
Conjugate-Computation Variational Inference : Converting Variational Inference in Non-Conjugate Models to Inferences in Conjugate Models

Mohammad Emtiyaz Khan

Center for Advanced Intelligence Project (AIP)
RIKEN, Tokyo

Wu Lin

Center for Advanced Intelligence Project (AIP)
RIKEN, Tokyo

Abstract

Variational inference is computationally challenging in models that contain both conjugate and non-conjugate terms. Methods specifically designed for conjugate models, even though computationally efficient, find it difficult to deal with non-conjugate terms. On the other hand, stochastic-gradient methods can handle the non-conjugate terms but they usually ignore the conjugate structure of the model which might result in slow convergence. In this paper, we propose a new algorithm called Conjugate-computation Variational Inference (CVI) which brings the best of the two worlds together – it uses conjugate computations for the conjugate terms and employs stochastic gradients for the rest. We derive this algorithm by using a stochastic mirror-descent method in the mean-parameter space, and then expressing each gradient step as a variational inference in a conjugate model. We demonstrate our algorithm’s applicability to a large class of models and establish its convergence. Our experimental results show that our method converges much faster than the methods that ignore the conjugate structure of the model.

1 Introduction

In this paper, we focus on designing efficient variational inference algorithms for models that contain both conjugate and non-conjugate terms, e.g., models such as Gaussian process classification (Kuss and Rasmussen, 2005), correlated topic models (Blei and Lafferty, 2007), exponential-family Probabilistic PCA (Mohamed et al., 2009), large-scale multi-class classification (Genkin et al.,

2007), Kalman filters with non-Gaussian likelihoods (Rue and Held, 2005), and deep exponential-family models (Ranganath et al., 2015). Such models are widely used in machine learning and statistics, yet variational inference on them remains computationally challenging.

The difficulty lies in the non-conjugate part of the model. In the traditional Bayesian setting, when the prior distribution is conjugate to the likelihood, the posterior distribution is available in closed-form and can be obtained through simple computations. For example, in a conjugate-exponential family, computation of the posterior distribution can be achieved by simply adding the sufficient statistics of the likelihood to the natural parameter of the prior. In this paper, we refer to such computations as *conjugate computations* (an example is included in the next section).

These types of conjugate computations have been used extensively in variational inference, primarily due to their computational efficiency. For example, the variational message-passing (VMP) algorithm proposed by Winn and Bishop (2005) uses conjugate computations within a message-passing framework. Similarly, stochastic variational inference (SVI) builds upon VMP and enables large-scale inference by employing stochastic methods (Hoffman et al., 2013).

Unfortunately, the computational efficiency of these methods is lost when the model contains non-conjugate terms. For example, the messages in VMP lose their convenient exponential-family form and become more complex as the algorithm progresses. Additional approximations for the non-conjugate terms can be used, e.g. those discussed by Winn and Bishop (2005) and Wang and Blei (2013), but such approximations usually result in a performance loss (Honkela and Valpola, 2004; Khan, 2012). Other existing alternatives, such as the non-conjugate VMP method of Knowles and Minka (2011) and the expectation-propagation method of Minka (2001), also require carefully designed quadrature methods to approximate the non-conjugate terms, and suffer from convergence problems and numerical issues.

Recently, many stochastic-gradient (SG) methods have

Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017, Fort Lauderdale, Florida, USA. JMLR: W&CP volume 54. Copyright 2017 by the author(s).

been proposed to deal with this issue (Ranganath et al., 2014; Salimans et al., 2013; Titsias and Lázaro-Gredilla, 2014). An advantage of these approaches is that they can be used as a black-box and applied to a wide-variety of inference problems. However, these methods usually do not directly exploit conjugacy, e.g., during the stochastic-gradient computation. This might lead to several issues, e.g., their updates may depend on the parameterization of the variational distribution, the number of variational parameters might be too large, and the updates might converge slowly.

In this paper, we propose an algorithm which brings the best of the two worlds together – it uses stochastic gradients for the non-conjugate terms, while retaining the computational efficiency of conjugate computations on the conjugate terms. We call our approach Conjugate-computation Variational Inference (CVI). Our main proposal is to use a stochastic *mirror-descent method* in the *mean-parameter space* which differs from many existing methods that use a stochastic *gradient-descent method* in the *natural-parameter space*. Our method has a natural advantage over these methods – gradient steps in our method can be implemented by using conjugate computations.

We demonstrate our approach on two classes of non-conjugate models. The first class contains models which can be split into a conjugate part and a non-conjugate part. For such models our gradient steps can be expressed as a Bayesian inference in a conjugate model. The second class of models additionally allows conditionally-conjugate terms. For this model class, our gradient steps can be written as a message passing algorithm where VMP or SVI is used for the conjugate part while stochastic gradients are employed for the rest. Our algorithm conveniently reduces to VMP when the model is conjugate.

We also prove convergence of our algorithm and establish its connections to many existing approaches. We apply our algorithm to many existing models and demonstrate that our updates can be implemented using variational inference in conjugate models. Empirical results on many models and datasets show that our method converges much faster than the methods that ignore the conjugate structure of the model. The code to reproduce results of this paper is available at <https://github.com/emtiyaz/cvi/>.

2 Conjugate Computations

Given a probabilistic graphical model $p(\mathbf{y}, \mathbf{z})$ with \mathbf{y} as the vector of observed variables and \mathbf{z} as the vector of latent variables, our goal in variational inference is to estimate a posterior distribution $p(\mathbf{z}|\mathbf{y})$. When the prior distribution $p(\mathbf{z})$ is *conjugate*¹ to the likelihood $p(\mathbf{y}|\mathbf{z})$, the

¹A prior distribution is said to be conjugate when it takes the same functional form as the likelihood. An exact definition is

posterior distribution is available in closed form and can be obtained through simple computations which we refer to as the *conjugate computations*. For example, consider the following exponential-family prior distribution: $p(\mathbf{z}) = h(\mathbf{z}) \exp[\langle \phi(\mathbf{z}), \boldsymbol{\eta} \rangle - A(\boldsymbol{\eta})]$, where $\boldsymbol{\eta}$ is the natural parameter, ϕ is the sufficient statistics, $\langle \cdot, \cdot \rangle$ is an inner product, $h(\mathbf{z})$ is the base measure, and $A(\boldsymbol{\eta})$ is the log-partition function. When the likelihood is conjugate to the prior, we can express the likelihood in the same form as the prior with respect to \mathbf{z} , as shown below:

$$p(\mathbf{y}|\mathbf{z}) = \exp[\langle \phi(\mathbf{z}), \boldsymbol{\eta}_{yz}(\mathbf{y}) \rangle - f_y(\mathbf{y})], \quad (1)$$

where $\boldsymbol{\eta}_{yz}$ and f_y are functions that depend on \mathbf{y} only. In such cases, the posterior distribution takes the same exponential form as $p(\mathbf{z})$ and its natural parameter can be obtained by simply adding the natural parameters $\boldsymbol{\eta}$ of the prior to the function $\boldsymbol{\eta}_{yz}(\mathbf{y})$ of the likelihood:

$$p(\mathbf{z}|\mathbf{y}) \propto h(\mathbf{z}) \exp[\langle \phi(\mathbf{z}), \boldsymbol{\eta} + \boldsymbol{\eta}_{yz}(\mathbf{y}) \rangle]. \quad (2)$$

This is a type of conjugate computation. Such conjugate computations are extensively used in Bayesian inference for conjugate models, as well as in variational inference for conditionally-conjugate models in algorithms such as variational message passing (Winn and Bishop, 2005) and expectation propagation (Minka, 2001).

3 Non-Conjugate Variational Inference

When the model also contains non-conjugate terms, variational inference becomes computationally challenging. In variational inference, we obtain a fixed-form variational approximation $q(\mathbf{z}|\boldsymbol{\lambda})$, where $\boldsymbol{\lambda}$ is the variational parameter, by maximizing a lower bound to the marginal likelihood:

$$\max_{\boldsymbol{\lambda} \in \Omega} \mathcal{L}(\boldsymbol{\lambda}) := \mathbb{E}_q[\log p(\mathbf{y}, \mathbf{z}) - \log q(\mathbf{z}|\boldsymbol{\lambda})], \quad (3)$$

where Ω is the set of valid variational parameters. Non-conjugate terms might make the lower-bound optimization challenging, e.g., by making it intractable. For example, Gaussian Process (GP) models usually employ a non-Gaussian likelihood and the variational lower bound becomes intractable, as discussed below.

GP Example: Consider a GP model for N input-output pairs $\{y_n, \mathbf{x}_n\}$ indexed by n . Let $z_n := f(\mathbf{x}_n)$ be the latent function drawn from a GP with mean 0 and covariance \mathbf{K} . Given z_n , we use a non-Gaussian likelihood $p(y_n|z_n)$ to model the output y_n . The joint distribution is shown below:

$$p(\mathbf{y}, \mathbf{z}) = \left[\prod_{n=1}^N p(y_n|z_n) \right] \mathcal{N}(\mathbf{z}|0, \mathbf{K}). \quad (4)$$

given in Appendix A.

It is a common practice to approximate the posterior distribution by a Gaussian distribution $q(\mathbf{z}|\boldsymbol{\lambda}) := \mathcal{N}(\mathbf{z}|\mathbf{m}, \mathbf{V})$ whose mean \mathbf{m} and covariance \mathbf{V} we need to estimate (Kuss and Rasmussen, 2005). By substituting the joint-distribution (4) in the lower bound (3), we get the following lower bound:

$$\sum_n \mathbb{E}_q [\log p(y_n|z_n)] - \mathbb{D}_{KL}[q(\mathbf{z}|\boldsymbol{\lambda}) \parallel \mathcal{N}(\mathbf{z}|0, \mathbf{K})], \quad (5)$$

where \mathbb{D}_{KL} is the Kullback-Leibler divergence. This lower bound is intractable for most *non-Gaussian* likelihoods because the expectation $\mathbb{E}_q [\log p(y_n|z_n)]$ usually does not have a closed-form expression, e.g., when $p(y_n|z_n)$ is a logistic or probit likelihood.

Despite its intractability, the lower bound can still be optimized by using a stochastic-gradient method, e.g., the following stochastic-gradient descent (SGD) algorithm:

$$\boldsymbol{\lambda}_{t+1} = \boldsymbol{\lambda}_t + \rho_t \widehat{\nabla}_{\boldsymbol{\lambda}} \mathcal{L}(\boldsymbol{\lambda}_t), \quad (6)$$

where t is the iteration number, ρ_t is a step size, and $\widehat{\nabla}_{\boldsymbol{\lambda}} \mathcal{L}(\boldsymbol{\lambda}_t) := \widehat{\partial} \mathcal{L} / \partial \boldsymbol{\lambda}$ is a stochastic gradient of the lower bound at $\boldsymbol{\lambda} = \boldsymbol{\lambda}_t$. The advantage of this approach is that it can be used as a black-box method and applied to a wide-variety of inference problems (Ranganath et al., 2014).

Despite its generality and scalability, there are major issues with the SG method. The conjugate terms in the lower bound might have a closed-form expression and may not require any stochastic approximations. A naive application of the SG method might ignore this. Another issue is that the efficiency and rate of convergence might depend on the parameterization used for the variational distribution $q(\mathbf{z}|\boldsymbol{\lambda})$. Some parameterizations might have simpler updates than others but it is usually not clear how to find the best one for a given model. We discuss these issues below for the GP example.

GP Example (issues with SGD): In GP models, it might seem that the number of variational parameters should be in $O(N^2)$, e.g., if we use $\{\mathbf{m}, \mathbf{V}\}$. However, as Opper and Archambeau (2009) show, there are only $O(N)$ free parameters. Therefore, choosing a naive parameterization might be an order of magnitude slower than the best option, e.g. the gradient updates with quadratic number of parameters will be much slower than the one using a linear number of parameters (see Appendix B for more details on inefficiency of SGD). In fact, as shown in Khan et al. (2012), choosing a good parameterization is a difficult problem in this case and a naive selection might make the problem more difficult than it needs to be.

Our algorithm, derived in the next sections, will fix these issues with SGD, e.g., for the GP example our algorithm will

conveniently express each gradient step as predictions in a GP regression model which naturally has an $O(N)$ number of free parameters. In the results section, we will see that this results in a fast convergent algorithm.

4 Conjugate-Computation Variational Inference (CVI)

We now derive the CVI algorithm that uses stochastic gradients for the non-conjugate terms, while retaining the computational efficiency of conjugate computations for the conjugate terms. We will use a stochastic *mirror-descent method* in the *mean-parameter space* and show that its gradient steps can be implemented by using conjugate computations. This will fix the issues of stochastic-gradient methods but maintain the computational efficiency of conjugate computations.

Our approach relies on the following two assumptions:

Assumption 1 [minimality] : *The variational distribution $q(\mathbf{z}|\boldsymbol{\lambda})$ is a “minimal” exponential-family distribution:*

$$q(\mathbf{z}|\boldsymbol{\lambda}) = h(\mathbf{z}) \exp \{ \langle \phi(\mathbf{z}), \boldsymbol{\lambda} \rangle - A(\boldsymbol{\lambda}) \}, \quad (7)$$

with $\boldsymbol{\lambda}$ as its natural parameters.

The minimal² representation implies that there is a one to one mapping between the mean parameter $\boldsymbol{\mu} := \mathbb{E}_q[\phi(\mathbf{z})]$ and the natural parameter $\boldsymbol{\lambda}$. Therefore, we can express the lower-bound optimization as a maximization problem over $\boldsymbol{\mu} \in \mathcal{M}$, where \mathcal{M} is the set of valid mean parameters. We denote the new objective function by $\widehat{\mathcal{L}}(\boldsymbol{\mu}) := \mathcal{L}(\boldsymbol{\lambda})$.

Assumption 2 [conjugacy] : *We assume that the joint distribution contains some terms, collectively denoted by \tilde{p}_c , which take the same form as q with respect to \mathbf{z} , i.e.,*

$$\tilde{p}_c(\mathbf{y}, \mathbf{z}) \propto h(\mathbf{z}) \exp \{ \langle \phi(\mathbf{z}), \boldsymbol{\eta} \rangle \}, \quad (8)$$

where $\boldsymbol{\eta}$ is a known parameter vector. We call \tilde{p}_c as the conjugate part of the model. We denote the non-conjugate terms by \tilde{p}_{nc} giving us the following partitioning of the joint distribution: $p(\mathbf{y}, \mathbf{z}) \propto \tilde{p}_{nc}(\mathbf{y}, \mathbf{z}) \tilde{p}_c(\mathbf{y}, \mathbf{z})$. These terms can be unnormalized with respect to \mathbf{z} .

We can always satisfy this assumption, e.g., by trivially setting $\boldsymbol{\eta} = 0$ and $\tilde{p}_{nc} = p(\mathbf{y}, \mathbf{z})/h(\mathbf{z})$. However, since there is no conjugacy in this formulation, our algorithm might not be able gain additional computational efficiency over the SG methods.

We now derive the CVI algorithm. We build upon an equivalent formulation of (6) which expresses the gradient step as the maximization of a local approximation:

$$\boldsymbol{\lambda}_{t+1} = \arg \max_{\boldsymbol{\lambda} \in \Omega} \left\langle \boldsymbol{\lambda}, \widehat{\nabla}_{\boldsymbol{\lambda}} \mathcal{L}(\boldsymbol{\lambda}_t) \right\rangle - \frac{1}{2\rho_t} \|\boldsymbol{\lambda} - \boldsymbol{\lambda}_t\|_2^2, \quad (9)$$

²A summary of exponential family is given in Appendix C.

where $\|\cdot\|_2$ is the Euclidean norm and Ω is the set of valid natural parameters. By taking the derivative and setting it to zero, we recover the SGD update of (6) which establishes the equivalence.

Instead of using the above SGD update in the natural-parameter space, we propose to use a stochastic mirror-descent update in the mean-parameter space. The mirror-descent algorithm (Nemirovskii et al., 1983) replaces the Euclidean geometry in (9) by a *proximity* function such as a Bregman divergence (Raskutti and Mukherjee, 2015). We propose the following mirror-descent algorithm:

$$\boldsymbol{\mu}_{t+1} = \arg \max_{\boldsymbol{\mu} \in \mathcal{M}} \left\langle \boldsymbol{\mu}, \widehat{\nabla}_{\boldsymbol{\mu}} \widetilde{\mathcal{L}}(\boldsymbol{\mu}_t) \right\rangle - \frac{1}{\beta_t} \mathbb{B}_{A^*}(\boldsymbol{\mu} \parallel \boldsymbol{\mu}_t), \quad (10)$$

where $A^*(\boldsymbol{\mu})$ is the convex-conjugate³ of the log-partition function $A(\boldsymbol{\lambda})$, \mathbb{B}_{A^*} is the Bregman divergence defined by A^* over \mathcal{M} , and $\beta_t > 0$ is the step-size. The following theorem establishes that (10) can be implemented by using a Bayesian inference in a conjugate model.

Theorem 1. *Under Assumption 1 and 2, the update (10) is equivalent to the Bayesian inference in the following conjugate model:*

$$q(\mathbf{z} \mid \boldsymbol{\lambda}_{t+1}) \propto e^{\langle \phi(\mathbf{z}), \widetilde{\boldsymbol{\lambda}}_t \rangle} \tilde{p}_c(\mathbf{y}, \mathbf{z}), \quad (11)$$

whose natural parameter can be obtained by conjugate computation: $\boldsymbol{\lambda}_{t+1} = \widetilde{\boldsymbol{\lambda}}_t + \boldsymbol{\eta}$ where $\widetilde{\boldsymbol{\lambda}}_t$ is the natural parameters of the exponential-family approximation to \tilde{p}_{nc} . The natural parameter $\boldsymbol{\lambda}_{t+1}$ can itself be obtained recursively by using stochastic gradients of \tilde{p}_{nc} :

$$\widetilde{\boldsymbol{\lambda}}_t = (1 - \beta_t) \widetilde{\boldsymbol{\lambda}}_{t-1} + \beta_t \widehat{\nabla}_{\boldsymbol{\mu}} \mathbb{E}_q [\log \tilde{p}_{nc}] \Big|_{\boldsymbol{\mu}=\boldsymbol{\mu}_t}, \quad (12)$$

with $\boldsymbol{\lambda}_0 = 0$ and $\boldsymbol{\lambda}_1 = \boldsymbol{\eta}$.

A proof is given in Appendix D. The update (11) replaces the non-conjugate term by an exponential-family approximation whose natural parameter $\widetilde{\boldsymbol{\lambda}}_t$ is a weighted sum of the gradients of the non-conjugate term \tilde{p}_{nc} . The resulting algorithm, which we refer to as the conjugate-computation variational inference (CVI) algorithm, is summarized in Algorithm 1. As desired, our algorithm computes stochastic-gradients only for the non-conjugate terms, as shown in Step 3. Given this gradient, in Step 4, the new variational parameter is obtained by using a conjugate computation by simply adding the natural parameters.

Note that, even though we proposed an update in the mean-parameter space, conjugate computations in Step 4 are performed in the natural-parameter space. The mean parameter is required only during the computation of stochastic gradients in Step 3. For the GP example, these updates are conveniently expressed as predictions in a GP regression model which naturally has an $O(N)$ number of free parameters, as discussed next.

³Definitions of convex-conjugate and Bregman divergence is given in Appendix C.

Algorithm 1 CVI for exponential-family approximations.

- 1: Initialize $\widetilde{\boldsymbol{\lambda}}_0 = 0$ and $\boldsymbol{\lambda}_1 = \boldsymbol{\eta}$.
 - 2: **for** $t = 1, 2, 3, \dots$ **do**
 - 3: $\widetilde{\boldsymbol{\lambda}}_t = (1 - \beta_t) \widetilde{\boldsymbol{\lambda}}_{t-1} + \beta_t \widehat{\nabla}_{\boldsymbol{\mu}} \mathbb{E}_q [\log \tilde{p}_{nc}] \Big|_{\boldsymbol{\mu}=\boldsymbol{\mu}_t}$.
 - 4: $\boldsymbol{\lambda}_{t+1} = \widetilde{\boldsymbol{\lambda}}_t + \boldsymbol{\eta}$.
 - 5: **end for**
-

GP Example (CVI updates): For the GP model, the non-conjugate part \tilde{p}_{nc} is $\prod_n p(y_n | z_n)$. Both Assumption 1 and 2 are satisfied since q is a Gaussian and the GP prior is conjugate to it.

For Step 3, we need to compute the gradient with respect to $\boldsymbol{\mu}$. Since \tilde{p}_{nc} factorizes, we can compute the gradient of each term $\mathbb{E}_q [\log p(y_n | z_n)]$ separately. This term depends on the marginal $q(z_n)$ which has two mean parameters: $\mu_n^{(1)} := m_n$ and $\mu_n^{(2)} := V_{nn} + m_n^2$, where V_{nn} is the n 'th diagonal element of \mathbf{V} . The gradients can be computed using the Monte Carlo (MC) approximation as shown by Opper and Archambeau (2009) (details are given in Appendix B.1). Let's denote these gradients at iteration t by $\hat{g}_{n,t}^{(1)}$ and $\hat{g}_{n,t}^{(2)}$. Using these, Step 3 of Algorithm 1 can be written as follows:

$$\widetilde{\boldsymbol{\lambda}}_{n,t}^{(i)} = (1 - \beta_t) \widetilde{\boldsymbol{\lambda}}_{n,t-1}^{(i)} + \beta_t \hat{g}_{n,t}^{(i)}, \quad (13)$$

for $i = 1, 2$ and $n = 1, 2, \dots, N$. These are the natural parameters for a Gaussian approximation of $p(y_n | z_n)$. Using them in (11) we obtain the following conjugate model:

$$q(\mathbf{z} \mid \boldsymbol{\lambda}_{t+1}) \propto \left[\prod_{n=1}^N e^{z_n \widetilde{\boldsymbol{\lambda}}_{n,t}^{(1)} + z_n^2 \widetilde{\boldsymbol{\lambda}}_{n,t}^{(2)}} \right] \mathcal{N}(\mathbf{z} \mid \mathbf{0}, \mathbf{K}).$$

This update can be done by using a conjugate computation which in this case corresponds to predictions in a GP regression model (since we only need to compute the mean parameter $\mu_n^{(1)}$ and $\mu_n^{(2)}$ for all n). We see that the only free parameters to be computed are $\widetilde{\boldsymbol{\lambda}}_{n,t}^{(1)}$ and $\widetilde{\boldsymbol{\lambda}}_{n,t}^{(2)}$, therefore the number of parameters is in $O(N)$. We naturally end up with the optimal number of parameters and avoid the computation of the full covariance matrix \mathbf{V} . Both of these give a huge computational saving over a naive SGD method.

The previous example shows that variational inference in non-conjugate GP models can be done by solving a sequence of GP regression problems. In Appendix E, we give many more such examples. In Appendix E.1, we show that the variational inference in a generalized linear model (GLM) can be implemented by using Bayesian inference in linear regression model. Similarly, in Appendix E.2, we show a Kalman filter model with GLM likelihood can be

implemented by using Bayesian inference in the standard Kalman filter. We also give examples involving Gamma variational distribution in Appendix E.3.

It is also possible to use a “doubly” stochastic approximation (Titsias and Lázaro-Gredilla, 2014) where, in addition to the MC approximation, we also randomly pick factors from both \tilde{p}_{nc} and \tilde{p}_c . As discussed below, this will result in a huge reduction in computation in Step 3 of Algorithm 1, e.g., bringing the number of stochastic-gradient computations to $O(1)$ from $O(N)$ in the GP example.

GP Example (doubly-stochastic CVI): We can use a doubly-stochastic scheme over \tilde{p}_{nc} since it factorizes over n . We sample one term (or pick a mini-batch) and compute its stochastic gradient. In our algorithm, this translates to modifying only the selected example’s $\tilde{\lambda}_{n,t}^{(i)}$. Denoting the selected sample index by n_t , this can be expressed as follows:

$$\tilde{\lambda}_{n,t}^{(i)} = (1 - \beta_t)\tilde{\lambda}_{n,t-1}^{(i)} + \delta_{n=n_t} \beta_t \hat{g}_{n,t}^{(i)}, \quad (14)$$

where $\delta_{n=n_t}$ is an indicator function which is 1 only when $n = n_t$. The number of stochastic gradient computation is therefore in $O(1)$ instead of $O(N)$. Computation is further reduced since at each iteration t we only need to compute *one* mean parameter corresponding to the marginal of z_{n_t} . This is much more efficient than a SGD update where we have to explicitly store \mathbf{V}_t . We can also reuse computations between iterations since updates in iteration t differs from iteration $t - 1$ only at one example $n = n_t$.

Another attractive property of our algorithm is that it can handle constraints on the variational parameters without much computational overhead, e.g., the covariance matrix in GP example will be positive definite as long as the Gaussian approximation of all non-conjugate term is valid. We can always make sure that this is the case by using rejection sampling inside the stochastic approximation.

One requirement for our algorithm is that we should be able to compute stochastic gradients with respect to the mean parameter. For distributions such as Gaussian and Multinomial, this can be done analytically, e.g. see Appendix B.1 for Gaussian distribution. For other distributions, such as Gamma, this quantity might be difficult to compute directly. We propose to build stochastic approximations by using the following identity:

$$\frac{\partial f}{\partial \boldsymbol{\mu}} = \left[\frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\lambda}} \right]^{-1} \frac{\partial f}{\partial \boldsymbol{\lambda}} = \mathbf{C}_\lambda^{-1} \frac{\partial f}{\partial \boldsymbol{\lambda}}, \quad (15)$$

where \mathbf{C}_λ is the Fisher information matrix and f is the function whose gradient we want to compute. We compute stochastic approximations of \mathbf{C}_λ and $\partial f / \partial \boldsymbol{\lambda}$ separately, and then solve the equation to get the gradients. More details are given in Appendix F. Our proposal is very similar

to the one discussed by Salimans et al. (2013) where the approximation is averaging over iterations as well. The advantage of our proposal is that we do not have to explicitly store or form the Fisher information matrix, rather only solve a linear system.

The convergence of our algorithm is guaranteed under mild conditions discussed in Khan et al. (2016). The update (10) converges to a local optimum of the lower bound $\mathcal{L}(\boldsymbol{\lambda})$ under the following conditions: $\tilde{\mathcal{L}}(\boldsymbol{\mu})$ is differentiable and its gradient is L -Lipschitz-continuous, the stochastic approximation is unbiased and has bounded variance, and the function $A^*(\boldsymbol{\mu})$ is continuously differentiable and strongly convex with respect to the L_2 norm. An exact statement of convergence is given in Proposition 3 of Khan et al. (2016).

5 CVI for Mean-Field Approximation

We now derive our algorithm for mean-field approximations on models that also allow *conditional-conjugacy*. This class of models is bigger than the one considered in the previous section, however the approximation is restricted to mean-field which is a stricter assumption than the one used in the previous section.

Consider a Bayesian network over $\mathbf{x} = \{\mathbf{y}, \mathbf{z}\}$ where \mathbf{y} is the vector of observed nodes \mathbf{y}_n for $n = 1, 2, \dots, N$, and \mathbf{z} is the vector of latent nodes \mathbf{z}_i for $i = 1, 2, \dots, M$. The joint distribution over \mathbf{x} is given as follows:

$$p(\mathbf{x}) = \prod_{a=1}^{M+N} p(\mathbf{x}_a | \mathbf{x}_{pa_a}), \quad (16)$$

where \mathbf{x}_{pa_a} is the set of parent nodes for variable \mathbf{x}_a . We will refer to a term $p(\mathbf{x}_a | \mathbf{x}_{pa_a})$ in $p(\mathbf{x})$ as the factor a .

Similar to the previous section, we make the following two assumptions for the CVI algorithm derived in this section.

Assumption 3 [mean-field + minimality] : *We assume that $q(\mathbf{z}) = \prod_i q_i(\mathbf{z}_i)$ with each factor being a minimal exponential-family distribution:*

$$q_i(\mathbf{z}_i | \boldsymbol{\lambda}_i) := h_i(\mathbf{z}_i) \exp [\langle \phi_i(\mathbf{z}_i), \boldsymbol{\lambda}_i \rangle - A_i(\boldsymbol{\lambda}_i)]. \quad (17)$$

We denote the vector of mean parameters $\boldsymbol{\mu}_i$ by $\boldsymbol{\mu}$ and the vector of natural parameters $\boldsymbol{\lambda}_i$ by $\boldsymbol{\lambda}$. Due to minimality, we can rewrite the lower bound in terms of $\boldsymbol{\mu}$, for which we use the same notation $\tilde{\mathcal{L}}(\boldsymbol{\mu})$ as in the previous section.

For the next assumption, we define \mathbb{N}_i to be the set containing the node \mathbf{z}_i and all its children. We define $\mathbf{x}_{/i}$ to be the set of all nodes \mathbf{x} except \mathbf{z}_i . Similarly, given a factor $p(\mathbf{x}_a | \mathbf{x}_{pa_a})$ for node $\mathbf{x}_a \in \mathbb{N}_i$, we define $\mathbf{x}_{a/i}$ to be the set of all nodes in the set $\mathbf{x}_a \cup \mathbf{x}_{pa_a}$ except \mathbf{z}_i .

Assumption 4 [conditional-conjugacy] : *For each node \mathbf{z}_i , we can split the following conditional distribution into*

a conjugate and a non-conjugate term as shown below:

$$\begin{aligned} p(\mathbf{z}_i|\mathbf{x}_{/i}) &\propto \prod_{a \in \mathbb{N}_i} p(\mathbf{x}_a|\mathbf{x}_{pa_a}) \\ &\propto h_i(\mathbf{z}_i) \prod_{a \in \mathbb{N}_i} \tilde{p}_{nc}^{a,i}(\mathbf{z}_i, \mathbf{x}_{a/i}) e^{\{\phi_i(\mathbf{z}_i), \boldsymbol{\eta}_{a,i}(\mathbf{x}_{a/i})\}}, \end{aligned} \quad (18)$$

where $\tilde{p}_{nc}^{a,i}$ is the non-conjugate part and $\boldsymbol{\eta}_{a,i}(\mathbf{x}_{a/i})$ is the natural parameter of the conjugate part for the factor a .

Similar to Assumption 2, we can always satisfy this assumption, but this may or may not guarantee the usefulness of our method.

Similar to the previous section, we can reparameterize the lower bound in terms of $\boldsymbol{\mu}$ to define $\tilde{\mathcal{L}}(\boldsymbol{\mu}) := \mathcal{L}(\boldsymbol{\lambda})$ and then use the update (10). Due to the mean-field approximation and linearity of the first term in (10), we can conveniently rewrite the objective as a sum over all nodes i :

$$\max_{\boldsymbol{\mu}} \sum_{i=1}^M \left[\left\langle \boldsymbol{\mu}_i, \widehat{\nabla}_{\boldsymbol{\mu}_i} \tilde{\mathcal{L}}(\boldsymbol{\mu}_t) \right\rangle - \frac{1}{\beta_t} \mathbb{B}_{A^*}(\boldsymbol{\mu}_i \| \boldsymbol{\mu}_{i,t}) \right], \quad (19)$$

where $\boldsymbol{\mu}_t$ and $\boldsymbol{\mu}_{i,t}$ denotes the value of $\boldsymbol{\mu}$ and $\boldsymbol{\mu}_i$, respectively, at iteration t . Therefore, we can either optimize all $\boldsymbol{\mu}_i$ parallelly or use a doubly-stochastic scheme by randomly picking a term in the sum.

The final algorithm is shown in Algorithm 2 and a detailed derivation is given in Appendix G. In Step 4, when combining all the messages received at a node i , the algorithm separates the conjugate computations from the non-conjugate ones. The first set of messages $\tilde{\boldsymbol{\eta}}_{a,i}$ (defined below) are obtained from the conjugate parts by taking the expectation over their natural parameters $\boldsymbol{\eta}_{a,i}(\mathbf{x}_{a/i})$:

$$\tilde{\boldsymbol{\eta}}_{a,i} := \mathbb{E}_{q_{/i,t}} \left[\boldsymbol{\eta}_{a,i}(\mathbf{x}_{a/i}) \right], \quad (20)$$

where $q_{/i,t}$ is the variational distribution at iteration t of all the nodes except \mathbf{z}_i . By comparing the above to Equation 17 in Winn and Bishop (2005), we can see that this operation is equivalent to a message-passing step in the VMP algorithm.

The second set of messages (the second term inside the summation) in Step 4 is simply the stochastic-gradient of the non-conjugate term in factor a . The two sets of messages are combined to get the resulting natural parameter. Finally, a convex combination is taken in Step 5 to get the natural parameter of $q_{i,t+1}$.

It is straightforward to see that in the absence of the second set of messages, our algorithm will reduce to VMP if we use a sequential or parallel updating scheme. However, an attractive property of our formulation is that we can also employ a doubly-stochastic update – we can randomly sample terms from (19), weight them appropriately to get an unbiased stochastic approximation to the gradient, and then take a step. This will correspond to updating

Algorithm 2 CVI for mean-field

- 1: Initialize $\boldsymbol{\lambda}_{i,0}$.
 - 2: **for** $t = 0, 1, 2, 3, \dots$, **do**
 - 3: **for** for all node \mathbf{z}_i (or a randomly sampled one) **do**
 - 4: $\tilde{\boldsymbol{\lambda}}_{i,t} = \sum_{a \in \mathbb{N}_i} \left[\tilde{\boldsymbol{\eta}}_{ai} + \widehat{\nabla}_{\boldsymbol{\mu}_i} \mathbb{E}_{q_t}(\log \tilde{p}_{nc}^{a,i}) |_{\boldsymbol{\mu}=\boldsymbol{\mu}_t} \right]$
 - 5: $\boldsymbol{\lambda}_{i,t+1} = (1 - \beta_t)\boldsymbol{\lambda}_{i,t} + \beta_t \tilde{\boldsymbol{\lambda}}_{i,t}$.
 - 6: **end for**
 - 7: **end for**
-

only a mini-batch of nodes in each iteration. If we use this type of updates on a conjugate-exponential model, our algorithm will be equivalent to SVI (given that we have local and global nodes and that we sample a local node followed by an update of the global node).

Convergence of Algorithm 2 is assured under the same conditions discussed earlier. Since the objective (19) can be expressed as a sum over all the nodes, our method converges under both stochastic updates (e.g. SVI like updates) and parallel updates (e.g. with one step of VMP). Note that we do *not* have to run these “inner loops” to convergence unlike many existing methods.

6 Related Work

One of the simplest method is to use local variational approximations to approximate the non-conjugate terms (Jaakkola and Jordan, 1996; Bouchard, 2007; Khan et al., 2010; Wang and Blei, 2013). Such approximations do not necessarily converge to a local maximum of the lower bound, leading to a performance loss (Kuss and Rasmussen, 2005; Marlin et al., 2011; Knowles and Minka, 2011; Khan, 2012). In contrast, our algorithm uses a stochastic-gradient step which is guaranteed to converge to the stationary point of the variational lower bound \mathcal{L} .

Another related approach is the Expectation-Propagation (EP) algorithm (Minka, 2001), which computes an exponential-family approximation (also called the site parameters) to the non-conjugate factors. The site parameter is very similar to $\tilde{\boldsymbol{\lambda}}_t$ in our algorithm, although our approximation is obtained by maximizing the lower bound unlike EP which uses moment matching. EP suffers from numerical issues and requires damping to ensure convergence, while our method has convergence guarantees.

The Non-conjugate variational message-passing (NC-VMP) algorithm (Knowles and Minka, 2011) is a variant of VMP for multinomial and binary regression. We can show that NC-VMP is a special case of our method under these conditions: gradients are exact and the step-size $\beta_t = 1$ (a formal proof is given in Appendix G.1). Therefore, our method is a stochastic version of NC-VMP with a principled way to do damping. Knowles and Minka (2011) also used damping in their experiment, although it was used as

a trick to make the method work.

Another related method is proposed by Salimans et al. (2013). They view the optimality condition as an instance of a linear regression and propose a stochastic-optimization method to solve it. This requires a stochastic estimate of the Fisher information matrix and the following gradient $\mathbf{g}_\lambda := \nabla_\lambda \mathbb{E}_q[\log p(\mathbf{y}, \mathbf{z})]$ with respect to the natural parameter. Denoting these two quantities at iteration t by $\widehat{\mathbf{C}}_{\lambda,t}$ and $\widehat{\mathbf{g}}_{\lambda,t}$, they do the following update: $\lambda_{t+1} = \widehat{\mathbf{C}}_{\lambda,t}^{-1} \widehat{\mathbf{g}}_{\lambda,t}$. By comparing this update to (15), we can see that the quantity in the right hand side of this update is similar to the gradient with respect to μ . However, in their method, the two stochastic gradients are maintained and averaged over iterations. Therefore, they need to store the Fisher information matrix explicitly which might be infeasible when the number of variational parameters is large (e.g. the GP model). We do not have this problem, because these gradients are required only when the gradient with respect to μ is not easy to compute directly, and can be computed on the fly at every iteration.

Our method is closely related the two existing works by Khan et al. (2015) and Khan et al. (2016) which use proximal-gradient methods for variational inference. Both of these works propose a splitting of the lower bound which is then optimized by using proximal-gradient methods in the natural-parameter space. Their update however does not directly correspond to an update in conjugate models, even though sometimes they can be obtained in closed-form. In contrast, we propose mirror-descent *without* any splitting and still obtain a closed-form update. In addition, our update corresponds to an update in a conjugate model. The key idea behind is to optimize in the mean-parameter space. Khan et al. (2016) speculate that their method could be generalized to a larger class of models. Our method fills this gap and derives a generalization to exponential-family models. Overall, our method is a subclass of the proximal-gradient methods discussed in Khan et al. (2016), but it provides a simpler way of applying it to non-conjugate exponential-family models.

7 Results

We present results on the following four models: Bayesian logistic regression, gamma factor model, gamma matrix factorization, and Gaussian process (GP) classification. Due to space constraints, details of the datasets are given in Appendix H. Additional results on Bayesian logistic regression and GP classification are in Appendix I.1.

We first discuss results for Bayesian logistic regression. We compare our method to the following four methods: explicit optimization with LBFGS method using Cholesky factorization and exact gradients (‘Chol’), proximal-gradient algorithm of Khan et al. (2015) (PG-

exact), Algorithm 2 of Salimans et al. (2013) (‘S&K-Alg2’), and Factorized-Gradient method of Salimans et al. (2013) (‘S&K-FG’).

The ‘Chol’ method does not exploit the structure of the problem and we expect it to be slow. The ‘S&K-FG’ works better than ‘S&K-Alg2’ when $D > N$ where N is the number of examples and D is the dimensionality. This is because S&K-Alg2 maintains an estimate of the Fisher information matrix which slows it down when D is large. However, the order is switched when $N > D$ with S&K-Alg2 performing better than S&K-FG. For these two methods we use the code provided by the authors.

Settings of various algorithmic parameters for these algorithms is given in Appendix I. For our method, we use CVI algorithm with stochastic gradient obtained using Monte Carlo (‘CVI’). Comparison to our algorithm with exact gradients is given included in Appendix I. Details of CVI update are given in Appendix E.1.

We compare the negative of the lower bound on the training set and the log-loss on the test set. The latter is computed as follows: given a test data with label $y_n \in \{0, 1\}$ and the estimate of $\hat{p}_n := \mathbb{E}_q[p(y_n = 1|\mathbf{z})]$ obtained by using a method, the log-loss is equal to the negative of $[y_n \log_2 \hat{p}_n + (1 - y_n) \log_2(1 - \hat{p}_n)]$. We report the average of the log-loss over all test points y_n . A value of 1 is equal to the performance of a random coin-toss, and any value lower than that is an improvement.

Figure (b) shows results on the ‘covtype’ dataset with $N = 581\text{K}$ and $D = 54$. The markers are drawn after iteration 1, 2, 3, 5, 8, and 10. We use 290K examples for training and leave the rest for testing. Chol is slowest as expected. Since $N > D$ for this dataset, S&K-Alg2 is faster than S&K-FG. CVI is as fast as S&K-Alg2 and PG-exact.

Figure (c) shows results on the Colon-Cancer dataset where $D > N$ ($N = 62$ and $D = 2000$). We use 31 observations for training. The markers are drawn after iteration 1, 2, 3, 5, 8, and 10. We can see that the S&K-Alg2 is much slower now, while S&K-FG is much faster. This is because the former computes and stores an estimate of the Fisher information matrix, which slows it down. In our approach, we completely avoid this computation and directly estimate the gradient of the mean (using Appendix B.1), that is why we are as fast as S&K-FG.

Overall, we see that CVI works well in both $N > D$ and $D > N$ regimes. The main reason behind is that CVI has closed-form updates which enables the application of matrix-inversion lemma to perform fast conjugate computations (Bayesian linear regression in this case).

Figure (c) shows the results for the gamma factor model discussed in Knowles (2015) (with $N = 522\text{K}$ and $D = 40$). Details of the model are given in Appendix J. We compare CVI to the algorithm proposed by Knowles (2015).

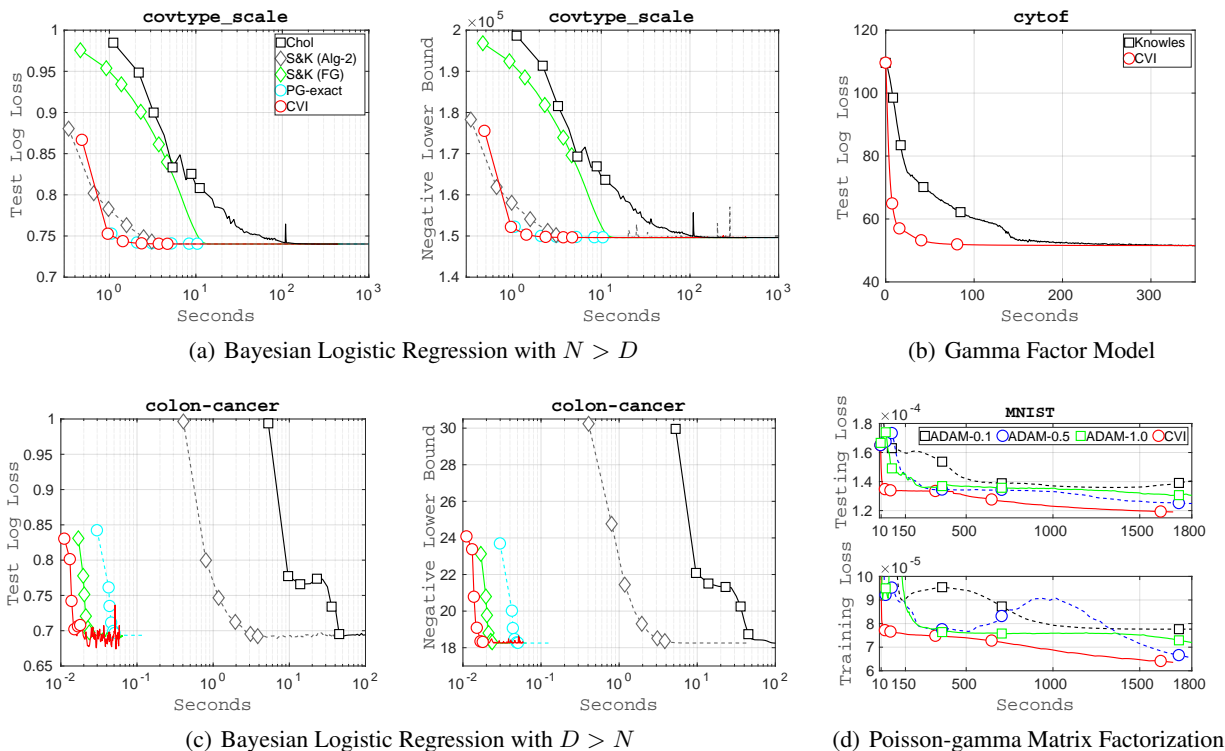


Figure 1: Comparison on Bayesian logistic regression, gamma factor model, and Poisson-gamma matrix factorization.

We use a constant step-size for CVI. For stochastic gradient computations, we use the implementation provided by the author. We use 40 latent factors and a fixed noise variance of 0.1. We compare the perplexity (average negative log-likelihood over test data points) using 2000 MC samples. Markers are shown after iteration 0, 9, 19, 49, and 99. Our method converges faster than the baseline, while achieving the same error.

Finally, Figure (d) shows the results for the Poisson-gamma matrix factorization model (Ranganath et al., 2015). Details are in Appendix K. We use the MNIST dataset of 70K images, each with 784 pixels. We use the provided train and test split, where 60K images are used for training. We use 100 latent factors and fixed model hyper-parameters. For the baseline, we use ADAM with the following transformation $\lambda = \log(1 + \exp(\lambda'))$ to satisfy the positivity constraint. We use many initial step-sizes for ADAM shown in Figure (d) (see ‘ADAM-x’ where x denotes the initial step-size). For CVI algorithm, a constant step size is used. For CVI, stochastic gradients sometimes violate the constraints (because of the reparameterization trick). To deal with this, we shrink the step size such that the steps are just within the constraints (similar to a method used in Khan et al. (2013)). Computation of stochastic gradients is based on the method of Knowles (2015) and is the same for both methods. We report the following reconstruction loss: $\sqrt{\sum_{ij}(y_{ij} - \hat{y}_{ij})^2 / (V \times N)}$ where N de-

notes the number of images and V is the number of pixels. Markers are drawn after iteration 50, 100, 500, 1000, and 2500. Our method converges faster than the baseline and also achieves a lower error.

8 Conclusions

In this paper, we proposed a new algorithm called the Conjugate-computation Variation Inference (CVI) algorithm for variational inference for models that contain both conjugate and non-conjugate terms. Our method is based on a stochastic mirror descent method in mean-parameter space. Each gradient step of our method can be expressed as a variational inference in a conjugate model. This leads to computationally efficient algorithm where stochastic gradients are employed only for the non-conjugate parts of the model, while using conjugate computations for the rest. Overall, CVI provides a general, modular, computationally efficient, scalable, and convergent approach for variational inference in non-conjugate models.

Acknowledgments: We would like to thank the anonymous reviewers for their feedback. Part of this work was done when MEK was a scientist in EPFL (Switzerland) and WL was a freelance researcher in Toronto (Canada).

References

- Blei, D. M. and Lafferty, J. D. (2007). A correlated topic model of science. *The Annals of Applied Statistics*, pages 17–35.
- Bouchard, G. (2007). Efficient bounds for the softmax and applications to approximate inference in hybrid models. In *NIPS 2007 Workshop on Approximate Inference in Hybrid Models*.
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2014). *Bayesian data analysis*, volume 2. Chapman & Hall/CRC Boca Raton, FL, USA.
- Genkin, A., Lewis, D. D., and Madigan, D. (2007). Large-scale bayesian logistic regression for text categorization. *Technometrics*, 49(3):291–304.
- Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. (2013). Stochastic variational inference. *The Journal of Machine Learning Research*, 14(1):1303–1347.
- Honkela, A. and Valpola, H. (2004). Unsupervised variational Bayesian learning of nonlinear models. In *Advances in Neural Information Processing Systems*, pages 593–600.
- Jaakkola, T. and Jordan, M. (1996). A variational approach to Bayesian logistic regression problems and their extensions. In *International conference on Artificial Intelligence and Statistics*.
- Khan, M. E. (2012). *Variational Learning for Latent Gaussian Models of Discrete Data*. PhD thesis, University of British Columbia.
- Khan, M. E., Aravkin, A. Y., Friedlander, M. P., and Seeger, M. (2013). Fast dual variational inference for non-conjugate latent Gaussian models. In *International Conference on Machine Learning*.
- Khan, M. E., Babanezhad, R., Lin, W., Schmidt, M., and Sugiyama, M. (2016). Faster Stochastic Variational Inference using Proximal-Gradient Methods with General Divergence Functions. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*.
- Khan, M. E., Baque, P., Flueret, F., and Fua, P. (2015). Kullback-Leibler Proximal Variational Inference. In *Advances in Neural Information Processing Systems*.
- Khan, M. E., Marlin, B., Bouchard, G., and Murphy, K. (2010). Variational Bounds for Mixed-Data Factor Analysis. In *Advances in Neural Information Processing Systems*.
- Khan, M. E., Mohamed, S., and Murphy, K. (2012). Fast Bayesian inference for non-conjugate Gaussian process regression. In *Advances in Neural Information Processing Systems*.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.
- Knowles, D. A. (2012). *Bayesian non-parametric models and inference for sparse and hierarchical latent structure*. PhD thesis, University of Cambridge.
- Knowles, D. A. (2015). Stochastic gradient variational Bayes for gamma approximating distributions. *arXiv preprint arXiv:1509.01631*.
- Knowles, D. A. and Minka, T. (2011). Non-conjugate variational message passing for multinomial and binary regression. In *Advances in Neural Information Processing Systems*, pages 1701–1709.
- Kuss, M. and Rasmussen, C. (2005). Assessing approximate inference for binary Gaussian process classification. *Journal of Machine Learning Research*, 6:1679–1704.
- Marlin, B., Khan, M., and Murphy, K. (2011). Piecewise bounds for estimating Bernoulli-logistic latent Gaussian models. In *International Conference on Machine Learning*.
- Minka, T. (2001). Expectation propagation for approximate Bayesian inference. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*.
- Mohamed, S., Ghahramani, Z., and Heller, K. A. (2009). Bayesian exponential family pca. In *Advances in neural information processing systems*, pages 1089–1096.
- Nemirovskii, A., Yudin, D. B., and Dawson, E. R. (1983). Problem complexity and method efficiency in optimization.
- Opper, M. and Archambeau, C. (2009). The Variational Gaussian Approximation Revisited. *Neural Computation*, 21(3):786–792.
- Ranganath, R., Gerrish, S., and Blei, D. M. (2014). Black box variational inference. In *International conference on Artificial Intelligence and Statistics*, pages 814–822.
- Ranganath, R., Tang, L., Charlin, L., and Blei, D. M. (2015). Deep exponential families. In *International conference on Artificial Intelligence and Statistics*.
- Raskutti, G. and Mukherjee, S. (2015). The information geometry of mirror descent. *IEEE Transactions on Information Theory*, 61(3):1451–1457.
- Rue, H. and Held, L. (2005). *Gaussian Markov random fields: theory and applications*. CRC Press.
- Salimans, T., Knowles, D. A., et al. (2013). Fixed-form variational posterior approximation through stochastic linear regression. *Bayesian Analysis*, 8(4):837–882.
- Titsias, M. and Lázaro-Gredilla, M. (2014). Doubly stochastic variational Bayes for non-conjugate inference. In *International Conference on Machine Learning*.
- Wainwright, M. J. and Jordan, M. I. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1–2:1–305.

- Wang, C. and Blei, D. M. (2013). Variational inference in nonconjugate models. *Journal of Machine Learning Research*, 14(1):1005–1031.
- Winn, J. and Bishop, C. M. (2005). Variational message passing. *Journal of Machine Learning Research*, 6(Apr):661–694.