Faster Relative Entropy Coding with Greedy Rejection Coding

Gergely Flamich*[1], Stratis Markou*[1], José Miguel Hernández Lobato[1],
{gf332, em626, jmh233}@cam.ac.uk, *Equal contribution, [1]University of Cambridge,

Paper

TLDR: We present greedy rejection coding, the first relative entropy coding algorithm whose runtime and codelength are optimal.

Problem Setup and Motivation

Setup: *Alice only:* target Q.

Alice & Bob: proposal P and **public** fair coin tosses $s_1, s_2, \ldots$.

Goal: Uniquely decodable code C which represents exact sample from Q.

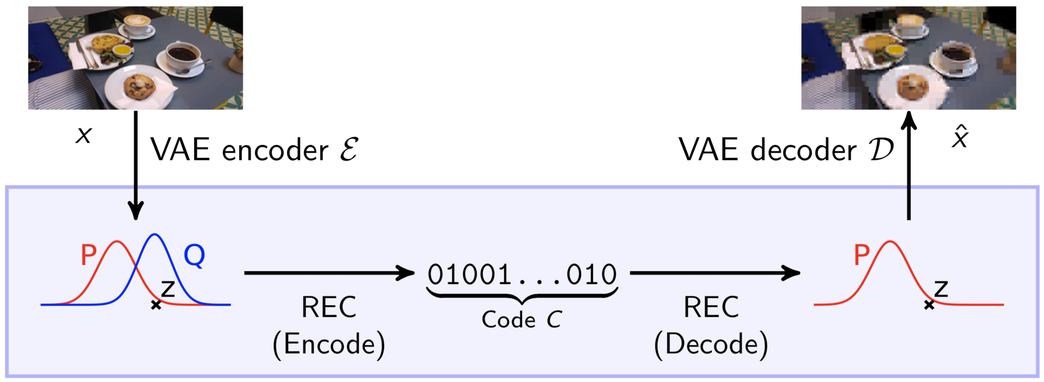

Challenge and our Solution

Runtime of general REC (Agustsson and Theis, 2020)

Without further assumptions, any REC scheme has $\Omega(\exp(D_{\mathrm{KL}}[Q\|P]))$ expected runtime.

Our Solution

Greedy Rejection Coding, an algorithm extending rejection sampling with partitions.

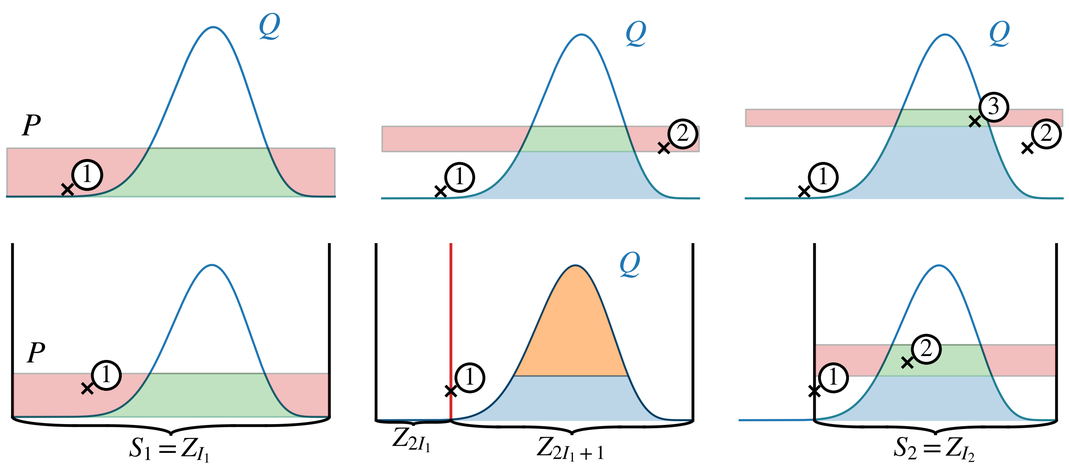

Theoretical results

Codelength of greedy rejection coding (informal)

Let C be the code returned by greedy rejection coding. Then

$$\mathbb{E}\left[|C|\right] = D_{\mathrm{KL}}[Q\|P] + 2\log(D_{\mathrm{KL}}[Q\|P] + 1) + \mathcal{O}(1).$$

Runtime of the on-sample splitting variant of GRC (informal)

For P, Q over $\mathbb{R}$ with unimodal q/p, the expected runtime of GRCS is

$$\mathbb{E}\left[T\right] = \mathcal{O}(D_{\mathrm{KL}}[Q\|P]).$$

Improves on A* coding (previous best): optimal codelength and order-optimal runtime.

Experiments

Synthetic experiments

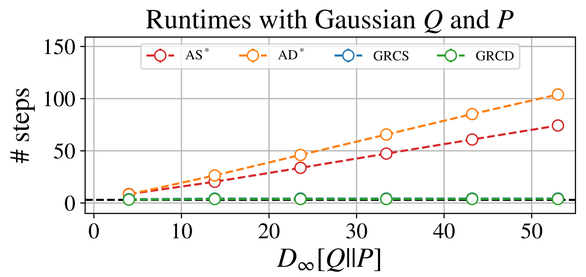

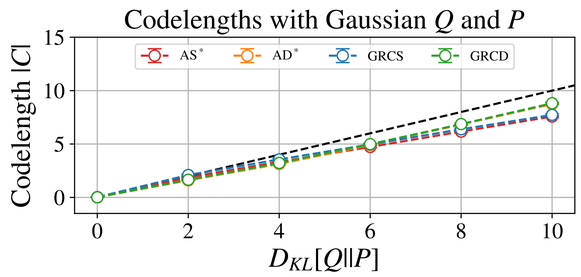

Compression with VAEs (MNIST)

# Latent	Neg. ELBO (bpd)	GRC (bpd)
20	1.391 ± 0.004	1.472 ± 0.004
50	1.357 ± 0.003	1.511 ± 0.003
100	1.362 ± 0.003	1.523 ± 0.003