Expectation Propagation performs a smoothed gradient descent

Guillaume Dehaene
Ecole Polytechnique Federale de Lausanne
guillaume.dehaene@gmail.com

Abstract

Bayesian inference is a popular method to build learning algorithms but it is hampered by the fact that its key object, the posterior probability distribution, is often uncomputable. Expectation Propagation (EP) (Minka [2001]) is a popular algorithm that solves this issue by computing a parametric approximation (e.g: Gaussian) to the density of the posterior. However, while it is known empirically to quickly compute fine approximations, EP is extremely poorly understood which prevents it from being adopted by a larger fraction of the community.

The object of the present article is to shed intuitive light on EP, by relating it to other better understood methods. More precisely, we link it to using gradient descent to compute the Laplace approximation of a target probability distribution. We show that EP is exactly equivalent to performing gradient descent on a smoothed energy landscape: i.e: the original energy landscape convoluted with some smoothing kernel. This also relates EP to algorithms that compute the Gaussian approximation which minimizes the reverse KL divergence to the target distribution, a link that has been conjectured before but has not been proved rigorously yet. These results can help practitioners to get a better feel for how EP works, as well as lead to other new results on this important method.

Throughout this article, we consider the task of approximating a probability density over a one-dimensional space:

$$p(\theta) = \exp(-\psi(\theta)) \tag{1}$$

where we will assume for simplicity that $\psi(\theta)$ is convex so that $\psi''(\theta) > 0$.

We will present various algorithms to compute Gaussian approximations of p. We will first seek to compute the "Laplace approximation". We then turn to computing the Gaussian which minimizes the reverse KL divergence to $p(\theta)$. Finally, we consider using Expectation Propagation (EP). These three approximations can be computed by similar algorithms: either by exact gradient descent or by variants to it which we refer to as "smoothed gradient descent". This new perspective on EP is useful in giving practitioners a more intuitive framework than current justifications of EP, but it can also be used to derive new theoretical results on the method. For example, we show that, in various asymptotes, the limit behavior of EP is simple.

To keep the presentation of these ideas sharp and light, we do not prove the results we present in this document. All proofs, which are stated in the high-dimensional case, can be found in the appendix.

1 The Laplace approximation

If we want to compute a Gaussian approximation of $p(\theta)$, a natural idea consists in performing a Taylor expansion to second order of $\psi(\theta)$. It seems natural to perform the expansion around the

global maximum $\theta^\star$ of the target $p(\theta)$ (which is unique because of the convexity of $\psi(\theta)$). This yields the "Laplace" approximation of the target (Murphy [2012] 8.4.1):

$$p(\theta) \approx p(\theta^\star) \exp\left(-\psi''(\theta^\star) \frac{(\theta^\star - \theta)^2}{2}\right) \tag{2}$$

This requires us to compute the global maximum $\theta^\star$ (which is unique under our convexity assumption). A nice solution for this is Newton's method (NT), which corresponds to gradient descent of $\psi(\theta)$ with a Hessian correction. This method is actually extremely closely linked to the "Laplace" approximation, since we can derive it as iterating over Gaussian approximations. Indeed, consider the following method:

- Start from some initial Gaussian approximation $q^{(0)}(\theta)$ to $p(\theta)$, with mean μ_0.
- Then loop until convergence:
  1. Compute the mean $\mu^{(t)}$ of the current approximation $q^{(t)}(\theta)$.
  2. Construct the new approximation using the second degree expansion of ψ around $\mu^{(t)}$:

$$q^{(t+1)}(\theta) \propto \exp\left(-\psi'\left(\mu^{(t)}\right)\left(\theta - \mu^{(t)}\right) - \psi''\left(\mu^{(t)}\right) \frac{\left(\theta - \mu^{(t)}\right)^2}{2}\right) \tag{3}$$

It is straightforward to check that the dynamics of $\mu^{(t)}$ in this algorithm match those in the classical NT algorithm. The fixed-point of the Gaussian-iterating algorithm is the Laplace approximations of the target $p(\theta)$.

Gradient descent is fairly easy to understand from an intuitive point of view. It corresponds roughly to the dynamics of a ball which we drop on some complicated energy landscape $\psi(\theta)$ and falls down along the valleys of the landscape until it reaches some local minimum.

In practice, the Laplace approximation does not provide a very good approximation of the target $p(\theta)$. Intuitively, this is because it depends on the value of ψ at a single point $\theta^\star$ so it naturally fails at providing a global account of the target distribution.

2 Smoothing the gradient minimizes the reverse KL divergence

In order to provide a Gaussian approximation that gives a more global fit to the target distribution, why not replace the point estimates $\psi'\left(\mu^{(t)}\right)$ and $\psi''\left(\mu^{(t)}\right)$ by averages over a large region of θ-space? We could use many densities for this, but it seems natural to use the current Gaussian approximation $q^{(t)}(\theta)$ which is indeed centered at $\mu^{(t)}$.

We can thus construct a "smoothed Newton method" by replacing eq. (3) with:

$$q^{(t+1)}\left(\tilde{\theta}\right) \propto \exp\left(-E_{q^{(t)}}\left[\psi'(\theta)\right]\left(\tilde{\theta} - \mu^{(t)}\right) - E_{q^{(t)}}\left[\psi''(\theta)\right] \frac{\left(\tilde{\theta} - \mu^{(t)}\right)^2}{2}\right) \tag{4}$$

This iteration has further interesting properties. First of all, the fixed-point of this iteration is unique under our assumption that $\psi(\theta)$ is convex (Challis and Barber [2011]). Furthermore, even when $\psi(\theta)$ is not convex, the ensemble of all fixed-points of the iteration is also the ensemble of extrema of the reverse Kullback-Leibler divergence on the space of Gaussians:

$$KL(q, p) = E_q\left(\log \frac{q(\theta)}{p(\theta)}\right) \tag{5}$$

In other words, the iterative algorithm we have defined computes a Gaussian Variational Bayes (VB) approximation of $p(\theta)$ (Hoffman et al. [2013]). This link between gradient descent and Variational Bayes was originally derived by Opper and Archambeau [2009].

Critically for the rest of this article, the computation of the expected value of the second derivative can be rewritten:

$$E_{q^{(t)}}\left[\psi^{''}\left(\theta\right)\right] = \left[\text{var}_{q^{(t)}}\right]^{-1} E_{q^{(t)}}\left[\psi^{'}\left(\theta\right)\left(\theta - \mu^{(t)}\right)\right] \tag{6}$$

This equality is found by integration by parts and is only true for a Gaussian distribution. When we consider non-Gaussian kernels in the following sections, we will adapt the updating equation (4) using this second form for the quadratic term.

3 Hybrid smoothing minimizes the alpha-divergence

It might seem weird to use a Gaussian smoothing: couldn't we use something that is closer to the target distribution? We will do so by geometrically mixing the target and the current Gaussian approximation, thus building a "hybrid" distribution. For some $0 < \alpha < 1$, construct the α-hybrid (with $Z_{\alpha}^{(t)}$ the normalizing constant):

$$h_{\alpha}^{(t)}\left(\theta\right) = \left(Z_{\alpha}^{(t)}\right)^{-1}\left(q^{(t)}\left(\theta\right)\right)^{\alpha}\left(p\left(\theta\right)\right)^{1-\alpha} \tag{7}$$

which we then use as the smoothing kernel instead of $q^{(t)}\left(\theta\right)$. We also change the centering point for the approximation: the VB update, eq. (4), uses the mean of the current Gaussian approximation; in our new update, we will use the mean of the hybrid distribution instead. The update corresponding to the α-hybrid is then, noting $\mu_h = E_{h_{\alpha}^{(t)}}\left[\theta\right]$ and $v_h = \text{var}_{h_{\alpha}^{(t)}}$:

$$q^{(t+1)}\left(\tilde{\theta}\right) \propto \exp\left(-E_{h_{\alpha}^{(t)}}\left[\psi^{'}\left(\theta\right)\right]\left(\tilde{\theta} - \mu_h\right) - v_h^{-1} E_{h_{\alpha}^{(t)}}\left[\psi^{'}\left(\theta\right)\left(\theta - \mu_h\right)\right]\frac{\left(\tilde{\theta} - \mu_h\right)^2}{2}\right) \tag{8}$$

Once again, this corresponds to both a smoothed Newton's method (by construction) but also, much more surprisingly, to an algorithm that minimizes a specific divergence. Indeed, all fixed-points of the α-hybrid-smoothing iteration eq. (8) are also extrema of the α-divergence (which represents a smooth interpolation between the reverse KL-divergence at $\alpha = 1$ and the forward KL-divergence $KL\left(p, q\right)$ at $\alpha = 0$; Minka [2005]).

4 Classical Expectation Propagation is a smoothed gradient method

In this article, we have shown that we can find extrema of all α-divergences and, critically, of the reverse KL-divergence by performing a form of smoothed gradient descent (with a Hessian correction). This might shed some light on these methods and help further theoretical investigation of these methods. However, the main interest of this approach consists in using this idea of smoothed gradient to give a justification of Expectation Propagation (EP, Minka [2001]) which is much more intuitively satisfying than current derivations of this method (which we will assume that the reader is familiar with due to space constraints; see Minka [2001], Seeger [2005], Bishop [2007]).

In order to apply EP, we have to assume that the target distribution factorizes into "simple" factor functions $f_i\left(\theta\right)$:

$$p\left(\theta\right) = \prod_{i=1}^{n} f_i\left(\theta\right) \tag{9}$$

We will note $\phi_i\left(\theta\right) = -\log\left(f_i\left(\theta\right)\right)$, thus the energy landscape for the gradient descent $\psi\left(\theta\right)$ has been split into n additive components:

$$\psi\left(\theta\right) = \sum_i \phi_i\left(\theta\right) \tag{10}$$

All algorithms we have presented so far compute a single Gaussian approximation for the whole target distribution. This next algorithm will compute for each time-step a "local" Gaussian approximation for each single factor $q_i^{(t)} \approx f_i$. We can then compute a global approximation of the target

by combining them multiplicatively:

$$q_{\text{global}}^{(t)} = \prod_{i=1}^{n} q_i^{(t)} \approx p \tag{11}$$

Each local Gaussian approximations $g_i^{(t)}$ is updated by a smoothed gradient descent on the corresponding local energy landscape ϕ_i. For the smoothing, we construct once more a hybrid distribution. The hybrid for the update of the i^{th} approximation is formed by multiplying the true factor f_i and the current local Gaussian approximation of all other factors $\left(g_i^{(t)}\right)_{j\neq i}$. I.e:

$$h_i^{(t)}(\theta) = \left(Z_i^{(t)}\right)^{-1} f_i(\theta) \prod_{j\neq i} g_j^{(t)}(\theta) \tag{12}$$

The new Gaussian approximation of the factor f_i is then given by, noting $\mu_{h_i} = E_{h_i^{(t)}}[\theta_h]$ and $v_{h_i} = \text{var}_{h_i^{(t)}}$:

$$g_i^{(t+1)}(\theta) \propto \exp\left(-E_{h_i^{(t)}}\left[\phi_i'(\theta_h)\right](\theta-\mu_{h_i}) - v_{h_i}^{-1} E_{h_i^{(t)}}\left[\phi_i'(\theta_h)(\theta_h-\mu_{h_i})\right]\frac{(\theta-\mu_{h_i})^2}{2}\right) \tag{13}$$

This iterating scheme is exactly the same as the classical EP update presented by Minka [2001]. We have thus rephrased the EP iteration from its original computationally-convenient but cryptic form into a smoothed gradient descent which is much more instructive for our intuitive understanding of EP.

5 Why this matters

5.1 Asymptotic behavior

This new perspective on VB, α-divergence minimization and EP can be used to derive several interesting results on the asymptotic behavior of these methods in a painless manner. Indeed, as we have shown, these algorithms all correspond to a smoothed gradient descent. Thus, in all limits in which the smoothing kernel is sufficiently concentrated, the dynamics of these algorithms asymptote to the dynamics of Newton's method.

Furthermore, in all limits in which the smoothing kernel of one method asymptotes to the smoothing kernel of a second, their dynamics asymptote to one another. A trivial example of this is the fact that the dynamics of α-hybrid smoothing (eq. (7)) asymptote to the dynamics of the Gaussian smoothing (eq. (4)) in the limit $\alpha \to 1$ (thus giving a much more intuitive understanding of the result of Dehaene and Barthelmé [2016]). A much more interesting example consists in proving a "folk theorem" on EP asserting that, in the limit of a large number of sites with each one having a negligible contribution to the ensemble, EP corresponds to minimizing the reverse KL divergence (eq. (5)). The present work shows painlessly that this is true in any limit in which all EP hybrids h_i (eq. (12)) asymptote to the current global approximation $q_{\text{global}} = \frac{q_i}{f_i} h_i$ (eq. (11)).

5.2 Intuitive understanding

In this article, we have unified several algorithms as performing smoothed gradient descent. While this might provide the path towards more useful algorithms for minimizing several divergence measures, we believe that the most useful contribution of this work is that it gives practitioners a more intuitive understanding of these algorithms. Indeed, gradient descent is both a fairly intuitive algorithm, matching our physical intuitions of an object sliding around on some energy landscape, and one which has been extremely extensively researched. The addition of smoothing complicates this picture, but only slightly. We thus hope that our results can help promote these methods by shining a new light on them. This contribution is probably the most helpful for EP, given that our iteration (eq. (13)) and classical EP (Minka [2001]) **have the exact same dynamics**, and that current presentations of EP give no intuitions about how the algorithm operates.

References

Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, 1st ed. 2006. corr. 2nd printing 2011 edition, October 2007. ISBN 0387310738. URL http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0387310738.

Edward Challis and David Barber. Concave gaussian variational approximations for inference in large-scale bayesian linear models. In *AISTATS*, pages 199–207, 2011.

Guillaume Dehaene and Simon Barthelmé. Expectation propagation in the large-data limit. *arXiv preprint arXiv:1503.08060*, 2016.

Matthew D Hoffman, David M Blei, Chong Wang, and John William Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, 2013.

T. Minka. Divergence Measures and Message Passing. Technical report, 2005. URL http://research.microsoft.com/en-us/um/people/minka/papers/message-passing/minka-divergence.pdf.

Thomas P. Minka. Expectation Propagation for approximate Bayesian inference. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, pages 362–369, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1-55860-800-1. URL http://portal.acm.org/citation.cfm?id=720257.

Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012. ISBN 0262018020, 9780262018029.

Manfred Opper and Cédric Archambeau. The variational gaussian approximation revisited. *Neural computation*, 21(3):786–792, 2009.

M. Seeger. Expectation Propagation for Exponential Families. Technical report, 2005. URL http://people.mmci.uni-saarland.de/~{}mseeger/papers/epexpfam.pdf.

Appendix of Expectation Propagation performs a smoothed gradient descent

Guillaume Dehaene
Ecole Polytechnique Federale de Lausanne
guillaume.dehaene@gmail.com

In this appendix, we will present all of the technical results underlying our main abstract, which was light on details so as to remain compact.

Throughout this document, we will present various methods to compute a Gaussian approximation of a multivariate target distribution:

$$p(\boldsymbol{\theta}) = \exp(-\psi(\boldsymbol{\theta})) \tag{1}$$

We will note d the dimensionality of the space. All vectors will be represented by bold letters (e.g: $\boldsymbol{\theta}, \boldsymbol{\mu}$) while matrices are bold capitalized letters (e.g: $\boldsymbol{S}, \boldsymbol{B}$).

In contexts where it matters to distinguish the inner and outer-product between vectors, we use the physicists Bra-Ket notation which makes it obvious whether a given vector should be considered as a $(1 * d)$ matrix (in which it is called a "Bra" $\langle\boldsymbol{\theta}|$) or a $(d * 1)$ matrix (in which case it is called a "Ket" $|\boldsymbol{\mu}\rangle$). In this notation, scalar products are represented as, for example: $\langle\boldsymbol{\theta}|\boldsymbol{\mu}\rangle = \sum_{i=1}^{d}\theta_i\mu_i$. A scalar product weighted by some matrix $\boldsymbol{B}$ is represented using: $\langle\boldsymbol{\theta}|\boldsymbol{B}|\boldsymbol{\mu}\rangle = \sum_{i,j}\theta_i\mu_j B_{i,j}$. Outer-products are noted as $|\boldsymbol{\theta}\rangle\langle\boldsymbol{\mu}|$ (which corresponds to the conventional notation $\boldsymbol{\theta}\boldsymbol{\mu}^T$) which gives a matrix: $(|\boldsymbol{\theta}\rangle\langle\boldsymbol{\mu}|)_{i,j} = \theta_i\mu_j$.

Finally, for a multivariate function such as $p(\boldsymbol{\theta})$, we note $\nabla_{\boldsymbol{\theta}}p(\boldsymbol{\theta})$ the gradient, i.e: the vector of the derivatives against each component of the variable. We note $\boldsymbol{H}p(\boldsymbol{\theta})$ the "Hessian" matrix of the second derivatives of the function.

1 The relationship between Variational Bayes and smoothed gradient descent

In this first section, we show that we can minimize the reverse KL divergence between a Gaussian approximation and the target distribution by performing smoothed gradient descent.

1.1 The Gaussian Variational Bayes approximation

First, let us consider computing a Gaussian approximation which minimizes the "reverse" KL divergence to the target. Noting $\mathcal{G}$ the space of all Gaussians, this Gaussian Variational Bayes approximation is defined as:

$$q_{\text{VB}}(\boldsymbol{\theta}) = \text{argmin}_{q\in\mathcal{G}}(q, p) \tag{2}$$

$$= \text{argmin}_{q\in\mathcal{G}}\int q(\boldsymbol{\theta})\log\left[\frac{q(\boldsymbol{\theta})}{p(\boldsymbol{\theta})}\right] \tag{3}$$

Let us now rewrite the objective function. The most common parameterization of a Gaussian distribution is via its mean and its covariance matrix. However, an alternative parameterization which is more relevant in this case is using the "matrix square root" of the covariance (which, in 1D, would

correspond to the standard-deviation). Noting $\boldsymbol{\Sigma}$ the covariance matrix, this "matrix square root" is a solution of:

$$\boldsymbol{S}\boldsymbol{S}^T = \boldsymbol{\Sigma} \tag{4}$$

This parameterization is useful as it enables us to write any Gaussian $\boldsymbol{g}_{\boldsymbol{\mu},\boldsymbol{S}}$with parameters $(\boldsymbol{\mu}, \boldsymbol{S})$ as a translated and shifted version of a Gaussian with mean 0 and covariance the identity matrix (which we note $\boldsymbol{g}_0$):

$$\boldsymbol{g}_{\boldsymbol{\mu},\boldsymbol{S}} = \boldsymbol{S}\boldsymbol{g}_0 + \boldsymbol{\mu} \tag{5}$$

However, note that this parameterization of the space of Gaussian is degenerate: the same Gaussian density corresponds to multiple values of $\boldsymbol{S}$.

Using this parameterization, it is easy to write down the reverse KL divergence as an expected value under the random variable $\boldsymbol{g}_0$:

$$KL\left(q_{\boldsymbol{\mu},\boldsymbol{S}}, p\right) = \int q(\boldsymbol{\theta}) \log\left[\frac{q\left(\boldsymbol{\theta}\right)}{p\left(\boldsymbol{\theta}\right)}\right] \tag{6}$$

$$= \int q(\boldsymbol{\theta})\psi\left(\boldsymbol{\theta}\right) + \int q\left(\boldsymbol{\theta}\right) \log\left(q\left(\boldsymbol{\theta}\right)\right) \tag{7}$$

$$= E\left[\psi\left(\boldsymbol{S}\boldsymbol{g}_0 + \boldsymbol{\mu}\right)\right] - \log\left(\left(2\pi\right)^{d/2} \left|\boldsymbol{S}\right|\right) \tag{8}$$

where $|\boldsymbol{S}|$ is the determinant of the $\boldsymbol{S}$ matrix.

Computing the gradient of the KL divergence against this parameterization of the Gaussians is then a straightforward exercise in vectorial and matrix derivatives (cf: the matrix cookbook Petersen and Pedersen [2012]) which yields:

$$\nabla_{\boldsymbol{\mu}} KL\left(q_{\boldsymbol{\mu},\boldsymbol{S}}, p\right) = E\left[\nabla\psi\left(\boldsymbol{S}\boldsymbol{g}_0 + \boldsymbol{\mu}\right)\right] \tag{9}$$

$$= E\left[\nabla\psi\left(\boldsymbol{g}_{\boldsymbol{\mu},\boldsymbol{S}}\right)\right] \tag{10}$$

$$\nabla_{\boldsymbol{S}} KL\left(q_{\boldsymbol{\mu},\boldsymbol{S}}, p\right) = E\left[\left|\nabla\psi\left(\boldsymbol{S}\boldsymbol{g}_0 + \boldsymbol{\mu}\right)\right\rangle \left\langle\boldsymbol{g}_0\right|\right] - \boldsymbol{S}^{-1} \tag{11}$$

We then integrate by parts eq. (11) (or, alternatively, we use Stein's lemma) to obtain an equation with the second derivative of ψ:

$$\nabla_{\boldsymbol{S}} KL\left(q_{\boldsymbol{\mu},\boldsymbol{S}}, p\right) = E\left[\boldsymbol{H}\psi\left(\boldsymbol{S}\boldsymbol{g}_0 + \boldsymbol{\mu}\right)\boldsymbol{S}^T\right] - \boldsymbol{S}^{-1} \tag{12}$$

$$= E\left[\boldsymbol{H}\psi\left(\boldsymbol{g}_{\boldsymbol{\mu},\boldsymbol{S}}\right)\right]\boldsymbol{S}^T - \boldsymbol{S}^{-1} \tag{13}$$

This gives us a simple characterization of all critical points of the function $\boldsymbol{\mu}, \boldsymbol{S} \rightarrow KL$: they obey the following equations:

$$E\left[\nabla\psi\left(\boldsymbol{g}_{\boldsymbol{\mu},\boldsymbol{S}}\right)\right] = \overrightarrow{0} \tag{14}$$

$$E\left[\boldsymbol{H}\psi\left(\boldsymbol{g}_{\boldsymbol{\mu},\boldsymbol{S}}\right)\right] = \boldsymbol{S}^{-1}\left(\boldsymbol{S}^T\right)^{-1} \tag{15}$$

$$= \boldsymbol{\Sigma}^{-1} \tag{16}$$

These conditions for the Gaussian Variational Bayes approximation of a target distribution were originally derived by Opper and Archambeau [2009]. They express that a Gaussian VB approximation must be such that the expected value of the log-gradient of the target distribution is 0 in all dimensions and that the inverse-variance of the approximation matches the expected value of the log-Hessian of the target distribution.

These conditions are quite naturally linked to the smoothed gradient descent algorithm which we now introduce.

1.2 Smoothed gradient descent

First, let us present the smoothed gradient descent (with a Hessian correction) in detail. It corresponds to the following algorithm.

- Initialize the algorithm with any Gaussian approximation $q^{(0)}$ of the target distribution
- Then loop until convergence:
  1. Compute the mean as well as the expected log-gradient and expected log-Hessian of the target distribution under the current Gaussian approximation $q^{(t)}(\boldsymbol{\theta})$:

$$\boldsymbol{\mu} = E_{q^{(t)}}[\boldsymbol{\theta}] \tag{17}$$

$$E\nabla = E_{q^{(t)}}[\nabla\psi(\boldsymbol{\theta})] \tag{18}$$

$$E\boldsymbol{H} = E_{q^{(t)}}[\boldsymbol{H}\psi(\boldsymbol{\theta})] \tag{19}$$

  2. Compute the new Gaussian approximation using the following formula:

$$q^{(t+1)} = \frac{1}{Z}\exp\left(-\langle E\nabla|\,|(\boldsymbol{\theta}-\boldsymbol{\mu})\rangle - \langle(\boldsymbol{\theta}-\boldsymbol{\mu})|\,\frac{E\boldsymbol{H}}{2}\,|(\boldsymbol{\theta}-\boldsymbol{\mu})\rangle\right) \tag{20}$$

     where Z is a normalizing constant.

We refer to this algorithm as "smoothed gradient descent" since this update exactly matches the update of gradient descent with a Hessian correction starting from $\boldsymbol{\mu} = E_{q^{(t)}}[\boldsymbol{\theta}]$ on the energy landscape given by a smoothing of ψ with a Gaussian kernel $q^{(t)}$ (in equations: $\boldsymbol{\theta} \to \left(\psi \otimes q^{(t)}\right)(\boldsymbol{\theta})$ is the new energy landscape).

This algorithm does not correspond to performing gradient descent on the space Gaussians parameterized by $(\boldsymbol{\mu}, \boldsymbol{S})$ of the reverse KL divergence, but it is closely related as these two algorithms share their fixed-points. Indeed, considering eq. (20) shows that this algorithm is stable if and only if:

- $E\nabla = 0$ so that the mean of the Gaussian approximation does not change.
- $E\boldsymbol{H} = (\mathrm{Cov}_q)^{-1}$ so that the variance of the distribution does not change either.

These stability conditions exactly match the characterization of critical points of $\boldsymbol{\mu}, \boldsymbol{S} \to KL$ (eqs. (14) and (16)). Thus, this smoothed gradient descent algorithm represents an iterative scheme to minimize the reverse KL divergence over the space of Gaussian distributions.

2 Minimizing the α-divergence

In this second section, we introduce the α-divergence, which interpolates between the forward KL divergence and the reverse KL divergence. We show that we can minimize the α-divergence between a Gaussian approximation and the target distribution by performing a smoothed gradient descent which uses an hybrid distribution as the smoothing kernel.

2.1 Defining the α-divergence

Consider two probability densities $(p(\boldsymbol{\theta}), q(\boldsymbol{\theta}))$. The α-divergence between these two distributions is given, for $\alpha \notin \{0, 1\}$:

$$D_\alpha(p, q) = \frac{1}{\alpha(1-\alpha)}\left[1 - \int (p(\boldsymbol{\theta}))^{1-\alpha}(q(\boldsymbol{\theta}))^{\alpha}\, d\boldsymbol{\theta}\right] \tag{21}$$

A few values of α correspond to interesting measures of the differences between p and q (Minka [2005]; note that we are not using the exact same definition of the α-divergence).

- For $\alpha = 1/2$, we recover the squared Hellinger distance:

$$D_{1/2} = 4\left[1 - \int \sqrt{pq}\right] \tag{22}$$

$$= 2\int (\sqrt{p} - \sqrt{q})^2 \tag{23}$$

- For $\alpha = -1$ and $\alpha = 2$, we recover interesting values:

$$D_{-1} = -\frac{1}{2}\left[1 - \int \frac{p^2}{q}\right] \tag{24}$$

$$= \frac{1}{2}\left[\int \frac{p^2}{q} + \int q - 2\int p\right] \tag{25}$$

$$= \frac{1}{2}\int \frac{(p-q)^2}{q} \tag{26}$$

$$D_2 = \frac{1}{2}\int \frac{(p-q)^2}{p} \tag{27}$$

  This corresponds to the two χ^2 distances between the two probability distributions

- Finally, and most interesting, in the limits $\alpha \to 0$ and $\alpha \to 1$, we recover the direct and reverse KL divergences:

$$D_\alpha(p,q) = \frac{1}{\alpha(1-\alpha)}\left[1 - \int p\left(\frac{p}{q}\right)^{-\alpha}\right] \tag{28}$$

$$= \frac{1}{\alpha(1-\alpha)}\left[1 - \int p \exp\left(-\alpha \log \frac{p}{q}\right)\right] \tag{29}$$

$$\approx \frac{1}{\alpha(1-\alpha)}\left[1 - \int p\left(1 + -\alpha \log \frac{p}{q} + \alpha^2 \dots\right)\right] \tag{30}$$

$$\approx \frac{1}{1-\alpha}\int p \log \frac{p}{q} + \alpha \dots \tag{31}$$

$$\lim_{\alpha\to 0} D_\alpha(p,q) = \int p \log \frac{p}{q} \tag{32}$$

$$= KL(p,q) \tag{33}$$

$$\lim_{\alpha\to 1} D_\alpha(p,q) = KL(q,p) \tag{34}$$

  This is the result we are most interested in, since it justifies our earlier remark that the α-divergences interpolate between the forward and reverse KL divergences.

2.2 Minimizing the α-divergence over the space of Gaussians

Once more, we will parametrize Gaussians using the mean and the square-root matrix of the covariance $\boldsymbol{S}$. We can then smartly rewrite our objective function, the α-divergence, by performing a change of variable $\tilde{\boldsymbol{\theta}} = \boldsymbol{S}^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu})$. In this reference frame, the Gaussian density is constant. This gives:

$$D_\alpha(p, q_{\boldsymbol{\mu},\boldsymbol{S}}) = \frac{1}{\alpha(1-\alpha)}\left[1 - \int p\left(\boldsymbol{S}\tilde{\boldsymbol{\theta}} + \boldsymbol{\mu}\right)^{1-\alpha} q_0\left(\tilde{\boldsymbol{\theta}}\right)^{\alpha} \frac{d\tilde{\boldsymbol{\theta}}}{|\boldsymbol{S}|^{1-\alpha}}\right] \tag{35}$$

We can then compute the gradients:

$$\nabla_{\boldsymbol{\mu}} D_\alpha = \frac{1}{\alpha(1-\alpha)}\int (1-\alpha)\nabla\psi\left(\boldsymbol{S}\tilde{\boldsymbol{\theta}} + \boldsymbol{\mu}\right) p\left(\boldsymbol{S}\tilde{\boldsymbol{\theta}} + \boldsymbol{\mu}\right)^{1-\alpha} q_0\left(\tilde{\boldsymbol{\theta}}\right)^{\alpha} \frac{d\tilde{\boldsymbol{\theta}}}{|\boldsymbol{S}|^{1-\alpha}} \tag{36}$$

By now performing the reverse change of variable and returning to $\boldsymbol{\theta}$, we get that the gradient corresponds to an expected value under the hybrid distribution: $h_\alpha(\boldsymbol{\theta}) = Z_\alpha^{-1}(p(\boldsymbol{\theta}))^{1-\alpha}(q_{\boldsymbol{\mu},\boldsymbol{S}}(\boldsymbol{\theta}))^{\alpha}$:

$$\nabla_{\boldsymbol{\mu}} D_\alpha = \frac{1}{\alpha}\int \nabla\psi(\boldsymbol{\theta})\, p(\boldsymbol{\theta})^{1-\alpha}\, q_{\boldsymbol{\mu},\boldsymbol{S}}(\boldsymbol{\theta})^{\alpha}\, d\boldsymbol{\theta} \tag{37}$$

$$= \frac{\left(\int (p(\boldsymbol{\theta}))^{1-\alpha}(q_{\boldsymbol{\mu},\boldsymbol{S}}(\boldsymbol{\theta}))^{\alpha}\, d\boldsymbol{\theta}\right)}{\alpha} E_{h_\alpha}(\nabla\psi(\boldsymbol{\theta})) \tag{38}$$

Similarly, the gradient to $\boldsymbol{S}$ gives another expected value under h_α:

$$\begin{aligned}
\nabla_{\boldsymbol{S}} D_\alpha =& \frac{1}{\alpha} \int \left|\tilde{\boldsymbol{\theta}}\right\rangle \left\langle \nabla\psi\left(\boldsymbol{S}\tilde{\boldsymbol{\theta}}+\boldsymbol{\mu}\right)\right| p\left(\boldsymbol{S}\tilde{\boldsymbol{\theta}}+\boldsymbol{\mu}\right)^{1-\alpha} q_0\left(\tilde{\boldsymbol{\theta}}\right)^{\alpha} \frac{d\tilde{\boldsymbol{\theta}}}{|\boldsymbol{S}|^{1-\alpha}} \\
& - \frac{1}{\alpha}\boldsymbol{S}^{-1} \frac{1}{|\boldsymbol{S}|^{1-\alpha}} \int p\left(\boldsymbol{S}\tilde{\boldsymbol{\theta}}+\boldsymbol{\mu}\right)^{1-\alpha} q_0\left(\tilde{\boldsymbol{\theta}}\right)^{\alpha} d\tilde{\boldsymbol{\theta}} && (39) \\
=& \frac{1}{\alpha}\boldsymbol{S}^{-1} \int |\boldsymbol{\theta}-\boldsymbol{\mu}\rangle \langle \nabla\psi(\boldsymbol{\theta})| \, p(\boldsymbol{\theta})^{1-\alpha} \, q_{\boldsymbol{\mu},\boldsymbol{S}}(\boldsymbol{\theta})^{\alpha} \, d\boldsymbol{\theta} \\
& - \frac{1}{\alpha}\boldsymbol{S}^{-1} \left(\int (p(\boldsymbol{\theta}))^{1-\alpha} \left(q_{\boldsymbol{\mu},\boldsymbol{S}}(\boldsymbol{\theta})\right)^{\alpha} d\boldsymbol{\theta} \right) && (40) \\
=& \frac{\left(\int (p(\boldsymbol{\theta}))^{1-\alpha} \left(q_{\boldsymbol{\mu},\boldsymbol{S}}(\boldsymbol{\theta})\right)^{\alpha} d\boldsymbol{\theta} \right)}{\alpha} \boldsymbol{S}^{-1} \left[E_{h_\alpha}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}\rangle \langle \nabla\psi(\boldsymbol{\theta})| \right) - \boldsymbol{I}_{d*d} \right] && (41)
\end{aligned}$$

Thus the critical points of the function $\boldsymbol{\mu}, \boldsymbol{S} \rightarrow D_\alpha\left(p, q_{\boldsymbol{\mu},\boldsymbol{S}}\right)$ obey the following two equations:

$$E_{h_\alpha}\left(\nabla\psi(\boldsymbol{\theta})\right) = 0 \tag{42}$$

$$E_{h_\alpha}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}\rangle \langle \nabla\psi(\boldsymbol{\theta})| \right) = \boldsymbol{I}_{d*d} \tag{43}$$

Furthermore, for any probability distribution (with fast decrease in the tails) by integration by parts we have that (illustrating the result with $p(\boldsymbol{\theta}) \propto \exp\left(-\psi(\boldsymbol{\theta})\right)$) $\forall \boldsymbol{v}$:

$$E_p\left(\nabla\psi(\boldsymbol{\theta})\right) = 0 \tag{44}$$

$$E_p\left(|\boldsymbol{\theta}-\boldsymbol{v}\rangle \langle \nabla\psi(\boldsymbol{\theta})| \right) = \boldsymbol{I}_{d*d} \tag{45}$$

By applying these two equations to the hybrid distribution at a critical point $h_\star = p^{1-\alpha}\left(q_{\boldsymbol{\mu}_\star,\boldsymbol{S}_\star}\right)^{\alpha}$ with the two eqs. (42) and (43), we get that at a critical point, the mean and variance of the hybrid and of the Gaussian approximation are identical. Noting $\boldsymbol{\Sigma}_\star = \boldsymbol{S}_\star \left(\boldsymbol{S}_\star\right)^T$ the covariance of the Gaussian distribution at the critical point:

$$\begin{aligned}
0 &= E_{h_\star}\left((1-\alpha)\nabla\psi(\boldsymbol{\theta}) + \alpha\boldsymbol{\Sigma}_\star^{-1}\left(\boldsymbol{\theta}-\boldsymbol{\mu}_\star\right) \right) && (46) \\
&= 0 + E_{h_\star}\left(\alpha\boldsymbol{\Sigma}_\star^{-1}\left(\boldsymbol{\theta}-\boldsymbol{\mu}_\star\right) \right) && (47) \\
&= \alpha\boldsymbol{\Sigma}_\star^{-1}\left(E_{h_\star}(\boldsymbol{\theta}) - \boldsymbol{\mu}_\star\right) && (48) \\
E_{h_\star}(\boldsymbol{\theta}) &= \boldsymbol{\mu}_\star && (49)
\end{aligned}$$

And (using $\boldsymbol{v} = \boldsymbol{\mu}_\star$):

$$\begin{aligned}
\boldsymbol{I}_{d*d} &= E_{h_\star}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}_\star\rangle \left\langle \left((1-\alpha)\nabla\psi(\boldsymbol{\theta}) + \alpha\boldsymbol{\Sigma}_\star^{-1}\left(\boldsymbol{\theta}-\boldsymbol{\mu}_\star\right) \right) \right| \right) && (50) \\
&= (1-\alpha) E_{h_\star}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}_\star\rangle \langle \nabla\psi(\boldsymbol{\theta})| \right) + \alpha E_{h_\star}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}_\star\rangle \langle \boldsymbol{\theta}-\boldsymbol{\mu}_\star| \, \boldsymbol{\Sigma}_\star^{-1} \right) && (51) \\
&= (1-\alpha)\boldsymbol{I}_{d*d} + \alpha E_{h_\star}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}_\star\rangle \langle \boldsymbol{\theta}-\boldsymbol{\mu}_\star| \, \boldsymbol{\Sigma}_\star^{-1} \right) && (52) \\
\alpha\boldsymbol{I}_{d*d} &= \alpha E_{h_\star}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}_\star\rangle \langle \boldsymbol{\theta}-\boldsymbol{\mu}_\star| \, \boldsymbol{\Sigma}_\star^{-1} \right) && (53) \\
\boldsymbol{I}_{d*d} &= \mathbf{Cov}_{h_\star}(\boldsymbol{\theta}) \, \boldsymbol{\Sigma}_\star^{-1} && (54) \\
\boldsymbol{\Sigma}_\star &= \mathbf{Cov}_{h_\star}(\boldsymbol{\theta}) && (55)
\end{aligned}$$

These last two points were already highlighted by Minka [2005].

Finally, we get a slight variant of the equations characterizing the critical points by combining eqs. (43) and (55):

$$\left(\mathrm{Cov}_{h_\star}\right)^{-1}(\boldsymbol{\theta}) \, E_{h_\star}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}_\star\rangle \langle \nabla\psi(\boldsymbol{\theta})| \right) = \boldsymbol{\Sigma}_\star^{-1} \tag{56}$$

2.3 A smoothed gradient descent

We will now propose a smoothed gradient descent algorithm such that all fixed-points of the algorithm will also be critical points of $\boldsymbol{\mu}, \boldsymbol{S} \rightarrow D_\alpha\left(p, q_{\boldsymbol{\mu},\boldsymbol{S}}\right)$.

Consider the following algorithm:

- Initialize the algorithm with any Gaussian approximation $q^{(0)}$ of the target distribution
- Then loop until convergence:
  1. Compute the current hybrid approximation of the target distribution:
  $$h_\alpha(\boldsymbol{\theta}) = Z_\alpha^{-1} p(\boldsymbol{\theta})^{1-\alpha} \left(q^{(t)}(\boldsymbol{\theta})\right)^\alpha \tag{57}$$
  2. Compute the following expected values under the current hybrid approximation $h_\alpha(\boldsymbol{\theta})$:
  $$\boldsymbol{\mu} = E_{h_\alpha}[\boldsymbol{\theta}] \tag{58}$$
  $$E\nabla = E_{h_\alpha}[\nabla\psi(\boldsymbol{\theta})] \tag{59}$$
  $$E\boldsymbol{H} = (\mathrm{Cov}_{h_\alpha})^{-1} E_{h_\alpha}[|\boldsymbol{\theta} - \mu\rangle \langle\nabla\psi(\boldsymbol{\theta})|] \tag{60}$$
  3. Compute the new Gaussian approximation using the following formula:
  $$q^{(t+1)} = \frac{1}{Z_q} \exp\left(-\langle E\nabla| |\boldsymbol{\theta} - \boldsymbol{\mu}\rangle - \langle\boldsymbol{\theta} - \boldsymbol{\mu}| \frac{E\boldsymbol{H}}{2} |\boldsymbol{\theta} - \boldsymbol{\mu}\rangle\right) \tag{61}$$
  where Z_q is a normalizing constant.

Any fixed-point of this iteration must obey the following equalities:

$$E_{h_\star}[\nabla\psi(\boldsymbol{\theta})] = 0 \tag{62}$$
$$(\mathbf{Cov}_{h_\star})^{-1} E_{h_\alpha}[|\boldsymbol{\theta} - \mu\rangle \langle\nabla\psi(\boldsymbol{\theta})|] = (\mathbf{Cov}_{q_\star})^{-1} \tag{63}$$

which is identical to the equations obeyed by critical points of $\boldsymbol{\mu}, \boldsymbol{S} \to D_\alpha$ (eqs. (42) and (56)).

3 Expectation Propagation

Finally, we come to Expectation Propagation (EP).

In order to be able to apply EP, we have to further specify a factorization of the target distribution:

$$p(\boldsymbol{\theta}) = \prod_{i=1}^{n} f_i(\boldsymbol{\theta}) \tag{64}$$

Given this factorization, we can then compute "the EP approximation of the target distribution of $p(\boldsymbol{\theta})$ factorized as $p = \prod f_i$" or, for short, the EP approximation of p.

We will note $\phi_i(\boldsymbol{\theta}) = -\log(f_i(\boldsymbol{\theta}))$.

3.1 The EP iteration

All algorithms we have presented so far seek a global Gaussian approximation of the target distribution. EP differs from this by seeking to find instead n local Gaussian approximation q_i to approximate each factor: $q_i \approx f_i$.

These approximations are improved iteratively according to:

- Initialize the algorithm with any n local Gaussian approximations $q_i^{(0)}(\boldsymbol{\theta})$
- Then loop until convergence:
  1. Select a subset $\mathcal{I} \subset [1, n]$ of indices[1].
  2. For all $i \in \mathcal{I}$ in parallel:

[1]Several choices are possible for the selection of the subset. The two most frequent are selecting a single index $\mathcal{I} = \{i_0\}$ (corresponding to a sequential variant of EP, as originally proposed in Minka [2001]) or the full range $\mathcal{I} = [1, n]$ (corresponding to the more modern parallel variant of EP). Minibatch or asynchronous variants of EP are also possible. Of course, the sequence of subsets should be such that each factor-approximation q_i is updated regularly.

(a) Compute the hybrid distribution:

$$h_i(\boldsymbol{\theta}) = Z_{h_i}^{-1} f_i(\boldsymbol{\theta}) \prod_{j \neq i} q_j^{(t)}(\boldsymbol{\theta}) \tag{65}$$

(b) Compute the mean and covariance of the hybrid: $(\boldsymbol{\mu}_{h_i}, \mathbf{Cov}_{h_i})$

(c) Compute a Gaussian distribution with that mean and variance. This Gaussian is the moment-matched Gaussian approximation of the hybrid:

$$q_{h_i}(\boldsymbol{\theta}) = Z_{q_{h_i}}^{-1} \exp\left(-\frac{1}{2} \langle \boldsymbol{\theta} - \boldsymbol{\mu}_{h_i} | \mathbf{Cov}_{h_i} | \boldsymbol{\theta} - \boldsymbol{\mu}_{h_i} \rangle\right) \tag{66}$$

(d) Compute the new local Gaussian approximation of $f_i(\boldsymbol{\theta})$ given by:

$$q_i^{(t+1)}(\boldsymbol{\theta}) \propto \frac{q_{h_i}(\boldsymbol{\theta})}{\prod_{j \neq i} q_j^{(t)}(\boldsymbol{\theta})} \tag{67}$$

For a more extensive presentation of EP and the EP iteration, we refer the interested reader to Minka [2001], Seeger [2005], Bishop [2007].

3.2 A smoothed gradient descent

Interestingly, the key step of the EP algorithm: the computation of the new local approximation $q_i^{(t+1)}(\boldsymbol{\theta})$ from its local target $f_i(\boldsymbol{\theta})$ and the current values of the other local approximations $\left(q_j^{(t)}\right)_{j \neq i}$, corresponds **exactly** to a smoothed gradient descent.

Theorem 1. *Smoothed gradient representation of the EP iteration.*

In the EP iteration, the new local approximation $q_i^{(t+1)}(\boldsymbol{\theta})$ can also be written as:

$$\boldsymbol{\mu} = E_{h_i}[\boldsymbol{\theta}] \tag{68}$$

$$E\nabla = E_{h_i}[\nabla \phi_i(\boldsymbol{\theta})] \tag{69}$$

$$E\boldsymbol{H} = (\boldsymbol{Cov}_{h_i})^{-1} E_{h_i}[|\boldsymbol{\theta} - \mu\rangle \langle \nabla \phi_i(\boldsymbol{\theta})|] \tag{70}$$

$$q_i^{(t+1)} \propto \exp\left(-\langle E\nabla | \boldsymbol{\theta} - \boldsymbol{\mu}\rangle - \langle \boldsymbol{\theta} - \boldsymbol{\mu}| \frac{E\boldsymbol{H}}{2} | \boldsymbol{\theta} - \boldsymbol{\mu}\rangle\right) \tag{71}$$

For thoroughness, it is also important to mention that:

$$E\boldsymbol{H} \approx E_{h_i}(\boldsymbol{H}\phi_i(\boldsymbol{\theta})) \tag{72}$$

in the limit in which the hybrid distribution is a strongly log-concave distribution with minimum curvature tending to $+\infty$ (Dehaene and Barthelmé [2016]).

This theorem represents the key innovation of the present work. This alternative formulation of the EP update might prove useful in deriving better computational variants of the EP iteration. However, we believe that this result is most important for theoretical investigations of the EP algorithm. Indeed, it provides a link between EP and the well-understood gradient descent algorithms. Furthermore, this result also enables users of EP to have a more intuitively satisfying presentation of how EP operates.

Proof. This result is actually fairly simple to prove. Indeed, by the definition of q_{h_i}, h_i and q_{h_i} have the same mean and variance: $(\boldsymbol{\mu}_{h_i}, \mathbf{Cov}_{h_i})$.

Now, note:

$$q_{-i}(\boldsymbol{\theta}) = \prod_{j \neq i} q_j^{(t)}(\boldsymbol{\theta}) \tag{73}$$

$$\propto \exp\left(-\langle \boldsymbol{r}_{-i} | \boldsymbol{\theta} - \boldsymbol{\mu}\rangle - \frac{1}{2} \langle \boldsymbol{\theta} - \boldsymbol{\mu} | \boldsymbol{B}_{-i} | \boldsymbol{\theta} - \boldsymbol{\mu}\rangle\right) \tag{74}$$

and:

$$q_i^{(t+1)} \propto \exp\left(-\langle \boldsymbol{r}_i| \, |\boldsymbol{\theta}-\boldsymbol{\mu}\rangle - \frac{1}{2}\langle \boldsymbol{\theta}-\boldsymbol{\mu}| \, \boldsymbol{B}_i \, |\boldsymbol{\theta}-\boldsymbol{\mu}\rangle\right) \tag{75}$$

Our objective is to find a simple expression for $\boldsymbol{r}_i$ and $\boldsymbol{B}_i$. This can be done by making use of what we call "Stein relationships" (in honor of Charles Stein and his Stein's lemma, which corresponds to the Gaussian case). Indeed, by integration by parts, we find that for any probability distribution (with fast decrease in the tails):

$$E_p\left(\nabla\psi\left(\boldsymbol{\theta}\right)\right) = 0 \tag{76}$$

$$E_p\left(|\boldsymbol{\theta}-\boldsymbol{\mu}\rangle\langle\nabla\psi\left(\boldsymbol{\theta}\right)|\right) = \boldsymbol{I}_{d*d} \tag{77}$$

Applying the first Stein relationship (eq. (76)) to the hybrid $h_i \propto f_i q_{-i}$, we get that:

$$\begin{aligned} 0 &= E_{h_i}\left(\nabla\left(\log h_i\left(\boldsymbol{\theta}\right)\right)\right) && (78)\\ &= E_{h_i}\left(\nabla\phi_i\left(\boldsymbol{\theta}\right) + \left(\boldsymbol{r}_{-i} + \boldsymbol{B}_{-i}\left(\boldsymbol{\theta}-\boldsymbol{\mu}\right)\right)\right) && (79)\\ &= E_{h_i}\left(\nabla\phi_i\left(\boldsymbol{\theta}\right)\right) + \boldsymbol{r}_{-i} + \boldsymbol{B}_{-i}\left(E_{h_i}\left(\boldsymbol{\theta}\right)-\boldsymbol{\mu}\right) && (80)\\ &= E_{h_i}\left(\nabla\phi_i\left(\boldsymbol{\theta}\right)\right) + \boldsymbol{r}_{-i} + 0 && (81) \end{aligned}$$

Now, we apply the first Stein relationship (eq. (76)) to the moment-matched approximation of h_i: $q_{h_i} \propto q_i q_{-i}$. We get:

$$\begin{aligned} 0 &= E_{q_{h_i}}\left(\nabla\left(\log q_{h_i}\left(\boldsymbol{\theta}\right)\right)\right) && (82)\\ &= E_{q_{h_i}}\left(\left(\boldsymbol{r}_i + \boldsymbol{B}_i\left(\boldsymbol{\theta}-\boldsymbol{\mu}\right)\right) + \left(\boldsymbol{r}_{-i} + \boldsymbol{B}_{-i}\left(\boldsymbol{\theta}-\boldsymbol{\mu}\right)\right)\right) && (83)\\ &= \boldsymbol{r}_i + \boldsymbol{r}_{-i} && (84) \end{aligned}$$

By combining eqs. (81) and (84), we get:

$$\boldsymbol{r}_i = E_{h_i}\left(\nabla\phi_i\left(\boldsymbol{\theta}\right)\right) \tag{85}$$

QED.

The proof for $\boldsymbol{B}_i = \left(\mathbf{Cov}_{h_i}\right)^{-1} E_{h_i}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}\rangle\langle\nabla\phi_i\left(\boldsymbol{\theta}\right)|\right)$ proceeds the exact same way, but by using the second Stein relationship (eq. (77)) to h_i and q_{h_i} yielding:

$$\begin{aligned} \boldsymbol{I}_{d*d} &= E_{h_i}\left(|\boldsymbol{\theta}-\boldsymbol{\mu}\rangle\langle\nabla\phi_i\left(\boldsymbol{\theta}\right)|\right) + \mathbf{Cov}_{h_i}\boldsymbol{B}_{-i} && (86)\\ &= \mathbf{Cov}_{h_i}\boldsymbol{B}_i + \mathbf{Cov}_{h_i}\boldsymbol{B}_{-i} && (87) \end{aligned}$$

Combining these equations yields the claimed result, concluding the proof. □

4 Why smoothed gradient descent leads to good approximations

The Stein relationships we have just used in the preceeding proof (eqs. 76 and 77) also provide an intuition as to why the EP and VB Gaussian approximations (as well as the D_α minima) provide good approximations of the target distribution. Indeed, these Stein relationships read:

$$E_p\left(\nabla\psi\left(\boldsymbol{\theta}\right)\right) = 0 \tag{88}$$

$$E_p\left(|\boldsymbol{\theta}-E_p\left(\boldsymbol{\theta}\right)\rangle\langle\nabla\psi\left(\boldsymbol{\theta}\right)|\right) = \boldsymbol{I}_{d*d} \tag{89}$$

A VB Gaussian approximation q_{VB} (i.e: a minimum of $KL\left(p,q\right)$) obeys very similar relationships:

$$E_{q_{VB}}\left(\nabla\psi\left(\boldsymbol{\theta}\right)\right) = 0 \tag{90}$$

$$E_{q_{VB}}\left(|\boldsymbol{\theta}-E_p\left(\boldsymbol{\theta}\right)\rangle\langle\nabla\psi\left(\boldsymbol{\theta}\right)|\right) = \boldsymbol{I}_{d*d} \tag{91}$$

Thus, q_{VB} and p have the same expected value for the functions: $\nabla\psi\left(\boldsymbol{\theta}\right)$ and $|\boldsymbol{\theta}-E_p\left(\boldsymbol{\theta}\right)\rangle\langle\nabla\psi\left(\boldsymbol{\theta}\right)|$ (which corresponds to $n\left(n+1\right)$ scalar equalities).

Meanwhile, a EP fixed-point with hybrids $h_i^\star$ obeys:

$$\sum_i E_{h_i^\star}\left(\nabla\phi_i\left(\boldsymbol{\theta}\right)\right) = 0 \tag{92}$$

$$\sum_i E_{h_i^\star}\left(\left|\boldsymbol{\theta} - E_p\left(\boldsymbol{\theta}\right)\right\rangle\left\langle\nabla\phi_i\left(\boldsymbol{\theta}\right)\right|\right) = \boldsymbol{I}_{d*d} \tag{93}$$

which means that the density defined by $n^{-1}\sum h_i^\star$ (where n is the number of factors/factor-approximations/hybrids) has the same expected value for the two functions we are considering.

Using this equality (or almost equality) of these expected values, we were able last year to show that, in the classical Bayesian large-data limit, EP and VB Gaussian approximations of a posterior distribution are:

- asymptotically valid: EP and VB both correctly estimate the mean and variance of the target distribution
- better than the alternative Laplace approximation: the error of the estimate of the mean is an order of magnitude better for EP and VB than for Laplace

This result was, however, derived under restrictive assumption and needs to be improved (Dehaene and Barthelmé [2015]).

Proof.

References

Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, 1st ed. 2006. corr. 2nd printing 2011 edition, October 2007. ISBN 0387310738. URL http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0387310738.

Guillaume Dehaene and Simon Barthelmé. Expectation propagation in the large-data limit. *arXiv preprint arXiv:1503.08060*, 2016.

Guillaume P Dehaene and Simon Barthelmé. Bounding errors of Expectation-Propagation. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 244–252. Curran Associates, Inc., 2015. URL http://papers.nips.cc/paper/5912-bounding-errors-of-expectation-propagation.pdf.

T. Minka. Divergence Measures and Message Passing. Technical report, 2005. URL http://research.microsoft.com/en-us/um/people/minka/papers/message-passing/minka-divergence.pdf.

Thomas P. Minka. Expectation Propagation for approximate Bayesian inference. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, pages 362–369, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1-55860-800-1. URL http://portal.acm.org/citation.cfm?id=720257.

Manfred Opper and Cédric Archambeau. The variational gaussian approximation revisited. *Neural computation*, 21(3):786–792, 2009.

K. B. Petersen and M. S. Pedersen. The matrix cookbook, nov 2012. URL http://www2.imm.dtu.dk/pubdb/p.php?3274. Version 20121115.

M. Seeger. Expectation Propagation for Exponential Families. Technical report, 2005. URL http://people.mmci.uni-saarland.de/~{}mseeger/papers/epexpfam.pdf.

□