
Pathwise Derivatives for Multivariate Distributions

Martin Jankowiak*

Uber AI Labs, San Francisco, CA

Theofanis Karaletsos

Abstract

We exploit the link between the transport equation and derivatives of expectations to construct efficient pathwise gradient estimators for multivariate distributions. We focus on two main threads. First, we use null solutions of the transport equation to construct adaptive control variates that can be used to construct gradient estimators with reduced variance. Second, we consider the case of multivariate mixture distributions. In particular we show how to compute pathwise derivatives for mixtures of multivariate Normal distributions with arbitrary means and diagonal covariances. We demonstrate in a variety of experiments in the context of variational inference that our gradient estimators can outperform other methods, especially in high dimensions.

1 Introduction

Stochastic optimization is a major component of many machine learning algorithms. In some cases—for example in variational inference—the stochasticity arises from an explicit, parameterized distribution $q_{\theta}(\mathbf{z})$. In these cases the optimization problem can often be cast as maximizing an objective

$$\mathcal{L} = \mathbb{E}_{q_{\theta}(\mathbf{z})} [f(\mathbf{z})] \quad (1)$$

where $f(\mathbf{z})$ is a (differentiable) test function and θ is a vector of parameters². In order to maximize \mathcal{L} we would like to follow gradients $\nabla_{\theta}\mathcal{L}$. If \mathbf{z} were a

*Correspondence to: jankowiak@uber.com

²Without loss of generality we assume that $f(\mathbf{z})$ has no explicit dependence on θ , since gradients of $f(\mathbf{z})$ are readily handled by bringing ∇_{θ} inside the expectation.

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

deterministic function of θ we would simply apply the chain rule; for any component θ_{α} of θ we have

$$\nabla_{\theta_{\alpha}} f(\mathbf{z}) = \nabla_{\mathbf{z}} f \cdot \frac{\partial \mathbf{z}}{\partial \theta_{\alpha}} \quad (2)$$

What is the generalization of Eqn. 2 to the stochastic case? As observed in [18], we can construct an unbiased estimator for the gradient of Eqn. 1 provided we can solve the following partial differential equation (the transport a.k.a. continuity equation):

$$\frac{\partial}{\partial \theta_{\alpha}} q_{\theta} + \nabla_{\mathbf{z}} \cdot (q_{\theta} \mathbf{v}^{\theta_{\alpha}}) = 0 \quad (3)$$

Here $\mathbf{v}^{\theta_{\alpha}}$ is a vector field defined on the sample space of $q_{\theta}(\mathbf{z})$ ³. Note that there is a vector field $\mathbf{v}^{\theta_{\alpha}}$ for each component θ_{α} of θ . We can then form the following *pathwise* gradient estimator:

$$\nabla_{\theta_{\alpha}} \mathcal{L} = \mathbb{E}_{q_{\theta}(\mathbf{z})} [\nabla_{\mathbf{z}} f \cdot \mathbf{v}^{\theta_{\alpha}}] \quad (4)$$

Pathwise gradient estimators are particularly appealing because, empirically, they generally exhibit lower variance than alternatives.

That the gradient estimator in Eqn. 4 is unbiased follows directly from the divergence theorem:

$$\begin{aligned} \nabla_{\theta_{\alpha}} \mathcal{L} &= \int d\mathbf{z} f(\mathbf{z}) \frac{\partial q_{\theta}(\mathbf{z})}{\partial \theta_{\alpha}} = - \int d\mathbf{z} f(\mathbf{z}) \nabla_{\mathbf{z}} \cdot (q_{\theta} \mathbf{v}^{\theta_{\alpha}}) \\ &= \int_V d\mathbf{z} \nabla_{\mathbf{z}} f \cdot (q_{\theta} \mathbf{v}^{\theta_{\alpha}}) - \oint_S f q_{\theta} \mathbf{v}^{\theta} \cdot \hat{\mathbf{n}} dS \\ &= \lim_{S \rightarrow \infty} \left\{ \int_V d\mathbf{z} \nabla_{\mathbf{z}} f \cdot (q_{\theta} \mathbf{v}^{\theta_{\alpha}}) - \oint_S f q_{\theta} \mathbf{v}^{\theta} \cdot \hat{\mathbf{n}} dS \right\} \\ &= \int d\mathbf{z} q_{\theta}(\mathbf{z}) \nabla_{\mathbf{z}} f \cdot \mathbf{v}^{\theta_{\alpha}} = \mathbb{E}_{q_{\theta}(\mathbf{z})} [\nabla_{\mathbf{z}} f \cdot \mathbf{v}^{\theta_{\alpha}}] \end{aligned}$$

Here we have substituted for $\frac{\partial}{\partial \theta_{\alpha}} q_{\theta}$ using Eqn. 3 in the first line, appealed to the divergence theorem⁴ in the second line, and taken the surface to infinity in the

³More explicitly: at each point \mathbf{z} in the sample space the vector field $\mathbf{v}^{\theta_{\alpha}}$ is specified by D components $\mathbf{v}_i^{\theta_{\alpha}}$, $i = 1, \dots, D$, where D is the dimension of the sample space.

⁴More explicitly: we apply the divergence theorem to the vector field $f q_{\theta} \mathbf{v}^{\theta}$ and expand the divergence in the volume integral using the product rule (cf. Green's identities).

third line. In the final step we assume that $f_{q\theta}v^\theta$ is sufficiently well-behaved that we can drop the surface integral as the surface goes to infinity.

Thus Eqn. 4 is a generalization of Eqn. 2 to the stochastic case, with the velocity field v^{θ_α} playing the role of $\frac{\partial z}{\partial \theta_\alpha}$ in the deterministic case. In contrast to the deterministic case where $\frac{\partial z}{\partial \theta_\alpha}$ is uniquely specified, Eqn. 3 generally admits an infinite dimensional space of solutions for v^{θ_α} .

The transport equation encodes an intuitive geometric picture. As we vary θ we move $q_\theta(z)$ along a curve in the space of distributions over the sample space. This curve corresponds to a time-varying cloud of particles; in this analogy the set of velocity fields $\{v^{\theta_\alpha}\}_\alpha$ describes the (infinitesimal) displacements that particles undergo as θ is varied. Because each velocity field obeys the corresponding transport equation, Eqn. 3, the total probability is conserved and the displaced particles are distributed according to the displaced distribution.

In this work we are interested in the multivariate case, where the solution space to Eqn. 3 is very rich, and where alternative gradient estimators tend to suffer from large variance, especially in high dimensions. Our contributions are two-fold. In Sec. 3 we exploit the fact that the transport equation admits a large space of solutions to construct gradient estimators that are *adapted* to the test function $f(z)$. In the context of the reparameterization trick (see Sec. 2.3) our approach is conceptually analogous to adaptively choosing a reparameterization \mathcal{T} from a family of transformations $\{\mathcal{T}_\lambda(\epsilon; \theta)\}$, with the difference that the transformation is never explicitly constructed; rather it is implicit as a solution to Eqn. 3. Second, in Sec. 4 we construct pathwise gradient estimators for mixtures of Normal distributions with diagonal covariance matrices. We emphasize that the resulting gradient estimators are the only known unbiased pathwise gradient estimators for this class of mixture distributions whose computational complexity is linear in the dimension of the sample space. As we show in Sec. 6 our gradient estimators often exhibit low variance and good performance as compared to alternative methods.

The rest of this paper is organized as follows. In Sec. 2 we give an overview of stochastic gradient variational inference (SGVI) and stochastic gradient estimators. In Sec. 3 we use parameterized families of solutions to Eqn. 3 to construct adaptive pathwise gradient estimators. In Sec. 4 we present solutions to Eqn. 3 for multivariate mixture distributions. In Sec. 5 we place our work in the context of recent research. In Sec. 6 we validate our proposed techniques with a variety of experiments in the context of variational inference. In Sec. 7 we discuss directions for future work.

2 Background

2.1 Stochastic Gradient Variational Inference

Given a probabilistic model $p(x, z)$ with observations x and latent random variables z , variational inference recasts posterior inference as an optimization problem. Specifically we define a family of variational distributions $q_\theta(z)$ parameterized by θ and seek to find a value of θ that minimizes the KL divergence between $q_\theta(z)$ and the (unknown) posterior $p(z|x)$ [19, 29, 15, 45, 31]. This is equivalent to maximizing the ELBO, defined as

$$\text{ELBO} = \mathbb{E}_{q_\theta(z)} [\log p(x, z) - \log q_\theta(z)] \quad (5)$$

This stochastic optimization problem will be the basic setting for most of our experiments in Sec. 6.

2.2 Score Function Estimator

The score function estimator, also referred to as REINFORCE [12, 44, 10], provides a simple and broadly applicable recipe for estimating gradients, with the simplest variant given by

$$\nabla_{\theta_\alpha} \mathcal{L} = \mathbb{E}_{q_\theta(z)} [f(z) \nabla_{\theta_\alpha} \log q_\theta(z)]$$

Although the score function estimator is very general (e.g. it can be applied to discrete random variables) it typically suffers from high variance, although this can be mitigated with the use of variance reduction techniques such as Rao-Blackwellization [5] and control variates [36].

2.3 Reparameterization Trick

The reparameterization trick (RT) is not as broadly applicable as the score function estimator, but it generally exhibits lower variance [30, 39, 22, 11, 42, 34]. It is applicable to continuous random variables whose probability density $q_\theta(z)$ can be reparameterized such that we can rewrite expectations $\mathbb{E}_{q_\theta(z)} [f(z)]$ as $\mathbb{E}_{q_0(\epsilon)} [f(\mathcal{T}(\epsilon; \theta))]$, where $q_0(\epsilon)$ is a fixed distribution and $\mathcal{T}(\epsilon; \theta)$ is a differentiable θ -dependent transformation. Since the expectation w.r.t. $q_0(\epsilon)$ has no θ dependence, gradients w.r.t. θ can be computed by pushing ∇_θ through the expectation. This reparameterization can be done for a number of distributions, including for example the Normal distribution. As noted in [18], in cases where the reparameterization trick is applicable, we have that

$$v^{\theta_\alpha} = \left. \frac{\partial \mathcal{T}(\epsilon; \theta)}{\partial \theta_\alpha} \right|_{\epsilon = \mathcal{T}^{-1}(z; \theta)} \quad (6)$$

is a solution to the transport equation, Eqn. 3.

3 Adaptive Velocity Fields

Whenever we have a solution to Eqn. 3 we can form an unbiased gradient estimator via Eqn. 4. An intriguing possibility arises if we can construct a parameterized *family* of solutions $\mathbf{v}_\lambda^{\theta_\alpha}$: as we take steps in θ -space we can simultaneously take steps in λ -space, thus adaptively choosing the form the θ_α gradient estimator takes. This can be understood as an instance of an adaptive control variate [36]; consider the solution $\mathbf{v}_\lambda^{\theta_\alpha}$ as well as a fixed reference solution $\mathbf{v}_0^{\theta_\alpha}$:

$$\begin{aligned} \nabla_{\theta_\alpha} \mathcal{L} &= \mathbb{E}_{q_\theta(\mathbf{z})} \left[\nabla_{\mathbf{z}} f \cdot \mathbf{v}_\lambda^{\theta_\alpha} \right] \\ &= \mathbb{E}_{q_\theta(\mathbf{z})} \left[\nabla_{\mathbf{z}} f \cdot \mathbf{v}_0^{\theta_\alpha} \right] + \mathbb{E}_{q_\theta(\mathbf{z})} \left[\nabla_{\mathbf{z}} f \cdot \left(\mathbf{v}_\lambda^{\theta_\alpha} - \mathbf{v}_0^{\theta_\alpha} \right) \right] \end{aligned} \quad (7)$$

The final expectation is identically equal to zero so the integrand is a control variate. If we choose λ appropriately we can reduce the variance of our gradient estimator. In order to specify a complete algorithm we need to answer two questions: i) how should we adapt λ ?; and ii) what is an appropriate family of solutions $\mathbf{v}_\lambda^{\theta_\alpha}$? We now address each of these in turn.

3.1 Adapting velocity fields

Whenever we compute a gradient estimate via Eqn. 4 we can use the same N_s sample(s) $\mathbf{z}_i \sim q_\theta$ to form an estimate of the variance of the gradient estimator:

$$\begin{aligned} \text{Var} \left(\frac{\partial \widehat{\mathcal{L}}}{\partial \theta} \right) &\equiv \sum_{\alpha} \text{Var} \left(\frac{\partial \widehat{\mathcal{L}}}{\partial \theta_\alpha} \right) = \\ &\frac{1}{N_s} \sum_{\alpha} \sum_{\mathbf{z}_i} \left[(\nabla_{\mathbf{z}} f \cdot \mathbf{v}_\lambda^{\theta_\alpha})^2 \right] - \sum_{\alpha} \left(\frac{1}{N_s} \sum_{\mathbf{z}_i} \left[\nabla_{\mathbf{z}} f \cdot \mathbf{v}_\lambda^{\theta_\alpha} \right] \right)^2 \end{aligned}$$

In direct analogy to the adaptation procedure used in (for example) [37] we can adapt λ by following noisy gradient estimates of this variance:

$$\nabla_{\lambda} \text{Var} \left(\frac{\partial \widehat{\mathcal{L}}}{\partial \theta} \right) = \frac{1}{N_s} \sum_{\alpha} \sum_{\mathbf{z}_i} \left[\nabla_{\lambda} (\nabla_{\mathbf{z}} f \cdot \mathbf{v}_\lambda^{\theta_\alpha})^2 \right] \quad (8)$$

In this way λ can be adapted to reduce the gradient variance during the course of optimization. See Algorithm 1 for a summary of the complete algorithm in the case where single-sample gradient estimates are used at each iteration.

3.2 Parameterizing velocity fields

Assuming we have obtained (at least) one solution \mathbf{v}_0^{θ} to the transport equation Eqn. 3 parameterizing a family of solutions $\mathbf{v}_\lambda^{\theta}$ is in principle easy. We simply solve the null equation, $\nabla_{\mathbf{z}} \cdot (q_\theta \tilde{\mathbf{v}}^{\theta}) = 0$, which is obtained

Algorithm 1: Stochastic optimization with Adaptive Velocity Fields. Note that steps 2 and 4 occur implicitly for the reparameterization trick estimator. What sets Adaptive Velocity Fields apart is steps 3 and 5.

Initialize θ, λ , and choose step sizes $\epsilon_\theta, \epsilon_\lambda$.

for $i = 1, 2, \dots, N_{\text{steps}}$ **do**

1. Sample $\mathbf{z}_i \sim q_\theta$
2. Compute $g(\theta)_i^{\theta_\alpha} \equiv \nabla_{\mathbf{z}} f(\mathbf{z}_i) \cdot \mathbf{v}_\lambda^{\theta_\alpha}$ for all α (cf. Eqn. 7)
3. Compute $\mathbf{g}(\lambda)_i \equiv \nabla_{\lambda} \text{Var} \left(\frac{\partial \mathcal{L}}{\partial \theta} \right)$ (cf. Eqn. 8)
4. Take a (noisy) gradient step $\theta \rightarrow \theta + \epsilon_\theta \mathbf{g}(\theta)_i$
5. Take a (noisy) gradient step $\lambda \rightarrow \lambda + \epsilon_\lambda \mathbf{g}(\lambda)_i$

end

from the transport equation by dropping the source term $\frac{\partial}{\partial \theta} q_\theta$. Then for a given family of solutions to the null equation, $\tilde{\mathbf{v}}_\lambda^{\theta}$, we get a solution $\mathbf{v}_\lambda^{\theta}$ to the transport equation by simple addition, i.e. $\mathbf{v}_\lambda^{\theta} = \mathbf{v}_0^{\theta} + \tilde{\mathbf{v}}_\lambda^{\theta}$. At first glance solving the null equation appears to be easy, since *any* divergence-free vector field immediately yields a solution:

$$\nabla_{\mathbf{z}} \cdot \mathbf{w}_\lambda = 0 \quad \text{and} \quad \tilde{\mathbf{v}}_\lambda^{\theta} = \frac{\mathbf{w}_\lambda}{q_\theta} \quad \Rightarrow \quad \nabla_{\mathbf{z}} \cdot (q_\theta \tilde{\mathbf{v}}_\lambda^{\theta}) = 0$$

However, in order to construct general purpose gradient estimators we need to impose appropriate boundary conditions on the velocity field; for more on this subtlety we refer the reader to the supplementary materials.

In order to demonstrate how to construct suitable null solutions we henceforth focus on a particular family of distributions $q_\theta(\mathbf{z})$, namely *elliptical* distributions. Elliptical distributions are probability distributions whose density is of the form $q(\mathbf{z}) \propto g(\mathbf{z}^T \Sigma^{-1} \mathbf{z})$ for some scalar density $g(\cdot)$ and positive definite symmetric matrix Σ . They include, for example, the multivariate Normal distribution and the multivariate t-distribution. For concreteness, in the following we focus exclusively on the multivariate Normal distribution. We stress, however, that the resulting adaptive gradient estimators are applicable to any elliptical distribution. We refer to the supplementary materials for a description of the specific case of the multivariate t-distribution.

3.3 Null Solutions for the Multivariate Normal Distribution

Consider the multivariate Normal distribution in D dimensions parameterized via a mean $\boldsymbol{\mu}$ and Cholesky factor \mathbf{L} . We would like to compute derivatives w.r.t. L_{ab} . While the reparameterization trick is applicable in this case, we can potentially get lower variance gradients by using Adaptive Velocity Fields. As we show in the supplementary materials, a simple solution to the

corresponding null transport equation is given by

$$\tilde{\mathbf{v}}_{\mathbf{A}}^{L_{ab}} = \mathbf{L}\mathbf{A}^{ab}\mathbf{L}^{-1}(\mathbf{z} - \boldsymbol{\mu}) \quad (9)$$

where $\boldsymbol{\lambda} \equiv \{\mathbf{A}^{ab}\}$ is an arbitrary collection of antisymmetric $D \times D$ matrices (one for each a, b). This is just a linear vector field and so it is easy to manipulate and (relatively) cheap to compute. Geometrically, for each entry of \mathbf{L} , this null solution corresponds to an infinitesimal rotation in whitened coordinates $\tilde{\mathbf{z}} = \mathbf{L}^{-1}(\mathbf{z} - \boldsymbol{\mu})$. Note that $\tilde{\mathbf{v}}_{\mathbf{A}}^{L_{ab}}$ is *added* to the reference solution $\mathbf{v}_0^{L_{ab}}$, and this is *not* equivalent to rotating $\mathbf{v}_0^{L_{ab}}$.

Since a, b runs over $\frac{D(D+1)}{2}$ pairs of indices and each \mathbf{A}^{ab} has $\frac{D(D-1)}{2}$ free parameters, this space of solutions is quite large. Since we will be adapting \mathbf{A}^{ab} via noisy gradient estimates—and because we would like to limit the computational cost—we choose to parameterize a smaller subspace of solutions. Empirically we find that the following parameterization works well:

$$A_{jk}^{ab} = \sum_{\ell=1}^M B_{\ell a} C_{\ell b} (\delta_{a\ell} \delta_{b\ell} - \delta_{a\ell} \delta_{b\ell}) \quad (10)$$

Here each $M \times D$ matrix \mathbf{B} and \mathbf{C} is arbitrary and the hyperparameter M allows us to trade off the flexibility of our family of null solutions with computational cost (typically $M \ll D$).

The computational complexity of using Adaptive Velocity Field gradients with this class of parameterized velocity fields (including the \mathbf{A} update equations) is $\mathcal{O}(D^3 + MD^2)$ per gradient step. This should be compared to the $\mathcal{O}(D^2)$ cost of the reparameterization trick gradient and the $\mathcal{O}(D^3)$ cost of the OMT gradient (see Sec. 5 for a discussion of the latter).⁵ We explore the performance of the resulting Adaptive Velocity Field (AVF) gradient estimator in Sec. 6.1.

4 Pathwise Derivatives for Multivariate Mixture Distributions

In this section we use the transport equation Eqn. 3 to construct pathwise gradient estimators for mixture distributions. Note that this section is logically independent from the previous section, although—as we discuss briefly in the supplementary materials—Adaptive Velocity Fields can be readily combined with the gradient estimators discussed here.

⁵Note, however, that the computational complexity of the AVF gradient is somewhat misleading in that the $\mathcal{O}(D^3)$ term arises from matrix multiplications, which tend to be quite fast. By contrast the OMT gradient estimator involves a singular value decomposition, which tends to be substantially more expensive than a matrix multiplication on modern hardware.

Consider a mixture of K multivariate distributions in D dimensions:

$$q_{\boldsymbol{\theta}}(\mathbf{z}) = \sum_{j=1}^K \pi_j q_{\boldsymbol{\theta}_j}(\mathbf{z}) \quad (11)$$

In order to compute pathwise derivatives of $\mathbf{z} \sim q_{\boldsymbol{\theta}}(\mathbf{z})$ we need to solve the transport equation w.r.t. the mixture weights π_j as well as the component parameters $\boldsymbol{\theta}_j$. For the derivatives w.r.t. $\boldsymbol{\theta}_j$ we can just repurpose the velocity fields of the individual components. That is, if $\mathbf{v}_{\text{single}}^{\boldsymbol{\theta}_j}$ is a solution of the single-component transport equation (i.e. $\frac{\partial q_{\boldsymbol{\theta}_j}}{\partial \boldsymbol{\theta}_j} + \nabla \cdot (q_{\boldsymbol{\theta}_j} \mathbf{v}_{\text{single}}^{\boldsymbol{\theta}_j}) = 0$), then

$$\mathbf{v}^{\boldsymbol{\theta}_j} = \frac{\pi_j q_{\boldsymbol{\theta}_j}}{q_{\boldsymbol{\theta}}} \mathbf{v}_{\text{single}}^{\boldsymbol{\theta}_j} \quad (12)$$

is a solution of the multi-component transport equation.⁶ See the supplementary materials for the (short) derivation. For example we can readily compute Eqn. 12 if each component distribution is reparameterizable. Note that a typical gradient estimator for, say, a mixture of Normal distributions first samples the discrete component $k \in [1, \dots, K]$ and then samples \mathbf{z} conditioned on k . This results in a pathwise derivative for \mathbf{z} that only depends on $\boldsymbol{\mu}_k$ and $\boldsymbol{\sigma}_k$. By contrast the gradient computed via Eqn. 12 will result in a gradient for all K component parameters for every sample \mathbf{z} . We explore this difference experimentally in Sec. 6.2.2.

4.1 Mixture weight derivatives

Unfortunately, solving the transport equation for the mixture weights is in general much more difficult, since the desired velocity fields need to coordinate mass transport among all K components.⁷ We now describe a formal construction for solving the transport equation for the mixture weights. For each pair of component distributions j, k consider the transport equation

$$q_{\boldsymbol{\theta}_j} - q_{\boldsymbol{\theta}_k} + \nabla_{\mathbf{z}} \cdot (q_{\boldsymbol{\theta}} \tilde{\mathbf{v}}^{jk}) = 0 \quad (13)$$

Intuitively, the velocity field $\tilde{\mathbf{v}}^{jk}$ moves mass from component k to component j . We can superimpose these solutions to form

$$\mathbf{v}^{\ell_j} = \pi_j \sum_{k \neq j} \pi_k \tilde{\mathbf{v}}^{jk} \quad (14)$$

As we show in the supplementary materials, the velocity field \mathbf{v}^{ℓ_j} yields an estimator for the derivative

⁶We note that the form of Eqn. 12 implies that we can easily introduce Adaptive Velocity Fields for each individual component (although the cost of doing so may be prohibitively high for a large number of components).

⁷For example, we expect this to be particularly difficult if some of the component distributions belong to different families of distributions, e.g. a mixture of Normal distributions and Wishart distributions.

w.r.t. the softmax logit ℓ_j that corresponds to component j .⁸ Thus provided we can find solutions $\tilde{\mathbf{v}}^{jk}$ to Eqn. 13 for each pair of components j, k we can form a pathwise gradient estimator for the mixture weights. We now consider one particular situation where this recipe can be carried out. We leave a discussion of other solutions—including a solution for a mixture of Normal distributions with arbitrary diagonal covariances—to the supplementary materials.

4.2 Mixture of Multivariate Normals with Shared Diagonal Covariance

Here each component distribution is specified by $q_{\theta_j}(\mathbf{z}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_j, \boldsymbol{\sigma})$, i.e. each component distribution has its own mean vector $\boldsymbol{\mu}_j$ but all K components share the same (diagonal) square root covariance $\boldsymbol{\sigma}$.⁹ We find that a solution to Eqn. 13 is given by

$$\tilde{\mathbf{v}}^{jk} = \frac{\left(\Phi(\tilde{\mathbf{z}}_{\parallel}^{jk} - \tilde{\boldsymbol{\mu}}_{\parallel}^{jk}) - \Phi(\tilde{\mathbf{z}}_{\parallel}^{jk} + \tilde{\boldsymbol{\mu}}_{\parallel}^{kj})\right) \phi(\|\tilde{\mathbf{z}}_{\perp}^{jk}\|^2)}{(2\pi)^{D/2-1} q_{\theta} \prod_{i=1}^D \sigma_i} \boldsymbol{\sigma} \hat{\boldsymbol{\mu}}^{jk} \quad (15)$$

where $\boldsymbol{\sigma} \hat{\boldsymbol{\mu}}^{jk}$ indicates matrix multiplication by $\text{diag}(\boldsymbol{\sigma})$, $\tilde{\mathbf{z}} \equiv \mathbf{z} \odot \boldsymbol{\sigma}^{-1}$, and where the quantities $\{\hat{\boldsymbol{\mu}}^j, \hat{\boldsymbol{\mu}}^{jk}, \tilde{\boldsymbol{\mu}}_{\parallel}^{jk}, \tilde{\boldsymbol{\mu}}_{\parallel}^{kj}, \tilde{\mathbf{z}}_{\parallel}^{jk}, \tilde{\mathbf{z}}_{\perp}^{jk}\}$ are defined in the supplementary materials. Here $\Phi(\cdot)$ is the CDF of the unit Normal distribution and $\phi(\cdot)$ is the corresponding probability density. Geometrically, the velocity field in Eqn. 15 moves mass along lines parallel to $\boldsymbol{\mu}_j - \boldsymbol{\mu}_k$. Note that the gradient estimator resulting from Eqn. 15 has cost $\mathcal{O}(DK^2)$.¹⁰ We explore the performance of the resulting pathwise gradient estimators in Sec. 6.2.

5 Related Work

There is a large body of work on constructing gradient estimators with reduced variance, much of which can be understood in terms of control variates [36]: for example, [27] constructs neural baselines for score function gradients; [40] discuss gradient estimators for stochastic computation graphs and their Rao-Blackwellization; and [43, 13] construct adaptive control variates for discrete random variables. Another example of this line of work is [25], where the authors construct control variates that are applicable when $q_{\theta}(\mathbf{z})$ is a *diagonal* Normal distribution and that rely on Taylor expansions

of the test function $f(\mathbf{z})$.¹¹ In contrast our AVF gradient estimator for the multivariate Normal has about half the computational cost, since it does not rely on second order derivatives. Another line of work constructs partially reparameterized gradient estimators for cases where the reparameterization trick is difficult to apply [38, 28]. In [14], the author constructs pathwise gradient estimators for mixtures of diagonal Normal distributions. Unfortunately, the resulting estimator is expensive, relying on a recursive computation that scales with the dimension of the sample space. Another approach to mixture distributions is taken in [7], where the authors use quadrature-like constructions to form continuous distributions that are reparameterizable. Ref. [35] considers a gradient estimator for mixture distributions in which all K components are summed out. Our work is closest to [18], which uses the transport equation to derive an ‘OMT’ gradient estimator for the multivariate Normal distribution whose velocity field is optimal in the sense of optimal transport. This gradient estimator can exhibit low variance but has $\mathcal{O}(D^3)$ computational complexity due to an (expensive) singular value decomposition. In addition the OMT gradient estimator is *not* adapted to the test function at hand; rather it has one fixed form.

As this manuscript was being completed, we became aware of reference [9], which has some overlap with this work. In particular, [9] describes an interesting recipe for computing pathwise gradients that can be applied to mixture distributions. A nice feature of this recipe (also true of [14]), is that it can be applied to mixtures of Normal distributions with arbitrary covariances. A disadvantage is that it can be expensive, exhibiting a $\mathcal{O}(D^2)$ computational cost even for a mixture of Normal distributions with diagonal covariances. In contrast the gradient estimators constructed in Sec. 4 by solving the transport equation are linear in D . We provide a comparison to their approach in Sec. 6.2.1.

6 Experiments


We conduct a series of experiments using synthetic and real world data to validate the performance of the gradient estimators introduced in Sec. 3 & 4. Throughout we use single sample estimators. See Sec. E in the supplementary materials for details on experimental setups. Open source implementations of all our gradient estimators are available at <https://git.io/fhbqH>. All variational inference experiments were implemented using the Pyro probabilistic programming language [2].

⁸That is with respect to ℓ_j where $\pi_j = e^{\ell_j} / \sum_k e^{\ell_k}$.


⁹To keep the equations compact we consider a diagonal covariance matrix. The resulting velocity field is readily transformed to a rotated coordinate system in which the (single shared) covariance matrix is arbitrary.

¹⁰As detailed in the supplementary materials, the pathwise gradient estimators for the families of mixture distributions we consider all have cost $\mathcal{O}(DK)$ or $\mathcal{O}(DK^2)$.


¹¹Also, in their approach variance reduction for scale parameter gradients $\nabla_{\boldsymbol{\sigma}}$ necessitates a multi-sample estimator (at least for high-dimensional models where computing the diagonal of the Hessian is expensive).



(a) We compare the OMT and AVF gradient estimators for the multivariate Normal distribution to the RT estimator for three test functions. The horizontal axis controls the magnitude of the off-diagonal elements of the Cholesky factor \mathbf{L} . The vertical axis depicts the ratio of the mean variance of the given estimator to that of the RT estimator for the off-diagonal elements of \mathbf{L} .



(b) We compare the variance of pathwise gradient estimators for various mixture distributions with $K = 10$ to the corresponding score function estimator. The vertical axis depicts the ratio of the logit gradient variance of the pathwise estimator to that of the score function estimator for the test function $f(\mathbf{z}) = \|\mathbf{z}\|^2$.



(c) We compare the variance of our pathwise gradient estimator for mixtures of diagonal Normal distributions to the one described in ref. [9]. Here $D = 2$, σ controls the covariance of each component distribution, and the vertical axis is the ratio of variances (with our estimator in the denominator). See supplementary materials for details.

Figure 1: We compare the gradient variances for different gradient estimators and synthetic test functions. See Sec. 6.1.1 and Sec. 6.2.1 for details.

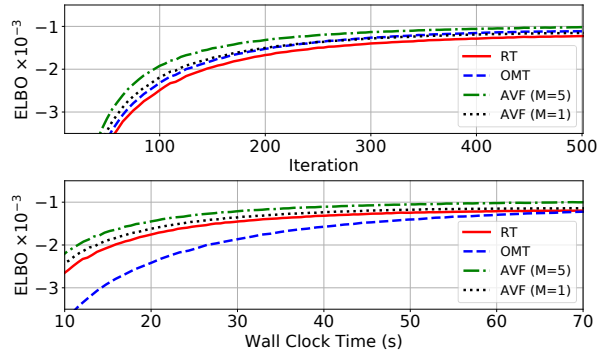


Figure 2: We compare the performance of various pathwise gradient estimators on the GP regression task in Sec 6.1.2. The AVF gradient estimator with $M = 5$ is the clear winner both in terms of wall clock time and the final ELBO achieved.

6.1 Adaptive Velocity Fields

6.1.1 Synthetic test functions

In Fig. 1a we use synthetic test functions to illustrate the variance reduction that is achieved with AVF gradients for the multivariate Normal distribution as compared to the reparameterization trick and OMT gradients from [18]. The dimension is $D = 50$; the results are qualitatively similar for different dimensions. Note that, as demonstrated here, whether AVF outperforms OMT is problem dependent.

6.1.2 Gaussian Process Regression

We investigate the performance of AVF gradients for the multivariate Normal distribution in the context of a Gaussian Process regression task reproduced from [18]. We model the Mauna Loa CO_2 data from [20] considered in [32]. We fit the GP using a single-sample Monte Carlo ELBO gradient estimator and all $N = 468$ data points (so $D = 468$). Both the OMT and AVF gradient estimators achieve higher ELBOs than the RT estimator (see Fig. 2); however, the OMT estimator does so at substantially increased computational cost ($\sim 1.9x$), while the AVF estimator for $M = 1$ ($M = 5$) requires only $\sim 6\%$ ($\sim 11\%$) more time per iteration. By iteration 250 the AVF gradient estimator with $M = 5$ has attained the same ELBO that the RT estimator attains at iteration 500.

6.2 Mixture Distributions

6.2.1 Synthetic test function

In Fig. 1b we depict the variance reduction of pathwise gradient estimators for various mixture distributions


compared to the corresponding score function estimator. We find that the typical variance reduction has magnitude $\mathcal{O}(D)$. The results are qualitatively similar for different numbers of components K . See the supplementary materials for details on the different families of mixture distributions. In Fig. 1c we compare the gradient estimator for mixtures of Normal distributions described in Sec. 4.2 to the estimator described in 9. We find that our estimator exhibits better performance when the mixture components have more overlap ($\sigma \gg 0$), while their estimator exhibits better performance in the opposite regime ($\sigma \rightarrow 0$). Since their estimator has a cost that is quadratic in the dimension D , while our estimator has a cost that is linear in the dimension D , which estimator is to be preferred will be problem specific.

6.2.2 Baseball Experiment


We consider a model for repeated binary trial data using the data in 8 and the modeling setup in 1 with partial pooling. The model has two global latent variables and 18 local latent variables so that the posterior is 20-dimensional. While mean field SGVI gives reasonable results for this model, it is not able to capture the detailed structure of the exact posterior as computed with the NUTS HMC implementation in Stan 16, 4. To go beyond mean field we consider variational distributions that are mixtures of K diagonal Normal distributions.¹² Adding more components gives a better approximation to the exact posterior, see Fig. 3a. Furthermore, using pathwise derivatives in the ELBO gradient estimator leads to faster convergence, see Fig. 3b. Here the hybrid estimator uses score functions gradients for the mixture logits but implements Eqn. 12 for gradients w.r.t. the component parameters. Since there is significant overlap among the mixture components, Eqn. 12 leads to substantial variance reduction.

6.2.3 Continuous state space model

We consider a non-linear state space model with observations \mathbf{x}_t and random variables \mathbf{z}_t at each time step (both two dimensional). Since the likelihood is of the form $p(\mathbf{x}_t|\mathbf{z}_t) = \mathcal{N}(\mathbf{x}_t|\boldsymbol{\mu}(\mathbf{z}_t), \sigma_x \mathbb{1}_2)$ with $\boldsymbol{\mu}(\mathbf{z}_t) = (z_{1t}^2, 2z_{2t})$, the posterior over $\mathbf{z}_{1:T}$ is highly multi-modal. We use SGVI to approximate the posterior over $\mathbf{z}_{1:T}$. To better capture the multi-modality we use a variational family that is a product of two-component mixture distributions of the form $q(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{x}_{t:T})$, one at each time step. As can be seen from Fig. 4 the variational family makes use of the mixture distributions at its disposal to model the multi-modal structure of the posterior, something that a variational family with a



(a) Two-dimensional cross-sections of three approximate posteriors; each cross-section includes one global latent variable, $\log(\kappa)$, and one local latent variable, $\text{logit}(\theta_0)$. The white density contours correspond to the different variational approximations, and the background depicts the approximate posterior computed with HMC. The quality of the variational approximation improves as we add more mixture components.



(b) ELBO training curves for three gradient estimators for $K = 20$. The lower variance of the pathwise and hybrid gradient estimators speeds up learning. The dotted lines denote mean ELBOs over 50 runs and the bands indicate 1- σ standard deviations. Note that the figure would be qualitatively similar if the ELBO were plotted against wall clock time, since the estimators have comparable computational cost (within 10%).

Figure 3: Variational approximations and ELBO training curves for the model in Sec. 6.2.2.

unimodal $q(\mathbf{z}_t|\cdot)$ struggles to do. In the next section we explore the use of mixture distributions in a much richer time series setup.

6.2.4 Deep Markov Model

We consider a variant of the Deep Markov Model (DMM) setup introduced in 23 on a polyphonic music dataset. We fix the form of the model and vary the variational family and gradient estimator used. We consider a variational family which includes a mixture distribution at each time step. In particular each factor $q(\mathbf{z}_t|\mathbf{z}_{t-1}, \mathbf{x}_{t:T})$ in the variational distribution is a mixture of K diagonal Normals. We compare the performance of the pathwise gradient estimator introduced in Sec. 4 to two variants of the Gumbel Softmax gradient estimator 17, 24, see Table 1. For a description of these two gradient estimators, which use the Gumbel Softmax trick to deal with mixture weight gradients, see Sec. E in the supplementary materials. Note that both Gumbel Softmax estimators are biased, while the pathwise estimator is unbiased. We find that the Gumbel Softmax estimators—which induce $\mathcal{O}(2 - 8)$ times higher gradient variance for the mixture logits

¹²Cf. the experiment in 26




Figure 4: Approximate marginal posteriors for the toy state space model in Sec. 6.2.3 for a randomly chosen test sequence for $t = 4, 5, 6$. The mixture distributions that serve as components of the variational family on the bottom help capture the multi-modal structure of the posterior.

as compared to the pathwise estimator—are unable to achieve competitive test ELBOs¹³

Note that while the numerical differences in Table 1 are not large (~ 0.05 nats), they are comparable to the quantitative improvement that one can achieve in this context when equipping the variational family with normalizing flows [41, 33]; for this particular task there appears to be little headroom above what can be attained with a (unimodal) Normal variational family.

Variational Family	Gradient Estimator		
	Pathwise	GS-Soft	GS-Hard
K=2	-6.84	-6.89	-6.88
K=3	-6.82	-6.87	-6.85

Table 1: Test ELBOs for the DMM in Sec. 6.2.4. Higher is better. For comparison: we achieve a test ELBO of -6.90 for a variational family in which each $q(z_t|\cdot)$ is a diagonal Normal.

6.2.5 VAE

We conduct a simple experiment with a VAE [22, 34] trained on MNIST. We fix the form of the model and vary the gradient estimator used. The variational family is of the form $q(z|\mathbf{x})$, where $q(z|\cdot)$ is a mixture of $K \in [3, 4, 5]$ diagonal Normals. We find that on

¹³We also found the Gumbel Softmax estimators to be less numerically stable, which prevented us from using more aggressive KL annealing schedules. This may have contributed to the performance gap in Table 1

this task our gradient estimator achieves similar performance to the two Gumbel Softmax estimators and the score function estimator, see Table 2. This is not too surprising, since VAE posteriors tend to be quite sharp for any given \mathbf{x} so that gradient variance is not the bottleneck for learning.

Var. Family	Gradient Estimator			
	Pathwise	GS-Soft	GS-Hard	SF
K=3	-95.88	-95.95	-95.96	-95.98
K=4	-95.89	-95.84	-95.88	-95.99
K=5	-96.04	-95.94	-95.89	-95.98

Table 2: Test ELBOs for the VAE in Sec. 6.2.5. Higher is better.

7 Discussion

We have seen that the link between the transport equation and derivatives of expectations offers a powerful tool for constructing pathwise gradient estimators. In the case of variational inference, we emphasize that—in addition to providing gradients with reduced variance—pathwise gradient estimators can enable quick iteration over complex models and variational families, since there is no need to reason about ELBO gradients. The modeler need only construct a Monte Carlo estimator of the ELBO; automatic differentiation will handle the rest. Our hope is that an expanding toolkit of pathwise gradient estimators can contribute to the development of innovative applications of variational inference.

One limitation of our approach is that solving the transport equation for arbitrary families of distributions presents a challenge. Of course, this limitation also applies to the reparameterization trick. While we acknowledge this limitation, we note that since: i) the Normal distribution is ubiquitous in machine learning; and ii) mixtures of Normals are universal approximators for continuous distributions, improving the gradient estimators that can be used in these two cases is already of considerable value.

Given the richness of the transport equation and the great variety of multivariate distributions, there are several exciting possibilities for future research. It would be of interest to identify more cases where we can construct null solutions to the transport equation, e.g. generalizing Eqn. 9 to higher orders. It could also be fruitful to consider other ways of adapting velocity fields. In some cases this could lead to improved variance reduction, especially in the case of mixture distributions, where the optimal transport setup in [6] might provide a path to constructing alternative gradient estimators.

References

- [1] Stan modeling language users guide and reference manual, version 2.17.0. <http://mc-stan.org/users/documentation/case-studies/pool-binary-trials.html>, 2017.
- [2] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, and N. D. Goodman. Pyro: Deep universal probabilistic programming. *Journal of Machine Learning Research*, 2018.
- [3] N. Boulanger-Lewandowski, Y. Bengio, and P. Vincent. Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription. *arXiv preprint arXiv:1206.6392*, 2012.
- [4] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell. Stan: A probabilistic programming language. *Journal of statistical software*, 76(1), 2017.
- [5] G. Casella and C. P. Robert. Rao-blackwellisation of sampling schemes. *Biometrika*, 83(1):81–94, 1996.
- [6] Y. Chen, T. T. Georgiou, and A. Tannenbaum. Optimal transport for gaussian mixture models. *arXiv preprint arXiv:1710.07876*, 2017.
- [7] J. Dillon and I. Langmore. Quadrature compound: An approximating family of distributions. *arXiv preprint arXiv:1801.03080*, 2018.
- [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. *arXiv preprint arXiv:1805.08498*, 2018.
- [10] M. C. Fu. Gradient estimation. *Handbooks in operations research and management science*, 13:575–616, 2006.
- [11] P. Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer Science & Business Media, 2013.
- [12] P. W. Glynn. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM*, 33(10):75–84, 1990.
- [13] W. Grathwohl, D. Choi, Y. Wu, G. Roeder, and D. Duvenaud. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. *arXiv preprint arXiv:1711.00123*, 2017.
- [14] A. Graves. Stochastic backpropagation through mixture density distributions. *arXiv preprint arXiv:1607.05690*, 2016.
- [15] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- [16] M. D. Hoffman and A. Gelman. The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- [17] E. Jang, S. Gu, and B. Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [18] M. Jankowiak and F. Obermeyer. Pathwise derivatives beyond the reparameterization trick. In *Proceedings of the 35th International Conference on Machine Learning*, pages 2235–2244, 2018.
- [19] 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.
- [20] C. D. Keeling and T. P. Whorf. Atmospheric co2 concentrations derived from flask air samples at sites in the sio network. *Trends: a compendium of data on Global Change*, 2004.
- [21] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [22] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [23] R. G. Krishnan, U. Shalit, and D. Sontag. Structured inference networks for nonlinear state space models. In *AAAI*, pages 2101–2109, 2017.
- [24] C. J. Maddison, A. Mnih, and Y. W. Teh. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- [25] A. Miller, N. Foti, A. D’Amour, and R. P. Adams. Reducing reparameterization gradient variance. In *Advances in Neural Information Processing Systems*, pages 3711–3721, 2017.

- [26] A. C. Miller, N. Foti, and R. P. Adams. Variational boosting: Iteratively refining posterior approximations. *arXiv preprint arXiv:1611.06585*, 2016.
- [27] A. Mnih and K. Gregor. Neural variational inference and learning in belief networks. *arXiv preprint arXiv:1402.0030*, 2014.
- [28] C. Naesseth, F. Ruiz, S. Linderman, and D. Blei. Reparameterization gradients through acceptance-rejection sampling algorithms. In *Artificial Intelligence and Statistics*, pages 489–498, 2017.
- [29] J. Paisley, D. Blei, and M. Jordan. Variational bayesian inference with stochastic search. *arXiv preprint arXiv:1206.6430*, 2012.
- [30] R. Price. A useful theorem for nonlinear devices having gaussian inputs. *IRE Transactions on Information Theory*, 4(2):69–72, 1958.
- [31] R. Ranganath, S. Gerrish, and D. Blei. Black box variational inference. In *Artificial Intelligence and Statistics*, pages 814–822, 2014.
- [32] C. E. Rasmussen. Gaussian processes in machine learning. In *Advanced lectures on machine learning*, pages 63–71. Springer, 2004.
- [33] D. J. Rezende and S. Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- [34] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- [35] 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.
- [36] S. M. Ross. *Simulation*. Academic Press, San Diego, 2006.
- [37] F. J. Ruiz, M. K. Titsias, and D. M. Blei. Overdispersed black-box variational inference. *arXiv preprint arXiv:1603.01140*, 2016.
- [38] F. R. Ruiz, M. T. R. AUEB, and D. Blei. The generalized reparameterization gradient. In *Advances in Neural Information Processing Systems*, pages 460–468, 2016.
- [39] T. Salimans, D. A. Knowles, et al. Fixed-form variational posterior approximation through stochastic linear regression. *Bayesian Analysis*, 8(4):837–882, 2013.
- [40] J. Schulman, N. Heess, T. Weber, and P. Abbeel. Gradient estimation using stochastic computation graphs. In *Advances in Neural Information Processing Systems*, pages 3528–3536, 2015.
- [41] E. Tabak and C. V. Turner. A family of non-parametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013.
- [42] M. Titsias and M. Lázaro-Gredilla. Doubly stochastic variational bayes for non-conjugate inference. In *International Conference on Machine Learning*, pages 1971–1979, 2014.
- [43] G. Tucker, A. Mnih, C. J. Maddison, J. Lawson, and J. Sohl-Dickstein. Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, pages 2624–2633, 2017.
- [44] R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [45] D. Wingate and T. Weber. Automated variational inference in probabilistic programming. *arXiv preprint arXiv:1301.1299*, 2013.