Compressing Images by Encoding Their Latent Representations with Relative Entropy Coding

UNIVERSITY OF CAMBRIDGE

NEURAL INFORMATION PROCESSING SYSTEMS

Gergely Flamich, Marton Havasi, José Miguel Hernández-Lobato
{gf332, mh740, jmh233}@cam.ac.uk

Introduction

1: We describe **Relative Entropy Coding (REC)**, a *stochastic* coding framework, and **index coding**, an REC coding scheme.
2: REC can be combined with Variational Auto-Encoders (VAEs) with **continuous latent distributions** to create efficient lossless and lossy compression codecs.

Relative Entropy Coding

Adapted from Havasi et al. [2019]. Given

1. images $\mathbf{x}$
2. generative model $p(\mathbf{x}, \mathbf{z})$
3. approximate posterior $q(\mathbf{z} \mid \mathbf{x})$

minimum information $T[\mathbf{x} : \mathbf{z}]$ required such that the decoder can sample $\mathbf{z} \sim q(\mathbf{z} \mid \mathbf{x})$ is bounded by

$$T[\mathbf{x} : \mathbf{z}] \leq \mathbb{I}[\mathbf{x} : \mathbf{z}] + 2\log(\mathbb{I}[\mathbf{x} : \mathbf{z}] + 1) + \mathcal{O}(1),$$

where $\mathbb{I}[\mathbf{x} : \mathbf{z}]$ denotes the mutual information between $\mathbf{x}$ and $\mathbf{z}$. We refer to any compression scheme that achieves this bound a **relative entropy coding (REC)** method.

Image Compression Architecture

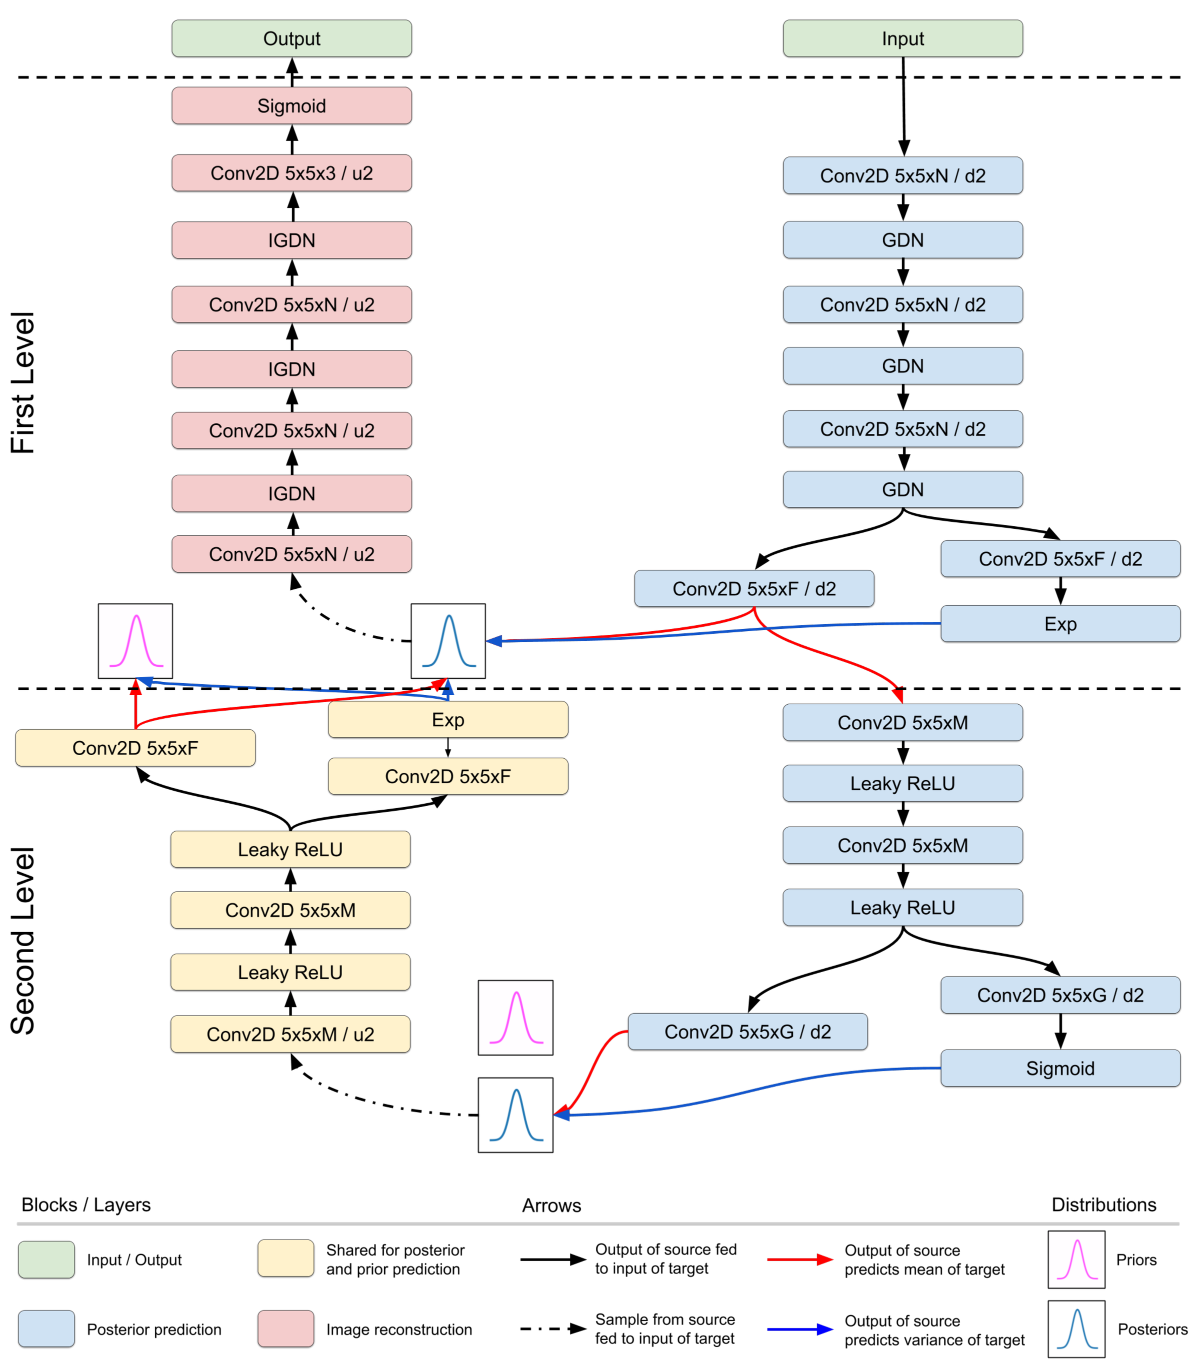

Index coding

Index coding is an approximate REC algorithm. Both parties generate the same sequence of samples from the coding distribution $p(\mathbf{z})$, by using the same random seed. Then the encoder communicates the sample in the sequence with the highest density ratio, by transmitting its index. The decoder can reconstruct the sample by indexing it's own sequence.

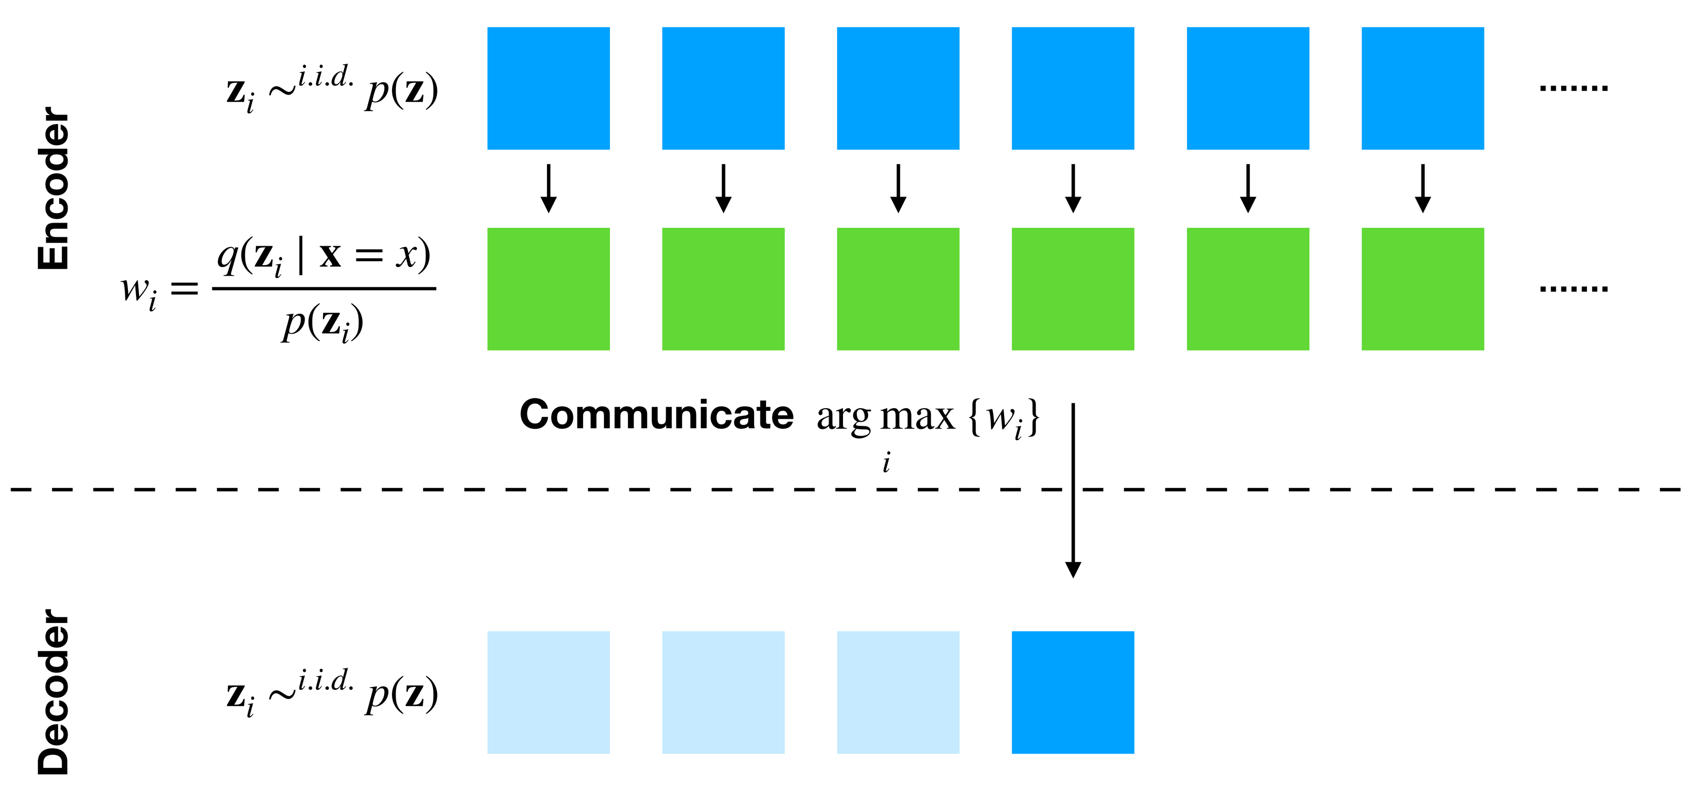

Figure 2: Schematic of index coding (iREC). We assume that $p(\mathbf{z})$ and the seed used to generate the samples are shared between the encoder and the decoder.

Lossless Image Compression

Images can be losslessly compressed by first coding a sample $\boldsymbol{z}$ from the latent posterior $p(\mathbf{z} \mid \boldsymbol{x})$ using REC, and then coding the original image $\boldsymbol{x}$ with the likelihood $p(\mathbf{x} \mid \boldsymbol{z})$. Our results using the ResNet VAE (RVAE) used by Townsend et al. [2020] on a few datasets are shown in Table 1.

		Cifar10 (32x32)	ImageNet32 (32x32)	Kodak (768x512)
Non bits-back	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	–
Bits-back	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)
REC	iREC (Ours)	**4.18**	**4.91**	**3.67**
	ELBO (RVAE)	[3.55]	[4.18]	[3.00]

Table 1: Single image, lossless compression performance in bits per dimension (lower is better). Best bits-back method shown in bold.

Lossy Image Compression

We apply REC to image compression using VAEs. Given a trained VAE, an image $\boldsymbol{x}$ can be lossily compressed by passing it through the encoder and using REC to code a sample from the latent posterior $q(\mathbf{z} \mid \boldsymbol{x})$ using the prior $p(\mathbf{z})$. Our results on the Kodak dataset are shown in Figure 3.

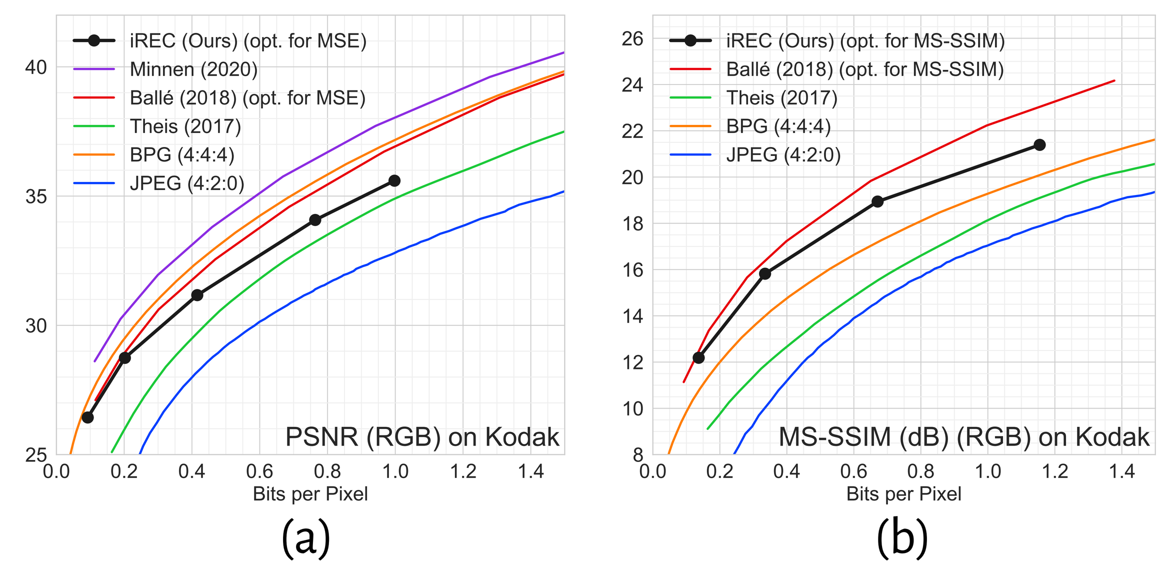

Figure 3: Comparison of REC against classical methods such as JPEG, BPG and competing ML-based methods. **(a)** PSNR comparisons **(b)** MS-SSIM comparisons in decibels, calculated using the formula $-10\log_{10}(1 - \text{MS-SSIM})$. See the supplementary material for more comparisons.

References

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`.

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`.

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`.