
Vprop: Variational Inference using RMSprop

Mohammad Emtiyaz Khan
RIKEN, Tokyo, Japan
emtiyaz.khan@riken.jp

Zuozhu Liu*
SUTD, Singapore
zuozhu_liu@mymail.sutd.edu.sg

Voot Tangkaratt
RIKEN, Tokyo, Japan
voot.tangkaratt@riken.jp

Yarin Gal
University of Oxford, UK
yarin.gal@cs.ox.ac.uk

Abstract

Many computationally-efficient methods for Bayesian deep learning rely on continuous optimization algorithms, but the implementation of these methods requires significant changes to existing code-bases. In this paper, we propose Vprop, a method for variational inference that can be implemented with two minor changes to the off-the-shelf RMSprop optimizer. Vprop also reduces the memory requirements of Black-Box Variational Inference by half. We derive Vprop using the conjugate-computation variational inference method, and establish its connections to Newton’s method, natural-gradient methods, and extended Kalman filters. Overall, this paper presents Vprop as a principled, computationally-efficient, and easy-to-implement method for Bayesian deep learning.

1 Introduction

Existing approaches for variational inference (VI), such as Black-Box Variational Inference (BBVI) (Ranganath et al., 2014), exploit stochastic-gradient methods to obtain simple and general implementations. Such approaches are widely applicable, but their implementation often requires significant changes to existing code-bases. For example, to implement BBVI for Bayesian neural networks, parameters have to be replaced with random variables, and the optimization objective is changed to the variational lower bound. In this paper we propose a method for Gaussian variational approximations which simplifies the above by exploiting a connection between VI and modern optimization literature. We are able to implement variational inference by making two minor changes to the off-the-shelf RMSprop optimizer. A summary is given in Figure 1.

Our approach enables the VI implementation to lie entirely within the optimization procedure, and allows a plug-and-play of deterministic models such as neural networks. By simply running the existing code-base of an optimizer, we surprisingly recover the optimum of the variational lower bound. Apart from this marvellous connection between variational inference and modern optimization literature, this view also reduces the memory requirement of BBVI by half. This work provides a new software design paradigm to the field of Bayesian deep learning, and extends on recent ideas for efficient Bayesian approximations using deep-learning methodologies (Gal, 2015, 2016; Mandt et al., 2017). For the latter part, we also establish connections of our method to Newton’s method, natural-gradient methods, and extended Kalman filters.

*Work done during an internship in RIKEN.

2 Optimization Algorithms for Variational Inference

Continuous optimization algorithms are extremely popular in machine learning, e.g., given a supervised-learning problem with output vector \mathbf{y} and input matrix \mathbf{X} , we can estimate model parameters $\boldsymbol{\theta}$ by minimizing the negative log-likelihood: $f(\boldsymbol{\theta}) := -\log p(\mathbf{y}|\mathbf{X}, \boldsymbol{\theta})$. This is the maximum-likelihood (ML) estimation, a popular method to fit complex models such as deep neural networks. This procedure scales well to large data and complex models, and the success of these models is partly due to the existence of efficient implementations of optimization methods such as RMSprop (Tieleman & Hinton, 2012), AdaGrad (Duchi et al., 2011), and Adam (Kingma & Ba, 2014).

Variational inference (VI) methods hope to exploit optimization methods to approximate the posterior distribution $p(\boldsymbol{\theta}|\mathbf{y}, \mathbf{X})$, which often involves cumbersome integration. The problem of integration is fundamentally more difficult than finding a point estimate, and the key idea in VI is to convert the integration problem to an optimization problem. A common approach is to approximate the normalizing constant $p(\mathbf{y}|\mathbf{X})$ of the posterior by finding an approximate distribution $q(\boldsymbol{\theta})$ that maximizes a lower bound to it. For example, if we assume a Gaussian prior $p(\boldsymbol{\theta}) := \mathcal{N}(\boldsymbol{\theta}|\mathbf{0}, \mathbf{I}/\lambda)$ with $\lambda > 0$, we can approximate the posterior by a Gaussian distribution, $q(\boldsymbol{\theta}) := \mathcal{N}(\boldsymbol{\theta}|\boldsymbol{\mu}, \text{diag}(\boldsymbol{\sigma}^2))$ with mean $\boldsymbol{\mu}$ and variance $\boldsymbol{\sigma}^2$. We do so by solving the following optimization problem:

$$\log p(\mathbf{y}|\mathbf{X}) \geq \max_{\boldsymbol{\mu}, \boldsymbol{\sigma}} \mathbb{E}_q [\log p(\mathbf{y}|\mathbf{X}, \boldsymbol{\theta}) + \log p(\boldsymbol{\theta}) - \log q(\boldsymbol{\theta})] := \mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2). \quad (1)$$

Optimization algorithms can now be applied to solve this problem, and their efficiency can be exploited to perform approximate Bayesian inference.

Despite this reformulation, the implementation of an optimization algorithm for VI differs significantly from those used for ML estimation. For example, consider the Black-Box Variational Inference (BBVI) method (Ranganath et al., 2014) which is one of the simplest approaches to optimize \mathcal{L} . BBVI employs the following simple stochastic-gradient update:

$$\text{BBVI : } \quad \boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t + \rho_t \left[\widehat{\nabla}_{\boldsymbol{\mu}} \mathcal{L}_t \right], \quad \boldsymbol{\sigma}_{t+1} = \boldsymbol{\sigma}_t + \rho_t \left[\widehat{\nabla}_{\boldsymbol{\sigma}} \mathcal{L}_t \right], \quad (2)$$

where $\rho_t > 0$ is a step size at iteration t , $\widehat{\nabla}$ denotes an unbiased stochastic-gradient estimate, and $\nabla \mathcal{L}_t$ denotes the gradient at the value of the iterate at iteration t . These updates are simple and general, but their implementation differs significantly from adaptive-gradient methods (e.g., RMSprop; see Figure 1(a) for a pseudo-code).

A major challenge with BBVI is the large number of parameters which need to be optimized. Compared to the ML estimate, BBVI doubles the number of parameters (in the best case; when a full covariance matrix is used this becomes quadratic in the original problem size). This is because BBVI optimizes not only a single point estimate, but the parameters of an entire distribution. If an adaptive optimization scheme (like RMSprop) were to be used for BBVI to adapt step-sizes, the memory requirements would have been doubled yet again in order to maintain a step-size for each of the distribution's parameters, i.e., two scaling vectors for $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ respectively, instead of just one vector \mathbf{s} in standard RMSprop as shown in Figure 1(a). This can become prohibitively expensive. Further, adaptation is also tricky since $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are two fundamentally different quantities with different units and their step-sizes might require different types of tuning. A final difference is that BBVI requires gradients with respect to $\boldsymbol{\sigma}$ which typically demands a different implementation than the naive one to avoid numerical issues.

In the next section we present perhaps surprising results, demonstrating a formulation of fully-factorized VI whose implementation is almost identical to that of the adaptive step-size RMSprop. Due to this similarity, we call our algorithm Vprop. As shown in Fig. 1, Vprop can be implemented with two minor changes to the implementation of RMSprop. We derive Vprop using a natural-gradient method of Khan & Lin (2017) called the conjugate-computation variational inference (CVI). Natural-gradient methods are preferable when optimizing parameters of a distribution (Hoffman et al., 2013) and our method has these desired theoretical properties as well. We establish additional connections to Newton's method, natural-gradient methods for continuous optimization, and extended Kalman filtering for approximate Bayesian inference. These connections make Vprop a principled approach for VI which is not only computationally-efficient but is also easy to implement. Empirical results on logistic regression and deep neural networks show that Vprop can indeed perform as well as BBVI but is much simpler to implement.

<pre> 1: $\theta \leftarrow \mu$ 2: $\mathbf{g} \leftarrow \nabla_{\theta} f(\theta)$ 3: $\mathbf{s} \leftarrow (1 - \beta)\mathbf{s} + \beta(\mathbf{g} * \mathbf{g})$ 4: $\mu \leftarrow \mu - \alpha(\mathbf{g} / \sqrt{\mathbf{s} + \delta})$ </pre> <p>(a) RMSprop update at $\theta = \mu$ to find the maximum-likelihood estimate by minimizing $f(\theta) := -\log p(\mathbf{y} \mathbf{X}, \theta)$</p>	<pre> 1: $\theta \leftarrow \mu + \epsilon / \sqrt{\mathbf{s} + \lambda}$ where $\epsilon \sim \mathcal{N}(\epsilon 0, \mathbf{I})$ 2: $\mathbf{g} \leftarrow \nabla_{\theta} f(\theta)$ 3: $\mathbf{s} \leftarrow (1 - \beta)\mathbf{s} + \beta(\mathbf{g} * \mathbf{g})$ 4: $\mu \leftarrow \mu - \alpha((\mathbf{g} + \lambda\mu) / (\mathbf{s} + \lambda))$ </pre> <p>(b) Vprop-1 update to find a Gaussian variational-distribution $q(\theta)$ that minimizes variational lower-bound (mean of q is at μ and its variance is equal to $1 / (\mathbf{s} + \lambda)$).</p>
--	---

Figure 1: This figure compares the pseudo-code of RMSprop (left) and Vprop, our RMSprop variant used for variational inference (right). Differences between the two are highlighted in red. RMSprop optimizes the log-likelihood to find a maximum-likelihood estimate. For RMSprop, μ is the current parameter vector, \mathbf{s} is the scaling vector, $\delta > 0$ is a small constant, and α and β are step-sizes. On the other hand, Vprop optimizes the variational lower bound (1) to estimate a Gaussian variational-distribution with mean μ and variance σ^2 . Here, λ is the prior precision and the variance σ^2 is obtained by setting $\sigma^2 \leftarrow 1 / (\mathbf{s} + \lambda)$. The code of Vprop differs from RMSprop in lines 1 and 4 (highlighted in red). In line 1 in Vprop we add noise to the current parameter μ which is equivalent to sampling μ from the variational distribution. In line 4, a term $\lambda\mu$ is added to the gradient, the scaling term is *not* raised to the power 1/2, and δ is set to be equal to λ . The code in the right is called Vprop-1 since it uses only one random sample. With just two lines of change in RMSprop, Vprop performs variational inference.

3 From Variational Inference to an RMSprop Variant

We will derive Vprop using the conjugate-computation variational inference (CVI) method proposed by Khan & Lin (2017). CVI is a natural-gradient method for VI and when applied to (1), results in the following update (see Appendix A for a proof):

$$\text{CVI: } \mu_{t+1} = \mu_t + \beta_t \sigma_{t+1}^2 \circ [\nabla_{\mu} \mathcal{L}], \quad \sigma_{t+1}^{-2} = \sigma_t^{-2} - 2\beta_t [\nabla_{\sigma^2} \mathcal{L}], \quad (3)$$

where \circ denotes element-wise multiplication of two vectors. These updates are natural-gradient updates and differ from BBVI update of (2) in two main aspects. First, these updates use gradients with respect to the variance σ^2 to update the precision σ^{-2} , while BBVI uses the gradient w.r.t. σ to update the standard-deviation σ . Second, the update for μ is an adaptive update because the step-size β_t is scaled by the variance. As we show next, these two differences enable the implementation of CVI using an RMSprop variant, which is not possible for BBVI.

Vprop can be derived from CVI in two steps. First, we use Bonnet’s and Price’s theorem (Rezende et al., 2014) to express the gradients with respect to μ and σ^2 in terms of gradient and Hessian of f respectively. Specifically, we use the following two identities (Opper & Archambeau, 2009):

$$\nabla_{\mu} \mathbb{E}_q [f(\theta)] = \mathbb{E}_q [\nabla_{\theta} f(\theta)], \quad \nabla_{\sigma^2} \mathbb{E}_q [f(\theta)] = \frac{1}{2} \mathbb{E}_q [\text{diag}(\nabla_{\theta\theta}^2 f(\theta))] \quad (4)$$

where $\text{diag}(\mathbf{A})$ extracts the diagonal of \mathbf{A} . Using these, we can rewrite the gradients as follows:

$$\nabla_{\mu} \mathcal{L} = \nabla_{\mu} \mathbb{E}_q [-f(\theta) + \log p(\theta) - \log q(\theta)] = -\mathbb{E}_q [\nabla_{\theta} f(\theta)] - \lambda\mu, \quad (5)$$

$$\nabla_{\sigma^2} \mathcal{L} = -\frac{1}{2} \mathbb{E}_q \{ \text{diag} [\nabla_{\theta\theta}^2 f(\theta)] \} - \frac{1}{2} \lambda \mathbf{1} + \frac{1}{2} \sigma^{-2}, \quad (6)$$

where $\mathbf{1}$ is a vector of ones and σ^{-2} denotes element-wise inverse square of the elements of the vector σ . Using the above, we can rewrite the CVI updates as the following:

$$\mu_{t+1} = \mu_t - \beta_t \sigma_{t+1}^2 \circ \{ \mathbb{E}_{q_t} [\nabla_{\theta} f(\theta)] + \lambda\mu_t \}, \quad (7)$$

$$\sigma_{t+1}^{-2} = (1 - \beta_t) \sigma_t^{-2} + \beta_t \{ \mathbb{E}_{q_t} [\text{diag} (\nabla_{\theta\theta}^2 f(\theta))] + \lambda \mathbf{1} \}, \quad (8)$$

where $q_t := \mathcal{N}(\theta|\mu_t, \sigma_t^2)$ is the variational distribution at iteration t .

These updates require computation of a Hessian which might be computationally difficult. The second step in Vprop derivation is to replace the Hessian by a Gauss-Newton approximation (Bertsekas, 1999). This approximation is also numerically useful when the Hessian is not positive semi-definite, which is typically the case when f is parameterized by a neural network. Using the Gauss-Newton approximation, we can simplify the precision update to the following:

$$\sigma_{t+1}^{-2} = (1 - \beta_t) \sigma_t^{-2} + \beta_t \left[\mathbb{E}_{q_t} [(\nabla_{\theta} f(\theta))^2] + \lambda \mathbf{1} \right]. \quad (9)$$

By defining $\mathbf{s}_t := \sigma_t^{-2} - \lambda \mathbf{1}$, we can rewrite the update as follows, which we call Vprop:

$$\text{Vprop: } \boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t - \beta_t (\mathbf{s}_{t+1} + \lambda \mathbf{1})^{-1} \circ \{\mathbb{E}_{q_t} [\nabla_{\theta} f(\boldsymbol{\theta})] + \lambda \boldsymbol{\mu}_t\}, \quad (10)$$

$$\mathbf{s}_{t+1} = (1 - \beta_t) \mathbf{s}_t + \beta_t \mathbb{E}_{q_t} [(\nabla_{\theta} f(\boldsymbol{\theta}))^2]. \quad (11)$$

Approximating the expectation using one sample $\boldsymbol{\theta}_t \sim \mathcal{N}(\boldsymbol{\theta} | \boldsymbol{\mu}_t, 1/(\mathbf{s}_t + \lambda \mathbf{1}))$, we get the variant of Vprop we call Vprop-1 (also shown in Fig. 1):

$$\text{Vprop-1: } \boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t - \beta_t (\mathbf{s}_{t+1} + \lambda \mathbf{1})^{-1} \circ [\nabla_{\theta} f(\boldsymbol{\theta}_t) + \lambda \boldsymbol{\mu}_t], \quad (12)$$

$$\mathbf{s}_{t+1} = (1 - \beta_t) \mathbf{s}_t + \beta_t [\nabla_{\theta} f(\boldsymbol{\theta}_t)]^2. \quad (13)$$

We can see the similarity to RMSprop by comparing its update to Vprop-1 as shown in Fig. 1. Both algorithms use a running sum of the square of the gradient to compute the scaling vector \mathbf{s}_t . The algorithms differ in only two lines of code with three major differences. First, Vprop uses samples from q to compute the gradient. This enables a local-exploration around the mean which is useful for uncertainty computation. Second, RMSprop raises the scaling vector to a power of $1/2$, while Vprop does not. Third, Vprop adds the term $\lambda \boldsymbol{\mu}_t$ to the gradient in the mean update. These three differences result in the surprising conversion of RMSprop, which maximizes the log-likelihood, to Vprop, which maximizes the variational lower bound.

We can also derive a deterministic version of Vprop called Vprop-0 which is a bit more similar to RMSprop than Vprop-1. In Vprop-0, instead of using MC samples, we approximate the expectation using a first-order delta approximation $\mathbb{E}_q [\mathbf{g}(\boldsymbol{\theta})] \approx \mathbf{g}(\boldsymbol{\mu})$ and $\mathbb{E}_q [(\mathbf{g}(\boldsymbol{\theta}))^2] \approx \mathbf{g}(\boldsymbol{\mu})^2$. This gives us the following update:

$$\text{Vprop-0: } \boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t - \beta_t (\mathbf{s}_{t+1} + \lambda \mathbf{1})^{-1} \circ [\nabla_{\theta} f(\boldsymbol{\mu}_t) + \lambda \boldsymbol{\mu}_t], \quad (14)$$

$$\mathbf{s}_{t+1} = (1 - \beta_t) \mathbf{s}_t + \beta_t [\nabla_{\theta} f(\boldsymbol{\mu}_t)]^2 \quad (15)$$

with $\mathbf{g}(\boldsymbol{\mu}_t)$ the gradient of f at $\boldsymbol{\mu}_t$. Our empirical results show that Vprop-0 performs worse than Vprop-1, establishing the importance of local exploration obtained by using samples from q .

4 Connections to Newton's Method and Natural-Gradient Methods

In this section, we consider extensions to non mean-field variational distribution, i.e., when the covariance $\boldsymbol{\Sigma}$ of $q(\boldsymbol{\theta})$ is not a diagonal matrix but a full matrix. For this case, we show that our algorithm is a second-order method and is related to an online version of Newton's method. By making a Gauss-Newton approximation to the Hessian, we establish connections to online natural-gradient method and extended Kalman filtering (EKF) method described in Ollivier (2017). The results presented in this section connect methods from three fields: variational inference, continuous optimization, and approximate Bayesian filtering.

The results derived in this section are similar to another work by Khan et al. (2017) who derive Newton-type methods for general purpose optimization. Our derivations are similar to theirs, but our results are about variational inference which is a different problem than the one considered in Khan et al. (2017).

The CVI algorithm for the fully-correlated case can be derived in a similar way to the derivation given in Appendix A. The resulting updates are shown below:

$$\text{CVI-Full: } \boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t + \beta_t \boldsymbol{\Sigma}_{t+1} [\nabla_{\boldsymbol{\mu}} \mathcal{L}], \quad \boldsymbol{\Sigma}_{t+1} = \boldsymbol{\Sigma}_t - 2\beta_t [\nabla_{\boldsymbol{\Sigma}} \mathcal{L}], \quad (16)$$

Using the identities given in (4), gradient expressions similar to (5) and (6), and one MC sample approximation, we can rewrite the updates in terms of the gradient and Hessian of $f(\boldsymbol{\theta})$ as shown below:

$$\text{VON-1: } \boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t - \beta_t (\mathbf{S}_{t+1} + \lambda \mathbf{I})^{-1} (\nabla_{\theta} f(\boldsymbol{\theta}_t) + \lambda \boldsymbol{\mu}_t), \quad (17)$$

$$\mathbf{S}_{t+1} = (1 - \beta_t) \mathbf{S}_t + \beta_t \nabla_{\theta\theta}^2 f(\boldsymbol{\theta}_t), \quad (18)$$

where $\mathbf{S}_t = \boldsymbol{\Sigma}_t^{-1} - \lambda \mathbf{I}$ is the scaling matrix and $\boldsymbol{\theta}_t$ is a sample from $q_t := \mathcal{N}(\boldsymbol{\theta} | \boldsymbol{\mu}_t, (\mathbf{S}_t + \lambda \mathbf{I})^{-1})$. We refer to this update as the Variational Online-Newton (VON) method because it resembles an online version of Newton's method where the scaling matrix is estimated online (the number 1 indicates

that expectations are approximated with one MC sample). We can clearly see this resemblance by comparing VON to the update for Newton’s method:

$$\text{Newton's method: } \boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \rho_t [\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 f(\boldsymbol{\theta}_t)]^{-1} [\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}_t)]. \quad (19)$$

In VON, the scaling matrix is a moving-average of the past Hessians and each Hessian is evaluated at a sample from $q(\boldsymbol{\theta})$. The scaling matrix maintains an online estimate of the past curvature information, making the update an online second-order method.

Like the Vprop derivation discussed in the previous section, if we use a Gauss-Newton approximation for the Hessian, the resulting updates are similar to the online natural-gradient descent:

$$\text{VONG: } \boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t - \beta_t (\mathbf{S}_{t+1} + \lambda \mathbf{I})^{-1} (\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}_t) + \lambda \boldsymbol{\mu}_t), \quad (20)$$

$$\mathbf{S}_{t+1} = (1 - \beta_t) \mathbf{S}_t + \beta_t \nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}_t) [\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}_t)]^T. \quad (21)$$

Due to this similarity, we call it the Variational Online Natural-Gradient (VONG) algorithm.

The VONG update is very similar to the regularized online natural-gradient step discussed in Proposition 4 in Ollivier (2017). There, the Gaussian prior over $\boldsymbol{\theta}$ is used as a Bayesian regularizer of the Fisher matrix. In VONG, the regularization natural arises as a result of approximate inference in a Bayesian model. Ollivier (2017) also show that their online natural-gradient algorithm is equivalent to extended Kalman filters (EKF). Therefore, VONG is also closely related to EKF.

VONG differs from these methods in two aspects. The first difference is that the gradients are evaluated at the samples from $q(\boldsymbol{\theta})$ instead of the current parameter. This enables a local exploration around the current parameter values which is expected to improve the performance by avoiding some local minima (see an example of such local-minima avoidance in Khan et al. (2017)). This difference shows the potential benefits obtained when using an approximate Bayesian method instead of a point estimate method.

The second difference is that VONG uses an empirical estimate of the Fisher matrix obtained by using the observed data \mathbf{y} . The method of Ollivier (2017), on the other hand, approximates the Fisher matrix by using an average over $p(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta})$. Traditionally, natural-gradient methods have relied on empirical estimates (Amari, 1998), but recent studies, such as Pascanu & Bengio (2013), have shown that using averaging leads to an unbiased estimate and performs better. Vprop updates are amenable to such modifications, although it is not clear if this is a valid step to perform variational inference.

The above similarities and differences are useful in connecting the two distinct fields of VI and continuous optimization. VI methods optimize in the space of variational distribution q and are fundamentally different from continuous optimization methods which optimize in the space of parameter $\boldsymbol{\theta}$. Yet, it is possible to design a VI methods as a modification of an existing optimization method, as we show in this paper. A few other recent works have also shown similar types of results (Gal, 2015, 2016; Mandt et al., 2017). Such results are very encouraging and further promote the use of optimization methods as a tool to design scalable algorithms for Bayesian deep learning.

5 Experimental Results

In this section, we present results to establish that Vprop gives comparable performance to existing VI methods. We show results logistic regression and multi-layer perceptron in Figure 2 and 3 respectively. Our results show that Vprop performs as well as CVI despite using the Gauss-Newton approximation. We also show that RMSprop overfits on these datasets, perhaps due to small data-size. Vprop-0 also performs badly but slightly better than RMSprop. We believe that the worse performance is because Vprop-0 does not use samples from q which might lead to overfitting. The slightly better performance of Vprop-0 compared to RMSprop is perhaps because it does not use the square root of the scaling vector.

6 Discussion and Future Works

We proposed Vprop, a Gaussian VI method that can be implemented with two minor changes to the off-the-shelf RMSprop optimizer. The memory requirement of Vprop is half of that required

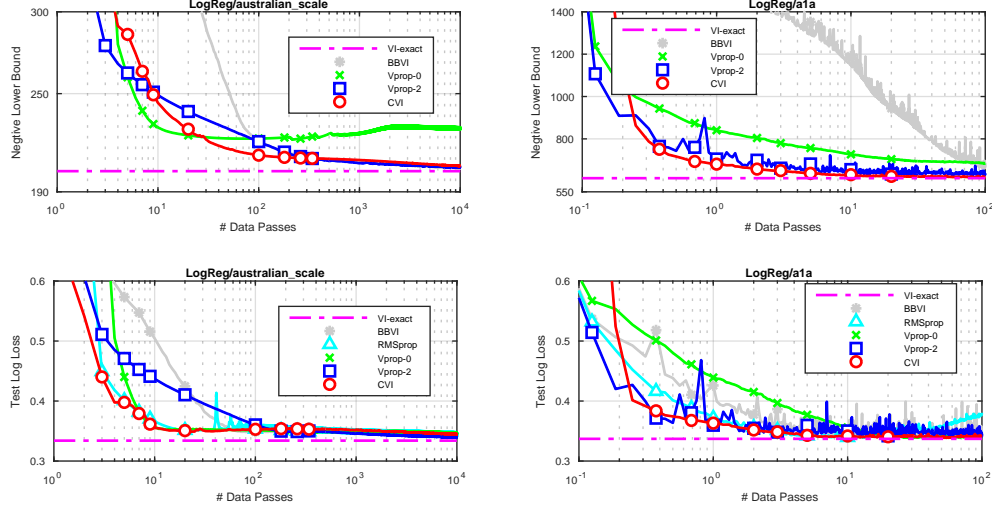


Figure 2: Results on logistic regression. Left column shows results for the Australian-Scale dataset ($N = 345, D = 15, \lambda = 10^{-5}$) where we plot ELBO on training data and log-loss on test data, respectively, versus number of data passes. Right column shows the same for the ‘a1a’ dataset ($N = 1605, D = 123, \lambda = 2.8072$). ‘VI-exact’ is the ground truth obtained by using LBFGS. BBVI is the update (2) with constant step-sizes. We also compare to ‘CVI’ using update (7)-(8) with 10 MC samples and exact Hessian computation. For our method, we use ‘Vprop-2’ using update (12)-(13) with 2 MC samples, and ‘Vprop-0’ implementing update (14)-(15). We see that they all converge either faster than BBVI or at the same rate, while enabling much simpler implementation. CVI uses exact Hessian computation, while Vprop-2 does not and still performs very similar. We also show the log-loss performance of the pure RMSprop method outlined in Fig. 1(a). This method does not optimize ELBO or compute uncertainty, but performs well.

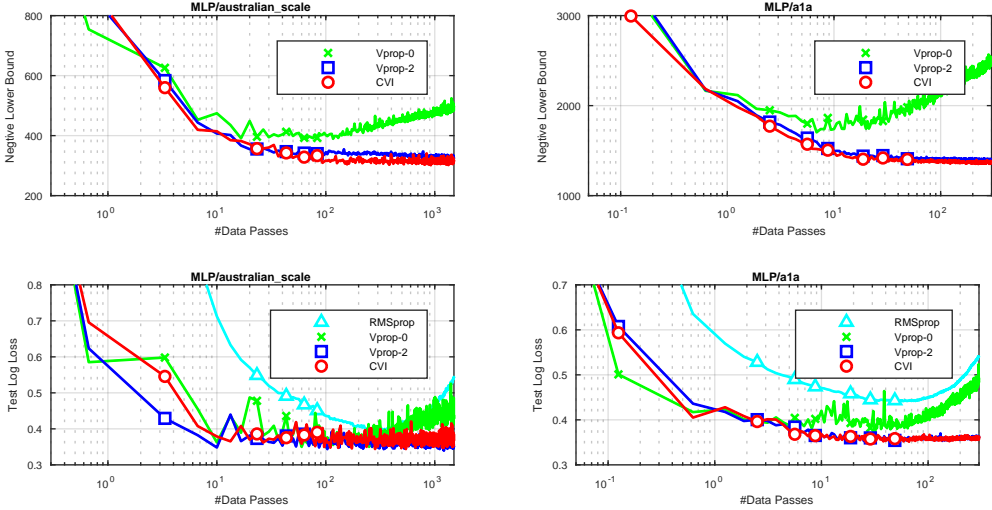


Figure 3: Same comparison as Figure 2 but by using Multi-Layer Perceptron (MLP) with two hidden layers, 10 units each ($\lambda = 1$). We do not have the ground truth for MLP since computing exact ELBO is difficult. For MLP, Vprop-2 performs almost the same as CVI which is expected since, for nonconvex objectives, Gauss-Newton is usually a good numerical approximation. Both Vprop-0 and RMSprop start overfitting after some iterations, but methods that use MC sampling, i.e., Vprop-2 and CVI, do not. We conjecture that this is because the MC sampling gives unbiased gradients that optimize ELBO, while the other two methods do not do so and overfit.

by BBVI. In addition, Vprop is an approximate natural-gradient VI method and inherits many good theoretical properties of natural-gradient methods. We show that Vprop is related to online versions of Newton’s method and the natural-gradient method, and also to extended Kalman filters. Vprop is a principled and computationally-efficient approach for VI, and is also an easy-to-implement method.

We have provided experimental evidence on small models and datasets. In our experiments, Vprop beats BBVI with other adaptive-gradient methods (not presented in the plots), but further experiments are required to confirm this. In the future, we plan to do extensive comparisons on larger problems and compare Vprop to many other existing methods. We also hope to compare to a version of Vprop with momentum and to RMSprop with momentum. We also plan to compare Vprop to other existing methods such as Bayesian Dropout and other black-box methods for VI.

Acknowledgement: We thank Wu Lin (RIKEN) and Didrik Nielsen (RIKEN) for useful discussions. We also thank anonymous reviewers for their useful feedback.

References

- Shun-Ichi Amari. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276, 1998.
- Dimitri P Bertsekas. *Nonlinear programming*. Athena Scientific, 1999.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.
- Yarin Gal. Rapid prototyping of probabilistic models: Emerging challenges in variational inference. In *Advances in Approximate Bayesian Inference workshop, NIPS*, 2015.
- Yarin Gal. *Uncertainty in Deep Learning*. PhD thesis, University of Cambridge, 2016.
- Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- Mohammad Emtiyaz Khan and Wu Lin. Conjugate-computation variational inference: Converting variational inference in non-conjugate models to inferences in conjugate models. *arXiv preprint arXiv:1703.04265*, 2017.
- Mohammad Emtiyaz Khan, Wu Lin, Voot Tangkaratt, Zuozhu Liu, and Didrik Nielsen. Variational Adaptive-Newton Method for Explorative Learning. *ArXiv e-prints*, November 2017.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Stephan Mandt, Matthew D Hoffman, and David M Blei. Stochastic gradient descent as approximate bayesian inference. *arXiv preprint arXiv:1704.04289*, 2017.
- Yann Ollivier. Online natural gradient as a kalman filter, 2017.
- M. Opper and C. Archambeau. The Variational Gaussian Approximation Revisited. *Neural Computation*, 21(3):786–792, 2009.
- Razvan Pascanu and Yoshua Bengio. Revisiting natural gradient for deep networks. *arXiv preprint arXiv:1301.3584*, 2013.
- Rajesh Ranganath, Sean Gerrish, and David M Blei. Black box variational inference. In *International conference on Artificial Intelligence and Statistics*, pp. 814–822, 2014.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-RMSprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning 4*, 2012.

A Derivation of CVI for Gaussian variational distribution

Denote the mean parameters of $q_t(\boldsymbol{\theta})$ by $\mathbf{m}_t := \{\boldsymbol{\mu}, \boldsymbol{\mu}^2 + \boldsymbol{\sigma}_t^2\}$. The mean parameter is equal to the expected value of the sufficient statistics $\boldsymbol{\phi}(\boldsymbol{\theta}) := \{\boldsymbol{\theta}, \boldsymbol{\theta}^2\}$, i.e., $\mathbf{m}_t := \mathbb{E}_{q_t}[\boldsymbol{\phi}(\boldsymbol{\theta})]$. The mirror descent update at iteration t is given by the solution to

$$\mathbf{m}_{t+1} = \underset{\mathbf{m}}{\operatorname{argmax}} \left\langle \mathbf{m}, \widehat{\nabla}_m \mathcal{L}_t \right\rangle - \frac{1}{\beta_t} \mathbb{D}_{KL}[q \parallel q_t] \quad (22)$$

$$= \underset{\mathbf{m}}{\operatorname{argmin}} \left\langle \mathbf{m}, -\widehat{\nabla}_m \mathcal{L}_t \right\rangle + \frac{1}{\beta_t} \mathbb{D}_{KL}[q \parallel q_t] \quad (23)$$

$$= \underset{\mathbf{m}}{\operatorname{argmin}} \mathbb{E}_q \left[\left\langle \boldsymbol{\phi}(\boldsymbol{\theta}), -\widehat{\nabla}_m \mathcal{L}_t \right\rangle + \log \left((q/q_t)^{1/\beta_t} \right) \right] \quad (24)$$

$$= \underset{\mathbf{m}}{\operatorname{argmin}} \mathbb{E}_q \left[\log \frac{\exp \left\langle \boldsymbol{\phi}(\boldsymbol{\theta}), -\widehat{\nabla}_m \mathcal{L}_t \right\rangle q^{1/\beta_t}}{q_t^{1/\beta_t}} \right] \quad (25)$$

$$= \underset{\mathbf{m}}{\operatorname{argmin}} \mathbb{E}_q \left[\log \left(\frac{q^{1/\beta_t}}{q_t^{1/\beta_t} \exp \left\langle \boldsymbol{\phi}(\boldsymbol{\theta}), \widehat{\nabla}_m \mathcal{L}_t \right\rangle} \right) \right] \quad (26)$$

$$= \underset{\mathbf{m}}{\operatorname{argmin}} \frac{1}{\beta_t} \mathbb{E}_q \left[\log \left(\frac{q}{q_t \exp \left\langle \boldsymbol{\phi}(\boldsymbol{\theta}), \beta_t \widehat{\nabla}_m \mathcal{L}_t \right\rangle} \right) \right] \quad (27)$$

$$= \underset{\mathbf{m}}{\operatorname{argmin}} \frac{1}{\beta_t} \mathbb{D}_{KL} \left[q \parallel q_t \exp \left\langle \boldsymbol{\phi}(\boldsymbol{\theta}), \beta_t \widehat{\nabla}_m \mathcal{L}_t \right\rangle / \mathcal{Z}_t \right]. \quad (28)$$

where \mathcal{Z} is the normalizing constant of the distribution in the denominator which is a function of the gradient and step size.

Minimizing this KL divergence gives the update

$$q_{t+1}(\boldsymbol{\theta}) \propto q_t(\boldsymbol{\theta}) \exp \left\langle \boldsymbol{\phi}(\boldsymbol{\theta}), \beta_t \widehat{\nabla}_m \mathcal{L}_t \right\rangle. \quad (29)$$

By rewriting this, we see that we get an update in the natural parameters $\boldsymbol{\lambda}_t$ of $q_t(\boldsymbol{\theta})$, i.e.

$$\boldsymbol{\lambda}_{t+1} = \boldsymbol{\lambda}_t + \beta_t \widehat{\nabla}_m \mathcal{L}_t. \quad (30)$$

Recalling that the mean parameters of a Gaussian $q(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta} | \boldsymbol{\mu}, \operatorname{diag}(\boldsymbol{\sigma}^2))$ are $\mathbf{m}^{(1)} = \boldsymbol{\mu}$ and $\mathbf{m}^{(2)} = \boldsymbol{\sigma}^2 + \boldsymbol{\mu}^2$ and using the chain rule, we can express the gradient $\widehat{\nabla}_m \mathcal{L}_t$ in terms of $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$,

$$\widehat{\nabla}_{m^{(1)}} \mathcal{L} = \widehat{\nabla}_{\boldsymbol{\mu}} \mathcal{L} - 2 \left[\widehat{\nabla}_{\boldsymbol{\sigma}^2} \mathcal{L} \right] \circ \boldsymbol{\mu} \quad (31)$$

$$\widehat{\nabla}_{m^{(2)}} \mathcal{L} = \widehat{\nabla}_{\boldsymbol{\sigma}^2} \mathcal{L}. \quad (32)$$

Finally, recalling that the natural parameters of a Gaussian $q(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta} | \boldsymbol{\mu}, \operatorname{diag}(\boldsymbol{\sigma}^2))$ are $\boldsymbol{\lambda}^{(1)} = \boldsymbol{\sigma}^{-2} \circ \boldsymbol{\mu}$ and $\boldsymbol{\lambda}^{(2)} = -\frac{1}{2} \boldsymbol{\sigma}^{-2}$, we can rewrite the CVI updates in terms of $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$,

$$\boldsymbol{\sigma}_{t+1}^{-2} = \boldsymbol{\sigma}_t^{-2} - 2\beta_t \left[\widehat{\nabla}_{\boldsymbol{\sigma}^2} \mathcal{L}_t \right], \quad (33)$$

$$\boldsymbol{\mu}_{t+1} = \boldsymbol{\sigma}_{t+1}^2 \circ \left[\boldsymbol{\sigma}_t^{-2} \circ \boldsymbol{\mu}_t + \beta_t \left(\widehat{\nabla}_{\boldsymbol{\mu}} \mathcal{L}_t - 2 \left[\widehat{\nabla}_{\boldsymbol{\sigma}^2} \mathcal{L}_t \right] \circ \boldsymbol{\mu}_t \right) \right], \quad (34)$$

$$= \boldsymbol{\sigma}_{t+1}^2 \circ \left(\boldsymbol{\sigma}_t^{-2} - 2\beta_t \left[\widehat{\nabla}_{\boldsymbol{\sigma}^2} \mathcal{L}_t \right] \right) \boldsymbol{\mu}_t + \beta_t \boldsymbol{\sigma}_{t+1}^2 \circ \left[\widehat{\nabla}_{\boldsymbol{\mu}} \mathcal{L}_t \right], \quad (35)$$

$$= \boldsymbol{\mu}_t + \beta_t \boldsymbol{\sigma}_{t+1}^2 \circ \left[\widehat{\nabla}_{\boldsymbol{\mu}} \mathcal{L}_t \right]. \quad (36)$$