\|_2^2. \quad (17)$$

The total loss for the first stage is:

$$\mathcal{L}_{1\text{st}} = \mathcal{L}_{LC} + \mathcal{L}_{\text{FaSE}} + \lambda_3 \mathcal{L}_{\text{diff}}. \quad (18)$$

In the second stage, our main focus is on training $F_\phi(\cdot, \cdot, \cdot)$ on the sampling process patterns. We freeze the parameters of the compressor and use $F_\phi(\cdot, \cdot, \cdot)$ for *two-step* sampling. We map the sampling results back to the image domain to calculate the reconstruction loss compared to the original image. Additionally, we must ensure that the control module and $F_\phi(\cdot, \cdot, \cdot)$ receive proper guidance. Specifically:

$$\mathcal{L}_{2\text{st}} = \mathcal{L}_{\text{per}}(\mathcal{D}(\hat{z}_0), \mathbf{x}) + \mathcal{L}_{\text{FaSE}} + \mathcal{L}_{\text{diff}}. \quad (19)$$

\mathcal{L}_{per} represents perceptual loss, and in this case, we are using LPIPS (Zhang et al. 2018).

Experiments

Experimental Setup

Datasets. During the validation phase, we referred to (Muckley et al. 2023; Yang and Mandt 2024) and utilized four widely recognized image compression benchmark datasets: CLIC20 (George Toderici 2020), DIV2K (Timofte et al. 2017), Kodak (Company 2013), and Tecnick (Asuni and Giachetti 2013). CLIC20 comprises 428 high-definition images. To ensure more reliable statistical fidelity validation within DIV2K, we selected the initial 400 images from DIV2K-train and 100 images from DIV2K-valid, resulting in a total of 500 images for testing.

Baselines. We shall compare our proposed DiffCR against a myriad of representative image compression methods, encompassing traditional compression standard such as BPG (Bellard 2014); VAE-based compression method ELIC (He et al. 2022); GAN-driven compression techniques including HiFiC (Mentzer et al. 2020), MS-ILLM (Muckley et al. 2023), and TACO (Lee et al. 2024); alongside diffusion-based approaches such as CDC (Yang and Mandt 2024), PerCo (Careil et al. 2024), DiffeIC (Li et al. 2024a), DiffPC (Xia et al. 2025), and ResULIC (Ke et al. 2025). Further details can be perused in the appendix. For certain baselines with undisclosed functionalities (Ke et al. 2025), we solely contrasted their publicly available data points on the Kodak.

Metrics. We employ a variety of widely acknowledged metrics to gauge the visual quality of reconstructed images, including perceptual metrics such as LPIPS (Zhang et al. 2018), DISTS (Ding et al. 2020), FID (Heusel et al. 2017), and KID (Bińkowski et al. 2018). Additionally, we utilize distortion metrics like PSNR and MS-SSIM (Wang, Simoncelli, and Bovik 2003) to assess the fidelity of reconstruction. Unless otherwise specified, LPIPS computations are performed using the Alex network, and we also present validation results of LPIPS-V metrics calculated using the VGG network. For the Kodak and Tecnick datasets, due to insuf-

Model	Encoding(s)	Decoding(s)	BD-rate(%)
			DISTS
ELIC	0.009	0.008	1017.3
HiFiC	0.012	0.062	376.04
MS-ILLM	0.069	0.068	125.25
CDC	0.007	3.081	116.47
Perco	0.131	7.455	90.757
DiffEIC	0.402	6.618	28.062
DiffPC	0.089	7.325	26.831
ResULIC	181.16*	1.021	16.904
Ours	0.076	0.481	0

Table 1: Encoding and decoding time on Kodak dataset. A higher BD-rate indicates poorer performance. * The values reported in the (Ke et al. 2025) paper are utilized.

efficient data for statistical fidelity metric calculations, we do not validate FID and KID metrics on these datasets. Further details can be found in the appendix.

Main Result

Figure 6 illustrates the quantitative comparison between DiffCR and baseline methods on CLIC20 and DIV2K. It is evident that our proposed approach excels in perceptual quality, matching the state-of-the-art solution MS-ILLM in terms of the LPIPS metric and outperforming all diffusion baselines consistently. Particularly, in the appendix, we showcase two metrics, KID and LPIPS-V, demonstrating the consistent superiority of DiffCR over all baselines. Noteworthy is the fact that the enhancement in perceptual metrics does not come at the cost of increased distortion: DiffCR significantly outperforms other diffusion baselines in terms of PSNR, even demonstrating competitiveness with HiFiC on the DIV2K dataset. This clearly indicates that DiffCR achieves a superior distortion-perception trade-off from a compression perspective, rather than solely relying on diffusion model priors for image enhancement. In the supplementary materials, we present qualitative visual comparisons to further elucidate the performance advantages of DiffCR.

Complexity Comparisons Table 1 presents the encoding and decoding latency as well as performance comparison between DiffCR and baseline methods on the Kodak dataset. The performance metric utilized is BD-rate (Bjontegaard 2001). Taking DiffCR as the reference model, a higher BD-rate indicates that the model requires more bits to achieve the same metric as the reference. DiffCR not only boasts satisfactory encoding speed but also, due to its requirement for only two sampling steps, exhibits decoding speeds over $10\times$ faster than the typical diffusion baselines that employ 50 steps of sampling and still outperforms the 4-step decoding of (Ke et al. 2025) by $2\times$.

S.Emb.	CRE	FDA	2-stage	BD-rate(%)	
				LPIPS	FID
✓	✓	✓		44.5	51.1
✓			✓	38.24	41.21
✓	✓		✓	13.52	15.64
	✓	✓	✓	8.32	10.11
✓	✓	✓	✓	0	0

Table 2: Ablation studies comparing each module on the CLIC20 dataset

Ablations

We conducted ablation experiments on the CLIC20, using DiffCR as a baseline to showcase the BD-rate from different ablations, as depicted in table 2.

Consistency Refinement Estimator (CRE). In fig. 4, we demonstrate the advantage of CRE in optimizing bit rate allocation. As shown in table 2, removing CRE leads to a performance degradation of up to 40%.

Frequency Decoupling Attention (FDA). We replaced the FDA module with cross-attention layers. As discussed in previous sections, standard cross-attention struggles to capture the intricate context of diffusion priors and compressed latents across different time steps, resulting in a performance decline of approximately 14%.

Second Stage Training (2-stage). Directly training the model in the second stage to learn sampling patterns from image-domain distortion. Omitting this training stage results in a performance drop of around 47%.

Semantic Embedding (S.Emb.). Replacing mixed semantic embeddings with pure textual semantic embeddings leads to a 9% performance decline and introduces slight color deviations in decoded images.

We also discussed the model’s sensitivity to text extracted by different image-captioning models and conducted ablation experiments on the sampling steps. Please refer to the appendix for further details.

Conclusions

In this paper, we introduce DiffCR, a high-fidelity, low-bitrate image compression framework that leverages consistency refinement diffusion priors. Through extensive experimentation, we demonstrate the significant contributions of the frequency-aware skip prediction module in optimizing bit-rate allocation and balancing perception-distortion trade-offs. DiffCR exhibits satisfactory perceptual quality and decoding efficiency in low-bitrate scenarios. Moreover, the flexible model architecture showcases promising prospects for practical applications.

Acknowledgements

This work is supported in part by the National Natural Science Foundation of China under grant 624B2088, 62301189, 62571298, and Shenzhen Science and Technology Program under Grant KJZD20240903103702004, SYSPG20241211173609009.

References

- Agustsson, E.; Minnen, D.; Toderici, G.; and Mentzer, F. 2023. Multi-realism image compression with a conditional generator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22324–22333.
- Asuni, N.; and Giachetti, A. 2013. TESTIMAGES: A large data archive for display and algorithm testing. *Journal of Graphics Tools*, 17(4): 113–125.
- Ballé, J.; Laparra, V.; and Simoncelli, E. P. 2016. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704*.
- Ballé, J.; Minnen, D.; Singh, S.; Hwang, S. J.; and Johnston, N. 2018. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*.
- Bellard, F. 2014. BPG image format. <https://bellard.org/bpg/>.
- Bińkowski, M.; Sutherland, D. J.; Arbel, M.; and Gretton, A. 2018. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*.
- Bjontegaard, G. 2001. Calculation of average PSNR differences between RD-curves. *ITU SG16 Doc. VCEG-M33*.
- Blau, Y.; and Michaeli, T. 2019. Rethinking lossy compression: The rate-distortion-perception tradeoff. In *International Conference on Machine Learning*, 675–685. PMLR.
- Careil, M.; Muckley, M. J.; Verbeek, J.; and Lathuilière, S. 2024. Towards image compression with perfect realism at ultra-low bitrates. In *The Twelfth International Conference on Learning Representations*.
- Cheng, Z.; Sun, H.; Takeuchi, M.; and Katto, J. 2020. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7939–7948.
- Company, E. K. 2013. Kodak lossless true color image suite.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Ding, K.; Ma, K.; Wang, S.; and Simoncelli, E. P. 2020. Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence*, 44(5): 2567–2581.
- George Toderici, W. S. 2020. Workshop and Challenge on Learned Image Compression (CLIC2020).
- He, D.; Yang, Z.; Peng, W.; Ma, R.; Qin, H.; and Wang, Y. 2022. Elic: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5718–5727.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hoogeboom, E.; Agustsson, E.; Mentzer, F.; Versari, L.; Toderici, G.; and Theis, L. 2023. High-fidelity image compression with score-based generative models. *arXiv preprint arXiv:2305.18231*.
- Jia, Z.; Li, J.; Li, B.; Li, H.; and Lu, Y. 2024. Generative latent coding for ultra-low bitrate image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26088–26098.
- Ke, A.; Zhang, X.; Chen, T.; Lu, M.; Zhou, C.; Gu, J.; and Ma, Z. 2025. Ultra Lowrate Image Compression with Semantic Residual Coding and Compression-aware Diffusion. *arXiv preprint arXiv:2505.08281*.
- Lee, H.; Kim, M.; Kim, J.-H.; Kim, S.; Oh, D.; and Lee, J. 2024. Neural Image Compression with Text-guided Encoding for both Pixel-level and Perceptual Fidelity. In *International Conference on Machine Learning*.
- Lei, E.; Uslu, Y. B.; Hassani, H.; and Bidokhti, S. S. 2023. Text+ sketch: Image compression at ultra low rates. *arXiv preprint arXiv:2307.01944*.
- Li, Z.; Zhou, Y.; Wei, H.; Ge, C.; and Jiang, J. 2024a. Towards Extreme Image Compression with Latent Feature Guidance and Diffusion Prior. *IEEE Transactions on Circuits and Systems for Video Technology*.
- Li, Z.; Zhou, Y.; Wei, H.; Ge, C.; and Mian, A. 2024b. Diffusion-based Extreme Image Compression with Compressed Feature Initialization. *arXiv preprint arXiv:2410.02640*.
- Lin, X.; He, J.; Chen, Z.; Lyu, Z.; Dai, B.; Yu, F.; Ouyang, W.; Qiao, Y.; and Dong, C. 2023. Diffbir: Towards blind image restoration with generative diffusion prior. *arXiv preprint arXiv:2308.15070*.
- Liu, X.; Chen, B.; Liu, Z.; Wang, Y.; and Xia, S.-T. 2025a. An Exploration with Entropy Constrained 3D Gaussians for 2D Video Compression. In *The Thirteenth International Conference on Learning Representations*.
- Liu, X.; Chen, J.; Chen, B.; Liu, Z.; An, B.; Xia, S.-T.; and Wang, Z. 2025b. An efficient implicit neural representation image codec based on mixed autoregressive model for low-complexity decoding. *IEEE Transactions on Multimedia*.
- Lu, L.; Xie, Y.; Jiang, W.; Wang, W.; Lin, X.; and Wang, Y. 2024. HybridFlow: Infusing Continuity into Masked Codebook for Extreme Low-Bitrate Image Compression. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 3010–3018.
- Luo, S.; Tan, Y.; Huang, L.; Li, J.; and Zhao, H. 2023. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*.
- Mao, Q.; Yang, T.; Zhang, Y.; Wang, Z.; Wang, M.; Wang, S.; Jin, L.; and Ma, S. 2024. Extreme image compression

- using fine-tuned vqgans. In *2024 Data Compression Conference (DCC)*, 203–212. IEEE.
- Mentzer, F.; Toderici, G. D.; Tschannen, M.; and Agustsson, E. 2020. High-fidelity generative image compression. *Advances in Neural Information Processing Systems*, 33: 11913–11924.
- Minnen, D.; Ballé, J.; and Toderici, G. D. 2018. Joint autoregressive and hierarchical priors for learned image compression. *Advances in neural information processing systems*, 31.
- Muckley, M. J.; El-Nouby, A.; Ullrich, K.; Jégou, H.; and Verbeek, J. 2023. Improving statistical fidelity for neural image compression with implicit local likelihood models. In *International Conference on Machine Learning*, 25426–25443. PMLR.
- Qian, Y.; Cai, Q.; Pan, Y.; Li, Y.; Yao, T.; Sun, Q.; and Mei, T. 2024. Boosting diffusion models with moving average sampling in frequency domain. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8911–8920.
- Qin, S.; Chen, B.; Huang, Y.; An, B.; Dai, T.; and Xia, S.-T. 2023. Perceptual image compression with cooperative cross-modal side information. *arXiv preprint arXiv:2311.13847*.
- Qin, S.; Wang, J.; Zhou, Y.; Chen, B.; Luo, T.; An, B.; Dai, T.; Xia, S.; and Wang, Y. 2024. Mambavc: Learned visual compression with selective state spaces. *arXiv preprint arXiv:2405.15413*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Shannon, C. E. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3): 379–423.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2023. Consistency Models. In *International Conference on Machine Learning*, 32211–32252. PMLR.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Timofte, R.; Agustsson, E.; Van Gool, L.; Yang, M.-H.; Zhang, L.; Lim, B.; et al. 2017. NTIRE 2017 Challenge on Single Image Super-Resolution: Methods and Results. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Wallace, G. K. 1992. The JPEG still picture compression standard. *IEEE transactions on consumer electronics*, 38(1): xviii–xxxiv.
- Wang, J.; Yue, Z.; Zhou, S.; Chan, K. C.; and Loy, C. C. 2024. Exploiting Diffusion Prior for Real-World Image Super-Resolution. *International Journal of Computer Vision*.
- Wang, P.; Yang, A.; Men, R.; Lin, J.; Bai, S.; Li, Z.; Ma, J.; Zhou, C.; Zhou, J.; and Yang, H. 2022. OFA: Unifying Architectures, Tasks, and Modalities Through a Simple Sequence-to-Sequence Learning Framework. *CoRR*, abs/2202.03052.
- Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; and Change Loy, C. 2018. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, 0–0.
- Wang, Z.; Simoncelli, E. P.; and Bovik, A. C. 2003. Multi-scale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, 1398–1402. Ieee.
- Xia, Y.; Zhou, Y.; Wang, J.; An, B.; Wang, H.; Wang, Y.; and Chen, B. 2025. DiffPC: Diffusion-based High Perceptual Fidelity Image Compression with Semantic Refinement. In *The Thirteenth International Conference on Learning Representations*.
- Yang, R.; and Mandt, S. 2024. Lossy image compression with conditional diffusion models. *Advances in Neural Information Processing Systems*, 36.
- Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Yuan, Y.; Guo, Y.; Wang, C.; Zhang, W.; Xu, H.; and Zhang, L. 2025. Freqprior: Improving video diffusion models with frequency filtering gaussian noise. *arXiv preprint arXiv:2502.03496*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.