Slice Sampling Reparameterization Gradients

Anonymous Author(s)
Affiliation
Address
email

Abstract

Many probabilistic modeling problems in machine learning use gradient-based optimization in which the objective takes the form of an expectation. These problems can be challenging when the parameters to be optimized determine the probability distribution under which the expectation is being taken, as the naïve Monte Carlo procedure is not differentiable. Reparameterization gradients make it possible to efficiently perform optimization of these Monte Carlo objectives by transforming the expectation to be differentiable, but the approach is typically limited to distributions with simple forms and tractable normalization constants. Here we describe how to differentiate samples from slice sampling to compute *slice sampling reparameterization gradients*, enabling a richer class of Monte Carlo objective functions to be optimized. Slice sampling is a Markov chain Monte Carlo algorithm for simulating samples from probability distributions; it only requires a density function that can be evaluated point-wise up to a normalization constant, making it applicable to a variety of inference problems and unnormalized models. Our approach is based on the observation that when the slice endpoints are known, the sampling path is a deterministic and differentiable function of the pseudo-random variables, since the algorithm is rejection-free. We evaluate the method on synthetic examples and apply it to a variety of applications with reparameterization of unnormalized probability distributions.

1 Introduction

Probabilistic modeling is a powerful approach to inferring latent structure in complex real-world processes, but often presents computational challenges for inference and further downstream tasks. In many modern probabilistic models, inference is recast as optimization of a probabilistic objective that takes the form of an expectation of a loss function with respect to some distribution. A salient example is variational inference [26, 3], where a simpler distribution is optimized to approximate a posterior distribution by minimizing the Kullback-Leibler (KL) divergence to the truth. In particular, Monte Carlo methods for estimating the gradient of the KL with respect to the parameters have extended the use of variational inference to non-conjugate Bayesian models [41, 42] and neural networks [30]. Such probabilistic objectives also appear in a number of other applications in machine learning and computational science, including the generator loss in generative adversarial networks [14], computing the sensitivity of expectations under the posterior distribution to prior hyperparameters [20, 18, 11], computing the Black-Scholes delta in computational finance [12], and choosing a design that maximizes the probability of improvement in an experiment [54].

Typically, the probabilistic objectives are comprised of expectations that cannot be computed in closed form, so gradients are often estimated via Monte Carlo sampling [37]. Two popular classes of Monte Carlo gradient estimators are the *score function* estimator [31, 13, 53, 41, 42] and the *pathwise* (or "reparameterization gradient") estimator [22, 23, 43, 30, 47]. Score function gradient estimators are general-purpose, as they apply when the underlying density is differentiable and can be sampled

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

from, even if the cost function is not differentiable. However, they often have high variance and are used with variance reduction techniques (e.g., Ranganath et al. [42]). In contrast, reparameterization gradients apply when the loss function is differentiable and samples from the underlying density can be generated by a known deterministic, differentiable transformation of samples from a simpler distribution that does not depend on the model parameters. Typically, reparameterization gradients have lower variance than score function gradients, provided the loss function is sufficiently smooth [37]. Therefore, developing effective reparameterized gradient estimators has been an active area of research in sensitivity and perturbation analysis [22, 23, 12] and stochastic backpropagation [30, 43].

However, reparameterization gradients are primarily limited to distributions with tractable normalizing constants. This precludes the use of complex models of interest such as energy-based models and non-conjugate Bayesian models. Thus, generalizing reparameterization with MCMC to unnormalized distributions is an important direction for developing effective gradient estimators for complicated models. Indeed, some recent work has focused on reparameterization gradients for unnormalized distributions in specialized problems, including Gibbs samplers with reparameterizable conditional sampling steps [51] and dynamics-based MCMC without accept/reject steps (e.g., Salimans et al. [48], Dai et al. [7]). However, developing estimators for general unnormalized distributions using reparameterized gradients and MCMC is not straightforward with existing approaches. For instance, not all Gibbs sampling steps are reparameterizable using current methods, due to, e.g., rejection sampling [10] or "Metropolis-within-Gibbs" sampling [6]. Next, dynamics-based MCMC samplers without accept/reject steps are approximate samplers with asymptotic bias. Finally, a key obstacle in applying MCMC methods with accept/reject steps, such as the Metropolis-Hastings (MH) [21], Metropolis-Adjusted Langevin (MALA) [15], and Hamiltonian Monte Carlo (HMC) [40] algorithms, is that they do not have differentiable sample paths.

An appealing alternative to MCMC methods currently used for reparameterized gradients is slice sampling [39], an auxiliary-variable MCMC method that can be applied to unnormalized probability distributions and does not require an accept/reject step or sensitive tuning parameters. Crucially, the lack of an accept/reject step leads to the key observation that for a fixed pseudo-random sequence, the realized slice sampling Markov chain is differentiable with respect to the model parameters. In this work, we develop reparameterization gradients for samples generated from slice sampling that apply to distributions known only up to a normalizing constant. Slice sampling reparameterization gradients are broadly applicable to complicated multivariate distributions, such as energy-based models (EBMs) [33] and non-conjugate Bayesian models. While the generated samples are correlated and the gradient estimates are biased because we simulate from a finite Markov chain, we demonstrate the efficacy of slice sampling reparameterization gradients in simulations, investigating the bias and variance properties of reparameterized slice sampling in comparison with existing gradient estimators. We then show applications of slice sampling reparameterized gradients to a number of problems in machine learning and statistics in the areas of deep generative modeling, approximate inference, and Bayesian sensitivity analysis.

2 Background

In this work we are concerned with probabilistic objectives that take the form of an expectation of a real-valued function ℓ with respect to a base density $p_\theta(x)$, i.e.,

$$\mathcal{L}(\theta) = \mathbb{E}_{p_\theta(x)}[\ell(x)]\,. \tag{1}$$

Typically the expectation cannot be computed in closed form, so computing the gradient of this objective with respect to the distributional parameters θ requires stochastic estimates of the gradient of the expectation. Two popular classes of stochastic gradient estimators are score function gradients and reparameterization (or pathwise) gradients; see Mohamed et al. [37] for a review.

2.1 Continuous reparameterization gradients

Reparameterization gradients are used for problems with a differentiable loss function ℓ and continuous density p_θ. They apply when samples from p_θ can be generated by a deterministic transformation f_θ of samples from a base distribution $p(\epsilon)$ that does not depend on θ. That is, if $\epsilon \sim p(\epsilon)$ then $x = f_\theta(\epsilon) \sim p_\theta(x)$. Applying this transformation, the gradient of the objective with respect to θ can then be expressed as an expectation with respect to a distribution that does not depend on θ.

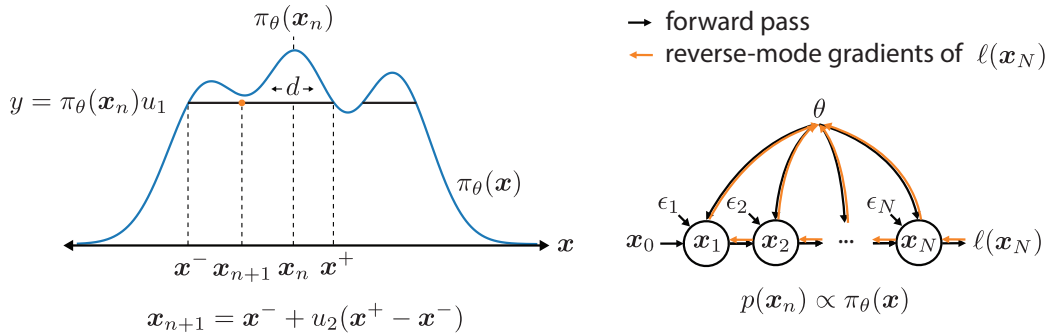

Figure 1: Slice sampling. *Left*: Visualization of one step of a 1D slice sampler. *Right*: Reparameterized slice sampling computational graph with reverse mode gradients shown for a loss computed on the final sample $\boldsymbol{x}_N$.

Assuming that the derivative and integral can be interchanged [37, Section 4.3.1]:

$$\nabla_\theta \mathcal{L}(\theta) = \nabla_\theta \mathbb{E}_{p(\epsilon)}[\ell(f_\theta(\epsilon))] = \mathbb{E}_{p(\epsilon)}[\nabla_\theta \ell(f_\theta(\epsilon))]. \tag{2}$$

Monte Carlo estimates of the gradient are then computed using samples from $p(\epsilon)$.

Many examples of reparameterization gradients for continuous distributions exist. "One-liner" reparameterization gradients use a simple function to transform the base distribution into the desired parameterized distribution [30, 43, 37]. Next, the inverse CDF can be used to transform uniform random variables into arbitrary 1D distributions, and gradients can be computed when evaluating and differentiating the inverse CDF are numerically tractable. Additional examples of reparameterization gradients are implicit reparameterization gradients [9], doubly reparameterized gradients [50], reparameterization of accept/reject sampling [38], and generalized reparameterization gradients [46].

2.2 Markov chain Monte Carlo via slice sampling

For complicated target distributions, direct sampling via, e.g., the inverse CDF, is often not feasible and moreover it is often the case that the target density function can only be evaluated up to an unknown normalization constant; this commonly occurs in Bayesian posterior inference where the marginal likelihood is unavailable. Markov chain Monte Carlo (MCMC) algorithms address both of these challenges. An MCMC algorithm is a recipe for constructing a Markov chain that is easy to simulate and that converges in distribution to the target, gaining computational tractability at the cost of initialization bias and correlation between the samples.

The most common recipe for constructing an MCMC transition operator is the Metropolis–Hastings algorithm [36], which combines an easy-to-simulate proposal distribution with a randomized accept/reject decision that depends only on a ratio of densities, thereby avoiding the need to know the normalization constant. Variations on this theme include Gibbs sampling, in which the proposal is chosen to be the conditional distribution for a subset of the variables, and Hamiltonian Monte Carlo (HMC) [40], in which the proposal is constructed from a fictive dynamical system.

Slice sampling [39] is an alternative approach to MCMC based on the observation that samples taken uniformly from the volume beneath the target PDF have the correct marginal distribution. Slice sampling defines a Markov transition operator that leaves this uniform distribution invariant, and is example of an *auxiliary variable* MCMC method (along with HMC, Swendsen-Wang [49], pseudo-marginal MCMC [1], and others). It is appealing because such transition operators can be constructed without having to identify a proposal distribution and the updates never get "stuck" as can happen with Metropolis–Hastings; it is also generally robust to choices of tuning parameters.

Consider a distribution with density $p(\boldsymbol{x}) = \frac{1}{Z}\pi(\boldsymbol{x})$ where the normalizing constant may not be known. Slice sampling typically proceeds in two steps. Starting from a point $\boldsymbol{x}_n$, a height y_{n+1} is sampled uniformly beneath the density at $\boldsymbol{x}_n$ such that $y_{n+1} \sim \mathcal{U}(0, \pi(\boldsymbol{x}_n))$. The height y_{n+1} defines a "slice" through the probability density given by $S = \{\boldsymbol{x} : y_{n+1} < \pi(\boldsymbol{x})\}$. The next point $\boldsymbol{x}_{n+1}$ is then sampled uniformly from the set S. There are different methods for sampling uniformly

from S, including component-wise updates and construction of hyper-rectangles; here we use *random-direction slice sampling* to sample from the set S (35, Chapter 29.7). A random direction $\boldsymbol{d}$ is generated from a uniform distribution over directions, and the direction defines a line segment through the slice with endpoints $\boldsymbol{x}^-$ and $\boldsymbol{x}^+$. Finally, a value $\boldsymbol{x}_{n+1}$ is sampled uniformly between the endpoints. The most common procedure for finding the endpoints, as proposed by Neal [39], is to use a reversible "stepping-out" procedure followed by "interval shrinking" to determine $\boldsymbol{x}_{n+1}$. In this work we will not use this stepping out procedure and will instead perform a direct search for the slice boundaries. This direct search is key to the present work, as it eliminates the non-differentiable accept/reject step that otherwise appears in most implementations of slice sampling.

3 Slice sampling reparameterization gradients

3.1 Random-direction slice sampling with numerical slice endpoints

Here we describe the steps of random-direction slice sampling in more detail. We consider a continuous, unnormalized density with parameters θ, $\pi_\theta(\boldsymbol{x})$, and start from a point $\boldsymbol{x}_n \in \mathbb{R}^d$. We sample three random quantities: two uniform random numbers $u_1, u_2 \in [0, 1]$ and a uniform random unit vector $\boldsymbol{d} \in \mathbb{R}^d$ ($\boldsymbol{d}$ is sampled by simulating d standard normal random variables and then normalizing). The value u_1 determines the height of the slice $u_1\pi_\theta(\boldsymbol{x}_n)$. The direction $\boldsymbol{d}$ induces a line through the slice. This line intersects with the density at points where $\pi_\theta(\boldsymbol{x}_n + \alpha\boldsymbol{d}) = u_1\pi_\theta(\boldsymbol{x}_n)$ for scalar values α. The intersecting points on the slice S induce the set of α values

$$\mathcal{A} := \{\alpha \in \mathbb{R} : \pi_\theta(\boldsymbol{x}_n + \alpha\boldsymbol{d}) = u_1\pi_\theta(\boldsymbol{x}_n)\}\,. \tag{3}$$

Define the slice endpoints representing the closest intersecting locations on the slice S in the positive and negative directions where $\pi_\theta(\boldsymbol{x}_n + \alpha\boldsymbol{d}) = u_1\pi_\theta(\boldsymbol{x}_n)$ as

$$\boldsymbol{x}^+ := \boldsymbol{x}_n + \alpha^+\boldsymbol{d}, \quad \boldsymbol{x}^- := \boldsymbol{x}_n + \alpha^-\boldsymbol{d}, \qquad \text{where} \quad \alpha^+ := \min_{\alpha>0} \mathcal{A}, \quad \alpha^- := \max_{\alpha<0} \mathcal{A}. \tag{4}$$

The points α^+ and α^- are identified using numerical root finding (Appendix E). The next sample $\boldsymbol{x}_{n+1}$ is then taken uniformly between the endpoints $\boldsymbol{x}^-$ and $\boldsymbol{x}^+$:

$$\boldsymbol{x}_{n+1} = \boldsymbol{x}^- + u_2(\boldsymbol{x}^+ - \boldsymbol{x}^-) = \boldsymbol{x}_n + u_2\alpha^+\boldsymbol{d} + (1 - u_2)\alpha^-\boldsymbol{d}\,. \tag{5}$$

Importantly, α^- and α^+ are implicit functions of $\boldsymbol{x}_n$, $\boldsymbol{d}$, u_1, and θ, denoted by $\alpha^+(\boldsymbol{x}_n, \boldsymbol{d}, u_1, \theta)$, and $\alpha^-(\boldsymbol{x}_n, \boldsymbol{d}, u_1, \theta)$. Via Equation (5), the location $\boldsymbol{x}_{n+1}$ is then a deterministic function of $\boldsymbol{x}_n$, θ, the random variables u_1, u_2, and $\boldsymbol{d}$:

$$\boldsymbol{x}_{n+1} = \boldsymbol{x}_n + u_2\,\alpha^+(\boldsymbol{x}_n, \boldsymbol{d}, u_1, \theta)\,\boldsymbol{d} + (1 - u_2)\,\alpha^-(\boldsymbol{x}_n, \boldsymbol{d}, u_1, \theta)\,\boldsymbol{d}\,. \tag{6}$$

In Figure 1 and later sections, we use ϵ to refer to the random draws u_1, u_2, and $\boldsymbol{d}$.

3.2 Reparameterization gradient

Recall that the goal is to estimate the gradient of an expectation taken with respect to $p_\theta(\boldsymbol{x}) \propto \pi_\theta(\boldsymbol{x})$ with an unknown normalization constant. Consider the following reparameterization of the probabilistic objective $\mathcal{L}(\theta)$ and its Monte Carlo estimator $L(\theta)$, respectively:

$$\begin{aligned}\mathcal{L}(\theta) = \mathbb{E}_{p_\theta(\boldsymbol{x})}[\ell(\boldsymbol{x})] &= \mathbb{E}_{p(\epsilon)}[\ell(\boldsymbol{x}(\epsilon;\theta))] \\ &\approx \frac{1}{N}\sum_{n=1}^{N} \ell(\boldsymbol{x}(\epsilon_n;\theta)) =: L(\theta), \quad \text{where } \epsilon^{(n)} \overset{\text{iid}}{\sim} p(\epsilon),\end{aligned} \tag{7}$$

and where $\boldsymbol{x}(\epsilon;\theta)$ is a deterministic function of ϵ and θ and $p(\epsilon)$ is the uniform distribution on $[0, 1]^2 \times \mathbb{R}^d$. Given an initial sample $\boldsymbol{x}_0$, the slice sampling algorithm generates the samples $\boldsymbol{x}_{1:N} := (\boldsymbol{x}_n)_{n=1}^N$, where $\boldsymbol{x}_n = \boldsymbol{x}(\epsilon^{(n)};\theta)$ is the value of the function established in Equation (5) evaluated at the previous point $\boldsymbol{x}_{n-1}$ and the sample $\epsilon^{(n)} = (u_1^{(n)}, u_2^{(n)}, \boldsymbol{d}^{(n)})$. After the forward pass of generating the samples, we evaluate the objective function and use reverse-mode automatic differentiation (AD) to compute gradients [16, Section 3.2]. In what follows, we derive the reverse-mode gradients, and for notational simplicity, we will drop the explicit dependence of $\boldsymbol{x}_n$ on ϵ and θ. Our derivation uses implicit differentiation to efficiently compute gradients of the slice endpoints, avoiding the need to compute gradients through the numerical root finding algorithm.

Applying the chain rule, we can compute the gradient of $L(\theta)$ with respect to θ via a recurrence relation between the sequentially-dependent samples $\boldsymbol{x}_n$:

$$\nabla_\theta L(\theta) = \sum_{n=1}^{N} [\mathcal{J}_\theta(\boldsymbol{x}_n)]^\mathsf{T} \nabla_{\boldsymbol{x}_n} L, \tag{8}$$

where $\mathcal{J}_\theta(\boldsymbol{x}_n) \in \mathbb{R}^{d\times m}$ is the Jacobian of $\boldsymbol{x}_n$ with respect to $\theta \in \mathbb{R}^m$. Here we observe that the gradient $\nabla_\theta L(\theta)$ depends on the gradients of the loss with respect to each sample, $\nabla_{\boldsymbol{x}_n} L \in \mathbb{R}^d$. For the final sample $\boldsymbol{x}_N$ this value is simply $\nabla_{\boldsymbol{x}_N} L = \frac{1}{N}\nabla_{\boldsymbol{x}_N}\ell(\boldsymbol{x}_N)$. Due to the sequential dependency, for earlier samples $\boldsymbol{x}_n$ where $n = 1, ..., N-1$ the loss gradients need to be backpropagated via

$$\nabla_{\boldsymbol{x}_n} L = \frac{1}{N}\nabla_{\boldsymbol{x}_n}\ell(\boldsymbol{x}_n) + [\mathcal{J}_{\boldsymbol{x}_n}(\boldsymbol{x}_{n+1})]^\mathsf{T}\nabla_{\boldsymbol{x}_{n+1}} L. \tag{9}$$

The computational graph demonstrates how each sample $\boldsymbol{x}_{n+1}$ depends on the previous samples $\boldsymbol{x}_{1:n}$ and the parameters θ (Figure 1). Therefore, to compute gradients with respect to the parameters and/or the initial condition, the Jacobians $\mathcal{J}_{\boldsymbol{x}_n}(\boldsymbol{x}_{n+1})$ and $\mathcal{J}_\theta(\boldsymbol{x}_{n+1})$ are needed. Using Equation (5) for $\boldsymbol{x}_{n+1}$, these Jacobians can be computed indirectly via adjoints, i.e.,

$$\begin{aligned}
\mathcal{J}_{\boldsymbol{x}_n}(\boldsymbol{x}_{n+1}) &= \mathcal{J}_{\boldsymbol{x}_n}\big(\boldsymbol{x}_n + \boldsymbol{d}u_2\alpha^+ + \boldsymbol{d}(1-u_2)\alpha^-\big) = \boldsymbol{I} + u_2\boldsymbol{d}\nabla_{\boldsymbol{x}_n}[\alpha^+]^\mathsf{T} + (1-u_2)\boldsymbol{d}\nabla_{\boldsymbol{x}_n}[\alpha^-]^\mathsf{T} \\
\mathcal{J}_\theta(\boldsymbol{x}_{n+1}) &= \mathcal{J}_\theta\big(\boldsymbol{x}_n + \boldsymbol{d}u_2\alpha^+ + \boldsymbol{d}(1-u_2)\alpha^-\big) = u_2\boldsymbol{d}\nabla_\theta[\alpha^+]^\mathsf{T} + (1-u_2)\boldsymbol{d}\nabla_\theta[\alpha^-]^\mathsf{T}.
\end{aligned} \tag{10}$$

We use implicit differentiation to compute $\nabla_{\boldsymbol{x}_n}[\alpha^+], \nabla_{\boldsymbol{x}_n}[\alpha^-], \nabla_\theta[\alpha^+], \nabla_\theta[\alpha^-]$. The values α are solutions to

$$f(\boldsymbol{x}_n, \boldsymbol{d}, \alpha, \theta) = \ln \pi_\theta(\boldsymbol{x}_n + \alpha\boldsymbol{d}) - \ln u_1 - \ln \pi_\theta(\boldsymbol{x}_n) = 0. \tag{11}$$

Applying implicit differentiation (16, Chap. 15) results in:

$$\begin{aligned}
\nabla_{\boldsymbol{x}_n}\alpha &= -\frac{\nabla_{\boldsymbol{x}_n} f}{\partial f/\partial\alpha} = -\frac{\nabla_{\boldsymbol{x}_n}\ln\pi_\theta(\boldsymbol{x}_n + \alpha\boldsymbol{d}) - \nabla_{\boldsymbol{x}_n}\ln\pi_\theta(\boldsymbol{x}_n)}{\boldsymbol{d}^\mathsf{T}\nabla_{\boldsymbol{x}_n}\ln\pi_\theta(\boldsymbol{x}_n + \alpha\boldsymbol{d})} \\
\nabla_\theta\alpha &= -\frac{\nabla_\theta f}{\partial f/\partial\alpha} = -\frac{\nabla_\theta\ln\pi_\theta(\boldsymbol{x}_n + \alpha\boldsymbol{d}) - \nabla_\theta\ln\pi_\theta(\boldsymbol{x}_n)}{\boldsymbol{d}^\mathsf{T}\nabla_{\boldsymbol{x}_n}\ln\pi_\theta(\boldsymbol{x}_n + \alpha\boldsymbol{d})}.
\end{aligned} \tag{12}$$

After computing $\{\nabla_{\boldsymbol{x}_1} L, \ldots, \nabla_{\boldsymbol{x}_N} L\}$, the gradient $\nabla_\theta L$ is then formed by computing by Equation (8). Importantly, we compute $\nabla_\theta L$ without ever fully representing either of the Jacobians by using vector-Jacobian products (Appendix D). We implemented the forward sampling and reverse mode AD in JAX [4].

3.3 Gradient of the ELBO

In our experiments, we use the reparameterized slice sampler to optimize the ELBO in variational inference with an unnormalized density $q_\phi(\boldsymbol{z}) = \frac{1}{Z(\phi)}e^{f_\phi(\boldsymbol{z})}$ with parameters ϕ as the approximate posterior. The reparameterized ELBO is

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{q_\phi(\boldsymbol{z})}[\log p_\theta(\boldsymbol{x}, \boldsymbol{z}) - \log q_\phi(\boldsymbol{z})] = \mathbb{E}_{p(\epsilon)}[\log p_\theta(\boldsymbol{x}, \boldsymbol{z}(\phi, \epsilon)) - \log q_\phi(\boldsymbol{z}(\phi, \epsilon))] \tag{13}$$

where $\boldsymbol{z}(\phi, \epsilon)$ is a sample generated from the reparameterized slice sampler with unnormalized density $f_\phi(\boldsymbol{z})$ and base random variables ϵ. Estimating the gradient appears to require the gradient of the normalizing constant $Z(\phi)$. However, applying the total derivative [44] yields a path derivative Monte Carlo estimator of the gradient for a sample ϵ that does not require the normalizing constant:

$$\hat{\nabla}_{\mathrm{PD}}(\epsilon, \phi) = \nabla_{\boldsymbol{z}}[\log p_\theta(\boldsymbol{x}, \boldsymbol{z}(\phi, \epsilon)) - f_\phi(\boldsymbol{z}(\phi, \epsilon))]\nabla_\phi \boldsymbol{z}(\phi, \epsilon). \tag{14}$$

See Appendix A for a full derivation. This estimator has favorable performance, and has zero variance when $p_\theta(\boldsymbol{z} \mid \boldsymbol{x}) = q_\phi(\boldsymbol{z} \mid \boldsymbol{x})$ [44].

4 Experiments

Here we apply slice sampling reparameterization gradients to several problems, demonstrating the generality of the approach and the potential appeal of reparameterization gradients for unnormalized distributions. For additional experimental details see the supplementary material (Section B).

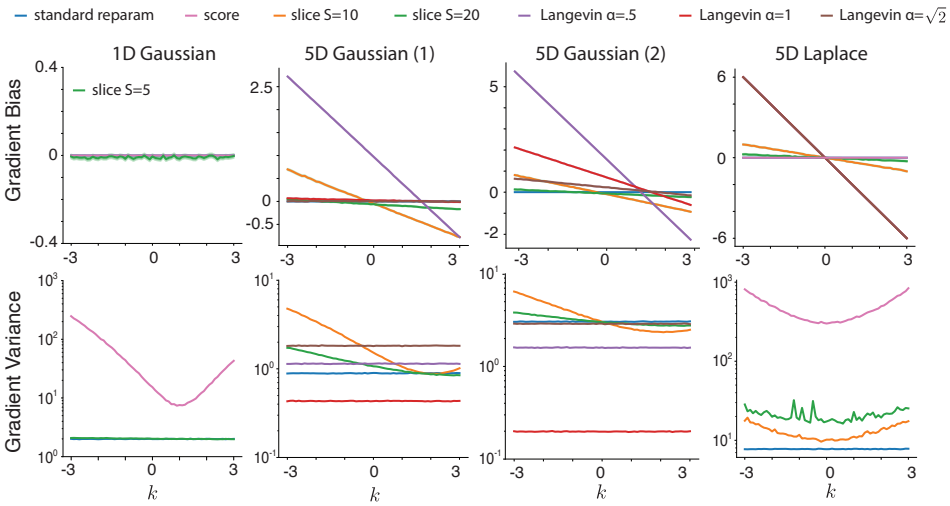

Figure 2: Estimator bias and variance for various Monte Carlo gradient estimators. The gradient is $\nabla_\theta \mathbb{E}_{q(x;\theta)}[\frac{1}{D}\sum_i (x_i - k)^2]$ for $x \in \mathbb{R}^D$. Error bars on bias are ± 2 standard error.

4.1 Bias and variance

Following Mohamed et al. [37], we quantify the empirical bias and variance of slice sampling reparameterization gradients in comparison to standard reparameterization gradients, score function gradients, and reparameterized Langevin dynamics (Figure 2). We estimate $\nabla_\theta \mathbb{E}_{q(\boldsymbol{x};\theta)}[\frac{1}{D}\sum_i (\boldsymbol{x}_i - k)^2]$ with $\boldsymbol{x} \in \mathbb{R}^D$ for multiple distributions q and various values of k. These experiments show that slice sampling gradient estimates have initialization bias for finite sample sizes, which is expected. Importantly, however, we also find that slice sampling reparameterization gradients can be lower variance than score function gradients computed from exact samples from the distribution.

The first example is a one-dimensional Gaussian with $q(x;\theta) = \mathcal{N}(x;\theta,1)$ and $\theta = 1$. In this case, slice sampling mixes well and the resulting reparameterization gradients are nearly unbiased with variance similar to the standard reparameterization estimator. We next focus on comparing slice sampling with Langevin dynamics on multi-dimensional examples. We consider two 5-dimensional multivariate Gaussian distributions $q(\boldsymbol{x};\theta) = \mathcal{N}(\boldsymbol{x};\theta\mathbf{1},\Psi_i)$ with diagonal covariances Ψ_1 and Ψ_2 and a Laplace distribution $q(x;\theta) = \prod_i \text{Laplace}(x_i;\theta,b_i)$. While Langevin with a tuned step size may outperform slice sampling in some cases, we find that the performance of Langevin dynamics varies considerably for the same set of step sizes across the two multivariate Gaussian distributions. On the other hand, slice sampling performance is consistent across each of the examples and improves with more samples. The versatility of slice sampling is highlighted by the Laplace example. In this case the gradient of the mean from reparameterized Langevin dynamics is zero. On the other hand, slice sampling reparameterization gradients have low bias and notably have substantially lower variance the score function gradients computed from exact samples from the distribution.

4.2 Minimizing KL-divergence

Minimizing the KL-divergence $D_{KL}(q_\theta||p)$ w.r.t. θ between a parameterized distribution q_θ and a target p is a core problem of approximate inference. We show that slice sampling reparameterization gradients can be used to successfully optimize this objective in a number of synthetic examples. We present one experiment in the main text with additional experiments in C.1.

To demonstrate the potential utility of unnormalized approximate posteriors, we choose a 2D "banana" distribution that has a complicated dependency between the two dimensions as the target

$$p(\boldsymbol{z}) = \mathcal{N}\left(\begin{bmatrix} \boldsymbol{z}_1 \\ \boldsymbol{z}_2 + \boldsymbol{z}_1^2 + 1 \end{bmatrix} \middle| \mathbf{0}, \begin{bmatrix} 1 & 0.9 \\ 0.9 & 1 \end{bmatrix}\right). \tag{15}$$

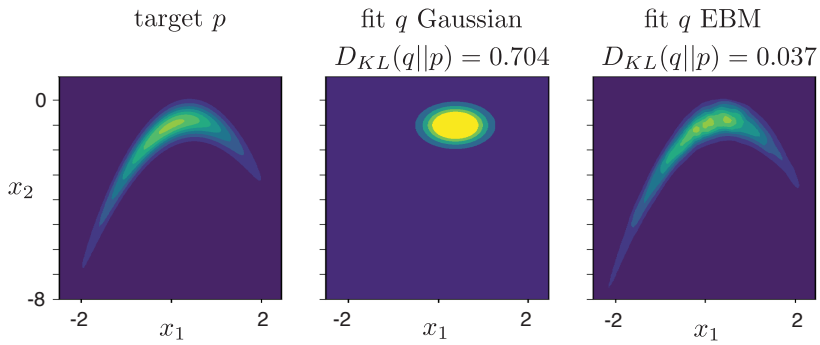

Figure 3: The target p (*left*) and fit mean-field Gaussian (*middle*) and EBM (*right*) approximate posteriors. Each panel is shown with the same color scale.

We use an EBM approximate posterior

$$q_\phi(\boldsymbol{z}) = \frac{1}{Z(\phi)} \exp(f_\phi(\boldsymbol{z}))\mathcal{N}(\boldsymbol{z}; \mu_\phi, \sigma_\phi^2) \tag{16}$$

where f_ϕ is a feed-forward neural network mapping from the two-dimensional input space to a scalar energy. The Gaussian term is included in the model to ensure integrability. The network had 2 hidden layers of 256 units and leaky ReLU nonlinearities. We optimized θ by minimizing $D_{KL}(q_\theta||p)$ using slice sampling reparameterization gradients and Adam [29].

The fit and target distributions are shown in Figure 3 along with a comparison to fitting a mean-field Gaussian to the same objective. Qualitatively, the EBM approximate posterior appears remarkably similar to the target. We estimated the KL divergence between the target p and the approximate posteriors q using Monte Carlo sampling. We find that the KL divergence between the EBM approximate posterior and the target approaches zero and is much lower than the KL divergence between the target and the mean-field Gaussian posterior.

4.3 Variational contrastive divergence with fully reparameterized gradients

Combining variational inference with MCMC methods is an active area of research [48, 25, 3, 45]. One approach to combine the methods is to use samples from a variational posterior to initialize an MCMC chain targeting the true posterior. The samples from the MCMC chain will be closer in KL divergence to the true posterior than the initial samples from the variational posterior, hence making the ELBO tighter. However, there are two challenges for this optimization. First, the dependence on the variational parameters decreases as the length of the MCMC chain decreases. Second, the distribution of the generated samples is not available in closed form.

Ruiz and Titsias [45] propose the variational contrastive divergence (VCD) objective for combining variational inference with MCMC to address these two problems. Let the approximate posterior be $q_\theta(\boldsymbol{z})$, and let $q_\theta^{(t)}(\boldsymbol{z})$ be the distribution given by sampling an initial value $\boldsymbol{z}^{(0)} \sim q_\theta(\boldsymbol{z})$ and then running t steps of an MCMC chain targeting the true posterior $p(\boldsymbol{z} \mid \boldsymbol{x})$. Then the VCD objective is

$$\mathcal{L}_{\text{VCD}}(\theta) = -\mathbb{E}_{q_\theta(\boldsymbol{z})}[\log p(\boldsymbol{x}, \boldsymbol{z}) - \log q_\theta(\boldsymbol{z})] + \mathbb{E}_{q_\theta^{(t)}(\boldsymbol{z})}[\log p(\boldsymbol{x}, \boldsymbol{z}) - \log q_\theta(\boldsymbol{z})]. \tag{17}$$

Notably, this does not require evaluating $q_\theta^{(t)}$ and instead only requires sampling from it. However, The second term of the VCD objective is an expectation with respect to MCMC samples, and Ruiz and Titsias [45] use a score function estimator to estimate this gradient. With reparameterized slice sampling we can use fully reparameterized gradients to optimize the VCD objective.

Table 1: Estimated heldout log likelihoods.

method	MNIST	Fashion MNIST
VCD [45]	-95.86	**-117.65**
VCD w/ Reparam (Ours)	**-93.99**	-118.02

We fit a variational autoencoder (VAE) [30, 43] using the VCD objective to MNIST [32] and Fashion MNIST [55]. Our network architectures matched those used in Ruiz and Titsias [45]: the latent dimensionality was 10 and both the encoder and decoder had two hidden layers with 200 units and ReLU nonlinearities. The model was optimized for 400K iterations using ADAM with batch sizes of 100. We also estimated the held-out log-likelihood via the importance sampling approach used in Ruiz and Titsias [45].

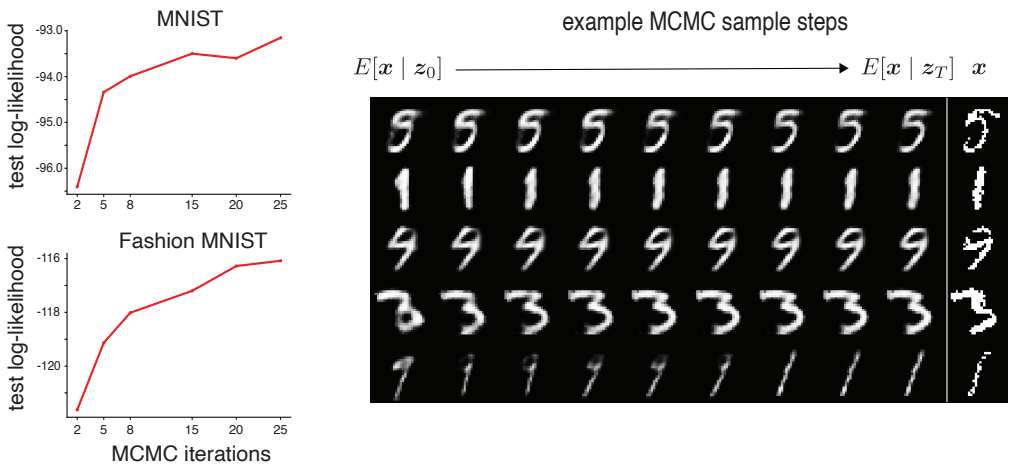

Figure 4: (Left) held out log-likelihood for MNIST and Fashion MNIST as a function of the number of MCMC iterations. (Right) The mean observation conditioned on sample $\boldsymbol{z}_i$ for $\boldsymbol{z}_0$ sampled from the variational posterior and for T samples targeting the true posterior.

In Table 1, we compare the held-out log likelihoods on MNIST and Fashion MNIST for our method with slice sampling reparameterization gradients and the results reported in Ruiz and Titsias [45] for $S = 8$ samples. We find that slice sampling reparameterization gradients perform comparably or better than the original VCD results, with each method having a higher test log likelihood on one of the datasets. However, as a function of the number of MCMC iterations (Figure 4), the test log likelihood for slice sampling reparameterization gradients appears to scale better than the results obtained in Ruiz and Titsias [45] (Figure 2 in that paper).

We also demonstrate that the learned model and slice sampling procedure provide a meaningful result. In Figure 4, we show the expected value of the observation given a sample from the variational posterior and how 8 steps of slice sampling towards the target posterior improve the latent sample. The resulting conditional mean estimate for $\boldsymbol{z}_T$ appears more similar to the true observation $\boldsymbol{x}$ than the estimate for $\boldsymbol{z}_0$. For example, in the bottom row the initial sample appears more like a nine but is transformed to a one, which is the true observation.

4.4 Conditional EBM approximate posterior

We next used slice sampling reparameterization gradients to fit a VAE with a conditional EBM approximate posterior:

$$q(\boldsymbol{z} \mid \boldsymbol{x}) = \frac{1}{Z(\theta, \boldsymbol{x})} \exp\{f_\phi(\boldsymbol{x})^\top g_\phi(\boldsymbol{z}) + \log p_0(\boldsymbol{z})\}. \tag{18}$$

The observations are $\boldsymbol{x}$, the latent variables are $\boldsymbol{z} \in \mathbb{R}^D$, and $\log p_0(\boldsymbol{z})$ is a prior. The functions $f_\phi(\boldsymbol{x})$ and $g_\phi(\boldsymbol{z})$ map to vectors with an embedding dimensionality D_e. This differs from the standard VAE in that the encoder f does not map to the mean of the latent space, but rather a separate embedding that is nonlinearly combined by $g(\boldsymbol{z})$ to give the energy. The form of the conditional EBM has been studied previously in Khemakhem et al. [27, 28]. We fit the VAE with a conditional EBM approximate posterior to MNIST data (Figure 5) with $D = 20$ and $D_e = 50$. Unconditional samples from the generative model look reasonable. We visualized the embedding of held out test images by the function f using t-SNE [34]. Interestingly, we see clusters in the embedding space corresponding to digits, with some overlap for similar digits.

4.5 Adversarial training of EBMs

We next apply our method to training generative adversarial networks (GANs) [14] with EBM generators. Let the generator distribution be $p_\theta(\boldsymbol{x}) \propto e^{f_\theta(\boldsymbol{x})}$. Standard GANs define an implicit distribution over the data points, whereas an EBM defines an unnormalized log-likelihood in the data space. We optimize the Wasserstein GAN [2] objective with the gradient penalty from [17]:

$$\min_\theta \max_{D \in \mathcal{D}} \mathbb{E}_{p_d(\boldsymbol{x})}[D(\boldsymbol{x})] - \mathbb{E}_{p_\theta(\boldsymbol{x})}[D(\boldsymbol{x})] + \lambda\, \mathbb{E}_{p(\hat{\boldsymbol{x}})}[(\|\nabla_{\hat{\boldsymbol{x}}} D(\hat{\boldsymbol{x}})\|_2 - 1)^2] \tag{19}$$

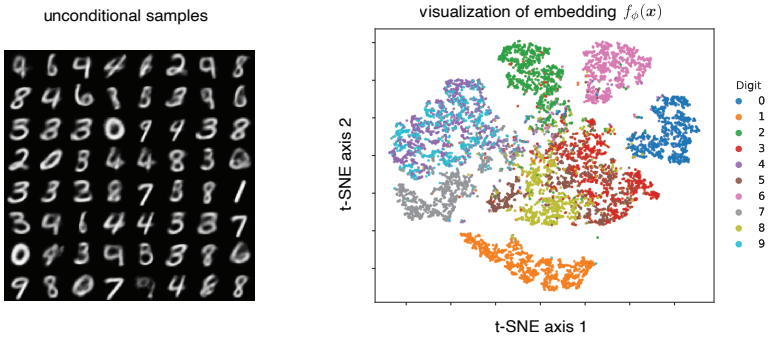

Figure 5: Unconditional samples (*left*) and embedding visualization (*right*) of a VAE fit to MNIST with an EBM approximate posterior.

where $\hat{\boldsymbol{x}}$ are sampled uniformly between pairs of real and generated data points [17] and D is the discriminator. We trained two dimensional EBMs using this approach on three synthetic datasets, and we find that the distribution of generated samples from the learned energy qualitatively match the training distribution (Figure 6).

4.6 Bayesian sensitivity analysis

The sensitivity of expectations under Bayesian posteriors to prior hyperparameters is a common quantity of interest in Bayesian statistics [20, 18]. One formalization of this sensitivity to the prior is as the gradient of the expectation with respect to the prior hyperparameters. Formally, the *local sensitivity* of an expectation $\mathbb{E}_{p_\alpha}[g(\theta)]$ at a point α_0 is $\nabla_\alpha \mathbb{E}_{p_\alpha}[g(\theta)]|_{\alpha=\alpha_0}$, where p_α is the posterior $p_\alpha = p(\theta \mid \mathcal{D}, \alpha)$ [19]. Typically the normalizing constant of the posterior is unknown and score function gradient approaches are needed to estimate this gradient.

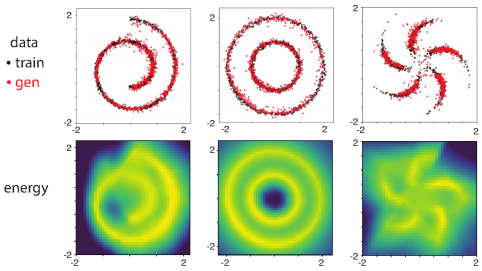

Figure 6: (top) Training and learned generative samples and (bottom) learned energy function for the Swiss roll, two circles, and spirals datasets.

Baseball Example We apply slice sampling reparameterization gradients to estimate the local sensitivity in a non-conjugate hierarchical model of baseball hit probabilities [8, 5]. The model is

$$\phi \sim U[0,1] \quad \kappa \sim \text{Pareto}(1,\alpha) \quad \theta_j \sim \text{Beta}(\phi \cdot \kappa, (1-\phi)\cdot\kappa) \quad y_j \sim \text{Binomial}(K_j, \theta_j), \quad (20)$$

where y_j is the number of successes (hits) for player j in K_j attempts (at bats). Using slice sampling reparameterization gradients, we estimate $\frac{d}{d\alpha}\mathbb{E}_{p_\alpha}[\phi]$ and $\frac{d}{d\alpha}\mathbb{E}_{p_\alpha}[\phi^2]$. Additionally, we estimate the gradient of the point-wise log predictive density on held-out data given the posterior samples [52], with the motivation that hyperparameters for Bayesian models could be learned via this approach. Due to space constraints these results are in the supplementary material Appendix B.

5 Discussion

In this work we developed slice sampling reparameterized gradients and demonstrated the method on a variety of applications involving approximate inference, energy-based models, GANs, and Bayesian statistics. In future work, we are interested in combining reparameterized slice sampling with normalizing flows to improve the geometry of the sampling space [24], developing reparameterization gradients for other classes of slice samplers, and continuing to explore applications of reparameterization gradients of unnormalized distributions. Finally, we would like to investigate how the Jacobian $\mathcal{J}_{\boldsymbol{x}_0}(\boldsymbol{x}_n)$ relates to convergence of the slice sampler.

Societal Impact The potential negative societal impacts of our work are broadly the harm that can be done using reparameterization gradient estimators to optimize models. One example would be optimizing GANs for harmful applications such as "DeepFakes."

References

[1] C. Andrieu and G. O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009.

[2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.

[3] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.

[4] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, and S. Wanderman-Milne. JAX: composable transformations of Python+NumPy programs, 2018. URL `http://github.com/google/jax`.

[5] B. Carpenter, J. Gabry, and B. Goodrich. Hierarchical partial pooling for repeated binary trials. *URL https://cran. r-project. org/web/packages/rstanarm/vignettes/pooling. html [Accessed 16 June 2017]*, 2017.

[6] S. Chib and E. Greenberg. Understanding the Metropolis-Hastings algorithm. *The American Statistician*, 49(4):327–335, 1995.

[7] B. Dai, Z. Liu, H. Dai, N. He, A. Gretton, L. Song, and D. Schuurmans. Exponential family estimation via adversarial dynamics embedding. In *Advances in Neural Information Processing Systems*, pages 10979–10990, 2019.

[8] B. Efron and C. Morris. Data analysis using Stein's estimator and its generalizations. *Journal of the American Statistical Association*, 70(350):311–319, 1975.

[9] M. Figurnov, S. Mohamed, and A. Mnih. Implicit reparameterization gradients. In *Advances in Neural Information Processing Systems*, pages 441–452, 2018.

[10] W. R. Gilks, N. G. Best, and K. Tan. Adaptive rejection Metropolis sampling within Gibbs sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 44(4):455–472, 1995.

[11] R. Giordano, T. Broderick, and M. I. Jordan. Covariances, robustness and variational Bayes. *Journal of Machine Learning Research*, 19(51), 2018.

[12] P. Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer Science & Business Media, 2013.

[13] P. W. Glynn. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM*, 33(10):75–84, 1990.

[14] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27:2672–2680, 2014.

[15] U. Grenander and M. I. Miller. Representations of knowledge in complex systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 56(4):549–581, 1994.

[16] A. Griewank and A. Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, 2008.

[17] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017.

[18] P. Gustafson. Local sensitivity of posterior expectations. *Annals of Statistics*, 24(1):174–195, 1996.

[19] P. Gustafson. *Local Robustness in Bayesian Analysis*, pages 71–88. Springer New York, New York, NY, 2000.

[20] P. Gustafson. Local robustness in Bayesian analysis. In *Robust Bayesian Analysis*, pages 71–88. Springer, 2000.

[21] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. 1970.

[22] Y. Ho and X. Cao. Optimization and perturbation analysis of queueing networks. *Journal of Optimization Theory and Applications*, 40(4):559–582, 1983.

[23] Y.-C. L. Ho and X.-R. Cao. *Perturbation analysis of discrete event dynamic systems*, volume 145. Springer Science & Business Media, 2012.

[24] M. Hoffman, P. Sountsov, J. V. Dillon, I. Langmore, D. Tran, and S. Vasudevan. Neutralizing bad geometry in Hamiltonian Monte Carlo using neural transport. *arXiv preprint arXiv:1903.03704*, 2019.

[25] M. D. Hoffman. Learning deep latent Gaussian models with Markov chain Monte Carlo. In *International Conference on Machine Learning*, pages 1510–1519, 2017.

[26] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.

[27] I. Khemakhem, D. Kingma, R. Monti, and A. Hyvarinen. Variational autoencoders and nonlinear ICA: A unifying framework. In *International Conference on Artificial Intelligence and Statistics*, pages 2207–2217. PMLR, 2020.

[28] I. Khemakhem, R. P. Monti, D. P. Kingma, and A. Hyvärinen. ICE-BeeM: Identifiable conditional energy-based deep models. *arXiv preprint arXiv:2002.11537*, 2020.

[29] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[30] D. P. Kingma and M. Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[31] J. P. Kleijnen and R. Y. Rubinstein. Optimization and sensitivity analysis of computer simulation models by the score function method. *European Journal of Operational Research*, 88(3):413–427, 1996.

[32] Y. LeCun and C. Cortes. MNIST handwritten digit database. 2010. URL `http://yann.lecun.com/exdb/mnist/`.

[33] Y. LeCun, S. Chopra, R. Hadsell, M. Ranzato, and F. Huang. A tutorial on energy-based learning. *Predicting Structured Data*, 1(0), 2006.

[34] L. v. d. Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

[35] D. J. C. MacKay. *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, 2002. ISBN 0521642981.

[36] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6): 1087–1092, 1953.

[37] S. Mohamed, M. Rosca, M. Figurnov, and A. Mnih. Monte carlo gradient estimation in machine learning. *Journal of Machine Learning Research*, 21(132):1–62, 2020.

[38] C. Naesseth, F. Ruiz, S. Linderman, and D. Blei. Reparameterization gradients through acceptance-rejection sampling algorithms. In *Artificial Intelligence and Statistics*, pages 489–498. PMLR, 2017.

[39] R. M. Neal. Slice sampling. *Annals of statistics*, pages 705–741, 2003.

[40] R. M. Neal et al. Mcmc using Hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11):2, 2011.

[41] J. Paisley, D. Blei, and M. Jordan. Variational Bayesian inference with stochastic search. *arXiv preprint arXiv:1206.6430*, 2012.

[42] R. Ranganath, S. Gerrish, and D. Blei. Black box variational inference. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, pages 814–822, 2014.

[43] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014.

[44] G. Roeder, Y. Wu, and D. K. Duvenaud. Sticking the landing: Simple, lower-variance gradient estimators for variational inference. In *Advances in Neural Information Processing Systems*, pages 6925–6934, 2017.

[45] F. Ruiz and M. Titsias. A contrastive divergence for combining variational inference and MCMC. In *International Conference on Machine Learning*, pages 5537–5545, 2019.

[46] F. J. Ruiz, M. K. Titsias, and D. M. Blei. The generalized reparameterization gradient. In *Advances in Neural Information Processing Systems*, pages 460–468, 2016.

[47] T. Salimans, D. A. Knowles, et al. Fixed-form variational posterior approximation through stochastic linear regression. *Bayesian Analysis*, 8(4):837–882, 2013.

[48] T. Salimans, D. Kingma, and M. Welling. Markov chain Monte Carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, pages 1218–1226, 2015.

[49] R. H. Swendsen and J.-S. Wang. Nonuniversal critical dynamics in Monte Carlo simulations. *Physical review letters*, 58(2):86, 1987.

[50] G. Tucker, D. Lawson, S. Gu, and C. J. Maddison. Doubly reparameterized gradient estimators for Monte Carlo objectives. In *International Conference on Learning Representations*, 2018.

[51] A. Vahdat, E. Andriyash, and W. G. Macready. Undirected graphical models as approximate posteriors. In *International Conference on Machine Learning (ICML)*, 2020.

[52] A. Vehtari, A. Gelman, and J. Gabry. Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and computing*, 27(5):1413–1432, 2017.

[53] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.

[54] J. Wilson, F. Hutter, and M. Deisenroth. Maximizing acquisition functions for Bayesian optimization. In *Advances in Neural Information Processing Systems*, pages 9884–9895, 2018.

[55] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] We note potential negative impacts in the discussion.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code is in the supplemental material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] In some cases error bars are included, in others not.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] In supplementary material.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We cited JAX and the MNIST and FashionMNIST datasets.
   (b) Did you mention the license of the assets? [No]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]