
Operator Variational Inference

Rajesh Ranganath
Princeton University

Jaan Altosaar
Princeton University

Dustin Tran
Columbia University

David M. Blei
Columbia University

Abstract

Variational inference is an umbrella term for algorithms which cast Bayesian inference as optimization. Classically, variational inference uses the Kullback-Leibler divergence to define the optimization. Though this divergence has been widely used, the resultant posterior approximation can suffer from undesirable statistical properties. To address this, we reexamine variational inference from its roots as an optimization problem. We use *operators*, or functions of functions, to design variational objectives. As one example, we design a variational objective with a Langevin-Stein operator. We develop a black box algorithm, operator variational inference (OPVI), for optimizing any operator objective. Importantly, operators enable us to make explicit the statistical and computational tradeoffs for variational inference. We can characterize different properties of variational objectives, such as objectives that admit *data subsampling*—allowing inference to scale to massive data—as well as objectives that admit *variational programs*—a rich class of posterior approximations that does not require a tractable density. We illustrate the benefits of OPVI on a mixture model and a generative model of images.

1 Introduction

Variational inference is an umbrella term for algorithms that cast Bayesian inference as optimization [12]. Originally developed in the 1990s, recent advances in variational inference have scaled Bayesian computation to massive data [9], provided black box strategies for generic inference in many models [24], and enabled more accurate approximations of a model’s posterior without sacrificing efficiency [26, 25]. These innovations have both scaled Bayesian analysis and removed the analytic burdens that have traditionally taxed its practice.

Given a model of latent and observed variables $p(\mathbf{x}, \mathbf{z})$, variational inference posits a family of distributions over its latent variables and then finds the member of that family closest to the posterior, $p(\mathbf{z} | \mathbf{x})$. This is typically formalized as minimizing a Kullback-Leibler (KL) divergence from the approximating family $q(\cdot)$ to the posterior $p(\cdot)$. However, while the $\text{KL}(q \| p)$ objective offers many beneficial computational properties, it is ultimately designed for convenience; it sacrifices many desirable statistical properties of the resultant approximation.

When optimizing KL, there are two issues with the posterior approximation that we highlight. First, it typically underestimates the variance of the posterior. Second, it can result in degenerate solutions that zero out the probability of certain configurations of the latent variables. While both of these issues can be partially circumvented by using more expressive approximating families, they ultimately stem from the choice of the objective. Under the KL divergence, we pay a large price when $q(\cdot)$ is big where $p(\cdot)$ is tiny; this price becomes infinite when $q(\cdot)$ has larger support than $p(\cdot)$.

In this paper, we revisit variational inference from its core principle as an optimization problem. We use *operators*—mappings from functions to functions—to design variational objectives, explicitly trading off computational properties of the optimization with statistical properties of the approximation. We use operators to formalize the basic properties needed for variational inference algorithms. We further outline how to use them to define new variational objectives; as one example, we design a variational objective using a Langevin-Stein operator.

We develop operator variational inference (OPVI), a black box algorithm that optimizes any operator objective. In the context of OPVI, we show that the Langevin-Stein objective enjoys two good properties. First, it is amenable to *data subsampling*, which allows inference to scale to massive data. Second, it permits rich approximating families, called *variational programs*, which do not require analytically tractable densities. This greatly expands the class of variational families and the fidelity of the resulting approximation. (We note that the traditional KL is not amenable to using variational programs.) We study OPVI with the Langevin-Stein objective on a mixture model and a generative model of images.

Related Work. There are several threads of research in variational inference with alternative divergences. An early example is expectation propagation (EP) [19]. EP promises approximate minimization of the inclusive KL divergence $\text{KL}(p||q)$ to find overdispersed approximations to the posterior. EP hinges on local minimization with respect to subsets of data and connects to work on α -divergence minimization [20, 8]. However, it does not have convergence guarantees and typically does not minimize KL or an α -divergence because it is not a global optimization method. We note that these divergences can be written as operator variational objectives, but they do not satisfy the tractability criteria and thus require further approximations. Li and Turner [17] present a variant of α -divergences that satisfy the full requirements of OPVI. Score matching [11], a method for estimating models by matching the score function of one distribution to another that can be sampled, also falls into the class of objectives we develop.

Here we show how to construct new objectives, including some not yet studied. We make explicit the requirements to construct objectives for variational inference. Finally, we discuss further properties that make them amenable to both scalable and flexible variational inference.

2 Operator Variational Objectives

We define operator variational objectives and the conditions needed for an objective to be useful for variational inference. We develop a new objective, the Langevin-Stein objective, and show how to place the classical KL into this class. In the next section, we develop a general algorithm for optimizing operator variational objectives.

2.1 Variational Objectives

Consider a probabilistic model $p(\mathbf{x}, \mathbf{z})$ of data \mathbf{x} and latent variables \mathbf{z} . Given a data set \mathbf{x} , approximate Bayesian inference seeks to approximate the posterior distribution $p(\mathbf{z} | \mathbf{x})$, which is applied in all downstream tasks. Variational inference posits a family of approximating distributions $q(\mathbf{z})$ and optimizes a divergence function to find the member of the family closest to the posterior.

The divergence function is the *variational objective*, a function of both the posterior and the approximating distribution. Useful variational objectives hinge on two properties: first, optimizing the function yields a good posterior approximation; second, the problem is tractable when the posterior distribution is known up to a constant.

The classic construction that satisfies these properties is the evidence lower bound (ELBO),

$$\mathbb{E}_{q(\mathbf{z})}[\log p(\mathbf{x}, \mathbf{z}) - \log q(\mathbf{z})]. \quad (1)$$

It is maximized when $q(\mathbf{z}) = p(\mathbf{z} | \mathbf{x})$ and it only depends on the posterior distribution up to a tractable constant, $\log p(\mathbf{x}, \mathbf{z})$. The ELBO has been the focus in much of the classical literature. Maximizing the ELBO is equivalent to minimizing the KL divergence to the posterior, and the expectations are analytic for a large class of models [5].

2.2 Operator Variational Objectives

We define a new class of variational objectives, *operator variational objectives*. An operator objective has three components. The first component is an operator $O^{p,q}$ that depends on $p(\mathbf{z} | \mathbf{x})$ and $q(\mathbf{z})$. (Recall that an operator maps functions to other functions.) The second component is a family of test functions \mathcal{F} , where each $f(\mathbf{z}) \in \mathcal{F}$ maps realizations of the latent variables to real vectors \mathbb{R}^d . In the objective, the operator and a function will combine in an expectation $\mathbb{E}_{q(\mathbf{z})}[(O^{p,q} f)(\mathbf{z})]$, designed such that values close to zero indicate that q is close to p . The third component is a distance

function $t(a) : \mathbb{R} \rightarrow [0, \infty)$, which is applied to the expectation so that the objective is non-negative. (Our example uses the square function $t(a) = a^2$.)

These three components combine to form the operator variational objective. It is a non-negative function of the variational distribution,

$$\mathcal{L}(q; O^{p,q}, \mathcal{F}, t) = \sup_{f \in \mathcal{F}} t(\mathbb{E}_{q(z)}[(O^{p,q} f)(z)]). \quad (2)$$

Intuitively, it is the worst-case expected value among all test functions $f \in \mathcal{F}$. Operator variational inference seeks to minimize this objective with respect to the variational family $q \in \mathcal{Q}$.

We use operator objectives for posterior inference. This requires two conditions on the operator and function family.

1. *Closeness*. The minimum of the variational objective is at the posterior, $q(z) = p(z | \mathbf{x})$. We meet this condition by requiring that $\mathbb{E}_{p(z | \mathbf{x})}[(O^{p,p} f)(z)] = 0$ for all $f \in \mathcal{F}$. Thus, optimizing the objective will produce $p(z | \mathbf{x})$ if it is the only member of \mathcal{Q} with zero expectation (otherwise it will produce a distribution in the equivalence class: $q \in \mathcal{Q}$ with zero expectation). In practice, the minimum will be the closest member of \mathcal{Q} to $p(z | \mathbf{x})$.
2. *Tractability*. We can calculate the variational objective up to a constant without involving the exact posterior $p(z | \mathbf{x})$. In other words, we do not require calculating the normalizing constant of the posterior, which is typically intractable. We meet this condition by requiring that the operator $O^{p,q}$ —originally in terms of $p(z | \mathbf{x})$ and $q(z)$ —can be written in terms of $p(\mathbf{x}, z)$ and $q(z)$. Tractability also imposes conditions on \mathcal{F} : it must be feasible to find the supremum. Below, we satisfy this by defining a parametric family for \mathcal{F} that is amenable to stochastic optimization.

[Equation 2](#) and the two conditions provide a mechanism to design meaningful variational objectives for posterior inference. Operator variational objectives try to match expectations with respect to $q(z)$ to those with respect to $p(z | \mathbf{x})$.

2.3 Understanding Operator Variational Objectives

Consider operators where $\mathbb{E}_{q(z)}[(O^{p,q} f)(z)]$ only takes positive values. In this case, distance to zero can be measured with the identity $t(a) = a$, so tractability implies the operator need only be known up to a constant. This family includes tractable forms of familiar divergences like the KL divergence (ELBO), Rényi’s α -divergence [17], and the χ -divergence [22].

When the expectation can take positive or negative values, operator variational objectives are closely related to Stein divergences [2]. Consider a family of scalar test functions \mathcal{F}^* that have expectation zero with respect to the posterior, $\mathbb{E}_{p(z | \mathbf{x})}[f^*(z)] = 0$. Using this family, a *Stein divergence* is

$$D_{\text{Stein}}(p, q) = \sup_{f^* \in \mathcal{F}^*} |\mathbb{E}_{q(z)}[f^*(z)] - \mathbb{E}_{p(z | \mathbf{x})}[f^*(z)]|.$$

Now recall the operator objective of [Equation 2](#). The closeness condition implies that

$$\mathcal{L}(q; O^{p,q}, \mathcal{F}, t) = \sup_{f \in \mathcal{F}} t(\mathbb{E}_{q(z)}[(O^{p,q} f)(z)] - \mathbb{E}_{p(z | \mathbf{x})}[(O^{p,p} f)(z)]).$$

In other words, operators with positive or negative expectations lead to Stein divergences with a more generalized notion of distance.

2.4 Langevin-Stein Operator Variational Objective

We developed the operator variational objective. It is a class of tractable objectives, each of which can be optimized to yield an approximation to the posterior. An operator variational objective is built from an operator, function class, and distance function to zero. We now use this construction to design a new type of variational objective.

An operator objective involves a class of functions that has known expectations with respect to an intractable distribution. There are many ways to construct such classes [1, 2]. Here, we construct an operator objective from the generator Stein’s method applied to the Langevin diffusion.

Let $\nabla^\top f$ denote the divergence of a vector-valued function f , that is, the sum of its individual gradients. Applying the generator method of Barbour [2] to Langevin diffusion or by simply summing the Stein equation one dimension at a time gives the operator

$$(O_{\text{LS}}^p f)(z) = \nabla_z \log p(x, z)^\top f(z) + \nabla^\top f. \quad (3)$$

We call this the Langevin-Stein (LS) operator. See also Mira et al. [21], Oates et al. [23], Gorham and Mackey [7] for related derivations. We obtain the corresponding variational objective by using the squared distance function and substituting Equation 3 into Equation 2,

$$\mathcal{L}(q; O_{\text{LS}}^p, \mathcal{F}) = \sup_{f \in \mathcal{F}} (\mathbb{E}_q[\nabla_z \log p(x, z)^\top f(z) + \nabla^\top f])^2. \quad (4)$$

The LS operator satisfies both conditions. First, it satisfies closeness because it has expectation zero under the posterior (Appendix A) and its unique minimizer is the posterior (Appendix B). Second, it is tractable because it requires only the joint distribution. The functions f will also be a parametric family, which we detail later.

Additionally, while the KL divergence finds variational distributions that underestimate the variance, the LS objective does not suffer from that pathology. The reason is that KL is infinite when the support of q is larger than p ; here this is not the case.

We provided one example of a variational objectives using operators, which is specific to continuous variables. In general, operator objectives are not limited to continuous variables; Appendix C describes an operator for discrete variables.

2.5 The KL Divergence as an Operator Variational Objective

Finally, we demonstrate how classical variational methods fall inside the operator family. For example, traditional variational inference minimizes the KL divergence from an approximating family to the posterior [12]. This can be construed as an operator variational objective,

$$(O_{\text{KL}}^{p,q} f)(z) = \log q(z) - \log p(z|x) \quad \forall f \in \mathcal{F}. \quad (5)$$

This operator does not use the family of functions—it trivially maps all functions f to the same function. Further, because KL is strictly positive, we use the identity distance $t(a) = a$.

The operator satisfies both conditions. It satisfies closeness because $\text{KL}(p||p) = 0$. It satisfies tractability because it can be computed up to a constant when used in the operator objective of Equation 2. Tractability comes from the fact that $\log p(z|x) = \log p(z, x) - \log p(x)$.

3 Operator Variational Inference

We described operator variational objectives, a broad class of objectives for variational inference. We now examine how it can be optimized. We develop a black box algorithm [33, 24] based on Monte Carlo estimation and stochastic optimization. Our algorithm applies to a general class of models and any operator objective.

Minimizing the operator objective involves two optimizations: minimizing the objective with respect to the approximating family \mathcal{Q} and maximizing the objective with respect to the function class \mathcal{F} (which is part of the objective).

We index the family \mathcal{Q} with *variational parameters* λ and require that it satisfies properties typically assumed by black box methods [24]: the variational distribution $q(z; \lambda)$ has a known and tractable density; we can sample from $q(z; \lambda)$; and we can tractably compute the score function $\nabla_\lambda \log q(z; \lambda)$. We index the function class \mathcal{F} with parameters θ , and require that $f_\theta(\cdot)$ is differentiable. In the experiments, we use neural networks, which are flexible enough to approximate a general family of test functions [10].

Given parameterizations of the variational family and test family, operator variational inference (OPVI) seeks to solve a minimax problem,

$$\lambda^* = \inf_{\lambda} \sup_{\theta} t(\mathbb{E}_\lambda[(O^{p,q} f_\theta)(z)]). \quad (6)$$

Algorithm 1: Operator variational inference

Input : Model $\log p(x, z)$, variational approximation $q(z; \lambda)$

Output: Variational parameters λ

Initialize λ and θ randomly.

while *not converged* **do**

 Compute unbiased estimates of $\nabla_{\lambda} \mathcal{L}_{\theta}$ from Equation 7.

 Compute unbiased estimates of $\nabla_{\theta} \mathcal{L}_{\lambda}$ from Equation 8.

 Update λ, θ with unbiased stochastic gradients.

end

We will use stochastic optimization [28, 15]. In principle, we can find stochastic gradients of λ by rewriting the objective in terms of the optimized value of θ , $\theta^*(\lambda)$. In practice, however, we simultaneously solve the maximization and minimization. Though computationally beneficial, this produces saddle points. In our experiments we found it to be stable enough. We derive gradients for the variational parameters λ and test function parameters θ . (We fix the distance function to be the square $t(a) = a^2$; the identity $t(a) = a$ also readily applies.)

Gradient with respect to λ . For a fixed test function with parameters θ , denote the objective

$$\mathcal{L}_{\theta} = t(\mathbb{E}_{\lambda}[(O^{p,q} f_{\theta})(z)]).$$

The gradient with respect to variational parameters λ is

$$\nabla_{\lambda} \mathcal{L}_{\theta} = 2 \mathbb{E}_{\lambda}[(O^{p,q} f_{\theta})(z)] \nabla_{\lambda} \mathbb{E}_{\lambda}[(O^{p,q} f_{\theta})(z)].$$

Now write the second expectation with the score function gradient [24]. This gradient is

$$\nabla_{\lambda} \mathcal{L}_{\theta} = 2 \mathbb{E}_{\lambda}[(O^{p,q} f_{\theta})(z)] \mathbb{E}_{\lambda}[\nabla_{\lambda} \log q(z; \lambda)(O^{p,q} f_{\theta})(z) + \nabla_{\lambda} (O^{p,q} f_{\theta})(z)]. \quad (7)$$

Equation 7 lets us calculate unbiased stochastic gradients. We first generate two sets of independent samples from q ; we then form Monte Carlo estimates of the first and second expectations. For the second expectation, we can use the variance reduction techniques developed for black box variational inference, such as Rao-Blackwellization [24].

We described the score gradient because it is general. An alternative is to use the reparameterization gradient for the second expectation [13, 27]. It requires that the operator be differentiable with respect to z and that samples from q can be drawn as a transformation r of a parameter-free noise source ϵ , $z = r(\epsilon, \lambda)$. In our experiments, we use the reparameterization gradient.

Gradient with respect to θ . Mirroring the notation above, the operator objective for fixed variational λ is

$$\mathcal{L}_{\lambda} = t(\mathbb{E}_{\lambda}[(O^{p,q} f_{\theta})(z)]).$$

The gradient with respect to test function parameters θ is

$$\nabla_{\theta} \mathcal{L}_{\lambda} = 2 \mathbb{E}_{\lambda}[(O^{p,q} f_{\theta})(z)] \mathbb{E}_{\lambda}[\nabla_{\theta} O^{p,q} f_{\theta}(z)]. \quad (8)$$

Again, we can construct unbiased stochastic gradients with two sets of Monte Carlo estimates. Note that gradients for the test function do not require score gradients (or reparameterization gradients) because the expectation does not depend on θ .

Algorithm. Algorithm 1 outlines OPVI. We simultaneously minimize the variational objective with respect to the variational family q_{λ} while maximizing it with respect to the function class f_{θ} . Given a model, operator, and function class parameterization, we can use automatic differentiation to calculate the necessary gradients [3]. Provided the operator does not require model-specific computation, this algorithm satisfies the black box criteria.

3.1 Data Subsampling and OPVI

With stochastic optimization, data subsampling scales up traditional variational inference to massive data [9, 31]. The idea is to calculate noisy gradients by repeatedly subsampling from the data set, without needing to pass through the entire data set for each gradient.

An as illustration, consider hierarchical models. Hierarchical models consist of global latent variables β that are shared across data points and local latent variables z_i each of which is associated to a data point x_i . The model’s log joint density is

$$\log p(x_{1:n}, z_{1:n}, \beta) = \log p(\beta) + \sum_{i=1}^n \left[\log p(x_i | z_i, \beta) + \log p(z_i | \beta) \right].$$

Hoffman et al. [9] calculate unbiased estimates of the log joint density (and its gradient) by subsampling data and appropriately scaling the sum.

We can characterize whether OPVI with a particular operator supports data subsampling. OPVI relies on evaluating the operator and its gradient at different realizations of the latent variables (Equation 7 and Equation 8). Thus we can subsample data to calculate estimates of the operator when it derives from linear operators of the log density, such as differentiation and the identity. This follows as a linear operator of sums is a sum of linear operators, so the gradients in Equation 7 and Equation 8 decompose into a sum. The Langevin-Stein and KL operator are both linear in the log density; both support data subsampling.

3.2 Variational Programs

Given an operator and variational family, Algorithm 1 optimizes the corresponding operator objective. Certain operators require the density of q . For example, the KL operator (Equation 5) requires its log density. This potentially limits the construction of rich variational approximations for which the density of q is difficult to compute.¹

Some operators, however, do not depend on having a analytic density; the Langevin-Stein (LS) operator (Equation 3) is an example. These operators can be used with a much richer class of variational approximations, those that can be sampled from but might not have analytically tractable densities. We call such approximating families *variational programs*.

Inference with a variational program requires the family to be reparameterizable [13, 27]. (Otherwise we need to use the score function, which requires the derivative of the density.) A reparameterizable variational program consists of a parametric deterministic transformation R of random noise ϵ . Formally, let

$$\epsilon \sim \text{Normal}(0, 1), \quad z = R(\epsilon; \lambda). \tag{9}$$

This generates samples for z , is differentiable with respect to λ , and its density may be intractable. For operators that do not require the density of q , it can be used as a powerful variational approximation. This is in contrast to the standard Kullback-Leibler (KL) operator.

As an example, consider the following variational program for a one-dimensional random variable. Let λ_i denote the i th dimension of λ and make the corresponding definition for ϵ :

$$z = (\epsilon_3 > 0)R(\epsilon_1; \lambda_1) - (\epsilon_3 \leq 0)R(\epsilon_2; \lambda_2). \tag{10}$$

When R outputs positive values, this separates the parametrization of the density to the positive and negative halves of the reals; its density is generally intractable. In Section 4, we will use this distribution as a variational approximation.

Equation 9 contains many densities when the function class R can approximate arbitrary continuous functions. We state it formally.

Theorem 1. *Consider a posterior distribution $p(z | x)$ with a finite number of latent variables and continuous quantile function. Assume the operator variational objective has a unique root at the posterior $p(z | x)$ and that R can approximate continuous functions. Then there exists a sequence of parameters $\lambda_1, \lambda_2, \dots$, in the variational program, such that the operator variational objective converges to 0, and thus q converges in distribution to $p(z | x)$.*

This theorem says that we can use variational programs with an appropriate q -independent operator to approximate continuous distributions. The proof is in Appendix D.

¹It is possible to construct rich approximating families with $\text{KL}(q||p)$, but this requires the introduction of an auxiliary distribution [18].

4 Empirical Study

We evaluate operator variational inference on a mixture of Gaussians, comparing different choices in the objective. We then study logistic factor analysis for images.

4.1 Mixture of Gaussians

Consider a one-dimensional mixture of Gaussians as the posterior of interest, $p(z) = \frac{1}{2}\text{Normal}(z; -3, 1) + \frac{1}{2}\text{Normal}(z; 3, 1)$. The posterior contains multiple modes. We seek to approximate it with three variational objectives: Kullback-Leibler (KL) with a Gaussian approximating family, Langevin-Stein (LS) with a Gaussian approximating family, and LS with a variational program.




Figure 1: The true posterior is a mixture of two Gaussians, in green. We approximate it with a Gaussian using two operators (in blue). The density on the far right is a variational program given in Equation 10 and using the Langevin-Stein operator; it approximates the truth well. The density of the variational program is intractable. We plot a histogram of its samples and compare this to the histogram of the true posterior.

Figure 1 displays the posterior approximations. We find that the KL divergence and LS divergence choose a single mode and have slightly different variances. These operators do not produce good results because a single Gaussian is a poor approximation to the mixture. The remaining distribution in Figure 1 comes from the toy variational program described by Equation 10 with the LS operator. Because this program captures different distributions for the positive and negative half of the real line, it is able to capture the posterior.

In general, the choice of an objective balances statistical and computational properties of variational inference. We highlight one tradeoff: the LS objective admits the use of a variational program; however, the objective is more difficult to optimize than the KL.

4.2 Logistic Factor Analysis

Logistic factor analysis models binary vectors \mathbf{x}_i with a matrix of parameters \mathbf{W} and biases \mathbf{b} ,

$$\begin{aligned} \mathbf{z}_i &\sim \text{Normal}(0, 1) \\ x_{i,k} &\sim \text{Bernoulli}(\sigma(\mathbf{w}_k^\top \mathbf{z}_i + b_k)), \end{aligned}$$

where \mathbf{z}_i has fixed dimension K and σ is the sigmoid function. This model captures correlations of the entries in \mathbf{x}_i through \mathbf{W} .

We apply logistic factor analysis to analyze the binarized MNIST data set [29], which contains 28x28 binary pixel images of handwritten digits. (We set the latent dimensionality to 10.) We fix the model parameters to those learned with variational expectation-maximization using the KL divergence, and focus on comparing posterior inferences.

We compare the KL operator to the LS operator and study two choices of variational models: a fully factorized Gaussian distribution and a variational program. The variational program generates samples by transforming a K -dimensional standard normal input with a two-layer neural network, using rectified linear activation functions and a hidden size of twice the latent dimensionality. Formally,

| Inference method | Completed data log-likelihood |
|--------------------------|-------------------------------|
| Mean-field Gaussian + KL | -59.3 |
| Mean-field Gaussian + LS | -75.3 |
| Variational Program + LS | -58.9 |

Table 1: Benchmarks on logistic factor analysis for binarized MNIST. The same variational approximation with LS performs worse than KL on likelihood performance. The variational program with LS performs better without directly optimizing for likelihoods.

the variational program we use generates samples of \mathbf{z} as follows:

$$\begin{aligned}
\mathbf{z}_0 &\sim \text{Normal}(0, I) \\
\mathbf{h}_0 &= \text{ReLU}(\mathbf{W}_0^q \mathbf{z}_0 + \mathbf{b}_0^q) \\
\mathbf{h}_1 &= \text{ReLU}(\mathbf{W}_1^q \mathbf{h}_0 + \mathbf{b}_1^q) \\
\mathbf{z} &= \mathbf{W}_2^q \mathbf{h}_1 + \mathbf{b}_2^q.
\end{aligned}$$

The variational parameters are the weights \mathbf{W}^q and biases \mathbf{b}^q . For f , we use a three-layer neural network with the same hidden size as the variational program and hyperbolic tangent activations where unit activations were bounded to have norm two. Bounding the unit norm bounds the divergence. We used the Adam optimizer [14] with learning rates 2×10^{-4} for f and 2×10^{-5} for the variational approximation.

There is no standard for evaluating generative models and their inference algorithms [30]. Following Rezende et al. [27], we consider a missing data problem. We remove half of the pixels in the test set (at random) and reconstruct them from a fitted posterior predictive distribution. Table 1 summarizes the results on 100 test images; we report the log-likelihood of the completed image. LS with the variational program performs best. It is followed by KL and the simpler LS inference. The LS performs better than KL even though the model parameters were learned with KL.

5 Summary

We present operator variational objectives, a broad yet tractable class of optimization problems for approximating posterior distributions. Operator objectives are built from an operator, a family of test functions, and a distance function. We outline the connection between operator objectives and existing divergences such as the KL divergence, and develop a new variational objective using the Langevin-Stein operator. In general, operator objectives produce new ways of posing variational inference.

Given an operator objective, we develop a black box algorithm for optimizing it and show which operators allow scalable optimization through data subsampling. Further, unlike the popular evidence lower bound, not all operators explicitly depend on the approximating density. This permits flexible approximating families, called variational programs, where the distributional form is not tractable. We demonstrate this approach on a mixture model and a factor model of images.

There are several possible avenues for future directions such as developing new variational objectives, adversarially learning [6] model parameters with operators, and learning model parameters with operator variational objectives.

Acknowledgments. This work is supported by NSF IIS-1247664, ONR N00014-11-1-0651, DARPA FA8750-14-2-0009, DARPA N66001-15-C-4032, Adobe, NSERC PGS-D, Porter Ogdan Jacobus Fellowship, Seibel Foundation, and the Sloan Foundation. The authors would like to thank Dawen Liang, Ben Poole, Stephan Mandt, Kevin Murphy, Christian Naesseth, and the anonymous reviews for their helpful feedback and comments.

References

- [1] Assaraf, R. and Caffarel, M. (1999). Zero-variance principle for monte carlo algorithms. In *Phys. Rev. Let.*
- [2] Barbour, A. D. (1988). Stein’s method and poisson process convergence. *Journal of Applied Probability*.

- [3] Carpenter, B., Hoffman, M. D., Brubaker, M., Lee, D., Li, P., and Betancourt, M. (2015). The Stan Math Library: Reverse-mode automatic differentiation in C++. *arXiv preprint arXiv:1509.07164*.
- [4] Cinlar, E. (2011). *Probability and Stochastics*. Springer.
- [5] Ghahramani, Z. and Beal, M. (2001). Propagation algorithms for variational Bayesian learning. In *NIPS 13*, pages 507–513.
- [6] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Neural Information Processing Systems*.
- [7] Gorham, J. and Mackey, L. (2015). Measuring sample quality with stein’s method. In *Advances in Neural Information Processing Systems*, pages 226–234.
- [8] Hernández-Lobato, J. M., Li, Y., Rowland, M., Hernández-Lobato, D., Bui, T., and Turner, R. E. (2015). Black-box α -divergence Minimization. *arXiv.org*.
- [9] Hoffman, M., Blei, D., Wang, C., and Paisley, J. (2013). Stochastic variational inference. *Journal of Machine Learning Research*, 14(1303–1347).
- [10] Hornik, K., Stinchcombe, M., and White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366.
- [11] Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(Apr):695–709.
- [12] Jordan, M., Ghahramani, Z., Jaakkola, T., and Saul, L. (1999). Introduction to variational methods for graphical models. *Machine Learning*, 37:183–233.
- [13] Kingma, D. and Welling, M. (2014). Auto-encoding variational bayes. In *ICLR*.
- [14] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- [15] Kushner, H. and Yin, G. (1997). *Stochastic Approximation Algorithms and Applications*. Springer New York.
- [16] Ley, C. and Swan, Y. (2011). Discrete stein characterizations and discrete information distances. *arXiv preprint arXiv:1201.0143*.
- [17] Li, Y. and Turner, R. E. (2016). Rényi divergence variational inference. *arXiv preprint arXiv:1602.02311*.
- [18] Maaløe, L., Sønderby, C. K., Sønderby, S. K., and Winther, O. (2016). Auxiliary deep generative models. *arXiv preprint arXiv:1602.05473*.
- [19] Minka, T. P. (2001). Expectation propagation for approximate Bayesian inference. In *UAI*.
- [20] Minka, T. P. (2004). Power EP. Technical report, Microsoft Research, Cambridge.
- [21] Mira, A., Solgi, R., and Imparato, D. (2013). Zero variance Markov chain Monte Carlo for Bayesian estimators. *Statistics and Computing*, pages 1–10.
- [22] Nielsen, F. and Nock, R. (2013). On the chi square and higher-order chi distances for approximating f-divergences. *arXiv preprint arXiv:1309.3029*.
- [23] Oates, C. J., Girolami, M., and Chopin, N. (2014). Control functionals for monte carlo integration. *arXiv preprint arXiv:1410.2392*.
- [24] Ranganath, R., Gerrish, S., and Blei, D. (2014). Black Box Variational Inference. In *AISTATS*.
- [25] Ranganath, R., Tran, D., and Blei, D. M. (2016). Hierarchical variational models. In *International Conference on Machine Learning*.
- [26] Rezende, D. J. and Mohamed, S. (2015). Variational inference with normalizing flows. In *Proceedings of the 31st International Conference on Machine Learning (ICML-15)*.
- [27] Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*.
- [28] Robbins, H. and Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):pp. 400–407.
- [29] Salakhutdinov, R. and Murray, I. (2008). On the quantitative analysis of deep belief networks. In *International Conference on Machine Learning*.
- [30] Theis, L., van den Oord, A., and Bethge, M. (2016). A note on the evaluation of generative models. In *International Conference on Learning Representations*.

- [31] Titsias, M. and Lázaro-Gredilla, M. (2014). Doubly stochastic variational bayes for non-conjugate inference. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 1971–1979.
- [32] Tran, D., Ranganath, R., and Blei, D. M. (2016). The variational Gaussian process. In *ICLR*.
- [33] Wingate, D. and Weber, T. (2013). Automated variational inference in probabilistic programming. *ArXiv e-prints*.

A Technical Conditions for Langevin-Stein Operators

Here we establish the conditions needed on the function class \mathcal{F} or the posterior distribution short-handed p for the operators to have expectation zero for all $f \in \mathcal{F}$. We derive properties using integration by parts for supports that are bounded open sets. Then we extend the result to unbounded supports using limits. We start with the Langevin-Stein operator. Let S be the set over which we integrate and let B be its boundary. Let v be the unit normal to the surface B , and v_i be the i th component of the surface normal (which is d dimensional). Then we have that

$$\begin{aligned}
\int_S p(O_{\text{LS}}^p f) dS &= \int_S p \nabla_z \log p^\top f + p \nabla^\top f dS \\
&= \sum_{i=1}^d \int_S \frac{\partial}{\partial z_i} [p] f_i + p \frac{\partial}{\partial z_i} [f_i] dS \\
&= \sum_{i=1}^d \int_S \frac{\partial}{\partial z_i} [p] f_i dS + \int_B f_i p v_i dB - \int_S \frac{\partial}{\partial z_i} [p] f_i dS \\
&= \int_B v^\top f p dB.
\end{aligned}$$

A sufficient condition for this expectation to be zero is that either p goes to zero at its boundary or that the vector field f is zero at the boundary.

For unbounded sets, the result can be written as a limit for a sequence of increasing sets $S_n \rightarrow S$ and a set of boundaries $B_n \rightarrow B$ using the dominated convergence theorem [4]. To use dominated convergence, we establish absolute integrability. Sufficient conditions for absolute integrability of the Langevin-Stein operator are for the gradient of $\log p$ to be bounded and the vector field f and its derivatives to be bounded. Via dominated convergence, we get that $\lim_n \int_{B_n} v^\top f p dB = 0$ for the Langevin-Stein operator to have expectation zero.

B Characterizing the zeros of the Langevin-Stein Operators

We provide analysis on how to characterize the equivalence class of distributions defined as $(O^{p,q} f)(z) = 0$. One general condition for equality in distribution comes from equality in probability on all Borel sets. We can build functions that have expectation zero with respect to the posterior that test this equality. Formally, for any Borel set A with δ_A being the indicator, these functions on A have the form:

$$\delta_A(z) - \int_A p(y) dy$$

We show that if the Langevin-Stein operator satisfies $\mathcal{L}(q; O_{\text{LS}}^p, \mathcal{F}) = 0$, then q is equivalent to p in distribution. We do this by showing the above functions are in the span of O_{LS}^p . Expanding the Langevin-Stein operator we have

$$(O_{\text{LS}}^p f) = p^{-1} \nabla_z p^\top f + \nabla^\top f = p^{-1} \sum_{i=1}^d \frac{\partial f_i p}{\partial z_i}.$$

Setting this equal to the desired function above yields the differential equation

$$\delta_A(z) - \int_A p(y) dy = p^{-1}(z) \sum_{i=1}^d \frac{\partial f_i p}{\partial z_i}(z).$$

To solve this, set $f_i = 0$ for all but $i = 1$. This yields

$$\delta_A(z) - \int_A p(y)dy = p^{-1}(z) \frac{\partial f_1 p}{\partial z_1}(z),$$

which is an ordinary differential equation with solution for f_1

$$f_1^A(z) = \frac{1}{p(z)} \int_{-\infty}^{z_1} p(a, z_2 \dots d) \left(\delta_A(a, z_2 \dots d) - \int_A p(y)dy \right) da.$$

This function is differentiable with respect to z_1 , so this gives the desired result. Plugging the function back into the operator variational objective gives

$$\mathbb{E}_q \left[\delta_A(z) - \int_A p(y)dy \right] = 0 \iff \mathbb{E}_q[\delta_A(z)] = \mathbb{E}_p[\delta_A(z)],$$

for all Borel measurable A . This implies the induced distance captures total variation.

C Operators for Discrete Variables

Some operators based on Stein’s method are applicable only for latent variables in a continuous space. There are Stein operators that work with discrete variables [1, 16]. We present one amenable to operator variational objectives based on a discrete analogue to the Langevin-Stein operator developed in [16]. For simplicity, consider a one-dimensional discrete posterior with support $\{0, \dots, c\}$. Let f be a function such that $f(0) = 0$, then an operator can be defined as

$$(O_{\text{DISCRETE}}^p f)(z) = \frac{f(z+1)p(z+1, \mathbf{x}) - f(z)p(z, \mathbf{x})}{p(z, \mathbf{x})}.$$

Since the expectation of this operator with respect to the posterior $p(z | x)$ is a telescoping sum with both endpoints 0, it has expectation zero.

This relates to the Langevin-Stein operator in the following. The Langevin-Stein operator in one dimension can be written as

$$(O_{\text{LS}}^p f) = \frac{d}{dz} [fp].$$

This operator is the discrete analogue as the differential is replaced by a discrete difference. We can extend this operator to multiple dimensions by an ordered indexing. For example, binary numbers of length n would work for n binary latent variables.

D Proof of Universal Representations

Consider the optimal form of R such that transformations of standard normal draws are equal in distribution to exact draws from the posterior. This means

$$R(\epsilon; \lambda) = P^{-1}(\Phi(\epsilon)),$$

where $\Phi(\epsilon)$ squashes the draw from a standard normal such that it is equal in distribution to a uniform random variable. The posterior’s inverse cumulative distribution function P^{-1} is applied to the uniform draws. The transformed samples are now equivalent to exact samples from the posterior. For a rich-enough parameterization of R , we may hope to sufficiently approximate this function.

Indeed, as in the universal approximation theorem of Tran et al. [32] there exists a sequence of parameters $\{\lambda_1, \lambda_2, \dots\}$ such that the operator variational objective goes to zero, but the function class is no longer limited to local interpolation. Universal approximators like neural networks [10] also work. Further, under the assumption that p is the unique root and by satisfying the conditions described in Section B for equality in distribution, this implies that the variational program given by drawing $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and applying $R(\epsilon)$ converges in distribution to $p(z | x)$.