
On the Convergence of Stochastic Variational Inference in Bayesian Networks

Ulrich Paquet
 Microsoft Research
 ulripa@microsoft.com

Abstract

We highlight a pitfall when applying stochastic variational inference to general Bayesian networks. For global random variables approximated by an exponential family distribution, natural gradient steps, commonly starting from a unit length step size, are averaged to convergence. This useful insight into the scaling of initial step sizes is lost when the approximation factorizes across a general Bayesian network, and care must be taken to ensure convergence in practice. We experimentally investigate how much of the baby (well-scaled steps) is thrown out with the bath water (exact gradients).

1 Introduction

Stochastic variational inference is framed as maximizing a global¹ variational parameter Λ , which is the natural parameter of a conjugate exponential distribution [2]. In this framework, stochastic gradient steps are taken along the natural gradient [1] to optimize for Λ . A pleasing property of stochastic variational inference on a conjugate exponential distribution and approximation $q(\Lambda)$ is that the gradient is automatically rescaled so that a unit length step size will minimize it. For a general Bayesian network, where the global variational parameters are subdivided to parameterize different factors q_i in the network’s variational approximation, the picture is less clear. Hoffman *et al.*’s appendix suggests a stochastic updating scheme like that of the global version [2]. We show here that the problem is more subtle in the general case, as component-wise noisy natural gradients can tightly couple variational parameters, and following the default recipe can sometimes lead to a scheme that “diverges” beyond recovery!

These remarks are of particular value to the Xbox recommender system, which uses stochastic variational inference in a Bayesian network on “worldwide” scale [4, 5]. Some of the results presented in Sec. 3 are preliminary investigations that were done when designing the system in 2012.

2 Variational Bayes

A Bayesian network between the variables $\mathbf{X} = \{\mathbf{x}_j\}$ defines the conditional dependency structure between them through their joint probability $p(\mathbf{X}) = \prod_j p(\mathbf{x}_j | \text{pa}_j)$. Following Fig. 1, let $\text{pa}(j)$ be the set of indexes of parents of random variable(s) \mathbf{x}_j ; for notational convenience we let $\text{pa}_j \doteq \{\mathbf{x}_k : k \in \text{pa}(j)\}$ denote the parent variables. The variables in the network can be hidden or observed, $\mathbf{X} = \{\mathbf{X}^h, \mathbf{X}^o\}$. Variational Bayes (VB) approximates the posterior $p(\mathbf{X}^h | \mathbf{X}^o)$ with $q(\mathbf{X}^h)$, by maximizing the evidence lower bound

$$\mathcal{L}[q] \doteq \int q(\mathbf{X}^h) \log \frac{p(\mathbf{X})}{q(\mathbf{X}^h)} d\mathbf{X}^h \leq \log p(\mathbf{X}^o).$$

¹The evidence lower bound is locally optimized with respect to local variational parameters.




Figure 1: A Bayesian network, indicating \mathbf{x}_i 's Markov blanket. The parents of \mathbf{x}_i are pa_i and its children $\mathbf{x}_k \in \text{ch}_i$. For a compact notation we also write $k \in \text{ch}(i)$ as the index set of children, where it is clear from context. Each child k has parents \mathbf{x}_i and cp_i (the co-parents with \mathbf{x}_i). The form of our notation loosely matches Winn and Bishop's in [6], as Alg. 1 can be interpreted as "stochastic variational message passing".

If i indexes the hidden variables, we factorize the approximation with

$$q(\mathbf{X}^h) = \prod_i q_i(\mathbf{x}_i).$$

Let $\mathbb{E}_{j \neq i}$ indicate the expectation taken over $\prod_{j \neq i} q_j(\mathbf{x}_j)$. The bound can be maximized in a component-wise fashion by iteratively setting each $q_i(\mathbf{x}_i)$ to the maximum

$$\log q_i^*(\mathbf{x}_i) = \mathbb{E}_{j \neq i} \left[\log p(\mathbf{x}_i | \text{pa}_i) \right] + \sum_{k \in \text{ch}(i)} \mathbb{E}_{j \neq i} \left[\log p(\mathbf{x}_k | \text{pa}_k) \right] + \text{const}. \quad (1)$$

In many practical networks there are some \mathbf{x}_i for which number of children $N_i \doteq |\text{ch}(i)|$ is large. In [2], \mathbf{x}_i is a topic-vocabulary distribution from which millions of documents are generated. In Sec. 3 and [4, 5] the interaction is bilinear, where user \mathbf{x}_i and item \mathbf{x}_j variables are combined to represent a user's affinity to an item. Rather than summing over all $\text{ch}(i)$ for each update in (1), we aim to stochastically approximate the expectations. It alleviates two problems: firstly, the sum contains many terms; secondly, the update depends on some $q(\mathbf{x}_j)$ which will be re-estimated, and the expense of fully estimating $q(\mathbf{x}_i)$ is lost as it too will be re-estimated.

2.1 Conditionally conjugate models

The updates in (1) are straightforward when the Bayesian network is conditionally conjugate; that is, when the distribution of \mathbf{x}_i , conditioned on pa_i , is (a) drawn from an exponential family, and (b) is conjugate with respect to the distribution of pa_i . We define the exponential family as

$$\log p(\mathbf{x}_i | \text{pa}_i) = \boldsymbol{\eta}_i(\text{pa}_i)^T \phi_i(\mathbf{x}_i) + f_i(\mathbf{x}_i) + g_i(\text{pa}_i) \quad (2)$$

where $\boldsymbol{\eta}_i(\text{pa}_i)$ is the natural parameter vector, $\phi_i(\mathbf{x}_i)$ forms the sufficient statistics, and $g_i(\text{pa}_i)$ defines the normalizing constant through $g_i(\text{pa}_i) = -\log \int \exp\{\boldsymbol{\eta}_i(\text{pa}_i)^T \phi_i(\mathbf{x}_i) + f_i(\mathbf{x}_i)\} d\mathbf{x}_i$. We can view (2) as a "prior" over \mathbf{x}_i . Now consider a node $\mathbf{x}_k \in \text{ch}_i$ in Fig. 1. We subdivide pa_k , the parents of \mathbf{x}_k , into \mathbf{x}_i and its co-parents cp_i :

$$\log p(\mathbf{x}_k | \mathbf{x}_i, \text{cp}_i) = \boldsymbol{\eta}_k(\mathbf{x}_i, \text{cp}_i)^T \phi_k(\mathbf{x}_k) + f(\mathbf{x}_k) + g(\mathbf{x}_i, \text{cp}_i).$$

We can view this as a contribution to the "likelihood" of \mathbf{x}_i . We include the co-parents as they are part of \mathbf{x}_i 's Markov blanket. Through conjugacy, $p(\mathbf{x}_i | \text{pa}_i)$ and $p(\mathbf{x}_k | \mathbf{x}_i, \text{cp}_i)$ have the same functional form with respect to \mathbf{x}_i , so that we can rewrite $p(\mathbf{x}_k | \mathbf{x}_i, \text{cp}_i)$ in terms of the sufficient statistics $\phi_i(\mathbf{x}_i)$ by defining some function $\boldsymbol{\eta}_{ki}$ with

$$\log p(\mathbf{x}_k | \mathbf{x}_i, \text{cp}_i) = \boldsymbol{\eta}_{ki}(\mathbf{x}_k, \text{cp}_i)^T \phi_i(\mathbf{x}_i) + h(\mathbf{x}_k, \text{cp}_i).$$

We furthermore parameterize the $q(\mathbf{x}_i)$ distributions in terms of their natural parameters. To distinguish them, we denote their natural parameters by $\boldsymbol{\lambda}_i$, and define $\boldsymbol{\Lambda} \doteq [\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_I]$:

$$\log q_i(\mathbf{x}_i | \boldsymbol{\lambda}_i) = \boldsymbol{\lambda}_i^T \phi_i(\mathbf{x}_i) + f_i(\mathbf{x}_i) + \tilde{g}_i(\boldsymbol{\lambda}_i). \quad (3)$$

2.2 Variational Bayes updates and their stochastic version

Returning to (1), we can write

$$\log q_i^*(\mathbf{x}_i) = \mathbb{E}_{j \neq i} \left[\boldsymbol{\eta}_i(\text{pa}_i) + \sum_{k \in \text{ch}(i)} \boldsymbol{\eta}_{ki}(\mathbf{x}_k, \text{cp}_i) \right]^T \phi_i(\mathbf{x}_i) + f_i(\mathbf{x}_i) + \text{const}, \quad (4)$$

from which we can directly read off the updated natural parameter λ_i^* through (3). Notice now that η_i is a multi-linear function of the random variables pa_i , i.e. it is linear in *each* parent random variable. In the same way η_{ki} is a multi-linear function of the random variables \mathbf{x}_k and cp_i . Furthermore, q factorizes over these variables (except where they are observed, of course). We can therefore reparameterize (4) in terms of expectations over q_j , $i \neq j$ with

$$\begin{aligned} \mathbb{E}_{j \neq i} [\eta_i(\text{pa}_i)] &= \tilde{\eta}_i \left(\left\{ \mathbb{E}_j [\phi_j(\mathbf{x}_j)] \right\}_{j \in \text{pa}(i)} \right) \doteq \tilde{\eta}_i \\ \mathbb{E}_{j \neq i} [\eta_{ki}(\mathbf{x}_k, \text{cp}_i)] &= \tilde{\eta}_{ki} \left(\mathbb{E}_k [\phi_k(\mathbf{x}_k)], \left\{ \mathbb{E}_j [\phi_j(\mathbf{x}_j)] \right\}_{j \in \text{cp}(i)} \right) \doteq \tilde{\eta}_{ki}. \end{aligned}$$

Algorithm 1: Stochastic Variational Bayes

```

1: for  $t = 1$  to  $t_{\max}$  or convergence do
2:    $\rho_t = (t + \tau)^{-\kappa}$ 
3:   for each hidden  $\mathbf{x}_i$  do
4:      $\mathcal{C} \leftarrow C$  random nodes from  $\text{ch}(i)$ 
5:      $\lambda_i^{\text{temp}} \leftarrow \tilde{\eta}_i + \frac{N_i}{C} \sum_{k \in \mathcal{C}} \tilde{\eta}_{ki}$ 
6:     option (a):  $\lambda_i \leftarrow (1 - \rho_t)\lambda_i + \rho_t \lambda_i^{\text{temp}}$ 
7:   end for
8:   option (b):  $\Lambda \leftarrow (1 - \rho_t)\Lambda + \rho_t \Lambda^{\text{temp}}$ 
9: end for

```

This is a key ingredient of algorithms like variational message passing [6]. (When \mathbf{x}_k is observed, $\phi_k(\mathbf{x}_k)$ is kept as is, as it is averaged over a delta function.) The variational update in (4) becomes

$$\lambda_i^* = \tilde{\eta}_i + \sum_{k \in \text{ch}(i)} \tilde{\eta}_{ki}. \quad (5)$$

The update is a step along the natural gradient [1], equivalent to setting the gradient to zero by solving for the zero of the derivative of \mathcal{L} with respect to λ_i . In particular, (5) updates

λ_i from its old value to λ_i^* using a *step of unit length* along the natural gradient. Sec. A.1 derives the gradient $\nabla_{\lambda_i} \mathcal{L}$, and Sec. A.2 states its natural form $\hat{\nabla}_{\lambda_i} \mathcal{L}$.

When $N_i \doteq |\text{ch}(i)|$ is large, not all the child nodes might be accessed in reasonable time. Furthermore, when $q(\mathbf{x}_i)$ is re-estimated, the (previous) large computation is discarded and recomputed. We may alternatively consider a subsample of nodes from $\text{ch}(i)$ to determine the sufficient statistics. By placing a uniform distribution \tilde{p}_i on the atoms $\tilde{\eta}_{ki}$, the update from (5) is equivalent to $\lambda_i = \tilde{\eta}_i + N_i \mathbb{E}_{\tilde{p}_i}[\tilde{\eta}]$. This expectation can be estimated in many ways. Let set \mathcal{C} be a sample of C children from $\text{ch}(i)$ without replacement and let

$$\lambda_i^{\text{temp}} = \tilde{\eta}_i + \frac{N_i}{C} \sum_{k \in \mathcal{C}} \tilde{\eta}_{ki}. \quad (6)$$

Taking expectations gives $\lambda_i^* = \mathbb{E}[\lambda_i^{\text{temp}}] = \tilde{\eta}_i + N_i \mathbb{E}_{\tilde{p}_i}[\tilde{\eta}]$. With $\rho_t \rightarrow 0$, $\sum_{t=1}^{\infty} \rho_t = \infty$ and $\sum_{t=1}^{\infty} \rho_t^2 < \infty$, these stochastic natural gradients are used in Alg. 1, which is a stochastic version of variational message passing. In Alg. 1, scalar $\kappa \in (\frac{1}{2}, 1]$ is a forgetting rate, while delay $\tau \geq 0$ discounts early iterations more.

There are two options in Alg. 1. For option (a), the mean value of the parameters of $q(\mathbf{X}^h)$ is periodic in I , the number of factors in q , and convergence to a local optimum can also be guaranteed for I -dependent mean values [3]. Option (b) is the update scheme given in [2].

3 Bayesian matrix factorization

To illustrate a general Bayesian network, we factorize a sparse matrix of a subsample of a million entries in the Netflix data set ($M = 4805$ users and $N = 16015$ items). Each entry r_{mn} is user m 's rating of movie n on a five-star rating scale. For illustrative purposes, consider a Gaussian bilinear ratings model

$$p(r_{mn} | \mathbf{u}_m, \mathbf{v}_n) = \mathcal{N}(r_{mn}; \mathbf{u}_m^T \mathbf{v}_n, 1)$$

for user parameter vector $\mathbf{u}_m \in \mathbb{R}^K$ and item trait vector $\mathbf{v}_n \in \mathbb{R}^K$. We place a factorized prior $\mathcal{N}(u_{mk}; 0, 1)$ on each of the entries of \mathbf{u}_m and \mathbf{v}_n . We choose a fully factorized Gaussian approximation $q(\mathbf{U}) = \prod_m \prod_k q(u_{mk})$, with a similar approximation for $q(\mathbf{V})$. The VB update for $q(u_{mk})$ therefore incorporates $2K - 1$ co-parents due to the inner product. With the Gaussian's natural parameters being its precision and mean-times-precision, it is

$$\lambda_{mk}^{\text{temp}} = \begin{pmatrix} \text{prec} \\ \text{mean} \cdot \text{prec} \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \frac{N_m}{C} \sum_{n \in \mathcal{C}} \begin{pmatrix} \text{var}_q[v_{nk}] + \mathbb{E}_q[v_{nk}]^2 \\ \mathbb{E}_q[v_{nk}](r_{mn} - \sum_{k' \neq k} \mathbb{E}_q[u_{mk'}] \mathbb{E}_q[v_{nk'}]) \end{pmatrix}. \quad (7)$$

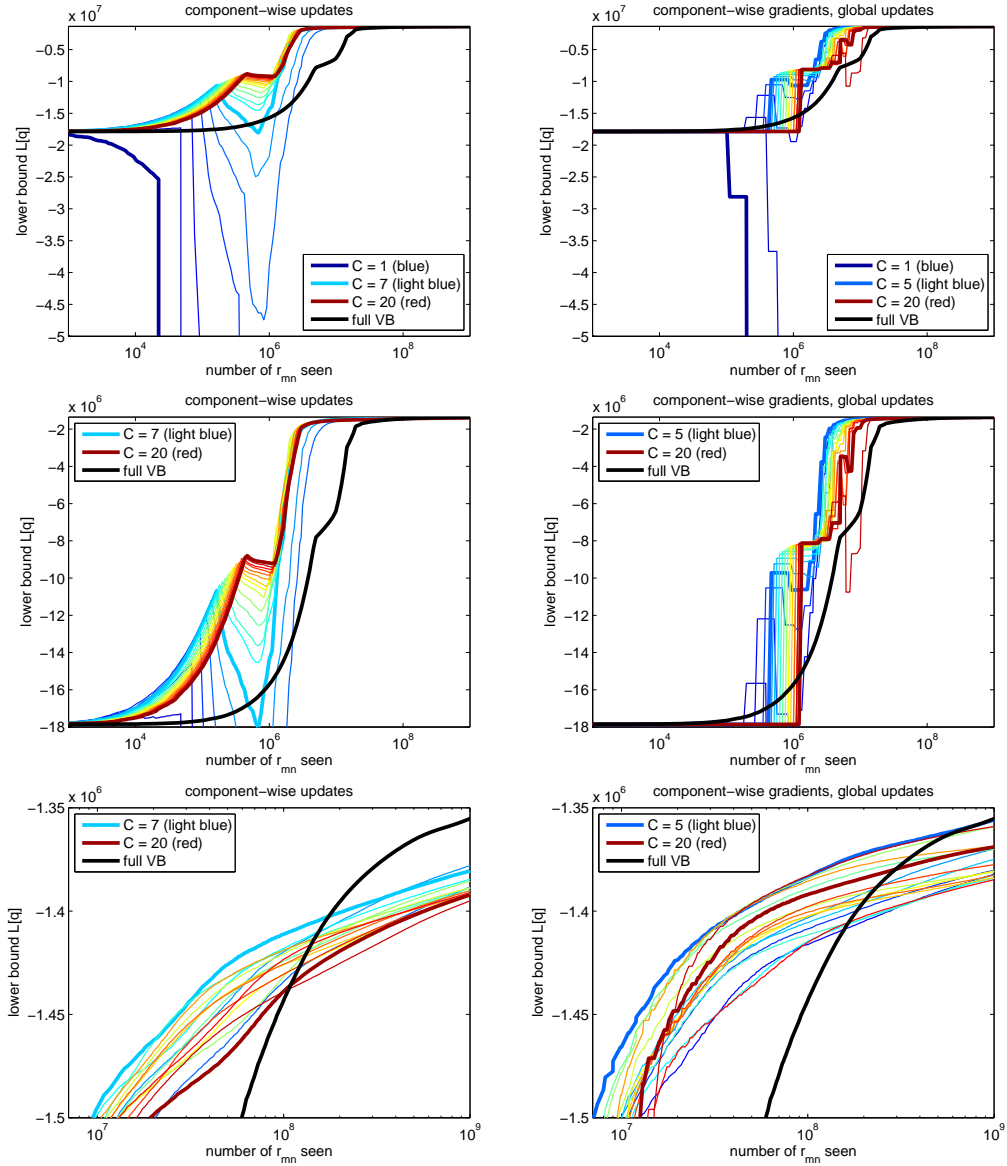


Figure 2: Convergence of $\mathcal{L}[q]$ with $\rho_t = t^{-0.6}$. Alg. 1’s option (a) is shown in the left column; option (b) is shown in the right column. The x-axes are on a logarithmic scale. The *global* stochastic gradient is not in its natural form, and the effect of a large variance in the gradient estimate and overshooting with too large step sizes of $\rho_t \in (0, 1]$ is clearly visible for small C . Note that r_{mn} ’s can be revisited over multiple loops in Alg. 1. Different magnifications of the same two convergence plots for options (a) and (b) are shown in the three rows of graphs.

Fig. 2 shows $\mathcal{L}[q]$ as a function of the number of times that individual ratings (observed nodes) r_{mn} are accessed or queried (using $K = 5$). The value of the bound is shown for the use of at most $C = 1, \dots, 20$ children when estimating the gradient of each random variable with (6) and (7). Both options (a) and (b) in Alg. 1 “diverge” in a numerically unrecoverable way when C is small.

Full VB, shown in black in Fig. 2, implicitly uses $\rho_t = 1$ in (5). As the stochastic natural gradient depends on other λ_j , the divergence can be attributed to simply using a step size of $\rho_t \in (0, 1]$ that is too big, and overshooting with too large gradient steps. Much smaller initial step sizes are required to not overshoot, as the variance of the gradient is simply too large compared to ρ_t . Fig. 2 illustrates this problem for *general Bayesian networks*; see especially the *top left* figure.

In practice, we can overcome this problem by starting with sufficiently small initial step sizes $\rho_1 \ll 1$. For $C = 1$ in option (a) this was starting from $\rho_1 = \frac{1}{512}$, and $\rho_1 = \frac{1}{64}$ for option (b). In [4, 5], which infers the posterior of a different model than the one stated here, the value of C varied depending on a user or item's usage, and there $\rho_t = 1$ was fixed for the first ten iterations before slowly decreasing it.

Have we thrown the baby (well-scaled steps) out with the bath water (exact gradients)? Maybe some. As shown by this short note, it is still an open question.

Acknowledgment

To the anonymous reviewer who pointed out that the Fisher information matrix of $q(\mathbf{X}^h|\Lambda)$ block-diagonal, having the Fisher information matrices of $q(\mathbf{x}_i|\lambda_i)$ along its diagonal: thank you!

References

- [1] S. Amari. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.
- [2] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14:1303–1347, 2013.
- [3] H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, 2003.
- [4] U. Paquet and N. Koenigstein. One-class collaborative filtering with random graphs. In *Proceedings of the 22nd international conference on World Wide Web, WWW '13*, pages 999–1008, 2013.
- [5] U. Paquet and N. Koenigstein. One-class collaborative filtering with random graphs: Annotated version. *arXiv:1309.6786*, 2013.
- [6] J. Winn and C.M. Bishop. Variational message passing. *Journal of Machine Learning Research*, 6:661–694, 2005.

A Gradients

In this appendix, we derive the component-wise gradients and their natural version, and present basic intuition for why steps along the stochastic gradient can be taken.

A.1 Component-wise gradients

The function that is minimized to find (5) is $\mathcal{L}(\lambda_i)$ below. It is a function of λ_i , whilst keeping all other λ_j for $j \neq i$ fixed:

$$\begin{aligned} \mathcal{L}(\lambda_i) &= \mathbb{E}_q \left[\log p(\mathbf{x}_i | \text{pa}_i) + \sum_{k \in \text{ch}(i)} \log p(\mathbf{x}_k | \mathbf{x}_i, \text{cp}_i) - \log q_i(\mathbf{x}_i | \lambda_i) \right] \\ &= \mathbb{E}_q \left[\boldsymbol{\eta}_i(\text{pa}_i)^T \phi_i(\mathbf{x}_i) + f_i(\mathbf{x}_i) + g_i(\text{pa}_i) \right. \\ &\quad \left. + \sum_{k \in \text{ch}(i)} \left(\boldsymbol{\eta}_k(\mathbf{x}_i, \text{cp}_i)^T \phi_k(\mathbf{x}_k) + f(\mathbf{x}_k) + g(\mathbf{x}_i, \text{cp}_i) \right) \right. \\ &\quad \left. - \lambda_i^T \phi_i(\mathbf{x}_i) - f_i(\mathbf{x}_i) - \tilde{g}_i(\lambda_i) \right]. \end{aligned}$$

The probability $p(\mathbf{x}_k | \mathbf{x}_i, \text{cp}_i)$ is stated in terms of the sufficient statistics $\phi_k(\mathbf{x}_k)$. Because of local conjugacy, it can be rewritten in terms of the sufficient statistics $\phi_i(\mathbf{x}_i)$ through a multi-linear function $\boldsymbol{\eta}_{ki}$ of the random variables \mathbf{x}_k and cp_i to yield

$$\begin{aligned} \mathcal{L}(\lambda_i) &= \mathbb{E}_q \left[\boldsymbol{\eta}_i(\text{pa}_i)^T \phi_i(\mathbf{x}_i) + f_i(\mathbf{x}_i) + g_i(\text{pa}_i) + \sum_{k \in \text{ch}(i)} \left(\boldsymbol{\eta}_{ki}(\mathbf{x}_k, \text{cp}_i)^T \phi_i(\mathbf{x}_i) + h(\mathbf{x}_k, \text{cp}_i) \right) \right. \\ &\quad \left. - \lambda_i^T \phi_i(\mathbf{x}_i) - f_i(\mathbf{x}_i) - \tilde{g}_i(\lambda_i) \right]. \end{aligned}$$

Taking expectations over q gives, as function of λ_i ,

$$\mathcal{L}(\lambda_i) = \left(\tilde{\eta}_i + \sum_{k \in \text{ch}(i)} \tilde{\eta}_{ki} \right)^T \mathbb{E}_i[\phi_i(\mathbf{x}_i)] - \lambda_i^T \mathbb{E}_i[\phi_i(\mathbf{x}_i)] - \tilde{g}_i(\lambda_i) + \text{const}, \quad (8)$$

with $\tilde{g}_i(\lambda_i) = -\log \int \exp\{\lambda_i^T \phi_i(\mathbf{x}_i) + f_i(\mathbf{x}_i)\} d\mathbf{x}_i$. We would like to compute the gradient $\nabla_{\lambda_i} \mathcal{L}(\lambda_i)$. The derivatives of the log partition function $-\tilde{g}_i(\lambda_i)$ with respect to λ_i give the expected sufficient statistics

$$\begin{aligned} -\nabla \tilde{g}_i(\lambda_i) &= \mathbb{E}_i[\phi_i(\mathbf{x}_i)] \\ -\nabla^2 \tilde{g}_i(\lambda_i) &= \nabla \mathbb{E}_i[\phi_i(\mathbf{x}_i)] \\ &= \mathbb{E}_i \left[(\phi_i(\mathbf{x}_i) - \mathbb{E}_i[\phi_i(\mathbf{x}_i)]) (\phi_i(\mathbf{x}_i) - \mathbb{E}_i[\phi_i(\mathbf{x}_i)])^T \right] \\ &= \text{cov}_i[\phi_i(\mathbf{x}_i)]. \end{aligned}$$

By using properties of the exponential family, the gradient of \mathcal{L} with respect to λ_i is therefore

$$\nabla_{\lambda_i} \mathcal{L}(\lambda_i) = \text{cov}_i[\phi_i(\mathbf{x}_i)] \left(\tilde{\eta}_i + \sum_{k \in \text{ch}(i)} \tilde{\eta}_{ki} - \lambda_i \right). \quad (9)$$

Solving for $\nabla_{\lambda_i} \mathcal{L}(\lambda_i) = \mathbf{0}$ yields the component-wise VB update $\lambda_i^* = \tilde{\eta}_i + \sum_{k \in \text{ch}(i)} \tilde{\eta}_{ki}$ that we find in (5). Gradient $\nabla_{\lambda_i} \mathcal{L}$ depends on λ_i through $\text{cov}_i[\phi_i(\mathbf{x}_i)]$ and λ_i , and in the next section we will show that the natural gradient $\tilde{\nabla}_{\lambda_i} \mathcal{L}$ removes the dependency on $\text{cov}_i[\phi_i(\mathbf{x}_i)]$, so that it is a linear function of λ_i , with the minimum being attained by taking a step of length one along it.

A.2 Component-wise natural gradients

The Fisher information matrix of q_i is

$$\begin{aligned} G(\lambda_i) &= \mathbb{E}_i \left[(\nabla_{\lambda_i} \log q(\mathbf{x}_i | \lambda_i)) (\nabla_{\lambda_i} \log q(\mathbf{x}_i | \lambda_i))^T \right] \\ &= \mathbb{E}_i \left[(\phi_i(\mathbf{x}_i) - \mathbb{E}_i[\phi_i(\mathbf{x}_i)]) (\phi_i(\mathbf{x}_i) - \mathbb{E}_i[\phi_i(\mathbf{x}_i)])^T \right] \\ &= \text{cov}_i[\phi_i(\mathbf{x}_i)], \end{aligned}$$

and the component-wise natural gradient is obtained by multiplying it with $\nabla_{\lambda_i} \mathcal{L}$, yielding

$$\widehat{\nabla}_{\lambda_i} \mathcal{L}(\lambda_i) \doteq G(\lambda_i)^{-1} \nabla_{\lambda_i} \mathcal{L}(\lambda_i) = \tilde{\eta}_i + \sum_{k \in \text{ch}(i)} \tilde{\eta}_{ki} - \lambda_i.$$

Take a gradient descent step along the natural gradient with step length $\rho > 0$. Starting at point $\lambda_i^{(t-1)}$, gradient descent updates it to $\lambda_i^{(t)}$ with

$$\begin{aligned} \lambda_i^{(t)} &\leftarrow \lambda_i^{(t-1)} + \rho \widehat{\nabla}_{\lambda_i^{(t-1)}} \mathcal{L}(\lambda_i) = \lambda_i^{(t-1)} + \rho \left(\tilde{\eta}_i + \sum_{k \in \text{ch}(i)} \tilde{\eta}_{ki} - \lambda_i^{(t-1)} \right) \\ &= (1 - \rho) \lambda_i^{(t-1)} + \rho \left(\tilde{\eta}_i + \sum_{k \in \text{ch}(i)} \tilde{\eta}_{ki} \right) \\ &= (1 - \rho) \lambda_i^{(t-1)} + \rho \lambda_i^*. \end{aligned}$$

When the above update is compared to (5), we see that the minimum $\lambda_i^{(t)} \leftarrow \lambda_i^*$ is obtained by applying a step size of $\rho = 1$ along the natural gradient.

A.3 Stochastic natural gradients: a bird's eye view

In this section an intuitive motivation will be provided for doing stochastic gradient descent using the natural gradient, as it was defined above in Sec. A.2. The explanation favours an intuitive

understanding above mathematical rigour. Imagine that instead of λ_i^* , we have access to a sequence of samples $\{\lambda_i^{\text{temp},\tau}\}_{\tau=1}^t$, so that $\mathbb{E}[\lambda_i^{\text{temp}}] = \lambda_i^*$. We can write the update $\lambda_i^{(t)} \leftarrow \lambda_i^*$ recursively using the sample average

$$\begin{aligned}\lambda_i^{(t)} &\leftarrow \frac{1}{t} \sum_{\tau=1}^t \lambda_i^{\text{temp},\tau} = \left(\frac{t-1}{t}\right) \left(\frac{1}{t-1} \sum_{\tau=1}^{t-1} \lambda_i^{\text{temp},\tau}\right) + \frac{1}{t} \lambda_i^{\text{temp},t} \\ &= \left(1 - \frac{1}{t}\right) \lambda_i^{(t-1)} + \frac{1}{t} \lambda_i^{\text{temp},t}.\end{aligned}$$

Define $\rho_t \doteq \frac{1}{t}$. In the running average, $\sum_{t=1}^{\infty} \frac{1}{t} = \infty$ and $\sum_{t=1}^{\infty} \left(\frac{1}{t}\right)^2 < \infty$, and therefore $\sum_{t=1}^{\infty} \rho_t = \infty$ and $\sum_{t=1}^{\infty} \rho_t^2 < \infty$. In the running average with $\rho_t \doteq \frac{1}{t}$, each gradient sample is treated equally. However, instead of incorporating a fraction $\frac{1}{t}$ of $\lambda_i^{\text{temp},t}$ into the running average, we may erase a bit more from the ‘‘past memory’’ $\lambda_i^{(t-1)}$ to include a bit *more* of the recent gradient $\lambda_i^{\text{temp},t}$. How much more is permissible?

Now define $\rho_t \doteq t^{-\kappa}$. For $\kappa = \frac{1}{2}$, the previous samples will be forgotten at a faster rate, and more of $\lambda_i^{\text{temp},t}$ will be included through $\lambda_i^{(t)} \leftarrow (1 - t^{-1/2})\lambda_i^{(t-1)} + t^{-1/2}\lambda_i^{\text{temp},t}$. However, at this rate past samples are forgotten too quickly, as both $\sum_{t=1}^{\infty} t^{-1/2} = \infty$ and $\sum_{t=1}^{\infty} (t^{-1/2})^2 = \infty$. For any $\kappa' > 1$, both infinite sums will be finite, e.g. $\sum_{t=1}^{\infty} t^{-\kappa'} < \infty$ and $\sum_{t=1}^{\infty} (t^{-\kappa'})^2 < \infty$, and the running average will cling on to old memories, and has too little capacity to incorporate recent gradient samples $\lambda_i^{\text{temp},t}$. Between forgetting too quickly or not at all, a setting of $\kappa \in (\frac{1}{2}, 1]$ in $\rho_t \doteq t^{-\kappa}$ is therefore permissible.

A.4 Converging with fickle neighbours

The running average in Sec. A.3 can boldly start at $\rho_t \doteq t^{-\kappa} = 1$ for $t = 1$, and from this unit length step along the natural gradient, it accumulates gradient samples until convergence. However, it rests on the premise that neighbours λ_j for $j \neq i$ from the Markov blanket of \mathbf{x}_i remain unchanged.

If this premise does not hold, much smaller steps $\rho_t \doteq (t + \tau)^{-\kappa}$ with delay $\tau \geq 0$ are required. This is indeed the case. As Sec. 3 shows, a delay $\tau \geq 0$ that is sufficiently large for the stochastic gradient scheme to converge in practice is not known *a priori*. By explicitly expanding the shorthand definitions of $\tilde{\eta}_i$ and $\tilde{\eta}_{ki}$ in (8), it is clear that the other λ_j appear through multi-linear functions in

$$\begin{aligned}\mathcal{L}(\lambda_i) &= \left(\tilde{\eta}_i \left(\{\mathbb{E}_j[\phi_j(\mathbf{x}_j)]\}_{j \in \text{pa}(i)} \right) + \sum_{k \in \text{ch}(i)} \tilde{\eta}_{ki} \left(\mathbb{E}_k[\phi_k(\mathbf{x}_k)], \{\mathbb{E}_j[\phi_j(\mathbf{x}_j)]\}_{j \in \text{cp}(i)} \right) \right)^T \mathbb{E}_i[\phi_i(\mathbf{x}_i)] \\ &\quad - \lambda_i^T \mathbb{E}_i[\phi_i(\mathbf{x}_i)] - \tilde{g}_i(\lambda_i) + \text{const}.\end{aligned}$$

If we now consider the global gradient $\nabla_{\Lambda} \mathcal{L}(\Lambda) = [\nabla_{\lambda_1} \mathcal{L}(\lambda_1), \dots, \nabla_{\lambda_I} \mathcal{L}(\lambda_I)]$, it is clear from the above form (multi-linear in all variables) that we can’t set the gradient to zero and solve for all Λ explicitly, as was done in (5). It is usually not even convex problem.

The gradient steps are along the global natural gradient. It is defined as

$$\tilde{\nabla}_{\Lambda} \mathcal{L}(\Lambda) \doteq G(\Lambda)^{-1} \nabla_{\Lambda} \mathcal{L}(\Lambda),$$

with $G(\Lambda)$ being the Fisher information matrix of q ,

$$G(\Lambda) = \mathbb{E}_q \left[(\nabla_{\Lambda} \log q(\mathbf{X}^h | \Lambda)) (\nabla_{\Lambda} \log q(\mathbf{X}^h | \Lambda))^T \right].$$

$G(\Lambda)$ is block-diagonal, as the covariance between $\phi_i(\mathbf{x}_i)$ and $\phi_j(\mathbf{x}_j)$ is zero for $i \neq j$, due to the factorization of q . Its inverse is therefore also block-diagonal, and the natural gradient has the form $\tilde{\nabla}_{\Lambda} \mathcal{L}(\Lambda) = [\tilde{\nabla}_{\lambda_1} \mathcal{L}(\lambda_1), \dots, \tilde{\nabla}_{\lambda_I} \mathcal{L}(\lambda_I)]$.

B Global batch samples

An alternative to Alg. 1 is to take a batch data sample \mathcal{D} of C_{global} observed variables at the start of each iteration, and follow either each $\tilde{\nabla}_{\lambda_i} \mathcal{L}(\lambda_i)$ or $\tilde{\nabla}_{\Lambda} \mathcal{L}(\Lambda)$. This is outlined in Alg. 2.

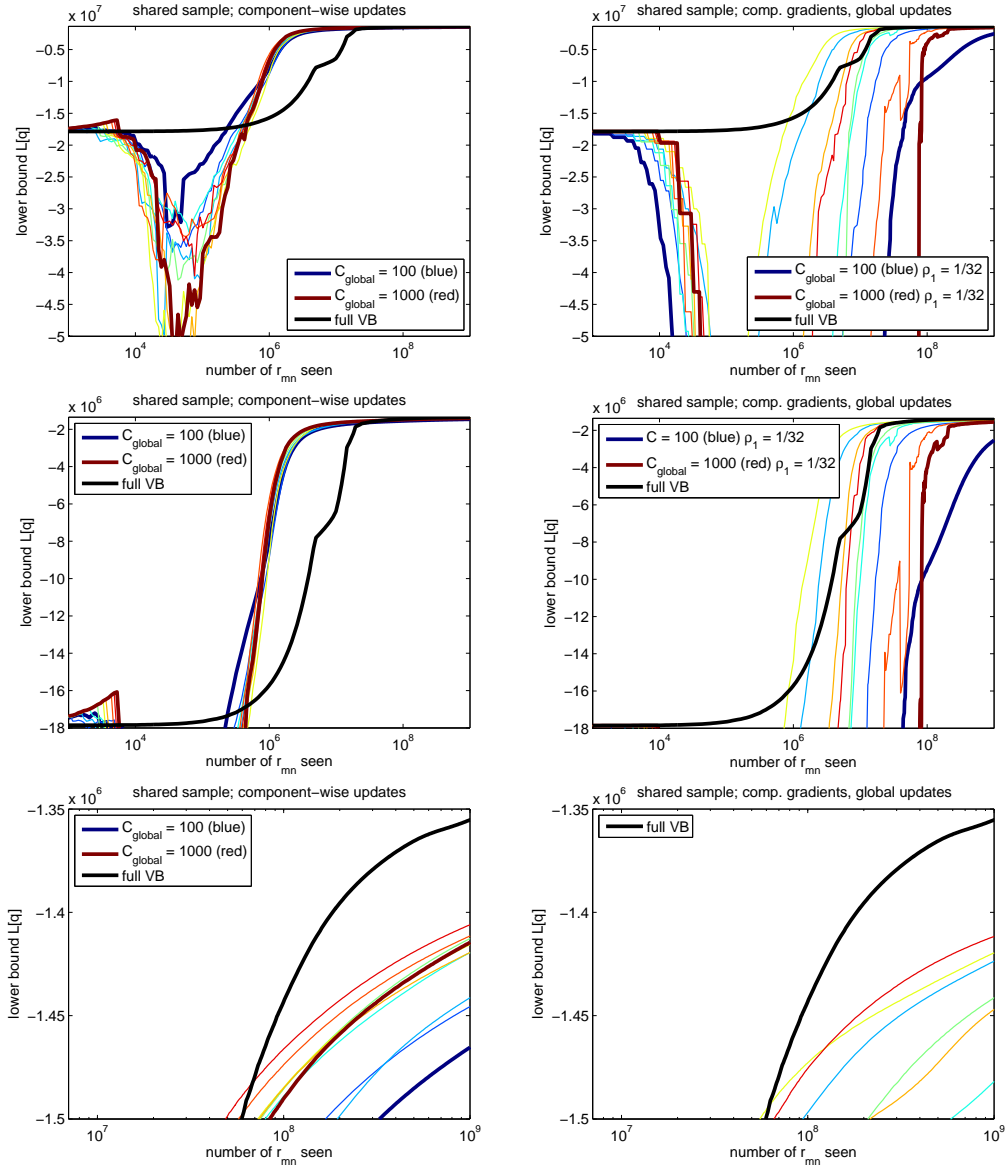


Figure 3: Convergence of $\mathcal{L}[q]$ with $\rho_t = t^{-0.6}$. Alg. 2's option (a) is shown in the left column; option (b) is shown in the right column. The x-axes are on a logarithmic scale.

Algorithm 2: Stochastic Variational Bayes

- 1: **for** $t = 1$ to t_{\max} or convergence **do**
 - 2: $\rho_t = (t + \tau)^{-\kappa}$
 - 3: $\mathcal{D} \leftarrow C_{\text{global}}$ random nodes from \mathbf{X}^o
 - 4: **for** each hidden $\mathbf{x}_i : i \in \text{pa}(\mathcal{D})$ **do**
 - 5: $\mathcal{D}_i \leftarrow \{\mathbf{x}_k \in \mathcal{D} \cap \text{ch}(i)\}$
 - 6: $\lambda_i^{\text{temp}} \leftarrow \tilde{\eta}_i + \frac{N_i}{|\mathcal{D}_i|} \sum_{k \in \mathcal{D}_i} \tilde{\eta}_{ki}$
 - 7: option (a): $\lambda_i \leftarrow (1 - \rho_t)\lambda_i + \rho_t\lambda_i^{\text{temp}}$
 - 8: **end for**
 - 9: // updates of $\text{pa}(i)$ etc.
 - 10: option (b): $\Lambda \leftarrow (1 - \rho_t)\Lambda + \rho_t\Lambda^{\text{temp}}$
 - 11: **end for**
-

Fig. 3 considers batch sizes of $C_{\text{global}} = 100$ to 1000, in intervals of 100. The convergence in Fig. 3 is much slower than that of Fig. 2. For the global update in option (b) in Alg. 2, the algorithm only converged on finite machine precision when $\rho_1 \leq \frac{1}{32}$ was chosen (smaller for some C_{global} settings), whereas for option (a) at least the algorithm converged from $\rho_1 = 1$ for all settings.

In Alg. 2, $\text{pa}(\mathcal{D})$ is the set of hidden variables that have a child node in \mathcal{D} . Line 9 makes provision for updating variables (like hyper-parameters) that don't have observed children; this was not required for our example.