
Undirected Graphical Models as Approximate Posteriors

Arash Vahdat¹ Evgeny Andriyash² William G. Macready²

Abstract

The representation of the approximate posterior is a critical aspect of effective variational autoencoders (VAEs). Poor choices for the approximate posterior have a detrimental impact on the generative performance of VAEs due to the mismatch with the true posterior. We extend the class of posterior models that may be learned by using undirected graphical models. We develop an efficient method to train undirected approximate posteriors by showing that the gradient of the training objective with respect to the parameters of the undirected posterior can be computed by backpropagation through Markov chain Monte Carlo updates. We apply these gradient estimators for training discrete VAEs with Boltzmann machines as approximate posteriors and demonstrate that undirected models outperform previous results obtained using directed graphical models. Our implementation is available [here](#).

1. Introduction

Training of likelihood-based deep generative models has advanced rapidly in recent years. These advances have been enabled by amortized inference (Hinton et al., 1995; Mnih & Gregor, 2014; Gregor et al., 2013), which scales up training of variational models, the reparameterization trick (Kingma & Welling, 2014; Rezende et al., 2014), which provides low-variance gradient estimates, and increasingly expressive neural networks. Combinations of these techniques have resulted in many different forms of variational autoencoders (VAEs) (Kingma et al., 2016; Chen et al., 2016).

It is also widely recognized that flexible approximate posterior distributions improve the generative quality of variational autoencoders (VAEs) by more faithfully modeling true posterior distributions. More accurate approxi-

mate posterior models have been realized by reducing the gap between true and approximate posteriors with tighter bounds (Burda et al., 2015), and using auto-regressive architectures (Gregor et al., 2013; 2015) or flow models (Rezende & Mohamed, 2015).

A complementary (but unexplored) direction for further improving the richness of approximate posteriors is available through the use of undirected graphical models (UGMs). In this approach, the approximate posterior over latent variables is formed using UGMs whose parameters are inferred for each observed sample in an amortized fashion similar to VAEs. This is compelling as UGMs can succinctly capture complex relationships between variables. However, there are three challenges when training UGMs as approximate posteriors: i) There is no known low-variance path-wise gradient estimator (like the reparameterization trick) for learning parameters of UGMs. ii) Sampling from general UGMs is intractable and approximate sampling is computationally expensive. iii) Evaluating the probability of a sample under a UGM can require an intractable partition function. These two latter costs are particularly acute when UGMs are used as posterior approximators because the number of UGMs grows with the size of the dataset.

However, we note that posterior¹ UGMs conditioned on observed data points are often simple as there are usually only a small number of modes in latent space explaining an observation. We expect then, that sampling and partition function estimation (challenges ii and iii) can be solved efficiently using parallelizable Markov chain Monte Carlo (MCMC) methods. In fact, we observe that UGM posteriors trained in a VAE model tend to be unimodal but with a mode structure that is not necessarily well-captured by mean-field posteriors. Nevertheless, in this case, MCMC-based sampling methods quickly equilibrate.

To address challenge i) we estimate the gradient by reparameterized sampling in the MCMC updates. Although an infinite number of MCMC updates are theoretically required to obtain an unbiased gradient estimator, we observe that a single MCMC update provides a low-variance but biased gradient estimation that is sufficient for training UGMs.

¹We may use “posterior” when referring to “approximate posterior” for brevity. To disambiguate, we always refer to “true posterior” explicitly.

¹NVIDIA, USA ²Sanctuary AI, Canada, Work done at D-Wave Systems. Correspondence to: Arash Vahdat <avahdat@nvidia.com>.

Binary UGMs in the form of Boltzmann machines (Ackley et al., 1985) have recently been shown to be effective as priors for VAEs allowing for discrete versions of VAEs (Rolfe, 2016; Vahdat et al., 2018b;a). However, previous work in this area has relied on directed graphical models (DGMs) for posterior approximation. In this paper, we replace the DGM posteriors of discrete VAEs (DVAEs) with Boltzmann machine UGMs and show that these posteriors provide a generative performance comparable to or better than previous DVAEs. We denote this model as DVAE##, where ## indicates the use of UGMs in the prior and the posterior.

We begin by summarizing related work on developing expressive posterior models including models trained nonvariationally. Nonvariational models employ MCMC to directly sample posteriors and differ from our variational approach which uses MCMC to sample from an amortized undirected posterior. Sec. 2 provides the necessary background on variational learning and MCMC to allow for the development in Sec. 3 of a gradient estimator for undirected posteriors. We provide examples for both Gaussian (Sec. 3.1) and Boltzmann machine (Sec. 3.2) UGMs. Experimental results are provided in Sec. 4 on VAEs (Sec. 4.1), importance-weighted VAEs (Sec. 4.2), and structured prediction (Sec. 4.3), where we observe consistent improvement using UGMs. We conclude in Sec. 5 with a list of future work.

1.1. Related Work

Inference gap reduction in VAEs: Previous work on reducing the gap between true and approximate posteriors can be grouped into three categories: i) Training objectives that replace the variational bound with tighter bounds (Burda et al., 2015; Li & Turner, 2016; Bornschein et al., 2016). ii) Autoregressive models (Hochreiter & Schmidhuber, 1997; Graves, 2013) that use DGMs to form more flexible distributions (Gregor et al., 2013; Gulrajani et al., 2016; Van Den Oord et al., 2016; Salimans et al., 2017; Chen et al., 2018). iii) Flow-based models (Rezende & Mohamed, 2015) that map data to a latent space using a class of invertible functions (Kingma et al., 2016; Kingma & Dhariwal, 2018; Dinh et al., 2014; 2016). To the best of our knowledge, UGMs have not been used as approximate posteriors.

Gradient estimation in latent variable models: REINFORCE (Williams, 1992) is the most generic approach for computing the gradient of the approximate posteriors but suffers from high-variance and must be augmented by variance reduction techniques. For many continuous latent-variable models the reparameterization trick (Kingma & Welling, 2014; Rezende et al., 2014) provides lower-variance gradient estimates. Reparameterization does not apply to discrete latent variables and recent methods for discrete variables have focused on REINFORCE with control variates (Mnih & Gregor, 2014; Gu et al., 2015; Mnih &

Rezende, 2016; Tucker et al., 2017; Grathwohl et al., 2018) or continuous relaxations (Maddison et al., 2016; Jang et al., 2016; Rolfe, 2016; Vahdat et al., 2018b;a). See (Andriyash et al., 2018) for a review.

Variational Inference and MCMC: A common alternative to variational inference is to use MCMC to sample directly from true posteriors in generative models (Welling & Teh, 2011; Salimans et al., 2015; Wolf et al., 2016; Hoffman, 2017; Li et al., 2017; Caterini et al., 2018). However, this approach is often computationally intensive, as it requires computing the prior and decoder distributions (often implemented by neural networks) many times at each parameter update for each training data point in a batch. Moreover, evaluating these models is more challenging, as there is no amortized approximating posterior for importance sampling (Burda et al., 2015). Our method differs from these techniques, as we use MCMC to sample from an amortized approximate posterior represented by UGMs.

2. Background

In this section we provide background for the topics discussed in this paper.

Undirected graphical models: A UGM represents the joint probability distribution for a set of random variables \mathbf{z} as $q(\mathbf{z}) = \exp(-E_\phi(\mathbf{z}))/Z_\phi$, where E_ϕ is a ϕ -parameterized energy function defined by an undirected graphical model, and $Z_\phi = \int d\mathbf{z} \exp(-E_\phi(\mathbf{z}))$ is the partition function. UGMs over binary variables \mathbf{z} with bipartite quadratic energy functions $E_\phi(\mathbf{z}_1, \mathbf{z}_2) = \mathbf{b}_1^T \mathbf{z}_1 + \mathbf{b}_2^T \mathbf{z}_2 + \mathbf{z}_1^T \mathbf{W} \mathbf{z}_2$ are called restricted Boltzmann machines (RBMs). Parameters $\phi = \{\mathbf{b}_1, \mathbf{b}_2, \mathbf{W}\}$ encode linear biases and pairwise interactions. RBMs permit parallel Gibbs sampling updates as $q_\phi(\mathbf{z}_1|\mathbf{z}_2)$ and $q_\phi(\mathbf{z}_2|\mathbf{z}_1)$ are factorial in the components of \mathbf{z}_1 and \mathbf{z}_2 .

Variational autoencoders: A VAE is a generative model factored as $p(\mathbf{x}, \mathbf{z}) = p(\mathbf{z})p(\mathbf{x}|\mathbf{z})$, where $p(\mathbf{z})$ is a prior distribution over latent variables \mathbf{z} and $p(\mathbf{x}|\mathbf{z})$ is a probabilistic decoder representing the conditional distribution over data variables \mathbf{x} given \mathbf{z} . VAEs are trained by maximizing a variational lower bound (ELBO) on $\log p(\mathbf{x})$:

$$\mathcal{L} = \mathbb{E}_{q(\mathbf{z}|\mathbf{x})} \left[\log \frac{p(\mathbf{x}, \mathbf{z})}{q(\mathbf{z}|\mathbf{x})} \right] \leq \log p(\mathbf{x}), \quad (1)$$

where $q(\mathbf{z}|\mathbf{x})$ is a probabilistic encoder that approximates the posterior over latent variables given a data point. A tighter importance-weighted bound is obtained with

$$\mathcal{L}_{IW} = \mathbb{E}_{\mathbf{z}_{1:K}} \left[\log \left(\frac{1}{K} \sum_{i=1}^K \frac{p(\mathbf{z}_i, \mathbf{x})}{q(\mathbf{z}_i|\mathbf{x})} \right) \right] \leq \log p(\mathbf{x}), \quad (2)$$

where $\mathbf{z}_{1:K} \sim \prod_i q(\mathbf{z}_i|\mathbf{x})$. For continuous latent variables, these bounds are optimized using the reparameterization

trick (Kingma & Welling, 2014; Rezende et al., 2014). However, as noted above, continuous relaxations are commonly employed with discrete latent variables.

MCMC: MCMC methods are used to draw approximate samples from a target distribution $q(\mathbf{z})$. Each MCMC method is characterized by a transition kernel $\mathcal{K}(\mathbf{z}|\mathbf{z}')$ designed so that

$$q(\mathbf{z}) = \int d\mathbf{z}' q(\mathbf{z}') \mathcal{K}(\mathbf{z}|\mathbf{z}') \quad \forall \mathbf{z}. \quad (3)$$

Samples from the fixed point $q(\mathbf{z})$ are found by sampling from an initial distribution $q_0(\mathbf{z})$ and iteratively updating the samples by drawing conditional samples using $\mathbf{z}_t|\mathbf{z}_{t-1} \sim \mathcal{K}(\mathbf{z}_t|\mathbf{z}_{t-1})$. Denoting the distribution of the samples after t iterations by $q_t = \mathcal{K}^t q_0$, the theory of MCMC shows that under some regularity conditions $q_t \rightarrow q$ as $t \rightarrow \infty$. In practice, $\mathcal{K}(\mathbf{z}|\mathbf{z}')$ is usually chosen to satisfy the detailed balance condition:

$$q(\mathbf{z}) \mathcal{K}(\mathbf{z}'|\mathbf{z}) = q(\mathbf{z}') \mathcal{K}(\mathbf{z}|\mathbf{z}') \quad \forall \mathbf{z}, \mathbf{z}' \quad (4)$$

which implies Eq. (3). \mathcal{K}^t itself is a valid transition kernel and satisfies Eq. (3) and Eq. (4).

Gibbs sampling is a particularly efficient MCMC method when variables $\mathbf{z} = [\mathbf{z}_1, \mathbf{z}_2]$ are partitioned as in an RBM. The transition kernel in this case may be written as

$$\mathcal{K}_{\text{Gibbs}}(\mathbf{z}_1, \mathbf{z}_2|\mathbf{z}'_1, \mathbf{z}'_2) = q(\mathbf{z}_2|\mathbf{z}_1) q(\mathbf{z}_1|\mathbf{z}'_2). \quad (5)$$

where all partitioned variables in \mathbf{z}_1 or \mathbf{z}_2 may be updated in parallel. Gibbs sampling is a special MCMC method that does not satisfy the detailed balance condition. However, it does admit a reverse kernel in the form:

$$\mathcal{K}_{\text{Reverse}}(\mathbf{z}'_1, \mathbf{z}'_2|\mathbf{z}_1, \mathbf{z}_2) = q(\mathbf{z}'_1|\mathbf{z}'_2) q(\mathbf{z}'_2|\mathbf{z}_1), \quad (6)$$

which samples from the variables in a reversed order. Thus, for Gibbs sampling on bipartite UGMs, we have:

$$q(\mathbf{z}') \mathcal{K}_{\text{Gibbs}}(\mathbf{z}|\mathbf{z}') = q(\mathbf{z}) \mathcal{K}_{\text{Reverse}}(\mathbf{z}'|\mathbf{z}) \quad \forall \mathbf{z}, \mathbf{z}', \quad (7)$$

where $\mathbf{z} = [\mathbf{z}_1, \mathbf{z}_2]$ and $\mathbf{z}' = [\mathbf{z}'_1, \mathbf{z}'_2]$.

3. Undirected Approximate Posteriors

In this section, we propose a gradient estimator for generic UGM approximate posteriors and discuss how the estimator is applied to RBM-based approximate posteriors.

The training objective for a probabilistic encoder network with $q_\phi(\mathbf{z})$ being a UGM can be written as $\max_{\phi} \mathbb{E}_{q_\phi(\mathbf{z})}[f(\mathbf{z})]$, where f is a differentiable function of \mathbf{z} . For simplicity of exposition, we assume that f does not depend on ϕ .² Using the fixed point equation in Eq. (3), we

²If f does depend on ϕ , then $\mathbb{E}_{q_\phi(\mathbf{z})}[\partial_\phi f(\mathbf{z})]$ is approximated with Monte Carlo sampling.




Figure 1: To estimate the gradient of $\mathcal{L} = \mathbb{E}_{q_\phi(\mathbf{z})}[f(\mathbf{z})]$ w.r.t the parameters of the UGM posterior (ϕ), we first sample from the approximating posterior ($\mathbf{z}' \sim q_\phi(\mathbf{z}')$), then, we apply an MCMC update via reparameterized sampling ($\mathbf{z} \sim K_\phi(\mathbf{z}|\mathbf{z}')$). Finally, we evaluate the function on the samples ($f(\mathbf{z})$). The gradient is computed automatically by **backpropagating** through the reparameterized samples while ignoring the dependency of \mathbf{z}' on ϕ .

have:

$$\mathbb{E}_{q_\phi(\mathbf{z})}[f(\mathbf{z})] = \mathbb{E}_{q_\phi(\mathbf{z}')} \left[\mathbb{E}_{\mathcal{K}_\phi^t(\mathbf{z}|\mathbf{z}')} [f(\mathbf{z})] \right], \quad (8)$$

where the right hand side implies sampling $\mathbf{z}' \sim q_\phi(\mathbf{z}')$ and applying MCMC updates t times. To maximize Eq. (8), we require its gradient:

$$\begin{aligned} \partial_\phi \mathbb{E}_{q_\phi(\mathbf{z})}[f(\mathbf{z})] &= \underbrace{\int d\mathbf{z}' \partial_\phi q_\phi(\mathbf{z}') \mathbb{E}_{\mathcal{K}_\phi^t(\mathbf{z}|\mathbf{z}')} [f(\mathbf{z})]}_I \\ &\quad + \underbrace{\mathbb{E}_{q_\phi(\mathbf{z}')} \left[\partial_\phi \mathbb{E}_{\mathcal{K}_\phi^t(\mathbf{z}|\mathbf{z}')} [f(\mathbf{z})] \right]}_{II}. \end{aligned} \quad (9)$$

The term marked as I is written as:

$$I = \mathbb{E}_{q_\phi(\mathbf{z}')} \mathcal{K}_\phi^t(\mathbf{z}|\mathbf{z}') [f(\mathbf{z}) \partial_\phi \log q_\phi(\mathbf{z}')] \quad (10)$$

$$= \mathbb{E}_{q_\phi(\mathbf{z})} \mathcal{K}_\phi^t(\mathbf{z}'|\mathbf{z}) [f(\mathbf{z}) \partial_\phi \log q_\phi(\mathbf{z}')] \quad (11)$$

$$= \mathbb{E}_{q_\phi(\mathbf{z})} \left[f(\mathbf{z}) \mathbb{E}_{\mathcal{K}_\phi^t(\mathbf{z}'|\mathbf{z})} [\partial_\phi \log q_\phi(\mathbf{z}')] \right], \quad (12)$$

where we have used detailed balance to get Eq. (11) from Eq. (10). The form of Eq. (12) makes it clear that I goes to zero as $t \rightarrow \infty$, because $\mathcal{K}_\phi^t(\mathbf{z}'|\mathbf{z}) \rightarrow q_\phi(\mathbf{z}')$ and $\mathbb{E}_{q_\phi(\mathbf{z}')} [\partial_\phi \log q_\phi(\mathbf{z}')] = 0$. However, similar to REINFORCE, Monte-Carlo estimate of I will have high variance due to the presence of $q_\phi(\mathbf{z}')$ in the denominator³.

The expectation in Eq. (9) marked as II involves the gradient of an expectation with respect to MCMC transition kernel. The reparameterization trick provides a low-variance gradient estimator for this term, and as $t \rightarrow \infty$, the II contribution approaches the full gradient because $I \rightarrow 0$.

Our key insight in this paper is that, given the high variance of I , it is beneficial to drop this term from the gradient (Eq. (9)) and use the biased but lower-variance estimate II .

³Technically, Gibbs sampling does not satisfy detailed balance. Instead, we can use the identity in Eq. (7) to derive the same argument for Gibbs sampling. In this case, the transition kernel in Eq. (11) and Eq. (12) will be the reverse kernel.

The choice of t trades increased computational complexity for decreased bias. We observe that increasing t has little effect on the optimization performance. Therefore, $t = 1$ is a good choice giving the smallest computational complexity. Consequently, we use the approximation:

$$\begin{aligned} \partial_{\phi} \mathbb{E}_{q_{\phi}(\mathbf{z})}[f(\mathbf{z})] &\approx \mathbb{E}_{q_{\phi}(\mathbf{z}')} \left[\partial_{\phi} \mathbb{E}_{\mathcal{K}_{\phi}(\mathbf{z}|\mathbf{z}')} [f(\mathbf{z})] \right] \quad (13) \\ &= \mathbb{E}_{q_{\phi}(\mathbf{z}')} \mathbb{E}_{\epsilon \sim p(\epsilon)} \left[(\partial_{\phi} \mathbf{z}) \partial_{\mathbf{z}} f(\mathbf{z}) \right], \end{aligned}$$

where $\mathbf{z} = \mathbf{z}(\epsilon, \phi, \mathbf{z}')$ is a reparameterized sample from $\mathcal{K}_{\phi}(\mathbf{z}|\mathbf{z}')$ where ϵ is a sample from a base distribution. Our gradient estimator is illustrated in Fig. 1. The extension to $t > 1$ Gibbs updates is straightforward.

In principle, MCMC methods such as Metropolis-Hastings or Hamiltonian Monte Carlo (HMC) can be used as the transition kernel in Eq. (13). However, unbiased reparameterized sampling for these methods is challenging.⁴ These complications are avoided when Gibbs sampling of $q_{\phi}(\mathbf{z})$ is possible. The Gibbs transition kernel $q_{\phi}(\mathbf{z}_2|\mathbf{z}_1)q_{\phi}(\mathbf{z}_1|\mathbf{z}'_2)$ of Eq. (5) does not contain an accept/reject step or any hyper-parameters. Assuming that the conditionals can be reparameterized, the gradient of the kernel is approximated efficiently using low-variance reparameterized sampling from each conditional. In this case, the gradient estimator for a single Gibbs update is:

$$\begin{aligned} \partial_{\phi} \mathbb{E}_q[f(\mathbf{z})] &\approx \mathbb{E}_{q_{\phi}(\mathbf{z}')} \left[\partial_{\phi} \mathbb{E}_{q_{\phi}(\mathbf{z}_2|\mathbf{z}_1)q_{\phi}(\mathbf{z}_1|\mathbf{z}'_2)} [f(\mathbf{z})] \right] \quad (14) \\ &= \mathbb{E}_{q_{\phi}(\mathbf{z}')} \mathbb{E}_{\epsilon_1, \epsilon_2} \left[\partial_{\phi} f(\mathbf{z}_1, \mathbf{z}_2) \right], \end{aligned}$$

where $\mathbf{z}_1(\epsilon_1, \phi, \mathbf{z}'_2) \sim q_{\phi}(\mathbf{z}_1|\mathbf{z}'_2)$ and $\mathbf{z}_2(\epsilon_2, \phi, \mathbf{z}_1) \sim q_{\phi}(\mathbf{z}_2|\mathbf{z}_1)$ are reparameterized samples.

Lastly, we note that the reparameterized \mathbf{z} sample has distribution $q_{\phi}(\mathbf{z})$ if \mathbf{z}' is equilibrated. So, if f has its own parameters (e.g., the parameters of the generative model in VAEs), the same sample can be used to compute an unbiased estimation of $\partial_{\theta} \mathbb{E}_{q_{\phi}(\mathbf{z})}[f_{\theta}(\mathbf{z})]$ where θ denotes the parameters of f .

3.1. Toy Example

We assess the efficacy of our approximations for a multivariate Gaussian. We consider learning the two-variable Gaussian distribution depicted in Fig. 2(b). The target distribution $p(\mathbf{z})$ has energy function $E(\mathbf{z}) = (\mathbf{z} - \boldsymbol{\mu})^T \boldsymbol{\Lambda} (\mathbf{z} - \boldsymbol{\mu})/2$, where $\boldsymbol{\mu} = [1, 1]^T$ and $\boldsymbol{\Lambda} = [1.1, 0.9; 0.9, 1.1]$. The approximating distribution $q_{\phi}(\mathbf{z})$ has the same form and we learn the parameters ϕ by minimizing the Kullback-Leibler (KL) divergence $\text{KL}[q_{\phi}(\mathbf{z})||p(\mathbf{z})]$. We compare Eq. (14) for $t = 1, 2, 4, 8$ with reparameterized sampling and with

⁴For example, Metropolis-Hastings contains a nondifferentiable operation in the accept/reject step, and HMC additionally requires backpropagating through gradients of the energy function and tuning hyper-parameters.

REINFORCE. Fig. 2(a) shows the KL divergence during training. For our method we consider two variants including with and without I in Eq. (9).

Our method significantly outperforms REINFORCE due to lower variance of the gradients as depicted in Fig. 2(c). There is little difference between different t until KL divergence becomes $\sim 10^{-4}$. Our method performs worse than reparameterized sampling, which is expected due to the bias introduced by neglecting I in Eq. (9). Including I negatively impacts optimization. The dashed lines of Fig. 2(a) show the KL objective when the I term is included. All curves lie atop one another and are noticeably slower to converge. This deterioration is driven by the noisier gradient estimates shown as dashed lines in Fig. 2(c).

Note that this experiment confirms that the reparameterization trick provides low-variance unbiased gradient estimation. However, this trick is not applicable to UGMs in general. In the case of UGMs, our gradient estimator provides gradient estimation with variance properties similar to the reparameterization trick.

3.2. Learning Boltzmann Machine Posteriors

Next, we consider training DVAE##, a DVAE (Rolfe, 2016; Vahdat et al., 2018b;a) where both approximating posterior and prior are RBM. The objective function of DVAE## is given by Eq. (1), where $q(\mathbf{z}|\mathbf{x}) \equiv q_{\phi(\mathbf{x})}(\mathbf{z})$ is an amortized RBM with neural-network-generated parameters $\phi(\mathbf{x}) = \{\mathbf{b}_1(\mathbf{x}), \mathbf{b}_2(\mathbf{x}), \mathbf{W}(\mathbf{x})\}$. To maximize \mathcal{L} , we use the gradient

$$\partial_{\phi(\mathbf{x})} \mathcal{L} = \partial_{\phi(\mathbf{x})} \mathbb{E}_{q_{\phi(\mathbf{x})}(\mathbf{z})} \left[\log \frac{p(\mathbf{x}, \mathbf{z})}{q_{\phi(\mathbf{x})}(\mathbf{z})} \right]. \quad (15)$$

at each training point \mathbf{x} . Although the RBM has a bipartite structure, applying the gradient estimator of Eq. (14) is challenging due to binary latent variables. To apply our estimator in Eq. (14), we relax the binary variables to continuous variables using Gumbel-Softmax or piece-wise linear relaxations (Andriyash et al., 2018).⁵ Representing the relaxations of $\mathbf{z}_1 \sim q_{\phi}(\mathbf{z}_1|\mathbf{z}'_2)$ by $\zeta_1 = \zeta_1(\epsilon_1, \mathbf{z}'_2, \phi)$ and $\mathbf{z}_2 \sim q_{\phi}(\mathbf{z}_2|\mathbf{z}_1)$ by $\zeta_2 = \zeta_2(\epsilon_2, \zeta_1, \phi)$, where $\epsilon_1, \epsilon_2 \sim \mathcal{U}[0, 1]$, we use the following estimator:

$$\begin{aligned} &\mathbb{E}_{q_{\phi}(\mathbf{z}'_1, \mathbf{z}'_2)} \left[\partial_{\phi} \mathbb{E}_{q_{\phi}(\zeta_1|\mathbf{z}'_2)} \mathbb{E}_{q_{\phi}(\zeta_2|\zeta_1)} [f(\zeta_1, \zeta_2)] \right] \\ &= \mathbb{E}_{q_{\phi}(\mathbf{z}'_1, \mathbf{z}'_2)} \mathbb{E}_{\epsilon_1, \epsilon_2} \left[\partial_{\phi} f(\zeta_1, \zeta_2) \right], \quad (16) \end{aligned}$$

$\partial_{\phi} f(\zeta_1, \zeta_2)$ is computed using the reparameterization trick.

Thus far, we have only accounted for the dependency of the objective f on ϕ through samples ζ . However, for VAEs, f also depends on ϕ through $\log q_{\phi(\mathbf{x})}(\mathbf{z})$. This dependency

⁵Another option is to use unbiased gradient estimators such as REBAR (Tucker et al., 2017) or RELAX (Grathwohl et al., 2018). However, with these approaches, the number of f evaluations increases with t .




Figure 2: (a) Our gradient estimator (for various t) compared with REINFORCE and reparameterized Gaussian samples for minimizing the KL divergence of a Gaussian to a target distribution (b). Dashed lines correspond to adding the term I of Eq. (9) to our gradient estimator.

can be ignored in VAEs as $\mathbb{E}_q[\partial_\phi \log q_\phi(\mathbf{z}|\mathbf{x})] = 0$ (Roeder et al., 2017), and it can be removed in importance weighted autoencoders using the doubly reparameterized gradient estimation (Tucker et al., 2018) (see Appendix C for more details). Note that Roeder et al. (2017) and Tucker et al. (2018) do not introduce any bias to the gradient and are known to reduce variance.

As t increases and each Gibbs update is relaxed, sampling from the relaxed chain diverges from the exact discrete chain resulting in increasingly biased gradient estimates. Thus, we use $t = 1$ in our experiments (see Appendices A and D for $t > 1$). Moreover, in Eq. (16), our estimator requires samples from the RBM posterior in the outer expectation. These samples are obtained by running persistent chains for each training datapoint.

Algorithm 1 summarizes training of DVAE##. We represent the Boltzmann prior using $p_\theta(\mathbf{z}) = \exp(-E_\theta(\mathbf{z}))/Z_\theta$. The number of Gibbs sweeps for generating \mathbf{z}' in Eq. (16) is denoted by s and the number of Gibbs sweeps for sampling ζ is denoted by t . The objective \mathcal{L} is defined so that automatic differentiation yields the gradient estimator in Eq. (16) for ϕ and $\partial_\theta f$ is evaluated using relaxed samples for θ . We use the method introduced by Vahdat et al. (2018b) Sec. E to obtain gradients of $\log Z_\theta$.

4. Experiments

We examine our proposed gradient estimator for training undirected posteriors on three tasks: variational autoencoders, importance weighted autoencoders, and structured prediction models using the binarized MNIST (Salakhutdinov & Murray, 2008) and OMNIGLOT (Lake et al., 2015) datasets. We follow the experimental setup in DVAE# (Vahdat et al., 2018a) with minor modifications outlined below.

4.1. Variational Autoencoders

We train a VAE of the form $p(\mathbf{z})p(\mathbf{x}|\mathbf{z})$, where $p(\mathbf{z})$ is an RBM and $p(\mathbf{x}|\mathbf{z})$ is a neural network. $p(\mathbf{x}|\mathbf{z})$ is represented

Algorithm 1 DVAE## with RBM prior and posterior

Input: training sample \mathbf{x} , number of Gibbs sweeps s , number of relaxed Gibbs sweeps t .

Output: training objective function $\mathcal{L}_\mathbf{x}$

$\phi = \text{encoder}(\mathbf{x})$

$\mathbf{z}'_{\text{old}} = \text{retrieve_persistent_states}(\mathbf{x})$

$\mathbf{z}'_{\text{new}} = \text{update_gibbs_samples}(\mathbf{z}'_{\text{old}}, \phi, s)$

$\mathbf{z}'_{sg} = \text{stop_gradient}(\mathbf{z}'_{\text{new}}) \triangleright \mathbf{z}'_1, \mathbf{z}'_2$ in Eq. 16

$\zeta = \text{relax_gibbs_sample}(\mathbf{z}'_{sg}, \phi, t) \triangleright \zeta_1, \zeta_2$ in Eq. 16

$\phi' = \text{stop_gradient}(\phi) \triangleright$ Roeder et al.

$\mathcal{L}_\mathbf{x} = \underbrace{-E_\theta(\zeta)}_{\text{prior}} - \underbrace{\log p_\theta(\mathbf{x}|\zeta)}_{\text{likelihood}} + \underbrace{E_{\phi'}(\zeta)}_{\text{approx. post}}$

using a fully-connected neural network having two 200-unit hidden layers, tanh activations, and batch normalization similar to DVAE++ (Vahdat et al., 2018b), DVAE# (Vahdat et al., 2018a), and GumBolt (Khoshaman & Amin, 2018). We compare generative models trained with an undirected posterior (DVAE##) to a directed posterior.

For DVAE##, $q(\mathbf{z}|\mathbf{x})$ is modeled with a neural network having two tanh hidden layers that predicts the parameters of the RBM (Fig. 3(a)). Training DVAE## is done using Algorithm 1 with $s = 10$ and $t = 1$ using a piece-wise linear relaxation (Andriyash et al., 2018) for relaxed Gibbs samples. We follow (Vahdat et al., 2018a) for batch size, learning rate schedule, and KL warm up parameters. During training, samples from the prior are required for computing the gradient of $\log Z_\theta$. This sampling is done using the QuPA library that offers population-annealing-based sampling and partition function estimation (using AIS) from within Tensorflow. We also use QuPA for sampling the undirected posteriors and estimating their partition function during evaluation. We explore VAEs with equally-sized RBM prior and posteriors consisting of either 200 (100+100) or 400 (200+200) latent variables. Test set negative log-likelihoods are computed using 4000 importance-weighted samples.

Baselines: We compare the RBM posteriors with directed posteriors where the posteriors are factored across groups of latent variables \mathbf{z}_i as $q(\mathbf{z}|\mathbf{x}) = \prod_i^L q_i(\mathbf{z}_i|\mathbf{x}, \mathbf{z}_{<i})$. Each

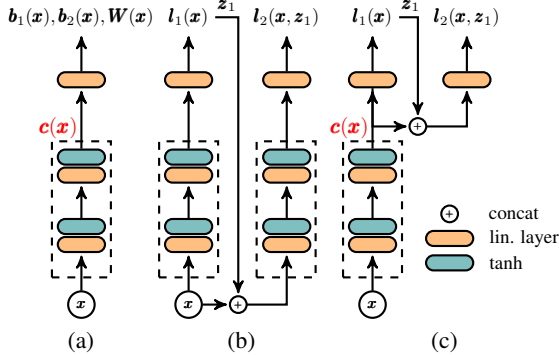


Figure 3: Neural networks representing $q(\mathbf{z}|\mathbf{x})$: (a) A 2-layer network predicts the parameters of RBM. (b) The directed posterior used in DVAE# consists of parallel 2-layer networks to successively predict, \mathbf{l} , the logits for each conditional in $q(\mathbf{z}|\mathbf{x}) = \prod_i q_i(z_i|\mathbf{x}, \mathbf{z}_{<i})$. (c) Our directed posterior differs from DVAE# and predicts the parameters of each conditional in $q(\mathbf{z}|\mathbf{x}) = \prod_i q_i(z_i|\mathbf{c}(\mathbf{x}), \mathbf{z}_{<i})$ using a linear transformation given the shared context feature $\mathbf{c}(\mathbf{x})$ and previous \mathbf{z} .

conditional, q_i , is factorial across the components of \mathbf{z}_i and L is the number of hierarchical layers. A fair comparison between directed and undirected posteriors is challenging because they differ in structure and in the number of parameters. We design a baseline so that the number of parameters and the number of nonlinearities is identical for a directed posterior with a single group ($L = 1$) and an undirected posterior with no pairwise interactions. This is reasonable as both cases reduce to a mean-field posterior. In Appendix B, we present a new structure for directed posteriors with shared context (Fig. 3(c)) that further improves the posteriors used in DVAE# (Vahdat et al., 2018a) and GumBolt (Khoshaman & Amin, 2018) (Fig. 3(b)).

We examine three recent methods for training VAEs with directed posteriors with shared context. These baselines are i) DVAE# (Vahdat et al., 2018a) that uses power-function distribution to relax the objective function, ii) Concrete relaxation used in GumBolt (Khoshaman & Amin, 2018), and iii) piece-wise linear (PWL) relaxation (Andriyash et al., 2018). All models are trained using $L = 1, 2$, or 4 hierarchical levels.

Results: The performance of DVAE## is compared against the baselines in Table 1. We make two observations. i) The baselines with $L = 1$ are equivalent to a mean-field posterior which is also the special case of an undirected posterior with no pairwise terms. Since PWL and DVAE## use the same continuous relaxation, we can compare PWL $L = 1$ and DVAE## to see how introducing pairwise interactions improve the quality of the DVAE models. We consistently observe ~ 0.5 -nat improvement arising from pairwise inter-




Figure 4: The difficulty of partition function estimation is visualized using the mean absolute difference between $\log Z$ estimates and its true value in top and the variance of the estimates for different number of temperatures in bottom. The true value is computed using 2^{18} temperatures. The number of temperatures required for achieving $\sigma = 10^{-2}$ for the RBM posterior is $\sim 10\times$ smaller than the number of temperatures required for the RBM prior.

actions. ii) DVAE## with undirected posteriors outperforms the directed baselines in most cases indicating that UGMs form more appropriate posteriors.

Varying the number of Gibbs steps: In Appendix D, we study the effect of changing s and t .

Using MCMC to sample from the true posterior: In Appendix E, we compare our DVAE## against MCMC methods that approximately sample from the true posterior.

4.1.1. EXAMINING UNDIRECTED POSTERiors

We also examine the posteriors trained with DVAE## on MNIST in terms of the number of modes and the difficulty of sampling and partition function estimation.

Multimodality: To estimate the number of modes, we find a variety of mean field solutions as follows: Given an RBM we draw samples using QuPA and use each sample to initialize the construction of a mean-field approximation. Specifically, we initialize the mean parameter of a factorial Bernoulli distribution to the sample value and then iteratively minimize the KL until we converge to a fixed point describing a mode. This way, the initial population of samples converges to a smaller number of unique modes.

We observe that the majority ($\sim 90\%$) of trained posteriors have a single mode. This is in contrast with the prior RBM that typically has 10 to 20 modes. However, we also note that the KL from the converged mean-field distributions to the RBM posteriors is typically in the range $[0.5, 2]$ indicating that, while most RBM posteriors are unimodal, the structure of the mode is not completely captured by a factorial distribution.

The difficulty of sampling and partition function estimation: Fig. 4 visualizes the deviation of log partition function estimations for an RBM posterior and prior for different numbers of AIS temperatures. As shown in the figure, the

Table 1: The performance of DVAE## is compared against directed posteriors, trained with the variational bound. Mean±standard deviation of the negative log-likelihood for 5 runs are reported. Boldface numbers indicate the best performing models per latent variable size and dataset.

		Prior RBM: 100+100				Prior RBM: 200+200			
		DVAE#	GumBolt	PWL	DVAE##	DVAE#	GumBolt	PWL	DVAE##
MNIST	1	84.97±0.03	84.66±0.04	84.62±0.05		83.21±0.03	83.19±0.05	83.22±0.05	
	2	84.96±0.05	84.71±0.03	84.50±0.07	84.09±0.06	83.13±0.04	83.04±0.02	82.99±0.05	82.75±0.05
	4	84.58±0.02	84.39±0.04	84.07±0.04		82.93±0.04	83.14±0.04	82.90±0.03	
OMNI.	1	101.64±0.05	101.41±0.06	101.24±0.02		99.51±0.03	99.39±0.04	99.32±0.04	
	2	101.75±0.04	101.39±0.06	101.14±0.05	100.69±0.04	99.40±0.05	99.12±0.07	99.10±0.04	98.61±0.08
	4	101.74±0.07	102.04±0.10	101.14±0.05		99.47±0.04	99.97±0.02	99.30±0.03	

Table 2: The performance of DVAE## is compared with directed posteriors with the PWL relaxation trained with the IW bound. Mean±standard deviation of the negative log-likelihood for 5 runs are reported. Boldface numbers indicate the best performing models per latent space size, dataset, and K .

		Prior RBM: 100+100				Prior RBM: 200+200			
		PWL $L=1$	PWL $L=2$	PWL $L=4$	DVAE##	PWL $L=1$	PWL $L=2$	PWL $L=4$	DVAE##
MNIST	1	84.60±0.04	84.49±0.03	84.04±0.03	84.06±0.03	83.27±0.06	82.99±0.04	82.89±0.03	82.76±0.05
	5	84.15±0.06	83.88±0.02	83.53±0.06	83.56±0.02	83.24±0.06	82.90±0.06	82.65±0.03	82.76±0.06
	25	83.96±0.05	83.70±0.03	83.38±0.02	83.44±0.04	83.38±0.04	83.05±0.03	82.84±0.02	82.97±0.11
OMNI.	1	101.21±0.06	101.14±0.07	101.12±0.03	100.68±0.05	99.32±0.02	99.11±0.02	99.28±0.09	98.61±0.06
	5	100.72±0.05	100.58±0.05	100.51±0.04	100.16±0.04	99.19±0.05	98.69±0.09	98.78±0.02	98.34±0.03
	25	100.58±0.01	100.48±0.05	100.37±0.03	100.02±0.04	99.10±0.03	98.70±0.07	98.86±0.06	98.34±0.04

number of temperatures required for achieving an acceptable precision (e.g., $\sigma = 10^{-2}$) for the RBM posterior is 10 times smaller than the number of temperatures required for the RBM prior. The main reason for this difference is that the RBM posteriors have strong linear biases. When annealing starts from a mean-field distribution containing only the biases, AIS requires fewer interpolating temperatures in order to accurately approximate the partition function. However, RBM priors which do not contain strong linear biases, do not benefit from this.

4.2. Importance Weighted Autoencoders

Next, we examine the generative model introduced in the previous section but using the importance-weighted (IW) bound. For comparison, we only use the PWL baseline as it achieves the best performance among directed posterior baselines in Table 1. For training, we use the same hyperparameters used in the previous section with an additional hyperparameter K representing the number of IW samples.

To sidestep computation of the gradient of the partition function for undirected posteriors in the IW bound, we use the path-wise gradient estimator introduced by (Tucker et al., 2018). However, we observe that an annealing scheme in IWAEs (similar to KL-warm up in VAEs) improves the performance of the generative model by preventing the latent variables from turning off. In Appendix C, we introduce this annealing mechanism for training IWAEs and show how the path-wise gradient can be computed while annealing the

objective function.

The experimental results for training DVAE## are compared against directed posteriors in Table 2. As can be seen, DVAE## achieves a comparable performance on MNIST, but outperforms the directed posteriors on OMNIGLOT.

4.3. Structured Output Prediction

Structured prediction is a form of conditional likelihood estimation (Sohn et al., 2015) concerned with modeling the distribution of a high-dimensional output given an input. Here, we predict the distribution of the bottom half \mathbf{x}_2 of an image given the top half \mathbf{x}_1 . For this, we follow the IW objective function proposed by (Raiko et al., 2015):

$$\mathbb{E}_{\mathbf{z}_1:K \sim \prod_i q(\mathbf{z}_i|\mathbf{x}_1)} \left[\log \left(\frac{1}{K} \sum_{i=1}^K p(\mathbf{x}_2|\mathbf{z}_i) \right) \right] + \lambda H(q(\mathbf{z}|\mathbf{x}_1)).$$

We have added an entropy term $H(q(\mathbf{z}|\mathbf{x}_1))$ to prevent the model from over-fitting the training data. This expectation (without the entropy term) is identical to the IW bound in Eq. (2), where the prior and the approximate posterior are both set to $q(\mathbf{z}|\mathbf{x}_1)$ and it can thus be considered as a lower bound on $\log p(\mathbf{x}_2|\mathbf{x}_1)$. The scalar λ is annealed during training from 1 to a small value (e.g., 0.05).

Experimental results for the structured prediction problem are reported in Table 3. The latent space is limited to 200 binary variables and $q(\mathbf{z}|\mathbf{x}_1)$ and $p(\mathbf{x}_2|\mathbf{z})$ are modeled by fully-connected networks with one 200-unit hidden layer

Table 3: The performance of DVAE## is compared against hierarchical posteriors with the PWL relaxation on the structured prediction problem.

		RBM Size: 100+100			
		PWL	PWL	PWL	DVAE##
		$L = 1$	$L = 2$	$L = 4$	
	K				
MNIST	1	60.82±0.17	59.54±0.12	60.22±0.12	57.13±0.18
	5	52.38±0.03	52.27±0.16	52.67±0.05	49.25±0.07
	25	48.30±0.08	48.41±0.05	48.61±0.11	45.75±0.10
OMNI.	1	67.06±0.04	67.11±0.11	67.35±0.08	63.74±0.08
	5	59.00±0.03	59.28±0.08	59.34±0.10	57.53±0.08
	25	54.79±0.04	54.83±0.04	54.88±0.06	54.32±0.04

to limit overfitting. Performance is assessed using the average log-conditional $\log p(\mathbf{x}_2|\mathbf{x}_1)$ measured using 4000 IW samples, and mean±standard deviation of the negative log-likelihood for five runs are reported. Undirected posteriors outperform the directed models by several nats on MNIST and OMNIGLOT.

4.4. Comparison against Previously Published Results

The experiments in this paper are all performed under an identical training framework for all the methods for a fair comparison. However, this requires reimplementing previous work. In order to better situate DVAE## against the state-of-the-art, we compare our performance with previously published results on training binary VAEs that use similar encoder/decoder. Here, for each previous work, we report the best number from the original paper for the latent variable size of 400 with the IW sample size of $K = 1$ and $K \sim 25$. Note that, for the IWAE experiments on DVAE##, K is set to 25, but for a previous work if $K = 25$ is not available, we report the closest K in $\{20, 50\}$.

As we can see in Table 4, our DVAE## model outperforms the previous methods that use similar encoder/decoder. We also observe that our implementation of GumBolt and DVAE# improves the performance of these models (see Appendix B for more details).

5. Conclusions and Future Directions

We have introduced a gradient estimator for stochastic optimization with undirected distributions and showed that VAEs and IWAEs with RBM posteriors can outperform similar models having directed posteriors. These encouraging results for UGMs over discrete variables suggest a number of promising research directions.

Exponential Family Harmoniums: The methods we have outlined also apply to UGMs defined over continuous variables (as in Fig 2). Exponential Family Harmoniums

Table 4: Comparison against previously published results on binary VAEs and IWAEs. Our DVAE## outperforms previous models that use similar encoder/decoder.

	VAE		IWAE	
	$K = 1$		$K \in \{20, 25, 50\}$	
	MNIST	OMNI.	MNIST	OMNI.
Concrete (Maddison et al.)	87.9	105.9	85.7	106.8
VIMCO (Mnih & Rezende)	88.4	111.7	85.5	113.2
DVAE++ (Vahdat et al.)	84.27	100.55	-	-
GumBolt (Khoshaman et al.)	83.28	99.83	82.75	98.81
GumBolt (our implement.)	83.04	99.12	-	-
DVAE# (Vahdat et al.)	83.18	99.65	82.82	98.88
DVAE# (our implement.)	82.93	99.40	-	-
DVAE##	82.75	98.61	82.97	98.34

(EFHs) (Welling et al., 2005) generalize RBMs and are composed of two disjoint groups of exponential-family random variables. EFHs retain the factorial structure of the conditionals $q(z_1|z_2)$ and $q(z_2|z_1)$ so that we can easily backpropagate through Gibbs updates. However, the exponential-family generalization allows for the mixing of different latent variable types and distributions.

Auxiliary Random Variables: An appealing property of RBMs (and EFHs) is that either group of variables can be analytically marginalized out. This property can be exploited to form more expressive approximate posteriors. Consider a generative model $p(\mathbf{z}, \mathbf{x}) = p(\mathbf{z})p(\mathbf{x}|\mathbf{z})$. To approximate the posterior we augment the latent space of \mathbf{z} with auxiliary variables \mathbf{h} , and form a UGM over the joint space. In this case, the marginal approximate posterior $q(\mathbf{z}|\mathbf{x}) = \sum_{\mathbf{h}} q(\mathbf{z}, \mathbf{h}|\mathbf{x})$ has more expressive power. Sampling from $q(\mathbf{z}|\mathbf{x})$ can be done by sampling from the joint $q(\mathbf{z}, \mathbf{h}|\mathbf{x})$ and our gradient estimator can be used for training the parameters of the UGM. The objective function of VAE can be optimized easily as the log-marginal $\log q(\mathbf{z}|\mathbf{x})$ has an analytic expression up to the normalization constant given a sample from the posterior.

Combining DGMs and UGMs: VAEs trained with DGM posteriors $q(\mathbf{z}) = \prod_i q_i(\mathbf{z}_i|\mathbf{z}_{<i})$ have shown promising results. However, each factor $q_i(\mathbf{z}_i|\mathbf{z}_{<i})$ in DGMs is typically assumed to be a product of independent distributions. We can build more powerful DGMs by modeling each $q_i(\mathbf{z}_i|\mathbf{z}_{<i})$ using a UGM.

Sampling and Computing Partition Function: A challenge when using continuous-variable UGMs as posteriors is the requirement for sampling and partition function estimation during evaluation if the test data log-likelihood is estimated using the IW bound. However, we note that approximate posteriors are only required for training VAEs. At evaluation, we can sidestep sampling and partition function estimation challenges using techniques such as AIS (Wu et al., 2017) starting from the prior distribution when the latent variables are continuous.

References

- Ackley, D. H., Hinton, G. E., and Sejnowski, T. J. A learning algorithm for Boltzmann machines. *Cognitive science*, 9(1): 147–169, 1985.
- Andriyash, E., Vahdat, A., and Macreedy, W. G. Improved gradient-based optimization over discrete distributions. *arXiv preprint arXiv:1810.00116*, 2018.
- Bornschein, J., Shabaniyan, S., Fischer, A., and Bengio, Y. Bidirectional Helmholtz machines. In *International Conference on Machine Learning*, pp. 2511–2519, 2016.
- Burda, Y., Grosse, R., and Salakhutdinov, R. Importance weighted autoencoders. *arXiv preprint arXiv:1509.00519*, 2015.
- Caterini, A. L., Doucet, A., and Sejdinovic, D. Hamiltonian variational auto-encoder. In *Neural Information Processing Systems*, 2018.
- Chen, X., Kingma, D. P., Salimans, T., Duan, Y., Dhariwal, P., Schulman, J., Sutskever, I., and Abbeel, P. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016.
- Chen, X., Mishra, N., Rohaninejad, M., and Abbeel, P. Pixel-SNAIL: An improved autoregressive generative model. In *International Conference on Machine Learning*, 2018.
- Dinh, L., Krueger, D., and Bengio, Y. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real NVP. *arXiv preprint arXiv:1605.08803*, 2016.
- Grathwohl, W., Choi, D., Wu, Y., Roeder, G., and Duvenaud, D. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. In *International Conference on Learning Representations (ICLR)*, 2018.
- Graves, A. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- Gregor, K., Danihelka, I., Mnih, A., Blundell, C., and Wierstra, D. Deep autoregressive networks. *arXiv preprint arXiv:1310.8499*, 2013.
- Gregor, K., Danihelka, I., Graves, A., Rezende, D., and Wierstra, D. Draw: A recurrent neural network for image generation. In *International Conference on Machine Learning*, pp. 1462–1471, 2015.
- Gu, S., Levine, S., Sutskever, I., and Mnih, A. MuProp: Unbiased backpropagation for stochastic neural networks. *arXiv preprint arXiv:1511.05176*, 2015.
- Gulrajani, I., Kumar, K., Ahmed, F., Taiga, A. A., Visin, F., Vazquez, D., and Courville, A. PixelVAE: A latent variable model for natural images. *arXiv preprint arXiv:1611.05013*, 2016.
- Hinton, G. E., Dayan, P., Frey, B. J., and Neal, R. M. The “wake-sleep” algorithm for unsupervised neural networks. *Science*, 268(5214):1158, 1995.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Hoffman, M. D. Learning deep latent Gaussian models with Markov chain Monte Carlo. In *International Conference on Machine Learning*, 2017.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with Gumbel-Softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Khoshaman, A. H. and Amin, M. H. GumBolt: Extending Gumbel trick to Boltzmann priors. In *Neural Information Processing Systems*, 2018.
- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 10236–10245. 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *The International Conference on Learning Representations (ICLR)*, 2014.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems*, pp. 4743–4751, 2016.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Li, Y. and Turner, R. E. Rényi divergence variational inference. In *Advances in Neural Information Processing Systems*, pp. 1073–1081, 2016.
- Li, Y., Turner, R. E., and Liu, Q. Approximate inference with amortised MCMC. *arXiv preprint arXiv:1702.08343*, 2017.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Mnih, A. and Gregor, K. Neural variational inference and learning in belief networks. *arXiv preprint arXiv:1402.0030*, 2014.
- Mnih, A. and Rezende, D. Variational inference for Monte Carlo objectives. In *International Conference on Machine Learning*, pp. 2188–2196, 2016.
- QuPA. Quadrant population annealing library. <https://try.quadrant.ai/qupa>. Accessed: 2019-01-01.
- Raiko, T., Berglund, M., Alain, G., and Dinh, L. Techniques for learning binary stochastic feedforward neural networks. In *International Conference on Learning Representations (ICLR)*, 2015.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pp. 1278–1286, 2014.

- Roeder, G., Wu, Y., and Duvenaud, D. K. Sticking the landing: Simple, lower-variance gradient estimators for variational inference. In *Advances in Neural Information Processing Systems*, pp. 6925–6934, 2017.
- Rolfe, J. T. Discrete variational autoencoders. *arXiv preprint arXiv:1609.02200*, 2016.
- Salakhutdinov, R. and Murray, I. On the quantitative analysis of deep belief networks. In *Proceedings of the 25th international conference on Machine learning*, pp. 872–879. ACM, 2008.
- Salimans, T., Kingma, D., and Welling, M. Markov chain Monte Carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, 2015.
- Salimans, T., Karpathy, A., Chen, X., and Kingma, D. P. Pixel-CNN++: A PixelCNN implementation with discretized logistic mixture likelihood and other modifications. In *ICLR*, 2017.
- Sohn, K., Lee, H., and Yan, X. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, 2015.
- Sønderby, C. K., Raiko, T., Maaløe, L., Sønderby, S. K., and Winther, O. Ladder variational autoencoders. In *Advances in neural information processing systems*, pp. 3738–3746, 2016.
- Tucker, G., Mnih, A., Maddison, C. J., Lawson, J., and Sohl-Dickstein, J. REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, pp. 2624–2633, 2017.
- Tucker, G., Lawson, D., Gu, S., and Maddison, C. J. Doubly reparameterized gradient estimators for Monte Carlo objectives. *arXiv preprint arXiv:1810.04152*, 2018.
- Vahdat, A., Andriyash, E., and Macreedy, W. G. DVAE#: Discrete variational autoencoders with relaxed Boltzmann priors. In *Neural Information Processing Systems*, 2018a.
- Vahdat, A., Macreedy, W. G., Bian, Z., Khoshaman, A., and Andriyash, E. DVAE++: Discrete variational autoencoders with overlapping transformations. In *International Conference on Machine Learning (ICML)*, 2018b.
- Van Den Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. Pixel recurrent neural networks. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, pp. 1747–1756. JMLR. org, 2016.
- Welling, M. and Teh, Y. W. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 681–688, 2011.
- Welling, M., Rosen-Zvi, M., and Hinton, G. E. Exponential family harmoniums with an application to information retrieval. In *Advances in neural information processing systems*, pp. 1481–1488, 2005.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*, pp. 5–32. Springer, 1992.
- Wolf, C., Karl, M., and van der Smagt, P. Variational inference with hamiltonian monte carlo. *arXiv preprint arXiv:1609.08203*, 2016.
- Wu, Y., Burda, Y., Salakhutdinov, R., and Grosse, R. On the quantitative analysis of decoder-based generative models. In *The International Conference on Learning Representations (ICLR)*, 2017.