
Amortized Population Gibbs Samplers with Neural Sufficient Statistics

Hao Wu¹ Heiko Zimmermann¹ Eli Sennesh¹ Tuan Anh Le² Jan-Willem van de Meent¹

Abstract

We develop amortized population Gibbs (APG) samplers, a class of scalable methods that frames structured variational inference as adaptive importance sampling. APG samplers construct high-dimensional proposals by iterating over updates to lower-dimensional blocks of variables. We train each conditional proposal by minimizing the inclusive KL divergence with respect to the conditional posterior. To appropriately account for the size of the input data, we develop a new parameterization in terms of neural sufficient statistics. Experiments show that APG samplers can train highly structured deep generative models in an unsupervised manner, and achieve substantial improvements in inference accuracy relative to standard autoencoding variational methods.

1. Introduction

Many inference tasks involve hierarchical structure. For example, when modeling a corpus of videos that contain moving objects, we may wish to reason about three levels of representation. At the corpus level, we can reason about the distribution over objects and dynamics of motion. At the instance level, i.e. for each video, we can reason the visual features and trajectories of individual objects. Finally, at a data point level, i.e. for each video frame, we can reason about the visibility of objects and their position.

In the absence of supervision, uncovering hierarchical structure from data requires inductive biases. Deep generative models let us incorporate biases in the form of priors that mirror the structure of the problem domain. By parameterizing conditional distributions with neural networks, we hope to learn models that use corpus-level characteristics (e.g. the appearance of objects and the motion dynamics) to

make data-efficient inferences about instance-level variables (e.g. object appearance and positions).

In recent years we have seen applications of amortized variational methods to inference in hierarchical domains, such as object detection (Eslami et al., 2016), modeling of user reviews (Esmaeili et al., 2019), and object tracking (Kosiorrek et al., 2018). These approaches build on the framework of variational autoencoders (VAEs) (Kingma & Welling, 2013; Rezende et al., 2014) to train structured deep generative models. However, scaling up these approaches to more complex domains still poses significant challenges. Whereas it is easy to train VAEs on large corpora of data, it is not easy to train structured encoders for models with a large number of (correlated) variables at the instance level.

In this paper, we develop methods for amortized inference that scale to structured models with hundreds of latent variables. Our approach takes inspiration from work by Johnson et al. (2016), which develops methods for efficient inference in deep generative models with conjugate-exponential family priors. In this class of models, we can perform inference using variational Bayesian expectation maximization (VBEM) (Beal, 2003; Bishop, 2006; Wainwright & Jordan, 2008), which iterates between closed-form updates to blocks of variables. VBEM is computationally efficient, often converges in a small number of iterations, and scales to a large number of variables. However, VBEM is also model-specific, difficult to implement, and only applicable to a restricted class of conjugate-exponential models.

To overcome these limitations, we develop a more general approach that frames variational inference as importance sampling with amortized proposals. We propose *amortized population Gibbs* samplers, a class of methods that iterate between conditional proposals to blocks of variables. To train these proposals, we minimize the inclusive KL divergence w.r.t. the conditional posterior. We use proposals in a sequential Monte Carlo (SMC) sampler to reduce the variance of importance weights, which improves efficiency during training and inference at test time.

Our experiments establish that APG samplers can capture the hierarchical structure in deep generative models in an unsupervised manner. In Gaussian mixture models, where Gibbs updates can be computed in closed form, the learned proposals converge to the conditional posteriors. To evalu-

¹Khoury College of Computer Sciences, Northeastern University, Boston, MA, USA ²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA. Correspondence to: Hao Wu <haowu@ccs.neu.edu>.

ate the performance on non-conjugate models, we design a deep generative mixture model and an unsupervised tracking model. We compare APG samplers to reweighted wake-sleep methods, a bootstrapped population sampler, and a HMC-augmented baseline. In all experiments APG samplers can capture the hierarchical characteristics of the problem and outperformed existing methods in terms of training efficiency and inference accuracy.

We summarize the contributions of this paper as follows:

1. We develop APG samplers, a new class of amortized inference methods that employ approximate Gibbs conditionals to iteratively improve samples.
2. We present a novel parameterization using neural sufficient statistics to learn proposals that generalize across input instances that vary in size.
3. We show that APG samplers can be used to train structured deep generative models with 100s of instance-level variables in an unsupervised manner.
4. We demonstrate substantial improvements in terms of test-time inference efficiency relative to baseline methods.

2. Background

In probabilistic machine learning, we commonly define a generative model in terms of a joint distribution $p_\theta(x, z)$ over data x and latent variables z . Given a model specification, we are interested in reasoning about the posterior distribution $p_\theta(z | x)$ conditioned on data x sampled from an (unknown) data distribution $p^{\text{DATA}}(x)$, which we in practice approximate using an empirical distribution over training data $\hat{p}(x)$. Amortized variational inference methods approximate the posterior $p_\theta(z | x)$ with a distribution $q_\phi(z | x)$ in some tractable variational family.

A widely used class of deep probabilistic models are variational autoencoders (Kingma & Welling, 2013; Rezende et al., 2014). These models are trained by maximizing the stochastic lower bound (ELBO) in Equation 1 on the log marginal likelihood, which is equivalent to minimizing the exclusive KL divergence $\text{KL}(q_\phi(z|x) || p_\theta(z|x))$.

$$\mathcal{L}(\phi) = \mathbb{E}_{\hat{p}(x) q_\phi(z|x)} \left[\log \frac{p_\theta(x, z)}{q_\phi(z | x)} \right] \quad (1)$$

When the distribution specified by the encoder is reparameterizable, we can compute Monte Carlo estimates of the gradient of this objective using pathwise derivatives. Non-reparameterizable cases, such as models with discrete variables, require likelihood-ratio estimators (Williams, 1992) which can have a high variance. A range of approaches have been proposed to reduce this variance, including reparameterized continuous relaxations (Maddison et al., 2017; Jang et al., 2017), credit assignment techniques (Weber et al.,

2019), and other control variates (Mnih & Rezende, 2016; Tucker et al., 2017; Grathwohl et al., 2018).

Rewighted wake-sleep (RWS) methods (Bornschein & Bengio, 2014) sidestep the need for reparameterization by minimizing the stochastic upper bound in Equation 2, which is equivalent to minimizing an inclusive KL divergence:

$$\mathcal{U}(\phi) = \mathbb{E}_{\hat{p}(x) p_\theta(z|x)} \left[\log \frac{p_\theta(x, z)}{q_\phi(z | x)} \right]. \quad (2)$$

Here we can compute a self-normalized gradient estimator using samples $z \sim q_\phi(z | x)$ and weights $w = \frac{p_\theta(x, z)}{q_\phi(z|x)}$

$$-\nabla_\phi \mathcal{U}(\phi) \simeq \sum_{l=1}^L \frac{w^l}{\sum_{l'} w^{l'}} \nabla_\phi \log q_\phi(z^l | x). \quad (3)$$

This gradient estimator has a number of advantages over the estimator used in standard VAEs (Le et al., 2019). Notably, it only requires that the *proposal density* is differentiable, whereas reparameterized estimators require that the *sample* itself is differentiable. This is a milder condition that holds for most distributions of interest, including those over discrete variables. Moreover, minimizing the inclusive KL reduces our risk of learning a proposal that collapses to a single mode of a multi-modal posterior (Le et al., 2019).

3. Amortized Population Gibbs Samplers

To generate high-quality samples from the posterior in an incremental manner, we develop methods that are inspired by EM and Gibbs sampling, both of which perform iterative updates to blocks of variables. Concretely, we assume that the latent variables in the generative model decompose into blocks $z = \{z_1, \dots, z_B\}$ and train proposals $q_\phi(z_b | x, z_{-b})$ that update the variables in a each block z_b conditioned on variables in the remaining blocks $z_{-b} = z \setminus \{z_b\}$.

Starting from a sample $q_\phi(z^1 | x)$ from a standard encoder, we generate a sequence of samples $\{z^1, \dots, z^K\}$ by performing a sweep of conditional updates to each block z_b

$$q_\phi(z^k | x, z^{k-1}) = \prod_{b=1}^B q_\phi(z_b^k | x, z_{\prec b}^k, z_{\succ b}^{k-1}), \quad (4)$$

where $z_{\prec b} = \{z_i | i < b\}$ and $z_{\succ b} = \{z_i | i > b\}$. Repeatedly applying sweep updates then yields a proposal

$$q_\phi(z^1, \dots, z^K | x) = q_\phi(z^1 | x) \prod_{k=2}^K q_\phi(z^k | x, z^{k-1}).$$

We want to train proposals that improve the quality of each sample z^k relative to that of the preceding sample z^{k-1} . There are two possible strategies for accomplishing this. One is to define an objective that minimizes the discrepancy

between the marginal $q_\phi(z^K | x)$ for the final sample and the posterior $p_\theta(z^K | x)$. This corresponds to learning a sweep update $q_\phi(z^k | x, z^{k-1})$ that transforms the initial proposal to the posterior in exactly K sweeps. An example of this type of approach, albeit one that does not employ block updates, is the recent work on annealing variational objectives (Huang et al., 2018).

In this paper, we will pursue a different approach. Instead of transforming the initial proposal in exactly K steps, we learn a sweep update that leaves the target density *invariant*

$$p_\theta(z^k | x) = \int p_\theta(z^{k-1} | x) q_\phi(z^k | x, z^{k-1}) dz^{k-1}.$$

When this condition is met, the proposal $q_\phi(z^1, \dots, z^K | x)$ is a Markov Chain whose stationary distribution is the posterior. This means a sweep update learned at training time can be applied at test time to iteratively improve sample quality, without requiring a pre-specified number of updates K .

When we require that each block update $q_\phi(z'_b | x, z_{-b})$ leaves the target density invariant,

$$\begin{aligned} p_\theta(z'_b, z_{-b} | x) &= \int p_\theta(z_b, z_{-b} | x) q_\phi(z'_b | x, z_{-b}) dz_b \\ &= p_\theta(z_{-b} | x) q_\phi(z'_b | x, z_{-b}), \end{aligned} \quad (5)$$

then each update must equal the exact conditional posterior $q_\phi(z'_b | x, z_{-b}) = p_\theta(z'_b | x, z_{-b})$. In other words, when the condition in Equation 5 is met, the proposal $q_\phi(z^1, \dots, z^K | x)$ is a Gibbs sampler.

3.1. Variational Objective

To learn each of the block proposals $q_\phi(z_b | x, z_{-b})$ we will minimize the inclusive KL divergence $\mathcal{K}_b(\phi)$

$$\mathbb{E}_{\hat{p}(x)p_\theta(z_{-b}|x)} \left[\text{KL} (p_\theta(z_b | x, z_{-b}) || q_\phi(z_b | x, z_{-b})) \right]$$

Unfortunately, this objective is intractable, since we are not able to evaluate the density of the true marginal $p_\theta(z_{-b} | x)$, nor that of the conditional $p_\theta(z_b | z_{-b}, x)$. As we will discuss in Section 6, this has implications for the evaluation of the learned proposals, since we cannot compute a lower or upper bound on the log marginal likelihood as in other variational methods. However, it is nonetheless possible to approximate the gradient of the objective

$$-\nabla_\phi \mathcal{K}_b(\phi) = \mathbb{E}_{\hat{p}(x)p_\theta(z_b, z_{-b}|x)} [\nabla_\phi \log q_\phi(z_b | x, z_{-b})].$$

We can estimate this gradient using any Monte Carlo method that generates samples from the posterior $z \sim p_\theta(z | x)$. In the next section, we will use the learned proposals to define an importance sampler, which we will use to obtain

a set of weighted samples $\{(w^l, z^l)\}_{l=1}^L$ to compute a self-normalized gradient estimate

$$-\nabla_\phi \mathcal{K}_b(\phi) \simeq \sum_{l=1}^L \frac{w^l}{\sum_{l'} w^{l'}} \nabla_\phi \log q_\phi(z_b^l | x, z_{-b}^l). \quad (6)$$

When we also want to learn a generative model $p_\theta(x, z)$, we can compute a similar self-normalized gradient estimate

$$\begin{aligned} \nabla_\theta \log p_\theta(x) &= \mathbb{E}_{p_\theta(z|x)} [\nabla_\theta \log p_\theta(x, z)], \\ &\simeq \sum_{l=1}^L \frac{w^l}{\sum_{l'} w^{l'}} \nabla_\theta \log p_\theta(x, z^l). \end{aligned} \quad (7)$$

to maximize the marginal likelihood. The identity in Equation 7 holds due to the identity $\mathbb{E}_{p_\theta(z|x)} [\nabla_\theta \log p_\theta(z|x)] = 0$ (see Appendix A for details).

In the following, we will describe an adaptive importance sampling scheme which generates high quality samples to accurately estimate gradients for the conditional proposal $q_\phi(z_b | x, z_{-b})$ and an initial one-shot encoder $q_\phi(z|x)$.

3.2. Generating High Quality Samples

A well-known limitation of self-normalized importance samplers is that weights can have a high variance when latent variables are high-dimensional and/or correlated, which is common in models with hierarchical structure. Moreover, this variance gives rise to bias in self-normalized estimators. This bias arises from the approximation of the normalizing factor $1/p_\theta(x) \simeq 1/\hat{Z}$ using the average weight $\hat{Z} = \frac{1}{L} \sum_l w^l$. The average weight is an unbiased estimate of the marginal likelihood $\mathbb{E}[\hat{Z}] = p_\theta(x)$. However $1/\hat{Z}$ is not an unbiased estimate of the normalizing factor, since by Jensen's inequality $\mathbb{E}[1/\hat{Z}] \geq 1/p_\theta(x)$, resulting in a bias that is directly tied to the variance of the weights.

In practice, limitations on GPU memory and computation time imply that the estimators in Equations 6 and 7 need to employ a small budget L . This means that the bias will likely be high, particularly during the early stages of training when the proposal poorly approximates the posterior. However, even a biased estimator can serve to improve the quality of a proposal in practice. A high-bias estimator will have a low effective sample size, which is to say that it will assign a normalized weight close to 1 to the best proposal in the set and weight 0 to all other samples. The gradient in Equation 6 will therefore initially serve to increase the probability of the highest-weight sample. As long as doing so improves the quality of the proposal, RWS-style optimization will gradually reduce the bias as training proceeds.

To help mitigate the bias of self-normalized estimators, we use a sequential Monte Carlo sampler (Del Moral et al., 2006) to reduce weight variance. SMC methods (Doucet

Algorithm 1 Sequential Monte Carlo sampler

```

1: Input Model  $\gamma^{1:K}$ , Proposals  $q^{1:K}$ , Number of particles  $L$ 
2: for  $l = 1$  to  $L$  do
3:    $z^{1,l} \sim q^1(\cdot)$  ▷ Propose
4:    $w^{1,l} = \frac{\gamma^1(z^{1,l})}{q^1(z^{1,l})}$  ▷ Weigh
5: end for
6: for  $k = 2$  to  $K$  do
7:    $z^{k-1,1:L}, w^{k-1,1:L} = \text{RESAMPLE}(z^{k-1,1:L}, w^{k-1,1:L})$ 
8:   for  $l = 1$  to  $L$  do
9:      $z^{k,l} \sim q^k(\cdot | z^{k-1,l})$  ▷ Propose
10:     $w^{k,l} = \frac{\gamma^k(z^{k,l}) r^{k-1}(z^{k-1,l} | z^{k,l})}{\gamma^{k-1}(z^{k-1,l}) q^k(z^{k,l} | z^{k-1,l})} w^{k-1,l}$  ▷ Weigh
11:   end for
12: end for
13: Output Weighted Samples  $\{z^{K,l}, w^{K,l}\}_{l=1}^L$ 
    
```

et al., 2001) decompose a high-dimensional sampling problem into a sequence of lower-dimensional problems. They do so by combining two ideas: (1) sequential importance sampling, which defines a proposal for a sequence of variables using a sequence of conditional proposals, (2) resampling, which selects intermediate proposals with probability proportional to their weights to improve sample quality.

SMC methods are most commonly used in state space models to generate proposals for a sequence of variables by proposing one variable at a time. We here consider SMC samplers (see Algorithm 1), a subclass of SMC methods that interleave resampling with the application of a transition kernel. The distinction between SMC methods for state space models and SMC samplers is subtle but important. Whereas the former generate proposals for a sequence of variables $z_{1:t}$ by proposing $z_t \sim q(z_t | z_{1:t-1})$ to *extend* the sample space at each iteration, SMC samplers can be understood as an importance sampling analogue to Markov chain Monte Carlo (MCMC) methods (Brooks et al., 2011), which construct a Markov chain $z^{1:K}$ by generating a proposal z^k from a transition kernel $q(z^k | z^{k-1})$ at each iteration.

To understand how approximate Gibbs proposals can be incorporated into a SMC sampler, we will first define a sequential importance sampler (SIS), which decomposes the importance weight into a sequence of *incremental* weights. SIS considers a sequence of unnormalized target densities $\gamma^1(z^1), \gamma^2(z^{1:2}), \dots, \gamma^K(z^{1:K})$. If we now consider an initial proposal $q^1(z^1)$, along with a sequence of conditional proposals $q^k(z^k | z^{1:k-1})$, then we can recursively construct a sequence of weights $w^k = v^k w^{k-1}$ by assuming $w^1 = \gamma^1(z^1)/q^1(z^1)$ and defining the incremental weight

$$v^k = \frac{\gamma^k(z^{1:k})}{\gamma^{k-1}(z^{1:k-1}) q^k(z^k | z^{1:k-1})}.$$

This construction ensures that, at step k in the sequence, we have a weight w^k relative to the intermediate density

$\gamma^k(z^{1:k})$ of the form (see Appendix B)

$$w^k = \frac{\gamma^k(z^{1:k})}{q^1(z^1) \prod_{k'=2}^k q^{k'}(z^{k'} | z^{1:k'-1})}.$$

We will now consider a specific sequence of intermediate densities that are defined using a *reverse kernel* $r(z' | z)$:

$$\gamma^k(z^{1:k}) = p_\theta(x, z^k) \prod_{k'=2}^k r(z^{k'-1} | z^{k'}).$$

This defines a density on an *extended space* which admits the marginal density

$$\gamma^k(z^k) = \int \gamma^k(z^{1:k}) dz^{1:k-1} = p_\theta(x, z^k).$$

This means that at each step k , we can treat the preceding samples $z^{1:k-1}$ as *auxiliary variables*; if we generate a proposal $z^{1:k}$ and simply disregard $z^{1:k-1}$, then the pair (w^k, z^k) is a valid importance sample relative to $p_\theta(z^k | x)$. If we additionally condition proposals on x , the incremental weight for this particular choice of target densities is

$$v^k = \frac{p_\theta(x, z^k) r(z^{k-1} | x, z^k)}{p_\theta(x, z^{k-1}) q(z^k | x, z^{k-1})}. \quad (8)$$

This construction is valid for any choice of *forward kernel* $q(z^k | z^{k-1})$ and reverse kernel $r(z^{k-1} | z^k)$, but the weight variance will depend on the choice of kernel. Given a forward kernel, the *optimal* reverse kernel is

$$r(z^{k-1} | x, z^k) = \frac{p_\theta(x, z^{k-1})}{p_\theta(x, z^k)} q(z^k | x, z^{k-1}).$$

For this choice of kernel, the incremental weights are 1, which minimizes the variance of the weights w^k .

We here propose to use the approximate Gibbs kernel $q_\phi(z^k | x, z^{k-1})$ as both the forward and the reverse kernel. When the approximate Gibbs kernel converges to the actual Gibbs kernel, this choice becomes optimal, since in limit the kernel satisfies detailed balance

$$p_\theta(x, z^k) q_\phi(z^{k-1} | x, z^k) = p_\theta(x, z^{k-1}) q_\phi(z^k | x, z^{k-1}).$$

In general, the weights w^k in SIS will have a high variance when z is high-dimensional or correlated. Moreover, we sample these variables K times. When the approximate Gibbs kernel converges to the true kernel, this will not increase the variance of weights (since $v^k = 1$ in limit). But during training variance of weights w^k will increase with k , since we are jointly sampling an entire Markov chain.

To overcome this problem, SMC samplers interleave applications of the transition kernel with a *resampling* steps (see Appendix D for Algorithm), which generates a new

Algorithm 2 Amortized Population Gibbs Sampler

```

1: Input Data  $x$ , Model  $p_\theta(x, z)$ 
2:   Proposals  $q_\phi(z|x), \{q_\phi(z_b|x, z_{-b})\}_{b=1}^B$ 
3:    $g_\phi = 0, g_\theta = 0$   $\triangleright$  Initialize gradient to 0
4:   for  $l = 1$  to  $L$  do  $\triangleright$  Initial proposal
5:      $z^{1,l} \sim q_\phi(\cdot | x)$   $\triangleright$  Propose
6:      $w^{1,l} = \frac{p_\theta(x, z^{1,l})}{q_\phi(z^{1,l}|x)}$   $\triangleright$  Weigh
7:   end for
8:    $g_\phi = g_\phi + \sum_{l=1}^L \frac{w^{1,l}}{\sum_{l'=1}^L w^{1,l'}} \nabla_\phi \log q_\phi(z^{1,l} | x)$ 
9:    $g_\theta = g_\theta + \sum_{l=1}^L \frac{w^{1,l}}{\sum_{l'=1}^L w^{1,l'}} \nabla_\theta \log p_\theta(x, z^{1,l})$ 
10:  for  $k = 2$  to  $K$  do  $\triangleright$  Gibbs sweeps
11:     $\tilde{z}^{1:L}, \tilde{w}^{1:L} = z^{k-1,1:L}, w^{k-1,1:L}$ 
12:    for  $b = 1$  to  $B$  do  $\triangleright$  Block updates
13:       $\tilde{z}^{1:L}, \tilde{w}^{1:L} = \text{RESAMPLE}(\tilde{z}^{1:L}, \tilde{w}^{1:L})$ 
14:      for  $l = 1$  to  $L$  do
15:         $\tilde{z}_b^{l,l} \sim q_\phi(\cdot | x, \tilde{z}_{-b}^l)$   $\triangleright$  Propose
16:         $\tilde{w}^l = \frac{p_\theta(x, \tilde{z}_b^{l,l}, \tilde{z}_{-b}^l) q_\phi(\tilde{z}_b^l | x, \tilde{z}_{-b}^l)}{p_\theta(x, \tilde{z}_b^l, \tilde{z}_{-b}^l) q_\phi(\tilde{z}_b^{l,l} | x, \tilde{z}_{-b}^l)} \tilde{w}^l$   $\triangleright$  Weigh
17:         $\tilde{z}_b^l = \tilde{z}_b^{l,l}$   $\triangleright$  Reassign
18:      end for
19:       $g_\phi = g_\phi + \sum_{l=1}^L \frac{\tilde{w}^l}{\sum_{l'=1}^L \tilde{w}^{l'}} \nabla_\phi \log q_\phi(\tilde{z}_b^l | x, \tilde{z}_{-b}^l)$ 
20:       $g_\theta = g_\theta + \sum_{l=1}^L \frac{\tilde{w}^l}{\sum_{l'=1}^L \tilde{w}^{l'}} \nabla_\theta \log p_\theta(x, \tilde{z}^l)$ 
21:    end for
22:     $z^{k,1:L}, w^{k,1:L} = \tilde{z}^{1:L}, \tilde{w}^{1:L}$ 
23:  end for
24: Output Gradient  $g_\phi, g_\theta$ , Weighted Samples  $\{z^{K,l}, w^{K,l}\}_{l=1}^L$ 
    
```

set of equally weighted samples by selecting samples from the current sample set with probability proportional to their weight. APG samplers employ resampling after each block update $z_b \sim q_\phi(z_b|x, z_{-b})$, which results in equal incoming weights in the subsequent block update. This allows us to compute the gradient estimate of block proposals based on the contribution of the incremental weight

$$v = \frac{p_\theta(x, z_b', z_{-b}) q_\phi(z_b | x, z_{-b})}{p_\theta(x, z_b, z_{-b}) q_\phi(z_b' | x, z_{-b})}.$$

These incremental weights will generally have a much lower variance than full weights. This can improve sample-efficiency for both gradient estimation and inference, particularly in models with high-dimensional correlated variables.

We refer to this implementation of a SMC sampler as an *amortized population Gibbs* sampler, and summarize all the steps of the computation in Algorithm 2. In Appendix E, we prove that this algorithm is correct using an argument based on proper weighting.

4. Neural Sufficient Statistics

Gibbs sampling strategies that sample from exact conditionals rely on conjugacy relationships. We assume a prior and

likelihood that can both be expressed as exponential families

$$p(x | z) = h(x) \exp\{\eta(z)^\top T(x) - \log A(\eta(z))\},$$

$$p(z) = h(z) \exp\{\lambda^\top T(z) - \log A(\lambda)\}.$$

where $h(\cdot)$ is a base measure, $T(\cdot)$ is a vector of sufficient statistics, and $A(\cdot)$ is a log normalizer. The two densities are jointly conjugate when $T(z) = (\eta(z), -\log A(\eta(z)))$. In this case, the posterior distribution lies in the same exponential family as the prior. Typically, the prior $p(z | \lambda)$ and likelihood $p(x | z)$ are not jointly conjugate, but it is possible to identify conjugacy relationships at the level of individual blocks of variables,

$$p(z_b | z_{-b}, x) \propto h(z_b) \exp\{(\lambda_{b,1} + T(x, z_{-b}))^\top T(z_b) - (\lambda_{b,2} + 1) \log A(\eta(z_b))\}.$$

In general, these conjugacy relationships will not hold. However, we can still take inspiration to design variational distributions that make use of conditional independencies in a model. We assume that each of the approximate Gibbs updates $q_\phi(z_b | x, z_{-b})$ is an exponential family, whose parameters are computed from a vector of prior parameters λ and a vector of neural sufficient statistics $T_\phi(x, z_{-b})$

$$q_\phi(z_b | x, z_{-b}) = p(z_b | \lambda + T_\phi(x, z_{-b})).$$

This parameterization has a number of desirable properties. Exponential families are the largest-entropy distributions that match the moments defined by the sufficient statistics (see e.g. [Wainwright & Jordan \(2008\)](#)), which is helpful when minimizing the inclusive KL divergence. In exponential families it is also more straightforward to control the entropy of the variational distribution. In particular, we can initialize $T_\phi(x, z_{-b})$ to output values close to zero in order to ensure that we initially propose from a prior and/or regularize $T_\phi(x, z_{-b})$ to help avoid local optima.

A useful case arises when the data $x = \{x_1, \dots, x_N\}$ are independent conditioned on z . In this setting, it is often possible to partition the latent variables $z = \{z^G, z^L\}$ into global (instance-level) variables z^G and local (datapoint-level) variables z^L . The dimensionality of global variables is typically constant, whereas local variables $z^L = \{z_1^L, \dots, z_N^L\}$ have a dimensionality that increases with the data size N . For models with this structure, the local variables are typically conditionally independent $z_n^L \perp z_{-n}^L | x, z^G$, which means that we can parameterize the variational distributions as

$$\tilde{\lambda}^G = \lambda^G + \sum_{n=1}^N T_\phi^G(x_n, z_n^L), \quad \tilde{\lambda}_n^L = \lambda_n^L + T_\phi^L(x_n, z_n^G).$$

The advantage of this parameterization is it allows us to train approximate Gibbs updates for global variables in a manner that scales dynamically with the size of the dataset,

and appropriately adjusts the posterior variance according to the amount of available data. See Appendix F for detailed explanations of how we use this parameterization to design the neural proposals in our experiments.

5. Related Work

Our work fits into a line of recent methods for deep generative modeling that seek to improve inference quality, either by introducing auxiliary variables (Maaløe et al., 2016; Ranganath et al., 2016), or by performing iterative updates (Marino et al., 2018). Work by Hoffman (2017) applies Hamiltonian Monte Carlo to samples that are generated from the encoder, which serves to improve the gradient estimate w.r.t. θ (Equation 7), while learning the inference network using a standard reparameterized ELBO objective. Li et al. (2017) also use MCMC to improve the quality of samples from an encoder, but additionally use these samples to train the encoder by minimizing the inclusive KL divergence relative to the filtering distribution of the Markov chain. As in our work, the filtering distribution after multiple MCMC steps is intractable. Li et al. (2017) therefore use an adversarial objective to minimize the inclusive KL. Neither of these methods consider block decompositions of latent variables, nor do they learn kernels.

Salimans et al. (2015) derive a stochastic lower bound for variational inference which uses an importance weight defined in terms of forward and reverse kernels in MCMC steps, similar to Equation 8. Caterini et al. (2018) extend this work by optimally selecting reverse kernels (rather than learning them) using inhomogeneous Hamiltonian dynamics. Huang et al. (2018) learn a sequence of transition kernels that performs annealing from the initial encoder to the posterior. Ruiz & Titsias (2019) define a contrastive divergence which can be tractably optimized w.r.t variational parameters ϕ . Since all of these methods minimize an exclusive KL, rather than an inclusive KL, the gradient estimates rely on reparameterization, which makes them inapplicable to models with discrete variables. Moreover, these methods perform a joint update on all variables, while we consider the task of learning conditional proposals.

Wang et al. (2018) develop a meta-learning approach to learn Gibbs block conditionals. This work assumes a setup in which it is possible to sample x, z from the true generative model $p(x, z)$. This circumvents the need for estimators in Equation 6 and 7, which are necessary when we wish to learn the generative model. Similar to our work, this approach minimizes the inclusive KL, but uses the learned conditionals to define an (approximate) MCMC sampler, rather than using them as proposals in a SMC sampler. This work also has a different focus from ours, in that it primarily seeks to learn block conditionals that have the potential to generalize to previously unseen graphical models.

6. Experiments

We evaluate APG samplers on three different tasks. We begin by analysing APG samplers on a Gaussian mixture model (GMM) as an exemplar of a model in the conjugate-exponential family. This experiment allows us to analytically compute the conditional posterior to verify if the learned proposals indeed converge to the Gibbs kernels. Here APG samplers outperform other iterative inference methods even when using a smaller computational budget. Next we consider a deep generative mixture model (DMM) that incorporates a neural likelihood and use APG samplers to jointly train the generative and inference model. We demonstrate the scalability of APG samplers by performing accurate inference at test-time on instances with 600 points. In our third experiment, we consider an unsupervised tracking model for multiple bouncing MNIST digits. We extend the task proposed by Srivastava et al. (2015) to up to five digits, and learn both a deep generative model for videos and an inference model that performs tracking. At test time we show that APG easily scales beyond previously reported results for a specialized recurrent architecture (Kosiorek et al., 2018). The results for each of these tasks constitute significant advances relative to the state of the art. APG samplers are not only able to scale to models with higher complexity, but also provides a general framework for performing inference in models with global and local variables, which can be adapted to a variety of model classes with comparative ease.

6.1. Baselines

We compare our APG sampler with three baseline methods. The first is RWS. The second is a bootstrapped population Gibbs (BPG) sampler, which uses exactly the same sampling scheme as in Algorithm 2, but proposes from the prior rather than from approximate Gibbs kernels. This serves to evaluate whether learning proposals improves the quality of inference results. The third is a method that augments RWS with Hamiltonian Monte Carlo (HMC) updates, which is analogous to the method proposed by Hoffman (2017). This serves to compare Gibbs updates to those that can be obtained with state-of-the-art MCMC methods.

In the HMC-RWS baseline, we employ the standard RWS gradient estimators from Equations 3 and 7, but update samples from the standard encoder using HMC. In the bouncing MNIST model, all variables are continuous and are updated jointly. In the GMM, we use HMC to sample from the marginal for continuous variables (by marginalizing over cluster assignments via direct summation), and then sample discrete variables from enumerated Gibbs conditionals to evaluate the full log joint $\log p_\theta(x, z)$. For the DMM we perform the same procedure, but use HMC to sample from the conditional distribution given discrete variables, because

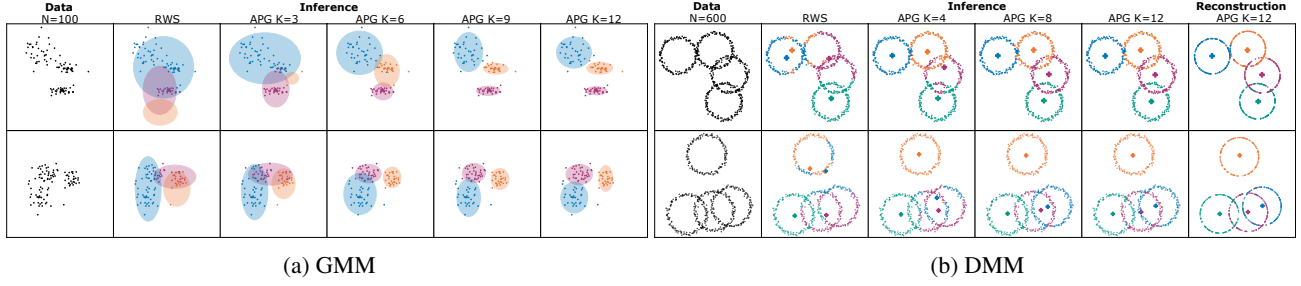


Figure 1. Visualization of a single sample for (a) GMM instances with $N = 100$ data points and (b) DMM instances with $N = 600$ data points. The left columns show the data instance, the second column the inference result from RWS, and the subsequent columns the inference result after K APG sweeps. In addition the rightmost column in (b) shows the reconstruction from the generative model. Unlike RWS, APG successfully learns the model parameters and an inference model and its performance improves with increasing K .

direct summation is not possible in this model.

Unless otherwise stated, we compare methods under an equivalent computational budget. We compare an APG sampler that performs K sweeps with L particles to RWS with $K \cdot L$ particles. When comparing to HMC-RWS methods, we perform K updates for L particles with a number of leapfrog steps LF in $\{1, 5, 10\}$, which corresponds to a 1x, 5x, and 10x computational budget as measured in terms of the number of evaluations of the log joint. We resort to report the log joint $\log p_\theta(x, z)$ as an evaluation metric, because the marginal $q_\phi(z^k | x)$ is intractable. A summary of these results can be found in Table 1.

6.2. Gaussian Mixture Model

To evaluate whether APG samplers can learn the exact Gibbs updates in conditionally conjugate models, we consider a 2D Gaussian mixture model

$$\begin{aligned} \mu_m, \tau_m &\sim \text{Normal-Gamma}(\mu_0, \nu_0, \alpha_0, \beta_0), \\ c_n &\sim \text{Cat}(\pi), x_n | c_n = m \sim \text{Normal}(\mu_m, 1/\tau_m) \\ m &= 1, 2, \dots, M, n = 1, 2, \dots, N. \end{aligned}$$

Each cluster has a vector-valued mean μ_m and a diagonal precision τ_m , for which we define an elementwise Normal-Gamma prior with $\mu_0 = 0, \nu_0 = 0.1, \alpha = 2, \beta = 2$.

In this model, the global variables $z^G = \{\mu_{1:M}, \tau_{1:M}\}$ are means and precisions of M mixture components; The local variables $z^L = \{c_{1:N}\}$ are cluster assignments of N points. Conditioned on cluster assignments, the Gaussian likelihood is conjugate to a normal-gamma prior with pointwise sufficient statistics $T(x_n, c_n)$

$$\left\{ \mathbb{I}[c_n = m], \mathbb{I}[c_n = m] x_n, \mathbb{I}[c_n = m] x_n^2 \mid m = 1, 2, \dots, M \right\}$$

where the identity function $\mathbb{I}[c_n = m]$ evaluates to 1 if the equality holds, and 0 otherwise.

We employ a block update strategy that iterates between blocks $\{\mu_{1:M}, \tau_{1:M}\}, \{c_{1:N}\}$ by learning neural proposals




Figure 2. $\log p_\theta(x, z)$ as a function of number of sweeps for a GMM test instance. We perform each method with L particles. For HMC-RWS, we perform multiple leapfrog steps LF in $\{1, 5, 10\}$. To achieve results comparable to APG and Gibbs samplers, HMC-RWS needs 10 times the computational budget, while BPG underperforms even with 100 times the computational budget.

$q_\phi(\mu, \tau \mid x, c)$ and $q_\phi(c \mid x, \mu, \tau)$. We use pointwise neural sufficient statistics modeled based on the ones in the analytic form. We train our model on 20,000 GMM instances, each containing $M = 3$ clusters and $N = 60$ points; We train with $K = 5$ sweeps, $L = 10$ particles, 20 instances per batch, learning rate 2.5×10^{-4} , and 2×10^5 gradient steps.

Figure 1a shows single samples in 2 test instances, each containing $N = 100$ points. Even when using a parameterization that employs neural sufficient statistics, the RWS encoder fails to propose reasonable clusters, whereas the APG sampler converges within 12 sweeps in GMMs with more variables than training instances.

We compare the APG samplers with true Gibbs samplers and the baselines in section 6.1. Figure 2 shows that the APG sampler converges to the $p_\theta(x, z)$ achieved the true Gibbs sampler; it outperforms HMC-RWS even when we run it with 10 times the computation budget; and learned proposals substantially improve on inference since BPG does not converge even with 100x computational budget. In addition, we verify the approximate Gibbs kernel converges to the true Gibbs kernel (see Appendix G).

Moreover, we compare joint block updates $\{\mu, \tau, c\}$ with de-

Table 1. $\log p_\theta(x, z)$ averaged over test corpus for RWS, BPG, HMC-RWS, and Gibbs Sampling for different numbers of sweeps K . We run HMC-RWS with 1, 5, and 10 leapfrog integration steps, which corresponds to 1, 5, and 10 times the computational budget used by APG with the same number of sweeps. The GMM test corpus contains 20,000 instances, each of which has $N = 100$ data points. The DMM test corpus contains 20,000 instances with $N = 600$ data points per instance. Bouncing MNIST’s test corpus contains 10,000 instances each containing $T = 100$ time steps and $D = 5$ digits. In all tasks APG with $K = 20$ performs the best.

	RWS	BPG K=20	HMC-RWS K=20, LF=1	HMC-RWS K=20, LF=5	HMC-RWS K=20, LF=10	APG K=5	APG K=10	APG K=20	GIBBS K=20
GMM ($\times 10^2$)	-7.651	-6.688	-6.594	-5.420	-4.673	-4.816	-4.556	-4.469	-4.573
DMM ($\times 10^3$)	-3.172	-4.204	-2.185	-2.184	-2.184	-2.050	-2.040	-2.035	-
BC-MNIST ($\times 10^5$)	-1.673	-1.884	-1.538	-1.405	-1.303	-0.706	-0.652	-0.620	-




Figure 3. $\log p(x, z)$ (left) and ESS (right) as a function of number of sweeps for APG samplers with different block update strategies and same overall sampling budget $K \cdot L = 1000$. A higher ESS is achieved by performing more sweeps with fewer samples.

composed blocks $\{\mu, \tau\}, \{c\}$, for varying number of sweeps (see Figure 3). We additionally assess sample quality using the effective sample size (ESS)

$$\frac{\text{ESS}}{L} = \frac{(\sum_{l=1}^L w^{k,l})^2}{L \sum_{l=1}^L (w^{k,l})^2}.$$

We can see that it is more effective to perform more sweeps K with a smaller number of particles L for a fixed computation budget. We also see that the decomposed block updates result in higher sample quality.

6.3. Deep Generative Mixture Model

In the deep generative mixture model (DMM) $p_\theta(x, z)$ our data is a corpus of 2D ring-shaped mixtures. The generative model takes the form

$$\begin{aligned} \mu_m &\sim \text{Normal}(\mu, \sigma_0^2 I), \\ c_n &\sim \text{Cat}(\pi), h_n \sim \text{Beta}(\alpha, \beta), \\ x_n | c_n = m &\sim \text{Normal}(g_\theta(h_n) + \mu_m, \sigma_\epsilon^2 I) \\ m &= 1, 2, \dots, M, \quad n = 1, 2, \dots, N. \end{aligned}$$

The global variables $z^G = \{\mu_{1:M}\}$ are the centers of M components; The local variables $z^L = \{c_{1:N}, h_{1:N}\}$ are cluster assignments and 1D embeddings of N points. Given an assignment c_n and an embedding h_n that goes into a MLP decoder g_θ , we sample a data point x_n with Gaussian noise. We choose $\mu = 0, \sigma_0 = 10, \alpha = 1, \beta = 1, \sigma_\epsilon = 0.1$.

We employ a block update strategy that iterates between blocks $\{\mu_{1:M}\}, \{c_{1:N}, h_{1:N}\}$ by learning neural proposals $q_\phi(\mu | x, c, h)$ and $q_\phi(c, h | x, \mu)$. We train our model on 20,000 instances, each containing $M = 4$ clusters and $N = 200$ points; We train our model with $K = 8$ sweeps, $L = 10$ particles, 20 instances per batch, learning rate 10^{-4} , and 3×10^5 gradient steps.

Once again, we compare the APG sampler with the encoders using RWS. Figure 1b shows visualization of single samples in 2 test instances, each containing $N = 600$ points. In both mixture model experiments, we can tell that the APG samplers scale to a much larger number of variables, whereas a standard encoder trained using RWS fails to produce reasonable proposals.

6.4. Time Series Model – Bouncing MNIST

In the bouncing MNIST model, our data is a corpus of video frames that contain multiple moving MNIST digits. To demonstrate the scalability of APG samplers, we train both a deep generative model and an inference model with 600,000 video instances, each of which contains 10 time steps and 3 digits. At test time we evaluate the model on instances that contain up to 100 time steps and 5 digits.

Consider a sequence of video frames $x_{1:T}$, which contains T time steps and D digits. We assume that the k th digit in the t th frame x_t can be represented by some digit feature z_d^{what} and a time-dependent position variable $z_{d,t}^{\text{where}}$. The deep generative model takes the form

$$\begin{aligned} z_d^{\text{what}} &\sim \text{Normal}(0, I), \\ z_{d,1}^{\text{where}} &\sim \text{Normal}(0, I), z_{d,t}^{\text{where}} \sim \text{Normal}(z_{d,t-1}^{\text{where}}, \sigma_0^2 I), \\ x_t &\sim \text{Bernoulli}\left(\sigma \left(\sum_d \text{ST}(g_\theta(z_d^{\text{what}}), z_{d,t}^{\text{where}}) \right)\right) \\ d &= 1, 2, \dots, D, t = 1, 2, \dots, T. \end{aligned}$$

In this model, the global variables $z^G = \{z_{1:D}^{\text{what}}\}$ are latent codes of the D digits, as those digits don’t change across the frames; The local variables $z^L = \{z_{1:D,1:T}^{\text{where}}\}$ are time-dependent positions. The hyper-parameter is $\sigma_0 = 0.1$. The dimensionalities are $z_d^{\text{what}} \in \mathbb{R}^{10}, z_{d,t}^{\text{where}} \in \mathbb{R}^2$.




Figure 4. A single sample from variational distribution in one video instance with $T = 100$ time steps and $D = 5$ digits. The visualization is truncated to the first 15 time steps due to limited space (see Appendix I for full time series and Appendix H for more examples). The sample is initialized from RWS (top) and updated by $K = 10$ APG sweeps (bottom). The APG sampler improves the inference results in tracking and learns better MNIST digit embeddings that improve reconstruction.

Table 2. Bouncing MNIST performance. Mean-squared error between test instance and its reconstruction across 5000 video instances with different number of time steps T and digits D . We use a fixed sampling budget of $K \cdot L = 1000$. APG samplers achieve substantial improvements relative to RWS and HMC-RWS.

	RWS	HMC-RWS K=20, LF=10	APG K=10	APG K=20
D=5, T=20	265.1	261.2	117.7	103.8
D=5, T=100	272.1	267.9	115.4	102.3
D=3, T=20	134.0	124.1	60.4	58.9
D=3, T=100	137.3	127.4	56.0	55.7

Given each digit feature z_d^{what} , a MLP decoder g_θ generates a MNIST digit image $g_\theta(z_d^{\text{what}})$ of the size 28×28 . Then a spatial transformer (Jaderberg et al., 2015) ST will map each digit image onto a 96×96 canvas using the corresponding position z_d^{where} . Finally we sample the frame x_t from a Bernoulli distribution, the parameter of which is the composition of that canvas with some activation function σ .

Compared with joint prediction $\{z_d^{\text{what}}, z_{d,1:T}^{\text{where}}\}$, it is easier to guess the features $\{z_d^{\text{what}}\}$ of digit d given its positions $\{z_{d,1:T}^{\text{where}}\}$ and vice versa. Moreover, we can predict the positions in a step-by-step manner and resample at each sub-step of $\{z_{d,t}^{\text{where}}\}$, which will reduce the variance of the importance weights. Thus we employ a block update strategy that iterates over $T + 1$ blocks

$$\{z_{1:D}^{\text{what}}\}, \{z_{1:D,1}^{\text{where}}\}, \{z_{1:D,2}^{\text{where}}\}, \dots, \{z_{1:D,T}^{\text{where}}\}.$$

In each block we sequentially predict variables for each digit $d = 1, 2, \dots, D$ (see Appendix F for detail).

We train our model on 60000 bouncing MNIST instances, each containing $T = 10$ time steps and $D = 3$ digits, with

$K = 5$ sweeps, $L = 10$ particles, 5 instances per batch, learning rate 10^{-4} , and 1.2×10^7 gradient steps.

Figure 4 shows that the APG sampler substantially improves inference and reconstruction results over RWS encoders, and scales to test instances with hundreds of latent variables. Moreover we compute the mean squared error between the video instances and their reconstructions. Table 2 shows that APG sampler can substantially improve the sample from RWS encoder, and scales up to 100 time steps and 5 digits.

7. Conclusion

We introduce APG samplers, a framework for amortized variational inference based on adaptive importance samplers that iterates between updates to blocks of variables. To appropriately account for the size of the input data we developed a novel parameterization in terms of neural sufficient statistics inspired by sufficient statistics in conjugate exponential families. We show that APG samplers can train structured deep generative models with hundreds of instance-level variables in an unsupervised manner, and compare favorably to existing amortized inference methods in terms of computational efficiency.

APG samplers offer a path towards the development of deep generative models that incorporate structured priors to provide meaningful inductive biases in settings where we have little or no supervision. These methods have particular strengths in problems with global variables, but more generally make it possible to design amortized approaches that exploit conditional independence. Our parameterization in terms of neural sufficient statistics makes it comparatively easy to design models that scale to much larger number of variables and thus generalize to datasets that vary in size.

Acknowledgements

We would like to thank our reviewers and the area chair for their thoughtful comments. This work was supported by the Intel Corporation, NSF award 1835309, startup funds from Northeastern University, the Air Force Research Laboratory (AFRL), and DARPA. Tuan Anh Le was supported by AFOSR award FA9550-18-S-0003.

References

- Beal, M. J. *Variational Algorithms for Approximate Bayesian Inference*. PhD thesis, 2003.
- Bishop, C. M. *Pattern Recognition and Machine Learning*. Springer, New York, 2006. ISBN 978-0-387-31073-2.
- Bornschein, J. and Bengio, Y. Reweighted wake-sleep. *arXiv preprint arXiv:1406.2751*, 2014.
- Brooks, S., Gelman, A., Jones, G., and Meng, X.-L. *Handbook of markov chain monte carlo*. CRC press, 2011.
- Caterini, A. L., Doucet, A., and Sejdinovic, D. Hamiltonian variational auto-encoder. In *Advances in Neural Information Processing Systems*, pp. 8167–8177, 2018.
- Del Moral, P., Doucet, A., and Jasra, A. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, June 2006. doi: 10.1111/j.1467-9868.2006.00553.x.
- Doucet, A., Freitas, N., and Gordon, N. (eds.). *Sequential Monte Carlo Methods in Practice*. Springer New York, New York, NY, 2001. ISBN 978-1-4419-2887-0 978-1-4757-3437-9. doi: 10.1007/978-1-4757-3437-9.
- Eslami, S. M. A., Heess, N., Weber, T., Tassa, Y., Szepesvari, D., kavukcuoglu, k., and Hinton, G. E. Attend, Infer, Repeat: Fast Scene Understanding with Generative Models. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 3225–3233. Curran Associates, Inc., 2016.
- Esmaili, B., Huang, H., Wallace, B., and van de Meent, J.-W. Structured Neural Topic Models for Reviews. In *Artificial Intelligence and Statistics*, pp. 3429–3439, April 2019. 00000.
- Grathwohl, W., Choi, D., Wu, Y., Roeder, G., and Duvenaud, D. Backpropagation through the Void: Optimizing control variates for black-box gradient estimation. *International Conference on Learning Representations*, 2018.
- Hoffman, M. D. Learning deep latent gaussian models with markov chain monte carlo. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1510–1519. JMLR. org, 2017.
- Huang, C.-W., Tan, S., Lacoste, A., and Courville, A. C. Improving explorability in variational inference with annealed variational objectives. In *Advances in Neural Information Processing Systems*, pp. 9701–9711, 2018.
- Jaderberg, M., Simonyan, K., Zisserman, A., et al. Spatial transformer networks. In *Advances in Neural Information Processing Systems*, pp. 2017–2025, 2015.
- Jang, E., Gu, S., and Poole, B. Categorical reparametrization with gumbel-softmax. In *International Conference on Learning Representations 2017*. OpenReviews. net, 2017.
- Johnson, M., Duvenaud, D. K., Wiltchko, A., Adams, R. P., and Datta, S. R. Composing graphical models with neural networks for structured representations and fast inference. In *Advances in Neural Information Processing Systems*, pp. 2946–2954, 2016.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *International Conference on Learning Representations*, 2013.
- Kosiorrek, A., Kim, H., Teh, Y. W., and Posner, I. Sequential Attend, Infer, Repeat: Generative Modelling of Moving Objects. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31*, pp. 8606–8616. Curran Associates, Inc., 2018.
- Le, T. A., Kosiorrek, A. R., Siddharth, N., Teh, Y. W., and Wood, F. Revisiting reweighted wake-sleep for models with stochastic control flow. In *Uncertainty in Artificial Intelligence*, 2019. Le and Kosiorrek contributed equally.
- Li, Y., Turner, R. E., and Liu, Q. Approximate inference with amortised mcmc. *arXiv preprint arXiv:1702.08343*, 2017.
- Maaløe, L., Sønderby, C. K., Sønderby, S. K., and Winther, O. Auxiliary deep generative models. In *International Conference on Machine Learning*, pp. 1445–1453, 2016.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *International Conference on Learning Representations*, 2017.
- Marino, J., Yue, Y., and Mandt, S. Iterative amortized inference. *International Conference on Machine Learning*, 2018.
- Mnih, A. and Rezende, D. Variational Inference for Monte Carlo Objectives. In *International Conference on Machine Learning*, pp. 2188–2196, June 2016.
- Naesseth, C., Lindsten, F., and Schon, T. Nested sequential monte carlo methods. In *International Conference on Machine Learning*, pp. 1292–1301, 2015.

- Ranganath, R., Tran, D., and Blei, D. Hierarchical variational models. In *International Conference on Machine Learning*, pp. 324–333, 2016.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In Xing, E. P. and Jebara, T. (eds.), *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 1278–1286, Beijing, China, June 2014. PMLR.
- Ruiz, F. J. and Titsias, M. K. A contrastive divergence for combining variational inference and mcmc. *arXiv preprint arXiv:1905.04062*, 2019.
- Salimans, T., Kingma, D., and Welling, M. Markov chain monte carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, pp. 1218–1226, 2015.
- Srivastava, N., Mansimov, E., and Salakhudinov, R. Unsupervised learning of video representations using lstms. In *International conference on machine learning*, pp. 843–852, 2015.
- 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.
- Wainwright, M. J. and Jordan, M. I. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008. doi: 10.1561/22000000001.
- Wang, T., Wu, Y., Moore, D., and Russell, S. J. Meta-learning mcmc proposals. In *Advances in Neural Information Processing Systems*, pp. 4146–4156, 2018.
- Weber, T., Heess, N., Buesing, L., and Silver, D. Credit Assignment Techniques in Stochastic Computation Graphs. *arXiv:1901.01761 [cs, stat]*, January 2019.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256, 1992. ISSN 0885-6125. doi: 10.1007/BF00992696.