
Rao-Blackwellized Stochastic Gradients for Discrete Distributions

Runjing Liu¹ Jeffrey Regier² Nilesh Tripuraneni² Michael I. Jordan^{1,2} Jon McAuliffe^{1,3}

Abstract

We wish to compute the gradient of an expectation over a finite or countably infinite sample space having $K \leq \infty$ categories. When K is indeed infinite, or finite but very large, the relevant summation is intractable. Accordingly, various stochastic gradient estimators have been proposed. In this paper, we describe a technique that can be applied to reduce the variance of any such estimator, without changing its bias—in particular, unbiasedness is retained. We show that our technique is an instance of Rao-Blackwellization, and we demonstrate the improvement it yields on a semi-supervised classification problem and a pixel attention task.

1. Introduction

Let z be a discrete random variable over $K \leq \infty$ categories, with distribution $q_\eta(z)$ parameterized by a real vector η and differentiable in η . We aim to minimize

$$\mathcal{L}(\eta) := \mathbb{E}_{z \sim q_\eta(z)} [f_\eta(z)] = \sum_{k=1}^K q_\eta(k) f_\eta(k), \quad (1)$$

where the real-valued integrand f_η also depends differentially on η . If K is finite and small enough, we can compute the exact gradient as

$$\begin{aligned} & \nabla_\eta \mathbb{E}_{q_\eta(z)} [f_\eta(z)] \\ &= \sum_{k=1}^K \left\{ [\nabla_\eta q_\eta(k)] f_\eta(k) + q_\eta(k) \nabla_\eta f_\eta(k) \right\}. \end{aligned} \quad (2)$$

On the other hand, K may be infinite, or large relative to the cost of evaluating $q_\eta \cdot f_\eta$. In either of these cases, which are

¹Department of Statistics, University of California, Berkeley
²Department of Electrical Engineering and Computer Sciences, University of California, Berkeley
³The Voleon Group. Correspondence to: Runjing Liu <runjing_liu@berkeley.edu>.

the focus of this paper, the exact gradient is computationally intractable. Thus, in order to optimize $\mathcal{L}(\eta)$, we seek low-variance stochastic approximations of the gradient.

The “reparametrization trick” (Spall, 2003; Kingma & Welling, 2014) provides efficient stochastic gradients when q_η is a continuous distribution, but it does not apply when z is discrete. Two well-known possibilities in the discrete case are continuous relaxation (Maddison et al., 2017; Jang et al., 2017) and REINFORCE (Williams, 1992) (also known as the score function estimator). The former replaces the discrete random variable with a continuous relaxation so that the reparametrization trick can be applied. However, it results in biased gradient estimates. The latter is impractical for most purposes due to its high variance.

Control variate methodology provides a general framework for variance reduction. Specific examples include RELAX (Grathwohl et al., 2018), REBAR (Tucker et al., 2017), NVIL (Mnih & Gregor, 2014), and MuProp (Gu et al., 2016). These methods provide a mechanism for reducing the variance of REINFORCE, but unfortunately they do not reduce the variance enough for many applications.

In the current paper, we show how to achieve further variance reduction via a meta-procedure that can be applied to any discrete-distribution stochastic-gradient procedure (e.g., REINFORCE or REINFORCE with control variate). Our framework reduces variance without changing the bias. In particular, an unbiased stochastic gradient remains unbiased after application of our approach. Further, our approach is “anytime” in the sense that it can reduce stochastic-gradient variances given any computational budget—the larger the budget, the greater the variance reduction. Hence it is well suited to our chosen setting, where K is infinite or very large, and/or $q_\eta \cdot f_\eta$ is expensive to evaluate.

Our method is particularly apt in the setting where the probability mass $q_\eta(z)$ is concentrated on only a few categories. For example, in extreme classification, only a few labels out of many are plausible. In reinforcement learning, only a few actions in the possible action space are advantageous. Neither control-variate methods nor continuous-relaxation techniques take advantage of this “sparsity,” and we show that the variance reduction of our method in this setting can be dramatic.

We show that our variance-reduction meta-procedure

is an instance of a general statistical method called Rao-Blackwellization (Casella & Robert, 1996). Rao-Blackwellization has been used in previous work to reduce the variance of stochastic gradients (Ranganath et al., 2014; Titsias K & Lázaro-Gredilla, 2015), but in a setting orthogonal to ours, one with multivariate discrete random variables. Our focus here is on a univariate discrete random variable with many categories. Our method can be applied in conjunction with the former work to extend to the case of multivariate discrete random variables, each with a large number of categories. This extension is not discussed in the present work, and we leave it as an avenue of future exploration.

The paper is organized as follows. We present our variance-reduction procedure in Section 2 and make the connection to Rao-Blackwellization in Section 3, demonstrating that our technique necessarily reduces stochastic-gradient variances. In Section 4 we discuss related work. In Section 5, we exhibit the benefits of our procedure on synthetic data, a semi-supervised classification problem, and a pixel attention task. We conclude in Section 6.

2. Method

We consider the situation where the number of categories K is infinite, or very large in the sense that computing the exact gradient in Equation (2) is intractable. One possible stochastic estimator for the gradient is the REINFORCE estimator,

$$f_\eta(z) \nabla_\eta \log q_\eta(z) + \nabla_\eta f_\eta(z) \quad z \sim q_\eta(z), \quad (3)$$

which one can check is unbiased for the true gradient in Equation (2).

In practice, the REINFORCE estimator often has variance too large to be useful. Control variates have been proposed to decrease the variance of the REINFORCE estimator. The key observation is that the score function $\nabla_\eta \log q_\eta(z)$ has zero expectation under $q_\eta(z)$, so

$$[f_\eta(z) - C] \nabla_\eta \log q_\eta(z) + \nabla_\eta f_\eta(z) \quad z \sim q_\eta(z) \quad (4)$$

is still unbiased for the true gradient. Several proposals have been put forth for choosing C to reduce the variance (Mnih & Gregor, 2014; Gu et al., 2016; Tucker et al., 2017).

In this paper, we present a meta-procedure that can be applied to any stochastic estimator for the gradient of a discrete expectation obtained by sampling from $q_\eta(z)$. Let $g(z)$ be any such estimator which is unbiased¹, i.e., satisfies $\mathbb{E}_{q_\eta(z)}[g(z)] = \nabla_\eta \mathbb{E}_{q_\eta(z)}[f_\eta(z)]$. An example is the REINFORCE estimator. We decompose the expectation

¹Our technique applies to biased estimators as well. For concreteness, we focus on the unbiased case.

$\mathbb{E}_{q_\eta(z)}[g(z)]$ into two components: one containing the high-probability atoms of q_η , and one containing the remaining atoms. We compute the exact contribution of the high-probability component to the expectation, and we use a stochastic estimator for the other component. The idea comes from observing that in many applications, $q_\eta(z)$ only puts significant mass on a few categories. If $g(z)$ is reasonably well behaved over z , then $q_\eta(z)g(z)$ is large when $q_\eta(z)$ attains its largest values and smaller elsewhere. By computing the high-probability component of the expectation exactly, we obtain a value already close to correct. A stochastic estimator is then added to correct, on average, for what error remains.

Formally, let \mathcal{C}_k be the set of z such that $q_\eta(z)$ assumes one of its k largest values. Ties may be broken arbitrarily. Let $\bar{\mathcal{C}}_k$ denote the complement of \mathcal{C}_k . Then

$$\nabla_\eta \mathbb{E}_{q_\eta(z)}[f_\eta(z)] = \mathbb{E}_{q_\eta(z)}[g(z)] \quad (5)$$

$$= \mathbb{E}_{q_\eta(z)}[g(z)\mathbf{1}\{z \in \mathcal{C}_k\} + g(z)\mathbf{1}\{z \in \bar{\mathcal{C}}_k\}] \quad (6)$$

$$= \sum_{z \in \mathcal{C}_k} q_\eta(z)g(z) + \mathbb{E}_{q_\eta(z)}[g(z)\mathbf{1}\{z \in \bar{\mathcal{C}}_k\}]. \quad (7)$$

It remains to approximate the expectation in the second term. We use an importance-sampling approximation based on a single draw from an importance distribution. We choose a simple importance distribution: the distribution of q_η conditional on the event $\bar{\mathcal{C}}_k$. We denote this importance distribution by $q_\eta|_{\bar{\mathcal{C}}_k}$. By construction, the importance weighting function is identically equal to $q_\eta(\bar{\mathcal{C}}_k)$, regardless of which $z \sim q_\eta|_{\bar{\mathcal{C}}_k}$ we draw. (Note that the indicator inside the second term of (7) always equals one, because we are only sampling from $z \in \bar{\mathcal{C}}_k$.)

Our estimator assumes that, given k , the set \mathcal{C}_k can be identified at little cost. This certainly holds in the case of inference: using variational Bayes, $q(z)$ is a variational approximate posterior chosen from a set we designate.

In summary, we estimate the gradient as

$$\hat{g}(v) = \sum_{z \in \mathcal{C}_k} q_\eta(z)g(z) + q_\eta(\bar{\mathcal{C}}_k)g(v) \quad (8)$$

$$v \sim q_\eta|_{\bar{\mathcal{C}}_k},$$

which also satisfies $\mathbb{E}_v[\hat{g}(v)] = \nabla_\eta \mathbb{E}_{q_\eta(z)}[f_\eta(z)]$.

We see that the first term of this estimator is deterministic and the second term is random, but scaled by $q_\eta(\bar{\mathcal{C}}_k)$, which is small when q_η is concentrated on a small number of atoms. Therefore, we intuitively expect this estimator to have smaller variance than the original estimator, $g(z)$.

In the next section, we confirm this intuition by interpreting the construction of the estimator $\hat{g}(v)$ as Rao-Blackwellization (which always reduces variance). Hence, we call $\hat{g}(v)$ the *Rao-Blackwellized gradient estimator*.

3. Theory

We begin by describing how a suitable representation of the original discrete variable $z \sim q_\eta(z)$ allows us to interpret our estimator as an instance of Rao-Blackwellization. Let $q_\eta|_{\mathcal{C}_k}$ denote the distribution of q_η conditional on the event \mathcal{C}_k . Consider the three independent random variables

$$u \sim q_\eta|_{\mathcal{C}_k}, \quad (9)$$

$$v \sim q_\eta|_{\bar{\mathcal{C}}_k}, \quad (10)$$

$$\text{and } b \sim \text{Bernoulli}(q_\eta(\bar{\mathcal{C}}_k)). \quad (11)$$

The triplet (u, v, b) provides a distributionally equivalent representation of z :

$$T(u, v, b) \stackrel{d}{=} z, \quad (12)$$

where

$$T(u, v, b) := u^{1-b}v^b. \quad (13)$$

The estimator in Equation (8) can then be written as

$$\hat{g}(v) = \mathbb{E}[g(T(u, v, b))|v], \quad (14)$$

where $g(z)$ is the original unbiased gradient estimator. To see this, break the right-hand side of (14) into two terms according to the value of b , then simplify. Equation (14) demonstrates directly that our estimator is an instance of Rao-Blackwellization. As such, it has the same expectation as the original estimator $g(z)$, a fact about Rao-Blackwellization that follows immediately from iterated expectation. In particular, if $g(z)$ is unbiased as we have assumed, so too is our estimator.

An application of the conditional variance decomposition gives

$$\mathbb{V}[g(z)] = \mathbb{V}[\hat{g}(v)] + \mathbb{E}\{\mathbb{V}[g(T(u, v, b))|v]\}, \quad (15)$$

showing that \hat{g} has lower variance than g , by at least as much as the last term in Equation (15). This too is a standard result about Rao-Blackwellization.

Proposition 1 further quantifies this variance reduction, showing the variance of $\hat{g}(v)$ must be less than the variance of $g(v)$ by the multiplicative factor $q_\eta(\bar{\mathcal{C}}_k)$.

Proposition 1. *Let $g(z)$ be an unbiased gradient estimator as in Equation (5) and $\hat{g}(v)$ denote the Rao-Blackwellized estimator defined in Equation (8). Then*

$$\mathbb{V}[\hat{g}(v)] \leq q_\eta(\bar{\mathcal{C}}_k)\mathbb{V}[g(z)]. \quad (16)$$

Proof. We apply the conditional variance decomposition. Let $\epsilon = q_\eta(\bar{\mathcal{C}}_k)$ and recall the Bernoulli random variable b defined in Equation (11). First,

$$\mathbb{V}[g(z)] = \mathbb{E}[\mathbb{V}[g(z)|b]] + \mathbb{V}[\mathbb{E}[g(z)|b]] \quad (17)$$

$$\geq \mathbb{E}[\mathbb{V}[g(z)|b]] \quad (18)$$

$$= \epsilon \mathbb{V}[g(z)|z \in \bar{\mathcal{C}}_k] + (1 - \epsilon) \mathbb{V}[g(z)|z \in \mathcal{C}_k]$$

$$\geq \epsilon \mathbb{V}[g(z)|z \in \bar{\mathcal{C}}_k].$$

But $\mathbb{V}[\hat{g}(v)] = \epsilon^2 \mathbb{V}[g(z)|z \in \bar{\mathcal{C}}_k]$, which in combination with the above yields the result. \square

The multiplicative factor of variance reduction guaranteed by Rao-Blackwellization can be significant if the probability mass of $q_\eta(z)$ is concentrated on just a few categories. But while Rao-Blackwellization reduces the variance of $g(z)$, this comes at the cost of evaluating $g(z)$ a total $k + 1$ times (cf. Equation (8)). An initial stochastic gradient $g(z)$ such as REINFORCE will only require a single evaluation of g .

There is an alternative approach to reducing the variance of an initial estimator $g(z)$ via multiple evaluations of $g(z)$: minibatching, i.e., simple Monte-Carlo averaging over independent draws of z . Thus, the question arises: given a budget of $N < K$ evaluations of $g(z)$, is it better to Rao-Blackwellize or minibatch? Computationally, our method is parallelizable in the same way that minibatching is parallelizable. The next proposition shows constructively that there is a choice of $k \leq N$ for which Rao-Blackwellization reduces variance at least as much as minibatching.

Proposition 2. *Suppose we have a budget of N evaluations of g . Consider the estimators*

$$\hat{g}_{N,k}(v) := \sum_{u \in \mathcal{C}_k} q_\eta(u)g(u) + \frac{q_\eta(\bar{\mathcal{C}}_k)}{N-k} \sum_{j=1}^{N-k} g(v_j), \quad (19)$$

$$v_1, \dots, v_{N-k} \stackrel{iid}{\sim} q_\eta|_{\bar{\mathcal{C}}_k}$$

and

$$g_N(z) := \frac{1}{N} \sum_{j=1}^N g(z_j), \quad z_1, \dots, z_N \stackrel{iid}{\sim} q_\eta. \quad (20)$$

If we choose

$$\hat{k} = \arg \min_{k \in \{0, \dots, N\}} \frac{q_\eta(\bar{\mathcal{C}}_k)}{N-k} \quad (21)$$

then $\mathbb{V}[\hat{g}_{N,\hat{k}}(v)] \leq \mathbb{V}[g_N(z)]$.

Proof. Let $V_1 = \mathbb{V}[g_1(z)]$. Then $\mathbb{V}[g_N(z)] = (1/N)V_1$, while Proposition 1 shows that $\mathbb{V}[\hat{g}_{N,k}(v)] \leq \frac{q_\eta(\bar{\mathcal{C}}_k)}{N-k} V_1$. Since $\frac{q_\eta(\bar{\mathcal{C}}_k)}{N-k} = \frac{1}{N}$ when $k = 0$, the result follows. \square

Together, Propositions 1 and 2 imply the following:

- Rao-Blackwellization leads to a significant variance reduction if the mass of $q_\eta(z)$ is concentrated.

- Even when restricting minibatched versions of the initial and Rao-Blackwellized estimators to an equal number of evaluations of g , Rao-Blackwellization yields equal or lower variance, for a computable choice of k .

4. Related Work

Methods to reduce the variance of stochastic gradients for discrete distributions generally fall into two broad categories: continuous relaxation methods and control variate methods.

In the first category, the Concrete distribution (Maddison et al., 2017) approximates the discrete random variable with a reparametrizable continuous random variable so that the standard reparametrization trick can be applied. While this continuous relaxation reduces the variance of the stochastic gradient, the resulting estimators are biased. Thus the Gumbel-softmax procedure (Jang et al., 2017) introduced an annealing step into the optimization whereby the continuous relaxation converges towards the discrete random variable as the optimization path moves forward.

In the second category, control variate methods include black-box variational inference (BBVI) (Ranganath et al., 2014), NVIL (Mnih & Gregor, 2014), DARN (Gregor et al., 2014), and MuProp (Gu et al., 2016). BBVI uses multiple samples at each step to estimate the ‘optimal’ control variate. NVIL introduces an observation dependent control variate learned by a separate neural network. DARN uses a Taylor expansion of $f_\eta(z)$ to compute a control variate, but this results in a biased estimator; MuProp proposes an estimate of this bias and corrects it.

Finally, RELAX (Grathwohl et al., 2018) and REBAR (Tucker et al., 2017) are a combination of the two broad methods and use a continuous relaxation to construct a control variate.

Section 5 compares both continuous relaxation and control variate methods to our Rao-Blackwellization.

A Rao-Blackwellization procedure for gradient estimation was also applied in BBVI and “local expectation gradients” (Titsias K & Lázaro-Gredilla, 2015), but for a different purpose. In their setting, the expectation is decomposed over a multivariate (discrete or continuous) random variable using iterated expectation. BBVI approximates each conditional expectation by sampling (with a control variate), while local expectation gradients compute each conditional expectation analytically. This Rao-Blackwellization is orthogonal to our approach: while they consider multiple discrete random variables, our approach focuses on a univariate discrete with many categories.

The process of summing out a few terms and sampling the remainder for gradient estimation has appeared in the con-

text of reinforcement learning (Titsias K, 2014; Liang et al., 2018), though to our knowledge we are the first to make the connection with Rao-Blackwellization. In MAPO (Liang et al., 2018), a procedure to create a memory buffer of trajectories for policy optimization, the terms with high rewards (or small loss) are kept and summed. In contrast, we choose to sum terms with high probability. In our setting, it is the loss $f_\eta(z)$, not the probability, $q_\eta(z)$, that is expensive to evaluate for all categories z .

Finally, the problem of having a large number of categories also manifests in language models, and methods such as noise contrastive estimation (Gutmann & Hyvärinen, 2012) and hierarchical softmax (Morin & Bengio, 2005) have been introduced. However, these methods are applied when the normalizing constant for $q_\eta(z)$ is intractable. In our work, we restrict ourselves to scenarios where $q_\eta(z)$ is normalized.

5. Experiments

In our experiments, we will consider applying the Rao-Blackwellization procedure to either the REINFORCE estimator,

$$g(z) = f_\eta(z) \nabla_\eta \log q_\eta(z) + \nabla_\eta f_\eta(z), \quad (22)$$

$$z \sim q_\eta(z),$$

or REINFORCE with a control variate C ,

$$g(z) = [f_\eta(z) - C] \nabla_\eta \log q_\eta(z) + \nabla_\eta f_\eta(z), \quad (23)$$

$$z \sim q_\eta(z).$$

A simple choice of control variate that works well in practice is to take $C = f_\eta(z')$ for an independent draw $z' \sim q_\eta$. We abbreviate this estimator as REINFORCE⁺.

Note that in both REINFORCE and REINFORCE⁺, $g(z)$ is unbiased for the true gradient. (In the second case, $g(z)$ is unbiased conditional on z' , and hence unconditionally unbiased as well.)

5.1. Bernoulli latent variables

We fix a vector $p = [0.6, 0.51, 0.48]^\top$ and seek to minimize the loss function

$$\mathbb{E}_{b_1, b_2, b_3 \stackrel{i.i.d.}{\sim} \text{Bern}(\sigma(\eta))} \left\{ \sum_{i=1}^3 (b_i - p_i)^2 \right\} \quad (24)$$

over $\eta \in \mathbb{R}$, where $\sigma(\eta)$ is the sigmoid function. Here, the discrete random vector $b = [b_1, b_2, b_3]^\top$ is supported over $K = 2^3 = 8$ categories. The optimal value of $\sigma(\eta)$ is 1, approached as $\eta \rightarrow \infty$.

Figure 1 shows the performance of Rao-Blackwellizing REINFORCE and REINFORCE⁺. We initialized η at $\eta = -4$, so the sampling distribution has large mass at $b = (0, 0, 0)$.

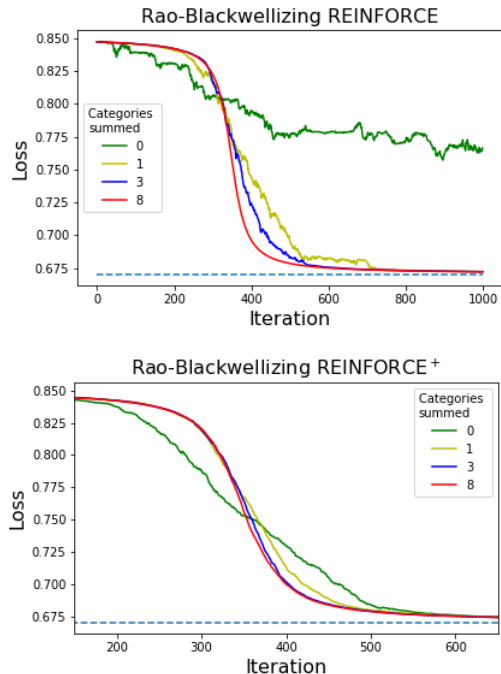


Figure 1. The loss function at each iteration in the Bernoulli experiments. Each line is an average over 20 trials from the same initialization. Zero categories summed is the original estimator, while eight categories summed returns the exact gradient.

The optimal distribution on the other hand should put all mass at $b = (1, 1, 1)$. In other words, we initialized the optimization procedure such that the mass is concentrated on the wrong point. The Rao-Blackwellized gradient is therefore initially slightly slower than the original gradient, since we are analytically summing the wrong category. However, Rao-Blackwellization improves the performance of both gradient estimators at the end of the path.

Figure 2 shows the variances of the gradient estimates at $\eta = 0$ and $\eta = -4$, as a function of k , the categories analytically summed. As expected, the variance decreases as more categories are analytically summed. At $\eta = 0$, the corresponding q_η distribution is uniform, i.e., maximally anti-concentrated, so the variance reduction of Rao-Blackwellization is not large. However, the gains are quite substantial at $\eta = -4$, where q_η is concentrated around the point $b = (0, 0, 0)$. In this case, analytically summing out one category removes nearly all the variance.

5.2. Gaussian mixture model

For our next experiment, we draw $N = 200$ observations (y_n) from a d -dimensional Gaussian mixture model with $K = 10$ components, taking $d = 2$.

$$z_n \stackrel{\text{iid}}{\sim} \text{Categorical}(\pi_{1:K}), n = 1, \dots, N, \quad (25)$$

$$\mu_k \stackrel{\text{iid}}{\sim} \mathcal{N}(0, \sigma_0^2 I_{d \times d}), k = 1, \dots, K, \quad (26)$$

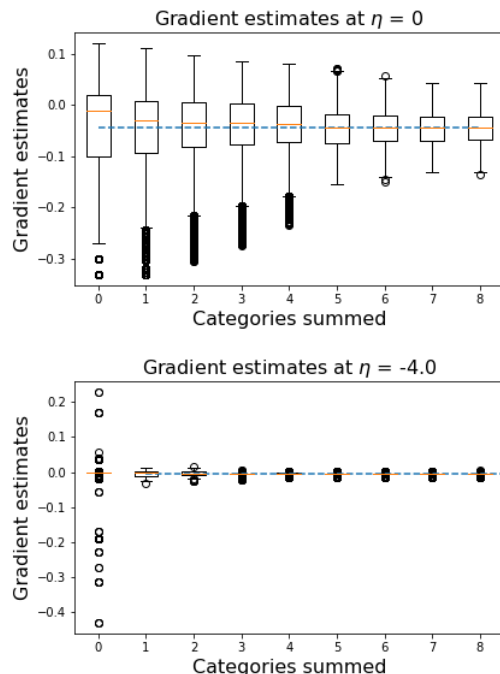


Figure 2. The distribution of gradient estimates from REINFORCE+ in the Bernoulli experiments. We examine the gradients at $\eta = 0$ and $\eta = -4$, as a function of k , the number of categories summed. Summing out categories reduces variance. The reduction is large at $\eta = -4$ where the variational distribution is concentrated on just one category. (Note there is still some random noise when we sum out all 8 categories here, because of the random control variate.)

$$y_n | z_n, \mu \stackrel{\text{iid}}{\sim} \mathcal{N}(\mu_{z_n}, \sigma_y^2 I_{d \times d}), n = 1, \dots, N. \quad (27)$$

Here each μ_k is a Gaussian centroid and each z_n is a cluster membership indicator.

As exact inference of the posterior $p(\mu, z | y)$ is intractable, we approximate it variationally (Blei et al., 2017) with the mean-field family

$$q(\mu, z) = \prod_{k=1}^K q(\mu_k) \prod_{n=1}^N q(z_n). \quad (28)$$

Here

$$q(\mu_k) = \delta\{\mu_k = \hat{\mu}_k\}, \quad (29)$$

$$q(z_n) = \text{Categorical}(\hat{\pi}_n), \quad (30)$$

where $\delta\{\cdot = \hat{\mu}_k\}$ is the Dirac-delta function.

We then seek to minimize $\text{KL}(q(\mu, z) \| p(\mu, z | y))$ over the variational parameters $\hat{\mu}$ and $\hat{\pi}$. This is equivalent to maxi-

mizing the ELBO

$$\sum_{n=1}^N \mathbb{E}_{q(z_n; \pi_n)} \left[\log \frac{p(y_n | \hat{\mu}, z_n) p(z_n)}{q(z_n)} \right] + \sum_{k=1}^K \log p(\hat{\mu}_k). \quad (31)$$

Note that the expectation over z_n is a summation over $K = 10$ categories. Figure 3 compares the performance of unbiased stochastic gradients produced from REINFORCE⁺ to the Rao-Blackwellization of REINFORCE⁺ for optimization of the ELBO in Equation (31).

Unlike the Bernoulli example, we are also optimizing parameters inside the expectation; specifically, in this case we are jointly optimizing the variational mean parameters $\hat{\mu}_k$ alongside the $\hat{\pi}_n$. We expect that more quickly learning the latent categories z_n aids the optimization process, since the mean parameters depend on the cluster memberships.

We initialized the optimization with K -means. Figure 3 shows that Rao-Blackwellization improves the convergence rate, with faster convergence when more categories are summed. With summing just three categories, we nearly recover the same ELBO trajectory of the exact gradient, which sums all ten categories. We chose $K = 10$ as an example so we can compare against the exact gradient; with larger K , computing the exact gradient will become intractable and stochastic methods such as ours will be required.

We also examine here the computational trade-off. Our Rao-Blackwellized estimator with k categories summed requires $k + 1$ evaluations of the original REINFORCE⁺ estimator. For a fairer comparison, we also consider the benefits of variance reduction obtained from simple Monte-Carlo sampling, where $k + 1$ samples of the REINFORCE⁺ estimator are averaged at each iteration. In this experiment, Rao-Blackwellization yields better performance than Monte-Carlo averaging. This is because for most observations, memberships are fairly unambiguous and so $q(z)$ is concentrated. This is the regime where our theory suggests significant variance reduction using Rao-Blackwellization.

5.3. Generative semi-supervised classification

5.3.1. SEMI-SUPERVISED MODELS

The goal of a semi-supervised classification task is to predict labels y from x , but where the training set consists of both labeled data $(x, y) \sim \mathcal{D}_L$ and unlabeled data $x \sim \mathcal{D}_U$. The approach proposed by Kingma et al. (2014) uses a variational autoencoder (VAE) whose latent space is joint over a Gaussian variable z and the discrete label y . The training objective is to learn a classifier $q_\phi(y|x)$, an inference model $q_\phi(z|x, y)$, and a generative model $p_\theta(x|y, z)$. On labeled data, the variational lower bound is

$$\log p_\theta(x, y) \geq \mathcal{L}^L(x, y) \quad (32)$$




Figure 3. Results for Gaussian mixture model experiment. (Top) Simulated data. (Bottom) Solid lines display the negative ELBO per iteration using REINFORCE⁺, for k categories summed. Zero categories summed is the original REINFORCE⁺ estimator, while 10 categories summed returns the analytic gradient. Dashed lines show performance when $n \in \{2, 4\}$ draws of the REINFORCE⁺ estimator are averaged at each iteration to reduce variance. Each line is an average over 20 trials from the same initialization.

$$:= \mathbb{E}_{q_\phi(z|x, y)} [\log p_\theta(x|y, z) + \log p_\theta(z) + \log p_\theta(y) - \log q_\phi(z|x, y)] \quad (33)$$

On unlabeled data, the unknown label y is treated as a latent variable and integrated out,

$$\log p_\theta(x) \geq \mathcal{L}^U(x) \quad (34)$$

$$:= \mathbb{E}_{q_\phi(z|x, y) q_\phi(y|x)} [\log p_\theta(x|y, z) + \log p_\theta(z) + \log p_\theta(y) - \log q_\phi(z|x, y) - \log q_\phi(y|x)] \quad (35)$$

$$= \mathbb{E}_{q_\phi(y|x)} [\mathcal{L}^L(x, y) - \log q_\phi(y|x)] \quad (36)$$

The full objective to be maximized is

$$\mathcal{J} = \mathbb{E}_{x \sim \mathcal{D}_U} [\mathcal{L}^U(x)] + \mathbb{E}_{(x, y) \sim \mathcal{D}_L} [\mathcal{L}^L(x, y)] + \alpha \mathbb{E}_{(x, y) \sim \mathcal{D}_L} [\log q_\phi(y|x)] \quad (37)$$

where the third term is added for the classifier $q_\phi(y|x)$ to also train on labeled data. α is a hyperparameter which we set to 1.0 in our experiments.

We take z to be a continuous random variable with a standard Gaussian prior. Hence, gradients can flow through z using the reparametrization trick. However, y is a discrete label. The original approach proposed by Kingma et al. (2014) computed the expectation in Equation (36) by




Figure 4. Results on the semisupervised MNIST task. Plotted is test set negative ELBO evaluated at the MAP label. Paths are averages over 10 runs from the same initialization. Vertical lines are standard errors. Our method (red) is comparable with summing out all ten categories (black).

exactly summing over the ten categories. However, most images are unambiguous in their classification, so $q_\phi(y|x)$ is often concentrated on just one category. We will show that applying our Rao-Blackwellization procedure with one category summed gives results comparable to computing the the full sum, more quickly.

5.3.2. EXPERIMENTAL RESULTS

We work with the MNIST dataset (Lecun et al., 1998). We used 50 000 MNIST digits in the training set, 10 000 digits in the validation set, and 10 000 digits in the test set. Among the 50 000 MNIST digits in the training set, 5 000 were randomly selected to be labeled, and the remaining 45 000 were unlabeled.

To optimize, we Rao-Blackwellized the REINFORCE estimator. We compared against REINFORCE without Rao-Blackwellization; the exact gradient with all 10 categories summed; REINFORCE⁺; Gumbel-softmax (Jang et al., 2017); NVIL (Mnih & Rezende, 2016); and RELAX (Grathwohl et al., 2018).

For all methods, we used performance on the validation set to choose step-sizes and other parameters. See Appendix for details concerning parameters and model architecture.

Figure 4 shows the negative ELBO, $-\mathcal{L}^L(x, y)$ from Equation (33), on the test set evaluated at the MAP label as a function of epoch. In this experiment, our Rao-Blackwellization with one category summed (RB-REINFORCE) achieves the same convergence rate as the original approach where all ten categories are analytically summed. Moreover, our method achieves comparable test accuracy, at 97%. Finally, our method requires about 18 seconds per epoch, compared to 31 seconds per epoch when using the full sum (Table 1).

In comparing with other approaches, we clearly improve

Table 1. Accuracies and timing results on semi-supervised MNIST classification. Standard errors of test accuracies are over 10 runs of each method. Standard deviations of timing are over the 100 epochs of 10 runs. Training was run on a p3.2xlarge instance on Amazon Web Services.

Method	test acc. (SE)	secs/epoch (SD)
RB-REINFORCE	0.965 (0.001)	17.5 (1.8)
Exact sum	0.966 (0.001)	31.4 (3.2)
REINFORCE	0.940 (0.002)	15.7 (1.6)
REINFORCE ⁺	0.953 (0.001)	17.2 (1.7)
RELAX	0.966 (0.001)	29.8 (3.0)
NVIL	0.956 (0.002)	17.5 (1.8)
Gumbel-softmax	0.954 (0.001)	16.4 (1.7)

upon the convergence rate of REINFORCE. We slightly improve on RELAX. On this example, REINFORCE⁺, NVIL, and Gumbel-softmax also give results comparable to ours.

5.4. Moving MNIST

In this section, we use a hard-attention mechanism (Mnih et al., 2014; Gregor et al., 2015) to model non-centered MNIST digits. We choose this problem because, as will be seen, the exact stochastic gradient is intractable due to the large number of categories. However, only a few of the categories will have significant probabilities.

Like the original VAE work (Kingma & Welling, 2014), we learn an inference model $q_\phi(z|x)$ and generative model $p_\theta(x|z)$, where z is a low-dimensional, continuous representation of the MNIST digit x . Unlike the previous section, we are no longer using the class label. However, we now work with a non-centered MNIST digit, and in order to train the inference and generative models, we must also infer the pixel at which the MNIST digit is centered.

More precisely, our generative model is as follows. For each image, we sample a two-vector representing the pixel at which to center the original 28×28 MNIST image:

$$\ell \sim \text{Categorical}(H \times W). \quad (38)$$

Here H and W are respectively the height and width, in pixels, of the larger image frame on which the MNIST digit will be placed. We take $H = W = 68$ in our experiments.

Next, we generate the non-centered MNIST digit as

$$z \sim \mathcal{N}(0, I_d), \quad (39)$$

$$x_{h,w} | \ell, z \stackrel{\text{ind}}{\sim} \text{Bernoulli}(\mu(z)[h - \ell_0, w - \ell_1]), \quad (40)$$

for $h \in \{0, \dots, H - 1\}$ and $w \in \{0, \dots, W - 1\}$. Here μ is a neural network that maps $z \in \mathbb{R}^d$ to a grid of mean parameters $\mu(z) \in \mathbb{R}^{28 \times 28}$. In Equation (40), we take $\mu(z)[a, b] = 0$ if $(a, b) \notin [0, 28]^2$.




Figure 5. Examples of non-centered MNIST digits

In this way, $x \in \mathbb{R}^{H \times W}$ is a random sample of an image containing a single non-centered MNIST digit on a blank background (Figure 5).

Hence, we need to learn not only the generative model for an MNIST digit, but also the pixel at which the digit is centered. Our two latent variables are z_n and ℓ_n . We find a variational approximation to the posterior using an approximating family of the form

$$\ell_n | x_n \sim \text{Categorical}(\zeta(x_n)), \quad (41)$$

$$z_n | x_n, \ell_n \sim \mathcal{N}(h_\mu(x_n, \ell_n), h_\Sigma(x_n, \ell_n)), \quad (42)$$

where ζ , h_μ , and h_Σ are neural networks. The appendix details the architecture for the neural networks.

REINFORCE was too high variance to be practical here, so we started with REINFORCE⁺ and its Rao-Blackwellization. Here, we chose to sum the top five categories. We again compare with NVIL, Gumbel-softmax, and RELAX. For all the methods, we use a validation set to tune step-sizes and other parameters.

Figure 6 shows the negative ELBO on the test set evaluated at the MAP pixel location as a function of epoch. RELAX converged to a similar ELBO as our method, but did so at a slower rate. While NVIL also converged quickly, it converged to a worse negative ELBO than our method.

Gumbel-softmax did not appear to converge to a reasonable ELBO. We believe that the bias of this procedure was too high in this application. In particular, because we are constrained to sampling discrete values for the pixel attention, we must use the straight-through version of Gumbel-softmax (Bengio et al., 2013; Jang et al., 2017), which suffers from even higher bias.

Our method is more computationally expensive per epoch than the others (Table 2). However, the gains in convergence are still substantive: for example, it takes about 44 seconds for our method to reach a negative ELBO of 500, while it takes RELAX about 110 seconds.

Our method performs best because it is the only one that takes advantage of the fact that only a few digit positions have high probabilities. Summing these positions analytically removes much of the variance.

6. Discussion

Efficient stochastic approximation of the gradient $\nabla_\eta \mathbb{E}_{q_\eta(z)}[f_\eta(z)]$, where z is discrete, is a basic problem that arises in many probabilistic modelling tasks. We




Figure 6. Results on the moving MNIST task. Plotted is test set negative ELBO evaluated at the MAP pixel location. Paths are averages over 10 runs from the same initialization. Vertical lines are standard errors. Our Rao-Blackwellization (red) with summing out the top five categories exhibits the fastest convergence and reaches a smaller negative ELBO than NVIL and REINFORCE⁺.

Table 2. Timing results on the moving MNIST task. Standard deviations of timing are over the 50 epochs of 10 runs. Training was run on a p3.2xlarge instance on Amazon Web Services.

Method	secs/epoch (SD)
RB-REINFORCE ⁺	15.4 (2.3)
REINFORCE ⁺	8.9 (1.3)
RELAX	11.1 (1.6)
NVIL	9.5 (1.4)
Gumbel-softmax	8.7 (1.2)

have presented a general method to reduce the variance of stochastic estimates of this gradient, without changing the bias. Our method is grounded in the classical technique of Rao-Blackwellization. Experiments on synthetic data and two large-scale MNIST modeling problems show the practical benefits of our variance-reduced estimators.

We have focused on the particular setting where z is a univariate discrete random variable, which is relevant for many applications. In other situations, multiple discrete variables will naturally appear in the expectations being optimized. Treating these as a single discrete variable over the Cartesian product of the sample spaces may make such problems amenable to our Rao-Blackwellization approach.

In addition, many multivariate discrete distributions arising in modeling applications will be structured (e.g., the discrete-space latent Markov chain of an HMM). Local expectation gradients (Titsias K & Lázaro-Gredilla, 2015) reduce high-dimensional expectations over these multivariate discrete distributions to iterated univariate expectations through appropriate conditioning on variable sets. Our technique can then be applied for variance reduction in computing the univariate expectations. This is an avenue of future research.

References

- Bengio, Y., Leonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. 2013. URL <https://arxiv.org/abs/1308.3432>.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Casella, G. and Robert, C. P. Rao-Blackwellisation of sampling schemes. *Biometrika*, 83(1):81–94, 1996.
- Grathwohl, W., Choi, D., Wu, Y., Roeder, G., and Duvenaud, D. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. In *International Conference on Learning Representations*, 2018.
- Gregor, K., Mnih, A., and Wierstra, D. Deep autoregressive networks. In *International Conference on Machine Learning*, 2014.
- Gregor, K., Danihelka, I., Graves, A., Rezende, D., and Wierstra, D. DRAW: a recurrent neural network for image generation. In *International Conference on Machine Learning*, 2015.
- Gu, S., Levine, S., Sutskever, I., and Mnih, A. MuProp: Unbiased backpropagation for stochastic neural networks. In *International Conference on Learning Representations*, 2016.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *International Conference on Artificial Intelligence and Statistics*, 2012.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with Gumbel-softmax. In *International Conference on Learning Representations*, 2017.
- Kingma, D. and Welling, M. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014.
- Kingma, D. P. and Ba, J. Adam: a method for stochastic optimization. In *International Conference for Learning Representations*, 2015.
- Kingma, D. P., Rezende, D. J., Mohamed, S., and Welling, M. Semi-supervised learning with deep generative models. *CoRR*, abs/1406.5298, 2014. URL <http://arxiv.org/abs/1406.5298>.
- Lecun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998.
- Liang, C., Norouzi, M., Berant, J., Le, Q., and Lao, N. Memory augmented policy optimization for program synthesis with generalization. In *Neural Information Processing Systems*, 2018.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*, 2017.
- Mnih, A. and Gregor, K. Neural variational inference and learning in belief networks. In *International Conference on Machine Learning*, 2014.
- Mnih, A. and Rezende, D. J. Variational inference for Monte Carlo objectives. In *International Conference on Machine Learning*, 2016.
- Mnih, V., Heess, N., Graves, A., et al. Recurrent models of visual attention. In *Advances in Neural Information Processing Systems*, 2014.
- Morin, F. and Bengio, Y. Hierarchical probabilistic neural network language model. In *International Conference on Artificial Intelligence and Statistics*, 2005.
- Ranganath, R., Gerrish, S., and Blei, D. M. Black box variational inference. In *International Conference on Artificial Intelligence and Statistics*, 2014.
- Royle, J. A. N-mixture models for estimating population size from spatially replicated counts. *Biometrics*, 60(1): 108–115, 2004.
- Spall, J. C. *Introduction to Stochastic Search and Optimization*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 2003.
- Titsias K, M. Combine Monte Carlo with exhaustive search: Effective variational inference and policy gradient reinforcement learning. In *NIPS Workshop: Advances in Approximate Inference*, 2014.
- Titsias K, M. and Lázaro-Gredilla, M. Local expectation gradients for black box variational inference. In *Neural Information Processing Systems*, 2015.
- Tucker, G., Mnih, A., Maddison, C. J., Lawson, J., and Sohl-Dickstein, J. REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Neural Information Processing Systems*, 2017.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.