Provable Gradient Variance Guarantees for Black-Box Variational Inference

Justin Domke College of Information and Computer Sciences University of Massachusetts Amherst domke@cs.umass.edu

Abstract

Recent variational inference methods use stochastic gradient estimators whose variance is not well understood. Theoretical guarantees for these estimators are important to understand when these methods will or will not work. This paper gives bounds for the common "reparameterization" estimators when the target is smooth and the variational family is a location-scale distribution. These bounds are unimprovable and thus provide the best possible guarantees under the stated assumptions.

1 Introduction

Take a distribution p(z, x) representing relationships between data x and latent variables z. After observing x, one might wish to approximate the marginal probability p(x) or the posterior p(z|x). Variational inference (VI) is based on the simple observation that for any distribution q(z),

$$\log p(\boldsymbol{x}) = \underbrace{\mathbb{E}}_{\substack{\boldsymbol{z} \sim q}} \log \frac{p(\boldsymbol{z}, \boldsymbol{x})}{q(\boldsymbol{z})}_{\text{ELBO}(q)} + KL(q(\boldsymbol{z}) \| p(\boldsymbol{z} | \boldsymbol{x})).$$
(1)

VI algorithms typically choose an approximating family q_w and maximize ELBO (q_w) over w. Since $\log p(x)$ is fixed, this simultaneously tightens a lower-bound on $\log p(x)$ and minimizes the divergence from $q_w(z)$ to the posterior p(z|x).

Traditional VI algorithms suppose p and q_w are simple enough for certain expectations to have closed forms, leading to deterministic coordinate-ascent type algorithms [6, 1, 20]. Recent work has turned towards stochastic optimization. There are two motivations for this. First, stochastic data subsampling can give computational savings [7]. Second, more complex distributions can be addressed if p is treated as a "black box", with no expectations available [9, 15, 19]. In both cases, one can still estimate a *stochastic* gradient of the ELBO [17] and thus use stochastic gradient optimization. It is possible to address very complex and large-scale problems using this strategy [10].

These improvements in scale and generality come at a cost: Stochastic optimization is typically less reliable than deterministic coordinate ascent. Convergence is often a challenge, and methods typically use heuristics for parameters like step-sizes. Failures do frequently occur in practice [22, 11, 4].

To help understand when black-box VI can be expected to work, this paper investigates the variance of gradient estimates. This is a major issue in practice, and many ideas have been proposed to attempt to reduce the variance [8, 5, 12, 2, 18, 13, 14, 16]. Despite all this, few rigorous guarantees on the variance of gradient estimators seem to be known (Sec. 5.1).

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

1.1 Contributions

This paper studies "reparameterization" (RP) or "path" based gradient estimators when q_w is in a multivariate location-scale family. We decompose ELBO $(q_w) = l(w) + h(w)$ where h(w) is the entropy of q_w (known in closed-form) and $l(w) = \mathbb{E}_{z \sim q_w} \log p(z, x)$. The key assumption is that $\log p(z, x)$ is (Lipschitz) *smooth* as a function of z, meaning that $\nabla_z \log p(z, x)$ can't change too quickly as z changes. Formally f(z) is M-smooth if $\|\nabla f(z) - \nabla f(z')\|_2 \le M \|z - z'\|_2$.

- **Bound for smooth target distributions:** If g is the RP gradient estimator of $\nabla l(w)$ and $\log p$ is M-smooth, then $\mathbb{E} \|\mathbf{g}\|^2$ is bounded by a quadratic function of w (Thm. 3). With a small relaxation, this is $\mathbb{E} \|\mathbf{g}\|^2 \leq aM^2 \|\mathbf{w} \bar{\mathbf{w}}\|^2$ (Eq. (3)) where \bar{w} are fixed parameters and a is determined by the location-scale family.
- **Generalized bound:** We extend this result to consider a more general notion of "matrix" smoothness (Thm. 5) reflecting that the sensitivity of $\nabla_{z} \log p(z, x)$ to changes in z may depend on the direction of change.
- **Data Subsampling:** We again extend this result to consider data subsampling (Thm. 6). In particular, we observe that *non-uniform* subsampling gives tighter bounds.

In all cases, we show that the bounds are unimprovable. We experimentally compare these bounds to empirical variance.

2 Setup

Given some "black box" function f, this paper studies estimating gradients of functions l of the form $l(w) = \mathbb{E}_{z \sim q_w} f(z)$. Now, suppose some base distribution s and mapping \mathcal{T}_w are known such that if $u \sim s$, then $\mathcal{T}_w(u) \sim q_w$. Then, l can be written as

$$l(\boldsymbol{w}) = \mathbb{E}_{\boldsymbol{w}} f(\mathcal{T}_{\boldsymbol{w}}(\boldsymbol{u})).$$

If we define $g = \nabla_{w} f(\mathcal{T}_{w}(u))$, then g is an unbiased estimate of ∇l , i.e. $\mathbb{E} g = \nabla l(w)$. The same idea can be used when f is composed as a finite sum as $f(z) = \sum_{n=1}^{N} f_n(z)$. If N is large, even evaluating f once might be expensive. However, take any positive distribution π over $n \in \{1, \dots, N\}$ and sample $n \sim \pi$ independently of u. Then, if we define $g = \nabla_{w} \pi(n)^{-1} f_n(\mathcal{T}_{w}(u))$, this is again an unbiased estimator with $\mathbb{E} g = \nabla l(w)$.

Convergence rates in stochastic optimization depend on the variability of the gradient estimator, typically either via the expected squared norm (ESN) $\mathbb{E} \|\mathbf{g}\|_2^2$ or the trace of the variance tr \mathbb{V} g. These are closely related, since $\mathbb{E} \|\mathbf{g}\|_2^2 = \operatorname{tr} \mathbb{V} \mathbf{g} + \|\mathbb{E} \mathbf{g}\|_2^2$.

The goal of this paper is to bound the variability of g for reparameterization / path estimators of g. This requires making assumptions about (i) the transformation function \mathcal{T}_w and base distribution s (which determine q_w) and (ii) the target function f.

Here, we are interested in the case of affine mappings. We use the mapping[17]

$$\mathcal{T}_{\boldsymbol{w}}(\boldsymbol{u}) = C\boldsymbol{u} + \boldsymbol{m},$$

where $\boldsymbol{w} = (\boldsymbol{m}, C)$ is a single vector of all parameters. This is the most common mapping used to represent location-scale families. That is, if $\boldsymbol{u} \sim s$ then $\mathcal{T}_{\boldsymbol{w}}(\boldsymbol{u})$ is equal in distribution to a location-scale family distribution. For example, if $s = \mathcal{N}(0, I)$ then $\mathcal{T}_{\boldsymbol{w}}(\boldsymbol{u})$ is equal in distribution to $\mathcal{N}(\boldsymbol{m}, CC^{\top})$.

We will refer to the base distribution as **standardized** if the components of $u = (u_1, \dots, u_d) \sim s$ are iid with $\mathbb{E} u_1 = \mathbb{E} u_1^3 = 0$ and $\mathbb{V} u_1 = 1$. The bounds will depend on the fourth moment $\kappa = \mathbb{E}[u_1^4]$, but are otherwise independent of *s*.

To apply these estimators to VI, choose $f(z) = \log(z, x)$. Then ELBO(w) = l(w) + h(w) where h is the entropy of q_w . Stochastic estimates of the gradient of l can be employed in a stochastic gradient method to maximize the ELBO. To model the stochastic setting, suppose that $X = (x_1, \dots, x_N)$ are iid and $p(z, X) = p(z) \prod_{n=1}^{N} p(x_n | z)$. Then, one may choose, e.g. $f_n(z) =$

 $\frac{1}{N}\log p(\boldsymbol{z}) + \log p(\boldsymbol{x}_n|\boldsymbol{z})$. The entropy h is related to the (constant) entropy of the base distribution as $h(\boldsymbol{w}) = \text{Entropy}(s) + \log |C|$.

The main bounds of this paper concern estimators for the gradient of l(w) alone, disregarding h(w). There are two reasons for this. First, in location-scale families, the exact gradient of h(w) is known. Second, if one uses a stochastic estimator for h(w), this can be "absorbed" into l(w) to some degree. This is discussed further in Sec. 5.

3 Variance Bounds

3.1 Technical Lemmas

We begin with two technical lemmas which will do most of the work in the main results. Both have (somewhat laborious) proofs in Sec. 7 (Appendix). The first lemma relates the norm of the *parameter* gradient of $f(\mathcal{T}_{w}(u))$ (with respect to w) to the norm of the gradient of f(z) itself, evaluated at $z = \mathcal{T}_{w}(u)$.

Lemma 1. For any *w* and *u*, $\|\nabla_{w} f(\mathcal{T}_{w}(u))\|_{2}^{2} = \|\nabla f(\mathcal{T}_{w}(u))\|_{2}^{2} (1 + \|u\|_{2}^{2})$.

The proof is tedious but essentially amounts to calculating the derivative with respect to each component of w (i.e. entries m_i and C_{ij}), summing the square of all entries, and simplifying. The second lemma gives a closed-form for the expectation of a closely related expression that will appear in the proof of Thm. 3 as a consequence of applying Lem. 1.

Lemma 2. Let $u \sim s$ for s standardized with $u \in \mathbb{R}^d$ and $\mathbb{E}_{u \sim s} u_i^4 = \kappa$. Then for any \bar{z} ,

$$\mathbb{E} \|\mathcal{T}_{\boldsymbol{w}}(\mathbf{u}) - \bar{\boldsymbol{z}}\|_{2}^{2} \left(1 + \|\mathbf{u}\|_{2}^{2}\right) = (d+1) \|\boldsymbol{m} - \bar{\boldsymbol{z}}\|_{2}^{2} + (d+\kappa) \|C\|_{F}^{2}.$$

Again, the proof is tedious but based on simple ideas: Substitute the definition of \mathcal{T}_w into the left-hand side and expand all terms. This gives terms between zeroth and fourth order (in u). Calculating the exact expectation of each and simplifying using the assumption that s is standardized gives the result.

3.2 Basic Variance Bound

Given these two lemmas, we give our major technical result, bounding the variability of a reparameterization-based gradient estimator. This will be later be extended to consider data subsampling, and a generalized notion of smoothness. Note that we do *not* require that f be convex.

Theorem 3. Suppose f is M-smooth, \bar{z} is a stationary point of f, and s is standardized with $u \in \mathbb{R}^d$ and $\mathbb{E} u_i^4 = \kappa$. Let $g = \nabla_w f(\mathcal{T}_w(u))$ for $u \sim s$. Then,

$$\mathbb{E} \|\mathbf{g}\|_{2}^{2} \leq M^{2} \left((d+1) \|\boldsymbol{m} - \bar{\boldsymbol{z}}\|_{2}^{2} + (d+\kappa) \|C\|_{F}^{2} \right).$$
⁽²⁾

Moreover, this result is unimprovable without further assumptions.

Proof. We expand the definition of g, and use the above lemmas and the smoothness of f.

$$\begin{split} \mathbb{E} \|\mathbf{g}\|_{2}^{2} &= \mathbb{E} \|\nabla_{\boldsymbol{w}} f(\mathcal{T}_{\boldsymbol{w}}(\mathbf{u}))\|_{2}^{2} & \text{(Definition of g)} \\ &= \mathbb{E} \|\nabla f(\mathcal{T}_{\boldsymbol{w}}(\mathbf{u}))\|_{2}^{2} \left(1 + \|\mathbf{u}\|_{2}^{2}\right) & \text{(Lem. 1)} \\ &= \mathbb{E} \|\nabla f(\mathcal{T}_{\boldsymbol{w}}(\mathbf{u})) - \nabla f(\bar{z})\|_{2}^{2} \left(1 + \|\mathbf{u}\|_{2}^{2}\right) & (\nabla f(\bar{z}) = 0) \\ &\leq \mathbb{E} M^{2} \|\mathcal{T}_{\boldsymbol{w}}(\mathbf{u}) - \bar{z}\|_{2}^{2} \left(1 + \|\mathbf{u}\|_{2}^{2}\right) & \text{(f is smooth)} \\ &= M^{2} \left((d+1) \|\boldsymbol{m} - \bar{z}\|_{2}^{2} + (d+\kappa) \|C\|_{F}^{2} \right). & \text{(Lem. 2)} \end{split}$$

To see that this is unimprovable without further assumptions, observe that the only inequality is using the smoothness on f to bound the norm of the difference of gradients at $\mathcal{T}_{w}(u)$ and at \bar{z} . But for $f(z) = \frac{M}{2} ||z - \bar{z}||_{2}^{2}$ this inequality is tight. Thus, for any M and \bar{z} , there is a function f satisfying the assumptions of the theorem such that Eq. (2) is an equality.

With a small amount of additional looseness, we can cast Eq. (2) into a more intuitive form. Define $\bar{w} = (\bar{z}, 0_{d \times d})$, where $0_{d \times d}$ is a $d \times d$ matrix of zeros. Then, $\|w - \bar{w}\|_2^2 = \|m - \bar{z}\|_2^2 + \|C\|_F^2$, so we can slightly relax Eq. (2) to the more user-friendly form of

$$\mathbb{E} \left\| \mathbf{g} \right\|_{2}^{2} \le (d+\kappa) M^{2} \left\| \boldsymbol{w} - \bar{\boldsymbol{w}} \right\|_{2}^{2}.$$
(3)

The only additional looseness is bounding $d+1 \le d+\kappa$. This is justified since when s is standardized, $\kappa = u_i^4$ is the kurtosis, which is at least one. Here, κ is determined by s and does not depend on the dimensionality. For example, if s is Gaussian, $\kappa = 3$. Thus, Eq. (3) will typically not be much looser than Eq. (2).

Intuitively, \bar{w} are parameters that concentrate q entirely at a stationary point of f. It is not hard to show that $\|w - \bar{w}\|^2 = \mathbb{E}_{z \sim q_w} \|z - \bar{z}\|^2$. Thus, Eq. (3) intuitively says that $\mathbb{E} \|g\|^2$ is bounded in terms of how far far the average point sampled from q_w is from \bar{z} . Since f need not be convex, there might be multiple stationary points. In this case, Thm. 3 holds simultaneously for all of them.

3.3 Generalized Smoothness

Since the above bound is not improvable, tightening it requires stronger assumptions. The tightness of Thm. 3 is determined by the smoothness condition that the difference of gradients at two points is bounded as $\|\nabla f(y) - \nabla f(z)\|_2 \leq M \|y - z\|_2$. For some problems, f may be much smoother in *certain directions* than others. In such cases, the smoothness constant M will need to reflect the worst-case direction. To produce a tighter bound for such situations, we generalize the notion of smoothness to allow M to be a symmetric matrix.

Definition 4. f is M-matrix-smooth if $\|\nabla f(y) - \nabla f(z)\|_2 \le \|M(y-z)\|_2$ (for symmetric M).

We can generalize the result in Thm. 3 to functions with this matrix-smoothness condition. The proof is very similar. The main difference is that after applying the smoothness condition, the matrix M needs to be "absorbed" into the parameters w = (m, C) before applying Lem. 2.

Theorem 5. Suppose f is M-matrix smooth, \bar{z} is a stationary point of f, and s is standardized with $u \in \mathbb{R}^d$ and $\mathbb{E} u_i^4 = \kappa$. Let $g = \nabla_w f(\mathcal{T}_w(u))$ for $u \sim s$. Then,

$$\mathbb{E} \|\mathbf{g}\|_{2}^{2} \leq (d+1) \|M(\boldsymbol{m} - \bar{\boldsymbol{z}})\|_{2}^{2} + (d+\kappa) \|MC\|_{F}^{2}.$$
(4)

Moreover, this result is unimprovable without further assumptions.

Proof. The proof closely mirrors that of Thm. 3. Here, given w = (m, C), we define v = (Mm, MC), to be w with M "absorbed" into the parameters.

$$\begin{split} \mathbb{E} \|\mathbf{g}\|_{2}^{2} &= \mathbb{E} \|\nabla_{\boldsymbol{w}} f(\mathcal{T}_{\boldsymbol{w}}(\mathbf{u}))\|_{2}^{2} & \text{Definition of g}) \\ &= \mathbb{E} \|\nabla f(\mathcal{T}_{\boldsymbol{w}}(\mathbf{u}))\|_{2}^{2} (1 + \|\mathbf{u}\|_{2}^{2}) & (\text{Lem. 1}) \\ &= \mathbb{E} \|\nabla f(\mathcal{T}_{\boldsymbol{w}}(\mathbf{u})) - \nabla f(\bar{\boldsymbol{z}})\|_{2}^{2} (1 + \|\mathbf{u}\|_{2}^{2}) & (\nabla f(\bar{\boldsymbol{z}}) = 0) \\ &\leq \mathbb{E} \|M(\mathcal{T}_{\boldsymbol{w}}(\mathbf{u}) - \bar{\boldsymbol{z}})\|_{2}^{2} (1 + \|\mathbf{u}\|_{2}^{2}) & (f \text{ is smooth}) \\ &= \mathbb{E} \|\mathcal{T}_{\boldsymbol{v}}(\mathbf{u}) - M(\bar{\boldsymbol{z}} - \boldsymbol{m})\|_{2}^{2} (1 + \|\mathbf{u}\|_{2}^{2}) & (\text{Absorb } M \text{ into } \boldsymbol{v}) \\ &= (d+1) \|M\boldsymbol{m} - M\bar{\boldsymbol{z}}\|_{2}^{2} + (d+\kappa) \|MC\|_{F}^{2}. & (\text{Lem. 2}) \end{split}$$

To see that this is unimprovable, observe that the only inequality is the matrix-smoothness condition on f. But for $f(z) = \frac{1}{2}(z - \bar{z})^{\top}M(z - \bar{z})$, the difference of gradients $\|\nabla f(y) - \nabla f(z)\|_2 =$ $\|M(y - z)\|_2$ is an equality. Thus, for any M and \bar{z} , there is an f satisfying the assumptions of the theorem such that the bound in Eq. (4) is an equality. \Box

It's easy to see that this reduces to Thm. 3 in the case that f is smooth in the standard sense– this corresponds to the situation where M is some constant times the identity. Alternatively, one can simply observe that the two results are the same if M is a scalar. Thus, going forward we will use Eq. (4) to represent the result with either type of smoothness assumption on f.

3.4 Subsampling

Often, the function f(z) takes the form of a sum over other functions $f_n(z)$, typically representing different data. Write this as

$$f(\boldsymbol{z}) = \sum_{n=1}^{N} f_n(\boldsymbol{z}).$$

To estimate the gradient of $\mathbb{E}_{u \sim s} f(\mathcal{T}_{w}(u))$, one can save time by using "subsampling": That is, draw a random n, and then estimate the gradient of $\mathbb{E}_{u \sim s} f_n(\mathcal{T}_w(u))$. The following result bounds this procedure. It essentially just takes a set of estimators, one corresponding to each function f_n , bounds their expected squared norm using the previous theorems, and then combines these.

Theorem 6. Suppose f_n is M_n -matrix-smooth, \bar{z}_n is a stationary point of f_n , and s is standardized with $\mathbf{u} \in \mathbb{R}^d$ and $\mathbb{E} \mathbf{u}_i^4 = \kappa$. Let $\mathbf{g} = \frac{1}{\pi(n)} \nabla f_n(\mathcal{T}_{\boldsymbol{w}}(\mathbf{u}))$ for $\mathbf{u} \sim s$ and $\mathbf{n} \sim \pi$ independent. Then,

$$\mathbb{E} \|\mathbf{g}\|_{2}^{2} \leq \sum_{n=1}^{N} \frac{1}{\pi(n)} \left((d+1) \|M_{n}(\boldsymbol{m}-\bar{\boldsymbol{z}}_{n})\|_{2}^{2} + (d+\kappa) \|M_{n}C\|_{F}^{2} \right).$$
(5)

Moreover, this result is unimprovable without further assumptions.

Proof. Consider a simple lemma: Suppose $a_1 \cdots a_N$ are independent random vectors and π is any distribution over $1 \cdots N$. Let $b = a_n/\pi(n)$ for $n \sim \pi$, where n is independent of a_n . It is easy to show that $\mathbb{E} b = \sum_{n=1}^{N} \mathbb{E} a_n$ and $\mathbb{E} \|b\|_2^2 = \sum_n \mathbb{E} \|a_n\|_2^2/\pi(n)$. The result follows from applying this with $a_n = \nabla_{\boldsymbol{w}} f_n(\mathcal{T}_{\boldsymbol{w}}(u))$, and then bounding $\mathbb{E} \|a_n\|_2^2$ using Thm. 5.

Again, in this result the only source of looseness is the use of the smoothness bound for the component functions f_n . Accordingly, the result can be shown to be unimprovable: For any set of stationary points \bar{z} and smoothness parameters M_n we can construct functions f_n (as in Thm. 5) for which the previous theorems are tight and thus this result is also tight.

This result generalizes all the previous bounds: Thm. 5 is the special case when N = 1, while Thm. 3 is the special-case when N = 1 and f_1 is M_1 -smooth (for a scalar M_1). The case where N > 1 but f_n is M_n -smooth (for scalar M_n) is also useful– the bound in Eq. (5) remains valid, but with a scalar M_n .

4 Empirical Evaluation

4.1 Model and Datasets

We consider Bayesian linear regression and logistic regression models on various datasets (Table 1). Given data $\{(\boldsymbol{x}_1, y_1), \cdots (\boldsymbol{x}_N, y_N)\}$, let \boldsymbol{y} be a vector of all y_n and X a matrix of all \boldsymbol{x}_n . We assume a Gaussian prior so that $p(\boldsymbol{z}, \boldsymbol{y}|X) = \mathcal{N}(\boldsymbol{z}|0, \sigma^2 I) \prod_{n=1}^{N} p(y_n|\boldsymbol{z}, \boldsymbol{x}_n)$. For linear regression, $p(y_n|\boldsymbol{z}, \boldsymbol{x}_n) = \mathcal{N}(y_n|\boldsymbol{z}^\top \boldsymbol{x}_i, \rho^2)$, while for logistic regression, $p(y_n|\boldsymbol{z}, \boldsymbol{x}_n) = \text{Sigmoid}(y_n \boldsymbol{z}^\top \boldsymbol{x}_n)$. For both models we use a prior of $\sigma^2 = 1$. For linear regression, we set $\rho^2 = 4$.

To justify the use of VI, apply the decomposition in Eq. (1) substituting p(z, y|X) in place of p(z, x) to get that

$$\log p(\boldsymbol{y}|X) = \mathbb{E}_{\mathbf{z} \sim q} \log \frac{p(\mathbf{z}, \boldsymbol{y}|X)}{q(\mathbf{z})} + KL(q(\mathbf{z}) || p(\mathbf{z}|\boldsymbol{y}, X)).$$

Thus, adjusting the parameters of q to maximize the first term on the right tightens a lower-bound on the conditional log likelihood log p(y|X) and minimizes the divergence from q to the posterior. So, we again take our goal as maximizing l(w) + h(w). In the batch setting, $f(z) = \log p(z, y|X)$, while with subsampling, $f_n(z) = \frac{1}{N} \log p(z) + \log p(y_n|z, x_n)$.



Figure 1: How loose are the bounds compared to reality? Odd Rows: Evolution of the ELBO during the single optimization trace used to compare all estimators. Even Rows: True and bounded variance with gradients estimated in "batch" (using the full dataset in each evaluation) and "uniform" (stochastically with $\pi(n) = 1/N$). The first two rows are for linear regression models, while the rest are for logistic regression. Key Observations: (i) Batch estimation is lower-variance but higher cost (ii) variance with stochastic estimation varies little over time (iii) using matrix smoothness significantly tightens bounds – and is exact for linear regression models.

Sec. 8 shows that if $0 \le \phi''(t) \le \theta$, then $\sum_{n=1}^{N} \phi(a_n^{\top} z + b_n)$ is *M*-matrix-smooth for $M = \theta \sum_{i=1}^{N} a_i a_i^{\top}$. Applying this¹ gives that f(z) and $f_n(z)$ are matrix-smooth for

$$M = rac{1}{\sigma^2}I + c\sum_{n=1}^N oldsymbol{x}_noldsymbol{x}_n^ op, ext{ and } M_n = rac{1}{N\sigma^2}I + c\,oldsymbol{x}_noldsymbol{x}_n^ op,$$

¹For linear regression, set $\phi(t) = -t^2/(2\rho^2)$, $a_n = x_n$ and $b_n = y_n$ and observe that $\phi'' = -1/\rho^2$. For logistic regression, set $\phi(t) = \log \text{Sigmoid}(t)$, $a_n = y_n x_n$ and $b_n = 0$ and observe that $\phi'' \leq 1/4$. Adding the prior and using the triangle inequality gives the result.



Figure 2: Tightening variance bounds reduces true variance. A comparison of the true (vertical bars) and bounded $\mathbb{E} \|\mathbf{g}\|^2$ values produced using five different gradient estimators. Batch does not use subsampling. Uniform uses subsampling $\pi(n) = 1/N$, proportional uses $\pi(n) \propto M_n$, opt (scalar) numerically optimizes $\pi(n)$ to tighten Eq. (5) with a scalar M_n and opt (matrix) tightens Eq. (5) with a matrix M_n . For each sampling strategy, we show the variance bound both with a scalar and matrix M_n . Uniform sampling has true and bounded values of $\mathbb{E} \|\mathbf{g}\|^2$ ranging between 1.5x and 10x higher than those for sampling with π numerically optimized.

where $c = 1/\rho^2$ for linear regression, and c = 1/4 for logistic regression. Taking the spectral norm of these matrices gives scalar smoothness constants. With subsampling, this is $||M_n||_2 = \frac{1}{\sigma^2 N} + c||\boldsymbol{x}_n||^2$.

4.2 Evaluation of Bounds

To enable a clear comparison of of different estimators and bounds, we generate a single optimization trace of parameter vectors w for each dataset. All comparisons use this same trace. These use a conservative optimization method: Find a maximum \bar{z} and then initialize to $w = (\bar{z}, 0)$. Then, optimization

Dataset	Type	# data	# dims
boston	r	506	13
fires	r	517	12
cpusmall	r	8192	13
ala	с	1695	124
ionosphere	с	351	35
australian	с	690	15
sonar	с	208	61
mushrooms	с	8124	113

Table 1: Regression (r) and classification (c) datasets

uses proximal stochastic gradient descent (with the proximal operator reflecting h) with a step size of 1/M (the scalar smoothness constant) and 1000 evaluations for each gradient estimate.

Fig. 1 shows the evolution of the ELBO along with the variance of gradient estimation either in batch or stochastically with a uniform distribution over data. For each iteration and estimator, we plot the empirical $\|\mathbf{g}\|^2$ along with this paper's bounds using either scalar or matrix smoothness.

4.3 Sampling distributions

With subsampling, variability depends on the sampling distribution π . We consider uniform sampling as well as three strategies that attempt to tighten the bound in Thm. 6. In general, $\sum_n f(n)^2/\pi(n)$ is minimized over distributions π by $\pi(n) \propto |f(n)|$. Thus, the tightest bound is given by

$$\pi_{\boldsymbol{w}}^{*}(n) \propto \sqrt{(d+1) \|M_{n}(\boldsymbol{m}-\bar{\boldsymbol{z}}_{n})\|_{2}^{2} + (d+\kappa) \|M_{n}C\|_{F}^{2}}.$$
(6)

We call this "opt (scalar)" or "opt (matrix)" when M_n is a scalar or matrix, respectively. We also consider a "proportional" heuristic with $\pi(n) \propto M_n$ for a scalar M_n . Sampling from Eq. (6) appears to require calculating the right-hand side for each n and then normalizing, which may not be practical for large datasets. While there are obvious heuristics for recursively approximating π^* during an optimization, to maintain focus we do not pursue these ideas here.

Fig. 2 shows the empirical and true variance at the final iteration of the optimization shown in Fig. 1. The basic conclusion is that using a more careful sampling distribution reduces both true and empirical variance.

5 Discussion

5.1 Related work

Xu et al. [21] compute the variance of a reparameterization estimator applied to a quadratic function, when the variational distribution is a fully-factorized Gaussian. This paper can be seen as extending this result to more general densities (full-rank location-scale families) and more general target functions (smooth functions).

Fan et al. [4] give an abstract variance bound for RP estimators. Essentially, they argue that if $g_i = \nabla_{w_i} f(\mathcal{T}_w(u))$ and $\nabla_{w_i} f(\mathcal{T}_w(u))$ is *M*-smooth as a function of u, then $\mathbb{V}[g_i] \leq M^2 \pi^2/4$ when $u \sim \mathcal{N}(0, I)$. While this result is fairly abstract – there is no proof that the smoothness assumption holds for any particular M with any particular f and \mathcal{T}_w – it is similar in spirit to the results in this paper.

5.2 Variance vs Expected Squared Norms

The above results are on the the expected squared norm (ESN) of the gradient $\mathbb{E} \|\mathbf{g}\|^2$. Some stochastic gradient convergence rates instead consider (the trace of) the variance $\mathbb{V}[\mathbf{g}]$. Since tr $\mathbb{V}[\mathbf{g}] = \mathbb{E} \|\mathbf{g}\|^2 - \|\mathbb{E}\mathbf{g}\|^2$, ESN bounds are valid as variance bounds. Still, one can ask if these bounds are loose. The following (proof in Sec. 7.3) gives a lower-bound that shows that there is not much to gain from a direct bound on the variance rather than just using the ESN bound from Thm. 6.

Theorem 7. For any symmetric matrices M_1, \dots, M_N and vectors $\bar{z}_1, \dots, \bar{z}_N$, there are functions f_1, \dots, f_N such that (1) f_n is M_n -matrix-smooth and has a stationary point at \bar{z}_n and (2) if s is standardized with $u \in \mathbb{R}^d$ and $\mathbb{E} u_i^4 = \kappa$, then for $g = \frac{1}{\pi(n)} \nabla f_n(\mathcal{T}_w(u))$,

$$\operatorname{tr} \mathbb{V} \|\mathbf{g}\|_{2}^{2} \geq \sum_{n=1}^{N} \frac{1}{\pi(n)} \left(d \|M_{n}(\boldsymbol{m} - \bar{\boldsymbol{z}}_{n})\|_{2}^{2} + (d + \kappa - 1) \|M_{n}C\|_{F}^{2} \right).$$

When $d \gg 1$ this lower-bound is very close to the upper-bound on $\mathbb{E} \|\mathbf{g}\|^2$ in Thm. 6. Thus, under this paper's assumptions, a variance bound cannot be significantly better than an ESN bound.

5.3 The Entropy Term

All discussion in this paper has been for gradient estimators for l, while the goal is of course to optimize l + h. For location-scale families, h is known in closed-form, meaning the exact gradient – or the proximal operator for h – can be computed exactly. Still, it has been observed that if q_w is very close to p(z|x), cancellations mean that estimating the gradient of h + l might have lower variance than the gradient of l alone [12].

With any variational family, it is well-known that the gradient of the entropy can be represented as $-\nabla_{\boldsymbol{w}} \mathbb{E}_{z \sim q_{\boldsymbol{w}}} \log q_{\boldsymbol{v}}(z)|_{\boldsymbol{v}=\boldsymbol{w}}$. That is, the dependence of $\log q_{\boldsymbol{w}}$ on \boldsymbol{w} can be neglected under differentiation. Thus, if one wishes to stochastically estimate the gradient of h, one can treat $\log q_{\boldsymbol{v}}$ in the same way as $\log p$ when calculating gradients. Then, one could apply the analysis in this paper to $f(\boldsymbol{z}) = \log p(\boldsymbol{z}, \boldsymbol{x}) - \log q_{\boldsymbol{v}}(\boldsymbol{z})$ rather than $f(\boldsymbol{z}) = \log p(\boldsymbol{z}, \boldsymbol{x})$ as done above. It is easy to imagine situations where subtracting $\log q_v$ (or a fraction of it) from $\log p$ would change M_n and \bar{z}_n in such a way as to produce a tighter bound. Thus, the bounds in this paper are consistent with practices [5, 12] where using $\log q_v$ as a control variate can reduce gradient variance.

5.4 Smoothness and Convergence Guarantees

At a very high level, convergence rates for stochastic gradient methods require both (1) control of the variability of the gradient estimator and (2) either convexity or Lipschitz smoothness of the objective. This paper is dedicated entirely to the first goal. Independent recent work has addressed at the second issue [3]. The basic summary is that if f(z) is smooth, then l(w) is smooth, and similarly if f(z) is strongly convex. However, full convergence guarantees for black-box VI remain an open research problem.

5.5 Prospects for Generalizing Bounds to Other Variational Families

The bounds given in this paper are closely tied to location-scale families: The exact form of the reparameterization function \mathcal{T}_w is used in Lem. 1 and Lem. 2, which underly the main results of Thm. 3, Thm. 5, and Eq. (4). Thus, extending our proof strategy to other variational families would require deriving new results analogous to Lem. 1 and Lem. 2 for the reparameterization function \mathcal{T}_w corresponding to those new variational families. Moreover, if the exact entropy is not available for a variational family, the analysis must address the variance of the entropy gradient estimator, as discussed in Sec. 5.3.

5.6 Limitations

This work has several limitations. First, it applies only to location-scale families, and requires that the target objective be smooth. Second, if $\log p$ is smooth, it may still be challenging in practice to establish what the smoothness constant is. Third, we observed that even with our strongest condition of matrix smoothness, the some looseness remains in the bounds with the logistic regression examples. Since the ESN bound is unimprovable, this looseness cannot be removed without using more detailed structure of the target $\log p$. It is not obvious what this structure would be, or how it would be obtained for practical black-box inference problems.

References

- David M. Blei, Alp Kucukelbir, and Jon D. McAuliffe. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- [2] Alexander Buchholz, Florian Wenzel, and Stephan Mandt. Quasi-Monte Carlo Variational Inference. In *ICML*, 2018.
- [3] Justin Domke. Provable Smoothness Guarantees for Black-Box Variational Inference. *arXiv:1901.08431 [cs, stat]*, 2019.
- [4] Kai Fan, Ziteng Wang, Jeff Beck, James Kwok, and Katherine Heller. Fast Second-Order Stochastic Backpropagation for Variational Inference. In *NeurIPS*, 2015.
- [5] Tomas Geffner and Justin Domke. Using Large Ensembles of Control Variates for Variational Inference. In *NeurIPS*, 2018.
- [6] Zoubin Ghahramani and Matthew Beal. Propagation Algorithms for Variational Bayesian Learning. In *NeurIPS*, 2001.
- [7] Matthew D. Hoffman, David M. Blei, Chong Wang, and John Paisley. Stochastic Variational Inference. *Journal of Machine Learning Research*, 14:1303–1347, 2013.
- [8] Andrew Miller, Nick Foti, Alexander D' Amour, and Ryan P Adams. Reducing Reparameterization Gradient Variance. In *NeurIPS*, 2017.
- [9] Rajesh Ranganath, Sean Gerrish, and David M. Blei. Black Box Variational Inference. In *AISTATS*, 2014.

- [10] Jeffrey Regier, Kiran Pamnany, Ryan Giordano, Rollin Thomas, David Schlegel, Jon McAuliffe, and Prabhat. Learning an Astronomical Catalog of the Visible Universe through Scalable Bayesian Inference. arXiv:1611.03404 [astro-ph, stat], 2016.
- [11] Jeffrey Regier, Michael I Jordan, and Jon McAuliffe. Fast Black-box Variational Inference through Stochastic Trust-Region Optimization. In *NeurIPS*, page 10, 2017.
- [12] Geoffrey Roeder, Yuhuai Wu, and David K Duvenaud. Sticking the Landing: Simple, Lower-Variance Gradient Estimators for Variational Inference. In *NeurIPS*, 2017.
- [13] Francisco J. R. Ruiz, Michalis K. Titsias, and David M. Blei. The Generalized Reparameterization Gradient. In *NeurIPS*, 2016.
- [14] Francisco J. R. Ruiz, Michalis K. Titsias, and David M. Blei. Overdispersed Black-Box Variational Inference. *arXiv:1603.01140 [stat]*, 2016.
- [15] Tim Salimans and David A. Knowles. Fixed-Form Variational Posterior Approximation through Stochastic Linear Regression. *Bayesian Analysis*, 8(4):837–882, 2013.
- [16] Linda S. L. Tan and David J. Nott. Gaussian variational approximation with sparse precision matrices. *Statistics and Computing*, 28(2):259–275, 2018.
- [17] Michalis Titsias and Miguel Lázaro-gredilla. Doubly Stochastic Variational Bayes for non-Conjugate Inference. In *ICML*, 2014.
- [18] Michalis K. Titsias and Miguel Lázaro-Gredilla. Local Expectation Gradients for Black Box Variational Inference. In *NeurIPS*, 2015.
- [19] David Wingate and Theophane Weber. Automated Variational Inference in Probabilistic Programming. arXiv:1301.1299 [cs, stat], 2013.
- [20] John Winn and Christopher M Bishop. Variational Message Passing. Journal of Machine Learning Research, 6:661–694, 2005.
- [21] Ming Xu, Matias Quiroz, Robert Kohn, and Scott A. Sisson. Variance reduction properties of the reparameterization trick. In *AISTATS*, 2019.
- [22] Yuling Yao, Aki Vehtari, Daniel Simpson, and Andrew Gelman. Yes, but Did It Work?: Evaluating Variational Inference. In *ICML*, 2018.