
Algorithms for the Communication of Samples

Lucas Theis
Google

Noureldin Yosri
Google

Abstract

We consider the problem of *reverse channel coding*, that is, how to simulate a noisy channel over a digital channel efficiently. We propose two new coding schemes with practical advantages over previous approaches. First, we introduce *ordered random coding* (ORC) which uses a simple trick to reduce the coding cost of previous approaches based on importance sampling. Our derivation also illuminates a connection between these schemes and the so-called *Poisson functional representation*. Second, we describe a hybrid coding scheme which uses dithered quantization to efficiently communicate samples from distributions with bounded support.

1 INTRODUCTION

Consider a problem where a sender has information \mathbf{x} and wants to communicate a noisy version of it over a digital channel,

$$\mathbf{Z} \sim \mathbf{x} + \mathbf{U}. \quad (1)$$

The sender does not care which value the noise \mathbf{U} takes as long as it follows a given distribution. For example, it may be desired that the noise is a fair sample from a Gaussian distribution. Can we exploit the sender’s indifference to the exact value of the noise to save bits in the communication? More generally, we may want to send a sample from a given distribution,

$$\mathbf{Z} \sim q_{\mathbf{x}}. \quad (2)$$

How can we communicate such a sample most efficiently? This is the problem of *reverse channel coding*. While channel coding tries to communicate digital information over a noisy channel with as few errors as

possible, reverse channel coding attempts to do the opposite, namely to simulate a noisy channel over a digital channel. This problem has therefore also been referred to as “channel simulation” (e.g., Cuff, 2008) and is closely related to “relative entropy coding” (Flamich et al., 2020).

The reverse channel coding problem occurs in many applications. In neural compression, differentiable channels enable gradient-based methods to optimize encoders but these are necessarily noisy if we want to limit their capacity. For example, it is common to approximate quantization with uniform noise when training neural networks for lossy compression (Ballé et al., 2017). Reverse channel coding allows us to implement such noisy channels at test time (Agustsson and Theis, 2020) and to use arbitrary distributions in place of uniform noise (Havasi et al., 2019).

In differential privacy, most mechanisms seek to limit the amount of sensitive information revealed to another party by adding noise to the data (Dwork et al., 2006). Efficiently communicating such private information is an active area of research (e.g., Chen et al., 2020) and the goal of reverse channel coding.

Quantum teleportation can be viewed as another instance of reverse channel coding where classical bits are used to communicate stochastic information in the form of a qubit. Some of the earliest results on reverse channel coding were obtained in quantum mechanics (Bennett and Shor, 2002).

The naive approach to our problem would be to let the sender generate a sample and then to encode this noisy data. If the data or the noise stems from a continuous distribution, lossless source coding is impossible as it would require an infinite number of bits. On the other hand, lossy source coding (by first quantizing) leads to further corruption of the data and still wastes bits on encoding noise. In contrast, efficient reverse channel coding techniques are able to communicate such stochastic information with a coding cost which is close to the mutual information between the data \mathbf{X} and the sample \mathbf{Z} (e.g., Harsha et al., 2007),

$$I[\mathbf{X}, \mathbf{Z}] = h[\mathbf{Z}] - h[\mathbf{Z} | \mathbf{X}] = \mathbb{E}[D_{\text{KL}}[q_{\mathbf{X}} \| p]], \quad (3)$$

where h is the differential entropy and p is the marginal distribution of \mathbf{Z} . Unlike the naive approach, the coding cost actually decreases as we introduce more noise, that is, when the (differential) entropy of $q_{\mathbf{X}}$ increases.

The present article aims to provide an overview of several useful algorithms for reverse channel coding and to compare them. For instance, we describe a practical algorithm based on the Poisson functional representation Li and El Gamal (2018). We then introduce new algorithms along with practical advantages and provide results on their theoretical properties. As a further contribution, we present a unifying view of some of these algorithms which helps to clarify the relationship between them and sheds light on their empirical behavior. Finally, we will provide the first direct empirical comparisons between different reverse channel coding algorithms. All proofs and additional empirical results are provided in the appendix.

2 RELATED WORK

The *reverse Shannon theorem* of Bennett and Shor (2002) shows that a sender who has access to \mathbf{X} can communicate an instance of \mathbf{Z} at a cost which is close to the two random variables’ mutual information. Many papers have considered problems related to reverse channel coding and derived bounds on its coding cost (e.g., Cover and Permuter, 2007; Harsha et al., 2007; Braverman and Garg, 2014). To our knowledge, the sharpest upper bound so far was provided by Li and Anantharam (2021) who showed that an optimal code does not require more than

$$I[\mathbf{X}, \mathbf{Z}] + \log(I[\mathbf{X}, \mathbf{Z}] + 1) + 4.732 \quad (4)$$

bits to communicate an exact sample. On the other hand, Li and El Gamal (2018) showed that distributions exist for which the coding cost is at least

$$I[\mathbf{X}, \mathbf{Z}] + \log(I[\mathbf{X}, \mathbf{Z}] + 1) - 1. \quad (5)$$

That is, the bound in Eq. 4 cannot be improved significantly without making additional assumptions about the distributions involved (see also Braverman and Garg, 2014). Note that the communication overhead (the second and third term) becomes relatively less important as the transmitted amount of information increases.

Most general reverse channel coding algorithms operate on the same basic principle. First, a potentially large number of candidates is generated from a fixed distribution which is known to the sender and receiver,

$$\mathbf{Z}_n \sim p, \quad (6)$$

where $n \in \mathbb{N}$ or $n \in \{1, \dots, N\}$. Both the sender and receiver are able to generate these candidates without

communication by using a shared source of randomness. In practice, this will typically be a pseudorandom number generator with a common seed. The sender selects an index N^* according to some distribution such that, at least approximately,

$$\mathbf{Z}_{N^*} \sim q_{\mathbf{X}}. \quad (7)$$

Note that only N^* needs to be communicated and this can be done efficiently if $H[N^*]$ is small. The main difference between algorithms is in how N^* is decided.

Li and El Gamal (2017) described an algorithm for communicating samples from distributions with log-concave PDFs without common randomness. Without a shared source of randomness, the number of bits required is at least *Wyner’s common information* (Wyner, 1975; Cuff, 2008), which can be significantly larger than the mutual information (Xu et al., 2011). In the following, we therefore focus on algorithms with access to common randomness.

Agustsson and Theis (2020) showed that there is no general algorithm whose computational complexity is polynomial in the communication cost. That is, as the amount of information transmitted increases, general purpose algorithms become prohibitively expensive. One solution to this problem is to split information into chunks and to encode them separately (Havasi et al., 2019; Flamich et al., 2020). However, this reduces statistical efficiency as each chunk will contribute its own overhead to the coding cost. We therefore typically find a tension between the computational efficiency and the coding efficiency of a scheme.

A more well-known idea in machine learning is *bits-back coding* (Wallace, 1990; Hinton and Van Camp, 1993), which at first sight appears closely related to reverse channel coding. Here, the goal is to losslessly compress a source \mathbf{X} using a model of its joint distribution with a set of latent variables \mathbf{Z} . Encoding an instance \mathbf{x} involves sampling $\mathbf{Z} \sim q_{\mathbf{Z}}$ while using previously encoded bits as a source of randomness. The data and latent variables are subsequently encoded using the model’s joint distribution (Townsend et al., 2019). Unlike reverse channel coding, however, bits-back coding necessarily transmits a perfect copy of the data, that is, it is an implementation of source coding. On the other hand, reverse channel coding can be viewed as a generalization of source coding which also supports lossy transmission. Source coding is recovered as a special case when choosing $q_{\mathbf{Z}}(\mathbf{z}) = \delta(\mathbf{z} - \mathbf{x})$.

3 ALGORITHMS

We will first continue the discussion of related work by introducing existing algorithms for the simulation

of noisy channels. New ideas are presented mainly in Sections 3.4, 3.5, and 3.7.

3.1 Rejection sampling

Rejection sampling (RS) is a method for generating a sample from one distribution given samples from another distribution. As an introductory example, we show how RS can be turned into a reverse channel coding scheme.

Let \mathbf{Z}_n be candidates drawn from a proposal distribution p . Further, let $U_n \sim \text{Uniform}([0, 1])$. RS selects the first index N_{RS}^* such that

$$U_{N_{\text{RS}}^*} \leq w_{\min} \frac{q_{\mathbf{x}}(\mathbf{Z}_{N_{\text{RS}}^*})}{p(\mathbf{Z}_{N_{\text{RS}}^*})} \quad (8)$$

where

$$w_{\min} \leq \inf_{\mathbf{z}} \frac{p(\mathbf{z})}{q_{\mathbf{x}}(\mathbf{z})} \quad (9)$$

ensures that the right-hand side in Eq. 8 is a probability. If $w_{\min} > 0$, then N^* will be finite and, crucially,

$$\mathbf{Z}_{N_{\text{RS}}^*} \sim q_{\mathbf{x}}. \quad (10)$$

A sender could thus communicate a sample from $q_{\mathbf{x}}$ by sending N_{RS}^* , assuming the receiver already has access to the candidates \mathbf{Z}_n . Note that this works even when the distribution is continuous since N_{RS}^* will still be discrete. While w_{\min} may depend on \mathbf{x} , in the following analysis we assume for simplicity that the same value is used for all target distributions.

Let us consider the coding cost of encoding N_{RS}^* . The average probability of accepting a candidate is

$$\int p(\mathbf{z}) w_{\min} \frac{q_{\mathbf{x}}(\mathbf{z})}{p(\mathbf{z})} d\mathbf{z} = w_{\min}. \quad (11)$$

The marginal distribution of N_{RS}^* is therefore a geometric distribution whose entropy can be bounded by

$$H[N_{\text{RS}}^*] \leq -\log w_{\min} + 1/\ln 2. \quad (12)$$

Rejection sampling is efficient if $H[N_{\text{RS}}^*]$ is not much more than the information contained in \mathbf{Z} . While it is easy to construct examples where RS is efficient, it is also easy to construct examples where $-\log w_{\min}$ is significantly larger than $I[\mathbf{X}, \mathbf{Z}]$. For instance, the density ratio in Eq. 9 may be unbounded. However, if we are willing to accept an approximate sample, then there are ways to limit the coding cost even then. For example, we may unconditionally accept the N th candidate if the first $N - 1$ candidates are rejected. In this case, the distribution of \mathbf{Z}_{N^*} will be a mixture distribution with density

$$\beta p(\mathbf{z}) + (1 - \beta) q_{\mathbf{x}}(\mathbf{z}), \quad (13)$$

where $\beta = (1 - w_{\min})^{N-1}$ is the probability of rejecting all $N - 1$ candidates. The quality of a sample is often measured in terms of the *total variation distance* (TVD),

$$D_{\text{TV}}[p, q] = \frac{1}{2} \int |p(\mathbf{z}) - q(\mathbf{z})| d\mathbf{z}. \quad (14)$$

When we measure the deviation of the mixture distribution from the target distribution $q_{\mathbf{x}}$, we obtain

$$D_{\text{TV}}[\beta p + (1 - \beta) q_{\mathbf{x}}, q_{\mathbf{x}}] = \beta D_{\text{TV}}[p, q_{\mathbf{x}}]. \quad (15)$$

That is, the divergence decays exponentially with N .

An alternative approach to limiting the coding cost is to choose an invalid but larger w_{\min} . Harsha et al. (2007) described a related approach which effectively uses $w_{\min} = 1$ but is nevertheless able to produce an exact sample by adjusting the target distribution after each candidate rejection (Appendix A). However, their approach is computationally expensive and generally infeasible for continuous distributions.

3.2 Minimal random coding

An approach closely related to *importance sampling* was first considered by Cuff (2008) and later dubbed *likelihood encoder* (Song et al., 2016). The approach was independently rediscovered in machine learning by Havasi et al. (2019) who referred to it as *minimal random coding* (MRC) and used it for model compression. It has since also been used for lossy image compression (Flamich et al., 2020). Unlike Havasi et al. (2019), Cuff (2008) only considered discrete distributions and assumed that p is the true marginal distribution of the data. But Cuff (2008) also described a more general approach where the amount of shared randomness between the sender and receiver is limited.

In MRC, the sender picks one of N candidates by sampling an index N_{MRC}^* from the distribution

$$\pi_{\mathbf{x}}(n) \propto q_{\mathbf{x}}(\mathbf{Z}_n) / p(\mathbf{Z}_n). \quad (16)$$

Unlike RS, the distribution of $\mathbf{Z}_{N_{\text{MRC}}^*}$ (call it $\tilde{q}_{\mathbf{x}}$) will in general only approximate $q_{\mathbf{x}}$. On the other hand, the coding cost of MRC can be significantly smaller. Havasi et al. (2019) showed that under reasonable assumptions, samples from $\tilde{q}_{\mathbf{x}}$ will be similar to samples from $q_{\mathbf{x}}$ if the number of candidates is

$$N = 2^{D_{\text{KL}}[q_{\mathbf{x}} \| p] + t} \quad (17)$$

for some $t > 0$. That is, the number of candidates required to guarantee a sample of high quality grows exponentially with the amount of information gained by the receiver.

Since acceptable candidates may appear anywhere in the sequence of candidates, each index is a priori equally likely to be picked. That is, the marginal distribution of N_{MRC}^* is uniform and its entropy is

$$H[N_{\text{MRC}}^*] = \log N. \quad (18)$$

3.3 Poisson functional representation

Li and El Gamal (2018) introduced the following *Poisson functional representation* (PFR) of a random variable. Let T_n be the arrival times of a homogeneous Poisson process on the non-negative real line such that $T_n \geq 0$ and $T_n \leq T_{n+1}$ for all $n \in \mathbb{N}$. Let

$$\mathbf{Z}_n \sim p, \quad N_{\text{PFR}}^* = \operatorname{argmin}_{n \in \mathbb{N}} T_n \frac{p(\mathbf{Z}_n)}{q_{\mathbf{x}}(\mathbf{Z}_n)}. \quad (19)$$

for all $n \in \mathbb{N}$. Then $\mathbf{Z}_{N_{\text{PFR}}^*}$ has the distribution $q_{\mathbf{x}}$. As in RS, the index N_{PFR}^* picks one of infinitely many candidates and we obtain an exact sample from the target distribution. However, Li and El Gamal (2018) provided much stronger guarantees for the coding cost of the N_{PFR}^* . In particular,

$$H[N_{\text{PFR}}^*] \leq I[\mathbf{X}, \mathbf{Z}] + \log(I[\mathbf{X}, \mathbf{Z}] + 1) + 4. \quad (20)$$

The distribution of N_{PFR}^* takes a more complicated form than in RS or MRC (Li and Anantharam, 2021, Eq. 29). Nevertheless, a coding cost corresponding to the bound above can be achieved by entropy encoding N_{PFR}^* while assuming a Zipf distribution $p_{\lambda}(n) \propto n^{-\lambda}$ with (Li and El Gamal, 2018)

$$\lambda = 1 + 1/(I[\mathbf{X}, \mathbf{Z}] + e^{-1} \log e + 1). \quad (21)$$

A downside of the PFR is that it depends on an infinite number of candidates. Unlike rejection sampling, we cannot consider the candidates in sequence but generally have to consider the scores of all candidates (Eq. 19). However, if we can bound the density ratio as in rejection sampling (Eq. 9), then we can terminate our search for the smallest score after considering a finite number of candidates. Let

$$S_n^* = \operatorname{argmin}_{i \leq n} T_i \frac{p(\mathbf{Z}_i)}{q_{\mathbf{x}}(\mathbf{Z}_i)} \quad (22)$$

be the smallest score observed after taking into account n candidates. Since $T_m \geq T_n$ for all $m > n$, all further scores will be at least $T_n w_{\min}$. Hence, if $S_n^* \leq T_n w_{\min}$, we can terminate the search for the best candidate. Algorithm 1 summarizes this idea. Here, `simulate`(n, \mathbf{p}) is a function which simulates a distribution \mathbf{p} by returning the n th pseudorandom sample derived from n and an implicit random seed. Similarly, `expon`($n, 1$) simulates an exponential distribution with rate 1.

Algorithm 1 PFR

Require: $\mathbf{p}, \mathbf{q}, w_{\min}$

```

1:  $t, n, s^* \leftarrow 0, 1, \infty$ 
2: repeat
3:    $z \leftarrow \text{simulate}(n, \mathbf{p})$    ▷ Candidate generation
4:    $t \leftarrow t + \text{expon}(n, 1)$    ▷ Poisson process
5:    $s \leftarrow t \cdot \mathbf{p}(z)/\mathbf{q}(z)$    ▷ Candidate's score
6:   if  $s < s^*$  then           ▷ Accept/reject candidate
7:      $s^*, n^* \leftarrow s, n$ 
8:   end if
9:    $n \leftarrow n + 1$ 
10: until  $s^* \leq t \cdot w_{\min}$ 
11: return  $n^*$ 
    
```

3.4 Ordered random coding

In typical applications of importance sampling we do not worry about coding costs and it may therefore not surprise that the entropy of N_{MRC}^* is large. We here show that a slight modification reduces the entropy of the selected index while keeping the distribution of the communicated sample exactly the same.

To generate a sample from $\pi_{\mathbf{x}}$ (Eq. 16), we can write

$$N^* = \operatorname{argmax}_{n \leq N} \log q_{\mathbf{x}}(\mathbf{Z}_n) - \log p(\mathbf{Z}_n) + G_n, \quad (23)$$

where the G_n are Gumbel distributed random variables (Gumbel, 1954) with scale parameter 1. This is the well-known Gumbel-max trick for sampling from a categorical distribution (Maddison et al., 2014). This trick still works if we permute the Gumbel random variables arbitrarily, as long as the permutation does not depend on the values of the candidates \mathbf{Z}_n .

Theorem 1. *Let \tilde{G}_n be the result of sorting the Gumbel random variables in decreasing order such that $\tilde{G}_1 \geq \dots \geq \tilde{G}_N$ and define*

$$\tilde{N}^* = \operatorname{argmax}_{n \leq N} \log q_{\mathbf{x}}(\mathbf{Z}_n) - \log p(\mathbf{Z}_n) + \tilde{G}_n. \quad (24)$$

Then $\mathbf{Z}_{N^} \sim \mathbf{Z}_{\tilde{N}^*}$.*

If the density ratio is bounded, the sample quality improves quickly once the coding cost exceeds the information contained in a sample. In particular, we have the following result.

Theorem 2. *Let \tilde{q} be the distribution of $\mathbf{Z}_{\tilde{N}^*}$. If the number of candidates is $N = 2^{D_{\text{KL}}[q \| p] + t}$ and $p(\mathbf{z})/q(\mathbf{z}) \geq w_{\min} > 0$ for all \mathbf{z} , then*

$$D_{\text{TV}}[\tilde{q}, q] = O(2^{-t/8}). \quad (25)$$

While the distribution of $\mathbf{Z}_{\tilde{N}^*}$ remains the same, the distribution of \tilde{N}^* is no longer uniform but biased towards smaller indices. Since N^* is uniform, we must have $H[\tilde{N}^*] \leq H[N^*]$. In the next section, we show that $H[\tilde{N}^*]$ is in fact on a par with $H[N_{\text{PFR}}^*]$. We dub the approach *ordered random coding* (ORC) and pseudocode for ORC is provided in Appendix B.

3.5 A unifying view

In this section we show a close connection between methods based on importance sampling and the PFR. First, we can rewrite Eq. 23 as follows,

$$N_{\text{MRC}}^* = \operatorname{argmin}_{n \leq N} S_n \frac{p(\mathbf{Z}_n)}{q_{\mathbf{x}}(\mathbf{Z}_n)} \quad (26)$$

where S_n are exponentially distributed with rate 1. This is true because $-\log S_n$ is Gumbel distributed. Note the similarity to the PFR,

$$N_{\text{PFR}}^* = \operatorname{argmin}_{n \in \mathbb{N}} T_n \frac{p(\mathbf{Z}_n)}{q_{\mathbf{x}}(\mathbf{Z}_n)}, \quad (27)$$

where $T_n \sim \sum_{m=1}^n S_m$. ORC, on the other hand, first sorts the exponential random variables, $\tilde{S}_1 \leq \dots \leq \tilde{S}_N$, before choosing

$$N_{\text{ORC}}^* = \operatorname{argmin}_{n \leq N} \tilde{S}_n \frac{p(\mathbf{Z}_n)}{q_{\mathbf{x}}(\mathbf{Z}_n)}. \quad (28)$$

It turns out that (Rényi, 1953)

$$\tilde{S}_n \sim \sum_{m=1}^n \frac{S_m}{N - m + 1}. \quad (29)$$

This allows us to generate the sorted exponential variables in $O(N)$ instead of $O(N \log N)$ time with sorting. More interestingly, Eq. 29 reveals a close connection to the PFR. Where the PFR uses cumulative sums of exponential random variables, ORC uses weighted sums. This representation allows us to arrive at the following result.

Theorem 3. *Let S_n be exponentially distributed RVs and $\mathbf{Z}_n \sim p$ for all $n \in \mathbb{N}$ (i.i.d.), and let*

$$T_n = \sum_{m=1}^n S_m, \quad \tilde{T}_{N,n} = \sum_{m=1}^n \frac{N}{N - m + 1} S_m \quad (30)$$

for $N \in \mathbb{N}$. Further, let

$$N_{\text{PFR}}^* = \operatorname{argmin}_{n \in \mathbb{N}} T_n \frac{p(\mathbf{Z}_n)}{q(\mathbf{Z}_n)}, \quad (31)$$

$$N_{\text{ORC}}^* = \operatorname{argmin}_{n \leq N} \tilde{T}_{N,n} \frac{p(\mathbf{Z}_n)}{q(\mathbf{Z}_n)}. \quad (32)$$

Then $N_{\text{ORC}}^* \leq N_{\text{PFR}}^*$. Further, there exists an $M \in \mathbb{N}$ such that for all $N \geq M$ we have $N_{\text{ORC}}^* = N_{\text{PFR}}^*$.

Note that the additional factor N in the definition of $\tilde{T}_{N,n}$ compared to \tilde{S}_n does not change the output of argmin . Using Theorem 3, it is not difficult to see that the bound on the coding of the PFR also applies to ORC, or the following result.

Corollary 1. *Let $C = \mathbb{E}_{\mathbf{x}}[D_{\text{KL}}[q_{\mathbf{x}} \parallel p]]$ and let N_{ORC}^* be defined as in Theorem 3. Then*

$$H[N_{\text{ORC}}^*] < C + \log(C + 1) + 4. \quad (33)$$

To achieve this bound, a Zipf distribution $p_{\lambda}(n) \propto n^{-\lambda}$ with $\lambda = 1 + 1/(C + e^{-1} \log e + 1)$ can be used to entropy encode the index, analogous to the PFR (Eq. 21).

The significance of these results is as follows. Li and El Gamal (2018) showed that the PFR is near-optimal in the sense that the entropy of N_{PFR}^* is close to the worst-case coding cost needed for communicating a perfect sample. However, the construction of the PFR relies on an infinite number of candidates and all theoretical guarantees are lost if we naively limit the number of candidates to N . In particular, we do not know how quickly the quality of a communicated sample deteriorates as we decrease N , or how large N should be. On the other hand, we do have some idea of the quality of a sample obtained via importance sampling from a finite number of candidates (e.g., Theorem 2 or the results of Cuff, 2008; Chatterjee and Diaconis, 2018; Havasi et al., 2019). But the coding cost of MRC is relatively large and continues to grow unbounded as N increases. ORC combines the best of both by inheriting the coding cost of the PFR and the sample quality of MRC.

Unlike the PFR, the guarantees of ORC hold for a finite number of samples. Unlike MRC, we can make N arbitrarily large without having to worry about an exploding coding cost, which also makes it easier to tune this parameter.

3.6 Dithered quantization

Dithered quantization, also known as *universal quantization* (Ziv, 1985), refers to quantization with a randomly shifted lattice. Consider a scalar $y \in \mathbb{R}$ and a random variable U uniform over any interval of length one, such as $[0, 1)$. Then (e.g., Schuchman, 1964)

$$\lfloor y - U \rfloor + U \sim y + U_0, \quad (34)$$

where U_0 is uniform over $[-0.5, 0.5)$. More generally, let Q be a quantizer which maps inputs to the nearest point on a lattice and let \mathbf{V} be a random variable which is uniformly distributed over an arbitrarily placed Voronoi cell of the lattice. Further, let \mathbf{V}_0 be

uniform over the Voronoi cell which contains the lattice point at zero. Then (Zamir, 2014)

$$Q(\mathbf{y} - \mathbf{V}) + \mathbf{V} \sim \mathbf{y} + \mathbf{V}_0. \quad (35)$$

Dithered quantization has been mainly used as a tool for studying quantization from a theoretical perspective. However, already Roberts (1962) considered some of its practical advantages over uniform quantization for compressing grayscale images, especially in terms of perceptual quality. Theis and Agustsson (2021) showed that universal quantization can outperform vector quantization where a realism constraint is considered. Universal quantization also recently started being used in neural compression (Choi et al., 2019), in particular to realize differentiable training losses at inference time (Agustsson and Theis, 2020).

To communicate a sample of a uniform distribution centered around y , the sender encodes $K = \lfloor y - U \rfloor$. The receiver decodes K and computes $Z = K + U$, which is distributed as $y + U_0$. Like the reverse channel coding schemes discussed so far, this requires a shared source of randomness in the form of U . Zamir and Feder (1992) showed that dithered quantization is statistically efficient in the sense that

$$H[K | U] = I[Y, Z]. \quad (36)$$

That is, the cost of encoding K is as close as can be to the amount of information contained in Z . Note that we can condition on U when encoding K since U is known to both sender and receiver. Another advantage of dithered quantization is that it is computationally highly efficient, at least for the simple case of the integer lattice.

3.7 Hybrid coding

While dithered quantization is computationally much more efficient than general reverse channel coding schemes, it is also much more limited in terms of the distributions it can simulate. Here we propose a hybrid coding scheme for continuous distributions which retains most of the flexibility of general purpose schemes but is computationally more efficient when the support of the target distribution is small.

The general idea is as follows. Instead of drawing candidates from a fixed distribution p , candidates \mathbf{Z}_n are drawn from a distribution $r_{\mathbf{x}}$ which acts as a bridge and more closely resembles the target distribution $q_{\mathbf{x}}$. Since $r_{\mathbf{x}}$ is closer to $q_{\mathbf{x}}$, we will require fewer candidates to find one that is suitable. Let

$$N^* = \operatorname{argmin}_{n \leq N} \frac{\tilde{T}_{N,n} r_{\mathbf{x}}(\mathbf{Z}_n)}{q_{\mathbf{x}}(\mathbf{Z}_n)}, \quad (37)$$

be the index of the selected candidate. Unlike before, only the sender has access to the candidates and so knowing N^* alone will not allow us to reconstruct \mathbf{Z}_{N^*} . Our hybrid coding scheme relies on two insights. First, the receiver does not require access to all candidates but only to the selected candidate. Second, the missing information can be encoded easily and efficiently if $r_{\mathbf{x}}$ can be simulated with dithered quantization.

We assume for now that there exist vectors $\mathbf{c}_{\mathbf{x}}$ such that the support of $q_{\mathbf{x}}$ is contained in the support of

$$r_{\mathbf{x}}(\mathbf{z}) = \begin{cases} 1 & \text{if } \mathbf{z} \in \mathbf{c}_{\mathbf{x}} + [-0.5, 0.5]^D, \\ 0 & \text{else} \end{cases} \quad (38)$$

which can be simulated via dithered quantization,

$$\mathbf{K}_n = \lfloor \mathbf{c}_{\mathbf{x}} - \mathbf{U}_n \rfloor, \quad \mathbf{Z}_n = \mathbf{K}_n + \mathbf{U}_n, \quad (39)$$

where $\mathbf{U}_n \sim \text{Uniform}([0, 1]^D)$. Define $\mathbf{K}^* = \mathbf{K}_{N^*}$. Hybrid coding transmits the pair (N^*, \mathbf{K}^*) , which the receiver uses to reconstruct the selected candidate via

$$\mathbf{Z}_{N^*} = \mathbf{K}^* + \mathbf{U}_{N^*}. \quad (40)$$

Theorem 4. *Let N^* and \mathbf{K}^* be defined as in Eq. 37 and below Eq. 39 and let p be the uniform distribution over $[0, M_1) \times \dots \times [0, M_D)$ for some $M_i \in \mathbb{N}$. Then*

$$H[N^*, \mathbf{K}^*] < C + \log(C - \sum_i \log M_i + 1) + 4,$$

where $C = \mathbb{E}_{\mathbf{X}}[D_{\text{KL}}[q_{\mathbf{X}} \parallel p]]$.

Theorem 4 shows that (N^*, \mathbf{K}^*) is an efficient representation if the marginal distribution of \mathbf{Z} is uniform over some box whose sides have lengths M_i , since then $C = I[\mathbf{X}, \mathbf{Z}]$. However, for continuous random variables this can always be achieved through a transformation Ψ . If $\tilde{q}_{\mathbf{x}}$ is the desired target distribution before the transformation, then

$$q_{\mathbf{x}}(\mathbf{z}) = \tilde{q}_{\mathbf{x}}(\Psi(\mathbf{z})) |D\Psi(\mathbf{z})| \quad (41)$$

is the target distribution in transformed space. Note that after the transformation, the support of $q_{\mathbf{x}}$ is always bounded. Moreover, for small enough M_i the support will be contained in the support of $r_{\mathbf{x}}$, satisfying our earlier assumption. After transmitting a sample from $q_{\mathbf{x}}$ via hybrid coding, it is assumed that the receiver applies the transformation Ψ to obtain a sample from $\tilde{q}_{\mathbf{x}}$.

To achieve the bound in Theorem 4, the sender first encodes N^* while assuming a Zipf distribution $p_{\lambda}(n) \propto n^{-\lambda}$ with

$$\lambda = 1 + 1/(C - \sum_i \log M_i + e^{-1} \log e + 1). \quad (42)$$

The K_i^* are subsequently added to the bit stream using a fixed rate of $\log M_i$ bits.

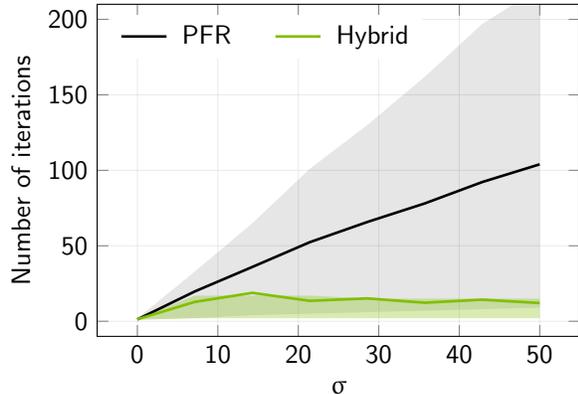


Figure 1: Computational cost of communicating a sample from a (truncated) Gaussian whose mean varies with standard deviation σ . The vertical axis measures the average number of candidates considered before an algorithm terminates. Shaded regions indicate the 10th and 90th percentile.

Algorithm 2 describes the algorithm used for encoding. Here, q is the transformed density and p is not needed as it is assumed to be uniform. Note that hybrid coding effectively reduces to ORC when the support is unconstrained ($M_i = 1$). For larger M_i , the bound on the coding cost improves only slightly but the computational cost reduces significantly. Since the number of candidates required for a sample of high quality grows exponentially in the KLD (Eq. 17) and

$$D_{\text{KL}}[q_{\mathbf{x}} \parallel r_{\mathbf{x}}] = D_{\text{KL}}[q_{\mathbf{x}} \parallel p] - \log M_i, \quad (43)$$

we should expect a speedup on the order of $\prod_i M_i$. We thus want to maximize M_i while making sure that the support of $q_{\mathbf{x}}$ is still contained within that of $r_{\mathbf{x}}$.

4 EXPERIMENTS

We run two sets of empirical experiments to compare the reverse channel coding schemes discussed above. We first investigate the effect of hybrid coding on the computational cost of communicating a (truncated) Gaussian sample. We then compare the performance of a wider set of algorithms for the task of approximately simulating a categorical distribution.

4.1 Gaussian distribution

Consider the task of communicating a sample from a D -dimensional Gaussian with random mean,

$$\mathbf{Z} \sim \mathcal{N}(\mathbf{X}, \mathbf{I}), \quad \mathbf{X} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}), \quad (44)$$

where \mathbf{I} is the identity matrix and the mean \mathbf{X} itself is Gaussian distributed with covariance $\sigma^2 \mathbf{I}$. The

Algorithm 2 Hybrid coding

Require: c, q, w_{\min}, N

```

1:  $t, n, s^* \leftarrow 0, 1, \infty$ 

2: repeat
3:    $u \leftarrow \text{uniform}(n, 0, 1)$   $\triangleright$  Candidate generation
4:    $k \leftarrow \text{round}(c - u)$ 
5:    $z \leftarrow k + u$ 

6:    $v \leftarrow N/(N - n + 1)$ 
7:    $t \leftarrow t + v \cdot \text{expon}(n, 1)$ 
8:    $s \leftarrow t/q(z)$   $\triangleright$  Candidate's score

9:   if  $s < s^*$  then  $\triangleright$  Accept/reject candidate
10:     $s^*, n^*, k^* \leftarrow s, n, k$ 
11:   end if

12:    $n \leftarrow n + 1$ 
13: until  $s^* \leq t \cdot w_{\min}$  or  $n > N$ 

14: return  $n^*, k^*$ 
    
```

marginal distribution of \mathbf{Z} is Gaussian with mean zero and covariance $\sigma^2 \mathbf{I} + \mathbf{I}$ and so we use this distribution as our candidate generating distribution p . The average information gained by obtaining a sample is

$$I[\mathbf{X}, \mathbf{Z}] = -\frac{D}{2} \log \left(1 - \frac{\sigma^2}{\sigma^2 + 1} \right). \quad (45)$$

To be able to apply the hybrid coding scheme, we slightly truncate the target distribution by assigning zero density to a small fraction θ of values with the lowest density. The TVD between the truncated Gaussian and the Gaussian distribution is θ . A classifier observing \mathbf{Z} would be able to distinguish between these two distributions with an accuracy of at most $1/2 + \theta/2$. In our experiments, we fix $\theta = 10^{-4}$ so that this accuracy is close to chance.

We compare hybrid coding (Algorithm 2) to ORC with $N = \infty$, which reduces to the PFR (Algorithm 1). Using an unlimited number of candidates allows us to avoid any further approximations and to focus on the computational cost. Appropriate values for w_{\min} and M are provided in Appendix H.

Figure 1 shows the average number of iterations an algorithm runs before identifying a suitable candidate of a 1-dimensional (truncated) Gaussian. The computational cost of the PFR grows exponentially with the amount of information transmitted, which is approximately $\log \sigma$. On the other hand, the computational cost of the hybrid coding scheme quickly saturates and remains low throughout, allowing for much quicker communication of the Gaussian sample.

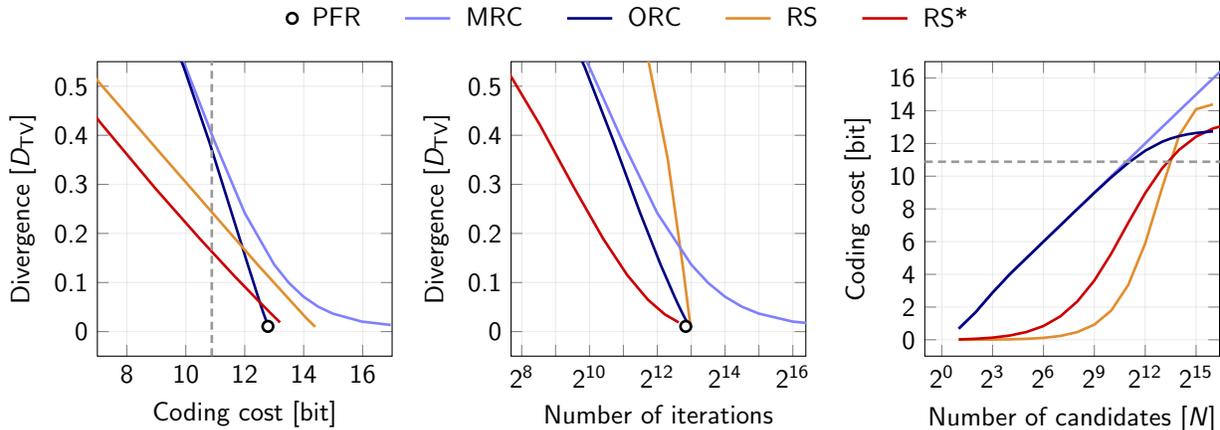


Figure 2: A comparison of various reverse channel coding algorithms for 2^{16} -dimensional categorical distributions which are themselves randomly distributed according to a Dirichlet distribution. Dashed lines indicate the average KLD between the target distribution and the uniform candidate generating distribution. *Left:* The sample quality as a function of the coding cost. *Middle:* The sample quality as a function of the computational cost (for which the number of iterations is a proxy, except for RS*). *Right:* The coding cost as a function of the maximum number of candidates considered.

4.2 Categorical distribution

As another example we consider D -dimensional categorical distributions distributed according to a Dirichlet distribution with concentration parameter α . We chose $D = 2^{16}$ and $\alpha_i = 3 \cdot 10^{-4}$ for all i , leading to sparse target distributions and a uniform marginal distribution. We include rejection sampling (RS) with an optimal choice for w_{\min} (a different value for each distribution) as well as the greedy rejection sampler (RS*) of Harsha et al. (2007) in the comparison. For each method and target distribution, we simulate 10^5 samples and measure the TVD between the resulting histogram and the target distribution. As a measure of the coding cost, we estimate the entropy of the index distribution obtained by averaging index histograms of 20 different target distributions.

We explore the effects of limiting the number of candidates available to an algorithm. We find that the sample quality of all methods deteriorates quickly as the coding cost drops below the information contained in exact samples (Fig. 2, left and middle). RS performed surprisingly well in the bit-rate constrained regime (left) but not as well when constraining computational cost (middle). RS* performed even better in the low bit-rate regime but we note that its iterations’ computational complexity is larger by a factor D compared to the other methods. PFR and ORC performed best for samples of high quality.

MRC (Cuff, 2008; Havasi et al., 2019) performed worse than the other methods mostly due to the coding and computational cost growing unboundedly with the

number of candidates (Fig. 2, right). ORC addresses this issue such that its coding cost converges to that of the PFR (Li and El Gamal, 2018).

5 DISCUSSION

We demonstrated a close connection between minimal random coding (MRC; Havasi et al., 2019), or likelihood encoding (Cuff, 2008; Song et al., 2016), and the Poisson functional representation (PFR; Li and El Gamal, 2018). Ordered random coding (ORC) occupies a space between the two and benefits from the theoretical guarantees of both. In practice, we found that ORC can significantly outperform MRC, especially as the desired sample quality increases (achieving a 20% reduction in coding cost at a TVD of about 0.02).

Our second coding scheme enables much more efficient communication of samples from distributions with concentrated support. When the target distributions’ support is unbounded, hybrid coding may still be applied after truncation, as in the Gaussian example. While the hybrid scheme is more efficient than other approaches for the Gaussian example, its cost still grows exponentially with dimensionality. A potential solution is to generate candidates using more sophisticated lattices than the integer lattice considered here (e.g., Leech, 1967; Zamir, 2014).

While some distributions are known to be hard to simulate (Long and Servedio, 2010), an important task for future research is to characterize other distributions which can be simulated efficiently and which are therefore of great interest for practical applications.

Acknowledgements

We would like to thank Abhin Shah, Johannes Ballé, Eirikur Agustsson, and Aaron B. Wagner for helpful discussions of the ideas presented in this manuscript.

References

- E. Agustsson and L. Theis. Universally Quantized Neural Compression. In *Advances in Neural Information Processing Systems 33*, 2020.
- J. Ballé, V. Laparra, and E. P. Simoncelli. End-to-end Optimized Image Compression. In *International Conference on Learning Representations*, 2017.
- C. H. Bennett and P. W. Shor. Entanglement-Assisted Capacity of a Quantum Channel and the Reverse Shannon Theorem. *IEEE Trans. Info. Theory*, 48(10), 2002.
- M. Braverman and A. Garg. Public vs private coin in bounded-round information. In J. Esparza, P. Fraigniaud, T. Husfeldt, and E. Koutsoupias, editors, *Automata, Languages, and Programming*, pages 502–513. Springer, 2014.
- S. Chatterjee and P. Diaconis. The sample size required in importance sampling. *The Annals of Applied Probability*, 28(2):1099–1135, 2018.
- W.-N. Chen, P. Kairouz, and A. Ozgur. Breaking the communication-privacy-accuracy trilemma. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 3312–3324. Curran Associates, Inc., 2020.
- Y. Choi, M. El-Khamy, and J. Lee. Variable rate deep image compression with a conditional autoencoder. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019.
- T. M. Cover and H. H. Permuter. Capacity of coordinated actions. In *2007 IEEE International Symposium on Information Theory*, pages 2701–2705, 2007. doi: 10.1109/ISIT.2007.4557184.
- P. Cuff. Communication requirements for generating correlated random variables. In *2008 IEEE International Symposium on Information Theory*, pages 1393–1397, 2008.
- C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In S. Halevi and T. Rabin, editors, *Theory of Cryptography*, pages 265–284. Springer Berlin Heidelberg, 2006.
- G. Flamich, M. Havasi, and J. M. Hernández-Lobato. Compressing Images by Encoding Their Latent Representations with Relative Entropy Coding, 2020. *Advances in Neural Information Processing Systems* 34.
- E. J. Gumbel. Statistical Theory of Extreme Values and Some Practical Applications. *U.S. Department of Commerce, National Bureau of Standards*, 33, 1954.
- P. Harsha, R. Jain, D. McAllester, and J. Radhakrishnan. The Communication Complexity of Correlation. In *Twenty-Second Annual IEEE Conference on Computational Complexity*, pages 10–23, 2007.
- M. Havasi, R. Peharz, and J. M. Hernández-Lobato. Minimal Random Code Learning: Getting Bits Back from Compressed Model Parameters. In *International Conference on Learning Representations*, 2019.
- G. E. Hinton and D. Van Camp. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 5–13, 1993.
- J. Leech. Notes on sphere packings. *Canadian Journal of Mathematics*, 19:251–267, 1967.
- C. T. Li and V. Anantharam. A unified framework for one-shot achievability via the poisson matching lemma. *IEEE Transactions on Information Theory*, 67(5):2624–2651, 2021. doi: 10.1109/TIT.2021.3058842.
- C. T. Li and A. El Gamal. Distributed simulation of continuous random variables. *IEEE Transactions on Information Theory*, 63(10):6329–6343, 2017. doi: 10.1109/TIT.2017.2735438.
- C. T. Li and A. El Gamal. Strong Functional Representation Lemma and Applications to Coding Theorems. *IEEE Transactions on Information Theory*, 64(11):6967–6978, 2018. doi: 10.1109/TIT.2018.2865570.
- P. M. Long and R. A. Servedio. Restricted Boltzmann Machines are Hard to Approximately Evaluate or Simulate. In *Proceedings of the 27th International Conference on Machine Learning*, 2010.
- C. J. Maddison, D. Tarlow, and T. Minka. A* Sampling. In *Advances in Neural Information Processing Systems*, volume 27, 2014.
- L. G. Roberts. Picture Coding Using Pseudo-Random Noise. *IRE Transactions on Information Theory*, 1962.
- A. Rényi. On the theory of order statistics. *Acta Mathematica Hungarica*, 4:191–231, 1953.
- L. Schuchman. Dither signals and their effect on quantization noise. *IEEE Transactions on Communication Technology*, 12(4):162–165, 1964.
- E. C. Song, P. Cuff, and H. V. Poor. The likelihood encoder for lossy compression. *IEEE Transactions*

- on *Information Theory*, 62(4):1836–1849, 2016. doi: 10.1109/TIT.2016.2529657.
- L. Theis and E. Agustsson. On the advantages of stochastic encoders. In *Neural Compression Workshop at ICLR 2021*, 2021.
- J. Townsend, T. Bird, and D. Barber. Practical lossless compression with latent variables using bits back coding. *arXiv preprint arXiv:1901.04866*, 2019.
- C. S. Wallace. Classification by minimum-message-length inference. In *International Conference on Computing and Information*, pages 72–81. Springer, 1990.
- A. Wyner. The common information of two dependent random variables. *IEEE Transactions on Information Theory*, 21(2):163–179, 1975. doi: 10.1109/TIT.1975.1055346.
- G. Xu, W. Liu, and B. Chen. Wyners common information for continuous random variables - a lossy source coding interpretation. In *45th Annual Conference on Information Sciences and Systems*, pages 1–6, 2011. doi: 10.1109/CISS.2011.5766249.
- R. Zamir. *Lattice Coding for Signals and Networks*. Cambridge University Press, 2014.
- R. Zamir and M. Feder. On universal quantization by randomized uniform/lattice quantizers. *IEEE Transactions on Information Theory*, 38(2):428–436, 1992.
- J. Ziv. On universal quantization. *IEEE Transactions on Information Theory*, 31(3):344–347, 1985.