

Laplacian-guided Entropy Model in Neural Codec with Blur-dissipated Synthesis

Atefeh Khoshkhahtinat, Ali Zafari, Piyush M. Mehta, Nasser M. Nasrabadi
 West Virginia University, WV, USA

{ak00043, az00004}@mix.wvu.edu, {piyush.mehta, nasser.nasrabadi}@mail.wvu.edu

Abstract

While replacing Gaussian decoders with a conditional diffusion model enhances the perceptual quality of reconstructions in neural image compression, their lack of inductive bias for image data restricts their ability to achieve state-of-the-art perceptual levels. To address this limitation, we adopt a non-isotropic diffusion model at the decoder side. This model imposes an inductive bias aimed at distinguishing between frequency contents, thereby facilitating the generation of high-quality images. Moreover, our framework is equipped with a novel entropy model that accurately models the probability distribution of latent representation by exploiting spatio-channel correlations in latent space, while accelerating the entropy decoding step. This channel-wise entropy model leverages both local and global spatial contexts within each channel chunk. The global spatial context is built upon the Transformer, which is specifically designed for image compression tasks. The designed Transformer employs a Laplacian-shaped positional encoding, the learnable parameters of which are adaptively adjusted for each channel cluster. Our experiments demonstrate that our proposed framework yields better perceptual quality compared to cutting-edge generative-based codecs, and the proposed entropy model contributes to notable bitrate savings. The code is available at <https://github.com/Atefeh-Khoshtinat/Blur-dissipated-compression>.

1. Introduction

Image compression is a crucial task in image processing which aims to decrease the amount of data required for storing or transmitting without significant loss of visual content. Recently, learning-based image compression methods [13, 25, 42, 45] have demonstrated the potential to surpass classical hand-engineered codecs in terms of rate-distortion performance. Learned image compression commonly comprises three steps: transformation, quantization, and lossless entropy coding, which resemble the components found in the traditional transform coding paradigm [44].

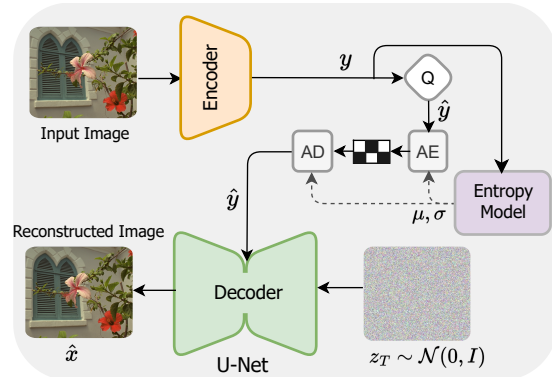


Figure 1. Overview of our proposed neural codec. The quantized semantic latent variable \hat{y} is utilized by a diffusion-based decoder to generate realistically reconstructed image.

Neural image compression methods are typically trained with two primary objectives: minimizing distortion between the original input image and its reconstructed version and reducing the bitrate for efficient data transmission or storage [38]. However, solely optimizing for rate-distortion may lead to blurry reconstructions. To address this issue, various generative-based codecs [4, 11, 28, 42] have emerged, aiming to generate reconstructions that are not only faithful to the original data but also visually realistic to human observers. [28] introduced HiFiC, an image compression network based on GANs that aims to achieve a balance between fidelity and perceptual quality. The network demonstrated substantial advantages for human observers, surpassing BPG even when using half the bits. However, HiFiC faces challenges due to the instability of GAN training and necessitates various design adjustments. Inspired by the achievements of Denoising Diffusion Probabilistic Models (DDPMs)[17] in generative modeling, authors in [42] replaced the Gaussian decoder of VAEs with a conditional diffusion model to mitigate the issue of blurriness in decoded images. Despite the impressive performance of diffusion-based network, its diffusion and denoising processes do not explicitly incorporate the inductive biases inherent to natural images. The imposing inductive biases,

such as the multi-scale nature of images, enable the generation of high-quality samples [20].

In addition to enhancing the quality of reconstructed images, accurate entropy estimation of latent representations is vital for boosting compression efficiency. The more the estimated entropy gets closer to the actual entropy of the latent, the lower the bitrate required for generating bitstream file. In this respect, several works have been proposed. Ballé *et al.*[5] introduced an entropy model that is conditioned on an additional latent variable, named the hyper-prior, to capture the existing spatial redundancies within the latent space. Studies [24, 30] introduced a sequential autoregressive block into the entropy model. This block leverages causally adjacent latent elements to estimate the distribution of the current latent, which results in slowing down of the decoding step. He *et al.*[13] utilized a two-way parallel context model as a substitute for the sequential autoregressive context model, aiming to accelerate the decoding of latent code. In another realm of studies, to mitigate the slow decoding times, authors in [29] introduced channel-wise context modeling as an alternative to serial spatial context. To expedite the extracting of the channel-conditional context model, He *et al.* [14] exploited an unevenly segmented channel-wise autoregressive model, wherein the channels are divided into varying sizes. Strengthening the entropy model through self-attention was initially proposed in [32]. This approach employed a global reference model to identify the most suitable spatial context, which doesn't have to be confined to the local scope. The idea of integrating global context into the entropy model was subsequently refined in a later work [33], where a specially crafted transformer-based entropy model was employed. This entropy model is equipped with the diamond-relative position encoding to effectively model long-distance dependencies. The diamond-shaped relative positional encoding is devised based on the clip function, which relies on the pre-defined threshold to characterize an effective receptive field. This approach is effective in computing context based on the positions of latent elements; yet, it cannot provide an optimal receptive field due to its reliance on the clip function. Recent works[21, 23] focused on developing adaptive context models to extract both channel and spatial correlations within the latent space while expediting the entropy decoding step. Authors in [21] proposed an entropy model named MEM, aiming to capture both local and global spatial correlations in each channel segment by using a parallel bidirectional context approach. However, despite MEM's attention to spatio-channel correlations, it appears to overlook the crucial role of positional encoding in computing long-range spatial dependencies.

To address the aforementioned challenges, we propose a conditional diffusion-based decoder in the neural compression pipeline to enhance the realism of reconstructed

images. The traditional diffusion models are generalized into a conditional non-isotropic diffusion model to incorporate an inductive bias, considering the relative importance of each frequency component of images. As a result, each frequency component of an image undergoes diffusion at distinct rates, leading to the generation of decoded images in a coarse-to-fine manner. Furthermore, a novel channel-conditional autoregressive entropy model is introduced to efficiently take into account both local and global spatial dependencies within each channel chunk of the latent space. To leverage global spatial context, a Transformer block is designed which incorporates a Laplacian-shaped positional encoding within checkerboard-shaped self-attention module. The effective Laplacian-shaped receptive field for each channel chunk is dynamically determined by learning the positional encoding parameters during the optimization of the entire neural codec.

2. Related Works

2.1. Learned Image Compression

Neural image compression methods adopt a non-linear transform coding paradigm which is founded on the variational autoencoders (VAEs) [12]. Within this scheme, the encoder initially maps the input image x to a compact latent representation y . Subsequently, quantization is applied to this latent representation, leading to discrete symbols that minimize the bit requirement. Afterward, entropy coding generates bit strings using an entropy model. Ultimately, the reconstructed image \hat{x} is generated by applying the decoder to the quantized latent representation \hat{y} . This framework is commonly trained with a trade-off between the rate and distortion, such as Mean Squared Error (MSE) [38].

Recent studies have attempted to enhance the realism of reconstructed images by optimizing neural codecs using a triple rate-distortion-perception loss [4, 28, 43]. Blau and Michaeli [7] explored the essential balance between distortion and realism, demonstrating that within a given rate, improving distortion comes at the cost of decreasing the perceptual quality of a reconstructed image. In line with this premise where perceptual quality is measured as the difference between the image distribution and the distribution of reconstructions, Mentzer *et al.* [28] employed a conditional GAN within an autoencoder-based compression model to enhance perceptual quality. In a similar vein, [15] introduced a model that aimed to improve the realism by enriching the loss function with an adversarial perceptual loss using LPIPS [46], along with a patch-based style loss [10]. Recently, Agustsson *et al.* [4] developed a Multi-Realism model that merges the ELIC codec [14] and PatchGAN [8], showing improved performance compared to HiFiC.

2.2. Diffusion Models

Diffusion models [17, 22, 37], which belong to the family of score-based generative models, acquire the data distribution through a gradual iterative denoising process, starting from a Gaussian distribution and progressing towards the actual data distribution. Recently, they have received considerable attention owing to their training stability and high quality image generation compared to GANs. Diffusion models are adopted in various domains, including image and video generation [18, 19, 34], super-resolution [36], inpainting [27], deblurring [41], compression [11, 39, 42].

In the domain of image compression, several diffusion-based codecs have been proposed. The research conducted by Theis *et al.* [39] involved employing a general unconditional diffusion model for transmitting Gaussian samples in a lossy manner. Their approach leverages a reverse channel coding concept, which stands apart from the conventional transform coding scheme. Although their methodology performs well, it suffers from high computational cost which makes it unfeasible for dealing with high-resolution images. Yang and Mandt [42] introduced a neural codec, called CDC, in which the decoder takes the form of a DDPM, conditioned on a quantized latent representation. Ghose *et al.* [11] began by optimizing an autoencoder network using a rate-distortion loss. Subsequently, they train a conditional diffusion model on the output of the decoder, with the goal of enhancing its perceptual quality.

2.3. Neural Entropy Model

The main purpose of the entropy model is to estimate the joint probability distribution over the quantized latent representation. When the learned entropy model precisely matches the true distribution of the latent representation, a lower bit-rate is required to generate a compressed file. Entropy estimation can take advantage of two key principles: backward and forward adaptation. Forward adaptation uses an extra latent variable, called hyper-prior [5], as a side information to capture spatial dependencies between elements of the latent representation. Backward adaptation, on the other hand, employs an autoregressive-based context model which leverages the previously decoded elements of the latent to predict the probability of the current element. The context model of backward methods can extract relationships between the current symbol and previously decoded symbols in various dimensions, including spatial and channel axes [13, 14, 21, 23, 29, 30, 32, 33, 47].

Minnen *et al.*[30] incorporated a local spatial context with a hyper-prior network to precisely predict distribution of the latent. This type of context model can exploit local correlations between the current latent and its neighboring causal elements using masked convolutions, which consequently necessitates the adoption of serial decoding. To expedite the decoding process, He *et al.*[13] evenly parti-

tioned the latent into two groups: anchors and non-anchors. The anchors are exclusively decoded using the hyper-prior, while the non-anchors utilize both the hyper-prior and the local context model. This approach allows for the parallel decoding of both anchors and non-anchors.

An alternative strategy for parallelizing the decoding process is to extract inter-channel dependencies. In [29] the latent code is divided into evenly sized chunks along the channel dimension, with each of these chunks being decoded using information from previously decoded ones. To speed up the channel-wise context extraction, He *et al.*[14] partitioned channels of the latent into an uneven set of groups, allocating fewer channels to the initial groups and more channels to the subsequent ones. This uneven division is motivated by the observation that the earlier channels possess higher entropy compared to the remaining channels.

Some works introduced adaptive context models which focus on capturing long-range spatial correlations. Qian *et al.*[32] devised a global reference context model which assesses throughout previously decoded elements to detect the most similar one to the current latent. The authors in [33] employed a Transformer-based entropy model, called Entroformer, for capturing long-range contexts. The Entroformer benefits from diamond-relative position encoding and a top-k self-attention mechanism, both of which contribute to providing an efficient receptive field.

3. Methods

In this work, we develop our decoder based on the blurring diffusion model [20], which is conditioned on the discrete latent representation and enables the differentiation among the frequency components of an image throughout the diffusion and denoising processes. The architecture of decoder will be explained in Appendix. Moreover, a novel entropy model is employed that efficiently encodes the quantized latent representation into a binary stream.

3.1. Blurring Diffusion Model

We can define the blurring diffusion model as a traditional diffusion model in the frequency space, but with distinct schedules for each dimension of data [20]. The specifics of the schedules for vectors α_t and σ_t are crucial in this type of model, which will be discussed.

Diffusion Process: The diffusion process [17] progressively degrades the clean image, transforming it into nearly pure Gaussian noise over the course of T time steps. Each step of the diffusion process in frequency space can be formulated as:

$$q(\mathbf{f}_t|\mathbf{f}_x) = \mathcal{N}(\mathbf{f}_t; \alpha_t \mathbf{f}_x, \sigma_t^2 \mathbf{I}), \quad (1)$$

where \mathbf{f}_t and \mathbf{f}_x are defined as $\mathbf{V}^T \mathbf{z}_t$ and $\mathbf{V}^T \mathbf{x}$, respectively. In such a model, each frequency component could

undergo diffusion at varying rates, with the modulation of these rates being determined by the values present in the vectors α_t and σ_t .

Denoising Process: The true denoising process in the frequency space, which is equal to the deblurring operation in the time domain, can be expressed as:

$$q(\mathbf{f}_{t-1}|\mathbf{f}_t, \mathbf{f}_x) = \mathcal{N}(\mathbf{f}_{t-1}; \boldsymbol{\mu}_{t \rightarrow t-1}, \boldsymbol{\sigma}_{t \rightarrow t-1}^2 \mathbf{I}), \quad (2)$$

where $\boldsymbol{\sigma}_{t \rightarrow t-1} = \boldsymbol{\sigma}_{t|t-1} \boldsymbol{\sigma}_{t-1} / \boldsymbol{\sigma}_t$ and $\boldsymbol{\mu}_{t \rightarrow t-1} = (\boldsymbol{\alpha}_{t|t-1} \boldsymbol{\sigma}_{t-1}^2 / \boldsymbol{\sigma}_t^2) \mathbf{f}_t + (\boldsymbol{\alpha}_{t-1} \boldsymbol{\sigma}_{t|t-1}^2 / \boldsymbol{\sigma}_t^2) \mathbf{f}_x$. The actual denoising process can be approximated using a learned denoising distribution to generate data. Thus, by considering the epsilon reparametrization of Eq. 2, we can derive the expression for the learned denoising process in frequency space:

$$p_\theta(\mathbf{f}_{t-1}|\mathbf{f}_t) = \mathcal{N}(\mathbf{f}_{t-1}; \boldsymbol{\mu}_{t \rightarrow t-1}(\hat{\mathbf{f}}_x, \mathbf{f}_t), \boldsymbol{\sigma}_{t \rightarrow t-1}^2 \mathbf{I})$$

$$\text{Re-parametrization : } \hat{\mathbf{f}}_x = (1/\alpha_t)(\mathbf{f}_t - \sigma_t \hat{\mathbf{f}}_\epsilon). \quad (3)$$

By substituting the approximation of $\hat{\mathbf{f}}_x$ into the true denoising distribution, the mean of the denoising network can be derived by $(\alpha_{t|t-1} \boldsymbol{\sigma}_{t-1}^2 / \boldsymbol{\sigma}_t^2) \mathbf{f}_t + (\alpha_{t-1} \boldsymbol{\sigma}_{t|t-1}^2 / \boldsymbol{\sigma}_t^2 \alpha_t)(\mathbf{f}_t - \sigma_t \hat{\mathbf{f}}_\epsilon)$.

Optimization: According to [17], the loss function is simplified as bellow:

$$\mathbf{E}_{t,x,\epsilon}[\|\mathbf{f}_\epsilon - \hat{\mathbf{f}}_\epsilon\|^2] = \mathbf{E}[\|\mathbf{V}^T(\epsilon - \hat{\epsilon})\|^2] \approx \mathbf{E}_{t,x,\epsilon}[\|\epsilon - \hat{\epsilon}\|^2], \quad (4)$$

where $\hat{\epsilon} = \phi_\theta(\mathbf{z}_t, t)$ and $\mathbf{z}_t = \mathbf{V}(\alpha_t \mathbf{V}^T \mathbf{x} + \sigma_t \mathbf{V}^T \epsilon)$. The $\hat{\epsilon} = \phi_\theta(\mathbf{z}_t, t)$ signifies that the neural network ϕ_θ takes \mathbf{z}_t as input and predicts $\hat{\epsilon}$. This formulation satisfies the requirement for neural networks to perform effectively in a standard pixel space. After prediction, it is sufficient to transition back and forth between frequency space and pixel space using the DCT matrix \mathbf{V}^T and its inverse \mathbf{V} .

3.2. Schedules of Blurring Diffusion Model

The schedules for the blurring diffusion model, denoted as α_t and σ_t , are obtained by combining a standard Gaussian noise diffusion schedule (with scalar values σ_t and α_t) along with a blurring schedule \mathbf{d}_t . Each element of the vector σ_t shares the same value, as identical noise is added to all frequency components. Therefore, it becomes adequate to present a schedule for a scalar value σ_t [20]. The noise schedule is chosen based on a variance-preserving cosine [31], specifically $\sigma_t^2 = 1 - \alpha_t^2$, where $\alpha_t = \cos(t\pi/2T)$ for $t \in [0, T]$. Following [35], the blurring schedule is then defined as:

$$\sigma_{B,t} = \sigma_{B,max} \sin(t\pi/2T)^2, \quad (5)$$

where $\sigma_{B,max}$ represents a hyperparameter equal to the maximum level of blur applied to the image. This schedule, in turn, corresponds to the dissipation time through

$\tau_t = \sigma_{B,t}^2/2$. Based on the formulation discussed in the Appendix, the blurring schedule \mathbf{d}_t , which is employed for α_t , is defined as follows:

$$\mathbf{d}_t = \exp(-\boldsymbol{\lambda}\tau_t), \quad (6)$$

where $\boldsymbol{\lambda}$ represents the vector containing squared frequencies, and τ_t corresponds to the dissipation time. To achieve a more gradual amplification of high frequencies during the denoising process, the blurring schedule \mathbf{d}_t is adjusted to $(1 - d_{min}) \exp(-\boldsymbol{\lambda}\tau_t) + d_{min}$, where d_{min} is set to 0.001. Finally, by combining the Gaussian noise schedule with the blurring schedule, resultant schedule is as follows:

$$\alpha_t = \alpha_t \cdot \mathbf{d}_t, \sigma_t = 1\sigma_t. \quad (7)$$

3.3. Blurring Diffusion Model for Compression

The rate-distortion objective in end-to-end compression resembles the loss function of a β -VAE, where a hyperparameter λ is utilized to balance the trade-off between the bit-rate (R) and distortion (D):

$$\mathbf{L} = D + \lambda R = \mathbf{E}_{\tilde{\mathbf{y}}}[-\log p_{\mathbf{x}|\tilde{\mathbf{y}}}(\mathbf{x}|\tilde{\mathbf{y}}) - \lambda \log p_{\tilde{\mathbf{y}}}(\tilde{\mathbf{y}})]. \quad (8)$$

While many neural codecs commonly employ Gaussian or Laplacian decoders, we introduce a novel approach using a conditional blurring diffusion model as the decoder. This approach aims to yield new distortion metrics that deviate from those based on Mean Squared Error (MSE) or Mean Absolute Error (MAE). In this framework, our proposed neural codec leverages two distinct types of latent variables: a **semantic** latent variable \mathbf{y} and **texture** latent variables $\mathbf{z}_{1:T}$. The semantic latent variable captures and encodes the overall content and meaning of the image. In contrast, the texture latent variables are tailored to carry finer details and intricate patterns that might not be fully represented by the semantic variable alone. Notably, unlike the semantic latent variable, the texture latent variables are not compressed but synthesized during the decoding phase.

The forward and backward processes of the conditional blurring diffusion model can be expressed as:

$$\begin{aligned} q(\mathbf{f}_t|\mathbf{f}_{t-1}) &= \mathcal{N}(\mathbf{f}_t; \alpha_{t|t-1} \mathbf{f}_{t-1}, \boldsymbol{\sigma}_{t|t-1}^2 \mathbf{I}), \\ p_\theta(\mathbf{f}_x, \mathbf{f}_{1:T}|\tilde{\mathbf{y}}) &= p(\mathbf{f}_T) p_\theta(\mathbf{f}_x|\mathbf{f}_1, \tilde{\mathbf{y}}) \prod_{t=2}^T p_\theta(\mathbf{f}_{t-1}|\mathbf{f}_t, \tilde{\mathbf{y}}), \\ p(\mathbf{f}_T) &= \mathcal{N}(\mathbf{f}_T; \mathbf{0}, \mathbf{I}) \\ p_\theta(\mathbf{f}_{t-1}|\mathbf{f}_t, \tilde{\mathbf{y}}) &= \mathcal{N}(\mathbf{f}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{f}_t, \tilde{\mathbf{y}}, t), \boldsymbol{\sigma}_{t \rightarrow t-1}^2 \mathbf{I}) \end{aligned} \quad (9)$$

where $\mathbf{f}_t = \mathbf{V}^T \mathbf{z}_t$ and $\mathbf{f}_x = \mathbf{V}^T \mathbf{x} = \mathbf{V}^T \mathbf{z}_o$. As shown in Eq. 8, distortion is equivalent to the negative marginal likelihood of input data $-\log p_{\mathbf{x}|\tilde{\mathbf{y}}}(\mathbf{x}|\tilde{\mathbf{y}})$, and minimizing it

is analogous to minimizing the negative marginal likelihood of the frequency representation of the data $p_{f_x|\tilde{y}}(f_x|\tilde{y}) = \int p(f_x, f_{1:T}|\tilde{y})df_{1:T}$. Since computing the marginal likelihood is intractable, we employ its ELBO with the specified diffusion and de-blurring distributions. Following this substitution, the rate-distortion objective, through the application of Jensen’s inequality, can be formulated as follows:

$$\begin{aligned} & \mathbf{E}_{\tilde{y}}[-\log p_{x|\tilde{y}}(x|\tilde{y}) - \lambda \log p_{\tilde{y}}(\tilde{y})] \\ & \leq \mathbf{E}_{\tilde{y}}[-\text{ELBO} - \lambda \log p_{\tilde{y}}(\tilde{y})], \end{aligned} \quad (10)$$

where $\text{ELBO} = \mathbf{E}_q[\frac{p_{\theta}(f_x, f_{1:T}|\tilde{y})}{q(f_{1:T}|f_{t-1})}]$. As explained in Appendix, the optimizing of ELBO can be simplified to:

$$\text{ELBO} \approx \mathbf{E}_{t,x,\epsilon,\tilde{y}}[\|\epsilon - \phi_{\theta}(z_t, t, \tilde{y})\|^2], \quad (11)$$

where $z_t = \mathbf{V}\alpha_t\mathbf{V}^T x + \sigma_t\epsilon$. Similar to [28], an LPIPS loss [46] is incorporated into the rate-distortion loss to enhance the perceptual quality of the reconstructed image. As the initial image at any time step can be decoded as a function of the texture latent variable z_t , the semantic latent variable y , and the time step t , i.e., $\hat{x}_t = \mathbf{V}(1/\alpha_t)(\mathbf{V}^T z_t - \sigma_t\mathbf{V}^T \hat{\epsilon})$, the total loss becomes as:

$$\begin{aligned} & \mathbf{E}_{t,x,\epsilon,\tilde{y}}[\|(1 - \beta)\|\epsilon - \phi_{\theta}(z_t, t, \tilde{y})\|^2 \\ & - \lambda \log p_{\tilde{y}}(\tilde{y}) + \beta d_{\text{LPIPS}}(x, \hat{x}_t)], \end{aligned} \quad (12)$$

where the λ and β represent hyperparameters that control the trade-off between rate, distortion, and perception.

Decoding Process: After the training, we utilize entropy decoding on \hat{y} through the entropy model which estimates the distribution $p_{\tilde{y}}(\hat{y})$. We then employ ancestral sampling to conditionally decode the image x , which is equal to z_0 , and starts from pure Gaussian noise $z_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$:

$$z_{t-1} \leftarrow \mathbf{V}(\hat{\mu}_{t \rightarrow t-1} + \sigma_{t \rightarrow t-1}\mathbf{V}^T\phi_{\theta}(z_t, t, \hat{y})) \quad (13)$$

where $\hat{\mu}_{t \rightarrow t-1} = \frac{\alpha_{t|t-1}\sigma_{t-1}^2}{\sigma_t^2}\mathbf{V}^T z_t + \frac{\sigma_{t|t-1}^2}{\alpha_{t|t-1}\sigma_t^2}(\mathbf{V}^T z_t - \sigma_t\mathbf{V}^T\phi_{\theta}(z_t, t, \hat{y}))$. The latent variables $z_{1:T}$ are not stored but produced during the decoding.

3.4. Proposed Entropy Model

The main goal of the proposed entropy model’s context is to exploit both channel-wise and spatial correlations, while expediting the decoding process. Inspired by the ELIC [14], we adopt an uneven grouping of latent channels, allowing most low entropy channels to depend on high entropy channels. The latent representation \hat{y} , with M channels, is clustered into five chunks along the channel dimension: 16, 16, 32, 64, and $M - 128$ channels, respectively. In this setup, each chunk depends on all its previous decoded chunks.

The Fig. 2(a) illustrates our proposed entropy model as applied in the j -th chunk. Within each chunk, we use spatial

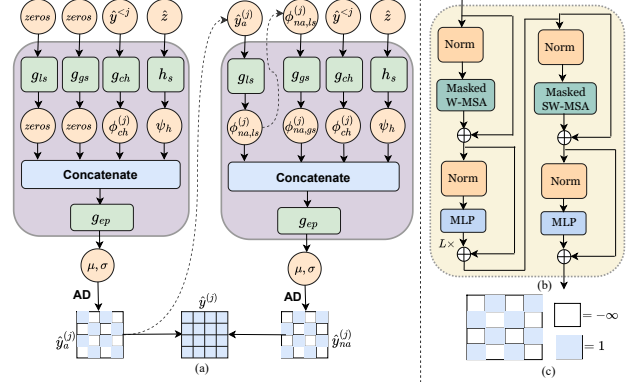


Figure 2. Diagram illustrating the application of the proposed entropy model for decoding the j -th chunk $\hat{y}^{(j)}$. (b) Global Spatial Context Block. (c) An example of a checkerboard-shaped mask.

context in conjunction with channel context to model correlations along both the channel and spatial dimensions. To expedite the decoding step, we employ a parallel bidirectional spatial context model which is capable of capturing both local and global spatial relationships. So, the anchor part is decoded in parallel by using solely the hyperprior and channel context, while the decoding of the non-anchor part relies on the hyperprior, as well as both the spatial and channel contexts.

3.4.1 Spatial Context

The spatial context design captures both local and global spatial correlations within the latent representation \hat{y} . Following the decoding of the anchor group, a checkerboard-shaped convolution is applied to this group, generating local context for all non-anchor elements in a parallel manner. Furthermore, the acquired local contexts of the non-anchor part are subsequently fed into a Transformer-based block to efficiently extract the global spatial context. This Transformer block takes advantage of positional encoding that is customized specifically for the compression task, along with a checkerboard-based attention mechanism.

3.4.2 Transformer-based Spatial Context

We adopt the Swin Transformer [26] blocks, including a masked window-based multi-head self-attention (W-MSA) and a masked shifted-window-based multi-head self-attention (SW-MSA), as our global spatial context model which shown in Fig. 2.(b). This selection enables us to strike a balance between computational efficiency and modeling capacity. The checkerboard-based self-attention in W-MSA and SW-MSA can be expressed:

$$\text{Atten}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}(\mathbf{Q}\mathbf{K}^T \odot \mathbf{M})\mathbf{V}, \quad (14)$$




Figure 3. The procedure for acquiring Laplacian relative position encoding for a window with a size of 2×2 .

where $Q, K, V \in R^{N^2 \times d}$ are the query, key, and value matrices, respectively. N^2 denotes the number of patches in a window and d is the dimension of query/key/value. $M \in R^{N^2 \times N^2}$ represents a checkerboard-shaped mask.

3.4.3 Relative Positional Encoding

The authors in [13, 33] have shown that the bitrate is impacted by the distance of neighboring latents within the spatial context modeling of current latent. Based on this insight, we introduce a receptive field-aware self-attention mechanism that employs a learnable Laplacian-shaped positional encoding for calculating spatial context. In our proposed approach, each channel chunk possesses its own receptive field during self-attention computation, which dynamically adjusts in response to changes in entropy of chunks. In the j -th chunk, the checkerboard-based attention of W-MSA and SW-MSA is refined to receptive field-aware self-attention as follows:

$$\text{Atten}(Q_j, K_j, V_j) = \text{softmax}(Q_j k_j^T \odot M + P_{Lap,j}) V_j, \quad (15)$$

where $P_{Lap,j}$ represents the learnable Laplacian relative position encoding for chunk j . To create the Laplacian relative position encoding for the j -th chunk, which comprises spatial windows with a size of $N \times N$ (i.e. including N^2 patches), three steps need to be taken. Firstly, We generate a 2D relative position table, where each coordinate of the relative position lies in the range $[-N + 1, N - 1]$. Subsequently, we derive the relative position matrix for each patch. As shown in Fig. 3, the first patch's relative distance coordinate of $(0, 0)$ (relative distance with its position) is located at the top-left of the orange box, while the last patch's relative distance of $(0, 0)$ is positioned at the bottom-right of the green box. Afterward, each relative position matrix is flattened and stacked together. Finally, we apply the 2D Laplacian function to each element of the resulting matrix to generate a Laplacian relative position encoding $P_{Lap,j}$ with learnable parameters A_j and σ_j , whose size is $R^{N^2 \times N^2}$. The 2D Laplacian function is defined as:

$$f_{Lap}(x, y) = A_j^2 \exp((-1/2\sigma_j^2)(|x| + |y|)), \quad (16)$$

where $x \in \{-N + 1, \dots, N - 1\}$, $y \in \{-N + 1, \dots, N - 1\}$. A_j and σ_j are learnable parameters which are determined

for each chunk through optimization. The value of them is associated with the wideness of the effective receptive field.

4. Experiments

4.1. Implementation Details

To train our learned image compression network, we utilize a merged dataset that includes high-resolution images from the DIV2K, Flickr2K, and CLIC [1] datasets. These images are randomly cropped to a size of 256×256 during the training phase. The model parameters of all architectures were optimized using the Adam optimizer for a total of 2.4 million steps, with a batch size of 8. The initial learning rate was configured to be 1×10^{-4} and was progressively reduced until the conclusion of training, reaching 1×10^{-7} . In order to cover a broad spectrum of bitrates, we selected a hyperparameter λ from the set $\{0.0004, 0.005, 0.01, 0.02, 0.04, 0.016\}$. The hyperparameter β which specifies the contribution of perceptual loss is considered 0.9 for all models. The parameter T , represents the required time steps for diffusion-based decoder, is consistently set to 500. To evaluate the performance of our compression approach, we select the Kodak dataset [2], which comprises 24 high-quality images with a resolution of 768×512 , and the CLIC2020 test set (428 images) with varying resolutions.

4.2. Comparison with the SOTA Methods

We compare our model with state-of-the-art generative based image compression networks, including both GAN-based and diffusion-based codecs, as well as a state-of-the-art hand-crafted codec such as VVC-Intra (VTM) [3]. Additionally, we choose two VAE codecs for comparison: DGML [9] and NSC [40]. Specifically, within the domain of GAN-based image compression approaches, we select HiFiC [28] and Multi-Realism [4] frameworks for comparison. Furthermore, we assess our method against diffusion-based models, namely DIRAC [11] and CDC [42]. This comparison is performed with respect to both rate-distortion and rate-perception tradeoff. The distortion is measured by the Peak Signal-to-Noise Ratio (PSNR) metric and perceptual quality are quantified via Fréchet Inception Distance (FID) [16] and Learned Perceptual Image Patch Similarity (LPIPS). As shown in Fig. 4, our model shows superior performance compared to all other codecs in terms of rate-perception. However, this heightened realism comes at the cost of distortion where other methods, except for the CDC model, either exhibit better performance or remain competitive.

4.3. Visual Quality

Fig. 5 illustrates reconstructed images (kodim20.png, kodim07.png) generated by our proposed model, as well as

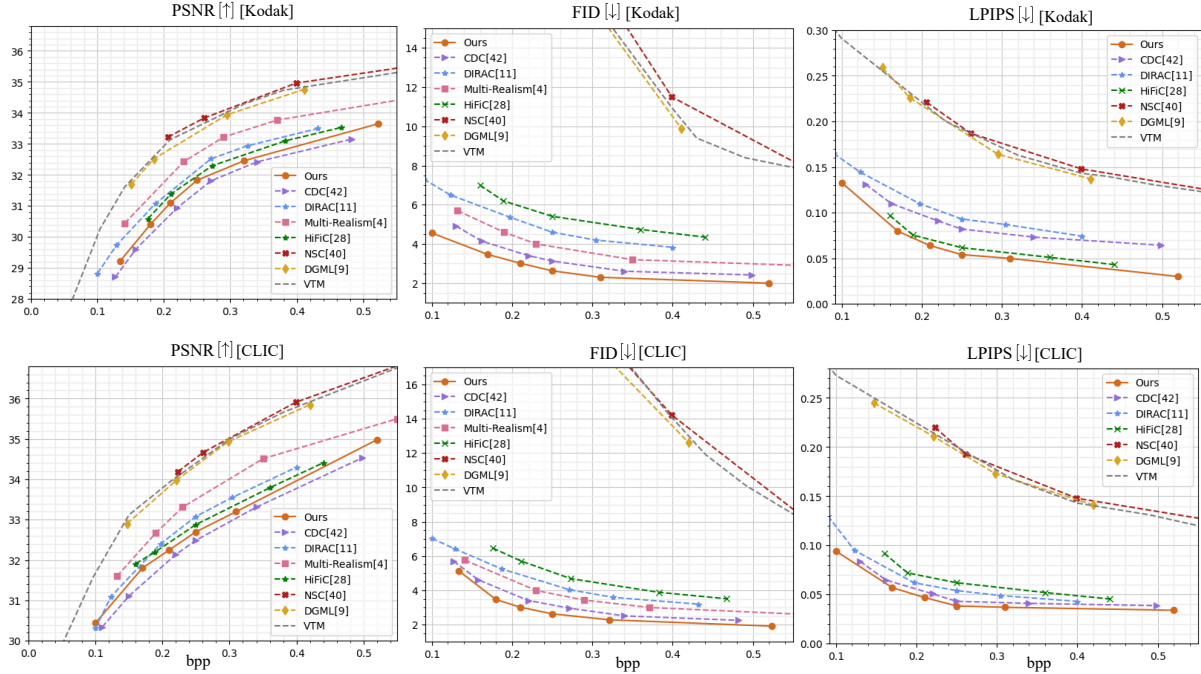


Figure 4. Comparison of our method with other codecs in terms of rate/distortion [bpp ↓ / PSNR ↑] and rate-perception, including [bpp ↓ / FID ↓] and [bpp ↓ / LPIPS ↓], for both the CLIC2020 test set and the Kodak dataset.

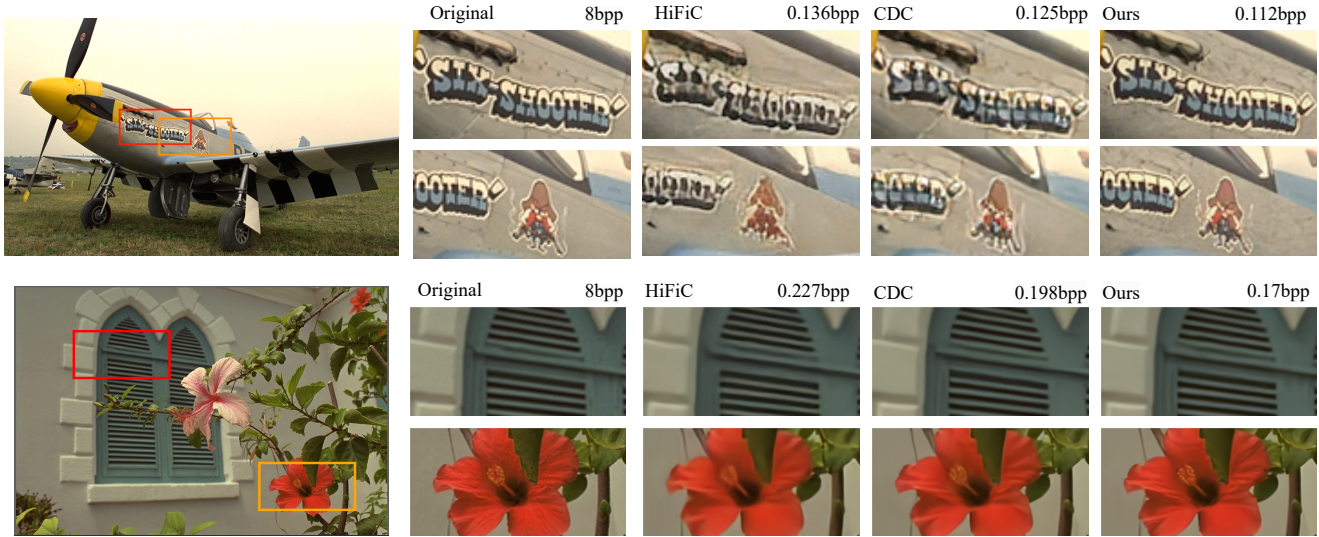


Figure 5. Visual comparison of our method to HiFiC and CDC models shows that our model achieves superior reconstruction quality, particularly at lower bit-rates. In addition, our model displays fewer artifacts compared to both the HiFiC and CDC models.

the CDC [42] and HiFiC [28] frameworks. Perceptually-oriented neural codecs, such as HiFiC and CDC, are criticized when their decoding process encounters semantic content, such as text within the compressed image. Our model effectively addresses this challenge, as depicted in Fig. 5. It enhances perceptual quality through a diffusion-based decoder and imposes a strict inductive bias via the

de-blurring diffusion process, resulting in both high perceptual quality and better text-preserving reconstructions. Additional qualitative comparisons are provided in the Appendix.

4.4. Ablation Study

Maximum Blurring: To explore the impact of the blurring schedule, we vary the values of $\sigma_{B,max}$ and compare the performance of the resulting models. It is evident that the model with $\sigma_{B,max}=0$ is equivalent to a standard denoising diffusion model. As indicated in Table 1a, blurring diffusion models with higher maximum blur levels $\sigma_{B,max} = 25$ generate higher-quality reconstructed images compared to other variant models in terms of FID score.

Table 1. Ablation Study: all models are optimized with $\lambda = 0.01$. (a) Investigating the impact of maximum blurring $\sigma_{B,max}$. (b) Exploring the effects of different types of positional encoding

$\sigma_{B,max}$	FID	Position Enc.	bpp
0	3.94	-	0.2811
5	3.78	Relative Pos.	0.2643
15	3.56	Diamond Relative Pos.	0.2512
25	3.45	Laplacian-shaped Pos.	0.2347

(a)

(b)

Positional Encoding: We investigate the influence of various types of positional encoding in the Transformer-based entropy model. As shown in Table 1b, the 2D diamond relative positional encoding, which is implemented using a clip function, demonstrates better performance than the relative position encoding. However, adopting the Laplacian-shaped positional encoding results in even more significant bitrate savings compared to the 2D diamond relative positional encoding. As depicted in Fig. 6, distinct values of A and σ are obtained for each channel chunk. As we progress towards the final chunks, the receptive field becomes narrower. As observed in [14], the last chunk of channels, conditioned on all previous chunks, exhibits lower entropy. Their entropy can be estimated by considering only a small neighborhood, corresponding to a narrower receptive field in our proposed relative positioning. Conversely, to provide a reasonable estimate of entropy of the first channel chunks, a wider receptive field gathers more information over a large context to compared to mentioned last chunks.

Analysis of Context Blocks: We conducted a comparison of the inference latency for entropy parameters during entropy decoding, as well as the bitrate savings of our proposed entropy model compared to other backward adaptation-based entropy models. For a fair comparison, all the models are equipped with a same encoder comprising of ResNet blocks and convolution layers to transform the input image $\mathbf{x} \in R^{H \times C \times 3}$ to a latent representation $\mathbf{y} \in R^{H/16 \times W/16 \times 256}$. As reported in Tabel 2, our entropy model’s speed is notably improved by employing unevenly grouped channels and implementing a bidirectional context model, as compared to the sequential spatial context modeling. Moreover, our findings clearly indicate that incorporating spatial global context results in superior per-




Figure 6. Receptive fields are extracted for each channel chunk.

Table 2. The Bjøntegaard-delta-rate (BD-rate) [6] and inference latency of entropy parameter estimation during entropy decoding (Dec.) for different context-based models on the Kodak dataset using a GPU (RTX A6000). The BD-Rate is computed relative to VVC. ([P]:parallel, [S]:Serial)

Method	Context Model		Dec.(ms)	BD-Rate
	Channel	Spatial		
Minnen <i>et. al.</i> [30]	-	Local[S]	$> 10^3$	-3.24
Minnen <i>et. al.</i> [29]	Even	-	67	-3.96
He <i>et. al.</i> [13]	-	Local[P]	28.2	-2.91
Qian <i>et. al.</i> [33]	-	Global[S]	$> 10^3$	-4.89
He <i>et. al.</i> [14]	Uneven	-	37.2	-3.12
He <i>et. al.</i> [14]	Uneven	Local[P]	156.9	-5.71
Jiang <i>et. al.</i> [21]	Uneven	Local[P]+Global[P]	207.4	-7.89
Ours	Uneven	Local[S]+Global[S]	$> 10^3$	-9.38
Ours	Uneven	Local[P]+Global[P]	196.3	-8.25
VVC	-	-	-	0.00

formance when contrasted with utilizing only local context. The evaluation emphasizes that the integration of both global and local context enhances the precision of capturing spatial correlations, ultimately leading to a more accurate entropy model. In addition, our results verify that incorporating Laplacian-shaped positional encoding enhances the compression efficiency compared to MEM, which does not consider positional encoding in computing global spatial context.

5. Conclusion

We developed a neural image compression model which improves perceptual image quality using a non-isotropic diffusion decoder. This decoder’s inductive bias effectively separates frequency components, leading to the creation of high-quality images. Moreover, we introduce an innovative entropy model that optimizes the trade-off between compression efficiency and decoding speed. This entropy model, founded on the Transformer architecture with Laplacian-shaped positional encoding, establishes a strong global spatial context. Our results underscore the efficacy of leveraging diffusion models and advanced entropy modeling to achieve outstanding image compression performance.

Acknowledgement: This research is based upon work supported by the National Aeronautics and Space Administration (NASA), via award number 80NSSC21M0322 under the title of *Adaptive and Scalable Data Compression for Deep Space Data Transfer Applications using Deep Learning*.

References

- [1] CLIC · challenge on learned image compression,. Available at <http://compression.cc/tasks/>, 2022. **6**
- [2] Kodak image dataset,. Available at <https://r0k.us/graphics/kodak/>, 2022. **6**
- [3] Versatile Video Coding Reference Software. Available at https://vcgit.hhi.fraunhofer.de/jvet/VVCSoftware_VTM, 2022. **6**
- [4] Eirikur Agustsson, David Minnen, George Toderici, and Fabian Mentzer. Multi-realism image compression with a conditional generator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22324–22333, 2023. **1, 2, 6**
- [5] Johannes Ballé, David Minnen, Saurabh Singh, Sung Jin Hwang, and Nick Johnston. Variational image compression with a scale hyperprior. In *International Conference on Learning Representations*, 2018. **2, 3**
- [6] Gisle Bjontegaard. Calculation of average psnr differences between rd-curves. *ITU SG16 Doc. VCEG-M33*, 2001. **8**
- [7] 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**
- [8] Ya-Liang Chang, Zhe Yu Liu, Kuan-Ying Lee, and Winston Hsu. Free-form video inpainting with 3d gated convolution and temporal patchgan. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9066–9075, 2019. **2**
- [9] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *CVPR*, pages 7939–7948, 2020. **6**
- [10] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. **2**
- [11] Noor Fathima Ghouse, Jens Petersen, Auke Wiggers, Tianlin Xu, and Guillaume Sautière. A residual diffusion model for high perceptual quality codec augmentation. *arXiv preprint arXiv:2301.05489*, 2023. **1, 3, 6**
- [12] Vivek K Goyal. Theoretical foundations of transform coding. *IEEE Signal Processing Magazine*, 18(5):9–21, 2001. **2**
- [13] Dailan He, Yaoyan Zheng, Baocheng Sun, Yan Wang, and Hongwei Qin. Checkerboard context model for efficient learned image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14771–14780, 2021. **1, 2, 3, 6, 8**
- [14] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5718–5727, 2022. **2, 3, 5, 8**
- [15] Dailan He, Ziming Yang, Hongjiu Yu, Tongda Xu, Jixiang Luo, Yuan Chen, Chenjian Gao, Xinjie Shi, Hongwei Qin, and Yan Wang. Po-elic: Perception-oriented efficient learned image coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1764–1769, 2022. **2**
- [16] 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. *Advances in neural information processing systems*, 30, 2017. **6**
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. **1, 3, 4**
- [18] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. **3**
- [19] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *The Journal of Machine Learning Research*, 23(1):2249–2281, 2022. **3**
- [20] Emiel Hoogeboom and Tim Salimans. Blurring diffusion models. *arXiv preprint arXiv:2209.05557*, 2022. **2, 3, 4**
- [21] Wei Jiang, Jiayu Yang, Yongqi Zhai, Peirong Ning, Feng Gao, and Ronggang Wang. Mlic: Multi-reference entropy model for learned image compression. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 7618–7627, 2023. **2, 3, 8**
- [22] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021. **3**
- [23] A Burakhan Koyuncu, Han Gao, Atanas Boev, Georgii Gaikov, Elena Alshina, and Eckehard Steinbach. Contextformer: A transformer with spatio-channel attention for context modeling in learned image compression. In *European Conference on Computer Vision*, pages 447–463. Springer, 2022. **2, 3**
- [24] Jooyoung Lee, Seunghyun Cho, and Seung-Kwon Beack. Context-adaptive entropy model for end-to-end optimized image compression. In *International Conference on Learning Representations*, 2018. **2**
- [25] 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. **1**
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. **5**
- [27] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022. **3**
- [28] Fabian Mentzer, George D Toderici, Michael Tschannen, and Eirikur Agustsson. High-fidelity generative image compression. *Advances in Neural Information Processing Systems*, 33:11913–11924, 2020. **1, 2, 5, 6, 7**

- [29] David Minnen and Saurabh Singh. Channel-wise autoregressive entropy models for learned image compression. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 3339–3343. IEEE, 2020. 2, 3, 8
- [30] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. *Advances in neural information processing systems*, 31, 2018. 2, 3, 8
- [31] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. 4
- [32] Yichen Qian, Zhiyu Tan, Xiuyu Sun, Ming Lin, Dongyang Li, Zhenhong Sun, Li Hao, and Rong Jin. Learning accurate entropy model with global reference for image compression. In *International Conference on Learning Representations*, 2020. 2, 3
- [33] Yichen Qian, Xiuyu Sun, Ming Lin, Zhiyu Tan, and Rong Jin. Entroformer: A transformer-based entropy model for learned image compression. In *International Conference on Learning Representations*, 2021. 2, 3, 6, 8
- [34] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 3
- [35] Severi Rissanen, Markus Heinonen, and Arno Solin. Generative modelling with inverse heat dissipation. *arXiv preprint arXiv:2206.13397*, 2022. 4
- [36] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4713–4726, 2022. 3
- [37] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015. 3
- [38] Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszár. Lossy image compression with compressive autoencoders. In *International Conference on Learning Representations*, 2016. 1, 2
- [39] Lucas Theis, Tim Salimans, Matthew D Hoffman, and Fabian Mentzer. Lossy compression with gaussian diffusion. *arXiv preprint arXiv:2206.08889*, 2022. 3
- [40] Dezhao Wang, Wenhan Yang, Yueyu Hu, and Jiaying Liu. Neural data-dependent transform for learned image compression. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 17379–17388, 2022. 6
- [41] Jay Whang, Mauricio Delbracio, Hossein Talebi, Chitwan Saharia, Alexandros G Dimakis, and Peyman Milanfar. Deblurring via stochastic refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16293–16303, 2022. 3
- [42] Ruihan Yang and Stephan Mandt. Lossy image compression with conditional diffusion models. *arXiv preprint arXiv:2209.06950*, 2022. 1, 3, 6, 7
- [43] Ren Yang, Luc Van Gool, and Radu Timofte. Perceptual learned video compression with recurrent conditional gan. *arXiv preprint arXiv:2109.03082*, 1, 2021. 2
- [44] Yibo Yang, Stephan Mandt, Lucas Theis, et al. An introduction to neural data compression. *Foundations and Trends® in Computer Graphics and Vision*, 15(2):113–200, 2023. 1
- [45] Ali Zafari, Atefeh Khoshkhahtinat, Piyush Mehta, Mohammad Saeed Ebrahimi Saadabadi, Mohammad Akyash, and Nasser M Nasrabadi. Frequency disentangled features in neural image compression. In *2023 IEEE International Conference on Image Processing (ICIP)*, pages 2815–2819. IEEE, 2023. 1
- [46] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 2, 5
- [47] Yin hao Zhu, Yang Yang, and Taco Cohen. Transformer-based transform coding. In *ICLR*, 2021. 3