Learning Deep Latent Gaussian Models with Markov Chain Monte Carlo: Additional Figures

Anonymous Authors[1]

Figure 1. Samples from the VAE trained using mean-field variational inference.

*Figure 2.* Samples from the DLGM trained with 20 steps of HMC.

Figure 3. Effects of keeping only the first C most important dimensions of z, marginalizing out the remaining dimensions. z vectors are randomly sampled from $\mathcal{N}(0, I)$.