

Multi-Realism Image Compression with a Conditional Generator

Eirikur Agustsson
 Google Research
 Reykjavík, Iceland
 eirikur@google.com

David Minnen
 Google Research
 Mountain View, USA
 dminnen@google.com

George Toderici
 Google Research
 Mountain View, USA
 gtoderici@google.com

Fabian Mentzer
 Google Research
 Zürich, Switzerland
 mentzer@google.com

Abstract

By optimizing the rate-distortion-realism trade-off, generative compression approaches produce detailed, realistic images, even at low bit rates, instead of the blurry reconstructions produced by rate-distortion optimized models. However, previous methods do not explicitly control how much detail is synthesized, which results in a common criticism of these methods: users might be worried that a misleading reconstruction far from the input image is generated. In this work, we alleviate these concerns by training a decoder that can bridge the two regimes and navigate the distortion-realism trade-off. From a single compressed representation, the receiver can decide to either reconstruct a low mean squared error reconstruction that is close to the input, a realistic reconstruction with high perceptual quality, or anything in between. With our method, we set a new state-of-the-art in distortion-realism, pushing the frontier of achievable distortion-realism pairs, *i.e.*, our method achieves better distortions at high realism and better realism at low distortion than ever before.

1. Introduction

Lossy image compression considers the trade-off between the number of bits used to store an input image and how close the reconstruction (that we obtain from the bits) is to that input image. As we use more bits, we will be able to get closer to the input. This idea is formalized in the fundamental rate-distortion trade-off [34], where “rate” stands for bit-rate, and “distortion” is formalized as a pairwise metric between the input image and the reconstruction (*e.g.*, the mean-squared error, MSE).

While minimizing this trade-off has been the focus of many works starting from JPEG [43] all the way to recent neural [13] and non-neural [11] codecs, there has been a surge of interest in additionally considering the “realism” or “perceptual quality” of the reconstructions [6, 14, 23, 25, 35, 37, 39, 41, 45, 48]. After all, as we move toward low rates, purely rate-distortion optimized systems will produce arti-




Figure 1. Decoding two reconstructions from the same representation \hat{y} , which takes 2345 bytes to store: We use a low realism weight $\beta = 0$ for the left reconstruction, and a high $\beta = 2.56$ for the right. Note that increasing β leads to a much sharper reconstruction, but the PSNR drops by 1.1dB, consistent with rate-distortion-realism theory [6]. We only show two reconstructions, but our generator G can produce any reconstruction in between by changing β . This allows the user to decide between viewing a reconstruction that is close to the input (left, *i.e.*, high PSNR), or that looks realistic (right).

facts in the reconstructions, such as the well known block artifacts of JPEG or blurry patches for neural approaches. There is simply not enough bitrate available to store all of the details, and if we target, *e.g.*, MSE, the best reconstruction is the average image over all images that map to the given representation since, inevitably, many images will map to the same representation at low rates. Intuitively, instead of an average image reconstruction, we could prefer a “realistic” reconstruction that is sharp and appropriately textured. This reconstruction might have worse MSE than the average image, but users might find it more perceptually pleasing and less artificial. We can see from this argument that there exists an additional trade-off here, between “realism” and “distortion”, and that distortion will increase as

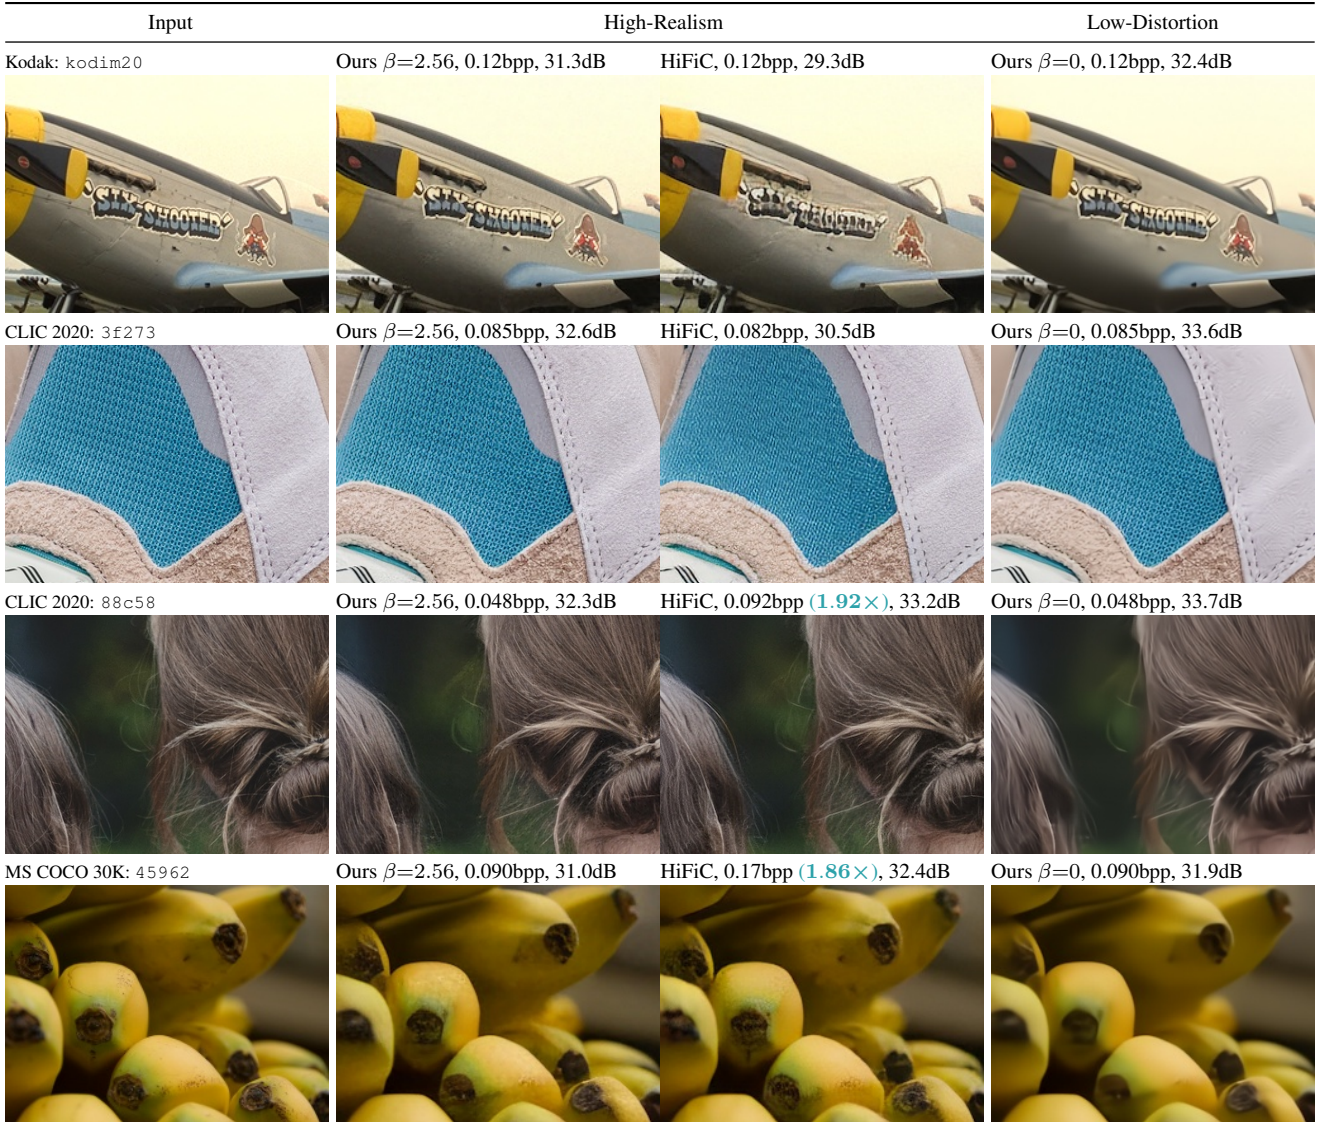


Figure 2. Comparing input images to reconstructions from our model at $\beta=2.56$, the generative state-of-the-art HiFiC, as well as our model at $\beta=0$. Note that both our models always have the same bits-per-pixel (bpp) per row, since for each row, the two reconstructions we show are obtained from the *same* representation—we simply vary β for the generator. Overall, we see how our high-realism reconstructions ($\beta=2.56$) closely match the input, more-so than HiFiC. On the airplane (first row), we can read the text in our reconstruction, in contrast to the one from HiFiC. In the second row, the texture of the sneaker is faithfully preserved. For the hair, we note that HiFiC uses **1.92** \times the bitrate of our model to achieve a similar reconstruction. In the last row, HiFiC uses **1.86** \times the rate. In the first two rows, where we have comparable bpp to HiFiC, both of our reconstructions have higher PSNR. In the rightmost column ($\beta = 0$) we can see the low-distortion reconstructions of our model. There we have near state-of-the-art PSNR at the cost of losing the (synthetic) detail.

we improve realism. Following Blau and Michaeli [6], we formalize “distortion” as a metric between pairs of images (e.g., MSE) that indicates how close is the reconstruction to the input, while “realism” indicates how realistic the reconstructions look (regardless of the input). We formalize the latter as a divergence $d(p_X, p_{\hat{X}})$ between the distribution of real images, X , and reconstructions, \hat{X} . Note that this can only be measured over a set of images since an accurate estimate of the distribution is needed. Throughout this text,

we use PSNR as a measure of distortion, and FID [16] as a measure of realism.

Previous work successfully optimized the triple rate-distortion-realism trade-off [1, 14, 25, 31, 33], however, there is one caveat. Since the realism constraint might produce reconstructions that are far away from the input, these systems might be looked at with suspicion because it is not clear which details are in the original and which were added by the architecture.

We address this caveat by training a decoder that, given a *single* compressed representation, either produces a reconstruction where little or no detail is generated (like rate-distortion optimized codecs), one where fine-grained detail is generated (like rate-distortion-realism optimized codecs), or anything in between (see Fig. 1). We emphasize that the *receiver* can decide how much detail to generate, because we condition the decoder, not the encoder, on a “realism factor”, β , and thus the receiver can produce the full spectrum of reconstructions from a single representation, \hat{y} .

Our main contributions are:

1. We bridge the generative and non-generative compression worlds, by navigating the trade-off between distortion and realism from a *single* representation using a conditional generator.
2. Our method sets a new state-of-the-art in terms of distortion-realism on high-resolution benchmark datasets, pushing the frontier of achievable distortion-realism pairs. Our method achieves better distortions at high realism (low FID) and better realism at low distortion (high PSNR) than ever before (Fig. 4).

2. Related Work

Zhang *et al.* [48] studied “universal” representations for rate-distortion-realism and provide a theoretical exploration of the bit-rate overhead incurred by using a *single* representations to obtain different points in the distortion-realism plane.¹ They formalize the rate overhead of a universal representation compared to different representations for each distortion-realism point, and show that this overhead is zero for scalar Gaussian sources. For general sources, they show that the rate-distortion-optimal representation can be used to meet any realism constraint by increasing the distortion by no more than a factor 2. They present empirical results on MNIST, showing that after training a single encoder-decoder pair, further decoders can be trained given the frozen encoder. We note that we jointly train an encoder and a conditional decoder that navigates the distortion-realism trade-off by adapting to a side input in the form of a single realism factor, β .

He *et al.* present the “ELIC” model [13], which is a state-of-the-art neural compression model for rate-distortion (MSE and MS-SSIM [44]) performance amongst practical methods. Some methods like Koyuncu *et al.* [20] outperform ELIC at some bit rates, but they use a serial autoregressive context model to improve entropy modeling. Such context models typically lead to 10x slower decode times

¹Note that the terms “realism” and “perception” can be used interchangeably. Both refer to a divergence between distributions over images. In contrast, “perceptual quality” is used more generally as a measure of subjective visual quality, which could be measured by “realism” or by an image quality metric like LPIPS or MS-SSIM.

due to underutilization of the parallel cores of GPUs and TPUs.

Various previous compression methods incorporated adversarial losses to boost realism. Mentzer *et al.* developed HiFiC, which combined a conditional GAN [29] with a hyperprior-based compression architecture [3] and showed rate savings of 50% for equal subjective quality compared to MSE-optimized and standard (non-neural) codecs. While ELIC was only optimized for rate-distortion, it was extended to create a “perceptually-oriented” variant called PO-ELIC [14]. This model focused on realism by augmenting the loss function with an adversarial term, a perceptual loss based on LPIPS [49], and a patch-based style loss [10]. Similarly, Li *et al.* [21] also combine multiple loss terms including a Laplacian loss, MSE, MAE, adversarial loss, and LPIPS, but they merge these terms in a spatially varying, content-adaptive manner based on different detectors (faces, edges, and structure) that run during training. The model is able to learn where to apply each type of loss based on image content, which boosts perceptual quality. Other methods utilize region-of-interest (ROI) or semantic maps to guide detail and texture synthesis [1, 22].

We emphasize that these methods target a single point on the distortion-realism tradeoff, and would require storing a different representation for each distortion-realism target. This is in contrast to our method, which only requires a single model and representation, yet can still generate reconstructions targeting any trade-off along the distortion-realism curve.

More related to our approach, Gao *et al.* [9] present an approach for targeting different multi-distortion trade-offs with a single model using semi-amortized inference: first, a model is trained for a single trade-off to predict a latent representation. This representation is then further optimized for a new trade-off at *encode time*. Although effective, this approach has several drawbacks: (1) the new trade-off parameters must be selected at encode time, not decode time and (2) encoding becomes very slow since hundreds or thousands of optimization steps must run for each image.

Iwai *et al.* [18] use network interpolation to achieve different distortion-realism targets, however their method operates in a different regime by targeting extremely low bitrates (<0.04bpp on Kodak), where visual quality clearly suffers. We also note that [18] is a very different approach: first they train a low-distortion encoder/decoder pair. Then the decoder is finetuned for a GAN loss to synthesize details. At inference time, the model weights of the decoders are interpolated post-hoc.

Theis *et al.* [38] show promising results for generative compression of small (64x64) images using gaussian diffusion and reverse channel coding, obtaining state-of-the-art results on ImageNet64. The approach is based on using reverse channel coding [36] to transmit samples, which

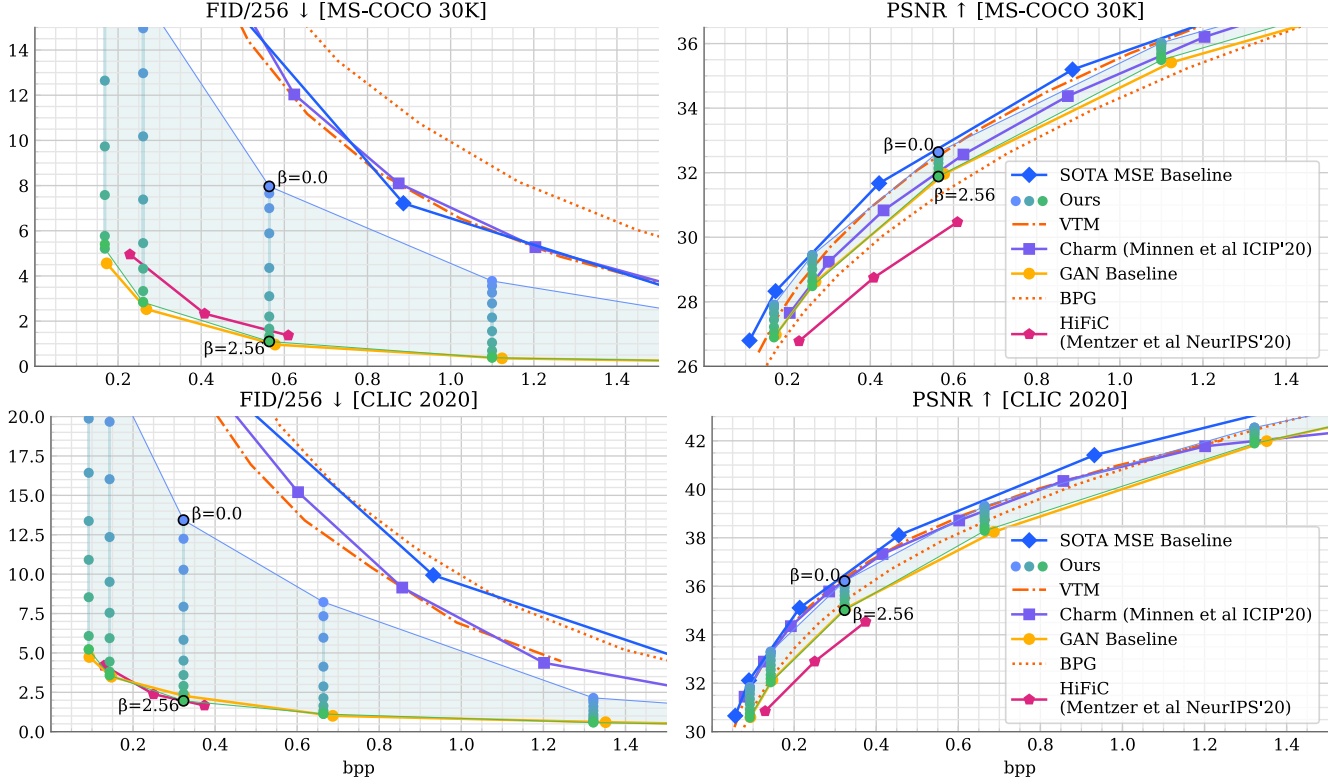


Figure 3. Results on MS-COCO (top) and CLIC 2020 (bottom). We show FID as a measure of *realism* (left, lower is better) and PSNR as a measure of *distortion* (right, higher is better). The chains of dots indicate the realism-distortion points our method can achieve simply by varying β on the receiver side, *i.e.*, we have single model per bitrate. We highlight two values of β to make this clear. We observe that for high β , we match or outperform HiFiC in FID (realism), while having a significantly better PSNR ($> 1\text{dB}$). On the low β side, our model outperforms Charm and matches VTm in terms of PSNR, while having a significantly better FID.

is under active research and currently computationally prohibitive for the large images we consider here.

3. Method

3.1. Background

Neural Image Compression We follow the commonly used [3, 8, 13, 15, 24, 27, 28, 50] non-linear transform coding approach [2] to do lossy image compression: We train an auto-encoder E, G , that maps an input image x to a *quantized* representation $\hat{y} = E(x)$ and back to a reconstruction $\hat{x} = G(\hat{y})$ (we use G for the decoder and call it “generator” to avoid confusion with the discriminator D we introduce below). Training E, G for reconstruction (*e.g.* with a MSE loss) already leads to a compression system, where the sender runs E , stores \hat{y} to disk, the receiver obtains \hat{y} from disk and runs G . However, naively storing \hat{y} to disk is expensive, as it takes $\log_2 |\mathcal{S}|$ bits per symbol (assuming $\hat{y}_i \in \mathcal{S}$). One can do better if the distribution of the symbols is known, as one can then assign shorter bitstrings to more likely symbols. Given a distribution $p(\hat{y})$ estimating the true (unknown) distribution $q(\hat{y})$ of the symbols in

\hat{y} , we can use entropy coding algorithms to store \hat{y} using $B(\hat{y}) = \mathbb{E}_{\hat{y} \sim q} -\log_2 p(\hat{y})$ bits, where B is the cross-entropy between the true data distribution q and our model p (see, *e.g.*, Yang *et al.* [46]). To minimize B , recent neural compression approaches train a separate “entropy model” P to predict $p(\hat{y})$. Using P , we can estimate a bitrate loss during training, $r(\hat{y}) = B(\hat{y})$ and thereby minimize the rate-distortion trade-off by minimizing

$$\mathcal{L}_{RD} = \mathbb{E}_{x \sim p_X} [r(\hat{y}) + \lambda \text{MSE}(x, \hat{x})], \quad (1)$$

where λ controls the trade-off. Typically, a set of models is trained by varying λ , which results in models covering different bitrates.

Generative Image Compression Inspired by the theoretical formalisation of “realism” as a divergence between the distribution of real images p_X and reconstructions $p_{\hat{X}}$ (see Introduction), previous works [1, 25] use a generative adversarial networks (GANs) [12]-based loss to estimate and minimize this divergence during training. We follow the formulation of Mentzer *et al.* [25], which is based on conditional GANs: In addition to E, G , a *conditional* discrim-

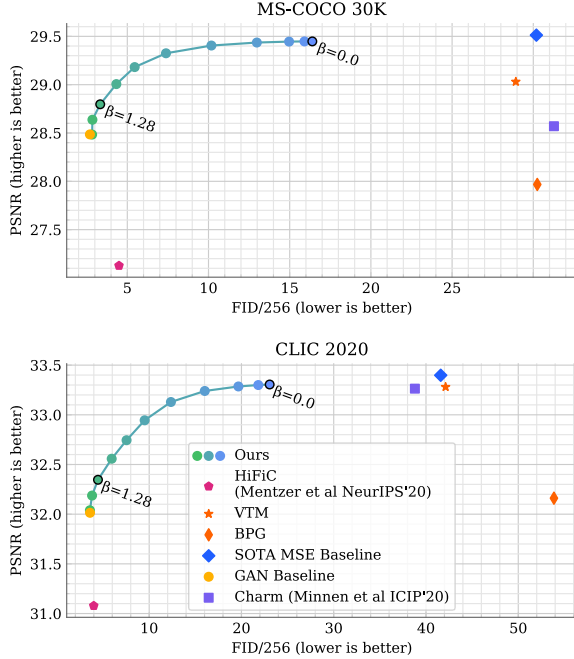


Figure 4. Distortion (PSNR) vs. realism (FID/256) trade-off on MS-COCO and CLIC 2020, for a *single bitrate*. This is obtained by slicing Fig. 3 at a single bitrate (0.26bpp for MS-COCO and 0.14bpp for CLIC 2020). In this figure, the optimal point is at the top left, and our method is closer to this optimum than the state-of-the-art generative method HiFiC, as well as the SOTA MSE Baseline. At high-realism (low FID), we outperform or match HiFiC in FID, while sporting a significantly higher PSNR. At low-distortion (high PSNR), we reach towards state-of-the-art, with significantly lower FID. We note that we get a chain of dots for our method because we can *decode multiple reconstructions from the same representation* using different realism weights β , whereas the shown baselines only have a single dot. We see similar results for other bitrates, see Sec A.5.

inator $D(\hat{y}, x)$ is trained to predict the probability that the given x is a realistic image corresponding to the representation \hat{y} . We use the patch discriminator from HiFiC [25]. Using D , we can formulate the GAN losses for G and D as follows:

$$\mathcal{L}_G = \mathbb{E}_{\hat{y} \sim p_Y} [-\log(D(\hat{y}, G(\hat{y})))] \quad (2)$$

$$\mathcal{L}_D = \mathbb{E}_{\hat{y} \sim p_Y} [-\log(1 - D(\hat{y}, G(\hat{y})))] + \quad (3)$$

$$\mathbb{E}_{x \sim p_X} [-\log D(E(x), x)],$$

where p_Y is the distribution of representations, induced by the encoder transform E .

3.2. Our Approach

Let $\beta \in [0, \beta_{\max}]$ be the “realism weight” that specifies whether our generator G should produce a low distortion or high realism reconstruction. Our goal is to train a single generator G to work well for any β . To this end, we base

our E, G on the “ELIC” architectures introduced by He *et al.* [13] to achieve state-of-the-art rate-distortion results, but we make G slightly wider, using $N=256$ channels instead of $N=192$.²

Additionally, and crucially, we condition G on β , obtaining a β -conditional generator $G(\hat{y}, \beta)$, as shown in Fig. 5 and described in Sec. 3.3. Thus, to obtain a reconstruction given a representation \hat{y} , the receiver chooses a $\beta \in [0, \beta_{\max}]$ and runs G to obtain $\hat{x}_\beta = G(\hat{y}, \beta)$.

We adopt the channel autoregressive entropy model (Charm) proposed by Minnen *et al.* [28] to minimize bitrate, using 10 slices. Combining the losses from the previous section, we obtain the overall loss for training E, G and D :

$$\mathcal{L}_{EGD}(\beta) = \mathbb{E}_{x \sim p_X} \left[\lambda' r(\hat{y}) + d(x, \hat{x}_\beta) + \underbrace{\beta (-\log(D(\hat{y}, \hat{x}_\beta) + C_P \mathcal{L}_P(x, \hat{x}_\beta))}_{\mathcal{L}_G} \right]. \quad (4)$$

We use $d(x, \hat{x}_\beta) = 1/100 \text{MSE}(x, \hat{x}_\beta)$ (where MSE is calculated on inputs and reconstructions scaled to $\{0, \dots, 255\}$), and $\lambda' = 100\lambda$. Following [25], we use $\mathcal{L}_P = \text{LPIPS}$ [49]. C_P is a hyper-parameter to weight \mathcal{L}_P relative to \mathcal{L}_G . We note that in contrast to Eq. 1, we have the rate parameter λ' on $r(\hat{y})$ instead of MSE. This formulation allows us to target different bitrates without changing the relative weight of the distortion compared to other terms. The factor $1/100$ in d and λ' is chosen such that this formulation is the same as Eq. 1 for $\lambda = 1/100$.

During training, we uniformly sample β and minimize $\mathbb{E}_{\beta \sim \mathcal{U}(0, \beta_{\max})} \mathcal{L}_{EGD}(\beta)$, using $\beta_{\max} = 5.12$ for training. During inference the receiver can choose β freely to navigate the distortion-realism trade-off, obtaining different \hat{x}_β from a fixed \hat{y} . As motivated in Sec. 5.2, we use $\beta_{\max} = 2.56$ for inference.

3.3. Beta Conditioning

To condition G , we use the β -conditioning scheme shown in Fig. 5, which we call **FourierCond**. It is inspired by how diffusion models condition on the timestep [17]: We first obtain global (*i.e.* shared for all layers) features $f(\beta)$ by calculating Fourier features [26, 42]. Here, we use the NeRF [26] approach, using $L=10$ in [26, Eq. 2]. Afterwards, we apply a 2-layer MLP (with ReLU activations and 512 channels for each of the two dense layers). We then learn a projection for each convolutional layer in each residual block in our G .

To explore whether it matters how exactly β is fed into G , we explore a second approach in the ablations, which we named **TableCond**. It is inspired by multi-rate image compression models [2], where we use a lookup table indexed

²On a one megapixel image using a NVidia V100 GPU, our G runs in 67ms, compared to 43ms for $N=192$, and 99ms for HiFiC [25].

by β to obtain scaling factors and biases which are applied to each of the (non-residual) convolutions in G .

We note that since we restrict our inference to a finite set of β s, both conditioning schemes lead to the same runtime during inference. After all, for any β , we can pre-compute the offset that gets applied to the residual blocks.

4. Experiments

4.1. Metrics

We evaluate our models for distortion and realism through PSNR and FID [16] respectively. PSNR in RGB is still the most widely used metric to assess distortion-optimized neural codecs, whereas FID is widely used to assess generative models in terms of realism [30, 32, 47, ...].

4.2. Datasets

We train our method on 256px crops extracted from a large set of high-resolution images, where each image is randomly resized such that the shorter side is between 500 and 1000 pixels. We evaluate on the following common benchmark datasets from *image compression*: **Kodak** [19], 24 images of resolution 512×768 or 768×512 and **CLIC 2020** [40], from which we use the `test` split with 428 high-resolution images. The shorter side is ≈ 1500 px for most images (see [25, Fig. A12] for more statistics). For Kodak, we only report PSNR since it has too few images to reliably estimate FID. For CLIC 2020, we follow HiFiC [25, Sec. A.7] and report *patched* FID, where we ex-

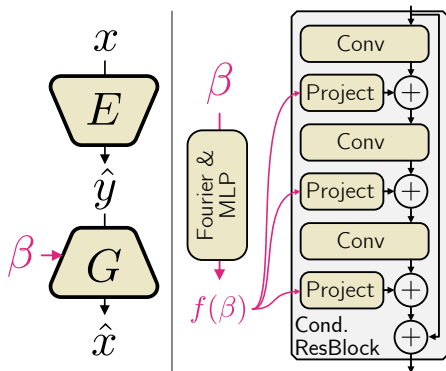


Figure 5. Overview of our architecture. Our encoder E and decoder/generator G are based on the ELIC [13] architecture, but we replace every residual block (RB) in G with our conditional RB shown on the right, to let G know which generative weight β to target. We first embed β using Fourier features into a d -dimensional space, then apply a 2-layer MLP to obtain features representing β , $f(\beta) = \text{MLP}(\text{Fourier}(\beta))$. Assuming the i -th conv. layer in the RB has C_i channels, we then project $f(\beta)$ to C_i channels with a learned (per layer) weight W_i , obtaining $W_i f(\beta)$. This is added to the output of the conv. layer.

tract 256×256 patches that cover all images (which we denote “FID/256”). This produces 30K overlapping patches (and 15K unique patches), which is of the order of magnitude required to measure FID. However, neither of these datasets are commonly used for evaluation in the generative modeling literature (e.g., DALL-E [30], Parti [47] and Imagen [32]), where **MS-COCO 30K** has become the main benchmark dataset, which we thus also use. The dataset consists of 30 000 256×256 images obtained from the MS-COCO 2014 validation set.³ We note that FID/256 is equivalent to vanilla FID on COCO.

4.3. Building Strong Baselines

No code or set of reconstructions is publicly available for the state-of-the-art non-generative image compression methods, so we aim to match the approach of He *et al.*, ELIC [13] in PSNR, as it is state-of-the-art while reporting fast inference on GPU (50ms on a Nvidia Titan XP for a 512×768 image).⁴ We use their E, G , but like for our method (Sec. 3.2), we use $N=256$ for G , and also use the Charm [28] entropy model (*i.e.*, in contrast to the paper by He *et al.*, we use equally sized slices and no checkerboard). The resulting model almost matches ELIC in PSNR (there is a ≈ 0.1 dB difference on Kodak, see Sec. 5.2.2), and we thus use it as a stand-in for state-of-the-art in PSNR, calling it **SOTA MSE Baseline**.

We combine this model with the discriminator and GAN formulation from Sec. 3.2 to form our **GAN baseline**. We train it for Eq. 4 using a fixed $\beta=2.56$, *i.e.*, this can be viewed as the same as our main model but using a non-conditional G that can only target a single realism weight. We use this baseline to tune the weights C_P for LPIPS and the GAN weight β , reported in Sec. 5.2, and then use the resulting C_P for our main models. The GAN baseline method outperforms HiFiC on COCO in FID, and nearly matches it on CLIC 2020, while being significantly stronger in terms of PSNR. We note that this is despite the fact that we do not port the multi-stage training or rate controller from HiFiC, *i.e.*, we train our models end-to-end from scratch.

Theoretically, a stochastic decoder is necessary to achieve high perceptual quality at low rates [41], so we explored concatenating noise to the representation before decoding in the GAN baseline. However, we found that this did not affect reconstructions at the rates we are interested in (intuitively stochasticity is crucial as the bitrate approaches zero).

4.4. Published Baselines

Since the publication of HiFiC [25] in 2020, there has been limited research in improving generative image com-

³Like Parti [47], we use the Dalle processing [30, Sec. A.2, Listing 1].

⁴Koyuncu *et al.* [20] report marginally better PSNR at a significant increase in inference time.

pression. The few methods that have been published [14, 18, 22] have very limited evaluation (usually only on the validation sets of CLIC’21 (41 images) or ’22 (30 images) which do not have enough images to estimate FID), and do not publish code to run on custom datasets. In contrast, **HiFiC** has code and reconstructions available. On high-resolution datasets large enough to estimate FID (CLIC 2020 and MS-COCO 30K), HiFiC remains state-of-the-art in terms of FID prior to the presented work.

On the MSE side, we compare to Minnen *et al.*’s **Charm** [28], since code is available for it and it is still close to state-of-the-art. We compare to **ELIC** [13] in terms of PSNR on Kodak, as well as visually to the two reconstructions they publish in Sec. A.4. We additionally compare to the non-learned **BPG** [5] (based on the HEVC standard) and **VTM** 17.1 [11] (the reference implementation of VVC [7]). VVC/VTM is the state-of-the-art among non-neural image codecs. In Sec. A.3 we zoom into the lowest bitrates to compare to **Iwai** et al. [18].

We detail how we run the publicly available methods in Sec. A.2.

4.5. Our Models

We train all baselines and ablations for 2M iterations at batch size 8 on 256px crops. For the multi-realism models, we train for 3M steps since the decoder needs to simultaneously learn to achieve high and low realism (we note that 3M is still less than training two models that target a single β , and our model can target infinitely many β s.) We use the Adam optimizer, with learning rate $1e^{-4}$ and default settings. As common in the literature, we train with a higher lambda (10x) in the first 15% steps, and decay the learning rate by a factor 10x in the last 15% steps. We did not tune these training parameters. We evaluate our model for $\beta \in \{0.0, \dots, 2.56\}$ (motivated in Sec. 5.2).

5. Results

5.1. Main Results

As mentioned in Sec. 4.3, the state-of-the-art in image compression [13, 20] in terms of MSE does not provide code or reconstructions. We thus use our “SOTA MSE Baseline” as a stand-in for the state-of-the-art in terms of PSNR. We establish its strength in Fig. 6, where we show that it is $\approx 0.0 - 0.2$ dB below ELIC [13].

In Fig. 3, we show that our model can achieve a new state-of-the-art in terms of distortion-realism: On the high-realism side ($\beta=2.56$), we match or outperform the state-of-the-art generative method HiFiC in FID (left plots, note the annotation of β), while also significantly outperforming it in PSNR (right plots). On the low-distortion side ($\beta=0$), we are strong in PSNR, reaching towards the SOTA MSE baseline in terms of PSNR (right plots), while significantly




Figure 6. Results on Kodak in terms of PSNR (note that Kodak is too small (24 images) to estimate FID). Our multi-realism model is strong in terms of PSNR in the low- β mode while also allowing for high realism reconstructions (see Fig 2). Additionally, here we can see that our “SOTA MSE Baseline” is competitive with the published (state-of-the-art) ELIC [13] model, we observe only a tiny gap (0.0 – 0.2dB).

outperforming it in FID (left plots). We emphasize that this means that *a*) our model is significantly closer to the input than HiFiC in the high- β mode (*i.e.*, it has higher PSNR, see right plots), leading to more faithful reconstructions, and also *b*) we have greater realism than state-of-the-art MSE models in the low- β mode.

This is even more apparent as we consider Fig. 4, which shows a single rate point and compares PSNR vs. FID. In this figure, it is best to be in the upper left, like in the common rate-distortion plots. *We can see that our approach reaches closer to this optimum than any previous method.* We now see more clearly that for low FID (high β), our method has significantly higher PSNR than HiFiC (≈ 1 dB), while for high PSNR (low β), our method has significantly lower FID (about 40%) than any of the non-generative methods (VTM, BPG, Charm, and the SOTA MSE baseline).

Comparing our model at $\beta = 0$ to the MSE baseline and our model at $\beta = 2.56$ to the GAN baseline, we might expect symmetric gaps on both sides. However, as we can see in Fig. 4, our GAN baseline leads to a similar operating point as our model set to $\beta=2.56$, while the MSE baseline has slightly better PSNR. It appears that the multi-task nature of our loss leads to models that slightly favor realism, perhaps not surprisingly, given that we randomly sample β during training and a large portion of the optimization thus uses $\beta>0$. Indeed, our model at $\beta=0$ actually has significantly better FID than the MSE baseline.

Additionally, in Fig. 6, we compare PSNR on Kodak, a commonly used benchmark in compression. Here we can see that our “SOTA MSE Baseline” is competitive with state-of-the-art ELIC [13]. Furthermore, despite the fact

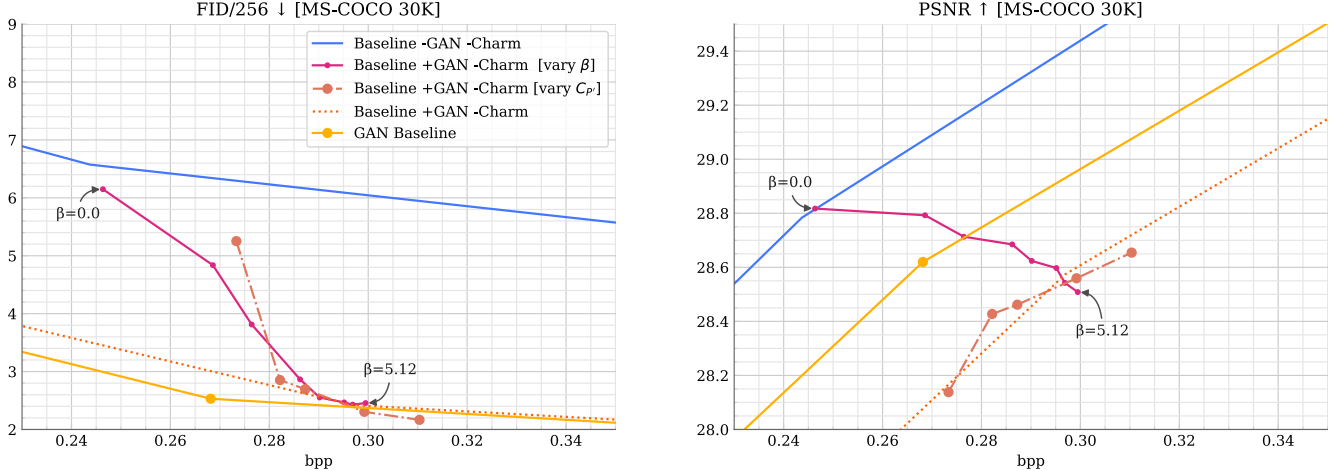


Figure 7. Ablations and tuning for our GAN Baseline. See Sec. 5.2 for details.

that we train a *single* model for the full spectrum from high-realism to low-distortion, our method at $\beta=0$ is competitive with VTM [7] and Charm [28] in terms of PSNR, and only $\approx 0.5dB$ below the state-of-the-art (with significantly better realism as discussed above).

Visual Comparison We compare visually to the generative method HiFiC, our strongest contender, in Fig. 2. We can observe that our reconstructions $\beta = 2.56$ are closer to the input than HiFiC when we compare at the same bitrate (first two rows), or we achieve the same visual quality at lower bitrates (third and fourth row). Our model achieves higher PSNR as we go towards $\beta=0$, but we lose detail in the reconstructions, as expected. For completeness, we compare to the two reconstructions available for the non-generative state-of-the-art ELIC in Sec. A.4. Visually, the reconstructions look similar to our model at $\beta=0$. Finally, we provide reconstructions on CLIC 2020 in Sec. A.7.

5.2. Ablations & Tuning

In Fig. 7 we show the results for our ablation experiments on MS-COCO. For the ablations, we decouple the weight on \mathcal{L}_P in Eq. 4 by setting $\beta C_P = C'_P$.

5.2.1 Loss Weights and Charm

Most ablation experiments are conducted using the Hyperprior [27] entropy model instead of Charm [28], which we denote with “-Charm”. We trained the baseline for various rate points with a fixed LPIPS weight $C'_P = 4.26$ (the default weight of HiFiC [25]) but without the GAN loss, which we refer to as “Baseline -GAN -Charm”. We then introduce the GAN loss, and vary first the weight of the GAN loss $\beta \in \{0.0, 0.08, 0.16, 0.32, \dots, 5.12\}$ (“Baseline +GAN -Charm [vary β]”). Here we find that the GAN loss lowers the FID up to $\beta = 2.56$ where it

starts saturating. With β fixed to this value, we now vary $C'_P \in \{0.0, 1.0, 2.0, 4.26, 8.0\}$ and find that the weight 4.26 remains a good choice (“Baseline +GAN -Charm [vary C'_P]”). The resulting model is called “Baseline +GAN -Charm”. Now we adopt Charm [28] entropy modeling which lowers the bitrate (resulting in “GAN Baseline”).

For the main model, we sampled β uniformly in the range $[0, 5.12]$ during training, but also found that the FID score at inference is lowest for $\beta=2.56$. We thus set $\beta_{\max}=2.56$ for our model during inference.

5.2.2 Conditioning

When comparing *FourierCond* and *TableCond* we found they give very similar results, see Sec. A.6. We choose *FourierCond* for our main model as it leads to a simpler implementation.

6. Conclusion

We have presented a method which is capable of outputting a single representation for compressed images, from which a receiver can either decode a high-realism reconstruction (high β) or a high-PSNR reconstruction (low β). We saw that in terms of distortion (PSNR) vs. realism (FID), our method can reach a new state-of-the-art.

To the best of our knowledge, this is the first single decoder method which allows the trade-off between realism and distortion to happen on the receiver side, with no change in the bitstream. This means that depending on the use-case, the users may choose to view reconstructions which are as close to the original as possible, or switch to view images with a better level of (generated) detail.

Somewhat surprisingly, we find that we can obtain high realism without sacrificing PSNR by more than $\approx 1 - 1.5dB$. We hope our findings inspire further work to push the boundary of the realism-distortion trade-off.

References

- [1] Eirikur Agustsson, Michael Tschannen, Fabian Mentzer, Radu Timofte, and Luc Van Gool. Generative adversarial networks for extreme learned image compression. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2, 3, 4
- [2] Johannes Ballé, Philip A Chou, David Minnen, Saurabh Singh, Nick Johnston, Eirikur Agustsson, Sung Jin Hwang, and George Toderici. Nonlinear transform coding. *IEEE Journal of Selected Topics in Signal Processing*, 15(2):339–353, 2020. 4, 5
- [3] 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 (ICLR)*, 2018. 3, 4
- [4] Johannes Ballé, Sung Jin Hwang, and Eirikur Agustsson. TensorFlow Compression: Learned data compression, 2022. 11
- [5] Fabrice Bellard. BPG Image format. <https://bellard.org/bpg/>, 2022. 7
- [6] Yochai Blau and Tomer Michaeli. Rethinking lossy compression: The rate-distortion-perception tradeoff. *arXiv preprint arXiv:1901.07821*, 2019. 1, 2
- [7] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer Ohm. Overview of the versatile video coding (vvc) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10):3736–3764, 2021. 7, 8
- [8] Zhengxue Cheng, Heming Sun, Masaru Takeuchi, and Jiro Katto. Learned image compression with discretized gaussian mixture likelihoods and attention modules. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7939–7948, 2020. 4
- [9] Chenjian Gao, Tongda Xu, Dailan He, Hongwei Qin, and Yan Wang. Flexible neural image compression via code editing. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 3
- [10] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2414–2423, 2016. 3
- [11] Fraunhofer Gesellschaft. VTM 17.1. https://vcgit.hhi.fraunhofer.de/jvet/VVCSsoftware_VTM/-/releases/VTM-17.1, 2022. 1, 7, 11
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 4
- [13] 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. 1, 3, 4, 5, 6, 7, 11, 13
- [14] 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 *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1763–1768, 2022. 1, 2, 3, 7
- [15] 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. 4
- [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. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017. 2, 6
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 5
- [18] Shoma Iwai, Tomo Miyazaki, Yoshihiro Sugaya, and Shinichiro Omachi. Fidelity-controllable extreme image compression with generative adversarial networks. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 8235–8242. IEEE, 2021. 3, 7, 11, 12
- [19] Kodak PhotoCD dataset. <http://r0k.us/graphics/kodak/>, 2022. 6
- [20] A Burakhan Koyuncu, Han Gao, and Eckehard Steinbach. Contextformer: A transformer with spatio-channel attention for context modeling in learned image compression. *arXiv preprint arXiv:2203.02452*, 2022. 3, 6, 7
- [21] Meng Li, Shangyin Gao, Yihui Feng, Yibo Shi, and Jing Wang. Content-oriented learned image compression. *arXiv preprint arXiv:2207.14168*, 2022. 3
- [22] Y. Ma, Y. Zhai, W. Jiang, I. Li, Z. Yang, and R. Wang. Roi image codec optimized for visual quality. In "2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)", 2022. 3, 7
- [23] Fabian Mentzer, Eirikur Agustsson, Johannes Ballé, David Minnen, Nick Johnston, and George Toderici. Towards generative video compression. *arXiv preprint arXiv:2107.12038*, 2021. 1
- [24] Fabian Mentzer, Eirikur Agustsson, Michael Tschannen, Radu Timofte, and Luc Van Gool. Conditional probability models for deep image compression. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 4
- [25] Fabian Mentzer, George D Toderici, Michael Tschannen, and Eirikur Agustsson. High-fidelity generative image compression. *Advances in Neural Information Processing Systems*, 33, 2020. 1, 2, 4, 5, 6, 8, 13
- [26] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 5
- [27] David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. In *Advances in Neural Information Processing Systems*, pages 10771–10780, 2018. 4, 8

- [28] David Minnen and Saurabh Singh. Channel-wise autoregressive entropy models for learned image compression. *arXiv preprint arXiv:2007.08739*, 2020. 4, 5, 6, 7, 8
- [29] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014. 3
- [30] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 6
- [31] Oren Rippel and Lubomir Bourdev. Real-time adaptive image compression. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2922–2930, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. 2
- [32] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. 6
- [33] Shibani Santurkar, David Budden, and Nir Shavit. Generative compression. *arXiv preprint arXiv:1703.01467*, 2017. 2
- [34] Claude E Shannon et al. Coding theorems for a discrete source with a fidelity criterion. *IRE Nat. Conv. Rec.*, 4(142-163):1, 1959. 1
- [35] Lucas Theis and Eirikur Agustsson. On the advantages of stochastic encoders. *arXiv preprint arXiv:2102.09270*, 2021. 1
- [36] Lucas Theis and Noureldin Y Ahmed. Algorithms for the communication of samples. In *International Conference on Machine Learning*, pages 21308–21328. PMLR, 2022. 3
- [37] Lucas Theis, Tim Salimans, Matthew D Hoffman, and Fabian Mentzer. Lossy compression with gaussian diffusion. *arXiv preprint arXiv:2206.08889*, 2022. 1
- [38] Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszar. Lossy image compression with compressive autoencoders. In *International Conference on Learning Representations (ICLR)*, 2017. 3
- [39] Lucas Theis and Aaron B Wagner. A coding theorem for the rate-distortion-perception function. *arXiv preprint arXiv:2104.13662*, 2021. 1
- [40] George Toderici, Lucas Theis, Nick Johnston, Eirikur Agustsson, Fabian Mentzer, Johannes Ballé, Wenzhe Shi, and Radu Timofte. *CLIC 2020: Challenge on Learned Image Compression*, 2020. <https://www.tensorflow.org/datasets/catalog/clic>. 6
- [41] Michael Tschannen, Eirikur Agustsson, and Mario Lucic. Deep generative models for distribution-preserving lossy compression. In *Advances in Neural Information Processing Systems*, pages 5929–5940, 2018. 1, 6
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 5
- [43] Gregory K Wallace. The JPEG still picture compression standard. *IEEE transactions on consumer electronics*, 38(1):xviii–xxxiv, 1992. 1
- [44] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multi-scale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402. Ieee, 2003. 3
- [45] Ren Yang, Luc Van Gool, and Radu Timofte. Perceptual learned video compression with recurrent conditional gan. *arXiv preprint arXiv:2109.03082*, 2021. 1
- [46] Y. Yang, S. Mandt, and L. Theis. An introduction to neural data compression. preprint, 2022. 4
- [47] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gungjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2022. 6
- [48] George Zhang, Jingjing Qian, Jun Chen, and Ashish Khisti. Universal rate-distortion-perception representations for lossy compression. *Advances in Neural Information Processing Systems*, 34:11517–11529, 2021. 1, 3
- [49] 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. 3, 5
- [50] Renjie Zou, Chunfeng Song, and Zhaoxiang Zhang. The devil is in the details: Window-based attention for image compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17492–17501, 2022. 4