

# Variational Russian Roulette for Deep Bayesian Nonparametrics

Kai Xu <sup>[1]</sup>

Joint work with Akash Srivastava <sup>[1,2]</sup> and Charles Sutton <sup>[1,3,4]</sup>

[1] University of Edinburgh

[2] MIT-IBM Watson AI Lab

[3] Google AI

[4] Alan Turing Institute

# tl;dr

- We train a **variational autoencoder** with **unbounded latent dimension**.
- The latent dimension is controlled by a **sparse binary matrix** with **infinitely many columns**, following an **Indian buffet process**.
- The actual dimensionality of the VAE is **inferred during training**.

# How an infinite binary matrix is useful for a VAE?

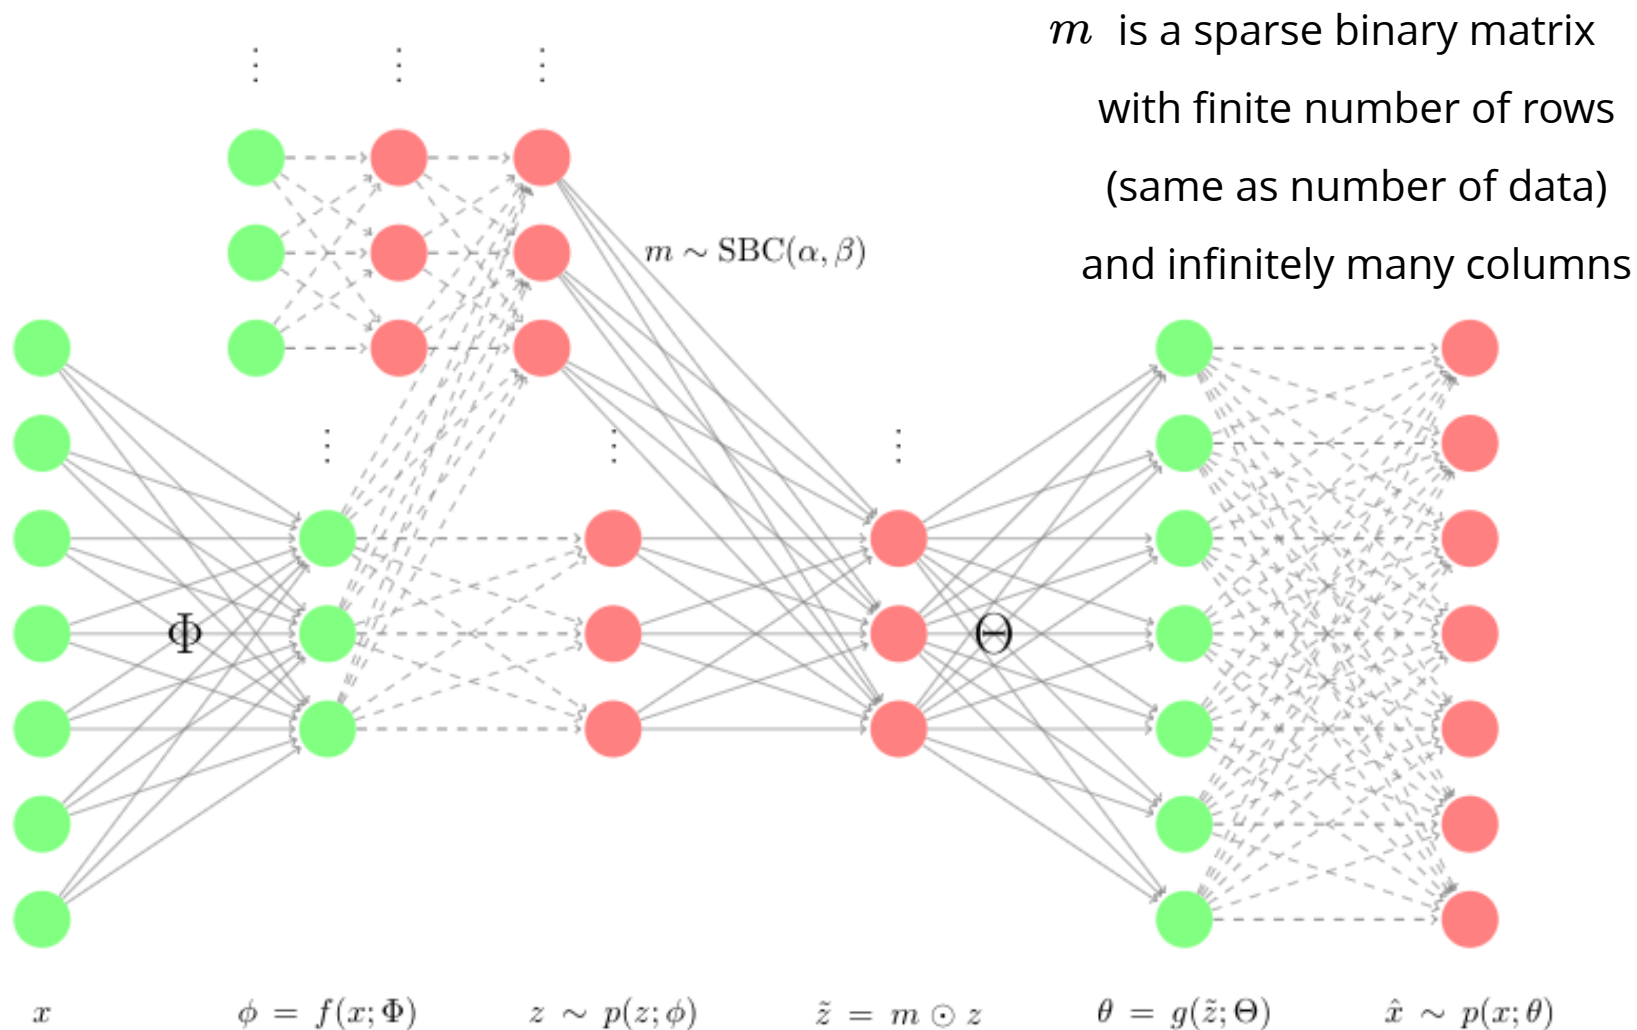


Fig. 1 Infinite VAE with an IBP prior (Chatzis, 2014; Singh et al. 2017)

# Truncation-free variational approximation

Previous work uses truncated variational approximation.  
Our method avoids using truncated approximation.

**Why** a truncated variational approximation is not ideal?

1. Truncation level is not easy to choose.
2. Poor interaction with amortised inference.

**How** do we avoid truncating variational posterior at all?

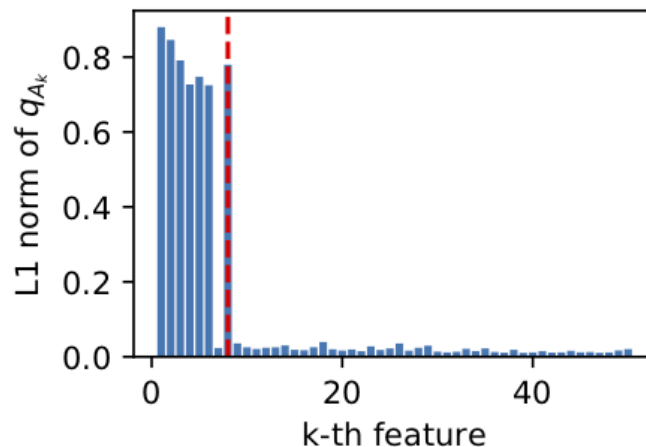
# RAVE: Roulette-based Amortized Variational Expectations

1. Introduce a new infinite variational approximation
  - Essentially an infinite mixture of truncated approximations
2. Derive a new tractable ELBO
  - Essentially an infinite mixture of truncated ELBOs
3. Compute an unbiased gradient estimate of the ELBO
  - Infinite summation is estimated by Russian roulette sampling


At any time, we only retain a **finite representation** in memory to compute the **unbiased gradient** estimate of the **infinite target**.

# Results

- Truncated approximation tends to activate collapsed component
  - Dimensions convey no information
- Russian roulette (marked by red vertical lines) automatically truncates at right dimension



Only first few components are informative



Non-informative components are still activated

Please come to check my poster @Pacific Ballroom #223