
Compressing Images by Encoding Their Latent Representations with Relative Entropy Coding

Gergely Flamich*

Department of Engineering
University of Cambridge
gf332@cam.ac.uk

Marton Havasi*

Department of Engineering
University of Cambridge
mh740@cam.ac.uk

José Miguel Hernández-Lobato

Department of Engineering
University of Cambridge,
Microsoft Research,
Alan Turing Institute
jmh233@cam.ac.uk

Abstract

Variational Autoencoders (VAEs) have seen widespread use in learned image compression. They are used to learn expressive latent representations on which downstream compression methods can operate with high efficiency. Recently proposed ‘bits-back’ methods can indirectly encode the latent representation of images with codelength close to the relative entropy between the latent posterior and the prior. However, due to the underlying algorithm, these methods can only be used for lossless compression, and they only achieve their nominal efficiency when compressing multiple images simultaneously; they are inefficient for compressing single images. As an alternative, we propose a novel method, Relative Entropy Coding (REC), that can directly encode the latent representation with codelength close to the relative entropy for single images, supported by our empirical results obtained on the Cifar10, ImageNet32 and Kodak datasets. Moreover, unlike previous bits-back methods, REC is immediately applicable to lossy compression, where it is competitive with the state-of-the-art on the Kodak dataset.

1 Introduction

The recent development of powerful generative models, such as Variational Autoencoders (VAEs) has caused a great deal of interest in their application to image compression, notably [Ballé et al. \(2016a, 2018\)](#); [Townsend et al. \(2020\)](#); [Minnen & Singh \(2020\)](#). The benefit of using these models as opposed to hand-crafted methods is that they can adapt to the statistics of their inputs more effectively, and hence allow significant gains in compression rate. A second advantage is their easier adaptability to new media formats, such as light-field cameras, 360° images, Virtual Reality (VR), video streaming, etc. for which classical methods are not currently applicable, or are not performant.

VAEs consist of two neural networks, the encoder and the decoder. The former maps images to their latent representations and the latter maps them back. Compression methods can operate very efficiently in latent space, thus realizing a non-linear transform coding method ([Goyal, 2001](#); [Ballé et al., 2016b](#)). The sender can use the encoder to obtain the latent posterior of an image and then use the compression algorithm to transmit a sample latent representation from the posterior. Then, the receiver can use the decoder to reconstruct the image from the latent representation they received. Note that this reconstruction contains small errors. In lossless compression, the sender must correct for this and has to transmit the residuals along with the latent code. In lossy compression, we omit the transmission of the residuals, as the model is optimized such that the reconstruction retains high perceptual quality.

*Equal contribution.

Bits-back methods for lossless compression have been at the center of attention recently. They realize the optimal compression rate postulated by the bits-back argument (Hinton & Van Camp, 1993): For a given image, the optimal compression rate using a latent variable model (such as a VAE) is given by the relative entropy between the latent posterior and the prior $\text{KL}[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]$ plus the expected residual error $\mathbb{E}[-\log P(\mathbf{x}|\mathbf{z})]$, where \mathbf{x} is the input image, \mathbf{z} denotes the stochastic latent representation, and $p(\mathbf{z})$ is the prior over the latent space. This quantity is also known as the negative Evidence Lower Bound (ELBO).

Current bits-back compression methods use variants of the Bits-Back with Asymmetric Numeral Systems (BB-ANS) algorithm (Townsend et al., 2019, 2020; Ho et al., 2019; Kingma et al., 2019). BB-ANS can achieve the bits-back compression rate asymptotically by allowing the codes in a sequence of images to overlap without losing information (hence getting the bits back). The issue is that the first image requires a string of auxiliary bits to start the sequence, which means that it is inefficient when used to compress a single image. Including the auxiliary bits, the compressed size of a single image is often 2-3 times the original size².

We introduce Relative Entropy Coding (REC), a lossless compression paradigm that subsumes bits-back methods. A REC method can encode a sample from the latent posterior $z \sim q(\mathbf{z}|\mathbf{x})$ with codelength close to the relative entropy $\text{KL}[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]$, given a shared source of randomness. Then the residuals are encoded using an entropy coding method such as arithmetic coding (Witten et al., 1987) with codelength $-\log P(\mathbf{x}|\mathbf{z})$. This yields a combined codelength close to the negative ELBO in expectation without requiring any auxiliary bits.

We propose a REC method, called index coding (iREC), based on importance sampling. To encode a sample from the posterior $q(\mathbf{z}|\mathbf{x})$, our method relies on a shared sequence of random samples from the prior $z_1, z_2, \dots \sim p(\mathbf{z})$, which in practice is realised using a pseudo-random number generator with a shared random seed. The algorithm selects an element z_i from the random sequence with high density under the posterior. Then, the code for the sample is simply ‘ i ’, its index in the sequence. Given i , the receiver can recover z_i by selecting the i th element from the shared random sequence. We show that the codelength of ‘ i ’ is close to the relative entropy $\text{KL}[q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})]$.

Apart from eliminating the requirement for auxiliary bits, REC offers a few further advantages. First, an issue concerning virtually every deep image compression algorithm is that they require a quantized latent space for encoding. This introduces an inherently non-differentiable step to training, which hinders performance and, in some cases, prevents model scaling (Hooeboom et al., 2019). Since our method relies on a shared sequence of prior samples, it is not necessary to quantize the latent space and it can be applied to off-the-shelf VAE architectures with continuous latent spaces. To our knowledge, our method is the first image compression algorithm that can operate in a continuous latent space.

Second, since the codelength scales with the relative entropy but not the number of latent dimensions, the method is unaffected by the pruned dimensions of the VAE, i.e. dimensions where the posterior collapses back onto the prior (Yang et al., 2020; Lucas et al., 2019) (this advantage is shared with bits-back methods, but not others in general).

Third, our method elegantly extends to lossy compression. If we choose to only encode a sample from the posterior using REC without the residuals, the receiver can use the decoder to reconstruct an image that is close to the original. In this setting, the VAE is optimized using an objective that allows the model to maintain high *perceptual quality* even at low bit rates. The compression rate of the model can be precisely controlled during training using a β -VAE-like training objective (Higgins et al., 2017). Our empirical results confirm that our REC algorithm is competitive with the state-of-the-art in lossy image compression on the Kodak dataset (Eastman Kodak Company, 1999).

The key contributions of this paper are as follows:

- iREC, a relative entropy coding method that can encode an image with codelength close to the negative ELBO for VAEs. Unlike prior bits-back methods, it does not require auxiliary bits, and hence it is efficient for encoding single images. We empirically confirm these findings on the Cifar10, ImageNet32 and Kodak datasets.
- Our algorithm forgoes the quantization of latent representations entirely, hence it is directly applicable to off-the-shelf VAE architectures with continuous latent spaces.

²Based on the number of auxiliary bits recommended by Townsend et al. (2020).

- The algorithm can be applied to lossy compression, where it is competitive with the state-of-the-art on the Kodak dataset.

2 Learned Image Compression

The goal of compression is to communicate some information between the sender and the receiver using as little bandwidth as possible. Lossless compression methods assign some *code* $C(x)$ to some input x , consisting of a sequence of bits, such that the original x can be always recovered from $C(x)$. The efficiency of the method is determined by the average length of $C(x)$.

All compression methods operate on the same underlying principle: *Commonly occurring patterns are assigned shorter codelengths while rare patterns have longer codelengths*. This principle was first formalized by Shannon (1948). Given an underlying distribution $P(x)$ ³ where x is the input taking values in some set \mathcal{X} , the rate of compression cannot be better than $H[P] = \sum_{x \in \mathcal{X}} -P(x) \log P(x)$, the Shannon-entropy of P . This theoretical limit is achieved when the codelength of a given x is close to the negative log-likelihood $|C(x)| \approx -\log P(x)$. Methods that get close to this limit are referred to as entropy coding methods, the most prominent ones being Huffman coding and arithmetic coding (Huffman, 1952; Witten et al., 1987).

The main challenge in image compression is that the distribution $P(x)$ is not readily available. Methods either have to hand-craft $P(x)$ or learn it from data. The appeal in using generative models to learn $P(x)$ is that they can give significantly better approximations than traditional approaches.

2.1 Image Compression Using Variational Autoencoders

A Variational Autoencoder (VAE, Kingma & Welling (2014)) is a generative model that learns the underlying distribution of a dataset in an unsupervised manner. It consists of a pair of neural networks called the *encoder* and *decoder*, that are approximate inverses of each other. The encoder network takes an input \mathbf{x} and maps it to a posterior distribution over the latent representations $q_\phi(\mathbf{z} | \mathbf{x})$. The decoder network maps a latent representation $\mathbf{z} \sim q_\phi(\mathbf{z} | \mathbf{x})$ to the conditional distribution $P_\theta(\mathbf{x} | \mathbf{z})$. Here, ϕ and θ denote the parameters of the encoder and the decoder, respectively. These two networks are trained jointly by maximizing a lower bound to the marginal log-likelihood $\log P(\mathbf{x})$, the Evidence Lower Bound $\mathcal{L}(\mathbf{x}, \phi, \theta)$ (ELBO):

$$\log P(\mathbf{x}) \geq \mathcal{L}(\mathbf{x}, \phi, \theta) = \underbrace{\mathbb{E}_{\mathbf{z} \sim q_\phi(\mathbf{z} | \mathbf{x})} [\log P_\theta(\mathbf{x} | \mathbf{z})]}_{\text{conditional log-likelihood}} - \underbrace{\text{KL} [q_\phi(\mathbf{z} | \mathbf{x}) || p(\mathbf{z})]}_{\text{relative entropy}}. \quad (1)$$

The VAE can be used to realize a non-linear form of transform coding (Goyal, 2001) to perform image compression. Given an image \mathbf{x} , the sender maps it to its latent posterior $q(\mathbf{z} | \mathbf{x})$, and communicates a sample $\mathbf{z} \sim q(\mathbf{z} | \mathbf{x})$ to the receiver (where we omit ϕ and θ for notational ease). Our proposed algorithm can accomplish this with communication cost close to the relative entropy $\text{KL} [q(\mathbf{z} | \mathbf{x}) || p(\mathbf{z})]$. Given \mathbf{z} , the receiver can use the decoder to obtain the conditional distribution $P(\mathbf{x} | \mathbf{z})$, which can be used for both lossless and lossy compression. Figure 1 depicts both processes.

For **lossless** compression, the sender can use an entropy coding method with $P(\mathbf{x} | \mathbf{z})$ to encode the residuals. The cost of communicating the residuals is the negative log-likelihood $-\log P(\mathbf{x} | \mathbf{z})$ which yields a combined codelength close to the negative ELBO. (See Figure 33c.)

For **lossy** compression, the most commonly used approach is to take the mean of the conditional distribution to be the approximate reconstruction $\tilde{\mathbf{x}} = \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x} | \mathbf{z})} [\mathbf{x}]$. This yields a reconstruction close to the original image, while only having to communicate \mathbf{z} . (See Figure 1b.)

The remaining question is how to communicate a sample \mathbf{z} from the posterior $q(\mathbf{z} | \mathbf{x})$ given the shared prior $p(\mathbf{z})$. The most widely given answer is the **quantization** of the latent representations, which are then encoded using entropy coding (Theis et al., 2017; Ballé et al., 2016a). This approach is simple to use but has two key weaknesses. First, because of the quantized latent space, the posterior is a discrete probability distribution, which is significantly more difficult to train with gradient descent than its continuous counterpart. It requires the use of gradient estimators and does not scale well with

³In this paper use roman letters (e.g. \mathbf{x}) for random variables and italicized letters (e.g. \mathbf{x}) for their realizations. Bold letters denote vectors. By a slight abuse of notation, we denote $p(\mathbf{x} = \mathbf{x})$ by $p(\mathbf{x})$.

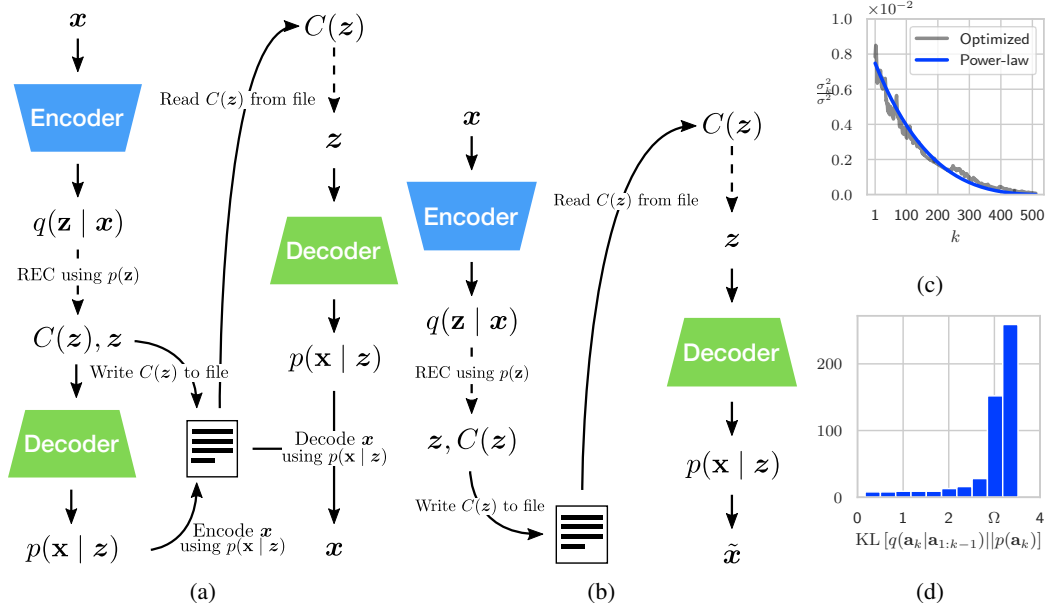


Figure 1: **(a)** Lossless compression using REC **(b)** Lossy compression using REC **(c)** The variances of the coding distributions of the auxiliary variables. We observe that the individually optimized values are well approximated by a power-law. **(d)** The relative entropies of the auxiliary variables are near or below Ω . ((c) and (d) depict statistics from the 23rd stochastic layer of a 24-layer ResNet VAE, since this layer contains the majority of the model’s total relative entropy.)

depth (Hooeboom et al., 2019). A second known shortcoming of this method is that the codelength scales with the number of latent dimensions even when those dimensions are pruned by the VAE, i.e. the posterior coincides with the prior and the dimension, therefore, carries no information (Yang et al., 2020; Lucas et al., 2019).

An alternative to quantization in lossless compression is **bits-back coding** (Townsend et al., 2019). It uses a string of auxiliary bits to encode z , which is then followed by encoding the residuals. When compressing multiple images, bits-back coding reuses the code of already compressed images as auxiliary bits to compress the remaining ones, bringing the asymptotic cost close to the negative ELBO. However, due to the use of auxiliary bits, it is inefficient to use for single images.

Relative Entropy Coding (REC), rectifies the aforementioned shortcomings. We propose a REC algorithm that can encode z with codelength close to the relative entropy, without requiring quantization or auxiliary bits. It is effective for both lossy and lossless compression.

3 Relative Entropy Coding

Relative Entropy Coding (REC) is a *lossless compression paradigm* that solves the problem of communicating a sample from the posterior distribution $q(\mathbf{z} | \mathbf{x})$ given the shared prior distribution $p(\mathbf{z})$. In more general terms, the sender wants to communicate a sample $z \sim q(\mathbf{z})$ to the receiver from a target distribution (the target distribution is only known to the sender) with a coding distribution $p(\mathbf{z})$ shared between the sender and the receiver, given a *shared source of randomness*. That is, over many runs, the empirical distribution of the transmitted z s converges to $q(\mathbf{z})$. Hence, REC is a *stochastic coding scheme*, in contrast with entropy coding which is fully deterministic.

We refer to algorithms that achieve communication cost provably close to $\text{KL}[q(\mathbf{z}) || p(\mathbf{z})]$ as REC algorithms. We emphasize the counter-intuitive notion, that communicating a stochastic sample from $q(\mathbf{z})$ can be much cheaper than communicating any *specific* sample z . For example, consider the case when $p(\mathbf{z}) = q(\mathbf{z}) = \mathcal{N}(\mathbf{z} | 0, I)$. First, consider the naive approach of sampling $z \sim q(\mathbf{z})$ and encoding it with entropy coding. The expected codelength of z is ∞ since the Shannon entropy of a Gaussian random variable is ∞ . Now consider a REC approach: the sender could simply indicate to the receiver that they should draw a sample from the shared coding distribution $p(\mathbf{z})$, which has

$\mathcal{O}(1)$ communication cost. This problem is studied formally in [Harsha et al. \(2007\)](#). They show that the relative entropy is a lower bound to the codelength under mild assumptions. As part of a formal proof, they present a rejection sampling algorithm that in practice is computationally intractable even for small problems, however, it provides a basis for our REC algorithm presented below.

3.1 Relative Entropy Coding with Index Coding

We present Index Coding (iREC), a REC algorithm that scales to the needs of modern image compression problems. The core idea for the algorithm is to rely on a shared source of randomness between the sender and the receiver, which takes the form of an infinite sequence of random samples from the prior $\mathbf{z}_1, \mathbf{z}_2, \dots \sim p(\mathbf{z})$. This can be practically realized using a pseudo-random number generator with a shared random seed. Then, to communicate an element \mathbf{z}_i from the sequence, it is sufficient to transmit its index, i.e. $C(\mathbf{z}_i) = i$. From i the receiver can reconstruct \mathbf{z}_i by using the shared source of randomness to generate $\mathbf{z}_1, \mathbf{z}_2, \dots$ and then selecting the i th element.

iREC is based on the importance sampling procedure proposed in [Havasi et al. \(2019\)](#). Let $M = \lceil \exp(\text{KL}[q(\mathbf{z}) \parallel p(\mathbf{z})]) \rceil$. Our algorithm draws M samples $\mathbf{z}_1, \dots, \mathbf{z}_M \sim p(\mathbf{z})$, then selects \mathbf{z}_i with probability proportional to the importance weights $P_{\text{accept}}(\mathbf{z}_m) \propto \frac{q(\mathbf{z}_m)}{p(\mathbf{z}_m)}$ for $m \in \{1, \dots, M\}$. Although \mathbf{z}_i is not an unbiased sample from $q(\mathbf{z})$, [Havasi et al. \(2019\)](#) show that considering M samples is sufficient to ensure that the bias remains low. The cost of communicating \mathbf{z}_i is simply $\log M \approx \text{KL}[q(\mathbf{z}) \parallel p(\mathbf{z})]$, since we only need to communicate i where $1 \leq i \leq M$.

Importance sampling has promising theoretical properties, however, it is still infeasible to use in practice because M grows exponentially with the relative entropy. To drastically reduce this cost, we propose sampling a sequence of auxiliary variables instead of sampling $q(\mathbf{z})$ directly.

3.1.1 Using Auxiliary Variables

We propose breaking \mathbf{z} up into a sequence of K auxiliary random variables $\mathbf{a}_{1:K} = \mathbf{a}_1, \dots, \mathbf{a}_K$ with independent coding distributions $p(\mathbf{a}_1) \dots, p(\mathbf{a}_K)$ such that they fully determine \mathbf{z} , i.e. $\mathbf{z} = f(\mathbf{a}_{1:K})$ for some function f . Our goal is to derive target distributions $q(\mathbf{a}_k \mid \mathbf{a}_{1:k-1})$ for each of the auxiliary variables given the previous ones, such that by sampling each of them via importance sampling, i.e. $\mathbf{a}_k \sim q(\mathbf{a}_k \mid \mathbf{a}_{1:k-1})$ for $k \in \{1, \dots, K\}$, we get a sample $\mathbf{z} = f(\mathbf{a}_{1:K}) \sim q(\mathbf{z})$. This implies the auxiliary coding distributions $p(\mathbf{a}_k)$ and target distributions $q(\mathbf{a}_k \mid \mathbf{a}_{1:k-1})$ must satisfy the marginalization properties

$$p(\mathbf{z}) = \int \delta(f(\mathbf{a}_{1:K}) - \mathbf{z}) p(\mathbf{a}_{1:K}) d\mathbf{a}_{1:K} \quad \text{and} \quad q(\mathbf{z}) = \int \delta(f(\mathbf{a}_{1:K}) - \mathbf{z}) q(\mathbf{a}_{1:K}) d\mathbf{a}_{1:K}, \quad (2)$$

where δ is the Dirac delta function, $p(\mathbf{a}_{1:K}) = \prod_{k=1}^K p(\mathbf{a}_k)$ and $q(\mathbf{a}_{1:K}) = \prod_{k=1}^K q(\mathbf{a}_k \mid \mathbf{a}_{1:k-1})$. Note that the coding distributions $p(\mathbf{a}_1) \dots, p(\mathbf{a}_K)$ and f can be freely chosen subject to Eq 2.

The cost of encoding each auxiliary variable using importance sampling is equal to the relative entropy between their corresponding target and coding distributions. Hence, to avoid introducing an overhead to the overall codelength, the targets must satisfy $\text{KL}[q(\mathbf{a}_{1:k}) \parallel p(\mathbf{a}_{1:k})] = \text{KL}[q(\mathbf{z}) \parallel p(\mathbf{z})]$ ⁴. To ensure this, for fixed f and joint auxiliary coding distribution $p(\mathbf{a}_{1:K})$, the joint auxiliary target distributions must have the form

$$q(\mathbf{a}_{1:k}) := \int p(\mathbf{a}_{1:k} \mid \mathbf{z}) q(\mathbf{z}) d\mathbf{z} \quad \text{for } k \in \{1 \dots K\}. \quad (3)$$

These are the only possible auxiliary targets that satisfy the condition on the sum of the relative entropies, which we formally show in the supplementary material.

Choosing the forms of the auxiliary variables. Factorized Gaussian priors are a popular choice for VAEs. For a Gaussian coding distribution $p(\mathbf{z}) = \mathcal{N}(0, \sigma^2 I)$, we propose $\mathbf{z} = \sum_{k=1}^K \mathbf{a}_k$, $p(\mathbf{a}_k) = \mathcal{N}(0, \sigma_k^2 I)$ for $k \in \{1, \dots, K\}$ such that $\sum_{k=1}^K \sigma_k^2 = \sigma^2$. The targets $q(\mathbf{a}_k \mid \mathbf{a}_{1:k-1})$ in this case turn out to be Gaussian as well, and their form is derived in the supplementary material.

⁴We use here the definition $\text{KL}[q(x|y) \parallel p(x)] = \mathbb{E}_{x,y \sim p(x,y)} \left[\log \frac{q(x|y)}{p(x)} \right]$ ([Cover & Thomas \(2012\)](#)).

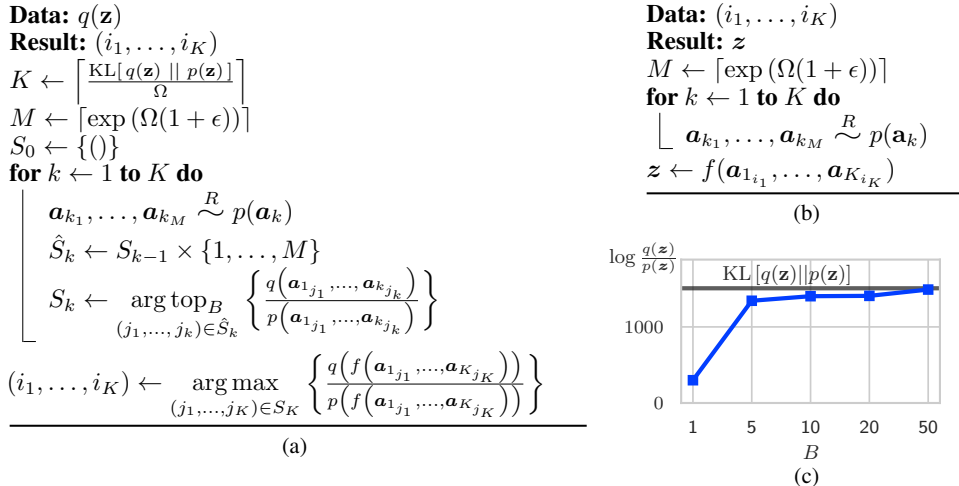


Figure 2: **(a)** iREC encoder **(b)** iREC decoder **(c)** Beam search ensures that $\log \frac{q(\mathbf{z})}{p(\mathbf{z})}$ is close to the relative entropy. B is the number of beams. Plotted using the 23rd stochastic layer of a 24-layer ResNet VAE, since this layer contains the majority of the model’s total relative entropy. Here, $\overset{R}{\sim}$ indicates sampling using a pseudo-random number generator with random seed R , and $\arg \text{top}_B$ selects the arguments of the top B ranking elements in a set.

To guarantee that every auxiliary variable can be encoded via importance sampling, the relative entropies should be similar across the auxiliary variables, i.e. $\text{KL}[q(\mathbf{a}_k | \mathbf{a}_{1:k-1}) \parallel p(\mathbf{a}_k)] \approx \Omega$, where Ω is a hyperparameter (we used $\Omega = 3$ in our experiments). This yields $K = \lceil \text{KL}[q(\mathbf{z}) \parallel p(\mathbf{z})] / \Omega \rceil$ auxiliary variables in total. We initially set the auxiliary coding distributions by optimizing their variances σ_k^2 on a small validation set to achieve relative entropies close to Ω . Later, we found that the ratio of the variance σ_k^2 of the k -th auxiliary variable to the remaining variance $\sigma^2 - \sum_{j=1}^{k-1} \sigma_j^2$ is well approximated by the power law $(K + 1 - k)^{-0.79}$, as shown in Figure 1c. In practice, we used this approximation to set each σ_k^2 . With these auxiliary variables, we fulfil the requirement of keeping the individual relative entropies near or below Ω as shown empirically in Figure 1d. To account for the auxiliary variables whose relative entropy slightly exceeds Ω , in practice, we draw $M = \lceil \exp(\Omega(1 + \epsilon)) \rceil$ samples, where ϵ is a small non-negative constant (we used $\epsilon = 0.2$ for lossless and $\epsilon = 0.0$ for lossy compression), leading to a $(1 + \epsilon)$ increase of the codelength.

3.1.2 Reducing the Bias with Beam Search

An issue with naively applying the auxiliary variable scheme is that the cumulative bias of importance sampling each auxiliary variable adversely impacts the compression performance, potentially leading to higher distortion and hence longer codelength.

To reduce this bias, we propose using a beam search algorithm (shown in Figure 2) to search over multiple possible assignments of the auxiliary variables. We maintain a set of the B lowest bias samples from the auxiliary variables and used the log-importance weight $\log \frac{q(\mathbf{z})}{p(\mathbf{z})}$ as a heuristic measurement of the bias. For an unbiased sample $\log \frac{q(\mathbf{z})}{p(\mathbf{z})} \approx \text{KL}[q(\mathbf{z}) \parallel p(\mathbf{z})]$, but for a biased sample, $\log \frac{q(\mathbf{z})}{p(\mathbf{z})} \ll \text{KL}[q(\mathbf{z}) \parallel p(\mathbf{z})]$. For each auxiliary variable \mathbf{a}_k in the sequence, we combine the B lowest bias samples for $\mathbf{a}_{1:k-1}$ with the M possible importance samples for \mathbf{a}_k . To choose the B lowest bias samples for the next iteration, we take the top B samples with the highest importance weights $\frac{q(\mathbf{a}_{1:k})}{p(\mathbf{a}_{1:k})}$ out of the $B \times M$ possibilities. At the end, we select $\mathbf{a}_{1:K}$ with the highest importance weight $\frac{q(f(\mathbf{a}_{1:K}))}{p(f(\mathbf{a}_{1:K}))} = \frac{q(\mathbf{z})}{p(\mathbf{z})}$.

3.1.3 Determining the hyperparameters

Finding good values for Ω , ϵ and B is crucial for the good performance of our method. We want as short a codelength as possible, while also minimizing computational cost. Therefore, we ran a grid search over a reasonable range of parameter settings for the lossless compression of a small number of ImageNet32 images. We compare the efficiency of settings by measuring the codelength overhead they produced in comparison to the ELBO, which represents optimal performance. We find that for reasonable settings of Ω (between 5-3) and for fixed ϵ , regular importance sampling ($B = 1$) gives between 25-80% overhead, whereas beam search with $B = 5$ gives 15-25%, and with $B = 20$ it gives 10-15%. Setting $B > 20$ does not result in significant improvements in overhead, while the computational cost is heavily increased. Thus, we find that 10-20 beams are sufficient to significantly reduce the bias as shown in Figure 2c. The details of our experimental setup and the complete report of our findings can be found in the supplementary material.

4 Experiments

We compare our method against state-of-the-art lossless and lossy compression methods. Our experiments are implemented in TensorFlow (Abadi et al., 2015) and are publicly available at <https://github.com/gergely-flamich/relative-entropy-coding>.

4.1 Lossless Compression

We compare our method on single image lossless compression (shown in Table 1) against PNG, WebP and FLIF, Integer Discrete-Flows (Hoogetboom et al., 2019) and the prominent bits-back approaches: Local Bits-Back Coding (Ho et al., 2019), BitSwap (Kingma et al., 2019) and HiLLOc (Townsend et al., 2020).

Our model for these experiments is a ResNet VAE (RVAE) (Kingma et al., 2016) with 24 Gaussian stochastic levels. This model utilizes skip-connections to prevent the posteriors on the higher stochastic levels to collapse onto the prior and achieves an ELBO that is competitive with the current state-of-the-art auto-regressive models. For better comparison, we used the exact model used by Townsend et al. (2020)⁵ trained on ImageNet32. The three hyperparameters of iREC are set to $\Omega = 3$, $\epsilon = 0.2$ and $B = 20$. Further details on our hyperparameter tuning process are included in the supplementary material.

We evaluated the methods on Cifar10 and ImageNet32 comprised of 32×32 images, and the Kodak dataset comprised of full-sized images. For Cifar10 and ImageNet32, we used a subsampled test set of size 1000 due to the speed limitation of our method (currently, it takes ~ 1 minute to compress a 32×32 image and 1-10 minutes to compress a large image). By contrast, decoding with our method is fast since it does not require running the beam search procedure.

iREC significantly outperforms other bits-back methods on all datasets since it does not require auxiliary bits, although it is still slightly behind non-bits-back methods as it has a $\sim 20\%$ overhead compared to the ELBO due to using $\epsilon = 0.2$.

4.2 Lossy Compression

On the lossy compression task, we present average rate-distortion curves calculated using the PSNR (Huynh-Thu & Ghanbari, 2008) and MS-SSIM (Wang et al., 2004) quality metrics on the Kodak dataset, shown in Figure 3. On both metrics, we compare against JPEG, BPG, Theis et al. (2017) and Ballé et al. (2018), with the first two being classical methods and the latter two being ML-based. Additionally, on PSNR we compare against Minnen & Singh (2020), whose work represents the current state-of-the-art to the best of our knowledge.⁶

⁵We used the publicly available trained weights published by the authors of Townsend et al. (2020).

⁶To overcome the issue of the inefficiency of bits-back methods for single or small-batch image compression, in practice an efficient single-image compression method (e.g. FLIF) is used to encode the first few images and only once the overhead of using bits-back methods becomes negligible do we switch to using them (see e.g. Townsend et al. (2020)).

⁷For all competing methods, we used publicly available data at https://github.com/tensorflow/compression/tree/master/results/image_compression.

Table 1: Single image, lossless compression performance in bits per dimension (lower is better). The best performing bits-back or REC method is highlighted for each dataset. The asymptotic rates are included in parenthesis where they are different from the single image case. To calculate the number of bits needed for single images, we added the number of auxiliary bits required to the asymptotic compression rate as reported in the respective papers.

		Cifar10 (32x32)	ImageNet32 (32x32)	Kodak (768x512)
<i>Non bits-back</i>	PNG	5.87	6.39	4.35
	WebP	4.61	5.29	3.20
	FLIF	4.19	4.52	2.90
	IDF	3.34	4.18	–
<i>Bits-back</i> ⁶	LBB	54.96 (3.12)	55.72 (3.88)	–
	BitSwap	6.53 (3.82)	6.97 (4.50)	–
	HiLLoC	24.51 (3.56)	26.80 (4.20)	17.5 (3.00)
<i>REC</i>	iREC (Ours)	4.18	4.91	3.67
	ELBO (RVAE)	[3.55]	[4.18]	[3.00]

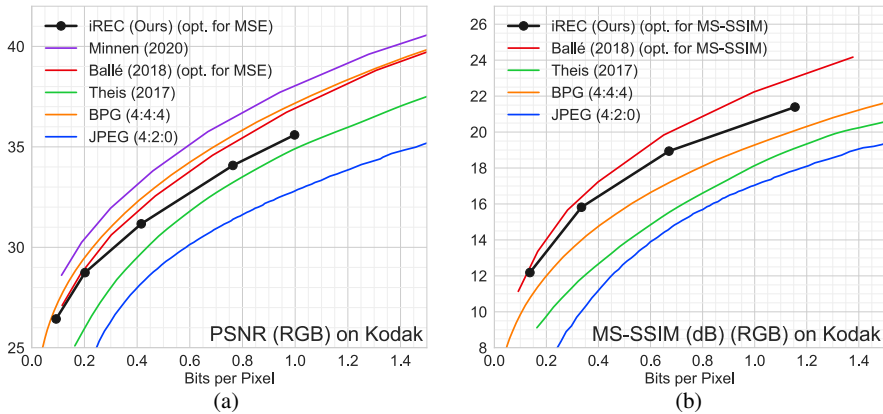


Figure 3: Comparison of REC against classical methods such as JPEG, BPG and competing ML-based methods. (a) PSNR comparisons (b) MS-SSIM comparisons in decibels, calculated using the formula $-10 \log_{10}(1 - \text{MS-SSIM})$. See the supplementary material for more comparisons.

We used the architecture presented in Ballé et al. (2018) with the latent distributions changed to Gaussians and a few small modifications to accommodate this; see the supplementary material for precise details. Following Ballé et al. (2018), we trained several models using $\mathcal{L}_\lambda(\mathbf{x}, \phi, \theta) = \lambda D(\mathbf{x}, \hat{\mathbf{x}}) - \text{KL}[q_\phi(\mathbf{z} | \mathbf{x}) || p(\mathbf{z})]$, where $\hat{\mathbf{x}}$ is the reconstruction of the image \mathbf{x} , and $D(\cdot, \cdot) \in \{\text{MSE}, \text{MS-SSIM}\}$ is a differentiable distortion metric.⁸ Varying λ in the loss yields models with different rate-distortion trade-offs. We optimized 5 models for MSE with $\lambda \in \{0.001, 0.003, 0.01, 0.03, 0.05\}$ and 4 models for MS-SSIM with $\lambda \in \{0.003, 0.01, 0.03, 0.08\}$. The hyperparameters of iREC were set this time to $\Omega = 3$, $\epsilon = 0$ and $B = 10$. As can be seen in Figure 3, iREC is competitive with the state-of-the-art lossy compression methods on both metrics.

5 Related Work

Multiple recent bits-back papers build on the work of Townsend et al. (2019), who first proposed BB-ANS. These works elevate the idea from a theoretical argument to a practically applicable algorithm. Kingma et al. (2019) propose BitSwap, an improvement to BB-ANS that allows it to be applied to

⁸In practice we used $1 - \text{MS-SSIM}$ as the loss function with power factors set to $\alpha = \beta = \gamma = 1$.

hierarchical VAEs. [Ho et al. \(2019\)](#) extend the work to general flow models, which significantly improves the asymptotic compression rate. Finally, [Townsend et al. \(2020\)](#) apply the original BB-ANS algorithm to full-sized images and improve its run-time by vectorizing its implementation.

All previous VAE-based lossy compression methods use entropy coding with quantized latent representations. They propose different approaches to circumvent the non-differentiability that this presents during training. Most prominently, [Ballé et al. \(2016a\)](#) describe a continuous relaxation of quantization based on dithering. Building on this idea, [Ballé et al. \(2018\)](#) introduce a 2-level hierarchical architecture, and [Minnen et al. \(2018\)](#); [Lee et al. \(2019\)](#); [Minnen & Singh \(2020\)](#) explore more expressive latent representations, such as using learned adaptive context models to aid their entropy coder, and autoregressive latent distributions.

6 Conclusion

This paper presents iREC, a REC algorithm, which extends the importance sampler proposed by [Havasi et al. \(2019\)](#). It enables the use of latent variable models (and VAEs in particular) with continuous probability distributions for both lossless and lossy compression.

Our method significantly outperforms bits-back methods on lossless single image compression benchmarks and is competitive with the asymptotic performance of competing methods. On the lossy compression benchmarks, our method is competitive with the state-of-the-art for both the PSNR and MS-SSIM perceptual quality metrics. Currently, the main practical limitation of bits-back methods, including our method, is the compression speed, which we hope to improve in future works.

7 Acknowledgements

Marton Havasi is funded by EPSRC. We would like to thank Stratis Markou and Ellen Jiang for the helpful comments on the manuscript.

8 Broader Impact

Our work presents a novel data compression framework and hence inherits both its up and downsides. In terms of positive societal impacts, data compression reduces the bandwidth requirements for many applications and websites, making them more inexpensive to access. This increases accessibility to online content in rural areas with limited connectivity or underdeveloped infrastructure. Moreover, it reduces the energy requirement and hence the environmental impact of information processing systems. However, care must be taken when storing information in a compressed form for long time periods, and backwards-compatibility of decoders must be maintained, as data may otherwise be irrevocably lost, leading to what has been termed the Digital Dark Ages ([Kuny, 1997](#)).

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimized image compression. *International Conference on Learning Representations*, 2016a. URL <https://openreview.net/pdf?id=rJxdQ3jeg>
- Johannes Ballé, Valero Laparra, and Eero P Simoncelli. End-to-end optimization of nonlinear transform codes for perceptual quality. In *2016 Picture Coding Symposium (PCS)*, pp. 1–5. IEEE, 2016b. URL <https://ieeexplore.ieee.org/iel7/7897355/7906305/07906310.pdf>

- 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. URL <https://openreview.net/forum?id=rkcQFMZRB>.
- Thomas M Cover and Joy A Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012.
- Eastman Kodak Company. Kodak lossless true color image suite. <http://r0k.us/graphics/kodak/>, 1999.
- Vivek K Goyal. Theoretical foundations of transform coding. *IEEE Signal Processing Magazine*, 18(5):9–21, 2001. URL <https://ieeexplore.ieee.org/iel5/79/20598/00952802.pdf>.
- Prahladh Harsha, Rahul Jain, David McAllester, and Jaikumar Radhakrishnan. The communication complexity of correlation. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pp. 10–23. IEEE, 2007. URL <https://ieeexplore.ieee.org/iel5/4262736/4262737/04262747.pdf>.
- Marton Havasi, Robert Peharz, and José Miguel Hernández-Lobato. Minimal random code learning: Getting bits back from compressed model parameters. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/pdf?id=r1f0YiCctm>.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/pdf?id=Sy2fzU9gl>.
- Geoffrey Hinton and Drew Van Camp. Keeping neural networks simple by minimizing the description length of the weights. In *in Proc. of the 6th Ann. ACM Conf. on Computational Learning Theory*. Citeseer, 1993. URL <https://dl.acm.org/doi/pdf/10.1145/168304.168306>.
- Jonathan Ho, Evan Lohn, and Pieter Abbeel. Compression with flows via local bits-back coding. In *Advances in Neural Information Processing Systems*, pp. 3874–3883, 2019. URL <https://papers.nips.cc/paper/8643-compression-with-flows-via-local-bits-back-coding.pdf>.
- Emiel Hoogeboom, Jorn Peters, Rianne van den Berg, and Max Welling. Integer discrete flows and lossless compression. In *Advances in Neural Information Processing Systems*, pp. 12134–12144, 2019. URL <https://papers.nips.cc/paper/9383-integer-discrete-flows-and-lossless-compression.pdf>.
- David A Huffman. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, 1952.
- Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of PSNR in image/video quality assessment. *Electronics letters*, 44(13):800–801, 2008.
- Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *International Conference on Learning Representations*, 2014. URL <https://arxiv.org/abs/1312.6114>.
- Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems*, pp. 4743–4751, 2016. URL <https://papers.nips.cc/paper/6581-improved-variational-inference-with-inverse-autoregressive-flow.pdf>.
- Friso Kingma, Pieter Abbeel, and Jonathan Ho. Bit-swap: Recursive bits-back coding for lossless compression with hierarchical latent variables. In *International Conference on Machine Learning*, pp. 3408–3417, 2019. URL <http://proceedings.mlr.press/v97/kingma19a.html>.
- Terry Kuny. A digital dark ages? challenges in the preservation of electronic information. *IFLA (International Federation of Library Associations and Institutions) Council and General Conference*, 64, 1997.
- Jooyoung Lee, Seunghyun Cho, and Seung-Kwon Beack. Context-adaptive entropy model for end-to-end optimized image compression. *International Conference on Learning Representations*, 2019. URL <https://openreview.net/pdf?id=HyxKIiAqYQ>.

- James Lucas, George Tucker, Roger Grosse, and Mohammad Norouzi. Understanding posterior collapse in generative latent variable models. *Deep Generative Models for Highly Structured Data (ICLR Workshop)*, 2019. URL <https://openreview.net/references/pdf?id=S13tjzrs4>.
- David Minnen and Saurabh Singh. Channel-wise autoregressive entropy models for learned image compression. *International Conference on Image Processing*, 2020. URL <https://ieeexplore.ieee.org/iel7/9184803/9190635/09190935.pdf>.
- David Minnen, Johannes Ballé, and George D Toderici. Joint autoregressive and hierarchical priors for learned image compression. In *Advances in Neural Information Processing Systems*, pp. 10771–10780, 2018. URL <http://papers.nips.cc/paper/8275-joint-autoregressive-and-hierarchical-priors-for-learned-image-compression.pdf>.
- Claude E Shannon. A mathematical theory of communication. *Bell system technical journal*, 27(3): 379–423, 1948.
- Lucas Theis, Wenzhe Shi, Andrew Cunningham, and Ferenc Huszár. Lossy image compression with compressive autoencoders. *International Conference on Learning Representations*, 2017. URL <https://openreview.net/pdf?id=rJiNwv9gg>.
- James Townsend, Tom Bird, and David Barber. Practical lossless compression with latent variables using bits back coding. *International Conference on Learning Representations*, 2019. URL <https://openreview.net/pdf?id=ryE98iR5tm>.
- James Townsend, Thomas Bird, Julius Kunze, and David Barber. Hilloc: Lossless image compression with hierarchical latent variable models. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=r1lZgyBYwS>.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, Eero P Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004.
- Ian H Witten, Radford M Neal, and John G Cleary. Arithmetic coding for data compression. *Communications of the ACM*, 30(6):520–540, 1987.
- Yibo Yang, Robert Bamler, and Stephan Mandt. Variational Bayesian quantization. In *International Conference on Machine Learning*, 2020.