

Perturbative Corrections for Approximate Inference in Gaussian Latent Variable Models

Manfred Opper

*Department of Computer Science
Technische Universität Berlin
D-10587 Berlin, Germany*

OPPERM@CS.TU-BERLIN.DE

Ulrich Paquet

*Microsoft Research Cambridge
Cambridge CB1 2FB, United Kingdom*

ULRIPA@MICROSOFT.COM

Ole Winther

*Informatics and Mathematical Modelling
Technical University of Denmark
DK-2800 Lyngby, Denmark*

OWI@IMM.DTU.DK

Editor: Neil Lawrence

Abstract

Expectation Propagation (EP) provides a framework for approximate inference. When the model under consideration is over a latent Gaussian field, with the approximation being Gaussian, we show how these approximations can systematically be corrected. A perturbative expansion is made of the exact but intractable correction, and can be applied to the model's partition function and other moments of interest. The correction is expressed over the higher-order cumulants which are neglected by EP's local matching of moments. Through the expansion, we see that EP is correct to first order. By considering higher orders, corrections of increasing polynomial complexity can be applied to the approximation. The second order provides a correction in quadratic time, which we apply to an array of Gaussian process and Ising models. The corrections generalize to arbitrarily complex approximating families, which we illustrate on tree-structured Ising model approximations. Furthermore, they provide a polynomial-time assessment of the approximation error. We also provide both theoretical and practical insights on the exactness of the EP solution.

Keywords: expectation consistent inference, expectation propagation, perturbation correction, Wick expansions, Ising model, Gaussian process

1. Introduction

Expectation Propagation (EP) (Opper and Winther, 2000; Minka, 2001a,b) is part of a rich family of variational methods, which approximate the sums and integrals required for exact probabilistic inference by an optimization problem. Variational methods are perfectly amenable to probabilistic graphical models, as the nature of the optimization problem often allows it to be distributed across a graph. By relying on local computations on a graph, inference in very large probabilistic models becomes feasible.

Being an approximation, some error may invariably be introduced. This paper is specifically concerned with the error that arises when a Gaussian approximating family is used, and lays a systematic foundation for examining and correcting these errors. It follows on earlier work by the

authors (Opper et al., 2009). The error that arises when the free energy (the negative logarithm of the partition function or normalizer of the distribution) is approximated, may for instance be written as a Taylor expansion (Opper et al., 2009; Paquet et al., 2009). A pleasing property of EP is that, at its stationary point, the first order term of such an expansion is zero. Furthermore, the quality of the approximation can then be ascertained in polynomial time by including corrections beyond the first order, or beyond the standard EP solution. In general, the corrections improve the approximation when they are comparatively small, but can also leave a question mark on the quality of approximation when the lower-order terms are large.

The approach outlined here is by no means unique in correcting the approximation, as is evinced by cluster-based expansions (Paquet et al., 2009), marginal corrections for EP (Cseke and Heskes, 2011) and the Laplace approximation (Rue et al., 2009), and corrections to Loopy Belief Propagation (Chertkov and Chernyak, 2006; Sudderth et al., 2008; Welling et al., 2012).

1.1 Overview

EP is introduced in a general way in Section 3, making it clear how various degrees of complexity can be included in its approximating structure. The partition function will be used throughout the paper to explain the necessary machinery for correcting any moments of interest. In the experiments, corrections to the marginal and predictive means and variances are also shown, although the technical details for correcting moments beyond the partition function are relegated to Appendix D. The Ising model, which is cast as a Gaussian latent variable model in Section 2, will furthermore be used as a running example throughout the paper.

The key to obtaining a correction lies in isolating the “intractable quantity” from the “tractable part” (or EP solution) in the true problem. This is done by considering the cumulants of both: as EP locally matches lower-order cumulants like means and variances, the “intractable part” exists as an expression over the higher-order cumulants which are neglected by EP. This process is outlined in Section 4, which concludes with two useful results: a shift of the “intractable part” to be an average over complex Gaussian variables with *zero* diagonal relation matrix, and Wick’s theorem, which allows us to evaluate the expectations of polynomials under centered Gaussian measures. As a last stage, the “intractable part” is expanded in Sections 5 and 7 to obtain corrections to various orders. In Section 6, we provide a theoretical analysis of the radius of convergence of these expansions.

Experimental evidence is presented in Section 8 on Gaussian process (GP) classification and (non-Gaussian) GP regression models. An insightful counterexample where EP diverges under increasing data, is also presented. Ising models are examined in Section 9.

Numerous additional examples, derivations, and material are provided in the appendices. Details on different EP approximations can be found in Appendix A, while corrections to tree-structured approximations are provided in Appendix B. In Appendix C we analytically show that the correction to a tractable example is zero. The main body of the paper deals with corrections to the partition function, while corrections to marginal moments are left to Appendix D. Finally, useful calculations of certain cumulants appear in Appendix E.

2. Gaussian Latent Variable Models

Let $\mathbf{x} = (x_1, \dots, x_N)$ be an unobserved random variable with an intractable distribution $p(\mathbf{x})$. In the Gaussian latent variable model (GLVM) considered in this paper, terms $t_n(x_n)$ are combined over a

quadratic exponential $f_0(\mathbf{x})$ to give

$$p(\mathbf{x}) = \frac{1}{Z} \prod_{n=1}^N t_n(x_n) f_0(\mathbf{x}) \quad (1)$$

with partition function (normalizer)

$$Z = \int \prod_{n=1}^N t_n(x_n) f_0(\mathbf{x}) \, d\mathbf{x} .$$

This model encapsulates many important methods used in statistical inference. As an example, f_0 can encode the covariance matrix of a Gaussian process (GP) prior on latent function observations x_n . In the case of GP classification with a class label $y_n \in \{-1, +1\}$ on a latent function evaluation x_n , the terms are typically probit link functions, for example

$$p(\mathbf{x}) = \frac{1}{Z} \prod_{n=1}^N \Phi(y_n x_n) \mathcal{N}(\mathbf{x}; \mathbf{0}, \mathbf{K}) . \quad (2)$$

The probit function is the standard cumulative Gaussian density $\Phi(x) = \int_{-\infty}^x \mathcal{N}(z; 0, 1) \, dz$. In this example, the partition function is not analytically tractable but for the one-dimensional case $N = 1$.

An Ising model can be constructed by letting the terms t_n restrict x_n to ± 1 (through Dirac delta functions). By introducing the symmetric coupling matrix \mathbf{J} and field $\boldsymbol{\theta}$ into f_0 , an Ising model can be written as

$$p(\mathbf{x}) = \frac{1}{Z} \prod_{n=1}^N \left[\frac{1}{2} \delta(x_n + 1) + \frac{1}{2} \delta(x_n - 1) \right] \exp \left\{ \frac{1}{2} \mathbf{x}^T \mathbf{J} \mathbf{x} + \boldsymbol{\theta}^T \mathbf{x} \right\} . \quad (3)$$

In the Ising model, the partition function Z is intractable, as it sums $f_0(\mathbf{x})$ over 2^N binary values of \mathbf{x} . In the variational approaches, the intractability is addressed by allowing approximations to Z and other marginal distributions, decreasing the computational complexity from being exponential to polynomial in N , which is typically cubic for EP.

3. Expectation Propagation

An approximation to Z can be made by allowing $p(\mathbf{x})$ in Equation (1) to factorize into a product of *factors* f_a . This factorization is not unique, and the structure of the factorization of $p(\mathbf{x})$ defines the complexity of the resulting approximation, resulting in different structures in the approximating distribution. Where GLVMs are concerned, a natural and computationally convenient choice is to use Gaussian factors g_a , and as such, the approximating distribution $q(\mathbf{x})$ in this paper will be Gaussian. Appendix A summarizes a number of factorizations for Gaussian approximations.

The tractability of the resulting inference method imposes a pragmatic constraint on the choice of factorization; in the extreme case $p(\mathbf{x})$ could be chosen as a single factor and inference would be exact. For the model in Equation (1), a three-term product may be factorized as $(t_1)(t_2)(t_3)$, which gives the typical GP setup. When a division is introduced and the term product factorizes as $(t_1 t_2)(t_2 t_3)/(t_2)$, the resulting free energy will be that of the tree-structured EC approximation (Opper and Winther, 2005). To therefore allow for regrouping, combining, splitting, and dividing terms, a power D_a is associated with each f_a , such that

$$p(\mathbf{x}) = \frac{1}{Z} \prod_a f_a(\mathbf{x})^{D_a} \quad (4)$$

with intractable normalization (or partition function) $Z = \int \prod_a f_a(\mathbf{x})^{D_a} d\mathbf{x}$.¹ Appendix A shows how the introduction of D_a lends itself to a clear definition of tree-structured and more complex approximations.

To define an approximation to p , terms g_a , which typically take an exponential family form, are chosen such that

$$q(\mathbf{x}) = \frac{1}{Z_q} \prod_a g_a(\mathbf{x})^{D_a} \tag{5}$$

has the same structure as p 's factorization. Although not shown explicitly, f_a and g_a have a dependence on the *same* subset of variables \mathbf{x}_a . The optimal parameters of the g_a -term approximations are found through a set of auxiliary *tilted* distributions, defined by

$$q_a(\mathbf{x}) = \frac{1}{Z_a} \left(\frac{q(\mathbf{x})f_a(\mathbf{x})}{g_a(\mathbf{x})} \right). \tag{6}$$

Here a *single* approximating term g_a is replaced by an original term f_a . Assuming that this replacement leaves q_a still tractable, the parameters in g_a are determined by the condition that $q(\mathbf{x})$ and all $q_a(\mathbf{x})$ should be made as similar as possible. This is usually achieved by requiring that these distributions share a set of generalised moments which usually coincide with the sufficient statistics of the exponential family. For example with sufficient statistics $\phi(\mathbf{x})$ we require that

$$\langle \phi(\mathbf{x}) \rangle_{q_a} = \langle \phi(\mathbf{x}) \rangle_q \quad \text{for all } a. \tag{7}$$

Note that those factors f_a in $p(\mathbf{x})$ which are already in the exponential family, such as the Gaussian terms in examples above, can trivially be solved for by setting $g_a = f_a$. The partition function associated with this approximation is

$$Z_{EP} = Z_q \prod_a Z_a^{D_a}. \tag{8}$$

Appendix A.2 shows that the moment-matching conditions must hold at a stationary point of $\log Z_{EP}$. The EP algorithm iteratively updates the g_a -terms by enforcing q to share moments with each of the tilted distributions q_a ; on reaching a fixed point all moments match according to Equation (7) (Minka, 2001a,b). Although Z_{EP} is defined in the terminology of EP, other algorithms may be required to solve for the fixed point, and Z_{EP} , as a free energy, can be derived from the saddle point of a set of self-consistent (moment-matching) equations (Oppel and Winther, 2005; van Gerven et al., 2010; Seeger and Nickisch, 2010). We next make EP concrete by applying it to the Ising model, which will serve as a running example in the paper. The section is finally concluded with a discussion of the interpretation of EP.

3.1 EP for Ising Models

The Ising model in Equation (3) will be used as a running example throughout this paper. To make the technical developments more concrete, we will consider both the N -variate and bivariate cases. The bivariate case can be solved analytically, and thus allows for a direct comparison to be made between the exact and approximate solutions.

We use the factorized approximation as a running example, dividing $p(\mathbf{x})$ in Equation (3) into $N + 1$ factors with $f_0(\mathbf{x}) = \exp\{\frac{1}{2}\mathbf{x}^T \mathbf{J} \mathbf{x} + \boldsymbol{\theta}^T \mathbf{x}\}$ and $f_n(x_n) = t_n(x_n) = \frac{1}{2}\delta(x_n + 1) + \frac{1}{2}\delta(x_n - 1)$, for

1. The factorization and EP energy function is expressed here in the form of Power EP (Minka, 2004).

$n = 1, \dots, N$ (see Appendix A for generalizations). We consider the Gaussian exponential family such that $g_n(x_n) = \exp\{\lambda_{n1}x_n - \frac{1}{2}\lambda_{n2}x_n^2\}$ and $g_0(\mathbf{x}) = f_0(\mathbf{x})$. The approximating distribution from Equation (5), $q(\mathbf{x}) \propto f_0(\mathbf{x}) \prod_{n=1}^N g_n(x_n)$, is thus a *full* multivariate Gaussian density, which we write as $q(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$.

3.1.1 MOMENT MATCHING

The moment matching condition in Equation (7) involves only the mean and variance if $q(\mathbf{x})$ fully factorizes according to $p(\mathbf{x})$'s terms. We therefore only need to match the mean and variances of marginals of $q(\mathbf{x})$ and the tilted distribution $q_n(\mathbf{x})$ in Equation (6). The tilted distribution may be decomposed into a Gaussian and a discrete part as $q_n(\mathbf{x}) = q_n(\mathbf{x}_{\setminus n}|x_n)q_n(x_n)$, where the vector $\mathbf{x}_{\setminus n}$ consists of all variables apart from x_n . We may marginalize out $\mathbf{x}_{\setminus n}$ and write $q_n(x_n)$ in terms of two factors:

$$q_n(x_n) \propto \underbrace{\frac{1}{2} [\delta(x_n + 1) + \delta(x_n - 1)]}_{f_n(\mathbf{x})=t_n(x_n)} \underbrace{\exp\left\{\gamma x_n - \frac{1}{2}\Lambda x_n^2\right\}}_{\propto \int d\mathbf{x}_{\setminus n} q(\mathbf{x})/g_n(\mathbf{x})}, \tag{9}$$

where we dropped the dependency of γ and Λ on n for notational simplicity. Through some manipulation, the tilted distribution is equivalent to

$$q_n(x_n) = \frac{1 + m_n}{2} \delta(x_n - 1) + \frac{1 - m_n}{2} \delta(x_n + 1), \quad m_n = \tanh(\gamma) = \frac{e^\gamma - e^{-\gamma}}{e^\gamma + e^{-\gamma}}. \tag{10}$$

This discrete distribution has mean m_n and variance $1 - m_n^2$. By adapting the parameters of $g_n(x_n)$ using for example the EP algorithm, we aim to match the mean and variance of the marginal $q(x_n)$ (of $q(\mathbf{x})$) to the mean and variance of $q_n(x_n)$. The reader is referred to Section 9 for benchmarked results for the Ising model.

3.1.2 ANALYTIC BIVARIATE CASE

Here we shall compare the exact result with EP and the correction for the simplest non-trivial model, the $N = 2$ Ising model with no external field

$$p(\mathbf{x}) = \frac{1}{4} \left(\delta(x_1 - 1) + \delta(x_1 + 1) \right) \left(\delta(x_2 - 1) + \delta(x_2 + 1) \right) e^{Jx_1x_2}.$$

In order to solve the moment matching conditions we observe that the mean values must be zero because the distribution is symmetric around zero. Likewise the linear term in the approximating factors disappears and we can write $g_n(x_n) = \exp\{-\lambda x_n^2/2\}$ and $q(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{0}, \boldsymbol{\Sigma})$ with $\boldsymbol{\Sigma} = \begin{bmatrix} \lambda & -J \\ -J & \lambda \end{bmatrix}^{-1}$. The moment matching condition for the variances, $1 = \Sigma_{nn}$, turns into a second order equation with solution $\lambda = \frac{1}{2} [J^2 + \sqrt{J^4 + 4}]$. We can now insert this solution into the expression for the EP partition function in Equation (8). By expanding the result to the second order in J^2 , we find that

$$\log Z_{\text{EP}} = -\frac{1}{2} + \frac{1}{2} \sqrt{1 + 4J^2} - \frac{1}{2} \log \left(\frac{1}{2} (1 + \sqrt{1 + 4J^2}) \right) = \frac{J^2}{2} - \frac{J^4}{4} + \dots$$

Comparing with the exact expression

$$\log Z = \log \cosh(J) = \frac{J^2}{2} - \frac{J^4}{12} + \dots$$

we see that EP gives the correct J^2 coefficient, but the J^4 coefficient comes out wrong. In Section 4 we investigate how cumulant corrections can correct for this discrepancy.

3.2 Two Explanations Why Gaussian EP is Often Very Accurate

EP, as introduced above, is an algorithm. The justification for the algorithm put forward by Minka and adopted by others (see for example recent textbooks by Bishop 2006, Barber 2012 and Murphy 2012) is useful for explaining the steps in the algorithm but may be misleading in order to explain why EP often provides excellent accuracy in estimation of marginal moments and Z .

The general justification for EP (Minka, 2001a,b) is based upon a minimization of Kullback-Leiber (KL) divergences. Ideally, one would determine the approximating distribution $q(\mathbf{x})$ as the minimizer of $\text{KL}(p||q)$ in an exponential family of (in our case, Gaussian) densities. Since this is not possible—it would require the computation of exact moments—we instead iteratively minimize “local” KL-divergences $\text{KL}(q_a||q)$, between the tilted distribution q_a and q , with respect to g_a (appearing in q). This leads to the moment matching conditions in Equation (7). The argument for this procedure is essentially that this will ensure that the approximation q will capture high density regions of the intractable posterior p . Obviously, this argument cannot be applied to Ising models because the exact and approximate distributions are very different, with the former being discrete due to the Dirac δ -functions that constrain $x_n = \pm 1$ to be binary variables. Even though the optimization still implies moment matching, this discrete-continuous discrepancy makes local KL-divergences $\text{KL}(q_a||q)$ *infinite!*

In order to justify the usefulness of EP for Ising models we therefore need an alternative argument. Our argument is entirely restricted to *Gaussian* EP for our extended definition of GLVMs and do not extend to approximations with other exponential families. In the following, we will discuss these assumptions in inference approximations that preceded the formulation of EP, in order to provide a possibly more relevant justification of the method. Although this justification is not strictly necessary for practically using EP nor corrections to EP, it nevertheless provides a good starting point for understanding both.

The argument goes back to the mathematical analysis of the Sherrington-Kirkpatrick (SK) model for a disordered magnet (a so-called spin glass) (Sherrington and Kirkpatrick, 1975). For this Ising model, the couplings \mathbf{J} are drawn at random from a Gaussian distribution. An important contribution in the context of inference for this model (the computations of partition functions and average magnetizations) was the work of Thouless et al. (1977) who derived *self-consistency equations* which are assumed to be valid with a probability (with respect to the drawing of random couplings) approaching one as the number of variables x_n grows to infinity. These so-called Thouless-Anderson-Palmer (TAP) equations are closely related to the EP moment matching conditions of Equation (7), but they differ by partly relying on the specific assumption of the randomness of the couplings. Self-consistency equations equivalent to the EP moment matching conditions which avoided such assumptions on the statistics of the random couplings were first derived by Oppor and Winther (2000) by using a so-called cavity argument (Mézard et al., 1987). A new important contribution of Minka (2001a) was to provide an efficient algorithmic recipe for solving these equations.

We will now sketch the main idea of the cavity argument for the GLVM. Let $\mathbf{x}_{\setminus n}$ (“ \mathbf{x} without n ”) denote the complement to x_n , that is $\mathbf{x} = \mathbf{x}_{\setminus n} \cup x_n$. Without loss of generality we will take the quadratic exponential term to be written as $f_0(\mathbf{x}) \propto \exp(-\mathbf{x}^T \mathbf{J} \mathbf{x} / 2)$. With similar definitions of $\mathbf{J}_{\setminus n}$,

the exact marginal distribution of x_n may be written as

$$\begin{aligned} p_n(x_n) &= \frac{1}{Z} t_n(x_n) \int \exp \left\{ -\frac{1}{2} \mathbf{x}^T \mathbf{J} \mathbf{x} \right\} \prod_{n' \neq n} t_{n'}(x_{n'}) \, d\mathbf{x}_{\setminus n} \\ &= \frac{t_n(x_n)}{Z} e^{-J_{nn} x_n^2 / 2} \int \exp \left\{ -x_n \sum_{n' \neq n} J_{nn'} x_{n'} - \frac{1}{2} \mathbf{x}_{\setminus n}^T \mathbf{J}_{\setminus n} \mathbf{x}_{\setminus n} \right\} \prod_{n' \neq n} t_{n'}(x_{n'}) \, d\mathbf{x}_{\setminus n} . \end{aligned}$$

It is clear that $p_n(x_n)$ depends entirely on the statistics of the random variable $h_n \equiv \sum_{n' \neq n} J_{nn'} x_{n'}$. This is the total ‘field’ created by all other ‘magnetic moments’ $x_{n'}$ in the ‘cavity’ opened once x_n has been removed from the system. In the context of densely connected models with weak couplings, we can appeal to the central limit theorem² to approximate h_n by a Gaussian random variable with mean γ_n and variance V_n . When looking at the influence of the remaining variables $\mathbf{x}_{\setminus n}$ on x_n , the non-Gaussian details of their distribution have been washed out in the marginalization. Integrating out the Gaussian random variable h_n gives the Gaussian cavity field approximation to the marginal distribution:

$$\begin{aligned} p_n(x_n) &\approx \text{const} \cdot t_n(x_n) e^{-J_{nn} x_n^2 / 2} \int e^{-x_n h} \mathcal{N}(h; \gamma_n, V_n) \, dh \\ &= \text{const} \cdot t_n(x_n) \exp \left\{ -x_n \gamma_n - \frac{1}{2} (J_{nn} - V_n) x_n^2 \right\} . \end{aligned}$$

This is precisely of the form of the marginal tilted distribution $q_n(x_n)$ of Equation (9) as given by Gaussian EP. In the cavity formulation, $q(\mathbf{x})$ is simply a *placeholder* for the sufficient statistics of the individual Gaussian cavity fields. So we may observe cases, with the Ising model or bounded support factors being the prime examples, where EP gives essentially correct results for the marginal distributions of the x_n and of the partition function Z , while $q(\mathbf{x})$ gives a poor or even meaningless (in the sense of KL divergences) approximation to the multivariate posterior. Note however, that the entire *covariance matrix* of the x_n can be computed simply from a derivative of the free energy (Opper and Winther, 2005) resulting in an approximation of this covariance by that of $q(\mathbf{x})$. This may indicate that a good EP approximation of the free energy may also result in a good approximation to the full covariance. The near exactness of EP (as compared to exhaustive summation) in Section 9 therefore shows the central limit theorem at work. Conversely, mediocre accuracy or even failure of Gaussian EP, as also observed in our simulations in Sections 8.3 and 9, may be attributed to breakdown of the Gaussian cavity field assumption. Exact inference on the strongest couplings as considered for the Ising model in Section 9 is one way to alleviate the shortcoming of the Gaussian cavity field assumption.

4. Corrections to EP

The Z_{EP} approximation can be corrected in a principled approach, which traces the following outline:

1. The exact partition function Z is re-written in terms of Z_{EP} , scaled by a correction factor $R = Z/Z_{\text{EP}}$. This correction factor R encapsulates the intractability in the model, and contains a ‘local marginal’ contribution by each f_a (see Section 4.1).

2. In the context of sparsely connected models, other cavity arguments lead to loopy belief propagation.

2. A “handle” on R is obtained by writing it in terms of the cumulants (to be defined in Section 4.2) of $q(\mathbf{x})$ and $q_a(\mathbf{x})$ from Equations (5) and (6). As $q_a(\mathbf{x})$ and $q(\mathbf{x})$ share their two first cumulants, the mean and covariance from the moment matching condition in Equation (7), a cumulant expansion of R will be in terms of *higher-order* cumulants (see Section 4.2).
3. R , defined in terms of cumulant differences, is written as a complex Gaussian average. Each factor f_a contributes a complex random variable \mathbf{k}_a in this average (see Section 4.3).
4. Finally, the cumulant differences are used as “small quantities” in a Taylor series expansion of R , and the leading terms are kept (see Sections 5 and 7).

The series expansion is in terms of a complex expectation with a *zero* “self-relation” matrix, and this has two important consequences. Firstly, it causes all first order terms in the Taylor expansion to disappear, showing that Z_{EP} is correct to first order. Secondly, due to Wick’s theorem (introduced in Section 4.4), these zeros will contract the expansion by making many other terms vanish.

The strategy that is presented here can be re-used to correct other quantities of interest, like marginal distributions or the predictive density of new data when $p(\mathbf{x})$ is a Bayesian probabilistic model. These corrections are outlined in Appendix D.

4.1 Exact Expression for Correction

We define the (intractable) correction R as $Z = RZ_{EP}$. We can derive a useful expression for R in a few steps as follows: First we solve for f_a in Equation (6), and substitute this into Equation (4) to obtain

$$\prod_a f_a(\mathbf{x})^{D_a} = \prod_a \left(\frac{Z_a q_a(\mathbf{x}) g_a(\mathbf{x})}{q(\mathbf{x})} \right)^{D_a} = Z_{EP} q(\mathbf{x}) \prod_a \left(\frac{q_a(\mathbf{x})}{q(\mathbf{x})} \right)^{D_a}. \quad (11)$$

We introduce $F(\mathbf{x})$

$$F(\mathbf{x}) \equiv \prod_a \left(\frac{q_a(\mathbf{x})}{q(\mathbf{x})} \right)^{D_a}$$

to derive the expression for the correction $R = Z/Z_{EP}$ by integrating Equation (11):

$$R = \int q(\mathbf{x}) F(\mathbf{x}) \, d\mathbf{x}, \quad (12)$$

where we have used $Z = \int \prod_a f_a(\mathbf{x})^{D_a} \, d\mathbf{x}$. Similarly we can write:

$$p(\mathbf{x}) = \frac{1}{Z} \prod_a f_a(\mathbf{x})^{D_a} = \frac{Z_{EP}}{Z} q(\mathbf{x}) F(\mathbf{x}) = \frac{1}{R} q(\mathbf{x}) F(\mathbf{x}). \quad (13)$$

Corrections to the marginal and predictive densities of $p(\mathbf{x})$ can be computed from this formulation. This expression will become especially useful because the terms in $F(\mathbf{x})$ turn out to be “local”, that is, they only depend on the marginals of the variables associated with factor a . Let $f_a(\mathbf{x})$ depend on the subset \mathbf{x}_a of \mathbf{x} , and let $\mathbf{x}_{\setminus a}$ (“ \mathbf{x} without a ”) denote the remaining variables. The distributions in Equations (5) and (6) differ only with respect to their marginals on \mathbf{x}_a , $q_a(\mathbf{x}_a)$ and $q(\mathbf{x}_a)$, and therefore

$$\frac{q_a(\mathbf{x})}{q(\mathbf{x})} = \frac{q(\mathbf{x}_{\setminus a} | \mathbf{x}_a) q_a(\mathbf{x}_a)}{q(\mathbf{x}_{\setminus a} | \mathbf{x}_a) q(\mathbf{x}_a)} = \frac{q_a(\mathbf{x}_a)}{q(\mathbf{x}_a)}.$$

Now we can rewrite $F(\mathbf{x})$ in terms of marginals:

$$F(\mathbf{x}) = \prod_a \left(\frac{q_a(\mathbf{x}_a)}{q(\mathbf{x}_a)} \right)^{D_a}. \quad (14)$$

The key quantity, then, is F , after which the key operation is to compute its expected value. The rest of this section is devoted to the task of obtaining a “handle” on F .

4.2 Characteristic Functions and Cumulants

The distributions present in each of the ratios in $F(\mathbf{x})$ in Equation (14) share their first two cumulants, mean and covariance. Cumulants and cumulant differences are formally defined in the next paragraph. This simple observation has a crucial consequence: As the $q(\mathbf{x}_a)$'s are Gaussian and do not contain any higher order cumulants (three and above), F can be expressed in terms of the higher cumulants of the *marginals* $q_a(\mathbf{x}_a)$. When the term-product approximation is fully factorized, these are simply cumulants of *one-dimensional* distributions.

Let N_a be the number of variables in subvector \mathbf{x}_a . In the examples presented in this work, N_a is one or two. Furthermore, let \mathbf{k}_a be an N_a -dimensional vector $\mathbf{k}_a = (k_1, \dots, k_{N_a})_a$. The characteristic function of q_a is

$$\chi_a(\mathbf{k}_a) = \int e^{i\mathbf{k}_a^T \mathbf{x}_a} q_a(\mathbf{x}_a) d\mathbf{x}_a = \langle e^{i\mathbf{k}_a^T \mathbf{x}_a} \rangle_{q_a}, \quad (15)$$

and is obtained through the Fourier transform of the density. Inversely,

$$q_a(\mathbf{x}_a) = \frac{1}{(2\pi)^{N_a}} \int e^{-i\mathbf{k}_a^T \mathbf{x}_a} \chi_a(\mathbf{k}_a) d\mathbf{k}_a. \quad (16)$$

The cumulants $c_{\alpha a}$ of q_a are the coefficients that appear in the Taylor expansion of $\log \chi_a(\mathbf{k}_a)$ around the zero vector,

$$c_{\alpha a} = \left[(-i)^l \left(\frac{\partial}{\partial \mathbf{k}_a} \right)^\alpha \log \chi_a(\mathbf{k}_a) \right]_{\mathbf{k}_a=\mathbf{0}}.$$

By this definition of $c_{\alpha a}$, the Taylor expansion of $\log \chi_a(\mathbf{k}_a)$ is

$$\log \chi_a(\mathbf{k}_a) = \sum_{l=1}^{\infty} i^l \sum_{|\alpha|=l} \frac{c_{\alpha a}}{\alpha!} \mathbf{k}_a^\alpha.$$

Some notation was introduced in the above two equations to facilitate manipulating a multivariate series. The vector $\alpha = (\alpha_1, \dots, \alpha_{N_a})$, with $\alpha_j \in \mathbb{N}_0$, denotes a multi-index on the elements of \mathbf{k}_a . Other notational conventions that employ α (writing k_j instead of $k_{a,j}$) are:

$$|\alpha| = \sum_j \alpha_j, \quad \mathbf{k}_a^\alpha = \prod_j k_j^{\alpha_j}, \quad \alpha! = \prod_j \alpha_j!, \quad \left(\frac{\partial}{\partial \mathbf{k}_a} \right)^\alpha = \prod_j \frac{\partial^{\alpha_j}}{\partial k_j^{\alpha_j}}.$$

For example, when $N_a = 2$, say for the edge-factors in a spanning tree, the set of multi-indices α where $|\alpha| = 3$ are $(3, 0)$, $(2, 1)$, $(1, 2)$, and $(0, 3)$.

There are two characteristic functions that come into play in $F(\mathbf{x})$ and R in Equation (13). The first is that of the tilted distribution, $\log \chi_a(\mathbf{k}_a)$, and the other is the characteristic function of the

EP marginal $q(\mathbf{x}_a)$, defined as $\chi(\mathbf{k}_a) = \langle e^{i\mathbf{k}_a^T \mathbf{x}_a} \rangle_q$. By virtue of matching the first two moments, and $q(\mathbf{x}_a)$ being Gaussian with cumulants $c'_{\alpha a}$,

$$\begin{aligned} r_a(\mathbf{k}_a) &= \log \chi_a(\mathbf{k}_a) - \log \chi(\mathbf{k}_a) = \sum_{l \geq 1} i^l \sum_{|\alpha|=l} \frac{c_{\alpha a} - c'_{\alpha a}}{\alpha!} \mathbf{k}_a^\alpha \\ &= \sum_{l \geq 3} i^l \sum_{|\alpha|=l} \frac{c_{\alpha a}}{\alpha!} \mathbf{k}_a^\alpha \end{aligned} \quad (17)$$

contains the remaining higher-order cumulants where the tilted and approximate distributions *differ*. All our subsequent derivations rest upon moment matching being attained. This especially means that one cannot use the derived corrections if EP has not converged.

4.2.1 ISING MODEL EXAMPLE

The cumulant expansion for the discrete distribution in Equation (10) becomes

$$\begin{aligned} \log \chi_n(k_n) &= \log \int dx_n e^{ik_n x_n} q_n(x_n) = \log \left(\frac{1+m}{2} e^{ik_n} + \frac{1-m}{2} e^{-ik_n} \right) \\ &= imk_n - \frac{1}{2!}(1-m^2)k_n^2 - \frac{i}{3!}(-2m+2m^3)k_n^3 + \frac{1}{4!}(-2+8m^2-6m^4)k_n^4 + \dots \end{aligned}$$

(we're compactly writing m for m_n), from which the cumulants are obtained as

$$\begin{aligned} c_{1n} &= m, & c_{4n} &= -2 + 8m^2 - 6m^4, \\ c_{2n} &= 1 - m^2, & c_{5n} &= 16m - 40m^3 + 24m^5, \\ c_{3n} &= -2m + 2m^3, & c_{6n} &= 16 - 136m^2 + 240m^4 - 120m^6. \end{aligned}$$

4.3 The Correction as a Complex Expectation

The expected value of F , which is required for the correction, has a dependence on a product of ratios of distributions $q_a(\mathbf{x}_a)/q(\mathbf{x}_a)$. In the preceding section it was shown that the contributing distributions share lower-order statistics, allowing a twofold simplification. Firstly, the ratio q_a/q will be written as a single quantity that depends on r_a , which was introduced above in Equation (17). Secondly, we will show that it is natural to shift integration variables into the complex plane, and rely on complex Gaussian random variables (meaning that both real and imaginary parts are jointly Gaussian). These complex random variables that define the r_a 's have a peculiar property: they have a zero self-relation matrix! This property has important consequences in the resulting expansion.

4.3.1 COMPLEX EXPECTATIONS

Assume that $q(\mathbf{x}_a) = \mathcal{N}(\mathbf{x}_a; \boldsymbol{\mu}_a, \boldsymbol{\Sigma}_a)$ and $q_a(\mathbf{x}_a)$ share the same mean and covariance, and substitute $\log \chi_a(\mathbf{k}_a) = r_a(\mathbf{k}_a) + \log \chi(\mathbf{k}_a)$ in the definition of q_a in Equation (16) to give

$$\frac{q_a(\mathbf{x}_a)}{q(\mathbf{x}_a)} = \frac{\int e^{-i\mathbf{k}_a^T \mathbf{x}_a + r_a(\mathbf{k}_a)} \chi(\mathbf{k}_a) d\mathbf{k}_a}{\int e^{-i\mathbf{k}_a^T \mathbf{x}_a} \chi(\mathbf{k}_a) d\mathbf{k}_a}. \quad (18)$$

Although the \mathbf{k}_a variables have not been introduced as random variables, we find it natural to *interpret them as such*, because the rules of expectations over Gaussian random variables will be

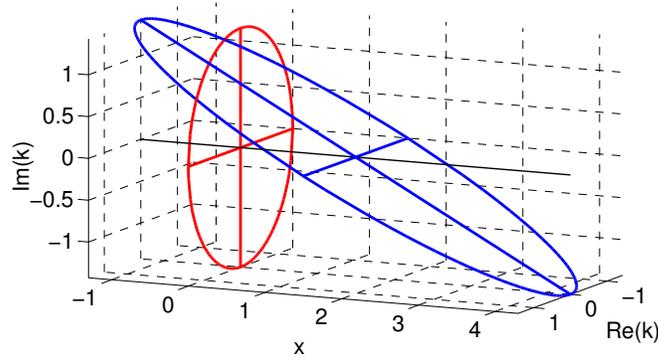


Figure 1: Equation (20) shifts \mathbf{k}_a to the complex plane. In the simplest case the joint density $p(k|x)q(x)$ is $x \sim \mathcal{N}(\mu, \sigma^2)$, $\Re(k) \sim \mathcal{N}(0, \sigma^{-2})$ and equality $\Im(k) = -\sigma^{-2}(x - \mu)$. Notice that $\Re(k)$'s variance is the *inverse* of that of x . The joint density is a two-dimensional flat ellipsoidal pancake that lives in three dimensions: x and the complex k plane (tilted ellipsoid). Integrating over x gives the marginal over a complex k , which is *still* a two-dimensional random variable (upright ellipsoid). The marginal has $\Im(k) \sim \mathcal{N}(0, \sigma^{-2})$, and hence k has relation $\langle (\Re(k) + i\Im(k))^2 \rangle = \sigma^{-2} - \sigma^{-2} = 0$ and variance $\langle k\bar{k} \rangle = 2\sigma^{-2}$.

extremely helpful in developing the subsequent expansions. We will therefore write $q_a(\mathbf{x}_a)/q(\mathbf{x}_a)$ as an expectation of $\exp r_a(\mathbf{k}_a)$ over a density $p(\mathbf{k}_a|\mathbf{x}_a) \propto e^{-i\mathbf{k}_a^T \mathbf{x}_a} \chi(\mathbf{k}_a)$:

$$\frac{q_a(\mathbf{x}_a)}{q(\mathbf{x}_a)} = \left\langle \exp r_a(\mathbf{k}_a) \right\rangle_{\mathbf{k}_a|\mathbf{x}_a}. \quad (19)$$

By substituting $\log \chi(\mathbf{k}_a) = i\boldsymbol{\mu}_a^T \mathbf{k}_a - \mathbf{k}_a^T \boldsymbol{\Sigma}_a \mathbf{k}_a / 2$ into Equation (18), we see that $p(\mathbf{k}_a|\mathbf{x}_a)$ can be viewed as Gaussian, but not for real random variables! We have to consider \mathbf{k}_a as Gaussian random variables with a real and an imaginary part with

$$\Re(\mathbf{k}_a) \sim \mathcal{N}\left(\Re(\mathbf{k}_a); \mathbf{0}, \boldsymbol{\Sigma}_a^{-1}\right), \quad \Im(\mathbf{k}_a) = -\boldsymbol{\Sigma}_a^{-1}(\mathbf{x}_a - \boldsymbol{\mu}_a).$$

For the purpose of computing the expectation in Equation (19), $\mathbf{k}_a|\mathbf{x}_a$ is a degenerate complex Gaussian that shifts the coefficients \mathbf{k}_a into the complex plane. The expectation of $\exp r_a(\mathbf{k}_a)$ is therefore taken over Gaussian random variables that have $q(\mathbf{x}_a)$'s *inverse* covariance matrix as their (real) covariance! As shorthand, we write

$$p(\mathbf{k}_a|\mathbf{x}_a) = \mathcal{N}\left(\mathbf{k}_a; -i\boldsymbol{\Sigma}_a^{-1}(\mathbf{x}_a - \boldsymbol{\mu}_a), \boldsymbol{\Sigma}_a^{-1}\right). \quad (20)$$

Figure 1 illustrates a simple density $p(\mathbf{k}_a|\mathbf{x}_a)$, showing that the imaginary component is a deterministic function of \mathbf{x}_a . Once \mathbf{x}_a is averaged out of the joint density $p(\mathbf{k}_a|\mathbf{x}_a)q(\mathbf{x}_a)$, a *circularly symmetric complex Gaussian* distribution over \mathbf{k}_a remains. It is circularly symmetric as $\langle \mathbf{k}_a \rangle = \mathbf{0}$, relation matrix $\langle \mathbf{k}_a \mathbf{k}_a^T \rangle = \mathbf{0}$, and covariance matrix $\langle \mathbf{k}_a \bar{\mathbf{k}}_a^T \rangle = 2\Sigma_a^{-1}$ (notation \bar{k} indicates the complex conjugate of k). For the purpose of computing the expected values with Wick's theorem (following in Section 4.4 below), we *only* need the relations $\langle \mathbf{k}_a \mathbf{k}_b^T \rangle$ for pairs of factors a and b . All of these will be derived next:

According to Equation (12), a further expectation over $q(\mathbf{x})$ is needed, after integrating over \mathbf{k}_a , to determine R . These variables will be *combined* into complex random variables to make the averages in the expectation easier to derive. By substituting Equation (19) into Equation (12), R is equal to

$$R = \langle F(\mathbf{x}) \rangle_{\mathbf{x} \sim q(\mathbf{x})} = \left\langle \prod_a \left\langle \exp r_a(\mathbf{k}_a) \right\rangle_{\mathbf{k}_a|\mathbf{x}_a}^{D_a} \right\rangle_{\mathbf{x}}. \tag{21}$$

When \mathbf{x} is given, the \mathbf{k}_a -variables are independent. However, when they are averaged over $q(\mathbf{x})$, the \mathbf{k}_a -variables become coupled. They are zero-mean complex Gaussians

$$\langle \mathbf{k}_a \rangle = \left\langle \langle \mathbf{k}_a \rangle_{\mathbf{k}_a|\mathbf{x}_a} \right\rangle_{\mathbf{x}} = \left\langle -i\Sigma_a^{-1}(\mathbf{x}_a - \boldsymbol{\mu}_a) \right\rangle_{\mathbf{x}} = \mathbf{0}$$

and are coupled with a zero self-relation matrix! In other words, if $\Sigma_{ab} = \text{cov}(\mathbf{x}_a, \mathbf{x}_b)$, the expected values $\langle \mathbf{k}_a \mathbf{k}_b^T \rangle$ between the variables in the set $\{\mathbf{k}_a\}$ are

$$\begin{aligned} \langle \mathbf{k}_a \mathbf{k}_b^T \rangle &= \left\langle \langle \mathbf{k}_a \mathbf{k}_b^T \rangle_{\mathbf{k}_{a,b}|\mathbf{x}} \right\rangle_{\mathbf{x}} + i^2 \Sigma_a^{-1} \left\langle (\mathbf{x}_a - \boldsymbol{\mu}_a)(\mathbf{x}_b - \boldsymbol{\mu}_b)^T \right\rangle_{\mathbf{x}} \Sigma_b^{-1} \\ &= \begin{cases} \mathbf{0} & \text{if } a = b \\ -\Sigma_a^{-1} \Sigma_{ab} \Sigma_b^{-1} & \text{if } a \neq b \end{cases}. \end{aligned} \tag{22}$$

Complex Gaussian random variables are additionally characterized by $\langle \mathbf{k}_a \bar{\mathbf{k}}_b^T \rangle$. However, these expectations are not required for computing and simplifying the expansion of $\log R$ in Section 5, and are not needed for the remainder of this paper. Figure 2 illustrates the structure of the resulting relation matrix $\langle \mathbf{k}_a \mathbf{k}_b^T \rangle$ for two different factorizations of the same distribution. Each factor f_a contributes a \mathbf{k}_a variable, such that the tree-structured approximation's relation matrix will be larger than that of the fully factorized one.

Section 5 shows that when $D_a = 1$, the above expectation can be written directly over $\{\mathbf{k}_a\}$ and expanded. In the general case, discussed in Section 7, the inner expectation is first expanded (to treat the D_a powers) before computing an expectation over $\{\mathbf{k}_a\}$. In both cases the expectation will involve polynomials in k -variables. The expected values of Gaussian polynomials can be evaluated with Wick's theorem.

4.4 Wick's Theorem

Wick's theorem provides a useful formula for mixed central moments of Gaussian variables. Let $k_{n_1}, \dots, k_{n_\ell}$ be real or complex centered jointly Gaussian variables, noting that they do not have to be different. Then

$$\langle k_{n_1} \cdots k_{n_\ell} \rangle = \sum_{\eta} \prod_{\eta} \langle k_{i_\eta} k_{j_\eta} \rangle, \tag{23}$$

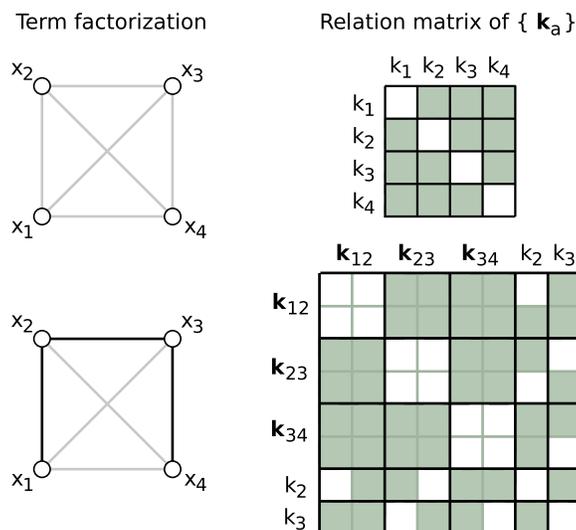


Figure 2: The relation matrices between \mathbf{k}_a for two factorizations of $\prod_{n=1}^4 t_n(x_n)$: the top illustration is for $t_1 t_2 t_3 t_4$, while the bottom illustration is of a tree structure $(t_1 t_2)(t_2 t_3)(t_3 t_4)/t_2/t_3$. The white squares indicate a zero relation matrix $\langle \mathbf{k}_a \mathbf{k}_b^T \rangle$, with the *diagonal* being zero. From the properties of Equation (22) there are additional zeros in the tree structure’s relation matrix, where edge and node factors share variables. The factor $f_0 = g_0$ is shadowed in grey in the left-hand figures, and can make $q(\mathbf{x})$ densely connected.

where the sum is over all partitions of $\{n_1, \dots, n_\ell\}$ into disjoint pairs $\{i_\eta, j_\eta\}$. If $\ell = 2m$ is even, then there are $(2m)!/(2^m m!) = (2m - 1)!!$ such partitions.³ If ℓ is odd, then there are none, and the expectation in Equation (23) is zero.

Consider the one-dimensional variable $k \sim \mathcal{N}(k; 0, \sigma^2)$. Wick’s theorem states that $\langle k^\ell \rangle = (\ell - 1)!! \sigma^\ell$ if ℓ is even, and $\langle k^\ell \rangle = 0$ if ℓ is odd. In other words, $\langle k^3 \rangle = 0$, $\langle k^4 \rangle = 3(\sigma^2)^2$, $\langle k^6 \rangle = 15(\sigma^2)^3$, and so forth.

5. Factorized Approximations

In the fully factorized approximation, with $f_n(x_n) = t_n(x_n)$, the exact distribution in Equation (13) depends on the *single node marginals* $F(\mathbf{x}) = \prod_n q_n(x_n)/q(x_n)$. Following Equation (21), the correction to the free energy

$$R = \left\langle \prod_n \left\langle \exp r_n(k_n) \right\rangle_{k_n|x_n} \right\rangle_{\mathbf{x}} = \left\langle \exp \left[\sum_n r_n(k_n) \right] \right\rangle_{\mathbf{k}} \tag{24}$$

is taken directly over the centered complex-valued Gaussian random variables $\mathbf{k} = (k_1, \dots, k_N)$, which have a relations

$$\langle k_m k_n \rangle = \begin{cases} 0 & \text{if } m = n \\ -\Sigma_{mn}/(\Sigma_{mm}\Sigma_{nn}) & \text{if } m \neq n \end{cases} \tag{25}$$

3. The double factorial is $(2m - 1)!! = (2m - 1) \times (2m - 3) \times (2m - 5) \times \dots \times 1$.

In the section to follow, all expectations shall be with respect to \mathbf{k} , which will be dropped where it is clear from the context.

Thus far, R is re-expressed in terms of site contributions. The expression in Equation (24) is exact, albeit still intractable, and will be treated through a power series expansion. Other quantities of interest, like marginal distributions or moments, can similarly be expressed exactly, and then expanded (see Appendix D).

5.1 Second Order Correction to $\log R$

Assuming that the r_n 's are small on average with respect to \mathbf{k} , Equation (24) is expanded and the lower order terms kept:

$$\begin{aligned} \log R &= \log \left\langle \exp \left[\sum_n r_n(k_n) \right] \right\rangle = \sum_n \langle r_n \rangle + \frac{1}{2} \left\langle \left(\sum_n r_n \right)^2 \right\rangle - \frac{1}{2} \left(\sum_n \langle r_n \rangle \right)^2 + \dots \\ &= \frac{1}{2} \sum_{m \neq n} \langle r_m r_n \rangle + \dots \end{aligned} \quad (26)$$

The simplification in the second line is a result of the variance terms being zero from Equation (25). The single marginal terms also vanish (and hence EP is correct to first order) because both $\langle k_n \rangle = 0$ and $\langle k_n^2 \rangle = 0$.

This result can give us a hint in which situations the corrections are expected to be small:

- Firstly, the r_n could be small for values of k_n where the density of \mathbf{k} is not small. For example, under a zero noise Gaussian process classification model, $q_n(x_n)$ equals a step function $t_n(x_n)$ times a Gaussian, where the latter often has small variance compared to the mean. Hence, $q_n(x_n)$ should be very close to a Gaussian.
- Secondly, for systems with weakly (posterior) dependent variables x_n we might expect that the log partition function $\log Z$ would scale approximately linearly with N , the number of variables. Since terms with $m = n$ vanish in the computation of $\ln R$, there are no corrections that are *proportional to* N when Σ_{mn} is sufficiently small as $N \rightarrow \infty$. Hence, the dominant contributions to $\log Z$ should already be included in the EP approximation. However, Section 8.3 illustrates an example where this need not be the case.

The expectation $\langle r_m r_n \rangle$, as it appears in Equation (26), is treated by substituting r_n with its cumulant expansion $r_n(k_n) = \sum_{l \geq 3} i^l c_{ln} k_n^l / l!$ from Equation (17). Wick's theorem now plays a pivotal role in evaluating the expectations that appear in the expansion:

$$\begin{aligned} \langle r_m(k_m) r_n(k_n) \rangle &= \sum_{l, s \geq 3} i^{l+s} \frac{c_{lm} c_{sn}}{l! s!} \langle k_m^s k_n^l \rangle \\ &= \sum_{l \geq 3} i^{2l} l! \frac{c_{lm} c_{ln}}{(l!)^2} \langle k_m k_n \rangle^l \\ &= \sum_{l \geq 3} \frac{c_{lm} c_{ln}}{l!} \left(\frac{\Sigma_{mn}}{\Sigma_{mm} \Sigma_{nn}} \right)^l. \end{aligned} \quad (27)$$

The second line above follows from contractions in Wick's theorem. All the *self-pairing terms*, when for example one of the l k_n 's is paired with another k_n in Equation (23), are zero because

$\langle k_n^2 \rangle = 0$. To therefore get a non-zero result for $\langle k_m^s k_n^l \rangle$, using Equation (23), *each* factor k_n has to be paired with some factor k_m , and this is possible only when $l = s$. Wick's theorem sums over all pairings, and there are $l!$ ways of pairing a k_n with a k_m , giving the result in Equation (27). Finally, plugging Equation (27) into Equation (26) gives the second order correction

$$\log R = \frac{1}{2} \sum_{m \neq n} \sum_{l \geq 3} \frac{c_{lm} c_{ln}}{l!} \left(\frac{\Sigma_{mn}}{\Sigma_{mm} \Sigma_{nn}} \right)^l + \dots \quad (28)$$

5.1.1 ISING EXAMPLE CONTINUED

We can now compute the second order $\log R$ correction for the $N = 2$ Ising model example of Section 3.1. The covariance matrix has $\Sigma_{nn} = 1$ from moment matching and $\Sigma_{12} = J/(\lambda^2 - J^2)$ with $\lambda = \frac{1}{2} \left[J^2 + \sqrt{J^4 + 4} \right]$. The uneven terms in the cumulant expansion derived in Section 4.2.1 disappear because $m = 0$. The first nontrivial term is therefore $l = 4$ which gives a contribution of $\frac{1}{2} \times 2 \times \frac{c_4^2}{4!} \Sigma_{12}^4 = \frac{(-2)^2}{4!} \Sigma_{12}^4 = \frac{1}{6} \Sigma_{12}^4$. In Section 3.1, we saw that $\log Z - \log Z_{EP} = \frac{J^4}{6}$ plus terms of order J^6 and higher. To lowest order in J we have $\Sigma_{12} = J$ and thus $\log R = \frac{J^4}{6}$ which exactly cancels the lowest order error of EP.

5.2 Corrections to Other Quantities

The schema given here is applicable to any other quantity of interest, be it marginal or predictive distributions, or the marginal moments of $p(\mathbf{x})$. The cumulant corrections for the marginal moments are derived in Appendix D; for example, the correction to the marginal mean μ_i of an approximation $q(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is

$$\langle x_i \rangle_{p(\mathbf{x})} - \mu_i = \sum_{l \geq 3} \sum_{j \neq n} \frac{\Sigma_{ij} c_{l+1,j} c_{ln}}{\Sigma_{jj}} \frac{1}{l!} \left(\frac{\Sigma_{jn}}{\Sigma_{jj} \Sigma_{nn}} \right)^l + \dots, \quad (29)$$

while the correction to the marginal covariance is

$$\begin{aligned} \langle (x_i - \mu_i)(x_{i'} - \mu_{i'}) \rangle_{p(\mathbf{x})} - \Sigma_{ii'} &= \sum_{l \geq 3} \sum_{j \neq n} \frac{\Sigma_{ij} \Sigma_{i'j} c_{l+2,j} c_{ln}}{\Sigma_{jj}^2} \frac{1}{l!} \left(\frac{\Sigma_{jn}}{\Sigma_{jj} \Sigma_{nn}} \right)^l \\ &+ \sum_{l \geq 3} \sum_{j \neq n} \frac{\Sigma_{ij} \Sigma_{i'n} c_{lj} c_{ln}}{\Sigma_{jj} \Sigma_{nn}} \frac{1}{l!} \left(\frac{\Sigma_{jn}}{\Sigma_{jj} \Sigma_{nn}} \right)^{l-1} + \dots \end{aligned} \quad (30)$$

5.3 Edgeworth-Type Expansions

To simplify the expansion of Equation (24), we integrated (combined) degenerate complex Gaussians $k_n |x_n$ over $q(\mathbf{x})$ to obtain fully complex Gaussian random variables $\{k_n\}$. We've then relied on $\langle k_n^2 \rangle = 0$ to simplify the expansion of $\log R$.

The expectations $\langle k_n^2 \rangle = 0$ are closely related to the orthogonality of Hermite polynomials, and this can be employed in an alternative derivation. In particular, one can *first* make a Taylor expansion of $\exp r_n(k_n)$ around zero, giving complex-valued polynomials in $\{k_n\}$. When the inner average in Equation (24) is then taken over $k_n |x_n$, a real-valued series of Hermite polynomials in $\{x_n\}$ arises. These polynomials are orthogonal under $q(\mathbf{x})$. The series that describes the tilted distribution $q_n(x_n)$ is equal to the product of $q(x_n)$ and an expansion of polynomials for the higher-cumulant *deviation*

from a Gaussian density. This line of derivation gives an Edgeworth expansion for each factor's tilted distribution.

As a second step, Equation (24) couples the product of separate Edgeworth expansions (one for each factor) together by requiring an outer average over $q(\mathbf{x})$. The orthogonality of Hermite polynomials under $q(\mathbf{x})$ now come into play: it allows products of orthogonal polynomials under $q(\mathbf{x})$ to integrate to zero. This is similar to contractions in Wick's theorem, where $\langle k_n^2 \rangle = 0$ allows us to simplify Equation (27). Although it is not the focus of this work, an example of such a derivation appears in Appendix C.1.

6. Radius of Convergence

We may hope that in practice the low order terms in the cumulant expansions will account already for the dominant contributions. But will such an expansion actually converge when extended to arbitrary orders? While we will leave a more general answer to future research, we can at least give a partial result for the example of the Ising model. Let $\mathbf{D} = \text{diag}(\Sigma)$, the diagonal of the covariance matrix of the EP approximation $q(\mathbf{x})$. We prove here that a cumulant expansion for R will converge when the eigenvalues of $\mathbf{D}^{-1/2}\Sigma\mathbf{D}^{-1/2}$ —which has diagonal values of one—are bounded between zero and two.

In practice we've found that even if the largest of these eigenvalues grows with N , the second-order correction gives a remarkable improvement. This, with the results in Figure 6, lead us to believe that the power series expansion is often divergent. It may well be that our expansions are only of an asymptotic type (Boyd, 1999) for which the summation of only a certain number of terms might give an improvement whereas further terms would lead to worse results. It leads to a paradoxical situation, which seems common when interesting functions are computed: On the one hand we may have a series which does not converge, but in many ways is more practical; on the other hand one might obtain an expansion that converges, but only impractically. Quoting George F. Carrier's rule from Boyd (1999):

Divergent series converge faster than convergent series because they don't have to converge.

For this, we do not yet have a clear-cut answer.

6.1 A Formal Expression for the Cumulant Expansion to All Orders

To discuss the question when our expansion will converge when extended to arbitrary orders, we introduce a single extra parameter λ into R , which controls the strength of the contribution of cumulants. Expanded into a series in powers of λ , contributions of cumulants of total order l are multiplied by a factor λ^l , for example $\lambda^l c_{nl}$ or $\lambda^{k+l} c_{nk} c_{nl}$. Of course, at the end of the calculation, we set $\lambda = 1$. This approach is obviously achieved by replacing

$$r_n(k_n) \rightarrow r_n(\lambda k_n)$$

in Equation (24). Hence, we define

$$R(\lambda) = \left\langle \exp \left[\sum_n r_n(\lambda k_n) \right] \right\rangle_{\mathbf{k}} = \left\langle \exp \left[\sum_n r_n(k_n) \right] \right\rangle_{\mathbf{k}'}$$

where

$$\langle k'_m k'_n \rangle = \begin{cases} 0 & \text{if } m = n \\ -\lambda^2 \Sigma_{mn} / (\Sigma_{mm} \Sigma_{nn}) & \text{if } m \neq n \end{cases} .$$

By working backwards, and expressing everything by the original densities over x_n , the correction can be written as

$$R(\lambda) = \left\langle \prod_n \frac{q_n(x_n)}{q(x_n)} \right\rangle_{q_\lambda(\mathbf{x})}, \quad (31)$$

where the density $q_\lambda(\mathbf{x})$ is a multivariate Gaussian with mean $\boldsymbol{\mu}$ and covariance given by

$$\boldsymbol{\Sigma}_\lambda = \mathbf{D} + z(\boldsymbol{\Sigma} - \mathbf{D}),$$

where $\mathbf{D} = \text{diag}(\boldsymbol{\Sigma})$ and $z = \lambda^2$. Hence, we see that the expansion in powers of λ is actually equivalent to an expansion in products of nondiagonal elements of $\boldsymbol{\Sigma}$.

Noticing that as $R(\lambda)$ depends on λ through the density $q_\lambda(\mathbf{x}) \propto |\boldsymbol{\Sigma}_\lambda|^{-1/2} e^{-\frac{1}{2} \mathbf{x}^\top \boldsymbol{\Sigma}_\lambda^{-1} \mathbf{x}}$, we can see by expressing $\boldsymbol{\Sigma}_\lambda^{-1}$ in terms of eigenvalues and eigenvectors that for any *fixed* \mathbf{x} , $q_\lambda(\mathbf{x})$ is an analytic function of the *complex variable* z as long as $\boldsymbol{\Sigma}_\lambda$ is positive definite. Since

$$\boldsymbol{\Sigma}_\lambda = \mathbf{D}^{1/2} \left\{ \mathbf{I} + z \left(\mathbf{D}^{-1/2} \boldsymbol{\Sigma} \mathbf{D}^{-1/2} - \mathbf{I} \right) \right\} \mathbf{D}^{1/2}$$

this is equivalent to the condition that the matrix $\mathbf{I} + z(\mathbf{D}^{-1/2} \boldsymbol{\Sigma} \mathbf{D}^{-1/2} - \mathbf{I})$ is positive definite. Introducing γ_i , the eigenvalues of $\mathbf{D}^{-1/2} \boldsymbol{\Sigma} \mathbf{D}^{-1/2}$, positive definiteness fails when for the first time $1 + z(\gamma_i - 1) = 0$. Thus the series for $q_\lambda(\mathbf{x})$ is convergent for

$$|z| < \min_i \frac{1}{|1 - \gamma_i|} .$$

Setting $z = 1$, this is equivalent to the condition

$$1 < \min_i \frac{1}{|1 - \gamma_i|} .$$

This means that the eigenvalues have to fulfil $0 < \gamma_i < 2$. Unfortunately, we can not conclude from this condition that pointwise convergence of $q_\lambda(\mathbf{x})$ for each \mathbf{x} leads to convergence of $R(\lambda)$ (which is an integral of $q_\lambda(\mathbf{x})$ over all \mathbf{x} !). However, in cases where the integral eventually becomes a finite sum, such as the Ising model, pointwise convergence in \mathbf{x} leads to convergence of $R(\lambda)$.

6.1.1 ISING MODEL EXAMPLE

From Section 4.2.1 the tilted distribution for the running example Ising model is $q_n(x_n) = \frac{1}{2} [\delta(x_n + 1) + \delta(x_n - 1)]$, and hence $q(x_n) = \frac{1}{(2\pi)^{1/2}} e^{-x_n^2/2}$. As each $q(x_n)$ is a unit-variance Gaussian, $\mathbf{D} = \text{diag}(\boldsymbol{\Sigma}) = \mathbf{I}$. Hence $\mathbf{D}^{-1/2} \boldsymbol{\Sigma} \mathbf{D}^{-1/2} = \boldsymbol{\Sigma}$ and

$$R(\lambda) = \frac{1}{\sqrt{|(1 - \lambda^2)\mathbf{I} + \lambda^2 \boldsymbol{\Sigma}|}} \frac{e^{N/2}}{2^N} \sum_{\mathbf{x} \in \{-1, 1\}^N} \exp \left[-\frac{1}{2} \mathbf{x}^\top \left((1 - \lambda^2)\mathbf{I} + \lambda^2 \boldsymbol{\Sigma} \right)^{-1} \mathbf{x} \right]$$

follows from Equation (31). The arguments of the previous section show that the *radius of convergence* of $R(\lambda)$ is determined by the condition that the matrix $\mathbf{I} + \lambda^2(\boldsymbol{\Sigma} - \mathbf{I})$ is positive definite or the eigenvalues l_i of $\boldsymbol{\Sigma}$ fulfil $|l_i - 1| \leq 1/\lambda^2$.

In the $N = 2$ case, $\Sigma = \begin{pmatrix} 1 & c \\ c & 1 \end{pmatrix}$ with $c = c(J) \in]-1, 1[$ which has eigenvalues $1 - c$ and $1 + c$, meaning that cumulant expansion for $R(\lambda)$ is convergent for the $N = 2$ Ising model. For $N > 2$, it is easy to show that this is not necessarily true. Consider the ‘isotropic’ Ising model with $J_{ij} = J$ and zero external field, then $\Sigma_{ii} = 1$ and $\Sigma_{ij} = c$ for $i \neq j$ with $c = c(J) \in]-1/(N-1), 1[$. The eigenvalues are now $1 + (N-1)c$ and $1 - c$ (the latter with degeneracy $N-1$). For finite c , the largest eigenvalue will scale with N and thus be larger than the upper value of two that would be required for convergence. Scaling with N for the largest eigenvalue of $\mathbf{D}^{-1/2} \Sigma \mathbf{D}^{-1/2}$ is also observed in the Ising model simulations Section 9.

We conjecture that convergence of the cumulant series for $R(\lambda)$ also implies convergence of the series for $\log R(\lambda)$ but leave an investigation of this point to future research. We only illustrate this point for the $N = 2$ Ising model case, where we have the explicit formula

$$\log R(\lambda) = 1 - \frac{1}{2} \log(1 - \lambda^4 c^2) - \frac{1}{1 - \lambda^4 c^2} + \log \cosh \left(\frac{\lambda^2 c}{1 - \lambda^4 c^2} \right).$$

As can be easily seen, an expansion in λ converges for $c^2 \lambda^4 < 1$ which gives the same radius of convergence $|c| < 1$ as for the expansion of R .

7. General Approximations

The general approximations differ from the factorized approximation in that an expansion in terms of expectations under $\{\mathbf{k}_a\}$ doesn’t immediately arise. Consider R in Equation (21): Its inner expectations are over $\mathbf{k}_a | \mathbf{x}$, and outer expectations are over \mathbf{x} . First take the binomial expansion of the *inner* expectation, and keep it to second order in r_a :

$$\begin{aligned} \left\langle e^{r_a(\mathbf{k}_a)} \right\rangle_{\mathbf{k}_a | \mathbf{x}}^{D_a} &= \left(1 + \langle r_a \rangle + \frac{1}{2} \langle r_a^2 \rangle + \dots \right)^{D_a} \\ &= 1 + D_a \left[\langle r_a \rangle + \frac{1}{2} \langle r_a^2 \rangle + \dots \right] + \frac{D_a(D_a-1)}{2} \left[\langle r_a \rangle + \frac{1}{2} \langle r_a^2 \rangle + \dots \right]^2 + \dots \\ &= 1 + D_a \langle r_a \rangle + \frac{D_a}{2} \langle r_a^2 \rangle + \frac{D_a(D_a-1)}{2} \langle r_a \rangle^2 + \dots \end{aligned}$$

Notice that $r_a(\mathbf{k}_a)$ can be complex, but $\langle r_a(\mathbf{k}_a) \rangle_{\mathbf{k}_a | \mathbf{x}}$, as it appears in the above expansion, is real-valued. Using this result, again expand $\langle \prod_a \langle e^{r_a} \rangle_{\mathbf{k}_a | \mathbf{x}}^{D_a} \rangle_{\mathbf{x}}$. The correction to $\log R$, up to second order, is

$$\begin{aligned} \log R &= \frac{1}{2} \sum_{a \neq b} D_a D_b \left\langle \langle r_a(\mathbf{k}_a) \rangle_{\mathbf{k}_a | \mathbf{x}} \langle r_b(\mathbf{k}_b) \rangle_{\mathbf{k}_b | \mathbf{x}} \right\rangle_{\mathbf{x}} \\ &\quad + \frac{1}{2} \sum_a D_a(D_a-1) \left\langle \langle r_a(\mathbf{k}_a) \rangle_{\mathbf{k}_a | \mathbf{x}}^2 \right\rangle_{\mathbf{x}} + \dots \end{aligned} \quad (32)$$

In the above relation the first-order terms all disappeared as $\langle \langle r_a(\mathbf{k}_a) \rangle \rangle = 0$. Terms involving $\langle \langle r_a(\mathbf{k}_a)^2 \rangle \rangle = 0$ similarly disappear, as every polynomial in the expansion $r_a(\mathbf{k}_a)^2$ averages to zero. This is a general case of Equation (26), in which $D_n = 1$ for all factors. In Appendix B we show how to use the general result for the case where the factorization is a tree and our factors are edges (pairs) and nodes (single variables).

8. Gaussian Process Results

One of the most important applications of EP is to statistical models with Gaussian process (GP) priors, where \mathbf{x} is a latent variable with Gaussian prior distribution with a kernel matrix \mathbf{K} as covariance $\mathbb{E}[\mathbf{xx}^T] = \mathbf{K}$.

It is well known that for many models, like GP classification, inference with EP is on par with MCMC ground truth (Kuss and Rasmussen, 2005). Section 8.1 underlines this case, and shows corrections to the partition function on the USPS data set over a range of kernel hyperparameter settings.

A common inference task is to predict the output for previously unseen data. Under a GP regression model, a key quantity is the predictive mean function. The predictive mean is analytically tractable when the latent function is corrupted with Gaussian noise to produce observations y_n . This need not be the case; in Section 8.2 we examine the problem of quantized regression, where the noise model is non-Gaussian with sharp discontinuities. We show practically how the corrections transfer to other moments, like the predictive mean. Through it, we arrive at a hypothetical rule of thumb: if the data isn't "sensible" under the (probabilistic) model of interest, there is no guarantee for EP giving satisfactory inference.

Armed with the rule of thumb, Section 8.3 constructs an insightful counterexample where the EP estimate diverges or is far from ground truth with more data. Divergence in the partition function is manifested in the initial correction terms, giving a test for the approximation accuracy that doesn't rely on any Monte Carlo ground truth.

8.1 Gaussian Process Classification

The GP classification model arises when we observe N data points \mathbf{s}_n with class labels $y_n \in \{-1, 1\}$, and model y through a latent function x with a GP prior. The likelihood terms for y_n are assumed to be $t_n(x_n) = \Phi(y_n x_n)$, where $\Phi(\cdot)$ denotes the cumulative Normal density.

An extensive MCMC evaluation of EP for GP classification on various data sets was given by Kuss and Rasmussen (2005), showing that the log marginal likelihood of the data can be approximated remarkably well. As shown by Opper et al. (2009), an even more accurate estimation of the approximation error is given by considering the second order correction in Equation (28). For GPC we generally found that the $l = 3$ term dominates $l = 4$, and we do not include any higher cumulants here.

Figure 3 illustrates the correction to $\log R$, with $l = 3, 4$, on the binary subproblem of the USPS 3's vs. 5's digits data set, with $N = 767$. This is the same set-up of Kuss and Rasmussen (2005) and Opper et al. (2009), using the kernel $k(\mathbf{s}, \mathbf{s}') = \sigma^2 \exp(-\frac{1}{2} \|\mathbf{s} - \mathbf{s}'\|^2 / \ell^2)$, and we refer the reader to both papers for additional and complimentary figures and results. We evaluated Equation (28) on a similar grid of $\log \ell$ and $\log \sigma$ values. For the same grid values we obtained Monte Carlo estimates of $\log Z$, and hence $\log R$. The correction, compared to the magnitude of the $\log Z$ grids by Kuss and Rasmussen (2005), is remarkably small, and underlines their findings on the accuracy of EP for GPC.

The correction from Equation (28), as computed here, is $O(N^2)$, and compares favorably to $O(N^3)$ complexity of EP for GPC.

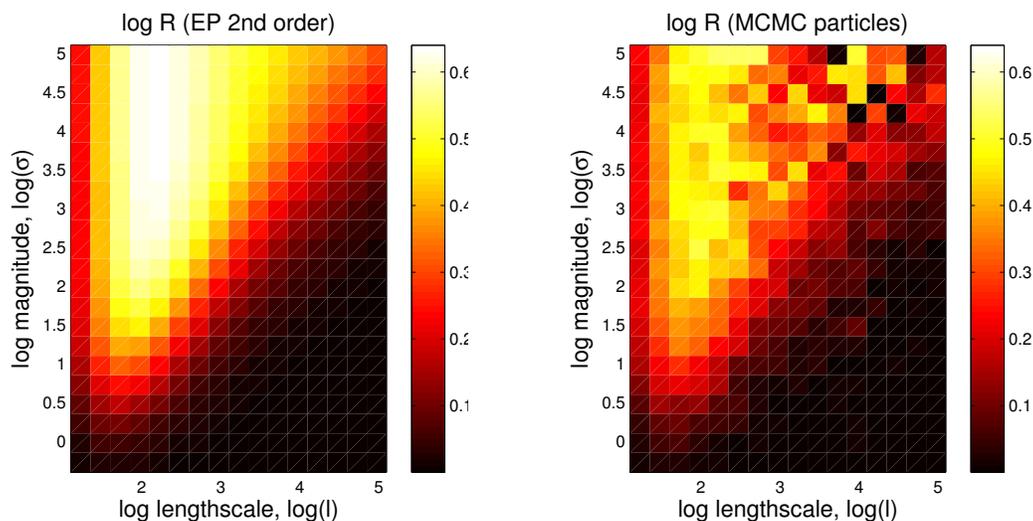


Figure 3: A comparison of $\log R$ using a perturbation expansion of Equation (28) against Monte Carlo estimates of $\log R$, using the USPS data set from Kuss and Rasmussen (2005). The second order correction to $\log R$, with $l = 3, 4$, is used on the *left*; the *right* plot uses a Monte Carlo estimate of $\log R$.

8.2 Uniform Noise Regression

We turn our attention to a regression problem, that of learning a latent function $x(\mathbf{s})$ from inputs $\{\mathbf{s}_n\}$ and matching real-valued observations $\{y_n\}$. A frequent nonparametric treatment assumes that $x(\mathbf{s})$ is *a priori* drawn from a GP prior with covariance function $k(\mathbf{s}, \mathbf{s}')$, from which a corrupted version y is observed. Analytically tractable inference is no longer possible in this model when the observation noise is non-Gaussian. Some scenarios include that of quantized regression, where y_n is formed by rounding $x(\mathbf{s}_n)$ to, say, the nearest integer, or where $x(\mathbf{s})$ indicates a robot’s path in a control problem, with conditions to stay within certain “wall” bounds. In these scenarios the latent function $x(\mathbf{s}_n)$ can be reconstructed from y_n by adding sharply discontinuous uniformly random $\mathcal{U}[-a, a]$ noise,

$$p(\mathbf{x}) = \frac{1}{Z} \prod_n \mathbb{I}[|x_n - y_n| < a] \mathcal{N}(\mathbf{x}; \mathbf{0}, \mathbf{K}) .$$

We now assume an EP approximation $q(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$, which can be obtained by using the moment calculations in Appendix E.2. To simplify the exposition of the predictive marginal, we follow the notation of Rasmussen and Williams (2005, Chapter 3) and let $\boldsymbol{\lambda}_n = (\tau_n, \nu_n)$, so that the final EP approximation multiplies g_n terms $\prod_n \exp\{-\frac{1}{2}\tau_n x_n^2 + \nu_n x_n\}$ into a joint Gaussian $\mathcal{N}(\mathbf{x}; \mathbf{0}, \mathbf{K})$.

8.2.1 MAKING PREDICTIONS FOR NEW DATA

The latent function $x(\mathbf{s}_*)$ at any new input \mathbf{s}_* is obtained by the predictive marginal $q(x_*)$ of $q(\mathbf{x}, x_*)$. The marginal $q(x_*)$ —given below in Equation (34)—is directly obtained from the EP approximation

$q(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$. However, the correction to its mean, as was given in Equation (29), requires covariances $\boldsymbol{\Sigma}_{*n}$, which are derived here.

Let $\boldsymbol{\kappa}_* = k(\mathbf{s}_*, \mathbf{s}_*)$, and \mathbf{k}_* be a vector containing the covariance function evaluations $k(\mathbf{s}_*, \mathbf{s}_n)$. Again following Rasmussen and Williams (2005)'s notation, let $\tilde{\boldsymbol{\Sigma}}$ be the diagonal matrix containing $1/\tau_n$ along its diagonal. The EP covariance, on the inclusion of x_* , is

$$\begin{aligned} \boldsymbol{\Sigma}_* &= \left(\begin{bmatrix} \mathbf{K} & \mathbf{k}_* \\ \mathbf{k}_*^T & \boldsymbol{\kappa}_* \end{bmatrix}^{-1} + \begin{bmatrix} \tilde{\boldsymbol{\Sigma}}^{-1} & \mathbf{0} \\ \mathbf{0}^T & 0 \end{bmatrix} \right)^{-1} \\ &= \begin{bmatrix} \boldsymbol{\Sigma} & \mathbf{k}_* - \mathbf{K}(\mathbf{K} + \tilde{\boldsymbol{\Sigma}})^{-1}\mathbf{k}_* \\ \mathbf{k}_*^T - \mathbf{k}_*^T(\mathbf{K} + \tilde{\boldsymbol{\Sigma}})^{-1}\mathbf{K} & \boldsymbol{\kappa}_* - \mathbf{k}_*^T(\mathbf{K} + \tilde{\boldsymbol{\Sigma}})^{-1}\mathbf{k}_* \end{bmatrix}, \end{aligned} \quad (33)$$

with $\boldsymbol{\Sigma} = \mathbf{K} - \mathbf{K}(\mathbf{K} + \tilde{\boldsymbol{\Sigma}})^{-1}\mathbf{K}$. There is no observation associated with \mathbf{s}_* , hence $\tau_* = 0$ in the first line above, and its inclusion has $c_{l*} = 0$ for $l \geq 3$. The second line follows by computing matrix partitioned inverses twice on $\boldsymbol{\Sigma}_*$. The joint EP approximation for any new input point \mathbf{s}_* is directly obtained as

$$q(\mathbf{x}, x_*) = \mathcal{N} \left(\begin{bmatrix} \mathbf{x} \\ x_* \end{bmatrix}; \begin{bmatrix} \boldsymbol{\mu} \\ \mathbf{k}_*^T \mathbf{K}^{-1} \boldsymbol{\mu} \end{bmatrix}, \boldsymbol{\Sigma}_* \right),$$

with the marginal $q(x_*)$ being

$$q(x_*) = \mathcal{N}(x_*; \mathbf{k}_*^T \mathbf{K}^{-1} \boldsymbol{\mu}, \boldsymbol{\kappa}_* - \mathbf{k}_*^T (\mathbf{K} + \tilde{\boldsymbol{\Sigma}})^{-1} \mathbf{k}_*) = \mathcal{N}(x_*; \mu_*, \sigma_*^2). \quad (34)$$

According to Equation (29), one needs the covariances $\boldsymbol{\Sigma}_{*j}$ to correct the marginal's mean; they appear in the last column of $\boldsymbol{\Sigma}_*$ in Equation (33). The correction is

$$\langle x_* \rangle_{p(\mathbf{x}, x_*)} - \mu_* = \sum_{l \geq 3} \sum_{j \neq n} \frac{\boldsymbol{\Sigma}_{*j}}{\boldsymbol{\Sigma}_{jj}} \frac{c_{l+1,j} c_{ln}}{l!} \left(\frac{\boldsymbol{\Sigma}_{jn}}{\boldsymbol{\Sigma}_{jj} \boldsymbol{\Sigma}_{nn}} \right)^l + \dots$$

The sum over pairs $j \neq n$ include the added dimension $*$, and thus pairs $(j, *)$ and $(*, n)$. The cumulants for this problem, used both for EP and correcting it, are derived in Appendix E.2.

8.2.2 PREDICTIVE CORRECTIONS

In Figure 4 we investigate the predictive mean correction for two cases, one where the data cannot realistically be expected to appear under the prior, and the other where the prior is reasonable. For $s \in \mathbb{R}$, the values of $x(s_*)$ are predicted using a GP with squared exponential covariance function $k(s, s') = \theta \exp(-\frac{1}{2}(s - s')^2/\ell)$.

In the first instance, the prior amplitude θ and lengthscale ℓ are deliberately set to values that are too big; in other words, a typical sample from the prior would not match the observed data. We illustrate the posterior marginal $q(x_*)$, and using Equations (29) and (30), show visible corrections to its mean and variance.⁴ For comparison, Figure 4 additionally shows what the predictive mean would have been were $\{y_n\}$ observed under Gaussian noise with the same mean and variance as $\mathcal{U}[-a, a]$: it is substantially different.

In the second instance, $\log Z_{\text{EP}}$ is maximized with respect to the covariance function hyperparameters θ and ℓ to get a kernel function that more reasonably describes the data. The correction

4. In the correction for the mean in Equation (29), we used $l = 3$ and $l = 4$ in the second order correction. For the correction to the variance in Equation (30), we used $l = 3$ in the first sum, and $l = 3$ and $l = 4$ in the second sum.

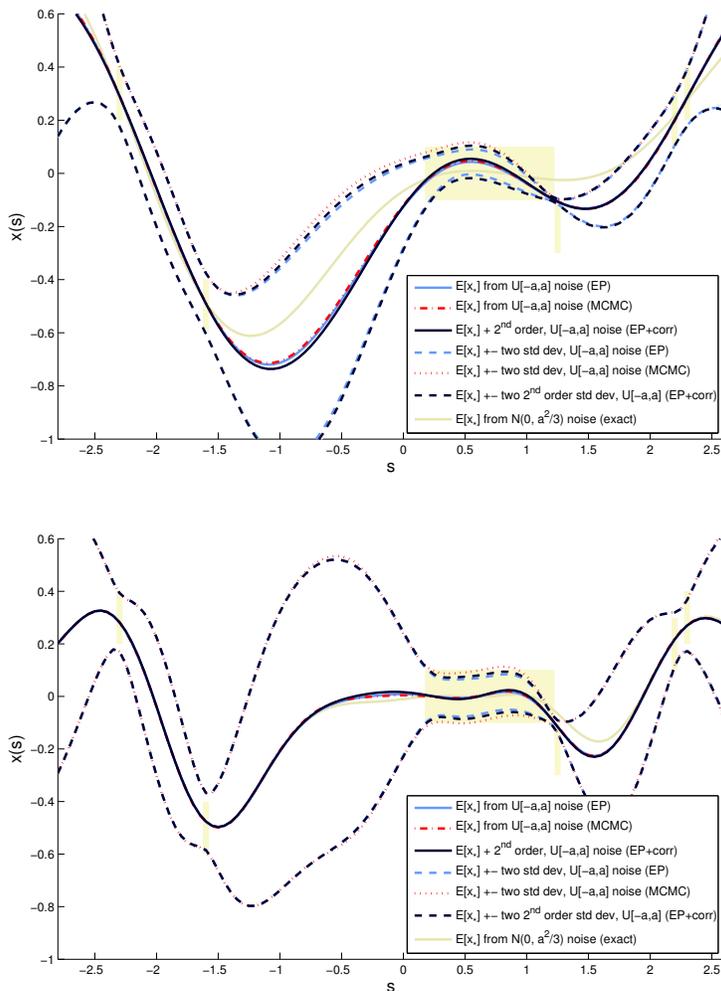


Figure 4: Predicting $x(s_*)$ with a GP. The “boxed” bars indicate the permissible $x(s_n)$ values; they are linked to observations y_n through the uniform likelihood $\mathbb{I}[|x_n - y_n| < a]$. Due to the $\mathcal{U}[-a, a]$ noise model, $q(x_*)$ is ambivalent to where in the “box” $x(s_*)$ is placed. A second order correction to the mean of $q(x_*)$ is shown in a dotted line. The lightly shaded function plots $p(x_*)$, if the likelihood was also Gaussian with variance matching that of the “box”. In the *top* figure both the prior amplitude θ and lengthscale ℓ are overestimated. In the *bottom* figure, θ and ℓ were chosen by maximizing $\log Z_{EP}$ with respect to their values. Notice the smaller EP approximation error.

to the mean of $q(s_*)$ is much smaller, and furthermore, generally follows the “Gaussian noise” posterior mean. When the observed data is not typical under the prior, the correction to $\langle x_* \rangle$ is substantially bigger than when the prior is representative of the data.

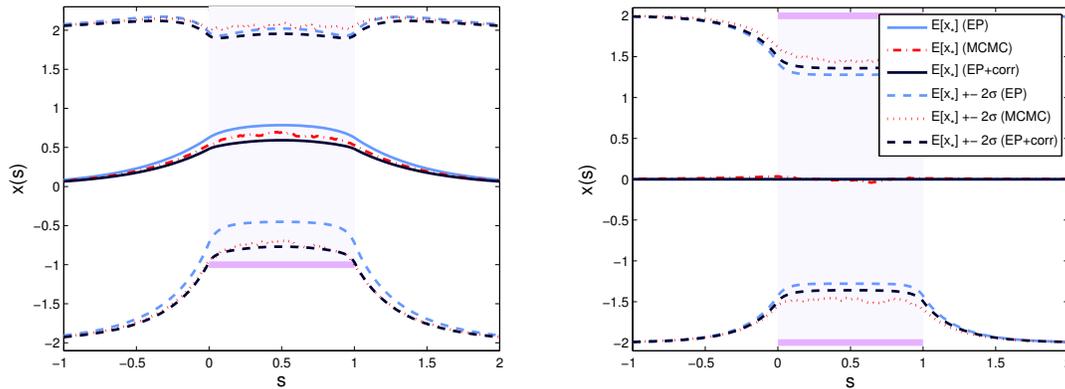


Figure 5: Predicting $x(s_*)$ with a GP with $k(s, s') = \exp\{-|s - s'|/2\ell\}$ and $\ell = 1$. In the *left* figure $\log R_{\text{MCMC}} = 0.41$, while the second order correction estimates it as $\log R \approx 0.64$. On the *right*, the correction to the variance is not as accurate as that on the *left*. The *right* correction is $\log R_{\text{MCMC}} = 0.28$, and its discrepancy with $\log R \approx 0.45$ (EP+corr) is much bigger.

8.2.3 UNDERESTIMATING THE TRUTH

Under closer inspection, the variance in Figure 4 is slightly underestimated in regions where there are many close box constraints $|x_n - y_n| < a$. However, under sparser constraints relative to the kernel width, EP accurately estimates the predictive mean and variance. In Figure 5 this is taken further: for $N = 100$ uniformly spaced inputs $s \in [0, 1]$, it is clear that $q(\mathbf{x})$ becomes too narrow. The second order correction, on the other hand, provides a much closer estimate to the ground truth.

One might inquire about the behavior of the EP estimate as $N \rightarrow \infty$ in Figure 5. In the next section, this will be used as a basis for illustrating a special case where $\log Z_{\text{EP}}$ *diverges*.

8.3 Gaussian Process in a Box

In the following insightful example—a special case of uniform noise regression— $\log Z_{\text{EP}}$ diverges from the ground truth with more data. Consider the ratio of functions $x(s)$ over $[0, 1]$, drawn from a GP prior with kernel $k(s, s')$, such that $x(s)$ lies within the $[-a, a]$ box. Figure 6 illustrates three random draws from a GP prior, two of which are not contained in the $[-a, a]$ interval. The ratio of functions contained in the interval is equal to the normalizing constant of

$$p(\mathbf{x}) = \frac{1}{Z} \prod_n \mathbb{I}[|x_n| < a] \mathcal{N}(\mathbf{x}; \mathbf{0}, \mathbf{K}). \quad (35)$$

The fraction of samples from the GP prior that lie inside $[-a, a]$ shouldn't change as the GP is sampled at increasing granularity of inputs s . As Figure 6 illustrates, the MCMC estimate of $\log Z$ converges to a constant as $N \rightarrow \infty$. The EP estimate $\log Z_{\text{EP}}$, on the other hand, diverges to $-\infty$. (The cumulants that are required for the correction in Equation (28), and recipes for deriving them, are given in Appendix E.1.) Of course the correction also depends on the value a chosen. Figure 7 shows that for both $a \rightarrow 0$ and $a \rightarrow \infty$ the correction is zero for large N .

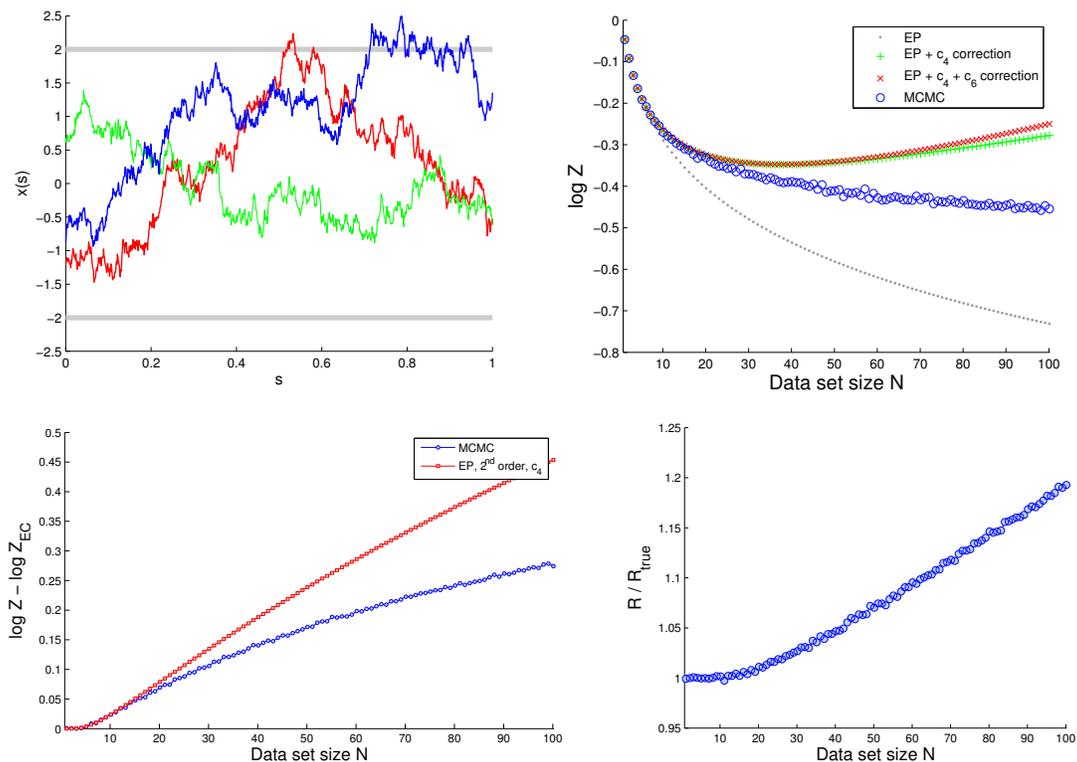


Figure 6: Samples from a GP prior with kernel $k(s, s') = \exp\{-|s - s'|/2\ell\}$ with $\ell = 1$, two of which are not contained in the $[-a, a]$ interval, are shown *top left*. As N increases in Equation (35), with $s_n \in [0, 1]$, $\log Z_{EP}$ diverges, while $\log Z$ converges to a constant. This is shown *top right*. The +’s and ×’s indicate the inclusion of the fourth (+) and fourth and sixth (×) cumulants from the 2nd order in Equation (28) (an arrangement by total order would include 3rd order $c_4 - c_4 - c_4$ in ×). *Bottom left* and *right* show the growth for 2nd order c_4 correction relative to the exact correction.

An intuitive explanation, due to Philipp Hennig, takes a one-dimensional model $p(x) = \mathbb{I}[|x| < a]^N \mathcal{N}(x; 0, 1)$. A fully-factorized approximation therefore has $N - 1$ *redundant* factors, as removing them doesn’t change $p(x)$. However, each additional $\mathbb{I}[|x| < a]$ truncates the estimate, forcing EP to further reduce the variance of $q(x)$. The EP estimate using N factors $\mathbb{I}[|x| < a]^{1/N}$ is correct (see Appendix C for a similar example and analysis), even though the original problem remains unchanged. Even though this immediate solution cannot be applied to Equation (35), the *redundancy* across factors could be addressed by a principled junction tree-like factorization, where tuples of “neighboring” factors can be co-treated. Although beyond the scope of this paper, Appendix A gives a guideline on how to structure such an approximation.

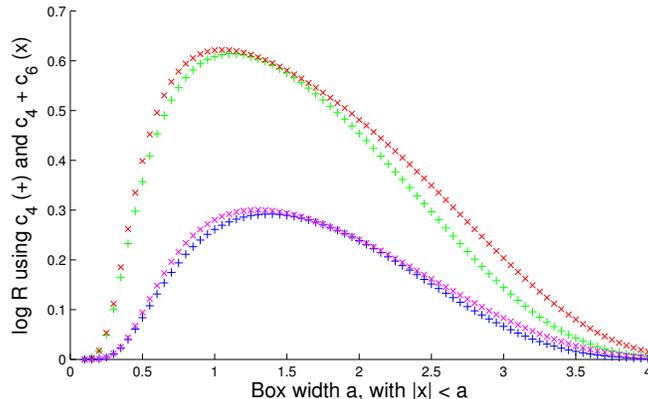


Figure 7: The accurateness of $\log Z_{EP}$ depends on the size of the $[-a, a]$ box relative to ℓ , with the estimation being exact as $a \rightarrow 0$ and $a \rightarrow \infty$. The second order correction for Figure 6’s kernel is illustrated here over varying a ’s. The $+$ ’s and \times ’s indicate the inclusion of the 4th ($+$) and 4th and 6th (\times) cumulants in Equation (28). Of these, the top pair of lines are for $N = 100$, and the bottom pair for $N = 50$.

9. Ising Model Results

This section discusses various aspects of corrections to EP as applied to the Ising model—a Bayesian network with binary variables and pairwise potentials—in Equation (3).

We consider the set-up proposed by Wainwright and Jordan (2006) in which $N = 16$ nodes are either fully connected or connected to their nearest neighbors in a 4-by-4 grid. The external field (observation) strengths θ_i are drawn from a *uniform* distribution $\theta_i \sim \mathcal{U}[-d_{\text{obs}}, d_{\text{obs}}]$ with $d_{\text{obs}} = 0.25$. Three types of coupling strength statistics are considered: repulsive (anti-ferromagnetic) $J_{ij} \sim \mathcal{U}[-2d_{\text{coup}}, 0]$, mixed $J_{ij} \sim \mathcal{U}[-d_{\text{coup}}, +d_{\text{coup}}]$, and attractive (ferromagnetic) $J_{ij} \sim \mathcal{U}[0, +2d_{\text{coup}}]$.

Previously we have shown (Opper and Winther, 2005) that EP/EC gives very competitive results compared to several standard methods. In Section 9.1 we are interested in investigating whether a further improvement is obtained with the cumulant expansion. In Section 9.2, we revisit the correction approach proposed in Paquet et al. (2009) and make an empirical comparison with the cumulant approach.

9.1 Cumulant Expansion

For the factorized approximation we use Equations (26) and (29) for the $\log Z$ and marginal corrections, respectively. The expression for the cumulants of the Ising model is given in Section 4.2.1. The derivation of the corresponding tree expressions may be found in Appendices B and E.4.

Table 1 gives the average absolute deviation (AAD) of marginals

$$\text{AAD} = \frac{1}{N} \sum_i \left| p(x_i = 1) - p(x_i = 1 | \text{method}) \right| = \frac{1}{2N} \sum_i |m_i - m_i^{\text{est}}|,$$

while Table 2 gives the absolute deviation of $\log Z$ averaged of 100 repetitions. In two cases (Grid, $d_{\text{coup}} = 2$ Repulsive and Attractive coupling) we observed some numerical problems with the EC

Problem type			AAD marginals				
Graph	Coupling	d_{coup}	LBP	LD	EC	EC c	EC t
Full	Repulsive	0.25	.037	.020	.003	.0006	.0017
		0.50	.071	.018	.031	.0157	.0143
	Mixed	0.25	.004	.020	.002	.0004	.0013
		0.50	.055	.021	.022	.0159	.0151
	Attractive	0.06	.024	.027	.004	.0023	.0025
		0.12	.435	.033	.117	.1066	.0211
Grid	Repulsive	1.0	.294	.047	.153	<i>.1693</i>	.0031
		2.0	.342	.041	.198	<i>.4244</i>	.0021
	Mixed	1.0	.014	.016	.011	<i>.0122</i>	.0018
		2.0	.095	.038	.082	<i>.0984</i>	.0068
	Attractive	1.0	.440	.047	.125	<i>.1759</i>	.0028
		2.0	.520	.042	.177	<i>.4730</i>	.0002

Table 1: Average absolute deviation (AAD) of marginals in a Wainwright-Jordan set-up, comparing loopy belief propagation (LBP), log-determinant relaxation (LD), EC, EC with $l = 4$ second order correction (EC c), and an EC tree (EC t). Results in bold face highlight best results, while italics indicate where the cumulant expression is less accurate than the original approximation.

tree solver. It might be some cases that a solution does not exist but we ascribe numerical instabilities in our implementation as the main cause for these problems. It is currently out of the scope of this work to come up with a better solver. We choose to report the average performance for those runs that could attain a high degree of expectation consistency: $\sum_{i=1}^N (\langle x_i \rangle_{q_i} - \langle x_i \rangle_q)^2 \leq 10^{-20}$. This was 69 out of 100 in the mentioned cases and 100 of 100 in the remaining.

We observe that for the Grid simulations, the corrected marginals in factorized approximation are less accurate than the original approximation. In Figure 8 we vary the coupling strength for a specific set-up (Grid Mixed) and observe a cross-over between the correction and original for the error on marginals as the coupling strength increases. We conjecture that when the error of the original solution is high then the number of terms needed in the cumulant correction increases. The estimation of the marginal seems more sensitive to this than the $\log Z$ estimate. The tree approximation is very precise for the whole coupling strength interval considered and the fourth order cumulant in the second order expansion is therefore sufficient to get often quite large improvements over the original tree approximation.

9.2 The ε -Expansion

In Paquet et al. (2009) we introduced an alternative expansion for R and applied it to Gaussian processes and mixture models. It is obtained from Equation (12) using a finite series expansion, where the normalized deviation

$$\varepsilon_n(x_n) = \frac{q_n(x_n)}{q(x_n)} - 1$$

Problem type			Absolute deviation log Z				
Graph	Coupling	d_{coup}	EC	EC c	EC ϵ c	EC t	EC tc
Full	Repulsive	0.25	.0310	.0018	.0061	.0104	.0010
		0.50	.3358	.0639	.0697	.1412	.0440
	Mixed	0.25	.0235	.0013	.0046	.0129	.0009
		0.50	.3362	.0655	.0671	.1798	.0620
	Attractive	0.06	.0236	.0028	.0048	.0166	.0006
		0.12	.8297	.1882	.2281	.2672	.2094
Grid	Repulsive	1.0	1.7776	.8461	.8124	.0279	.0115
		2.0	4.3555	2.9239	3.4741	.0086	.0077
	Mixed	1.0	.3539	.1443	.0321	.0133	.0039
		2.0	1.2960	.7057	.4460	.0566	.0179
	Attractive	1.0	1.6114	.7916	.7546	.0282	.0111
		2.0	4.2861	2.9350	3.4638	.0441	.0433

Table 2: Absolute deviation log partition function in a Wainwright-Jordan set-up, comparing EC, EC with $l = 4$ second order correction (EC c), EC with a full second order ϵ expansion (EC ϵ c), EC tree (EC t) and EC tree with $l = 4$ second order correction (EC tc). Results in bold face highlight best results. The cumulant expression is consistently more accurate than the original approximation.

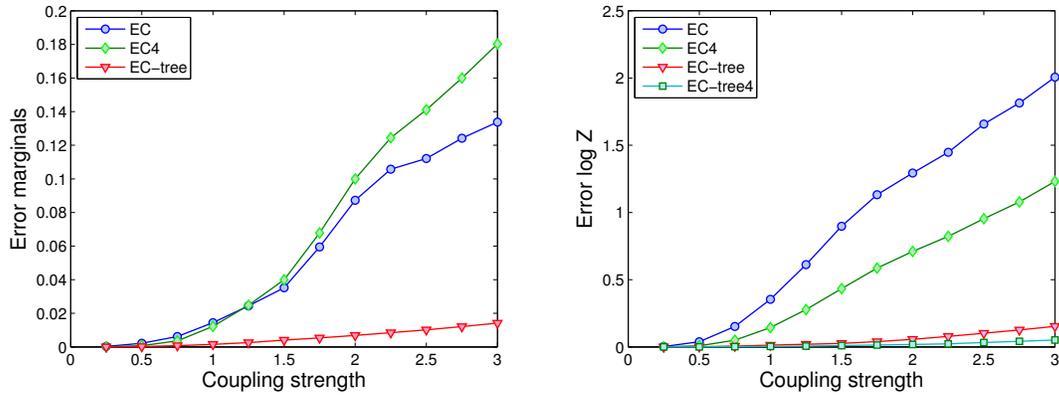


Figure 8: Error on marginal (*left*) and log Z (*right*) for grid and mixed couplings as a function of coupling strength.

is treated as the small quantity instead of higher order cumulants. R has an exact representation with 2^N terms that we may truncate at lowest non-trivial order:

$$R = \left\langle \prod_n (1 + \epsilon_n(x_n)) \right\rangle_{q(\mathbf{x})} \approx 1 + \sum_{m < n} \langle \epsilon_m(x_m) \epsilon_n(x_n) \rangle + O(\epsilon^3).$$

The linear terms are all equal to one because $\left\langle \frac{q_n(x_n)}{q(x_n)} \right\rangle_q = \int q(x_n) \frac{q_n(x_n)}{q(x_n)} dx_n = 1$ and since $q_n(x_n)$ is a binary distribution the quadratic term becomes a weighted sum of ratios of Normal distributions:

$$\left\langle \frac{q_m(x_m)}{q(x_m)} \right\rangle_{q(\mathbf{x})} = \sum_{x_n, x_m = \pm 1} \frac{1 + x_m m_m}{2} \frac{1 + x_n m_n}{2} \frac{q(x_m, x_n)}{q(x_m)q(x_n)}.$$

The final expression for the lowest order approximation to R is then

$$R \approx 1 + \sum_{m < n} \sum_{x_n, x_m = \pm 1} \frac{1 + x_m m_m}{2} \frac{1 + x_n m_n}{2} \frac{q(x_m, x_n)}{q(x_m)q(x_n)} - \frac{N(N-1)}{2}.$$

From Table 2 we observe an improvement over the original factorized approximation and results similar to the cumulant correction to the factorized approximation for all settings. The ϵ -expansion may also be used to calculate marginals and applied to generalized factorizations. These topics will be studied elsewhere.

10. Future Directions

Corrections to Gaussian EP approximations were examined in this paper. The Gaussian measure allowed for a convenient set of mathematical tools to be employed, mostly because it admits orthogonality of a set of polynomials, the Hermite polynomials, which allowed a clean simplification of many expressions. So far we have restricted ourselves to expansions to low orders in cumulants. Our results indicate that these first corrections to EP can already provide useful information about the quality of the EP solution. Small corrections typically show that EP is fairly accurate and the corrections improve on that. On the other hand, large corrections indicate that the EP approximation performs poorly. The low order corrections can yield a step in the right direction but in general their result may not be trusted and alternatives to the Gaussian EP approximation should be considered. It will be interesting to develop similar expansions to EP approximations with other exponential families besides the Gaussian one.

Can we expect that higher order terms in the cumulant expansion will give more reliable approximations? Before such a question could be attacked one first would need to decide in which order the terms of the expansion should be evaluated in order to obtain the most dominant contributions. For example, we might think of trying to first compute all terms in the second order expansion of the exponential in Equation (26), and then move on to higher orders. An alternative is to sort the expansion by the total sum of the orders of cumulants involved. This is in fact possible by introducing a suitable expansion parameter (which is later set equal to one) such that the formal Taylor series with respect to this parameter yields the desired expansion. However, it is not clear yet if and when such a power series expansion would actually converge. It may well be that our expansions are only of an asymptotic type (Boyd, 1999) for which the summation of only a certain number of terms might give an improvement whereas further terms would lead to worse results.

We expect that such questions could at least be answered for toy models such as the *Gaussian process in a box* model of Section 8.3. Our results for the latter example (together with the related *uniform noise regression* case) indicates that EP may not be understood as an off the shelf method for approximately calculating arbitrary high dimensional sums or integrals. One may conjecture that its quality strongly depends on the fact that such sums or integrals may or may not have an interpretation in terms of a proper statistical inference model which contain data that are highly

probable with respect to the model. It would be interesting to see if one can develop a theory for the average case performance of EP under such statistical assumptions of the data.

Appendix A. Factorizations: Gaussian Examples

As $p(\mathbf{x})$ is a latent Gaussian model, the g -terms in Equation (5) are chosen in this paper to give a Gaussian approximation

$$q(\mathbf{x}) = \frac{1}{Z_q} \exp\{\boldsymbol{\lambda}^T \phi(\mathbf{x})\} = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}).$$

The sufficient statistics $\phi(\mathbf{x})$ and natural parameters $\boldsymbol{\lambda}$ of the Gaussian are defined as

$$\phi(\mathbf{x}) = (\mathbf{x}, -\frac{1}{2}\mathbf{x}\mathbf{x}^T) \quad \text{and} \quad \boldsymbol{\lambda} = (\boldsymbol{\gamma}, \boldsymbol{\Lambda}),$$

where $\boldsymbol{\lambda}^T \phi(\mathbf{x}) = \boldsymbol{\gamma}^T \mathbf{x} - \frac{1}{2} \text{tr}[\boldsymbol{\Lambda} \mathbf{x} \mathbf{x}^T] = \boldsymbol{\gamma}^T \mathbf{x} - \frac{1}{2} \mathbf{x}^T \boldsymbol{\Lambda} \mathbf{x}$. There exists a bijection between the canonical parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ and natural parameters, such that the mean and covariance can be determined with $\boldsymbol{\Sigma} = \boldsymbol{\Lambda}^{-1}$ and $\boldsymbol{\mu} = \boldsymbol{\Sigma} \boldsymbol{\gamma}$.

In Equation (1) we can define $g_0(\mathbf{x}) = \exp\{\boldsymbol{\lambda}_0^T \phi(\mathbf{x})\}$, where $\boldsymbol{\lambda}_0 = (\boldsymbol{\gamma}^{(0)}, \boldsymbol{\Lambda}^{(0)})$, such that it is essentially a rescaling of factor f_0 . In the Ising model in Equation (3), this means that $\boldsymbol{\Lambda}^{(0)} = -\mathbf{J}$ and $\boldsymbol{\gamma}^{(0)} = \boldsymbol{\theta}$. In the Gaussian process classification model in Equation (2), this implies that $\boldsymbol{\Lambda}^{(0)} = \mathbf{K}^{-1}$ and $\boldsymbol{\gamma}^{(0)} = \mathbf{0}$.

A.1 Term-Wise Factorizations

It remains to define a suitable factorization for the term-product $\prod_n t_n(x_n)$. This factorization can be fully factorized, factorized over disjoint sets of variables, factorized as a tree, or follow more arbitrary factorizations (see the simple example in Appendix C). A few such factorizations are given below in increasing orders of complexity. *In each case we do not include the f_0 factor for clarity.* Furthermore, even though the term factorization may be chosen to fully factorize, $q(\mathbf{x})$ may be fully connected through the inclusion of f_0 .

A.1.1 FULLY FACTORIZED

A common factorization of $\prod_n t_n(x_n)$ is to set $f_n(\mathbf{x}) = t_n(x_n)$. The natural parameters of $g_n(\mathbf{x}) = \exp\{\boldsymbol{\lambda}_n^T \phi(\mathbf{x})\}$ are chosen to be $\boldsymbol{\lambda}_n = (\boldsymbol{\gamma}_n^{(n)}, \boldsymbol{\Lambda}_{nn}^{(n)})$, corresponding to $\phi_n(x_n) = (x_n, -\frac{1}{2}x_n^2)$. For clarity the other $\boldsymbol{\gamma}$ and $\boldsymbol{\Lambda}$ parameters in $\boldsymbol{\lambda}_n$ are not shown, as they are clamped at zero. This gives an approximation $q(\mathbf{x})$ that is defined by $\boldsymbol{\lambda} = \boldsymbol{\lambda}_0 + \sum_n \boldsymbol{\lambda}_n$.

A.1.2 FACTORIZATION INTO DISJOINT PAIRS

As a second step the N variables can be subdivided into *disjoint* pairs $\mathbf{x}_\pi = (x_m, x_n)$. The factorization over terms couples pairs of variables through

$$\prod_n t_n(x_n) = \prod_{\pi=(m,n)} [t_m(x_m)t_n(x_n)] = \prod_{\pi} f_{\pi}(\mathbf{x}).$$

In this case each factor will have a contribution $g_{\pi}(\mathbf{x}) = \exp\{\boldsymbol{\lambda}_{\pi}^T \phi(\mathbf{x})\}$ to the overall approximation, and, as g_{π} is a function of two variables, it is parameterized by the ‘‘correlated Gaussian form’’ $\boldsymbol{\lambda}_{\pi} = (\boldsymbol{\gamma}_m^{(\pi)}, \boldsymbol{\gamma}_n^{(\pi)}, \boldsymbol{\Lambda}_{mm}^{(\pi)}, \boldsymbol{\Lambda}_{nn}^{(\pi)}, \boldsymbol{\Lambda}_{mn}^{(\pi)})$. By symmetry $\boldsymbol{\Lambda}_{nm}^{(\pi)} = \boldsymbol{\Lambda}_{mn}^{(\pi)}$. The resulting $q(\mathbf{x})$ is defined in terms of these disjoint sets with $\boldsymbol{\lambda} = \boldsymbol{\lambda}_0 + \sum_{\pi} \boldsymbol{\lambda}_{\pi}$.

A.1.3 TREE-STRUCTURED FACTORIZATION

A tree structure factorization can be defined by extending the above “disjoint pairs” case to allow for overlaps between terms. Let \mathcal{G} define a spanning tree structure over all \mathbf{x} , and let $\tau = (m, n) \in \mathcal{G}$ define the edges in the tree. Let d_n be the number of edges emanating from node x_n in the graph. Through a clever regrouping of terms into a “junction tree” form with

$$\prod_n t_n(x_n) = \frac{\prod_{\tau=(m,n)} [t_m(x_m)t_n(x_n)]}{\prod_n t_n(x_n)^{d_n-1}} = \frac{\prod_{\tau} f_{\tau}(\mathbf{x})}{\prod_n f_n(\mathbf{x})^{d_n-1}},$$

the term-approximation will be tree-structured. In this example the D_a powers are 1 for edge factors f_{τ} and $(1 - d_n)$ for node factors f_n . Let $g_{\tau}(\mathbf{x})$ and $g_n(\mathbf{x})$ be parameterized by λ_{τ} and λ_n , as was done in the two examples above. Using

$$\frac{\prod_{\tau} g_{\tau}(\mathbf{x})}{\prod_n g_n(\mathbf{x})^{d_n-1}} = \frac{\prod_{\tau} \exp\{\lambda_{\tau}^T \phi(\mathbf{x})\}}{\prod_n \exp\{\lambda_n^T \phi(\mathbf{x})\}^{d_n-1}},$$

the resulting $q(\mathbf{x})$ has parameter vector $\lambda = \lambda_0 + \sum_{\tau} \lambda_{\tau} - \sum_n (d_n - 1)\lambda_n$.

It is useful to note that the form of the tree-structured approximation given here is that used by Oppor and Winther (2005); it approximates the “junction tree” form using a Power EP factorization (Minka, 2004). The factorization and stationary condition is *different* from that of Tree EP (Minka and Qi, 2004).

A.2 Stationary Point

The EP moment matching conditions from Equation (7) are uniquely met at the stationary point of $\log Z_{\text{EP}}$ in Equation (8), and are shown here. Consider the logarithm of the normalizer,

$$\log Z_{\text{EP}} = \log Z_q + \sum_a D_a \log Z_a. \quad (36)$$

Using the sufficient statistics and natural parameters defined above, the two normalizers that constitute Equation (36) are

$$Z_q = \int e^{\sum_a D_a \lambda_a^T \phi(\mathbf{x})} d\mathbf{x},$$

$$Z_a = \frac{1}{Z_q} \int e^{\sum_b D_b \lambda_b^T \phi(\mathbf{x}) - \lambda_a^T \phi(\mathbf{x})} f_a(\mathbf{x}) d\mathbf{x}.$$

Using these definitions, the derivatives of the terms in Equation (36) with respect to some EP factor c 's parameters λ_c are

$$\frac{\partial \log Z_q}{\partial \lambda_c} = D_c \langle \phi(\mathbf{x}) \rangle_q,$$

$$\frac{\partial \log Z_a}{\partial \lambda_c} = \begin{cases} D_c \langle \phi(\mathbf{x}) \rangle_{q_a} - D_c \langle \phi(\mathbf{x}) \rangle_q & \text{if } c \neq a \\ (D_c - 1) \langle \phi(\mathbf{x}) \rangle_{q_c} - D_c \langle \phi(\mathbf{x}) \rangle_q & \text{if } c = a \end{cases}.$$

When $\partial \log Z_{\text{EP}} / \partial \lambda_c = \mathbf{0}$ for any c , the following therefore holds:

$$\mathbf{0} = (D_c - 1)(\langle \phi(\mathbf{x}) \rangle_{q_c} - \langle \phi(\mathbf{x}) \rangle_q) + \sum_{a \neq c} D_a (\langle \phi(\mathbf{x}) \rangle_{q_a} - \langle \phi(\mathbf{x}) \rangle_q).$$

Let \mathbf{D} be a square matrix where the values in column a are D_a ; all the rows in \mathbf{D} are equal and it is singular. Furthermore, let $\psi_a = \langle \phi(\mathbf{x}) \rangle_{q_a} - \langle \phi(\mathbf{x}) \rangle_q$. By stacking all the ψ_a 's into a column vector ψ , the above set of equalities lead to a system of equations

$$\mathbf{0} = ((\mathbf{D} - \mathbf{I}) \otimes \mathbf{I}_{\text{dim}}) \psi .$$

(The Kronecker product is only required as the sufficient statistics' differences ψ_a have dimensionality "dim", usually larger than one.) As $\mathbf{D} - \mathbf{I}$ is nonsingular, it is solved by $\psi = \mathbf{0}$, and hence $\langle \phi(\mathbf{x}) \rangle_{q_a} = \langle \phi(\mathbf{x}) \rangle_q$ for all a .

The choice of parameterization of λ_a might give an overcomplete representation, and the exact moment-matching conditions $\langle \phi(\mathbf{x}) \rangle_{q_a} = \langle \phi(\mathbf{x}) \rangle_q$ might have more than one unique solution. However, this does not invalidate that at the stationary point of Equation (36), all moment-matching conditions must hold.

Appendix B. Tree-Structured Approximation

Let the factorization of the term-product $\prod_n t_n(x_n)$ take the form of a tree \mathcal{G} with edges $\tau = (m, n) \in \mathcal{G}$, as is described in Appendix A.1.3. The number connections to a node or vertex n shall be denoted by d_n . From Equation (32) the second order expansion is

$$\begin{aligned} \log R &= \frac{1}{2} \sum_{\tau \neq \tau'} \langle \langle r_\tau \rangle \langle r_{\tau'} \rangle \rangle + \frac{1}{2} \sum_{m \neq n} (1 - d_m)(1 - d_n) \langle \langle r_m \rangle \langle r_n \rangle \rangle \\ &\quad + \sum_{\tau, n} (1 - d_n) \langle \langle r_\tau \rangle \langle r_n \rangle \rangle + \frac{1}{2} \sum_n (1 - d_n)(-d_n) \langle \langle r_n \rangle^2 \rangle + \dots , \end{aligned} \quad (37)$$

where the inner expectations are over $\mathbf{k}_\tau | \mathbf{x}$ and $k_n | \mathbf{x}$, while the outer expectations are over \mathbf{x} .⁵ The edge-edge, edge-node, and node-node expectations that are needed in Equation (37) are given in the following three sections.

B.1 Edge-Edge Expectations

The edge-edge expectation provides a beautiful illustration of the combinatorics that may be involved in Wick's theorem. For $\tau \neq \tau'$, the following expectation needs to be evaluated:

$$\begin{aligned} &\langle \langle r_\tau(\mathbf{k}_\tau) \rangle \langle r_{\tau'}(\mathbf{k}_{\tau'}) \rangle \rangle \\ &= \left\langle \sum_{l \geq 3} \sum_{s \geq 3} i^{l+s} \left\{ \sum_{|\alpha|=l} \frac{c_{\alpha\tau}}{\alpha!} \langle \mathbf{k}_\tau^\alpha \rangle_{\mathbf{k}_\tau | \mathbf{x}} \right\} \left\{ \sum_{|\alpha'|=s} \frac{c_{\alpha'\tau'}}{\alpha'!} \langle \mathbf{k}_{\tau'}^{\alpha'} \rangle_{\mathbf{k}_{\tau'} | \mathbf{x}} \right\} \right\rangle_{\mathbf{x}} . \end{aligned} \quad (38)$$

The vectors α that are summed over to get $|\alpha| = l$ are $\alpha = (0, l), (1, l-1), \dots, (l, 0)$; let $\alpha = (\alpha_1, l - \alpha_1)$ when $|\alpha| = l$. From the independence of $\mathbf{k}_\tau | \mathbf{x}$ and $\mathbf{k}_{\tau'} | \mathbf{x}$,

$$\left\langle \langle \mathbf{k}_\tau^\alpha \rangle_{\mathbf{k}_\tau | \mathbf{x}} \langle \mathbf{k}_{\tau'}^{\alpha'} \rangle_{\mathbf{k}_{\tau'} | \mathbf{x}} \right\rangle_{\mathbf{x}} = \left\langle \langle \mathbf{k}_\tau^\alpha \mathbf{k}_{\tau'}^{\alpha'} \rangle_{\mathbf{k}_\tau, \mathbf{k}_{\tau'} | \mathbf{x}} \right\rangle_{\mathbf{x}} = \left\langle k_{\tau_1}^{\alpha_1} k_{\tau_2}^{l-\alpha_1} k_{\tau'_1}^{\alpha'_1} k_{\tau'_2}^{s-\alpha'_1} \right\rangle_{\mathbf{k}_\tau, \mathbf{k}_{\tau'}} , \quad (39)$$

5. Some readers might wonder why there is no $\frac{1}{2}$ associated with the sum over (τ, n) in Equation (37). In the other quadratic sums, for example over $m \neq n$, each (m, n) pair appears twice, as $r_m r_n$ and as $r_n r_m$. Each edge-node pair makes only one appearance in the sum; if the sum double-counted by including node-edge pairs, a division by two would have been necessary.

and therefore $\langle\langle r_\tau \rangle\rangle \langle\langle r_{\tau'} \rangle\rangle = \langle\langle r_\tau r_{\tau'} \rangle\rangle$ whenever $\tau \neq \tau'$.

Wick's theorem is again instrumental in computing $\langle\mathbf{k}_\tau^\alpha \mathbf{k}_{\tau'}^{\alpha'}\rangle$, as all possible pairings of the random variables $\mathbf{k}_\tau = (k_{\tau_1}, k_{\tau_2})$ and $\mathbf{k}_{\tau'} = (k_{\tau'_1}, k_{\tau'_2})$ need to be included. As $\langle k_{\tau_1}^2 \rangle = 0$, $\langle k_{\tau_1} k_{\tau_2} \rangle = 0$, $\langle k_{\tau'_1}^2 \rangle = 0$, and $\langle k_{\tau'_1} k_{\tau'_2} \rangle = 0$, the only non-zero expectations in the Wick expansion of Equation (39) occur when *all* the variables in \mathbf{k}_τ and $\mathbf{k}_{\tau'}$ are paired. This immediately means that $\langle k_{\tau_1}^{\alpha_1} k_{\tau_2}^{l-\alpha_1} k_{\tau'_1}^{\alpha'_1} k_{\tau'_2}^{s-\alpha'_1} \rangle = 0$ whenever $l \neq s$, as there will be some remaining variables in \mathbf{k}_τ (or $\mathbf{k}_{\tau'}$) that can't be paired and have to be self-paired with zero expectation.

Given $l = s$, evaluate the expectation in Equation (39). We introduce the ‘‘pairing count’’ vector β with elements $\beta_j \in \mathbb{N}_0$ and constraint $\sum_{j=1}^4 \beta_j = l$. Let β_1 count the number of pairings of k_{τ_1} with $k_{\tau'_1}$, and β_2 count the number of pairings of k_{τ_1} with $k_{\tau'_2}$. As there are α_1 k_{τ_1} terms, the sum of its outgoing pairings should equal α_1 with

$$\beta_1 + \beta_2 = \alpha_1 .$$

A furthermore requirement is that

$$\beta_1 + \beta_3 = \alpha'_1 , \quad \beta_3 + \beta_4 = \alpha_2 , \quad \beta_2 + \beta_4 = \alpha'_2 ,$$

where $\alpha_2 = l - \alpha_1$ and $\alpha'_2 = l - \alpha'_1$, and β_3 and β_4 be as in the Wick expansion below. Define \mathcal{B} to be the set of all such β 's, and let $C(\beta)$ count the number of permuted configurations for a given pairing β . From Wick's theorem the expected value is equal to the sum over all possible pairings β :

$$\left\langle k_{\tau_1}^{\alpha_1} k_{\tau_2}^{\alpha_2} k_{\tau'_1}^{\alpha'_1} k_{\tau'_2}^{\alpha'_2} \right\rangle_{\mathbf{k}_\tau, \mathbf{k}_{\tau'}} = \sum_{\beta \in \mathcal{B}} C(\beta) \langle k_{\tau_1} k_{\tau'_1} \rangle^{\beta_1} \langle k_{\tau_1} k_{\tau'_2} \rangle^{\beta_2} \langle k_{\tau_2} k_{\tau'_1} \rangle^{\beta_3} \langle k_{\tau_2} k_{\tau'_2} \rangle^{\beta_4} .$$

A simple scheme to enumerate all $\beta \in \mathcal{B}$ is to let

$$\beta = \left[\beta_1, \alpha_1 - \beta_1, \alpha'_1 - \beta_1, (l + \beta_1) - (\alpha_1 + \alpha'_1) \right],$$

so that $\beta \in \mathcal{B}$ for each $\beta_1 \in \{\max(0, (\alpha_1 + \alpha'_1) - l), \dots, \min(\alpha_1, \alpha'_1)\}$. The remaining components of β are uniquely determined from β_1 .

B.1.1 COUNTING PAIRINGS

How many permuted pairings $C(\beta)$ are there?

1. There are $\binom{\alpha_1}{\beta_1}$ ways of choosing β_1 k_{τ_1} 's, and then $\frac{\alpha'_1!}{(\alpha'_1 - \beta_1)!}$ ways of choosing $k_{\tau'_1}$ to pair with.
2. This leaves a remaining $(\alpha_1 - \beta_1)$ k_{τ_1} 's, that need to be paired with $(l - \alpha'_1)$ $k_{\tau'_2}$'s. There are $\frac{(l - \alpha'_1)!}{((l - \alpha'_1) - (\alpha_1 - \beta_1))!}$ such pairings.
3. There are also $\alpha'_1 - \beta_1$ remaining $k_{\tau'_1}$'s, that need to be paired with k_{τ_2} variables. There are $\binom{l - \alpha_1}{\alpha'_1 - \beta_1}$ ways of picking a k_{τ_2} , and a further $(\alpha'_1 - \beta_1)!$ ways of arranging the remaining $k_{\tau'_1}$.
4. Finally, the $(l - \alpha'_1) - (\alpha_1 - \beta_1)$ remaining k_{τ_2} 's need to be coupled with the remaining $k_{\tau'_2}$'s, and there are $((l - \alpha'_1) - (\alpha_1 - \beta_1))!$ such arrangements.

Multiplying the possible pairings from the four steps above gives

$$\begin{aligned} C(\beta) &= \binom{\alpha_1}{\beta_1} \frac{\alpha'_1!}{(\alpha'_1 - \beta_1)!} \frac{(l - \alpha'_1)!}{((l + \beta_1) - (\alpha_1 + \alpha'_1))!} \cdots \\ &\quad \cdots \times \binom{l - \alpha_1}{\alpha'_1 - \beta_1} (\alpha'_1 - \beta_1)! ((l + \beta_1) - (\alpha_1 + \alpha'_1))! \\ &= \binom{\alpha_1}{\beta_1} \alpha'_1! (l - \alpha'_1)! \binom{l - \alpha_1}{\alpha'_1 - \beta_1}, \end{aligned}$$

which adds up to the total number of possible pairings $\sum_{\beta \in \mathcal{B}} C(\beta) = l!$. A further useful simplification is $C(\beta)/\alpha! \alpha'! = 1/\beta!$ when $|\alpha| = |\alpha'| = l$, and is used below.

B.1.2 EDGE-EDGE EXPECTATION

The absence of any self-interacting loops from Wick's theorem lets the $\sum_{s \geq 3}$ drop away in Equation (38), as all terms are zero except for when $l = s$. Substituting $\langle \mathbf{k}_\tau^\alpha \mathbf{k}_{\tau'}^{\alpha'} \rangle$ and $C(\beta)$ into Equation (38) gives the final result,

$$\begin{aligned} &\langle \langle r_\tau(\mathbf{k}_\tau) \rangle \langle r_{\tau'}(\mathbf{k}_{\tau'}) \rangle \rangle \\ &= \sum_{l \geq 3} (-1)^l \sum_{|\alpha|=l} \sum_{|\alpha'|=l} c_{\alpha\tau} c_{\alpha'\tau'} \left\{ \sum_{\beta \in \mathcal{B}} \frac{1}{\beta!} \langle k_{\tau_1} k_{\tau'_1} \rangle^{\beta_1} \langle k_{\tau_1} k_{\tau'_2} \rangle^{\beta_2} \langle k_{\tau_2} k_{\tau'_1} \rangle^{\beta_3} \langle k_{\tau_2} k_{\tau'_2} \rangle^{\beta_4} \right\}. \end{aligned}$$

B.2 Edge-Node Expectations

The derivation for the edge-node expectations is similar to that of the edge-edge case,

$$\begin{aligned} \langle \langle r_\tau(\mathbf{k}_\tau) \rangle \langle r_n(k_n) \rangle \rangle &= \left\langle \sum_{l \geq 3} \sum_{s \geq 3} i^{l+s} \sum_{|\alpha|=l} \frac{c_{\alpha\tau} c_{sn}}{\alpha! s!} \langle \mathbf{k}_\tau^\alpha \rangle_{\mathbf{k}_\tau | \mathbf{x}} \langle k_n^s \rangle_{k_n | \mathbf{x}} \right\rangle_{\mathbf{x}} \\ &= \sum_{l \geq 3} (-1)^l \sum_{|\alpha|=l} \frac{c_{\alpha\tau} c_{ln}}{\alpha!} \langle k_{\tau_1} k_n \rangle^{\alpha_1} \langle k_{\tau_2} k_n \rangle^{l - \alpha_1}, \end{aligned}$$

where the expectations in the last line are again over $\{\mathbf{k}_\tau, k_n\}$. When $\langle \mathbf{k}_\tau^\alpha k_n^s \rangle$ is evaluated with Wick's theorem, there are α_1 copies of k_{τ_1} , $l - \alpha_1$ copies of k_{τ_2} , and s copies of k_n . The zero relation of \mathbf{k}_τ and k_n ensures that the only non-zero terms in the Wick sum are those where all the k_τ 's are paired with k_n 's; in other words, when $l = s$. There are $l!$ possible pairings, which cancels $l!$ in the denominator.

The above edge-node expectation is for any edge and node in the tree, but notice that it simplifies greatly when the edge τ is a connection to node n . Say τ_1 is the edge variable corresponding to x_n . In this case the covariance with respect to the *opposite* pair is zero, with $\langle k_{\tau_2}, k_n \rangle = 0$ (see Figure 2) and only *one* of the α 's will have a non-zero contribution to the sum, namely when $\alpha = (l, 0)$.

B.3 Node-Node Expectations

The node-node expectation is given in Equation (27), and is also used for $\langle \langle r_n \rangle^2 \rangle$.⁶

6. Due to the square in $\langle \langle r_n \rangle_{k_n | \mathbf{x}}^2 \rangle_{\mathbf{x}}$, the inner average $\langle r_n \rangle_{k_n | \mathbf{x}}$ should first be computed to give an expansion over Hermite polynomials in $x_n - \mu_n$. An example of such a result is given Appendix C. The orthogonality of these polynomials over $q(\mathbf{x} - \boldsymbol{\mu})$ allows $\langle \langle r_n \rangle_{k_n | \mathbf{x}}^2 \rangle_{\mathbf{x}}$ to also reduce to Equation (27).

Appendix C. A Tractable, One-Dimensional Example

The following example illustrates a tractable one-dimensional model with two factors. It is shown analytically that the correction to $\log Z_{EP}$ must be zero, and that the result is reflected in the higher-order terms in Equation (32), which are also zero.

Consider the factorization of a probit term with a Gaussian prior into

$$p(x) = \frac{1}{Z} \Phi(x) \mathcal{N}(x; 0, 1) = \frac{1}{Z} f_a(x)^{1/2} f_b(x)^{1/2} \mathcal{N}(x; 0, 1) ,$$

where $\Phi(x)$ is the cumulative Gaussian density function, and $f_a(x) = f_b(x) = \Phi(x)$. Z can be computed exactly, but for the sake of example $p(x)$ will be approximated with

$$q(x) = \frac{1}{Z_q} g_a(x)^{1/2} g_b(x)^{1/2} \mathcal{N}(x; 0, 1) = \mathcal{N}(x; \mu, \sigma^2) .$$

Choose $g_a(x) = \exp\{\phi(x)^T \lambda_a\}$, and $g_b(x) = \exp\{\phi(x)^T \lambda_b\}$. The q approximation has parameter vector $\lambda = \lambda_0 + \frac{1}{2} \lambda_a + \frac{1}{2} \lambda_b$. The EP fixed point is defined by $\lambda_a = \lambda_b$ and $Z_a = Z_b$. (For example, subtracting λ_a at the fixed point will leave $\lambda_a = \lambda_0 + \mathbf{0}$, which is equal to a scaled version of the prior $f_0(x)$. The factor $f_a(x) = \Phi(x)$ is hence incorporated into the prior, giving Z_a . By a symmetric argument, $Z_a = Z_b$.) Although it is trivial to show that $Z_{EP} = Z_q Z_a^{1/2} Z_b^{1/2}$ will be equal to the true partition function Z , we shall prove it by showing that the correction term is $\log R = 0$.

C.1 Analytic Correction

In this section a transformation of variables from x to $y \sim \mathcal{N}(y; 0, 1)$, with $y = (x - \mu)/\sigma$, will be used to make the derivation slightly simpler, and therefore

$$k_a|y \sim \mathcal{N}\left(k_a; -\frac{iy}{\sigma}, \sigma^{-2}\right) , \quad k_b|y \sim \mathcal{N}\left(k_b; -\frac{iy}{\sigma}, \sigma^{-2}\right) .$$

Below we analytically show that the correction $\log R$ is zero, and hence that

$$R = \left\langle \left\langle e^{r_a(k_a)} \right\rangle_{k_a|y}^{1/2} \left\langle e^{r_b(k_b)} \right\rangle_{k_b|y}^{1/2} \right\rangle_y = \left\langle \sqrt{\mathcal{F}_a(y)} \sqrt{\mathcal{F}_b(y)} \right\rangle_y = 1 , \quad (40)$$

where $\mathcal{F}_a(y)$ is a shorthand for $\langle e^{r_a(k_a)} \rangle_{k_a|y}$ and

$$r_a(k_a) = \sum_{l \geq 3} i^l \frac{c_{al}}{l!} k_a^l , \quad r_b(k_b) = \sum_{l \geq 3} i^l \frac{c_{bl}}{l!} k_b^l .$$

Because $f_a = f_b$, the cumulants will be the same for all l , hence $c_{al} = c_{bl}$. Furthermore, $k_a|y$ and $k_b|y$ are both distributed according to the *same* density. Now define, using $e^{r_a} = 1 + r_a + \frac{1}{2} r_a^2 + \dots$,

$$\begin{aligned} \mathcal{F}_a(y) &= \left\langle 1 + \sum_{l \geq 3} i^l \frac{c_{al}}{l!} k_a^l + \frac{1}{2} \sum_{l,s \geq 3} i^{l+s} \frac{c_{al} c_{as}}{l! s!} k_a^{l+s} + \dots \right\rangle_{k_a|y} \\ &= \left\langle 1 + \sum_{l \geq 3} \frac{c_{al}}{l!} \left(\frac{1}{\sigma}\right)^l (y + iu)^l + \frac{1}{2} \sum_{l,s \geq 3} \frac{c_{al} c_{as}}{l! s!} \left(\frac{1}{\sigma}\right)^{l+s} (y + iu)^{l+s} + \dots \right\rangle_u \\ &= 1 + \sum_{l \geq 3} \frac{c_{al}}{l!} \left(\frac{1}{\sigma}\right)^l H_l(y) + \frac{1}{2} \sum_{l,s \geq 3} \frac{c_{al} c_{as}}{l! s!} \left(\frac{1}{\sigma}\right)^{l+s} H_{l+s}(y) + \dots \end{aligned} \quad (41)$$

In the second line above a transformation of variables was made in the integral, with $u = \sigma k_a + iy$, such that $k_a = (u - iy)/\sigma$. The Jacobian $1/\sigma$ ensures proper normalization so that the average is over $u \sim \mathcal{N}(u; 0, 1)$. In the last line $H_l(y)$ is the *Hermite polynomial* of degree l ,

$$\begin{aligned} H_0(y) &= 1, & H_1(y) &= y, & H_2(y) &= y^2 - 1, \\ H_3(y) &= y^3 - 3y, & H_4(y) &= y^4 - 6y^2 + 3, & H_5(y) &= y^5 - 10y^3 + 15y \quad \dots \end{aligned}$$

which can be obtained for any real y and integer $l = 0, 1, 2, \dots$ from the average $H_l(y) = \langle (y + iu)^l \rangle_u$ over $u \sim \mathcal{N}(u; 0, 1)$.⁷

The remarkable property $\langle H_l(y) \rangle_y = 0$ for all l , ensures that $\langle \mathcal{F}_a(y) \rangle_y = 1$ in Equation (41). Furthermore, $\mathcal{F}_a(y) = \mathcal{F}_b(y)$ follows from the equivalence in cumulants $c_{al} = c_{bl}$; the roots in Equation (40) disappear to give $\langle \mathcal{F}_a(y) \rangle_y$, proving that $R = 1$ in Equation (40).

C.2 Second Order Correction

The second order expansion in Equation (32) in Section 7 evaluates to zero, as the matching cumulants $c_{al} = c_{bl}$ and equal distributions of $k_a|x$ and $k_b|x$ ensure that $\langle r_a(k_a) \rangle_{k_a|x} = \langle r_b(k_b) \rangle_{k_b|x}$:

$$\begin{aligned} \log R &= \frac{1}{4} \left\langle \langle r_a(k_a) \rangle_{k_a|x} \langle r_b(k_b) \rangle_{k_b|x} \right\rangle_x - \frac{1}{8} \left(\left\langle \langle r_a(k_a) \rangle_{k_a|x}^2 \right\rangle_x + \left\langle \langle r_b(k_b) \rangle_{k_b|x}^2 \right\rangle_x \right) + \dots \\ &= \frac{1}{4} \left\langle \langle r_a(k_a) \rangle_{k_a|x}^2 \right\rangle_x - \frac{1}{8} \left(2 \left\langle \langle r_a(k_a) \rangle_{k_a|x}^2 \right\rangle_x \right) + \dots \\ &= 0 + \dots \end{aligned}$$

Appendix D. Corrections to Marginals Distributions

Corrections to the marginal distributions follow from a similar derivation to that of the normalizing constant. As a simplification, let the Gaussian approximation be centred with $\mathbf{y} = \mathbf{x} - \boldsymbol{\mu}$, so that $q(\mathbf{y}) = \mathcal{N}(\mathbf{y}; \mathbf{0}, \boldsymbol{\Sigma})$, and assume that $q(\mathbf{x})$ arises from the fully factorized approximation in Section 5. In this appendix corrections will be computed for the mean $\langle x_i - \mu_i \rangle_{p(\mathbf{x})} = \langle y_i \rangle_{p(\mathbf{y})}$, and variance $\langle (x_i - \mu_i)(x_j - \mu_j) - \Sigma_{ij} \rangle_{p(\mathbf{x})} = \langle y_i y_j \rangle_{p(\mathbf{y})} - \Sigma_{ij}$.

A further simplification that will be employed in the following section is a change of variables $\boldsymbol{\eta}_n = k_n + i\Sigma_{nn}^{-1}y_n$, so that $\boldsymbol{\eta}_n \sim \mathcal{N}(\boldsymbol{\eta}_n; \mathbf{0}, \Sigma_{nn}^{-1})$. Let

$$z_n = \boldsymbol{\eta}_n - i\Sigma_{nn}^{-1}y_n,$$

which is zero-mean complex Gaussian random variable with a relation $\langle z_n^2 \rangle = 0$ and $\langle z_m z_n \rangle = -\Sigma_{mn}/(\Sigma_{mm}\Sigma_{nn})$ when $m \neq n$. Following Equation (24), the correction reads

$$R = \left\langle \prod_n \langle r_n(k_n) \rangle_{k_n|y_n} \right\rangle_{\mathbf{y}} = \left\langle \prod_n \langle r_n(\boldsymbol{\eta}_n - i\Sigma_{nn}^{-1}y_n) \rangle_{\boldsymbol{\eta}_n} \right\rangle_{\mathbf{y}} = \left\langle \exp \left[\sum_n r_n(z_n) \right] \right\rangle_{\mathbf{z}}.$$

7. When $\mathcal{F}(y)$ in Equation (41) is rearranged as a power series in σ^l , we obtain an Edgeworth expansion to arbitrary order l . The deviation from the Gaussian $q(y)$ is thereby factorized out of tilted distribution with $q_a(y) = q(y)\mathcal{F}(y)$. The interested reader is pointed to Blinnikov and Moessner (1998).

D.1 The Marginal Mean

The lowest order correction to the EP marginal’s mean follows from the result in Equation (13):

$$\begin{aligned} \langle y_i \rangle_{p(\mathbf{y})} &= \frac{1}{R} \left\langle y_i e^{\sum_n r_n(z_n)} \right\rangle_{\mathbf{z}} \\ &= \frac{1}{R} \sum_j \Sigma_{ij} \left\langle \frac{\partial}{\partial y_j} e^{\sum_n r_n(z_n)} \right\rangle \\ &= \frac{1}{R} \sum_j \Sigma_{ij} \left\langle \frac{\partial}{\partial y_j} \left(1 + \sum_n r_n(z_n) + \frac{1}{2} \sum_{m,n} r_m(z_m) r_n(z_n) + \dots \right) \right\rangle \\ &= \frac{1}{R} \sum_j \Sigma_{ij} \left\langle \frac{\partial r_j(z_j)}{\partial y_j} + \sum_n r_n(z_n) \frac{\partial r_j(z_j)}{\partial y_j} + \dots \right\rangle. \end{aligned}$$

In the above expansion the first order term is $\frac{\partial r_j(z_j)}{\partial y_j} = \frac{\partial r_j(z_j)}{\partial z_j} \frac{\partial z_j}{\partial y_j} = -i \Sigma_{jj}^{-1} \frac{\partial r_j(z_j)}{\partial z_j}$, and disappears as $\left\langle \frac{\partial r_j(z_j)}{\partial z_j} \right\rangle = 0$. The $j = n$ second order term also disappears as $\left\langle r_j(z_j) \frac{\partial r_j(z_j)}{\partial z_j} \right\rangle = 0$. These equivalences can be seen by taking $r_j(z_j)$ (and also its derivative) as a expansion over powers of z_j ; as $\langle z_j^2 \rangle = 0$, Wick’s theorem states that every expectation of powers of z_j should be zero. Hence

$$\langle y_i \rangle_{p(\mathbf{y})} = -\frac{i}{R} \sum_{j \neq n} \frac{\Sigma_{ij}}{\Sigma_{jj}} \left\langle r_n(z_n) \frac{\partial r_j(z_j)}{\partial z_j} \right\rangle_{\mathbf{z}} + \dots \tag{42}$$

The derivative of the characteristic function, as required in Equation (42), is

$$\frac{\partial r_j(z_j)}{\partial z_j} = \frac{\partial}{\partial z_j} \left[\sum_{l \geq 3} i^l \frac{c_{lj}}{l!} z_j^l \right] = i \sum_{l \geq 3} i^{l-1} \frac{c_{lj}}{(l-1)!} z_j^{l-1} = i \sum_{l \geq 2} i^l \frac{c_{l+1,j}}{l!} z_j^l.$$

The expectations for $j \neq n$ in Equation (42) evaluate to

$$\begin{aligned} \left\langle r_n(z_n) \frac{\partial r_j(z_j)}{\partial z_j} \right\rangle_{\mathbf{z}} &= i \sum_{s,l \geq 3} i^{s+l} \frac{c_{l+1,j} c_{sn}}{l! s!} \left\langle z_j^l z_n^s \right\rangle + i \sum_{s \geq 3, l=2} i^{s+l} \frac{c_{l+1,j}}{2! s!} \left\langle z_j^l z_n^s \right\rangle \\ &= i \sum_{l \geq 3} i^{2l} \frac{c_{l+1,j} c_{ln}}{(l!)^2} \left\langle z_j^l z_n^l \right\rangle, \end{aligned} \tag{43}$$

with the second term disappearing as $s > l = 2$ ensures that some z_n is always self-paired in Wick’s theorem. Finally, by substituting Equation (43) into (42), the correction to the mean is

$$\langle y_i \rangle_{p(\mathbf{y})} = \sum_{l \geq 3} \sum_{j \neq n} \frac{\Sigma_{ij}}{\Sigma_{jj}} \frac{c_{l+1,j} c_{ln}}{l!} \left(\frac{\Sigma_{jn}}{\Sigma_{jj} \Sigma_{nn}} \right)^l \pm \dots$$

D.2 The Marginal Covariance

The correction to the second moments follow the same recipe as that of the marginal mean in Appendix D.1. We proceed by first treating y_i with

$$\begin{aligned} \langle y_i y_j \rangle_{p(\mathbf{y})} &= \frac{1}{R} \left\langle y_i \left\{ y_j e^{\sum_n r_n(z_n)} \right\} \right\rangle_{\mathbf{z}} \\ &= \frac{1}{R} \sum_k \Sigma_{ik} \left\langle \frac{\partial}{\partial y_k} \left\{ y_j e^{\sum_n r_n(z_n)} \right\} \right\rangle \\ &= \frac{1}{R} \sum_k \Sigma_{ik} \left\langle \delta_{jk} e^{\sum_n r_n(z_n)} + y_j \frac{\partial}{\partial y_k} e^{\sum_n r_n(z_n)} \right\rangle \\ &= \Sigma_{ij} + \frac{1}{R} \sum_k \Sigma_{ik} \left\langle y_j \frac{\partial}{\partial y_k} e^{\sum_n r_n(z_n)} \right\rangle. \end{aligned}$$

Reapplying the recipe gives the correction to the covariance:

$$\begin{aligned} \langle y_i y_j \rangle_{p(\mathbf{y})} - \Sigma_{ij} &= \frac{1}{R} \sum_{kl} \Sigma_{il} \Sigma_{jk} \left\langle \frac{\partial^2}{\partial y_k \partial y_l} e^{\sum_n r_n(z_n)} \right\rangle_{\mathbf{z}} \\ &= -i \sum_{kl} \frac{\Sigma_{il}}{\Sigma_{ll}} \Sigma_{jk} \left\langle \frac{\partial}{\partial y_k} \frac{\partial r_l(z_l)}{\partial z_l} e^{\sum_n r_n(z_n)} \right\rangle + \dots \\ &= - \sum_{kl} \frac{\Sigma_{il}}{\Sigma_{ll}} \frac{\Sigma_{jk}}{\Sigma_{kk}} \left\langle \left[\delta_{kl} \frac{\partial^2 r_l(z_l)}{\partial z_l^2} + \frac{\partial r_k(z_k)}{\partial z_k} \frac{\partial r_l(z_l)}{\partial z_l} \right] e^{\sum_n r_n(z_n)} \right\rangle \\ &= \sum_{s \geq 3} \sum_{k \neq l} \frac{\Sigma_{il} \Sigma_{jl}}{\Sigma_{ll}^2} \frac{c_{sk} c_{s+2,l}}{s!} \left(\frac{\Sigma_{kl}}{\Sigma_{kk} \Sigma_{ll}} \right)^s \\ &\quad + \sum_{s \geq 3} \sum_{k \neq l} \frac{\Sigma_{il}}{\Sigma_{ll}} \frac{\Sigma_{jk}}{\Sigma_{kk}} \frac{c_{sk} c_{sl}}{s!} \left(\frac{\Sigma_{kl}}{\Sigma_{kk} \Sigma_{ll}} \right)^{s-1} + \dots \end{aligned}$$

Appendix E. Higher Order Cumulants

Much of this paper hinges on cumulants beyond the second order. These are frequently more cumbersome to obtain than the initial moments that are required by EP. This appendix provides details of the cumulants used in this paper.

The cumulants of a distribution $q_n(x)$ can be obtained from its moments through

$$\begin{aligned} c_3 &= \langle x^3 \rangle - 3 \langle x^2 \rangle \langle x \rangle + 2 \langle x \rangle^3, \\ c_4 &= \langle x^4 \rangle - 4 \langle x^3 \rangle \langle x \rangle - 3 \langle x^2 \rangle^2 + 12 \langle x^2 \rangle \langle x \rangle^2 - 6 \langle x \rangle^4, \\ c_5 &= \langle x^5 \rangle - 5 \langle x^4 \rangle \langle x \rangle - 10 \langle x^3 \rangle \langle x^2 \rangle + 20 \langle x^3 \rangle \langle x \rangle^2 + 30 \langle x^2 \rangle^2 \langle x \rangle - 60 \langle x^2 \rangle \langle x \rangle^3 + 24 \langle x \rangle^5; \end{aligned}$$

they are derived for doubly-truncated Gaussian distributions in Appendices E.1 and E.2. One might also directly take derivatives of the cumulant generating function, and the cumulants of a Probit-times-Gaussian distribution, common to GP classification models, are derived this way in Appendix E.3.

The tree-structured approximation in Sections 7 and 9.1, and Appendices A.1.3 and B, require cumulants over two variables. They are presented in Appendix E.4 for the Ising model.

E.1 Doubly Truncated Centered Gaussian

Consider the centered distribution $q_n(x_n) \propto \mathbb{I}[|x_n| < a] \mathcal{N}(x_n; 0, \lambda_n^{-1})$. The odd moments of this tilted distributions are, by symmetry, $\langle x_n \rangle = \langle x_n^3 \rangle = \langle x_n^5 \rangle = 0$. Let

$$Z_n = 2\sqrt{\frac{\lambda}{2\pi}} \int_0^a e^{-\frac{1}{2}\lambda x^2} dx = 2\Phi(z) - 1, \quad z = \sqrt{\lambda}a,$$

with the Probit function being $\Phi(x) = \int_{-\infty}^x \mathcal{N}(z; 0, 1) dz$. Subscripts n are dropped where they are clearly implied by their context. To get the even moments, consider

$$A_1 = \partial_\lambda \log Z_n = \partial_\lambda \log \left(\sqrt{\lambda} \int_{-a}^a dx e^{-\frac{1}{2}\lambda x^2} \right) = \frac{1}{2\lambda} - \frac{1}{2} \langle x^2 \rangle,$$

$$A_2 = \partial_\lambda^2 \log Z_n = -\frac{1}{2\lambda^2} + \frac{1}{4} \left(\langle x^4 \rangle - \langle x^2 \rangle^2 \right).$$

Using the partition function, we get

$$A_1 = \partial_\lambda \log (2\Phi(z) - 1) = \frac{a}{\sqrt{\lambda}} \left(\frac{\mathcal{N}(z)}{2\Phi(z) - 1} \right),$$

$$A_2 = \frac{a^2}{2\lambda} \left(\frac{z\mathcal{N}(z)}{2\Phi(z) - 1} \right) - \frac{a}{2\lambda^{3/2}} \left(\frac{\mathcal{N}(z)}{2\Phi(z) - 1} \right) - \frac{a^2}{\lambda} \left(\frac{\mathcal{N}(z)}{2\Phi(z) - 1} \right)^2,$$

and thus

$$\langle x^2 \rangle = \frac{1}{\lambda} - 2A_1,$$

$$\langle x^4 \rangle = \frac{2}{\lambda^2} + \langle x^2 \rangle^2 + 4A_2.$$

We can further determine $A_3 = \partial_\lambda^3 \log Z_n$ using the partition function, giving

$$A_3 = \frac{3a}{4\lambda^{5/2}} \left(\frac{\mathcal{N}(z)}{2\Phi(z) - 1} \right) + \frac{3a^2}{4\lambda^2} \left(\frac{z\mathcal{N}(z)}{2\Phi(z) - 1} \right) + \frac{3a^2}{2\lambda^2} \left(\frac{\mathcal{N}(z)}{2\Phi(z) - 1} \right)^2 + \frac{2a^3}{\lambda^{3/2}} \left(\frac{\mathcal{N}(z)}{2\Phi(z) - 1} \right)^3$$

$$+ \frac{3a^3}{2\lambda^{3/2}} \left(\frac{\mathcal{N}(z)}{2\Phi(z) - 1} \right) \left(\frac{z\mathcal{N}(z)}{2\Phi(z) - 1} \right) + \frac{a^3}{4\lambda^{3/2}} \left(\frac{(z^2 - 1)\mathcal{N}(z)}{2\Phi(z) - 1} \right).$$

Therefore

$$\langle x^6 \rangle = \frac{8}{\lambda^3} + \langle x^2 \rangle \langle x^4 \rangle - 2\langle x^2 \rangle^3 + 2\langle x^4 \rangle \langle x^2 \rangle - 8A_3.$$

E.2 Doubly Truncated Non-Centered Gaussian

The same calculation from Appendix E.1 can be repeated to get the moments of the non-centered truncated Gaussian $q_n(x_n) \propto \mathbb{I}[|x_n| < a] \mathcal{N}(x_n; \mu, \lambda_n^{-1})$. The subscripts n are dropped where evident. The partition function is

$$Z(\lambda, \mu) = \sqrt{\frac{\lambda}{2\pi}} \int_{-a}^a e^{-\frac{1}{2}\lambda(x-\mu)^2} dx = \Phi(z_{\max}) - \Phi(z_{\min}),$$

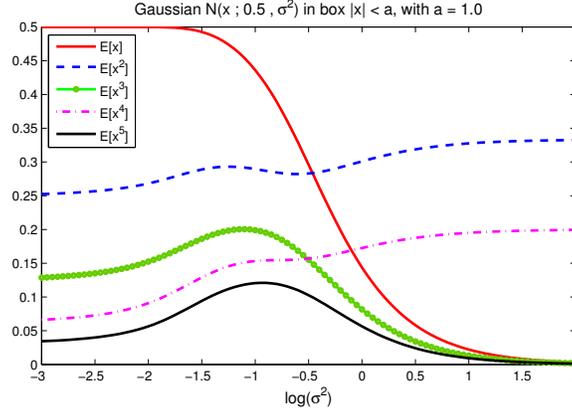


Figure 9: The moments of $q_n(x) \propto \mathbb{I}[|x| < a] \mathcal{N}(x; \mu, \sigma^2)$, as a function of σ^2 . As the Gaussian variance $\sigma^2 \rightarrow \infty$, the moments converge to that of a uniform $\mathcal{U}[-a, a]$ distribution.

where

$$z_{\max} = \sqrt{\lambda}(\mu + a), \quad z_{\min} = \sqrt{\lambda}(\mu - a).$$

By again taking increasing derivatives of $Z(\lambda, \mu)$ with respect to μ and λ , the moments solved for are

$$\begin{aligned} \langle x \rangle &= \mu + \frac{1}{\sqrt{\lambda}} \frac{\mathcal{N}(z_{\max}) - \mathcal{N}(z_{\min})}{\Phi(z_{\max}) - \Phi(z_{\min})}, \\ \langle x^2 \rangle &= 2 \langle x \rangle \mu + \frac{1}{\lambda} - \mu^2 - \frac{1}{\lambda} \frac{z_{\max} \mathcal{N}(z_{\max}) - z_{\min} \mathcal{N}(z_{\min})}{\Phi(z_{\max}) - \Phi(z_{\min})}, \\ \langle x^3 \rangle &= 3 \langle x^2 \rangle \mu + \langle x \rangle \left[\frac{3}{\lambda} - 3\mu^2 \right] - \frac{3}{\lambda} \mu + \mu^3, \\ &\quad - \frac{1}{\lambda^{3/2}} \frac{(1 - z_{\max}^2) \mathcal{N}(z_{\max}) - (1 - z_{\min}^2) \mathcal{N}(z_{\min})}{\Phi(z_{\max}) - \Phi(z_{\min})}, \\ \langle x^4 \rangle &= 4 \langle x^3 \rangle \mu + \langle x^2 \rangle \left[\frac{2}{\lambda} - 6\mu^2 \right] + \langle x \rangle \left[4\mu^3 - \frac{4}{\lambda} \mu \right] + \frac{2}{\lambda} \mu^2 - \mu^4 + \frac{1}{\lambda^2} \\ &\quad - \frac{1}{\lambda^2} \frac{z_{\max}(1 + z_{\max}^2) \mathcal{N}(z_{\max}) - z_{\min}(1 + z_{\min}^2) \mathcal{N}(z_{\min})}{\Phi(z_{\max}) - \Phi(z_{\min})}. \end{aligned}$$

Finally,

$$\begin{aligned} \langle x^5 \rangle &= 5 \langle x^4 \rangle \mu + \langle x^3 \rangle \left[\frac{6}{\lambda} - 10\mu^2 \right] + \langle x^2 \rangle \left[10\mu^3 - \frac{18}{\lambda} \mu \right] + \langle x \rangle \left[\frac{18}{\lambda} \mu^2 - 5\mu^4 - \frac{3}{\lambda^2} \right] + \frac{3}{\lambda^2} \mu \\ &\quad - \frac{6}{\lambda} \mu^3 + \mu^5 - \frac{1}{\lambda^{5/2}} \frac{(1 + 2z_{\max}^2 - z_{\max}^4) \mathcal{N}(z_{\max}) - (1 + 2z_{\min}^2 - z_{\min}^4) \mathcal{N}(z_{\min})}{\Phi(z_{\max}) - \Phi(z_{\min})}. \end{aligned}$$

As Figure 9 illustrates, these moments will converge to that of a uniform distribution as the Gaussian's variance grows large.

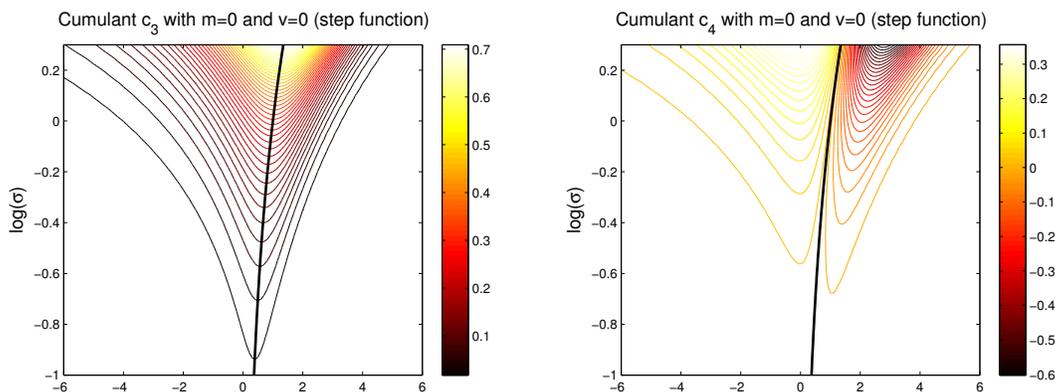


Figure 10: The third and fourth cumulants of the density $q_n(x) \propto \Phi((x - m)/v) \mathcal{N}(x; \mu, \sigma^2)$ in Appendix E.3. The step function $\Theta(x)$, with $m = v = 0$, is taken as an example here. The third cumulant is always positive, while the fourth cumulant is positive only when $\sigma > \mu$.

E.3 Probit Link Cumulants

EP approximations to Probit regression models, and Gaussian process classification models in general (see Section 8.1), depend on the moments of $q_n(x) \propto \Phi((x - m)/v) \mathcal{N}(x; \mu, \sigma^2)$. We introduce $v \geq 0$ so that the likelihood can become a step function at $v = 0$, for example. We shall obtain the cumulants by taking derivatives of the characteristic function. The characteristic function of $q_n(x)$, as described by Equation (15), is

$$\chi_n(k) = \langle e^{ikx} \rangle_{q_n(x)} = \exp \left\{ ik\mu - \frac{1}{2} k^2 \sigma^2 \right\} \frac{\Phi(z_k)}{\Phi(z)},$$

with

$$z = \frac{\mu - m}{\sqrt{v^2 + \sigma^2}}, \quad z_k = \frac{\mu + ik\sigma^2 - m}{\sqrt{v^2 + \sigma^2}}.$$

The cumulants c_{ln} are determined from the derivatives of $\log \chi_n(k)$ at zero; a lengthy calculation shows that they are

$$\begin{aligned} c_{3n} &= \alpha^3 \beta [2\beta^2 + 3z\beta + z^2 - 1], \\ c_{4n} &= -\alpha^4 \beta [6\beta^3 + 12z\beta^2 + 7z^2\beta + z^3 - 4\beta - 3z], \end{aligned}$$

where $\alpha = \sigma^2 / \sqrt{v^2 + \sigma^2}$ and $\beta = \mathcal{N}(z; 0, 1) / \Phi(z)$.

E.4 Two-Variable Ising Model Cumulants

We need some third and fourth order two-variable cumulants and thus generalize the results of Section 4.2 to the bivariate case. To do this we can exploit the cumulant generating property of $\log \chi_a(\mathbf{k}_a)$. Let $c_{(l,l')}$ denote the joint l, l' order cumulant of variable one and two, respectively. We can generate this cumulant from derivatives of $\log \chi_a(\mathbf{k}_a)$:

$$c_{(l,l')} = \left(\frac{\partial}{\partial ik_1} \right)^l \left(\frac{\partial}{\partial ik_2} \right)^{l'} \log \chi_a(\mathbf{k}_a) \Big|_{\mathbf{k}=\mathbf{0}}.$$

We can also express this as a recursion in terms of cumulants:

$$c_{(l+n, l'+n')} = \left(\frac{\partial}{\partial i k_1} \right)^n \left(\frac{\partial}{\partial i k_2} \right)^{n'} c_{(l, l')}(\mathbf{k}) \Big|_{\mathbf{k}=\mathbf{0}} .$$

By explicit calculation for a bivariate binary distribution we get the first two orders' cumulants: $c_{(1,0)} = m_1$, $c_{(0,1)} = m_2$, $c_{(2,0)} = 1 - m_1^2$, $c_{(0,2)} = 1 - m_2^2$ and $c_{(1,1)}$ is equal to the covariance between the two variables (to be matched with $q(x)$). The fact that we can write $c_{(2,0)}$ in terms of the first order cumulant shows that we can express all order cumulants in terms of the first and second order cumulant for example:

$$c_{(2,1)} = \frac{\partial}{\partial i k_2} c_{(2,0)}(\mathbf{k}) \Big|_{\mathbf{k}=\mathbf{0}} = \frac{\partial}{\partial i k_2} (1 - c_{(1,0)}^2(\mathbf{k})) \Big|_{\mathbf{k}=\mathbf{0}} = -2c_{(1,0)}c_{(1,1)} .$$

Using the same recursion it is easy to show: $c_{(3,0)} = -2c_{(1,0)}c_{(2,0)}$, $c_{(4,0)} = -2c_{(2,0)}^2 - 2c_{(1,0)}c_{(3,0)}$, $c_{(3,1)} = -2c_{(2,0)}c_{(1,1)} - 2c_{(1,0)}c_{(2,1)}$ and $c_{(2,2)} = -2c_{(1,1)}^2 - 2c_{(1,0)}c_{(1,2)} = -2c_{(1,1)}^2 + 4c_{(1,0)}c_{(0,1)}c_{(1,1)}$.

References

- D. Barber. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.
- C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- S. Blinnikov and R. Moessner. Expansions for nearly Gaussian distributions. *Astronomy and Astrophysics Supplement Series*, 130:193–205, 1998.
- J. P. Boyd. The devil's invention: Asymptotic, superasymptotic and hyperasymptotic series. *Acta Applicandae Mathematicae*, 56:1–98, 1999.
- M. Chertkov and V. Y. Chernyak. Loop series for discrete statistical models on graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 2006:P06009, 2006.
- B. Cseke and T. Heskes. Approximate marginals in latent Gaussian models. *Journal of Machine Learning Research*, 12:417–457, 2011.
- M. Kuss and C. E. Rasmussen. Assessing approximate inference for binary Gaussian process classification. *Journal of Machine Learning Research*, 6:1679–1704, 2005.
- M. Mézard, G. Parisi, and M. A. Virasoro. *Spin Glass Theory and Beyond*, volume 9 of *Lecture Notes in Physics*. World Scientific, 1987.
- T. P. Minka. Expectation propagation for approximate Bayesian inference. In *UAI 2001*, pages 362–369, 2001a.
- T. P. Minka. *A family of algorithms for approximate Bayesian inference*. PhD thesis, MIT Media Lab, 2001b.
- T. P. Minka. Power EP. Technical Report MSR-TR-2004-149, Microsoft Research Ltd, 2004.

- T. P. Minka and Y. Qi. Tree-structured approximations by expectation propagation. In *Advances in Neural Information Processing Systems 16*. 2004.
- K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- M. Opper and O. Winther. Gaussian processes for classification: Mean field algorithms. *Neural Computation*, 12:2655–2684, 2000.
- M. Opper and O. Winther. Expectation consistent approximate inference. *Journal of Machine Learning Research*, 6:2177–2204, 2005.
- M. Opper, U. Paquet, and O. Winther. Improving on expectation propagation. In *Advances in Neural Information Processing Systems 21*, pages 1241–1248. 2009.
- U. Paquet, M. Opper, and O. Winther. Perturbation corrections in approximate inference: Mixture modelling applications. *Journal of Machine Learning Research*, 10:935–976, 2009.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2005.
- H. Rue, S. Martino, and N. Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(2):319–392, 2009.
- M. W. Seeger and H. Nickisch. Fast convergent algorithms for expectation propagation approximate Bayesian inference. *Arxiv preprint arXiv:1012.3584*, 2010.
- D. Sherrington and S. Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35(26):1792–1796, December 1975.
- E. Sudderth, M. Wainwright, and A. Willsky. Loop series and Bethe variational bounds in attractive graphical models. In *Advances in Neural Information Processing Systems 20*, pages 1425–1432. 2008.
- D. J. Thouless, P. W. Anderson, and R. G. Palmer. Solution of a ‘solvable model of a spin glass’. *Phil. Mag.*, 35:593, 1977.
- M. A. J. van Gerven, B. Cseke, F. P. de Lange, and T. Heskes. Efficient Bayesian multivariate fMRI analysis using a sparsifying spatio-temporal prior. *NeuroImage*, 50(1):150–161, 2010.
- M. J. Wainwright and M. I. Jordan. Log-determinant relaxation for approximate inference in discrete Markov random fields. *IEEE Transactions on Signal Processing*, 54(6):2099–2109, 2006.
- M. Welling, A. Gelfand, and A. Ihler. A cluster-cumulant expansion at the fixed points of belief propagation. In *Uncertainty in Artificial Intelligence (UAI)*. 2012.