Fast Relative Entropy Coding with A* 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 A* coding, the first relative entropy coding algorithm whose runtime is practically fast.

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.

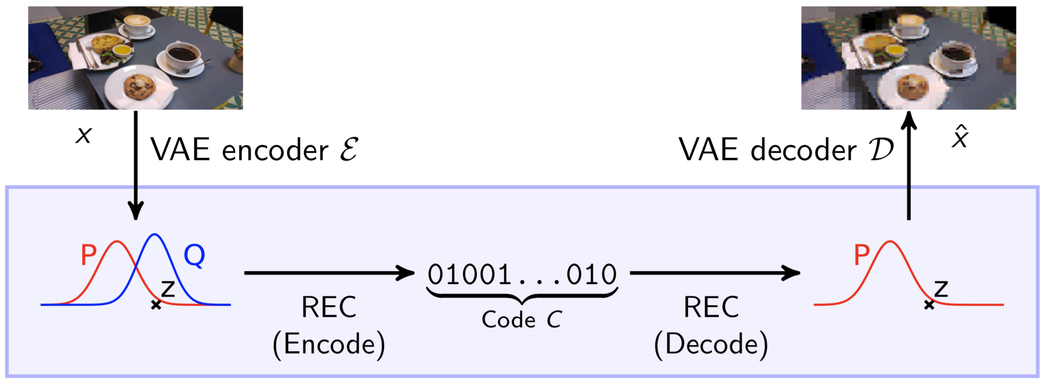

Challenge and our Solution

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

A* coding, a REC algorithm based on A* sampling, binary space partition trees.

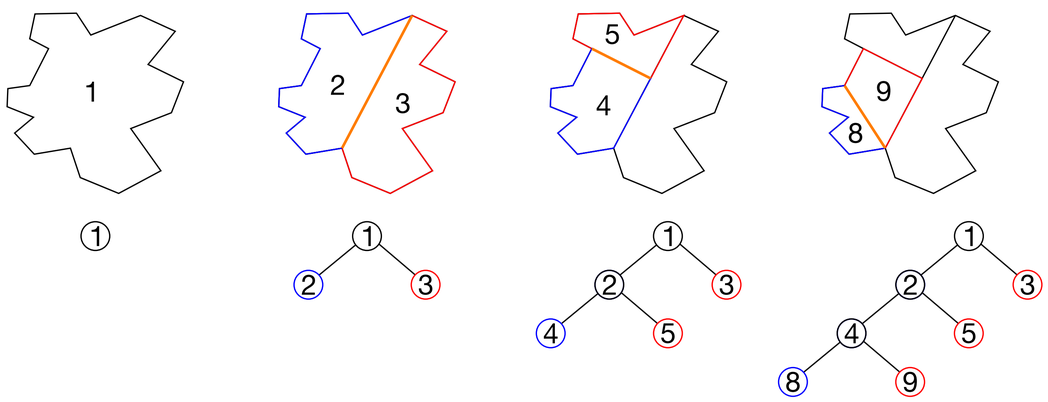

Theoretical results

Codelength of A* coding (informal)

Let C be the code returned by A* coding. Then

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

Runtime of AS* coding (informal)

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

$$\mathbb{E}\left[T\right] = \mathcal{O}(D_{\infty}[Q\|P]) = \mathcal{O}\left(\log \sup_{z \in \mathbb{R}} \frac{q(z)}{p(z)}\right).$$

Experiments

Synthetic experiments

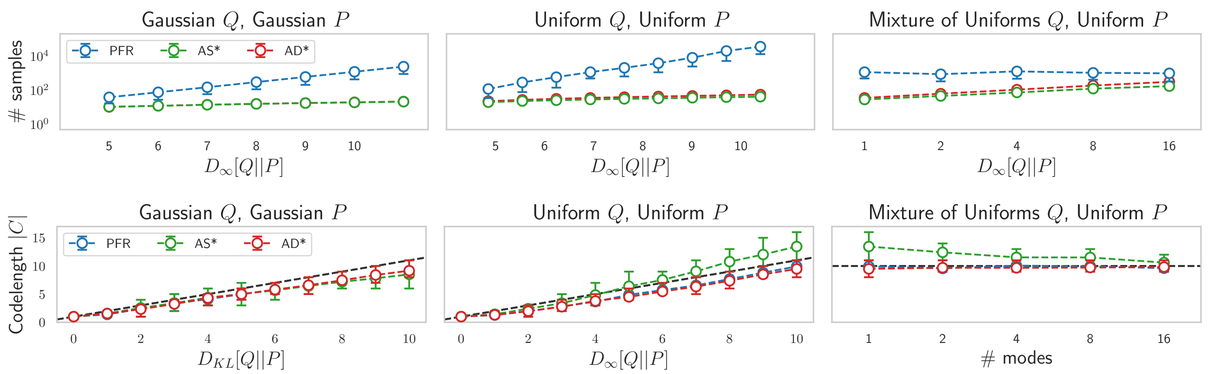

Compression with VAEs (MNIST)

# Latent	Neg. ELBO	A* Coding
20	1.43 ± 0.01	1.53 ± 0.01
50	1.40 ± 0.01	1.66 ± 0.01