

ML 2019

Tightening bounds for variational inference by revisiting perturbation theory^{*}

To cite this article: Robert Bamler *et al* *J. Stat. Mech.* (2019) 124004

View the [article online](#) for updates and enhancements.



IOP | ebooks™

Bringing you innovative digital publishing with leading voices
to create your essential collection of books in STEM research.

Start exploring the **collection** - **download the first chapter of
every title for free.**

Tightening bounds for variational inference by revisiting perturbation theory*

Robert Bamler^{1,4}, Cheng Zhang^{2,4}, Manfred Opper³
and Stephan Mandt^{1,4}

¹ University of California, Irvine, CA, United States of America

² Microsoft Research, Cambridge, United Kingdom

³ Technical University of Berlin, Berlin, Germany

E-mail: rbamler@uci.edu, Cheng.Zhang@microsoft.com,
manfred.opper@tu-berlin.de and mandt@uci.edu

Received 19 June 2019

Accepted for publication 3 September 2019

Published 20 December 2019



Online at stacks.iop.org/JSTAT/2019/124004
<https://doi.org/10.1088/1742-5468/ab43d3>

Abstract. Variational inference has become one of the most widely used methods in latent variable modeling. In its basic form, variational inference employs a fully factorized variational distribution and minimizes its Kullback–Leibler divergence to the posterior. As the minimization can only be carried out approximately, this approximation induces a bias. In this paper, we revisit perturbation theory as a powerful way of improving the variational approximation. Perturbation theory relies on a form of Taylor expansion of the log marginal likelihood, vaguely in terms of the log ratio of the true posterior and its variational approximation. While first order terms give the classical variational bound, higher-order terms yield corrections that tighten it. However, traditional perturbation theory does not provide a lower bound, making it inapt for stochastic optimization. In this paper, we derive a similar yet alternative way of deriving corrections to the evidence lower bound that resemble perturbation theory, but that result in a valid bound. We show in experiments on Gaussian processes and variational autoencoders that the new

⁴ These authors contributed equally.

* This original manuscript is based on the article presented at NIPS 2017: Bamler R, Zhang C, Opper M and Mandt S 2017 Perturbative black box variational inference *Advances in Neural Information Processing Systems 30* ed I Guyon, U V Luxburg, S Bengio, H Wallach, R Fergus, S Vishwanathan and R Garnett (Red Hook, NY: Curran Associates, Inc.) pp 5079–5088

bounds are more mass covering, and that the resulting posterior covariances are closer to the true posterior and lead to higher likelihoods on held-out data.

Keywords: machine learning

Contents

1. Introduction	2
2. Background: perturbation theory for variational inference	4
3. Perturbative black box variational inference	6
3.1. Variational inference with generalized lower bounds.....	6
3.1.1. Variational inference and importance sampling.	7
3.1.2. Bias variance trade-off.	7
3.2. Perturbative lower bounds	7
3.3. Connection to variational perturbation theory.....	10
3.4. Proofs of the theorems.....	11
4. Experiments	13
4.1. Mass covering effect.....	13
4.2. GP regression on synthetic data.....	13
4.3. Gaussian process classification.....	14
4.3.1. Data.	14
4.3.2. Convergence speed comparison.	15
4.4. Variational autoencoder.....	15
5. Related work	17
5.1. Black box variational inference (BBVI).....	17
5.2. Generalized divergences measures	17
5.3. Variational perturbation theory	18
6. Conclusion	18
References	19

1. Introduction

Bayesian inference is the task of reasoning about random variables that can only be measured indirectly. Given a probabilistic model and a data set of observations, Bayesian inference seeks the posterior probability distribution over the remaining, i.e. unobserved or ‘latent’ variables. The computational bottleneck is calculating the

marginal likelihood, the probability of the data with latent variables being integrated out, which is also an important quantity for model selection. It is the objective function of the expectation-maximization (EM) algorithm [8], which has seen a recent resurgence with the popularity of variational autoencoders (VAEs) [16].

Variational inference (VI) [13, 14, 46] scales up approximate Bayesian inference to large data sets by framing inference as an optimization problem. VI derives and maximizes a lower bound to the log marginal likelihood, termed ‘evidence’, and uses this so-called ‘evidence lower bound’ (ELBO) as a proxy for the true evidence. In its original formulation, VI was limited to a restricted class of so-called conditionally conjugate models [13], for which a closed-form expression for the ELBO can be derived by solving known integrals. More recently, black box variational inference (BBVI) approaches have become more popular [34, 37], which lift this restriction by approximating the gradient of the ELBO by Monte Carlo samples. BBVI also enables the optimization of alternative bounds that do not possess closed-form integrals. With Monte Carlo gradient estimation, the focus shifts from tractability of these integrals to the variance of the Monte Carlo gradient estimators, asking for bounds with low-variance gradients.

In this article, we propose a family of new lower bounds that are tighter than the standard ELBO while having lower variance than previous alternatives and admitting unbiased Monte Carlo estimation of the gradients of the bound. We derive these new bounds by introducing ideas from so-called variational perturbation theory into BBVI.

Variational perturbation theory provides an alternative to VI for approximating the evidence [29–31, 41]. It is based on a Taylor expansion of the evidence around the variational distribution, i.e. a parameterized approximate posterior distribution. The lowest order of the Taylor expansion recovers the standard ELBO, and higher order terms correct for the mismatch between the variational distribution and true posterior.

Variational perturbation theory based on so-called cumulant expansions [29] requires a fair amount of manual derivations and puts strong constraints on the tractability of certain integrals. A cumulant expansion also generally does not result in a lower bound, making it impossible to minimize the approximation error with stochastic gradient descent. In this article, we propose a new variant of perturbation theory that addresses these issues. Our proposed perturbative expansion leads to a family of lower bounds on the marginal likelihood that can be optimized with the same black box techniques as the standard ELBO. The proposed bounds $\mathcal{L}^{(K)}$ are enumerated by an odd integer K , the order of the perturbative expansion, and are given by

$$\mathcal{L}^{(K)}(\lambda, V_0) = e^{-V_0} \sum_{k=0}^K \frac{1}{k!} \mathbb{E}_{\mathbf{z} \sim q_\lambda} \left[(V_0 + \log p(\mathbf{x}, \mathbf{z}) - \log q_\lambda(\mathbf{z}))^k \right] \leq p(\mathbf{x}). \quad (1)$$

Here, $p(\mathbf{x}, \mathbf{z})$ is the joint distribution of the probabilistic model with observed variables \mathbf{x} and latent variables \mathbf{z} , $q_\lambda(\mathbf{z})$ is the variational distribution with variational parameters λ , and $p(\mathbf{x})$ is the marginal likelihood. Further, the reference energy $V_0 \in \mathbb{R}$ is an additional free parameter over which we optimize jointly with λ .

The proposed bounds in equation (1) generalize the standard ELBO, which one obtains for $K = 1$. Higher order terms make the bound successively tighter while guaranteeing that it still remains a lower bound for all odd orders K . In the limit $K \rightarrow \infty$, equation (1) becomes an asymptotic series to the exact marginal likelihood $p(\mathbf{x})$.

Our main contributions are as follows.

- We revisit variational perturbation theory [29–31] based on cumulant expansions, which does not result in a bound. By introducing a generalized VI framework, we propose a similar yet alternative construction to cumulant expansions which results in a valid lower bound of the evidence.
- We furthermore show that the new bound admits unbiased Monte Carlo gradients with low variance, making it well suited to stochastic optimization. Unbiasedness is an important property and is necessary to guarantee convergence [38]. Biased gradients may destroy the bound and may lead to divergence problems [7, 10, 11].
- We evaluate the proposed method experimentally both for inference and for model selection. Already the lowest nonstandard order $K=3$ is less prone to underestimating posterior variances than standard VI, and it fits better models in a VAE on small data sets. While this reduction in bias comes at the cost of larger gradient variance, we show that the gradient variance is still smaller than in previously proposed alternatives [10, 12, 27] to the standard ELBO, resulting in faster convergence.

This paper is an extended version of a proceeding conference paper by the same authors [4], which also gave the first reference to VI as a form of biased importance sampling, where the tightness of the bound is traded off against a low-variance stochastic gradient.

Our paper is structured as follows. In section 2, we revisit the cumulant expansion for VI. In section 3, we derive a new family of bounds with similar properties, but which is amenable to stochastic optimization. We then discuss theoretical properties of the bounds. Experiments are presented in section 4. Finally, we discuss related work in section 5 and open questions in section 6.

2. Background: perturbation theory for variational inference

We begin by reviewing variational perturbation theory as introduced in [29]. We consider a probabilistic model with data \mathbf{x} , latent variables \mathbf{z} , and joint distribution $p(\mathbf{x}, \mathbf{z})$. The goal of Bayesian inference is to find the posterior distribution,

$$p(\mathbf{z}|\mathbf{x}) = \frac{p(\mathbf{x}, \mathbf{z})}{p(\mathbf{x})} \quad \text{where} \quad p(\mathbf{x}) = \int p(\mathbf{x}, \mathbf{z}) d\mathbf{z}. \quad (2)$$

Exact posterior inference is impossible in all but the most simple models because of the intractable integral defining the marginal likelihood $p(\mathbf{x})$ in equation (2). Variational perturbation theory approximates the log marginal likelihood, or evidence, $\log p(\mathbf{x})$, via a Taylor expansion. We introduce a so-called variational distribution $q_\lambda(\mathbf{z})$, parameterized by variational parameters λ , and we write the evidence as follows:

$$\log p(\mathbf{x}) = \log \left(\mathbb{E}_{\mathbf{z} \sim q_\lambda} \left[\frac{p(\mathbf{x}, \mathbf{z})}{q_\lambda(\mathbf{z})} \right] \right) = \log \left(\mathbb{E}_{\mathbf{z} \sim q_\lambda} \left[e^{-\beta V(\mathbf{x}, \mathbf{z})} \right] \right) \Big|_{\beta=1}. \quad (3)$$

Above, the ‘interaction energy’ V is defined as

$$V(\mathbf{x}, \mathbf{z}) \equiv \log q_\lambda(\mathbf{z}) - \log p(\mathbf{x}, \mathbf{z}), \quad (4)$$

and the notation $(\cdots)|_{\beta=1}$ on the right-hand side of equation (3) denotes that we introduced an auxiliary real parameter β , which we set to one at the end of the calculation.

The reason for this notation is as follows. We approximate the right-hand side of equation (3) by a finite order Taylor expansion in β before setting $\beta = 1$. While β is not a small parameter, it is a placeholder that keeps track of appearances of $V(\mathbf{x}, \mathbf{z})$ which we consider to be small for all \mathbf{z} in the support of q_λ . More precisely, it is enough to demand that $V(\mathbf{x}, \mathbf{z})$ is almost constant in \mathbf{z} with small deviations, as is the case when $q_\lambda(\mathbf{z})$ is close to the posterior $p(\mathbf{z}|\mathbf{x})$. (This will become clear in equation (6) below.) Hence, we can think of this approximation informally as an expansion in terms of the difference between the log variational distribution and the log posterior.

The first order term of the Taylor expansion is the evidence lower bound, or ELBO,

$$\mathcal{L}(\lambda) \equiv \mathbb{E}_{\mathbf{z} \sim q_\lambda} [-V(\mathbf{x}, \mathbf{z})] = \mathbb{E}_{\mathbf{z} \sim q_\lambda} [\log p(\mathbf{x}, \mathbf{z}) - \log q_\lambda(\mathbf{z})]. \quad (5)$$

As is well known in VI [5, 14, 46], the ELBO is a lower bound on the evidence, i.e. $\mathcal{L}(\lambda) \leq \log p(\mathbf{x})$ for all λ . VI maximizes the ELBO over λ , and thus tries to find the best approximation of the evidence that is possible within a first order Taylor expansion in V .

Higher order terms of the Taylor expansion lead to corrections to the standard ELBO. Suppressing the dependence of V on \mathbf{x} and \mathbf{z} to simplify the notation, we obtain

$$\begin{aligned} \log p(\mathbf{x}) \approx & \mathbb{E}_{q_\lambda} [-V] + \frac{1}{2} \mathbb{E}_{q_\lambda} [(V - \mathbb{E}_{q_\lambda}[V])^2] - \frac{1}{3!} \mathbb{E}_{q_\lambda} [(V - \mathbb{E}_{q_\lambda}[V])^3] \\ & + \frac{1}{4!} \mathbb{E}_{q_\lambda} [(V - \mathbb{E}_{q_\lambda}[V])^4] - \frac{1}{2} \left(\frac{1}{2} \mathbb{E}_{q_\lambda} [(V - \mathbb{E}_{q_\lambda}[V])^2] \right)^2 \\ & + \dots \end{aligned} \quad (6)$$

Equation (6) is called the *cumulant expansion* of the evidence [29]. Due to the higher order correction terms, it can potentially approximate the true evidence better than the ELBO.

Variational perturbation theory has not found its way into mainstream machine learning, which arguably has to do with two major problems of the cumulant expansion and related approaches. First, in contrast to the ELBO, a higher order cumulant expansion does not result in a lower bound on the evidence. This prevents the usage of gradient-based optimization methods (as opposed to, e.g. coordinate updates), which is currently the mainstream approach to VI. A second drawback that will be discussed below is that the cumulant expansion cannot be estimated efficiently using Monte Carlo sampling, as we discuss in section 3.3 below. In the following sections, we will present a similar construction of an improved variational objective that does not suffer from these shortcomings, and that is compatible with black-box optimization approaches. To this end, we have to start from a generalized formulation of VI.

3. Perturbative black box variational inference

This section presents our main results. We derive a family of new objective functions for BBVI which are amenable to stochastic optimization. We begin by deriving a generalized ELBO for VI based on concave functions. We then show that, in a special case, we obtain a strict bound with resemblance to variational perturbation theory.

3.1. Variational inference with generalized lower bounds

VI approximates the evidence $\log p(\mathbf{x})$ of the model in equation (2) by maximizing the ELBO $\mathcal{L}(\lambda)$ (equation (5)) over variational parameters λ . A more common approach to derive the ELBO is through Jensen's inequality [5, 14, 46]. Here, we present a similar derivation for a broader family of lower bounds.

To this end, we consider an arbitrary concave function over the positive reals, $f: \mathbb{R}_{>0} \rightarrow \mathbb{R}$. As an example, f could be the logarithm. We now consider the marginal likelihood $p(\mathbf{x}) = \mathbb{E}_{\mathbf{z} \sim q_\lambda} [p(\mathbf{x}, \mathbf{z})/q_\lambda(\mathbf{z})]$, and we derive a lower bound on $f(p(\mathbf{x}))$ using Jensen's inequality:

$$f(p(\mathbf{x})) \geq \mathbb{E}_{\mathbf{z} \sim q_\lambda} \left[f \left(\frac{p(\mathbf{x}, \mathbf{z})}{q_\lambda(\mathbf{z})} \right) \right] =: \mathcal{L}_f(\lambda). \quad (7)$$

As stated above, for $f(\cdot) = \log(\cdot)$, the bound in equation (7) is the standard ELBO from equation (5). In this case, we refer to the corresponding VI scheme as KLVI, as maximizing the standard ELBO is equivalent to minimizing the Kullback–Leibler (KL) divergence from $q_\lambda(\mathbf{z})$ to the true posterior $p(\mathbf{z}|\mathbf{x})$.

The above class of generalized bounds $\mathcal{L}_f(\lambda)$ is compatible with current mainstream optimization schemes for VI, oftentimes summarized as BBVI [34, 37]. For the reader's convenience, we briefly review these techniques. BBVI optimizes the ELBO via stochastic gradient descent (SGD) based on noisy estimates of its gradient. Crudely speaking, BBVI obtains its gradient estimates in three steps: (1) by drawing Monte Carlo (MC) samples $\mathbf{z} \sim q_\lambda$, (2) by averaging the bound over these samples, and (3) by computing the gradient on the resulting average. A slight complication arises as both the expression inside the expectation in equation (7) as well as the distribution of MC samples itself depend on λ . The two main solutions to this problem are as follows.

- Score function gradients [34], or the REINFORCE method, use the chain rule to relate the change in MC samples to a change in the log variational distribution,

$$\nabla_\lambda \mathcal{L}_f(\lambda) = \mathbb{E}_{\mathbf{z} \sim q_\lambda} \left[(\nabla_\lambda \log q_\lambda(\mathbf{z})) f \left(\frac{p(\mathbf{x}, \mathbf{z})}{q_\lambda(\mathbf{z})} \right) + \nabla_\lambda f \left(\frac{p(\mathbf{x}, \mathbf{z})}{q_\lambda(\mathbf{z})} \right) \right]. \quad (8)$$

- Reparameterization gradients [16] can be used if the samples $\mathbf{z} \sim q_\lambda$ can be expressed as some deterministic differentiable transformation $\mathbf{z} = g(\boldsymbol{\epsilon}; \lambda)$ of auxiliary random variables $\boldsymbol{\epsilon}$ drawn from some λ -independent distribution $\tilde{q}(\boldsymbol{\epsilon})$. Under these conditions, the gradient of the bound can be written as

$$\nabla_{\lambda} \mathcal{L}_f(\lambda) = \mathbb{E}_{\epsilon \sim \tilde{q}} \left[\nabla_{\lambda} f \left(\frac{p(\mathbf{x}, g(\epsilon; \lambda))}{q_{\lambda}(g(\epsilon; \lambda))} \right) \right]. \quad (9)$$

One obtains an unbiased estimate of $\nabla_{\lambda} \mathcal{L}_f(\lambda)$ by approximating the expectation on the right-hand side of equations (8) or (9) via MC samples $\mathbf{z} \sim q_{\lambda}$ or $\epsilon \sim \tilde{q}$, respectively. Empirically, reparameterization gradients often have smaller variance than score function gradients.

Based on these new bounds $\mathcal{L}_f(\lambda)$, we make the following observations.

3.1.1. Variational inference and importance sampling. The generalized family of bounds, equation (7), shows that VI and importance sampling are closely connected [4, 11]. For $f(\cdot)$ being the identity the variational bound becomes independent of λ and instead becomes an unbiased importance sampling estimator of $p(\mathbf{x})$, with the proposal distribution $q_{\lambda}(\mathbf{z})$ and the importance ratio $p(\mathbf{x}, \mathbf{z})/q_{\lambda}(\mathbf{z}) = e^{-V(\mathbf{x}, \mathbf{z})}$. For other choices of f , the bound is no longer an unbiased estimator of $p(\mathbf{x})$. This way, we can identify BBVI as a form of biased importance sampling.

3.1.2. Bias variance trade-off. Importance sampling suffers from high variance in high dimensions because the log importance ratio $V(\mathbf{x}, \mathbf{z})$ (equation (4)) scales approximately linearly with the size of the model. (To see this, consider a setup where both p and q factorize over all N data points, in which case V is proportional to N .)

The choice of function f in equation (7) trades off some of this variance for a bias. KLVI uses $f(\cdot) = \log(\cdot)$, which leads to low variance gradient estimates in equations (8) and (9) that depend only linearly on V , rather than exponentially. This makes the KLVI bound easy to optimize, at the cost of some bias. An alternative to KLVI that has been explored in the literature is α -VI [10, 12, 21, 27], which corresponds to setting $f(\xi) = \xi^{1-\alpha}$ with some $\alpha \in (0, 1)$, and thus $f(e^{-V}) = e^{-(1-\alpha)V}$. This choice has an alternative bias-variance trade-off, as V enters in the exponent, leading to more variance than KLVI. Our empirical results in figure 4, discussed in section 4.3, confirm this picture by showing that α -VI suffers from slow convergence due to large gradient variance.

According to our new findings, a good alternative bound should perform favorably on the bias-variance trade-off. It should ideally depend on V only polynomially as opposed to exponentially. Such a family of bounds will be introduced next. We also show how they connect to variational perturbation theory (section 2).

3.2. Perturbative lower bounds

We now propose a family of concave functions f for the bound in equation (7) that is inspired by ideas from perturbation theory (section 2). The choice of the function f determines the tightness of the bound $\mathcal{L}_f(\lambda)$. A tight bound, i.e. a good approximation of the marginal likelihood, is important for the variational expectation-maximization (EM) algorithm, which uses the bound as a proxy for the marginal likelihood. The bound $\mathcal{L}_f(\lambda)$ becomes perfectly tight if we set $f = \text{id}$ to the identity function, $\text{id}(\xi) = \xi$ (black line in figure 1). However, as we discussed, this is a singular limit in which the




Figure 1. Different choices for the concave function $f(\xi)$ in equation (7). For $f(\xi) = \xi$ (black), the bound is tight but independent of λ . KLVI corresponds to $f(\xi) = \log(\xi) + \text{const.}$ (pink). Our proposed PBBVI bound uses $f_{V_0}^{(K)}(\xi)$ (green, equation (10)), which is tighter than KLVI for $\xi \approx e^{-V_0}$ (we set $K=3$ and $V_0=0$ for PBBVI here).

bound $\mathcal{L}_{\text{id}}(\lambda) = p(\mathbf{x})$ does not depend on λ , and therefore does not provide any gradient signal to learn an optimal variational distribution q_λ .

Ideally, the bound should be as tight as possible if $q_\lambda(\mathbf{z})$ is close to the posterior $p(\mathbf{z}|\mathbf{x})$. In order to obtain a gradient signal, the bound should be lower if $q_\lambda(\mathbf{z})$ is very different from the posterior. To achieve this, we recall that the argument ξ of the function $f(\xi)$ in equation (7) is the fraction $p(\mathbf{x}, \mathbf{z})/q_\lambda(\mathbf{z})$, which equals $p(\mathbf{x})$ if $q_\lambda(\mathbf{z})$ is the true posterior. Thus, the concave function f should be close to the identity function for arguments close to $p(\mathbf{x})$. Since the only things we know about $p(\mathbf{x})$ are that it is positive and independent of \mathbf{z} , we parameterize it as e^{-V_0} with a real parameter V_0 , over which we optimize.

For any $V_0 \in \mathbb{R}$, and any odd integer $K \geq 1$, we propose the following concave function:

$$f_{V_0}^{(K)}(\xi) = e^{-V_0} \sum_{k=0}^K \frac{1}{k!} (V_0 + \log \xi)^k \xrightarrow{K \rightarrow \infty} \text{id}(\xi) \quad (\text{pointwise } \forall \xi). \quad (10)$$

The pink and green lines in figure 1 plot $f_{V_0}^{(K)}(\xi)$ for $V_0=0$ and for $K=1$ and $K=3$, respectively. The function $f_{V_0}^{(K)}$ converges pointwise to the identity function (black line in figure 1) for $K \rightarrow \infty$ because it is the K th order Taylor expansion of $\text{id}(\xi) = e^{-V_0} e^{V_0 + \log \xi}$ in $\log \xi$ around the reference point $-V_0$. For any finite K , the Taylor expansion approximates the identity function (black line in figure 1) most closely if the expansion parameter $|V_0 + \log \xi|$ is small. We provide further intuition for equation (10) in section 3.3 below.

An important result of this section is that the functions $f_{V_0}^{(K)}$ define a lower bound on the marginal likelihood for any choice of λ and V_0 ,

$$p(\mathbf{x}) \geq f(p(\mathbf{x})) \geq \mathcal{L}_{f_{V_0}^{(K)}}(\lambda) =: \mathcal{L}^{(K)}(\lambda, V_0) \quad (\text{for odd } K). \quad (11)$$

Here, the last equality merely introduces a more convenient notation for the bound defined by equations (7) and (10). An explicit form for $\mathcal{L}^{(K)}(\lambda, V_0)$ was given in equation (1) in the introduction. Equation (11) follows from equation (7) and from the fact that $f_{V_0}^{(K)}$ is concave and lies below the identity function, which we prove formally in section 3.4 below. Since $\mathcal{L}^{(K)}(\lambda, V_0)$ is a lower bound on $p(\mathbf{x})$ for all V_0 , we can find the optimal reference point V_0 that leads to the tightest bound by maximizing $\mathcal{L}^{(K)}(\lambda, V_0)$ jointly over both λ and V_0 . We also prove in section 3.4 that the bound on $p(\mathbf{x})$ is nontrivial, i.e. that $\mathcal{L}^{(K)}(\lambda, V_0)$ is positive at its maximum. (A negative bound would be vacuous since $p(\mathbf{x})$ is trivially larger than any negative number.)

Algorithm 1. Perturbative black box variational inference (PBBVI).

Input: joint probability $p(\mathbf{x}, \mathbf{z})$; order of perturbation K (odd integer); learning rate schedule ρ_t ; number of Monte Carlo samples S ; number of training iterations T ; variational family $q_\lambda(\mathbf{z})$ that admits reparameterization gradients as in equation (9).

Output: fitted variational parameters λ^* .

```

1 Initialize  $\lambda$  randomly and  $V_0 \leftarrow 0$ ;
2 for  $t \leftarrow 1$  to  $T$  do
3   Draw  $S$  samples  $\epsilon_1, \dots, \epsilon_S \sim \tilde{q}$  from the noise distribution (see equation (9));
   // Obtain gradient estimates of surrogate objective  $\tilde{\mathcal{L}}^{(K)}(\lambda, V_0)$ , see equation (12):
4    $g_\lambda \leftarrow \nabla_\lambda \left[ \frac{1}{S} \sum_{s=1}^S \sum_{k=0}^K \frac{1}{k!} (V_0 + \log p(\mathbf{x}, g(\epsilon_s; \lambda)) - \log q_\lambda(g(\epsilon_s; \lambda)))^k \right]$ ;
5    $g_{V_0} \leftarrow \nabla_{V_0} \left[ \frac{1}{S} \sum_{s=1}^S \sum_{k=0}^K \frac{1}{k!} (V_0 + \log p(\mathbf{x}, g(\epsilon_s; \lambda)) - \log q_\lambda(g(\epsilon_s; \lambda)))^k \right]$ ;
   // Perform update steps with rescaled gradients, see equation (13):
6    $\lambda \leftarrow \lambda + \rho_t g_\lambda$ ;
7    $V_0 \leftarrow V_0 + \rho_t \left[ g_{V_0} - \frac{1}{S} \sum_{s=1}^S \sum_{k=0}^K \frac{1}{k!} (V_0 + \log p(\mathbf{x}, g(\epsilon_s; \lambda)) - \log q_\lambda(g(\epsilon_s; \lambda)))^k \right]$ ;
end
```

To optimize the bound, we use SGD with reparameterization gradients, see section 3.1. Score function gradients can be used in a similar way. algorithm 1 explains the optimization in detail. We name the method ‘perturbative black box variational inference’ (PBBVI). The algorithm implicitly scales all gradients by a factor of e^{V_0} to avoid a potential exponential explosion or decay of gradients due to the factor e^{-V_0} in the bound, equation (1). We calculate these rescaled gradients in a numerically stable way by considering the surrogate objective function

$$\begin{aligned} \tilde{\mathcal{L}}^{(K)}(\lambda, V_0) &\equiv e^{V_0} \mathcal{L}^{(K)}(\lambda, V_0) \\ &= \sum_{k=0}^K \frac{1}{k!} \mathbb{E}_{\mathbf{z} \sim q_\lambda} \left[(V_0 + \log p(\mathbf{x}, \mathbf{z}) - \log q_\lambda(\mathbf{z}))^k \right]. \end{aligned} \quad (12)$$

Table 1. Results for Gaussian process (GP) experiments. (a) Average marginal posterior variances at the positions of data points for GP regression with synthetic data (see figure 3). The proposed PBBVI is closer to the analytic solution than standard KLVI. (b) Error rate of GP classification on the test set. The lower the better. Our proposed PBBVI consistently obtains better classification results.

(a) GP regression		(b) GP classification				
Method	Average variances	Data set	Crab	Pima	Heart	Sonar
Analytic	0.0415	KLVI	0.22	0.245	0.148	0.212
KLVI	0.0176	PBBVI	0.11	0.240	0.133	0.173
PBBVI	0.0355					

The rescaled gradients are thus

$$\begin{aligned} e^{V_0} \nabla_{\lambda} \mathcal{L}^{(K)}(\lambda, V_0) &= \nabla_{\lambda} \tilde{\mathcal{L}}^{(K)}(\lambda, V_0); \\ e^{V_0} \nabla_{V_0} \mathcal{L}^{(K)}(\lambda, V_0) &= \nabla_{V_0} \tilde{\mathcal{L}}^{(K)}(\lambda, V_0) - \tilde{\mathcal{L}}^{(K)}(\lambda, V_0). \end{aligned} \quad (13)$$

Equations (12) and (13) allow us to estimate the rescaled gradients (see lines 4–7 in algorithm 1) without having to evaluate any potentially overflowing or underflowing expressions $e^{\pm V_0}$.

3.3. Connection to variational perturbation theory

We now provide more intuition for the choice of the functions $f_{V_0}^{(K)}$ in equation (10) by comparing the resulting bounds, equation (1), to the cumulant expansion, equation (6). We will show that the bounds enjoy similar benefits as the cumulant expansion while being valid lower bounds and providing unbiased Monte Carlo gradients.

Both the cumulant expansion and the proposed bounds $\mathcal{L}^{(K)}(\lambda, V_0)$ are Taylor expansions in $V(\mathbf{x}, \mathbf{z}) = \log q_{\lambda}(\mathbf{z}) - \log p(\mathbf{x}, \mathbf{z})$. The cumulant expansion in equation (6) is a Taylor expansion of the evidence $\log p(\mathbf{x})$. By contrast, the bound $\mathcal{L}^{(K)}(\lambda, V_0)$ in equation (1) is a Taylor expansion of the marginal likelihood $p(\mathbf{x})$, and the expansion starts from a reference point V_0 over which we optimize.

Despite this difference, the two expansions are remarkably similar up to order $K = 3$. Up to this order, the cumulant expansion, equation (6), coincides with a Taylor expansion of $p(\mathbf{x})$ in $(V - \mathbb{E}_{q_{\lambda}}[V])$. It is thus a special case of the proposed lower bound $\mathcal{L}^{(3)}(\lambda, V_0)$ if we set $V_0 = \mathbb{E}_{q_{\lambda}}[V]$. For $K > 3$, the cumulant expansion contains additional terms that do not appear in the bound $\mathcal{L}^{(K)}(\lambda, V_0)$, such as the last spelled out term on the right-hand side of equation (6). These terms contain products of expectations under q_{λ} , which are difficult to estimate with Monte Carlo techniques without introducing an additional bias. The bounds $\mathcal{L}^{(K)}(\lambda, V_0)$, by contrast, depend only linearly on expectations under q_{λ} , which means that they can be estimated without bias with a single sample from q_{λ} .

In addition, the main advantage of $\mathcal{L}^{(K)}(\lambda, V_0)$ over the cumulant expansion is that $\mathcal{L}^{(K)}(\lambda, V_0)$ is a lower bound on the marginal likelihood for all λ and all V_0 . This allows us to make the bound as tight as possible by maximizing it over λ and V_0 using stochastic gradient estimates. By contrast, variational perturbation theory with the cumulant

expansion requires either optimizing towards saddle points, or optimizing a separate objective function, such as the ELBO, to find a good variational distribution [29].

The lack of a lower bound in the cumulant expansion is a particular limitation when performing model selection with the variational expectation-maximization (EM) algorithm. Variational EM fits a model by maximizing an approximation of the marginal likelihood or the model evidence over model parameters. Such an optimization can diverge if the approximation is not a lower bound. This makes perturbative BBVI an interesting alternative for improved evidence approximations [7, 10, 11, 21, 33].

3.4. Proofs of the theorems

We conclude the presentation of the proposed bounds $\mathcal{L}^{(K)}(\lambda, V_0)$ by providing proofs of the claims made in section 3.2.

Theorem 1. *For all $V_0 \in \mathbb{R}$ and all odd integers $K \geq 1$, the function $f_{V_0}^{(K)}$ defined in equation (10) satisfies the following properties:*

- (i) $f_{V_0}^{(K)}$ is concave; and
- (ii) $f_{V_0}^{(K)}(\xi) \leq \xi \quad \forall \xi \in \mathbb{R}_{>0}$.

Proof.

- (i) The first derivative of $f_{V_0}^{(K)}$ is

$$f_{V_0}^{(K)'}(\xi) = e^{-V_0} \sum_{k=0}^{K-1} \frac{1}{k!} \frac{(V_0 + \log \xi)^k}{\xi}. \quad (14)$$

For the second derivative, the two contributions from the denominator and the enumerator cancel for all but the highest order term, and we obtain

$$f_{V_0}^{(K)''}(\xi) = -\frac{e^{-V_0}}{(K-1)!} \frac{(V_0 + \log \xi)^{K-1}}{\xi^2} \leq 0 \quad (15)$$

which is nonnegative everywhere for odd K . Therefore, $f_{V_0}^{(K)}$ is concave.

- (ii) Consider the function $g(\xi) = f_{V_0}^{(K)}(\xi) - \xi$, which is also concave since it has the same second derivative as $f_{V_0}^{(K)}$. Further, g has a stationary point at $\xi = e^{-V_0}$ since $f_{V_0}^{(K)'}(e^{-V_0}) = 1$, which can be verified by inserting into equation (14). For a concave function, a stationary point is a global maximum. Therefore, we have

$$g(\xi) \leq g(e^{-V_0}) = 0 \quad \forall \xi \in \mathbb{R}_{>0} \quad (16)$$

which is equivalent to the proposition. \square

Theorem 2. *The bound on the positive quantity $p(\mathbf{x})$ is nontrivial, i.e. the maximum value of the bound, $\max_{\lambda, V_0} \mathcal{L}^{(K)}(\lambda, V_0)$, is positive for all odd integers $K \geq 1$.*

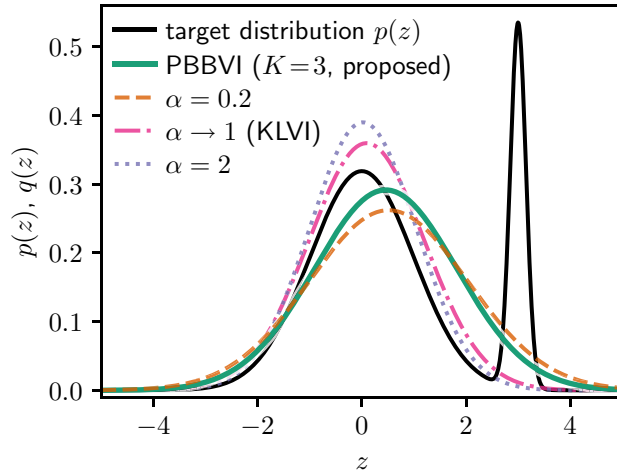


Figure 2. Behavior of different VI methods on fitting a univariate Gaussian to a bimodal target distribution (black). PBBVI (proposed, green) covers more of the mass of the entire distribution than the traditional KLVI (pink). α -VI is mode seeking for large α (blue) and mass covering for smaller α (orange).

Proof. At the maximum position (λ^*, V_0^*) of $\mathcal{L}^{(K)}(\lambda, V_0)$, its gradient is zero. Taking the gradient with respect to V_0 in equation (1) and using the product rule for the prefactor e^{-V_0} and the remaining expression, we find that all terms except the contribution from $k = K$ cancel. We thus obtain, at the maximum position (λ^*, V_0^*) ,

$$\mathbb{E}_{\mathbf{z} \sim q_{\lambda^*}} \left[(V_0^* + \log p(\mathbf{x}, \mathbf{z}) - \log q_{\lambda^*}(\mathbf{z}))^K \right] = 0. \quad (17)$$

Thus, when we evaluate $\mathcal{L}^{(K)}$ at (λ^*, V_0^*) , the term with $k = K$ on the right-hand side of equation (1) evaluates to zero, and we are left with

$$\mathcal{L}^{(K)}(\lambda^*, V_0^*) = e^{-V_0^*} \mathbb{E}_{\mathbf{z} \sim q_{\lambda^*}} \left[h(V_0^* + \log p(\mathbf{x}, \mathbf{z}) - \log q_{\lambda^*}(\mathbf{x}, \mathbf{z})) \right] \quad (18)$$

with

$$h(u) = \sum_{k=0}^{K-1} \frac{u^k}{k!} \quad (19)$$

where the sum runs only to $K - 1$.

We now show that $h(u)$ is positive for all u and all odd K . If $K = 1$, then $h(u) = 1$ is trivially positive. For $K \geq 3$, $h(u)$ is a polynomial in u of even order $K - 1$, whose highest order term has a positive coefficient $\frac{1}{(K-1)!}$. Therefore, $h(u)$ goes to positive infinity for both $u \rightarrow \infty$ and $u \rightarrow -\infty$. It thus has a global minimum at some value $\tilde{u} \in \mathbb{R}$. At the global minimum, its derivative is zero, i.e.

$$0 = h'(\tilde{u}) = \sum_{k=0}^{K-2} \frac{\tilde{u}^k}{k!}. \quad (20)$$

Subtracting equation (20) from equation (19), we find that the value of h at its global minimum \tilde{u} is

$$h(\tilde{u}) = \frac{\tilde{u}^{K-1}}{(K-1)!} \geq 0 \quad (21)$$

which is nonnegative because $K-1$ is even. Further, $h(\tilde{u})$ is not zero since this would imply $\tilde{u} = 0$ by equation (21), but $h(0) = 1$ by equation (19). Therefore, $h(\tilde{u})$ is strictly positive, and since \tilde{u} is a global minimum of h , we have $h(u) > 0$ for all $u \in \mathbb{R}$. Inserting into equation (18) concludes the proof that the lower bound at the optimum is positive. \square

4. Experiments

We evaluate PBBVI with different models. We begin with a qualitative comparison on a simple one-dimensional toy problem (section 4.1). We then investigate the behavior of BBVI in a controlled setup of Gaussian processes on synthetic data (section 4.2). Next, we evaluate PBBVI based on a classification task using Gaussian processes classifiers, where we use data from the UCI machine learning repository (section 4.3). This is a Bayesian non-conjugate setup where black box inference is required. Finally, we use an experiment with the VAE to explore our approach on a deep generative model (section 4.4). This experiment is carried out on MNIST data. We use the perturbative order $K=3$ for all experiments with PBBVI. This corresponds to the lowest order beyond standard VI with the KL divergence (KLVI), since K has to be an odd integer, and PBBVI with $K=1$ is equivalent to KLVI. Across all the experiments, PBBVI demonstrates advantages based on different metrics.

4.1. Mass covering effect

In figure 2, we fit a Gaussian distribution to a one-dimensional bimodal target distribution (black line), using different divergences. Compared to BBVI with the standard KL divergence (KLVI, pink line), α -divergences are more mode-seeking (purple line) for large values of α , and more mass-covering (orange line) for small α [21] (the limit $\alpha \rightarrow 1$ recovers KLVI [27]). Our PBBVI bound ($K=3$, green line) achieves a similar mass-covering effect as in α -divergences, but with associated low-variance reparameterization gradients. This is seen in figure 4(b), discussed in section 4.3 below, which compares the gradient variances of α -VI and PBBVI as a function of latent dimensions.

4.2. GP regression on synthetic data

In this section, we inspect the inference behavior using a synthetic data set with Gaussian processes (GP). We generate the data according to a Gaussian noise distribution centered around a mixture of sinusoids, and we sample 50 data points (green dots in figure 3). We then use a GP to model the data, thus assuming the generative process $f \sim \mathcal{GP}(0, \Lambda)$ and $y_i \sim \mathcal{N}(f_i, \epsilon)$.




Figure 3. Gaussian process regression on synthetic data (green dots). Three standard deviations are shown in varying shades of orange. The blue dashed lines show three standard deviations of the true posterior. The red dashed lines show the inferred three standard deviations using KLVI (a) and PBBVI (b). We see that the results from our proposed PBBVI are close to the analytic solution while traditional KLVI underestimates the variances. (a) KLVI. (b) PBBVI with $K = 3$.

In this experiment, the model admits an analytic solution (three standard deviations shown in blue dashed lines in figure 3). We compare the analytic solution to approximate posteriors using a fully factorized Gaussian variational distribution obtained by KLVI (figure 3(a)) and by the proposed PBBVI (figure 3(b)). The results from PBBVI are almost identical to the analytic solution. In contrast, KLVI underestimates the posterior variance. This is consistent with table 1(a), which shows the average marginal variances. PBBVI results are much closer to the analytic results.

4.3. Gaussian process classification

We evaluate the performance of PBBVI and KLVI on a GP classification task. Since the model is non-conjugate, no analytical baseline is available in this case. We model the data with the following generative process:

$$f \sim \mathcal{GP}(0, \Lambda), \quad z_i = \sigma(f_i), \quad y_i \sim \text{Bern}(z_i). \quad (22)$$

Above, Λ is the GP kernel, σ indicates the sigmoid function, and Bern indicates the Bernoulli distribution. We furthermore use the Matern-32 kernel,

$$\Lambda_{ij} = s^2 \left(1 + \frac{\sqrt{3} r_{ij}}{l} \right) \exp \left(-\frac{\sqrt{3} r_{ij}}{l} \right), \quad r_{ij} = \|x_i - x_j\|_2. \quad (23)$$

4.3.1. Data. We use four data sets from the UCI machine learning repository, suitable for binary classification: Crab (200 datapoints), Pima (768 datapoints), Heart (270 datapoints), and Sonar (208 datapoints). We randomly split each of the data sets into two halves. One half is used for training and the other half is used for testing. We set the hyper parameters $s = 1$ and $l = \sqrt{D}/2$ throughout all experiments, where D is the dimensionality of the input x .

Table 1(b) shows the classification performance (error rate) for these data sets. Our proposed PBBVI consistently performs better than the traditional KLVI.

4.3.2. Convergence speed comparison. We also carry out a comparison in terms of speed of convergence, focusing on PBBVI and α -VI. Our results indicate that the smaller variance of the reparameterization gradient in PBBVI leads to faster convergence of the optimization algorithm.

We train the GP classifier from equations (22) and (23) on the Sonar UCI data set using a constant learning rate. Figure 4(a) shows the test log-likelihood under the posterior mean as a function of training iterations. We split the data set into equally sized training, validation, and test sets. We then tune the learning rate and the number of Monte Carlo samples per gradient step to obtain optimal performance on the validation set after optimizing the α -VI bound with a fixed budget of random samples. We use $\alpha = 0.5$ here; smaller values of α (i.e. stronger mass-covering effect) leads to even slower convergence. We then optimize the PBBVI lower bound with the same learning rate and number of Monte Carlo samples. The final test error rate is 22% (the data set has binary labels and is approximately balanced). Although the hyperparameters are tuned for α -VI, PBBVI converges an order of magnitude faster (figure 4(a)).

Figure 4(b) provides more insight in the scaling of the gradient variance. Here, we fit GP regression models on synthetically generated data by maximizing the PBBVI lower bound and the α -VI lower bound with $\alpha \in \{0.2, 0.5, 2\}$. We generate a separate synthetic data set for each $N \in \{1, \dots, 200\}$ by drawing N random data points around a sinusoidal curve. For each N , we fit a one-dimensional GP regression with PBBVI and α -VI, respectively, using the same data set for both methods. The variational distribution is a fully factorized Gaussian with N latent variables. After convergence, we estimate the sampling variance of the gradient of each lower bound with respect to the posterior mean. We calculate the empirical variance of the gradient based on 10^5 samples from q_λ , and we average over the N coordinates. Figure 4 shows the average sampling variance as a function of N on a logarithmic scale. The variance of the gradient of the α -VI bound grows exponentially in the number of latent variables. By contrast, we find only algebraic growth for PBBVI.

4.4. Variational autoencoder

We experiment on VAEs, and we compare the PBBVI and the KLVI bound in terms of predictive likelihoods on held-out data [16]. Autoencoders compress unlabeled training data into low-dimensional representations by fitting it to an encoder-decoder model that maps the data to itself. These models are prone to learning the identity function when the hyperparameters are not carefully tuned, or when the network is too expressive, especially for a moderately sized training set. VAEs are designed to partially avoid this problem by estimating the uncertainty that is associated with each data point in the latent space. It is therefore important that the inference method does not underestimate posterior variances. We show that, for small data sets, training a VAE by maximizing the PBBVI lower bound leads to higher predictive likelihoods than maximizing the KLVI lower bound.

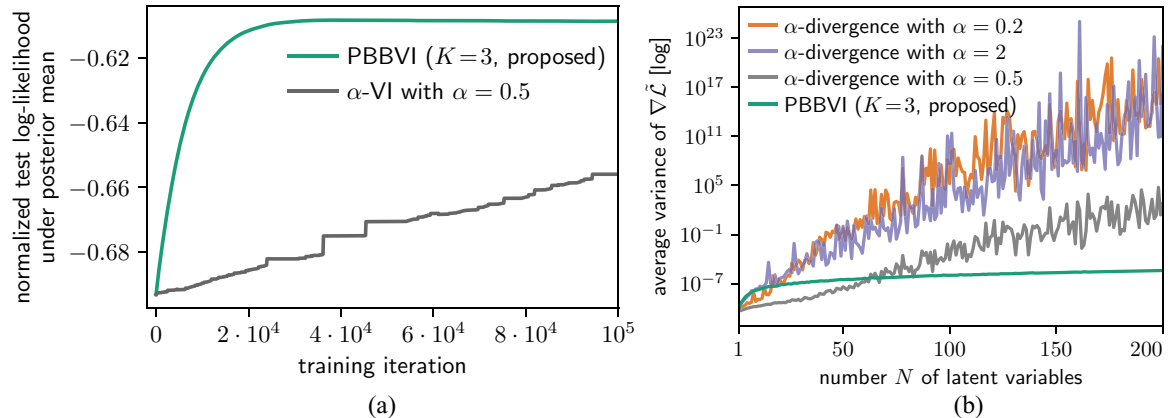


Figure 4. Comparisons between PBBVI (proposed) and α -VI. (a) Training curves (test log-likelihood per data point) for GP classification on the Sonar data set. PBBVI converges faster than α -VI even though we tuned the number of MC samples per training step (100) and the constant learning rate (10^{-5}) so as to maximize the performance of α -VI on a validation set. (b) Sampling variance of the stochastic gradient (averaged over its components) in the optimum of a GP regression model with synthetic data, for α -divergences (orange, purple, gray), and the proposed PBBVI (green). The variance grows exponentially with the latent dimension N for α -VI, and only algebraically for PBBVI. (a) Speed of convergence. (b) Gradient variance.

We train the VAE on the MNIST data set of handwritten digits [19]. We build on the publicly available implementation by [7] and we also use the same architecture and hyperparameters, with $L = 2$ stochastic layers and $S = 5$ samples from the variational distribution per gradient step. The model has 100 latent units in the first stochastic layer and 50 latent units in the second stochastic layer.

The VAE model factorizes over all data points. We train it by stochastically maximizing the sum of the PBBVI lower bounds for all data points using a minibatch size of 20. The VAE amortizes the gradient signal across data points by training inference networks. The inference networks express the mean and variance of the variational distribution as a function of the data point. We add an additional inference network that learns the mapping from a data point to the reference energy V_0 . Here, we use a network with four fully connected hidden layers of 200, 200, 100, and 50 units, respectively.

MNIST contains 60 000 training images. To test our approach on smaller-scale data where Bayesian uncertainty matters more, we evaluate the test likelihood after training the model on randomly sampled fractions of the training set. We use the same training schedules as in the publicly available implementation, keeping the total number of training iterations independent of the size of the training set. Different to the original implementation, we shuffle the training set before each training epoch as this turns out to increase the performance for both our method and the baseline.

Figure 5 shows the predictive log-likelihood of the whole test set, where the VAE is trained on random subsets of different sizes of the training set. We use the same subset to train with PBBVI and KLVI for each training set size. PBBVI leads to a higher predictive likelihood than traditional KLVI on subsets of the data. We explain this finding with our observation that the variational distributions obtained from PBBVI capture

Tightening bounds for variational inference by revisiting perturbation theory




Figure 5. Predictive likelihood of a VAE trained on different sizes of the data. The training data are randomly sampled subsets of the MNIST training set. The higher value the better. Our proposed PBBVI method outperforms KLVI mainly when the size of the training data set is small. The fewer the training data, the more advantage PBBVI obtains.

more of the posterior variance. As the size of the training set grows—and the posterior uncertainty decreases—the performance of KLVI catches up with PBBVI.

As a potential explanation why PBBVI converges to the KLVI result for large training sets, we note that $\mathbb{E}_{q_{\lambda^*}}[(V_0^* - V)^3] = 0$ at the optimal variational distribution q_{λ^*} and reference energy V_0^* (see section 3.4). If V becomes a symmetric random variable (such as a Gaussian) in the limit of a large training set, then this implies that $\mathbb{E}_{q_{\lambda^*}}[V] = V_0^*$, and PBBVI with $K = 3$ reduces to KLVI for large training sets.

5. Related work

Our approach is related to BBVI, VI with generalized divergences, and variational perturbation theory. We thus briefly discuss related work in these three directions.

5.1. Black box variational inference (BBVI)

BBVI has already been addressed in the introduction [16, 34, 37, 39, 40]; it enables VI for many models [2, 3, 9, 20, 23, 26, 35]. Recent developments include variance reduction and improvements in reparameterization and amortization [6, 17, 24, 25, 36, 39, 45] which are all compatible with our approach. Our work builds upon BBVI in that BBVI makes a large class of new divergence measures between the posterior and the approximating distribution tractable. Depending on the divergence measure, BBVI may suffer from high-variance stochastic gradients. This is a practical problem that we aim to improve in this paper.

5.2. Generalized divergences measures

Our work connects to generalized information-theoretic divergences [1]. Minka [27] introduced a broad class of divergences for VI, including α -divergences. Most of these divergences have been intractable in large-scale applications until the advent of BBVI. In this context, α -divergences were first suggested by [12] for local divergence

minimization, and later for global minimization by [21] and [10]. As we show in this paper, α -divergences have the disadvantage of inducing high-variance gradients, since the ratio between posterior and variational distribution enters the bound polynomially instead of logarithmically. In contrast, our approach leads to a more stable inference scheme in high dimensions.

5.3. Variational perturbation theory

Perturbation theory refers to methods that aim to truncate a typically divergent power series to a finite series. In machine learning, these approaches have been addressed from an information-theoretic perspective by [42, 43]. Thouless–Anderson–Palmer equations [44] are a form of second-order perturbation theory. They were originally developed in statistical physics to include perturbative corrections to the mean-field solution of Ising models. They have been adopted into Bayesian inference in [32] and were advanced by many authors [15, 28, 30, 31]. In VI, perturbation theory yields extra terms to the mean-field variational objective which are difficult to calculate analytically. This may be a reason why the methods discussed are not widely adopted by practitioners. In this paper, we emphasize the ease of including perturbative corrections in a BBVI framework. Furthermore, in contrast to earlier formulations, our approach yields a strict lower bound to the marginal likelihood which can be conveniently optimized. Our approach is different from the traditional variational perturbation formulation [18, 41], which generally does not result in a bound.

6. Conclusion

We first presented a view on black box variational inference as a form of biased importance sampling, where we can trade-off bias versus variance by the choice of divergence. Bias refers to the deviation of the bound from the true marginal likelihood, and variance refers to its reparameterization gradient estimator. We then proposed a family of new variational bounds that connect to variational perturbation theory, and which include corrections to the standard Kullback–Leibler bound. Our proposed PBBVI bound converges to the true marginal likelihood for large order K of the perturbative expansion, and we showed both theoretically and experimentally that it has lower-variance reparameterization gradients compared to α -VI. In order to scale up our method to massive data sets, future work will explore stochastic versions of PBBVI. Since the PBBVI bound contains interaction terms between all data points, breaking it up into mini-batches is not straightforward. Besides, while our experiments used a fixed perturbative order of $K = 3$, it could be beneficial to increase the perturbative order at some point during the training cycle once an empirical estimate of the gradient variance drops below a certain threshold. It would also be interesting to investigate a K -independent formulation of PBBVI using Russian roulette estimates [22]. Furthermore, the PBBVI and α -bounds can also be combined, such that PBBVI further approximates α -VI. This could lead to promising results on large data sets where traditional α -VI is hard to optimize due to its variance, and traditional PBBVI

converges to KLVI. As a final remark, a tighter variational bound is not guaranteed to always result in a better posterior approximation since the variational family limits the quality of the solution. However, in the context of variational EM, where one performs gradient-based hyperparameter optimization on the log marginal likelihood, our bound gives more reliable results since higher orders of K can be assumed to approximate the marginal likelihood better.

References

- [1] Amari S 2012 *Differential-Geometrical Methods in Statistics* vol 28 (Berlin: Springer)
- [2] Bamler R and Mandt S 2017 Dynamic word embeddings *Int. Conf. on Machine Learning* pp 380–9
- [3] Bamler R and Mandt S 2018 Improving optimization in models with continuous symmetry breaking *Int. Conf. on Machine Learning* pp 432–41
- [4] Bamler R, Zhang C, Oppor M and Mandt S 2017 Perturbative black box variational inference *Advances in Neural Information Processing Systems* (Red Hook, NY: Curran Associates, Inc.) pp 5079–88
- [5] Blei D M, Kucukelbir A and McAuliffe J D 2017 Variational inference: a review for statisticians *J. Am. Stat. Assoc.* **112** 859–77
- [6] Buchholz A, Wenzel F and Mandt S 2018 Quasi-monte carlo variational inference *Int. Conf. on Machine Learning* pp 667–76
- [7] Burda Y, Grosse R and Salakhutdinov R 2016 Importance weighted autoencoders *Int. Conf. on Learning Representations* pp 1–9
- [8] Dempster A P, Laird N M and Rubin D B 1977 Maximum likelihood from incomplete data via the EM algorithm *J. R. Stat. Soc. Ser. B* **39** 1–22
- [9] Deng Z, Navarathna R, Carr P, Mandt S, Yue Y, Matthews I and Mori G 2017 Factorized variational autoencoders for modeling audience reactions to movies *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* pp 2577–86
- [10] Dieng A B, Tran D, Ranganath R, Paisley J and Blei D M 2017 Variational inference via χ upper bound minimization *Advances in Neural Information Processing Systems* (Red Hook, NY: Curran Associates, Inc.) pp 2732–41
- [11] Domke J and Sheldon D R 2018 Importance weighting and variational inference *Advances in Neural Information Processing Systems* (Red Hook, NY: Curran Associates, Inc.) pp 4470–9
- [12] Hernandez-Lobato J, Li Y, Rowland M, Bui T, Hernández-Lobato D and Turner R 2016 Black-box alpha divergence minimization *Int. Conf. on Machine Learning* pp 2256–73
- [13] Hoffman M D, Blei D M, Wang C and Paisley J W 2013 Stochastic variational inference *J. Mach. Learn. Res.* **14** 1303–47
- [14] Jordan M I, Ghahramani Z, Jaakkola T S and Saul L K 1999 An introduction to variational methods for graphical models *Mach. Learn.* **37** 183–233
- [15] Kappen H J and Wiegierinck W 2001 *Second Order Approximations for Probability Models* pp 238–44
- [16] Kingma D P and Welling M 2014 Auto-encoding variational Bayes *Int. Conf. on Learning Representations* pp 1–9
- [17] Kingma D P, Salimans T and Welling M 2015 Variational dropout and the local reparameterization trick *Advances in Neural Information Processing Systems* (Red Hook, NY: Curran Associates, Inc.) pp 2575–83
- [18] Kleinert H 2009 *Path Integrals in Quantum Mechanics, Statistics, Polymer Physics, and Financial Markets* (Singapore: World scientific)
- [19] LeCun Y, Bottou L, Bengio Y and Haffner P 1998 Gradient-based learning applied to document recognition *Proc. of the IEEE* **86** 2278–324
- [20] Li Y and Mandt S 2018 Disentangled sequential autoencoder *Int. Conf. on Machine Learning* pp 5656–65
- [21] Li Y and Turner R E 2016 Rényi divergence variational inference *Advances in Neural Information Processing Systems* (Red Hook, NY: Curran Associates, Inc.) pp 1073–81
- [22] Lyne A M *et al* 2015 On russian roulette estimates for bayesian inference with doubly-intractable likelihoods *Stat. Sci.* **30** 443–67
- [23] Ma C, Tschitschek S, Palla K, Lobato J M H, Nowozin S and Zhang C 2019 EDDI: efficient dynamic discovery of high-value information with partial VAE *Int. Conf. on Machine Learning* pp 1–8
- [24] Mandt S and Blei D 2014 Smoothed gradients for stochastic variational inference *Advances in Neural Information Processing Systems* (Red Hook, NY: Curran Associates, Inc.) pp 2438–46
- [25] Marino J, Yue Y and Mandt S 2018 Iterative amortized inference *Int. Conf. on Machine Learning* pp 3400–9

- [26] Mescheder L, Nowozin S and Geiger A 2017 Adversarial variational bayes: unifying variational autoencoders and generative adversarial networks *Int. Conf. on Machine Learning* pp 2391–400
- [27] Minka T 2005 Divergence measures and message passing *Technical Report* Microsoft Research
- [28] Oppor M 2015 Expectation propagation *Statistical Physics, Optimization, Inference and Message-Passing Algorithms* (Oxford: Oxford University Press) ch 9 pp 263–92
- [29] Oppor M, Fraccaro M, Paquet U, Susemihl A and Winther O 2015 Perturbation theory for variational inference *Advances in Neural Information Processing Systems Workshop on Advances in Approximate Bayesian Inference*
- [30] Oppor M, Paquet U and Winther O 2013 Perturbative corrections for approximate inference in gaussian latent variable models *J. Mach. Learn. Res.* **14** 2857–98
- [31] Paquet U, Winther O and Oppor M 2009 Perturbation corrections in approximate inference: mixture modelling applications *J. Mach. Learn. Res.* **10** 1263–304
- [32] Plefka T 1982 Convergence condition of the TAP equation for the infinite-ranged ising spin glass model *J. Phys. A: Math. Gen.* **15** 1971
- [33] Rainforth T, Kosiorek A, Le T A, Maddison C, Igl M, Wood F and Teh Y W 2018 Tighter variational bounds are not necessarily better *Int. Conf. on Machine Learning* pp 4277–85
- [34] Ranganath R, Gerrish S and Blei D M 2014 Black box variational inference *Int. Conf. on Artificial Intelligence and Statistics* pp 814–22
- [35] Ranganath R, Tang L, Charlin L and Blei D 2015 Deep exponential families *Artificial Intelligence and Statistics* pp 762–71
- [36] Rezende D J and Mohamed S 2015 Variational inference with normalizing flows *Int. Conf. on Machine Learning* pp 1530–8
- [37] Rezende D J, Mohamed S and Wierstra D 2014 Stochastic backpropagation and approximate inference in deep generative models *Int. Conf. on Machine Learning* pp 1278–86
- [38] Robbins H and Monro S 1951 A stochastic approximation method *Ann. Math. Stat.* **22** 400–7
- [39] Ruiz F, Titsias M and Blei D 2016 The generalized reparameterization gradient *Advances in neural information processing systems* pp 460–8
- [40] Salimans T and Knowles D A 2013 Fixed-form variational posterior approximation through stochastic linear regression *Bayesian Anal.* **8** 837–82
- [41] Schwartz M and Katzav E 2008 The ideas behind self-consistent expansion *J. Stat. Mech.* **23**
- [42] Tanaka T 1999 A theory of mean field approximation *Advances in Neural Information Processing Systems* (Cambridge, MA: MIT Press) pp 1–8
- [43] Tanaka T 2000 Information geometry of mean-field approximation *Neural Comput.* **12** 1951–68
- [44] Thouless D J, Anderson P W and Palmer R G 1977 Solution of solvable model of a spin glass *Phil. Mag.* **35** 593–601
- [45] Tucker G, Mnih A, Maddison C J, Lawson J and Sohl-Dickstein J 2017 Rebar: low-variance, unbiased gradient estimates for discrete latent variable models *Advances in Neural Information Processing Systems* (Red Hook, NY: Curran Associates, Inc.) pp 2627–36
- [46] Zhang C, Butepage J, Kjellstrom H and Mandt S 2018 Advances in variational inference *IEEE Trans. Pattern Anal. Mach. Intell.* **41** 1–20