
Neural Joint Source-Channel Coding

Kristy Choi¹ Kedar Tatwawadi² Aditya Grover¹ Tsachy Weissman² Stefano Ermon¹

Abstract

For reliable transmission across a noisy communication channel, classical results from information theory show that it is asymptotically optimal to separate out the source and channel coding processes. However, this decomposition can fall short in the finite bit-length regime, as it requires non-trivial tuning of hand-crafted codes and assumes infinite computational power for decoding. In this work, we propose to jointly learn the encoding and decoding processes using a new discrete variational autoencoder model. By adding noise into the latent codes to simulate the channel during training, we learn to both compress and error-correct given a fixed bit-length and computational budget. We obtain codes that are not only competitive against several separation schemes, but also learn useful robust representations of the data for downstream tasks such as classification. Finally, inference amortization yields an extremely fast neural decoder, almost an order of magnitude faster compared to standard decoding methods based on iterative belief propagation.

1. Introduction

We consider the problem of encoding images as bit-strings so that they can be reliably transmitted across a noisy communication channel. Classical results from Shannon (1948) show that for a memoryless communication channel, as the image size goes to infinity, it is optimal to separately: 1) compress the images as much as possible to remove redundant information (source coding) and 2) use an error-correcting code to re-introduce redundancy, which allows for reconstruction in the presence of noise (channel coding). This *separation theorem* has been studied in a wide variety of contexts and has also enjoyed great success in the development of new coding algorithms.

¹Department of Computer Science, Stanford University

²Department of Electrical Engineering, Stanford University. Correspondence to: Kristy Choi <kechoi@cs.stanford.edu>.

However, this elegant decomposition suffers from two critical limitations in the finite bit-length regime. First, without an infinite number of bits for transmission, the overall distortion (reconstruction quality) is a function of both source and channel coding errors. Thus optimizing for both: (1) the number of bits to allocate to each process and (2) the code designs themselves is a difficult task. Second, maximum-likelihood decoding is in general NP-hard (Berlekamp et al. (1978)). Thus obtaining the accuracy that is possible *in theory* is contingent on the availability of infinite computational power for decoding, which is impractical for real-world systems. Though several lines of work have explored various relaxations of this problem to achieve tractability (Koetter & Vontobel (2003), Feldman et al. (2005), Vontobel & Koetter (2007)), many decoding systems rely on heuristics and early stopping of iterative approaches for scalability, which can be highly suboptimal.

To address these challenges we propose Neural Error Correcting and Source Trimming (NECST) codes, a deep learning framework for jointly learning to compress and error-correct an input image given a fixed bit-length budget. Three key steps are required. First, we use neural networks to encode each image into a suitable bit-string representation, sidestepping the need to rely on hand-designed coding schemes that require additional tuning for good performance. Second, we simulate a discrete channel *within* the model and inject noise directly into the latent codes to enforce robustness. Third, we amortize the decoding process such that after training, we obtain an extremely fast decoder at test time. These components pose significant optimization challenges due to the inherent non-differentiability of discrete latent random variables. We overcome this issue by leveraging recent advances in unbiased low-variance gradient estimation for variational learning of discrete latent variable models, and train the model using a variational lower bound on the mutual information between the images and their binary representations to obtain good, *robust* codes. At its core, NECST can also be seen as an implicit deep generative model of the data derived from the perspective of the joint-source channel coding problem (Bengio et al. (2014), Sohl-Dickstein et al. (2015), Goyal et al. (2017)).

In experiments, we test NECST on several grayscale and RGB image datasets, obtaining improvements over industry-standard compression (e.g, WebP (Google, 2015)) and error-

correcting codes (e.g., low density parity check codes). We also learn an extremely fast neural decoder, yielding almost an order of magnitude (two orders for magnitude for GPU) in speedup compared to standard decoding methods based on iterative belief propagation (Fossorier et al. (1999), Chen & Fossorier (2002)). Additionally, we show that in the process we learn discrete representations of the input images that are useful for downstream tasks such as classification.

In summary, our contributions are: 1) a novel framework for the JSCC problem and its connection to generative modeling with discrete latent variables; 2) a method for robust representation learning with latent additive discrete noise.




Figure 1. Illustration of NECST. An image is encoded into a binary bit-string, where it is corrupted by a discrete channel that is simulated *within* the model. Injecting noise into the latent space encourages NECST to learn robust representations and also turns it into an implicit generative model of the data distribution.

2. Source and Channel Coding

2.1. Preliminaries

We consider the problem of reliably communicating data across a noisy channel using a system that can detect and correct errors. Let X be a space of possible inputs (e.g., images), and $p_{\text{data}}(x)$ be a *source* distribution defined over $x \in X$. The goal of the communication system is to encode samples $x \sim p_{\text{data}}(x)$ as messages in a codeword space $\hat{Y} = \{0, 1\}^m$. The messages are transmitted over a noisy channel, where they are corrupted to become *noisy* codewords in Y . Finally, a decoder produces a reconstruction $\hat{x} \in X$ of the original input x from a received *noisy* code y .

The goal is to minimize the overall *distortion* (reconstruction error) $\|x - \hat{x}\|$ in ℓ_1 or ℓ_2 norm while keeping the message length m short (*rate*). In the absence of channel noise, low reconstruction errors can be achieved using a compression scheme to encode inputs $x \sim p_{\text{data}}(x)$ as succinctly as possible (e.g., WebP/JPEG). In the presence of noise, longer messages are typically needed to redundantly encode the information and recover from errors (e.g., parity check bits).

2.2. The Joint Source-Channel Coding Problem

Based on Shannon (1948)’s fundamental results, existing systems adhere to the following pipeline. A *source en-*

coder compresses the source image into a bit-string with the minimum number of bits possible. A *channel encoder* re-introduces redundancies, part of which were lost during source coding to prepare the code \hat{y} for transmission. The *decoder* then leverages the channel code to infer the original signal \hat{y} , producing an approximate reconstruction \hat{x} .

In the separation theorem, Shannon proved that the above scheme is optimal in the limit of infinitely long messages. That is, we can minimize distortion by optimizing the source and channel coding processes independently. However, given a finite bit-length budget, the relationship between rate and distortion becomes more complex. The more bits that we reserve for compression, the fewer bits we have remaining to construct the best error-correcting code and vice versa. Balancing the two sources of distortion through the optimal bit allocation, in addition to designing the best source and channel codes, makes this *joint source-channel coding* (JSCC) problem challenging.

In addition to these design challenges, real-world communication systems face computational and memory constraints that hinder the straightforward application of Shannon’s results. Specifically, practical decoding algorithms rely on approximations that yield suboptimal reconstruction accuracies. The information theory community has studied this problem extensively, and proposed different bounds for finite bit-length JSCC in a wide variety of contexts (Pile (1967), Csiszar (1982), Kostina & Verdú (2013)). Additionally, the issue of packet loss and queuing delay in wireless video communications suggest that modeling the forward-error correction (FEC) process may prove beneficial in improving communication systems over digital packet networks (Zhai & Katsaggelos (2007), Fouladi et al. (2018)). Thus rather than relying on hand-crafted coding schemes that may require additional tuning in practice, we propose to *learn* the appropriate bit allocations and coding mechanisms using a flexible deep learning framework.

3. Neural Source and Channel Codes

Given the above considerations, we explore a learning-based approach to the JSCC problem. To do so, we consider a flexible *class* of codes parameterized by a neural network and jointly optimize for the encoding and decoding procedure. This approach is inspired by recent successes in training (discrete) latent variable generative models (Neal (1992), Rolfe (2016), van den Oord et al. (2017)) as well as discrete representation learning in general (Hu et al. (2017), Kaiser & Bengio (2018), Roy et al. (2018)). We also explore the connection to different types of autoencoders in section 4.1.

3.1. Coding Process

Let X, \hat{Y}, Y, \hat{X} be random variables denoting the inputs, codewords, noisy codewords, and reconstructions respectively. We model their joint distribution $p(x, \hat{y}, y, \hat{x})$ using the following graphical model $X \rightarrow \hat{Y} \rightarrow Y \rightarrow \hat{X}$ as:

$$p(x, \hat{y}, y, \hat{x}) = p_{\text{data}}(x) p_{\text{enc}}(\hat{y}|x; \phi) p_{\text{channel}}(y|\hat{y}; \epsilon) p_{\text{dec}}(\hat{x}|y; \theta) \quad (1)$$

In Eq. 1, $p_{\text{data}}(x)$ denotes the distribution over inputs X . It is not known explicitly in practice, and only accessible through samples. $p_{\text{channel}}(y|\hat{y}; \epsilon)$ is the channel model, where we specifically focus on two widely used discrete channel models: (1) the binary erasure channel (BEC); and (2) the binary symmetric channel (BSC).

The BEC erases each bit \hat{y}_i into a corrupted symbol $?$ (e.g., $0 \rightarrow ?$) with some i.i.d probability ϵ , but faithfully transmits the correct bit otherwise. Therefore, \hat{Y} takes values in $\hat{Y} = \{0, 1, ?\}^m$, and p_{BEC} can be described by the following transition matrix for each bit:

$$P_{\text{BEC}} = \begin{bmatrix} 1 & \epsilon & 0 \\ 0 & 1 & \epsilon \\ \epsilon & \epsilon & \epsilon \end{bmatrix}$$

where the first column denotes the transition dynamics for $y_i = 0$ and the second column is that of $y_i = 1$. Each element per column denotes the probability of the bit remaining the same, flipping to a different bit, or being erased.

The BSC, on the other hand, independently flips each bit in the codeword with probability ϵ (e.g., $0 \rightarrow 1$). Therefore, \hat{Y} takes values in $\hat{Y} = Y = \{0, 1\}^m$ and

$$p_{\text{BSC}}(y|\hat{y}; \epsilon) = \prod_{i=1}^m \epsilon^{y_i \oplus \hat{y}_i} (1 - \epsilon)^{y_i \oplus \hat{y}_i \oplus 1}$$

where \oplus denotes addition modulo 2 (i.e., an eXclusive OR). All experiments in this paper are conducted with the BSC channel. We note that the BSC is a more difficult communication channel to work with than the BEC (Richardson & Urbanke (2008)).

A *stochastic* encoder $q_{\text{enc}}(\hat{y}|x; \phi)$ generates a codeword \hat{y} given an input x . Specifically, we model each bit \hat{y}_i in the code with an independent Bernoulli random vector. We model the parameters of this Bernoulli with a neural network $f_\phi(\cdot)$ (an MLP or a CNN) parameterized by ϕ . Denoting $\sigma(z)$ as the sigmoid function:

$$q_{\text{enc}}(\hat{y}|x, \phi) = \prod_{i=1}^m \sigma(f_\phi(x_i))^{\hat{y}_i} (1 - \sigma(f_\phi(x_i)))^{(1 - \hat{y}_i)}$$

Similarly, we posit a probabilistic decoder $p_{\text{dec}}(\hat{x}|y; \theta)$ parameterized by an MLP/CNN that, given y , generates a

decoded image \hat{x} . For real-valued images, we model each pixel as a factorized Gaussian with a fixed, isotropic covariance to yield: $\hat{x}|y \sim \mathcal{N}(f_\theta(y), \sigma^2 I)$, where $f_\theta(\cdot)$ denotes the decoder network. For binary images, we use a Bernoulli decoder: $\hat{x}|y \sim \text{Bern}(f_\theta(y))$.

4. Variational Learning of Neural Codes

To learn an effective coding scheme, we **maximize the mutual information** between the input X and the corresponding *noisy* codeword Y (Barber & Agakov (2006)). That is, the code \hat{y} should be robust to partial corruption; even its noisy instantiation y should preserve as much information about the original input x as possible (MacKay (2003)).

First, we note that we can analytically compute the encoding distribution $q_{\text{noisy_enc}}(y|x; \epsilon, \phi)$ *after* it has been perturbed by the channel by marginalizing over \hat{y} :

$$q_{\text{noisy_enc}}(y|x; \epsilon, \phi) = \sum_{\hat{y} \in \mathcal{Z}^m} q_{\text{enc}}(\hat{y}|x; \phi) p_{\text{channel}}(y|\hat{y}; \epsilon)$$

The BEC induces a 3-way Categorical distribution for $q_{\text{noisy_enc}}$ where the third category refers to bit-erasure:

$$y|x \sim \text{Cat}(1 - \epsilon - \sigma(f_\phi(x)) + \sigma(f_\phi(x)), \sigma(f_\phi(x)), \sigma(f_\phi(x)), \epsilon, \epsilon)$$

while for the BSC, it becomes:

$$q_{\text{noisy_enc}}(y|x; \phi, \epsilon) = \prod_{i=1}^m (\sigma(f_\phi(x_i)) - 2\sigma(f_\phi(x_i))\epsilon + \epsilon)^{y_i} (1 - \sigma(f_\phi(x_i)) + 2\sigma(f_\phi(x_i))\epsilon - \epsilon)^{(1 - y_i)}$$

We note that although we have outlined two specific channel models, it is relatively straightforward for our framework to handle other types of channel noise.

Finally, we get the following optimization problem:

$$\begin{aligned} \max_{\phi} I(X, Y; \phi, \epsilon) &= H(X) - H(X|Y; \phi, \epsilon) \\ &= \mathbb{E}_x p_{\text{data}}(x) \mathbb{E}_y q_{\text{noisy_enc}}(y|x; \epsilon, \phi) [\log p(x|y; \epsilon, \phi)] + \text{const.} \\ &\quad \mathbb{E}_x p_{\text{data}}(x) \mathbb{E}_y q_{\text{noisy_enc}}(y|x; \epsilon, \phi) [\log p_{\text{dec}}(x|y; \theta)] + \text{const.} \end{aligned} \quad (2)$$

where $p(x|y; \epsilon, \phi)$ is the true (and intractable) posterior from Eq. 1 and $p_{\text{dec}}(x|y; \theta)$ is an amortized variational approximation. The true posterior $p(x|y; \epsilon, \phi)$ — the posterior probability over possible inputs x given the received noisy codeword y — is the best possible decoder. However, it is also often intractable to evaluate and optimize. We therefore use a *tractable* variational approximation $p_{\text{dec}}(\hat{x}|y; \theta)$. Crucially, this variational approximation is *amortized* (Kingma & Welling (2013)), and is the inference distribution that will

actually be used for decoding at test time. Because of amortization, *decoding is guaranteed to be efficient*, in contrast with existing error correcting codes, which typically involve NP-hard MPE inference queries.

Given any encoder (ϕ), we can find the best amortized variational approximation by maximizing the lower bound (Eq. 2) as a function of θ . Therefore, the NECST training objective is given by:

$$\max_{\theta, \phi} \mathbb{E}_x \mathbb{E}_{p_{\text{data}}(x)} \mathbb{E}_y \mathbb{E}_{q_{\text{noisy_enc}}(y|x; \epsilon, \phi)} [\log p_{\text{dec}}(x|y; \theta)]$$

In practice, we approximate the expectation of the data distribution $p_{\text{data}}(x)$ with a finite dataset D :

$$\max_{\theta, \phi} \sum_{x \in D} \mathbb{E}_y \mathbb{E}_{q_{\text{noisy_enc}}(y|x; \epsilon, \phi)} [\log p_{\text{dec}}(x|y; \theta)] \quad L(\phi, \theta; x, \epsilon) \quad (3)$$

Thus we can jointly learn the encoding and decoding scheme by optimizing the parameters ϕ, θ . In this way, the encoder (ϕ) is “aware” of the computational limitations of the decoder (θ), and is optimized accordingly (Shu et al., 2018). However, a main learning challenge is that we are not able to backpropagate directly through the *discrete* latent variable y ; we elaborate upon this point further in Section 5.1.

For the BSC, the solution will depend on the number of available bits m , the noise level ϵ , and the structure of the input data x . For intuition, consider the case $m = 0$. In this case, no information can be transmitted, and the best decoder will fit a single Gaussian to the data. When $m = 1$ and $\epsilon = 0$ (noiseless case), the model will learn a mixture of $2^m = 2$ Gaussians. However, if $m = 1$ and $\epsilon = 0.5$, again no information can be transmitted, and the best decoder will fit a single Gaussian to the data. Adding noise forces the model to decide how it should effectively partition the data such that (1) similar items are grouped together in the same cluster; and (2) the clusters are “well-separated” (even when a few bits are flipped, they do not “cross over” to a different cluster that will be decoded incorrectly). Thus, from the perspective of unsupervised learning, NECST attempts to learn *robust binary representations* of the data.

4.1. NECST as a Generative Model

The objective function in Eq. 3 closely resembles those commonly used in generative modeling frameworks; in fact, NECST can also be viewed as a generative model. In its simplest form, our model with a noiseless channel (i.e., $\epsilon = 0$), deterministic encoder, and deterministic decoder is identical to a traditional autoencoder (Bengio et al. (2007)). Once channel noise is present ($\epsilon > 0$) and the encoder/decoder become probabilistic, NECST begins to more closely resemble other variational autoencoding frameworks. Specifically, NECST is similar to Denoising Autoencoders (DAE) (Vincent et al. (2008)), except that it is explicitly trained for

robustness to partial destruction of *the latent space*, as opposed to the *input space*. The stacked DAE (Vincent et al., 2010) also denoises the latent representations, but does not explicitly inject noise into each layer.

NECST is also related to the VAE, with two nuanced distinctions. While both models learn a joint distribution over the observations and latent codes, NECST: (1) optimizes a variational lower bound to the *mutual information* $I(X, Y)$ as opposed to the marginal log-likelihood $p(X)$, and (2) *does not posit* a prior distribution $p(Y)$ over the latent variables. Their close relationship is evidenced by a line of work on rate-distortion optimization in the context of VAEs (Ballé et al. (2016), Higgins et al. (2016), Zhao et al. (2017), Alemi et al. (2017), Zhao et al. (2018)), as well as other information-theoretic interpretations of the VAE’s information preference (Hinton & Van Camp (1993), Honkela & Valpola (2004), Chen et al. (2016b)). Although NECST’s method of representation learning is also related to that of the Deep Variational Information Bottleneck (VIB) (Alemi et al., 2016), VIB *minimizes* the MI between the observations and latents while maximizing the MI between the observations and targets in a supervised setting. Also, existing autoencoders are aimed at compression, while for sufficiently large m NECST will attempt to learn lossless compression with added redundancy for error correction.

Finally, NECST can be seen as a discrete version of the Uncertainty Autoencoder (UAE) (Grover & Ermon (2019)). The two models share identical objective functions with two notable differences: For NECST, (1) the latent codes Y are discrete random variables, and (2) the noise model is no longer continuous. The special properties of the continuous UAE carry over directly to its discrete counterpart. Grover & Ermon (2019) proved that under certain conditions, the UAE specifies an implicit generative model (Diggle & Gratton (1984), Mohamed & Lakshminarayanan (2016)). We restate their major theorem and extend their results here.

Starting from any data point $x^{(0)} \in X$, we define a Markov chain over $X \times Y$ with the following transitions:

$$y^{(t)} \sim q_{\text{noisy_enc}}(y|x^{(t)}; \phi, \epsilon), \quad x^{(t+1)} \sim p_{\text{dec}}(x|y^{(t)}; \theta) \quad (4)$$

Theorem 1. *For a fixed value of $\epsilon > 0$, let θ, ϕ denote an optimal solution to the NECST objective. If there exists a ϕ such that $q_{\phi}(x|y; \epsilon) = p_{\theta}(x|y)$, then the Markov Chain (Eq. 4) with parameters ϕ and θ is ergodic and its stationary distribution is given by $p_{\text{data}}(x)q_{\text{noisy_enc}}(y|x; \epsilon, \phi)$.*

Proof. We follow the proof structure in Grover & Ermon (2019) with minor adaptations. First, we note that the Markov chain defined in Eq. 4 is ergodic due to the BSC noise model. This is because the BSC defines a distribution over all possible (finite) configurations of the latent




Figure 2. For each dataset, we fix the number of bits m for NECST and compute the resulting distortion. Then, we plot the *additional* number of bits the WebP + ideal channel code system would need in theory to match NECST’s distortion at different levels of channel noise. NECST’s ability to jointly source and channel code yields much greater bitlength efficiency as compared to the separation scheme baseline, with the discrepancy becoming more dramatic as we increase the level of channel noise.

code Y , and the Gaussian decoder posits a non-negative probability of transitioning to all possible reconstructions X . Thus given any (X, Y) and (X^0, Y^0) such that the density $p(X, Y) > 0$ and $p(X^0, Y^0) > 0$, the probability density of transitioning $p(X^0, Y^0 | X, Y) > 0$.

Next, we can rewrite the objective in Eq. 3 as the following:

$$\begin{aligned} & \mathbb{E}_{p(x,y;\epsilon,\phi)} [\log p_{\text{dec}}(x|y; \theta)] = \\ & \mathbb{E}_{q(y;\epsilon,\phi)} \left[\int q(x|y; \phi) \log p_{\text{dec}}(x|y; \theta) \right] = \\ & H(X|Y; \phi) - \mathbb{E}_{q(y;\epsilon,\phi)} [KL(q(X|y; \phi) || p_{\text{dec}}(X|y; \theta))] \end{aligned}$$

Since the KL divergence is minimized when the two argument distributions are equal, we note that for the optimal value of $\theta = \theta^*$, if there exists a ϕ in the space of encoders being optimized that satisfies $q_{\phi}(x|y; \epsilon) = p_{\theta^*}(x|y)$ then $\phi = \phi^*$. Then, for any ϕ , we note that the following Gibbs chain converges to $p(x, y)$ if the chain is ergodic:

$$y^{(t)} \sim q_{\text{noisy_enc}}(y|x^{(t)}; \phi, \epsilon), \quad x^{(t+1)} \sim p(x|y^{(t)}; \theta) \quad (5)$$

Substituting $p(x|y; \theta^*)$ for $p_{\text{dec}}(x|y; \theta)$ finishes the proof. \square

Hence NECST has an intractable likelihood but can generate samples from p_{data} by running the Markov Chain above. Samples obtained by running the chain for a trained model, initialized from both random noise and test data, can be found in the Supplementary Material. Our results indicate that although we do not meet the theorem’s conditions in practice, we are still able to learn a reasonable model.

5. Experimental Results

We first review the optimization challenges of training discrete latent variable models and elaborate on our training procedure. Then to validate our work, we first assess NECST’s compression and error correction capabilities against a combination of two widely-used compression

(WebP, VAE) and channel coding (ideal code, LDPC (Galager (1962))) algorithms. We experiment on randomly generated length-100 bitstrings, MNIST (LeCun (1998)), Omniglot (Lake et al. (2015)), Street View Housing Numbers (SVHN) (Netzer et al. (2011)), CIFAR10 (Krizhevsky (2009)), and CelebA (Liu et al. (2015)) datasets to account for different image sizes and colors. Next, we test the decoding speed of NECST’s decoder and find that it performs upwards of *an order of magnitude faster* than standard decoding algorithms based on iterative belief propagation (and two orders of magnitude on GPU). Finally, we assess the quality of the latent codes after training, and examine interpolations in latent space as well as how well the learned features perform for downstream classification tasks.

5.1. Optimization Procedure

Recall the NECST objective in equation 3. While obtaining Monte Carlo-based gradient estimates with respect to θ is easy, gradient estimation with respect to ϕ is challenging because these parameters specify the Bernoulli random variable $q_{\text{noisy_enc}}(y|x; \epsilon, \phi)$. The commonly used *reparameterization trick* cannot be applied in this setting, as the discrete stochastic unit in the computation graph renders the overall network non-differentiable (Schulman et al. (2015)).

An alternative is to use the *score function estimator* in place of the gradient, as defined in the REINFORCE algorithm (Williams (1992)). However, this estimator suffers from high variance, and several others have explored different formulations and control variates to mitigate this issue (Wang et al. (2013), Gu et al. (2015), Ruiz et al. (2016), Tucker et al. (2017), Grathwohl et al. (2017), Grover et al. (2018)). Others have proposed a continuous relaxation of the discrete random variables, as in the Gumbel-softmax (Jang et al. (2016)) and Concrete (Maddison et al. (2016)) distributions.

Empirically, we found that using that using a continuous relaxation of the discrete latent features led to worse perfor-

	0.0	0.1	0.2	0.3	0.4	0.5
Binary MNIST						
100-bit VAE+LDPC	0.047	0.064	0.094	0.120	0.144	0.150
100-bit NECST	0.039	0.052	0.066	0.091	0.125	0.131
Binary Omniglot						
200-bit VAE+LDPC	0.048	0.058	0.076	0.092	0.106	0.100
200-bit NECST	0.035	0.045	0.057	0.070	0.080	0.080
Random bits						
50-bit VAE+LDPC	0.352	0.375	0.411	0.442	0.470	0.502
50-bit NECST	0.291	0.357	0.407	0.456	0.500	0.498




Figure 3. (Left) Reconstruction error of NECST vs. VAE + rate-1/2 LDPC codes. We observe that NECST outperforms the baseline in almost all settings except for the random bits, which is expected due to its inability to leverage any underlying statistical structure. (Right) Average decoding times for NECST vs. 50 iterations of LDPC decoding on Omniglot. One forward pass of NECST provides 2x orders of magnitude in speedup on GPU for decoding time when compared to a traditional algorithm based on iterative belief propagation.

mance at test time when the codes were forced to be discrete. Therefore, we used VIMCO (Mnih & Rezende (2016)), a multi-sample variational lower bound objective for obtaining low-variance gradients. VIMCO constructs leave-one-out control variates using its samples, as opposed to the single-sample objective NVIL (Mnih & Gregor (2014)) which requires learning additional baselines during training. Thus, we used the 5-sample VIMCO objective in subsequent experiments for the optimization procedure, leading us to our final multi-sample ($K = 5$) objective:

$$L^K(\phi, \theta; x, \epsilon) = \max_{\theta, \phi} \sum_{x \in D} E_{y^{1:K} \sim q_{\text{noisy-enc}}(y|x; \epsilon, \phi)} \left[\log \frac{1}{K} \sum_{i=1}^K p_{\text{dec}}(x_j y^i; \theta) \right] \quad (6)$$

5.2. Fixed distortion: WebP + Ideal channel code

In this experiment, we compare the performances of: (1) NECST and (2) WebP + ideal channel code in terms of compression on 3 RGB datasets: CIFAR-10, CelebA, and binarized SVHN. We note that before the comparing our model with WebP, we remove the headers from the compressed images for a fair comparison.

Specifically, we fix the number of bits m used by NECST to source and channel code, and obtain the corresponding distortion levels (reconstruction errors) at various noise levels ϵ . For fixed m , distortion will increase with ϵ . Then, for each noise level we estimate the number of bits an alternate system using WebP and an ideal channel code — the best that is theoretically possible — would have to use to match NECST’s distortion levels. As WebP is a variable-rate compression algorithm, we compute this quantity for each image, then average across the entire dataset. Assuming an ideal channel code implies that all messages will be transmitted across the noisy channel at the highest possible communication rate (i.e. the channel capacity); thus the re-

sulting distortion will only be a function of the compression mechanism. We use the well-known channel capacity formula for the BSC channel $C = 1 - H_b(\epsilon)$ to compute this estimate, where $H_b(\cdot)$ denotes the binary entropy function.

We find that NECST excels at compression. Figure 1 shows that in order to achieve the same level of distortion as NECST, the WebP-ideal channel code system requires a much greater number bits *across all noise levels*. In particular, we note that NECST is slightly more effective than WebP at pure compression ($\epsilon = 0$), and becomes significantly better at higher noise levels (e.g., for $\epsilon = 0.4$, NECST requires 20 less bits).

5.3. Fixed Rate: source coding + LDPC

Next, we compare the performances of: (1) NECST and (2) discrete VAE + LDPC in terms of distortion. Specifically, we fix the bit-length budget for both systems and evaluate whether learning compression and error-correction jointly (NECST) improves reconstruction errors (distortion) for the same rate. We do not report a comparison against JPEG + LDPC as we were often unable to obtain valid JPEG files to compute reconstruction errors after imperfect LDPC decoding. However, as we have seen in the JPEG/WebP + ideal channel code experiment, we can reasonably conclude that the JPEG/WebP-LDPC system would have displayed worse performance, as NECST was superior to a setup with the best channel code theoretically possible.

We experiment with a discrete VAE with a uniform prior over the latent codes for encoding on binarized MNIST and Omniglot, as these datasets are settings in which VAEs are known to perform favorably. We also include the random bits dataset to assess our channel coding capabilities, as the dataset lacks any statistical structure that can be exploited at source coding time. To fix the total number of bits that are transmitted through the channel, we double the length of the VAE encodings with a rate-1/2 LDPC channel code


(a) MNIST: 6 \rightarrow 3 leads to an intermediate 5(b) MNIST: 8 \rightarrow 4 leads to an intermediate 2 and 9

Figure 4. Latent space interpolation, where each image is obtained by randomly flipping one additional latent bit and passing through the decoder. We observe that the NECST model has indeed learned to encode redundancies in the compressed representations of the images, as it takes a significant number of bit flips to alter the original identity of the digit.

to match NECST’s latent code dimension m . We found that fixing the number of LDPC bits and varying the number of bits in the VAE within a reasonable interval did not result in a significant difference. We note from Figure 2a that although the resulting distortion is similar, NECST outperforms the VAE across all noise levels for image datasets. These results suggest that NECST’s learned coding scheme performs at least as well as LDPCs, an industrial-strength, widely deployed class of error correcting codes.

5.4. Decoding Time

Another advantage of NECST over traditional channel codes is that after training, the amortized decoder can very efficiently map the transmitted code into its best reconstruction at test time. Other sparse graph-based coding schemes such as LDPC, however, are more time-consuming because decoding involves an NP-hard optimization problem typically approximated with multiple iterations of belief propagation.

To compare the speed of our neural decoder to LDPC’s belief propagation, we fixed the number of bits that were transmitted across the channel and averaged the total amount of time taken for decoding across ten runs. The results are shown below in Figure 2b for the statically binarized Omniglot dataset. On CPU, NECST’s decoder averages an order of magnitude faster than the LDPC decoder running for 50 iterations, which is based on an optimized C implementation. On GPU, NECST displays *two orders of magnitude in speedup*. We also conducted the same experiment without batching, and found that NECST still remains an order of magnitude faster than the LDPC decoder.

6. Robust Representation Learning

In addition to its ability to source and channel code effectively, NECST is also an implicit generative model that yields robust and interpretable latent representations and realistic samples from the underlying data distribution. Specifically, in light of Theorem 1, we can think of $q_{\text{enc}}(\hat{y}|x; \phi)$ as mapping images x to latent representations \hat{y} .

6.1. Latent space interpolation

To assess whether the model has: (1) injected redundancies into the learned codes and (2) learned interesting features, we interpolate between different data points in latent space and qualitatively observe whether the model captures semantically meaningful variations in the data. We select two test points to be the start and end, sequentially flip one bit at a time in the latent code, and pass the altered code through the decoder to observe how the reconstruction changes. In Figure 3, we show two illustrative examples from the MNIST digits. From the starting digit, each bit-flip slowly alters characteristic features such as rotation, thickness, and stroke style until the digit is reconstructed to something else entirely. We note that because NECST is trained to encode redundancies, we do not observe a drastic change per flip. Rather, it takes a significant level of corruption for the decoder to interpret the original digit as another. Also, due to the i.i.d. nature of the channel noise model, we observe that the sequence at which the bits are flipped do not have a significant effect on the reconstructions. Additional interpolations may be found in the Supplementary Material.

6.2. Downstream classification

To demonstrate that the latent representations learned by NECST are useful for downstream classification tasks, we extract the binary features and train eight different classification algorithms: k -nearest neighbors (KNN), decision trees (DT), random forests (RF), multilayer perceptron (MLP), AdaBoost (AdaB), Quadratic Discriminant Analysis (QDA), and support vector machines (SVM). As shown on the left-most numbers in Table 1, the simple classifiers perform reasonably well in the digit classification task for MNIST ($\epsilon = 0$). We observe that with the exception of KNN, all other classifiers achieve higher levels of accuracy when using features that were trained with simulated channel noise ($\epsilon = 0.1, 0.2$).

To further test the hypothesis that training with simulated channel noise yields *more robust* latent features, we freeze the pre-trained NECST model and evaluate classification accuracy using a “noisy” MNIST dataset. We synthesized

Table 1. Classification accuracy on MNIST/noisy MNIST using 100-bit features from NECST.

Noise ϵ	KNN	DT	RF	MLP	AdaB	NB	QDA	SVM
0	0.95/0.86	0.65/0.54	0.71/0.59	0.93/0.87	0.72/0.65	0.75/0.65	0.56/0.28	0.88/0.81
0.1	0.95/0.86	0.65/0.59	0.74/0.65	0.934/0.88	0.74/0.72	0.83/0.77	0.94/0.90	0.92/0.84
0.2	0.94/0.86	0.78/0.69	0.81/0.76	0.93/0.89	0.78/0.80	0.87/0.81	0.93/0.90	0.93/0.86

noisy MNIST by adding $\epsilon \sim N(0, 0.5)$ noise to all pixel values in MNIST, and ran the same experiment as above. As shown in the rightmost numbers of Table 1, we again observe that most algorithms show improved performance with added channel noise. Intuitively, when the latent codes are corrupted by the channel, the codes will be "better separated" in latent space such that the model will still be able to reconstruct accurately despite the added noise. Thus NECST can be seen as a "denoising autoencoder"-style method for learning more robust latent features, with the twist that the noise is injected into the latent space.

7. Related work

There has been a recent surge of work applying deep learning and generative modeling techniques to lossy image compression, many of which compare favorably against industry standards such as JPEG, JPEG2000, and WebP (Toderici et al. (2015), Ballé et al. (2016), Toderici et al. (2017)). Theis et al. (2017) use compressive autoencoders that learn the optimal number of bits to represent images based on their pixel frequencies. Ballé et al. (2018) use a variational autoencoder (VAE) (Kingma & Welling (2013) Rezende et al. (2014)) with a learnable scale hyperprior to capture the image's partition structure as side information for more efficient compression. Santurkar et al. (2017) use adversarial training (Goodfellow et al. (2014)) to learn neural codecs for compressing images and videos using DCGAN-style ConvNets (Radford et al. (2015)). Yet these methods focus on source coding only, and do not consider the setting where the compression must be robust to channel noise.

In a similar vein, there has been growing interest on leveraging these deep learning systems to sidestep the use of hand-designed codes. Several lines of work train neural decoders based on known coding schemes, sometimes learning more general decoding algorithms than before (Nachmani et al. (2016), Gruber et al. (2017), Cammerer et al. (2017), Dorner et al. (2017)). (Kim et al. (2018a), Kim et al. (2018b)) parameterize sequential codes with recurrent neural network (RNN) architectures that achieve comparable performance to capacity-approaching codes over the additive white noise Gaussian (AWGN) and bursty noise channels. However, source coding is out of the scope for these methods that focus on learning good channel codes.

The problem of end-to-end transmission of structured data, on the other hand, is less well-studied. Zarcone et al. (2018)

utilize an autoencoding scheme for data storage. Farsad et al. (2018) use RNNs to communicate text over a BEC and are able to preserve the words' semantics. The most similar to our work is that of Bourtsoulatzé et al. (2018), who use autoencoders for transmitting images over the AWGN and slow Rayleigh fading channels, which are continuous. We provide a holistic treatment of various discrete noise models and show how NECST can also be used for unsupervised learning. While there is a rich body of work on using information maximization for representation learning (Chen et al. (2016a), Hu et al. (2017), Hjelm et al. (2018)), these methods do not incorporate the addition of discrete latent noise in their training procedure.

8. Conclusion

We described how NECST can be used to learn an efficient joint source-channel coding scheme by simulating a noisy channel during training. We showed that the model: (1) is competitive against a combination of WebP and LDPC codes on a wide range of datasets, (2) learns an extremely fast neural decoder through amortized inference, and (3) learns a latent code that is not only robust to corruption, but also useful for downstream tasks such as classification.

One limitation of NECST is the need to train the model separately for different code-lengths and datasets; in its current form, it can only handle fixed-length codes. Extending the model to streaming or adaptive learning scenarios that allow for learning variable-length codes is an exciting direction for future work. Another direction would be to analyze the characteristics and properties of the learned latent codes under different discrete channel models. We provide reference implementations in Tensorflow (Abadi et al., 2016), and the codebase for this work is open-sourced at <https://github.com/ermongroup/necst>.

Acknowledgements

We are thankful to Neal Jean, Daniel Levy, Rui Shu, and Jiaming Song for insightful discussions and feedback on early drafts. KC is supported by the NSF GRFP and Stanford Graduate Fellowship, and AG is supported by the MSR Ph.D. fellowship, Stanford Data Science scholarship, and Lieberman fellowship. This research was funded by NSF (#1651565, #1522054, #1733686), ONR (N00014-19-1-2145), AFOSR (FA9550-19-1-0024), and Amazon AWS.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (fOSDIg 16)*, pp. 265–283, 2016.
- Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.
- Alemi, A. A., Poole, B., Fischer, I., Dillon, J. V., Saurous, R. A., and Murphy, K. An information-theoretic analysis of deep latent-variable models. *arXiv preprint arXiv:1711.00464*, 2017.
- Ballé, J., Laparra, V., and Simoncelli, E. P. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704*, 2016.
- Ballé, J., Minnen, D., Singh, S., Hwang, S. J., and Johnston, N. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*, 2018.
- Barber, D. and Agakov, F. V. Kernelized infomax clustering. In *Advances in neural information processing systems*, pp. 17–24, 2006.
- Bengio, Y., LeCun, Y., et al. Scaling learning algorithms towards ai. 2007.
- Bengio, Y., Laufer, E., Alain, G., and Yosinski, J. Deep generative stochastic networks trainable by backprop. In *International Conference on Machine Learning*, pp. 226–234, 2014.
- Berlekamp, E., McEliece, R., and Van Tilborg, H. On the inherent intractability of certain coding problems (corresp.). *IEEE Transactions on Information Theory*, 24(3): 384–386, 1978.
- Bourtsoulatze, E., Kurka, D. B., and Gunduz, D. Deep joint source-channel coding for wireless image transmission. *arXiv preprint arXiv:1809.01733*, 2018.
- Cammerer, S., Gruber, T., Hoydis, J., and ten Brink, S. Scaling deep learning-based decoding of polar codes via partitioning. In *GLOBECOM 2017-2017 IEEE Global Communications Conference*, pp. 1–6. IEEE, 2017.
- Chen, J. and Fossorier, M. P. Near optimum universal belief propagation based decoding of low-density parity check codes. *IEEE Transactions on communications*, 50(3): 406–414, 2002.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, pp. 2172–2180, 2016a.
- Chen, X., Kingma, D. P., Salimans, T., Duan, Y., Dhariwal, P., Schulman, J., Sutskever, I., and Abbeel, P. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016b.
- Csiszar, I. Linear codes for sources and source networks: Error exponents, universal coding. *IEEE Transactions on Information Theory*, 28(4):585–592, 1982.
- Diggle, P. J. and Gratton, R. J. Monte carlo methods of inference for implicit statistical models. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 193–227, 1984.
- Dorner, S., Cammerer, S., Hoydis, J., and ten Brink, S. On deep learning-based communication over the air. In *Signals, Systems, and Computers, 2017 51st Asilomar Conference on*, pp. 1791–1795. IEEE, 2017.
- Farsad, N., Rao, M., and Goldsmith, A. Deep learning for joint source-channel coding of text. *arXiv preprint arXiv:1802.06832*, 2018.
- Feldman, J., Wainwright, M. J., and Karger, D. R. Using linear programming to decode binary linear codes. *IEEE Transactions on Information Theory*, 51(3):954–972, 2005.
- Fossorier, M. P., Mihaljevic, M., and Imai, H. Reduced complexity iterative decoding of low-density parity check codes based on belief propagation. *IEEE Transactions on communications*, 47(5):673–680, 1999.
- Fouladi, S., Emmons, J., Orbay, E., Wu, C., Wahby, R. S., and Winstein, K. Salsify: Low-latency network video through tighter integration between a video codec and a transport protocol. 2018.
- Gallager, R. Low-density parity-check codes. *IRE Transactions on information theory*, 8(1):21–28, 1962.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Google. Webp compression study. https://developers.google.com/speed/webp/docs/webp_study, 2015. [Online; accessed 22-January-2019].

- Goyal, A. G. A. P., Ke, N. R., Ganguli, S., and Bengio, Y. Variational walkback: Learning a transition operator as a stochastic recurrent net. In *Advances in Neural Information Processing Systems*, pp. 4392–4402, 2017.
- Grathwohl, W., Choi, D., Wu, Y., Roeder, G., and Duvenaud, D. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. *arXiv preprint arXiv:1711.00123*, 2017.
- Grover, A. and Ermon, S. Uncertainty autoencoders: Learning compressed representations via variational information maximization. In *International Conference on Artificial Intelligence and Statistics*, 2019.
- Grover, A., Gummadi, R., Lazaro-Gredilla, M., Schuurmans, D., and Ermon, S. Variational rejection sampling. In *Proc. 21st International Conference on Artificial Intelligence and Statistics*, 2018.
- Gruber, T., Cammerer, S., Hoydis, J., and ten Brink, S. On deep learning-based channel decoding. In *Information Sciences and Systems (CISS), 2017 51st Annual Conference on*, pp. 1–6. IEEE, 2017.
- Gu, S., Levine, S., Sutskever, I., and Mnih, A. Muprop: Unbiased backpropagation for stochastic neural networks. *arXiv preprint arXiv:1511.05176*, 2015.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.
- Hinton, G. E. and Van Camp, D. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pp. 5–13. ACM, 1993.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Trischler, A., and Bengio, Y. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*, 2018.
- Honkela, A. and Valpola, H. Variational learning and bits-back coding: an information-theoretic view to bayesian learning. *IEEE Transactions on Neural Networks*, 15(4): 800–810, 2004.
- Hu, W., Miyato, T., Tokui, S., Matsumoto, E., and Sugiyama, M. Learning discrete representations via information maximizing self-augmented training. *arXiv preprint arXiv:1702.08720*, 2017.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Kaiser, Ł. and Bengio, S. Discrete autoencoders for sequence models. *arXiv preprint arXiv:1801.09797*, 2018.
- Kim, H., Jiang, Y., Kannan, S., Oh, S., and Viswanath, P. Deepcode: Feedback codes via deep learning. *arXiv preprint arXiv:1807.00801*, 2018a.
- Kim, H., Jiang, Y., Rana, R., Kannan, S., Oh, S., and Viswanath, P. Communication algorithms via deep learning. *arXiv preprint arXiv:1805.09317*, 2018b.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Koetter, R. and Vontobel, P. O. Graph-covers and iterative decoding of finite length codes. Citeseer, 2003.
- Kostina, V. and Verdú, S. Lossy joint source-channel coding in the finite blocklength regime. *IEEE Transactions on Information Theory*, 59(5):2545–2575, 2013.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- LeCun, Y. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- MacKay, D. J. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Mnih, A. and Gregor, K. Neural variational inference and learning in belief networks. *arXiv preprint arXiv:1402.0030*, 2014.
- Mnih, A. and Rezende, D. J. Variational inference for monte carlo objectives. *arXiv preprint arXiv:1602.06725*, 2016.
- Mohamed, S. and Lakshminarayanan, B. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.
- Nachmani, E., Be’ery, Y., and Burshtein, D. Learning to decode linear codes using deep learning. In *Communication, Control, and Computing (Allerton), 2016 54th Annual Allerton Conference on*, pp. 341–346. IEEE, 2016.
- Neal, R. M. Connectionist learning of belief networks. *Artificial intelligence*, 56(1):71–113, 1992.

- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*. NIPS, 2011.
- Pilc, R. J. *Coding theorems for discrete source-channel pairs*. PhD thesis, Massachusetts Institute of Technology, 1967.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Richardson, T. and Urbanke, R. *Modern coding theory*. Cambridge university press, 2008.
- Rolfe, J. T. Discrete variational autoencoders. *arXiv preprint arXiv:1609.02200*, 2016.
- Roy, A., Vaswani, A., Neelakantan, A., and Parmar, N. Theory and experiments on vector quantized autoencoders. *arXiv preprint arXiv:1805.11063*, 2018.
- Ruiz, F. R., AUEB, M. T. R., and Blei, D. The generalized reparameterization gradient. In *Advances in neural information processing systems*, pp. 460–468, 2016.
- Santurkar, S., Budden, D., and Shavit, N. Generative compression. *arXiv preprint arXiv:1703.01467*, 2017.
- Schulman, J., Heess, N., Weber, T., and Abbeel, P. Gradient estimation using stochastic computation graphs. In *Advances in Neural Information Processing Systems*, pp. 3528–3536, 2015.
- Shannon, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27:623–656, 1948.
- Shu, R., Bui, H. H., Zhao, S., Kochenderfer, M. J., and Ermon, S. Amortized inference regularization. In *NIPS*, 2018.
- Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. *arXiv preprint arXiv:1503.03585*, 2015.
- Theis, L., Shi, W., Cunningham, A., and Huszár, F. Lossy image compression with compressive autoencoders. *arXiv preprint arXiv:1703.00395*, 2017.
- Toderici, G., O’Malley, S. M., Hwang, S. J., Vincent, D., Minnen, D., Baluja, S., Covell, M., and Sukthankar, R. Variable rate image compression with recurrent neural networks. *arXiv preprint arXiv:1511.06085*, 2015.
- Toderici, G., Vincent, D., Johnston, N., Hwang, S. J., Minnen, D., Shor, J., and Covell, M. Full resolution image compression with recurrent neural networks. In *CVPR*, pp. 5435–5443, 2017.
- Tucker, G., Mnih, A., Maddison, C. J., Lawson, J., and Sohl-Dickstein, J. Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, pp. 2627–2636, 2017.
- van den Oord, A., Vinyals, O., et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pp. 6306–6315, 2017.
- Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pp. 1096–1103. ACM, 2008.
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Manzagol, P.-A. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research*, 11(Dec):3371–3408, 2010.
- Vontobel, P. O. and Koetter, R. On low-complexity linear-programming decoding of ldpc codes. *European transactions on telecommunications*, 18(5):509–517, 2007.
- Wang, C., Chen, X., Smola, A. J., and Xing, E. P. Variance reduction for stochastic gradient optimization. In *Advances in Neural Information Processing Systems*, pp. 181–189, 2013.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Zarcone, R., Paiton, D., Anderson, A., Engel, J., Wong, H. P., and Olshausen, B. Joint source-channel coding with neural networks for analog data compression and storage. In *2018 Data Compression Conference*, pp. 147–156. IEEE, 2018.
- Zhai, F. and Katsaggelos, A. Joint source-channel video transmission. *Synthesis Lectures on Image, Video, and Multimedia Processing*, 3(1):1–136, 2007.
- Zhao, S., Song, J., and Ermon, S. Infovae: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*, 2017.
- Zhao, S., Song, J., and Ermon, S. A lagrangian perspective on latent variable generative models. In *Proc. 34th Conference on Uncertainty in Artificial Intelligence*, 2018.