Journal of Machine Learning Research X (2011) XX-XX Submitted 04/11; Published X/11

Gaussian Probabilities and Expectation Propagation

John P. Cunningham JPC74@CAM.AC.UK
Department of Engineering
University of Cambridge
Cambridge, UK

Philipp Hennig PHENNIG@TUEBINGEN.MPG.DE
Max Planck Institute for Intelligent Systems
Tübingen, Germany

Simon Lacoste-Julien SL522@CAM.AC.UK
Department of Engineering
University of Cambridge
Cambridge, UK

arXiv:1111.6832v2 [stat.ML] 28 Nov 2013

Editor: Editor Name

Abstract

While Gaussian probability *densities* are omnipresent in applied mathematics, Gaussian cumulative *probabilities* are hard to calculate in any but the univariate case. We study the utility of Expectation Propagation (EP) as an approximate integration method for this problem. For rectangular integration regions, the approximation is highly accurate. We also extend the derivations to the more general case of polyhedral integration regions. However, we find that in this polyhedral case, EP's answer, though often accurate, can be almost arbitrarily wrong. We consider these unexpected results empirically and theoretically, both for the problem of Gaussian probabilities and for EP more generally. These results elucidate an interesting and non-obvious feature of EP not yet studied in detail.

Keywords: Multivariate Normal Probabilities, Gaussian Probabilities, Expectation Propagation, Approximate Inference

1. Introduction

This paper studies approximations to definite integrals of Gaussian (also known as normal or central) probability distributions. We define the Gaussian distribution $p_0(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{m}, K)$ as

$$p_0(\mathbf{x}) = \frac{1}{(2\pi)^{\frac{n}{2}}|K|^{\frac{1}{2}}} \exp\left\{-\frac{1}{2}(\mathbf{x}-\mathbf{m})^T K^{-1}(\mathbf{x}-\mathbf{m})\right\}, \tag{1}$$

where $\mathbf{x} \in \mathbb{R}^n$ is a vector with n real valued elements, $\mathbf{m} \in \mathbb{R}^n$ is the mean vector, and $K \in \mathbb{R}^{n\times n}$ is the symmetric, positive semidefinite covariance matrix. The Gaussian is perhaps the most widely used distribution in science and engineering. Whether this popularity

©2011 John P. Cunningham, Philipp Hennig, and Simon Lacoste-Julien.

is reflective of a fundamental character (often afforded to this distribution because of its role in the Central Limit Theorem) is debatable, but at any rate, a battery of convenient analytic characteristics make it an indispensible tool for many applications. Multivariate Gaussians induce Gaussian marginals and conditionals on all linear subspaces of their domain, and are closed under linear transformations of their domain. Under Gaussian beliefs, expectations can be evaluated for a number of important function classes, such as linear functions, quadratic forms, exponentials of the former, trigonometric functions, and linear combinations of such functions. Gaussians form an exponential family (which also implies closure, up to normalisation, under multiplication and exponentiation, and the existence of an analytic conjugate prior—the normal-Wishart distribution). The distribution's moment generating function and characteristic function have closed-form. Historically, much importance has been assigned to the fact that the Gaussian maximises differential entropy for given mean and covariance. More recently, the fact that the distribution can be analytically extended to the infinite-dimensional limit—the Gaussian Process (Rasmussen and Williams, 2006)—has led to a wealth of new applications. Yet, despite all these great aspects of Gaussian *densities*, Gaussian *probabilities* are difficult to calculate: even the cumulative distribution function (cdf) has no closed-form expression and is numerically challenging in high-dimensional spaces. More generally, we consider the probability that a draw from $p(\mathbf{x})$ falls in a region $\mathcal{A} \subseteq \mathbb{R}^n$, which we will denote as

$$F(\mathcal{A}) = \text{Prob}\left\{\mathbf{x} \in \mathcal{A}\right\} = \int_{\mathcal{A}} p(\mathbf{x})\,\mathrm{d}\mathbf{x} = \int_{\ell_1(\mathbf{x})}^{u_1(\mathbf{x})} \cdots \int_{\ell_n(\mathbf{x})}^{u_n(\mathbf{x})} p(\mathbf{x})\,\mathrm{d}x_n \cdots \mathrm{d}x_1, \tag{2}$$

where $\ell_1, \ldots, \ell_n$ and $u_1, \ldots, u_n$ denote the upper and lower bounds of the region $\mathcal{A}$. In general, $\mathbf{u}(\mathbf{x})$ and $\mathbf{l}(\mathbf{x})$ may be nonlinear functions of $\mathbf{x}$ defining a valid region $\mathcal{A}$, but the methods presented here discuss the case where $\mathbf{u}$ and $\mathbf{l}$ are linear functions of $\mathbf{x}$, meaning that $\mathcal{A}$ forms a (possibly unbounded) polyhedron. The probability $F(\mathcal{A})$ generalises the cdf, which is recovered by setting the upper limits $\mathbf{u} = (u_i)_{i=1,\ldots,n}$ to a single point in $\mathbb{R}^n$ and the lower limits to $l_1 = \ldots = l_n = -\infty$. Applications of these multivariate Gaussian probabilities are widespread. They include statistics (Genz, 1992; Joe, 1995; Hothorn et al., 2005) (where a particularly important use case is probit analysis (Ashford and Sowden, 1970; Sickles and Taubman, 1986; Bock and Gibbons, 1996)), economics (Boyle et al., 2005), mathematics (Hickernell and Hong, 1999), biostatistics (Thiebaut and Jacqmin-Gadda, 2004; Zhao et al., 2005), medicine (Lesaffre and Molenberghs, 1991), environmental science (Buccianti et al., 1998), computer science (de Klerk et al., 2000), neuroscience (Pillow et al., 2004), machine learning (Liao et al., 2007), and more (see for example the applications listed in Gassman et al. (2002)). For the machine learning researcher, two popular problems that can be cast as Gaussian probabilities include the Bayes Point Machine (Herbrich, 2002) and Gaussian Process classification (Rasmussen and Williams, 2006; Kuss and Rasmussen, 2005), which will we discuss more specifically later in this work.

Univariate Gaussian probabilities can be so quickly and accurately calculated (*e.g.* Cody, 1969) that the univariate cumulative density function is available with machine-level precision in many statistical computing packages (*e.g.*, `normcdf` in MATLAB, `CDFNORM` in SPSS, `pnorm` in R, to name a few). Unfortunately, no similarly powerful algorithm exists for the multivariate case. There are some known analytic decompositions of Gaussian integrals into lower-dimensional forms (Plackett, 1954; Curnow and Dunnett, 1962; Lazard-Holly and

Holly, 2003), but such decompositions have very high computational complexity (Huguenin et al., 2009). In statistics, sampling based methods have found application in specific use cases (Lerman and Manski, 1981; McFadden, 1989; Pakes and Pollard, 1989), but the most efficient general, known method is numerical integration (Genz, 1992; Drezner and Wesolowsky, 1989; Drezner, 1994; Genz and Kwong, 2000; Genz and Brentz, 1999, 2002; Genz, 2004). A recent book (Genz and Bretz, 2009) gives a good overview. These algorithms make a series of transformations to the Gaussian $p_0(\mathbf{x})$ and the region $\mathcal{A}$, using the Cholesky factor of the covariance K, the univariate Gaussian cdf and its inverse, and randomly generated points. These methods aim to restate the calculation of $F(\mathcal{A})$ as a problem that can be well handled by quasi-random or lattice point numerical integration. Many important studies across many fields have been critically enabled by these and other algorithms (Joe, 1995; Hothorn et al., 2005; Boyle et al., 2005; Hickernell and Hong, 1999; Thiebaut and Jacqmin-Gadda, 2004; Zhao et al., 2005; Buccianti et al., 1998; de Klerk et al., 2000; Pillow et al., 2004; Liao et al., 2007). In particular, the methods of Genz represent the state of the art, and we will use this method as a basis for comparison. We note that the numerical accuracy of these methods is generally high and can be increased by investing additional computational time (more integration points), though by no means to the numerical precision of the univariate case. As such, achieving high accuracy invokes substantial computational cost. Also, applications such as Bayesian model selection require *analytic* approximations of $F(\mathcal{A})$, usually because the goal is to optimise $F(\mathcal{A})$ with respect to the parameters $\{\mathbf{m}, K\}$ of the Gaussian, which requires the corresponding derivatives. These derivatives and other features (to be discussed) are not currently offered by numerical integration methods. Thus, while there exist sophisticated methods to compute Gaussian probabilities with high accuracy, there remains significant work to be done to address this important problem.

In Section 2, we will develop an analytic approximation to $F(\mathcal{A})$ by using *Expectation Propagation (EP)* (Minka, 2001a,b, 2005; Opper and Winther, 2001) as an approximate integration method. We first give a brief introduction to EP (Section 2.1), then develop the necessary methods for integration over *hyperrectangular* regions $\mathcal{A}$ in Section 3.1, which is the case most frequently computed (and which includes the cdf). In Section 3.2, we then describe how to generalise these derivations to *polyhedral* $\mathcal{A}$. However, while these polyhedral extensions are conceptually elegant, it turns out that they do not always lead to an accurate algorithm. In Section 6, we compare EP's analytic approximations to numerical results. For rectangular $\mathcal{A}$, the approximations are of generally high quality. For polyhedral $\mathcal{A}$, they are often considerably worse. This shortcoming of EP will come as a surprise to many readers, because the differences to the rectangular case are inconspiciously straightforward. Indeed, we will demonstrate that hyperrectangles are in fact not a fundamental distinction, and we will use intuition gained from the polyhedral case to build pathological cases for hyperrectangular cases also. We study this interesting result in Section 7 and give some insight into this problem. Our overall empirical result is that EP provides reliable analytic approximations for rectangular Gaussian integrals, but should be used only with caution on more general regions. This has implications for some important applications, such as Bayesian generalised regression, where EP is often used.

This work bears connection to previous literature in machine learning. Most obviously, when considering the special case of Gaussian probabilities over hyperrectangular regions

(Section 3.1), our algorithm is derived directly from EP, yet distinct in two ways. First, conceptually, we do not use EP for approximate inference, but instead we put integration bounds in the place of likelihood terms, thereby using EP as a high dimensional integration scheme. Second, because of this conceptual change, we are dealing with unnormalised EP (as the "likelihood" factors are not distributions). From this special case, we extend the method to more general integration regions, which involves rank one EP updates. Rank one EP updates have been previously discussed, particularly in connection to the Bayes Point Machine (Minka, 2001b; Herbrich, 2002; Minka, 2008), but again their use for Gaussian probabilities has not been investigated.

Thus, our goal here is to build from EP and the general importance of the Gaussian probability problem, to offer three contributions: first, we give a full derivation that focuses EP on the approximate integration problem for multivariate Gaussian probabilities. Second, we perform detailed numerical experiments to benchmark our method and other methods, so that this work may serve as a useful reference for future researchers confronted with this ubiquitous and challenging computation. Third, we discuss empirically and theoretically some features of EP that have not been investigated in the literature, which has importance for EP well beyond Gaussian probabilities.

The remainder of this paper is laid out as follows. In Section 2, we discuss the EP algorithm in general and its specific application when the approximating distribution is Gaussian, as is our case of interest here. In Section 3, we apply EP specifically to Gaussian probabilities, describing the specifics and features of our algorithm: Expectation Propagation for Multivariate Gaussian Probabilities (EPMGP). In Section 4, we describe other existing methods for calculating Gaussian probabilities, as they will be compared to EPMGP in the results section. In Section 5, we describe the probabilities that form the basis of our numerical computations. Section 6 gives the results of these experiments, and Section 7 discusses the implications of this work and directions for future work.

2. Approximate Integration Through Approximate Moment Matching: Expectation Propagation

The key step to motivate this approach is to note that we can cast the Gaussian probability problem as one of integrating an intractable and unnormalised distribution. Defining $p(\mathbf{x})$ as

$$p(\mathbf{x}) = \begin{cases} p_0(\mathbf{x}) & \mathbf{x} \in \mathcal{A} \\ 0 & \text{otherwise,} \end{cases} \tag{3}$$

we see that the normaliser of this distribution is our probability of interest $F(\mathcal{A}) = \text{Prob}\left\{\mathbf{x} \in \mathcal{A}\right\} = \int_{\mathcal{A}} p_0(\mathbf{x})\,\mathrm{d}\mathbf{x} = \int p(\mathbf{x})\,\mathrm{d}\mathbf{x}$.

Expectation Propagation (Minka, 2001a,b, 2005; Opper and Winther, 2001) is a method for finding an approximate unnormalised distribution $q(\mathbf{x})$ to replace an intractable true unnormalised distribution $p(\mathbf{x})$. The EP objective function is involved, but EP is *motivated* by the idea of minimising Kullback-Leibler (KL) divergence (Cover and Thomas, 1991) from the true distribution to the approximation. We define this KL-divergence as

$$D_{\text{KL}}(p\|q) = \int p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})}\,\mathrm{d}\mathbf{x} + \int q(\mathbf{x})\,\mathrm{d}\mathbf{x} - \int p(\mathbf{x})\,\mathrm{d}\mathbf{x} \tag{4}$$

which is a form of the divergence that allows unnormalised distributions and reverts to the more popular form if p and q have the same indefinite integral.

KL-divergence is a natural measure for the quality of the approximation $q(\mathbf{x})$. If we choose $q(\mathbf{x})$ to be a high-dimensional Gaussian, then the choice of $q(\mathbf{x})$ that minimises $D_{\mathrm{KL}}(p \parallel q)$ is the $q(\mathbf{x})$ with the same zeroth, first, and second moments as $p(\mathbf{x})$. If we are trying to calculate $F(\mathcal{A})$, then we equivalently seek the zeroth moment of $p(\mathbf{x})$. As such, trying to minimise this global KL divergence is an appropriate and sensible method to calculate Gaussian probabilities. Appendix E reviews the equivalence between moment matching and minimising KL-divergence in this problem.

Unfortunately, minimising global KL-divergence directly is in many cases intractable. This fact motivated the creation of the EP algorithm, which seeks to approximately do this global KL minimisation by iteratively minimising the KL-divergence of local, single factors of $q(\mathbf{x})$ with respect to $p(\mathbf{x})$. Since $F(\mathcal{A})$ is the zeroth moment of $p(\mathbf{x})$ as defined in Equation (2), any algorithm trying to minimise this global KL-divergence objective, or an approximation thereof, is also a candidate method for the (approximate) calculation of Gaussian probabilities.

2.1 Expectation Propagation

The following sections review EP and introduce its use for Gaussian probability calculations. It is a prototypical aspect of Bayesian inference that an intractable distribution $p(\mathbf{x})$ is a product of a prior distribution $p_0(\mathbf{x})$ and one or more likelihood functions or factors $t_i(\mathbf{x})$:

$$p(\mathbf{x}) = p_0(\mathbf{x}) \prod_i t_i(\mathbf{x}). \tag{5}$$

Note that our Gaussian probability problem has this unnormalised form, as it can be written as:

$$F(\mathcal{A}) = \int_{\mathcal{A}} p_0(\mathbf{x}) \, \mathrm{d}\mathbf{x} = \int p(\mathbf{x}) \, \mathrm{d}\mathbf{x} = \int p_0(\mathbf{x}) \prod_i t_i(\mathbf{x}) \, \mathrm{d}\mathbf{x} \tag{6}$$

where $t_i(\mathbf{x})$ is an indicator function defined in a particular direction, namely a "box function":

$$t_i(\mathbf{x}) = \mathbb{I}\{l_i < \mathbf{c}_i^T \mathbf{x} < u_i\} = \begin{cases} 1 & l_i < \mathbf{c}_i^T \mathbf{x} < u_i \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

The above form is our intractable problem of interest[1], and we assume without loss of generality that the $\mathbf{c}_i$ have unit norm. An important clarification for the remainder of this work is to note that most of the distributions discussed will be unnormalised. By the previous construction, the factors $t_i(\mathbf{x})$ are unnormalised, and thus so is $p(\mathbf{x})$ (as will be the EP approximation $q(\mathbf{x})$). Though the steps we present will all be valid for unnormalised distributions, we make this note to call out this somewhat atypical problem setting.

1. This box function notation may seem awkward vs. the more conventional halfspace definitions of polyhedra, but we will make use of this definition. Further, halfspaces can be recovered by setting any $u_i = \infty$, so this definition is general.

The approach of the EP algorithm is to replace each intractable $t_i(\mathbf{x})$ with a tractable unnormalised Gaussian $\tilde{t}_i(\mathbf{x})$. Our specific algorithm thus yields a nice geometric interpretation. We want to integrate a Gaussian over a polyhedron defined by several box functions, but this operation is intractable. Instead, EP allows us to replace each of those intractable box truncation functions with soft Gaussian truncations $\tilde{t}_i(\mathbf{x})$. Then, since we know that these exponential family distributions are simple to multiply (multiplication essentially amounts to summing the natural parameters), our problem reduces to finding $\tilde{t}_i(\mathbf{x})$.

This perspective motivates the EP algorithm, which approximately minimises $D_{\mathrm{KL}}(p\|q)$ by iteratively constructing an unnormalised approximation $q(\mathbf{x})$. At any point in the iteration, EP tracks a current approximate $q(\mathbf{x})$ in an exponential family, and a set of *approximate factors* (or *messages*) $\tilde{t}_i$, also in the family. The factors are updated by constructing a *cavity* distribution

$$q^{\backslash i}(\mathbf{x}) = \frac{q(\mathbf{x})}{\tilde{t}_i(\mathbf{x})} \tag{8}$$

(this division operation on unnormalised distributions is also well-defined for members of exponential families - it amounts to subtracting natural parameters), and then *projecting* into the exponential family

$$\tilde{t}_i(\mathbf{x})q^{\backslash i}(\mathbf{x}) = \mathrm{proj}[t_i(\mathbf{x})q^{\backslash i}(\mathbf{x})]. \tag{9}$$

where this projection operation (Minka (2004), an M-projection from information geometry (Koller and Friedman, 2009)) is defined as setting $\tilde{t}_i(\mathbf{x})$ to the unnormalised member of the exponential family minimizing $D_{\mathrm{KL}}(t_i q^{\backslash i}\|\tilde{t}_i q^{\backslash i})$. Intuitively, to update an approximate factor $\tilde{t}_i(\mathbf{x})$, we first remove its effect from the current approximation (forming the cavity), and then we include the effect of the true factor $t_i(\mathbf{x})$ by an M-projection, updating $\tilde{t}_i(\mathbf{x})$ accordingly. As derived in Appendix E, this projection means matching the sufficient statistics of $\tilde{t}_i(\mathbf{x})q^{\backslash i}(\mathbf{x})$ to those of $t_i(\mathbf{x})q^{\backslash i}(\mathbf{x})$. In particular for Gaussian $\tilde{t}_i(\mathbf{x})$, matching sufficient statistics is equivalent to matching *zeroth*, first and second moments. We now restrict this general prescription of EP to the Gaussian case, as it will clarify the simplicity of our Gaussian probability algorithm.

2.1.1 Gaussian EP with rank-one factors

In cases where the approximating family is Gaussian, the EP derivations can be specified further without having to make additional assumptions about the true $t_i(\mathbf{x})$. Because it will lend clarity to Sections 3.1 and 3.2, we give that analysis here. To orient the reader, it is important to be explicit about the numerous uses of "Gaussian" here. When we say "Gaussian EP," we mean *any* EP algorithm where the approximating $q(\mathbf{x})$ (the prior $p_0(\mathbf{x})$ and approximate factors $\tilde{t}_i(\mathbf{x})$) are Gaussian. The true factors $t_i(\mathbf{x})$ can be arbitrary. In particular, here we discuss Gaussian EP, and only in Section 3 do we discuss Gaussian EP as used for multivariate Gaussian probability calculations (with box function factors $t_i(\mathbf{x})$). It is critical to note that all steps in this next section are invariant to changes in the form of the true $t_i(\mathbf{x})$, so what follows is a general description of Gaussian EP. The unnormalised Gaussian approximation is

$$q(\mathbf{x}) = p_0(\mathbf{x}) \prod_i \tilde{t}_i(\mathbf{x}) = p_0(\mathbf{x}) \prod_i \tilde{Z}_i \mathcal{N}(\mathbf{x}; \tilde{\mu}_i, \tilde{\sigma}_i^2) = Z\mathcal{N}(\mathbf{x}; \mu, \Sigma). \tag{10}$$

We will require the cavity distribution for the derivation of the updates, which is defined as

$$q^{\backslash i}(\mathbf{x}) = \frac{q(\mathbf{x})}{\tilde{t}_i(\mathbf{x})} = Z^{\backslash i}\mathcal{N}(\mathbf{x}; \mathbf{u}^{\backslash i}, V^{\backslash i}) \tag{11}$$

The above step is the *cavity* step of EP (Equation 8). We now must do the projection operation of Equation 9, which involves moment matching the approximation $\tilde{t}_i(\mathbf{x})q^{\backslash i}(\mathbf{x})$ to the appropriate moments of $t_i(\mathbf{x})q^{\backslash i}(\mathbf{x})$:

$$\begin{aligned}
\hat{Z}_i &= \int t_i(\mathbf{x}) q^{\backslash i}(\mathbf{x}) \, \mathrm{d}\mathbf{x} && (12)\\
\hat{\mathbf{u}}_i &= \frac{1}{\hat{Z}_i} \int \mathbf{x} t_i(\mathbf{x}) q^{\backslash i}(\mathbf{x}) \, \mathrm{d}\mathbf{x} && (13)\\
\hat{V}_i &= \frac{1}{\hat{Z}_i} \int (\mathbf{x} - \hat{\mathbf{u}}_i)(\mathbf{x} - \hat{\mathbf{u}}_i)^T t_i(\mathbf{x}) q^{\backslash i}(\mathbf{x}) \, \mathrm{d}\mathbf{x}. && (14)
\end{aligned}$$

These updates are useful and tractable if the individual $t_i(\mathbf{x})$ have some simplifying structure such as low rank, as in Equation 7. In particular, as is the case of interest here, if the $t_i(\mathbf{x})$ have rank one structure, then these moments have simple form:

$$\begin{aligned}
\hat{Z}_i &= \int t_i(\mathbf{x}) q^{\backslash i}(\mathbf{x}) \, \mathrm{d}\mathbf{x} \\
\hat{\mathbf{u}}_i &= \hat{\mu}_i \mathbf{c}_i \\
\hat{V}_i &= \hat{\sigma}_i^2 \mathbf{c}_i \mathbf{c}_i^T, && (15)
\end{aligned}$$

where $\{\hat{Z}_i, \hat{\mu}_i, \hat{\sigma}_i^2\}$ depend on the factor $t_i(\mathbf{x})$ and $\mathbf{c}_i$ is the rank one direction as in Equation 7. The important thing to note is that these parameters $\{\hat{Z}_i, \hat{\mu}_i, \hat{\sigma}_i^2\}$ are the only place in which the actual form of the $t_i(\mathbf{x})$ enter in. This can also be derived from differential operators, as was done in a short technical report (Minka, 2008).

To complete the projection step, we must update the approximate factors $\tilde{t}_i(\mathbf{x})$ such that the moments of $\tilde{t}_i(\mathbf{x})q^{\backslash i}(\mathbf{x})$ match $\{\hat{Z}_i, \hat{\mu}_i, \hat{\sigma}_i^2\}$, as doing so is the KL minimisation over this particular factor $\tilde{t}_i(\mathbf{x})$. We derive the simple approximate factor updates using basic properties of the Gaussian (again just subtracting natural parameters). We have

$$\begin{aligned}
\tilde{t}_i(\mathbf{x}) &= \tilde{Z}_i \mathcal{N}(\mathbf{c}_i^T \mathbf{x}; \tilde{\mu}_i, \tilde{\sigma}_i^2), && (16)\\
\text{where} \quad \tilde{\mu}_i &= \tilde{\sigma}_i^2 (\hat{\sigma}_i^{-2} \hat{\mu}_i - \sigma_{\backslash i}^{-2} \mu_{\backslash i}), \qquad \tilde{\sigma}_i^2 = (\hat{\sigma}_i^{-2} - \sigma_{\backslash i}^{-2})^{-1}, \\
\tilde{Z}_i &= \hat{Z}_i \sqrt{2\pi} \sqrt{\sigma_{\backslash i}^2 + \tilde{\sigma}_i^2} \exp\Big\{ \frac{1}{2} (\mu_{\backslash i} - \tilde{\mu}_i)^2 / (\sigma_{\backslash i}^2 + \tilde{\sigma}_i^2) \Big\},
\end{aligned}$$

and where cavity parameters $\{\sigma_{\backslash i}^2, \mu_{\backslash i}\} = \{\mathbf{c}_i^T V^{\backslash i} \mathbf{c}_i \, , \, \mathbf{c}_i^T \mathbf{u}^{\backslash i}\}$ are fully derived in Appendix A as:

$$\mu_{\backslash i} = \sigma^2_{\backslash i}\Big(\frac{\mathbf{c}_i^T\boldsymbol{\mu}}{\mathbf{c}_i^T\Sigma\mathbf{c}_i} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\Big), \quad \text{and} \quad \sigma^2_{\backslash i} = \big((\mathbf{c}_i^T\Sigma\mathbf{c}_i)^{-1} - \tilde{\sigma}_i^{-2}\big)^{-1}. \tag{17}$$

These equations are all standard manipulations of Gaussian random variables. Now, by the definition of the approximation, we can calculate the new approximation $q(\mathbf{x})$ as the product $q(\mathbf{x}) = p_0(\mathbf{x})\prod_i \tilde{t}_i(\mathbf{x})$:

$$\begin{aligned} q(\mathbf{x}) &= Z\mathcal{N}(\boldsymbol{\mu}, \Sigma), \\ \text{where} \quad \boldsymbol{\mu} &= \Sigma\Big(K^{-1}\mathbf{m} + \sum_{i=1}^m \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\mathbf{c}_i\Big), \qquad \Sigma = \Big(K^{-1} + \sum_{i=1}^m \frac{1}{\tilde{\sigma}_i^2}\mathbf{c}_i\mathbf{c}_i^T\Big)^{-1}, \end{aligned} \tag{18}$$

where $p_0(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{m}, K)$ as in Equation 1. These forms can be efficiently and numerically stably calculated in quadratic computation time via steps detailed in Appendix B. Lastly, once the algorithm has converged, we can calculate the normalisation constant of $q(\mathbf{x})$. While this step is again general to Gaussian EP, we highlight it because it is our quantity of interest - the approximation of the probability $F(\mathcal{A})$:

$$\begin{aligned} \log Z &= -\frac{1}{2}\left(\mathbf{m}^T K^{-1}\mathbf{m} + \log|K|\right) \\ &\quad + \sum_{i=1}^m \left(\log \tilde{Z}_i - \frac{1}{2}\left(\frac{\tilde{\mu}_i^2}{\tilde{\sigma}_i^2} + \log\tilde{\sigma}_i^2 + \log(2\pi)\right)\right) \\ &\quad + \frac{1}{2}\left(\boldsymbol{\mu}^T\Sigma^{-1}\boldsymbol{\mu} + \log|\Sigma|\right) \end{aligned} \tag{19}$$

In the above we have broken this equation up into three lines to clarify that the normalisation term $\log Z$ has contribution from the prior $p_0(\mathbf{x})$ (first line), the approximate factors (the second line), and the full approximation $q(\mathbf{x})$ (third line). An intuitive interpretation is that the integral of $q(\mathbf{x})$ is equal to the product of all the contributing factor masses (the prior and the approximate factors $\tilde{t}_i$), divided by what is already counted by the normalisation constant of the final approximation $q(\mathbf{x})$. Similar reparameterisations of these calculations help to ensure numerical stability and quicker, more accurate computation. We derive that reparameterisation in the second part of Appendix B. Taking stock, we have derived Gaussian EP in a way that requires only rank one forms. These factors may be of any functional form, but in the following section we show the box function form of the Gaussian probability problem leads to a particularly fast and accurate computation of Gaussian EP.

At this point it may be helpful to clarify again that EP has two levels of approximation. First, being an approximate inference framework, EP would ideally choose an approximating $q(\mathbf{x})$ from a tractable family (such as the Gaussian) such that $D_{\mathrm{KL}}(p\|q)$ is minimised. That global problem being intractable, EP makes a second level of approximation in not actually minimising global $D_{\mathrm{KL}}(p\|q)$, but rather minimising the local $D_{\mathrm{KL}}(t_i q^{\backslash i}\|\tilde{t}_i q^{\backslash i})$. Because there is no obvious connection between these global and local divergences, it is difficult to make quantitative statements about the quality of the approximation provided by EP

without empirical evaluation. However, EP has been shown to provide excellent approximations in many applications (Minka, 2001a; Kuss and Rasmussen, 2005; Herbrich et al., 2007; Stern et al., 2009), often outperforming competing methods. These successes have deservedly established EP as a useful tool for applied inference, but they may also mislead the applied researcher into trusting its results more than is warranted by theoretical underpinnings. This paper, aside from developing a tool for approximate Gaussian integration, also serves to demonstrate a non-obvious way in which EP can yield poor results. The following two sections introduce our use of EP for approximate Gaussian integration that performs very well (Sections 3.1 and 6), followed by an ostensibly straightforward extension (Section 3.2), which turns out not to perform particularly well (Section 6), for a subtle reason (Section 7).

3. EP Multivariate Gaussian Probability (EPMGP) Algorithm

This section introduces the EPMGP algorithm, which uses Expectation Propagation to calculate multivariate Gaussian probabilities. We consider two cases: axis-aligned hyperrectangular integration regions $\mathcal{A}$ and more general poyhedral integration regions $\mathcal{A}$. Because of our general Gaussian EP derivation above, for each case we need only calculate the parameters that are specific to the choices of factor $t_i(\mathbf{x})$, namely $\{\hat{Z}_i, \hat{\mu}_i, \hat{\sigma}_i^2\}$ from Equation 15. This approach simplifies the presentation of this section and clarifies what is specific to Gaussian EP and what is specific to the Gaussian probability problem.

3.1 Rectangular Integration Regions

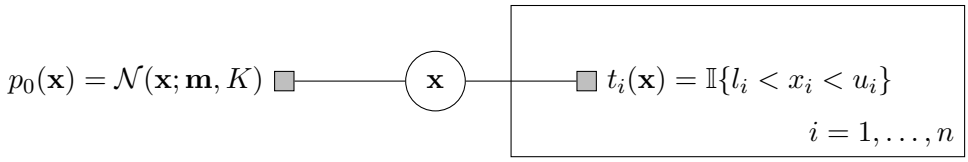

Figure 1: Graphical model (factor graph) for the rectangular Gaussian integration problem. In the language of Bayesian inference, the unconstrained Gaussian forms a "prior", the integration limits form "likelihood" factors, which give a "posterior" distribution on $\mathbf{x}$, and we are interested in the normalisation constant of that distribution. Note, however, that this inference terminology is used here without the usual philosophical interpretations of prior knowledge and subsequently aquired data. Note also that in the rectangular integration case there are exactly n such factors, corresponding to each dimension of the problem.

This section derives an EP algorithm for Gaussian integration over rectangular regions. Without loss of generality, we will assume that the rectangular region $\mathcal{A}$ is axis-aligned. Otherwise, there exists a unitary map U such that in the rotated coordinate system $\mathbf{x} \to U\mathbf{x}$ the region is axis aligned and the transformation morphs the Gaussian distribution $p_0(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{m}, K)$ into another Gaussian $p_0'(\mathbf{x}) = \mathcal{N}(U\mathbf{x}; U\mathbf{m}, UKU^T)$. Hence, the integration

boundaries are not functions of $\mathbf{x}$, and the integral $F(\mathcal{A})$ can be written using a product of n factors $t_i(x_i) = \mathbb{I}\{l_i < x_i < u_i\}$:

$$F(\mathcal{A}) = \int_{l_1}^{u_1} \cdots \int_{l_n}^{u_n} \mathcal{N}(\mathbf{x}; \mathbf{m}, K)\, \mathrm{d}\mathbf{x} = \int \mathcal{N}(\mathbf{x}; \mathbf{m}, K) \prod_{i=1}^{n} t_i(x_i)\, \mathrm{d}x_i. \tag{20}$$

EP's iterative steps thus correspond to the incorporation of individual box functions of upper and lower integration limits into the approximation, where the incorporation of a factor implies approximating the box function with a univariate Gaussian aligned in that direction. Thanks to the generality of the derivations in Section 2.1.1, the remaining derivations for this setting are conveniently short. The required moments for Equation 15 (the EP projection step) are calculated exactly (Jawitz, 2004) using the error function[2]:

$$\hat{Z}_i = \frac{1}{2}\Big(\mathrm{erf}(\beta) - \mathrm{erf}(\alpha)\Big), \tag{21}$$

$$\hat{\mu}_i = \mu_{\backslash i} + \frac{1}{\hat{Z}_i}\frac{\sigma_{\backslash i}}{\sqrt{2\pi}}\big(\exp\{-\alpha^2\} - \exp\{-\beta^2\}\big), \quad \text{and} \tag{22}$$

$$\hat{\sigma}_i^2 = \mu_{\backslash i}^2 + \sigma_{\backslash i}^2 + \frac{1}{\hat{Z}_i}\frac{\sigma_{\backslash i}}{\sqrt{2\pi}}\big((l_i + \mu_{\backslash i})\exp\{-\alpha^2\} - (u_i + \mu_{\backslash i})\exp\{-\beta^2\}\big) - \hat{\mu}_i^2, \tag{23}$$

where we have everywhere above used the shorthand $\alpha = \frac{l_i - \mu_{\backslash i}}{\sqrt{2}\sigma_{\backslash i}}$ and $\beta = \frac{u_i - \mu_{\backslash i}}{\sqrt{2}\sigma_{\backslash i}}$, and $\{\mu_{\backslash i}, \sigma_{\backslash i}^2\}$ are:

$$\mu_{\backslash i} = \sigma_{\backslash i}^2\Big(\frac{\mu_i}{\Sigma_{ii}} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\Big), \quad \text{and} \quad \sigma_{\backslash i}^2 = \big(\Sigma_{ii}^{-1} - \tilde{\sigma}_i^{-2}\big)^{-1}. \tag{24}$$

The above is simply the form of Equation 17 with the vectors $\mathbf{c}_i$ set to the cardinal axes $\mathbf{e}_i$. These parameters are the forms required for Gaussian EP to be used for hyperrectangular Gaussian probabilities.

3.2 Polyhedral Integration Regions

The previous section assumed axis-aligned integration limits. While a fully general set of integration limits is difficult to analyse, the derivations of the previous section can be extended to *polyhedral* integration regions. This extension is straightforward: each factor truncates the distribution along the axis orthogonal to the vector $\mathbf{c}_i$ (which we assume has unit norm without loss of generality) instead of the cardinal axes as in the previous section. The notational overlap of $\mathbf{c}_i$ with the Gaussian EP derivation of Section 2.1.1 is intentional, as these factors of Figure 2 are precisely our general rank one factors for EP. An important difference to note in Figure 2 is that there are m factors, where m can be larger or smaller (or equal to) the problem dimensionality n, unlike in the hyperrectangular case where $m = n$. To implement EP, we find the identical moments as shown in Equation

2. These erf(·) computations also require a numerically stable implementation, though the details do not add insight to this problem and are not specific to this problem (unlike Appendix B). Thus we defer those details to the EPMGP code package available at `<url to come with publication>`.

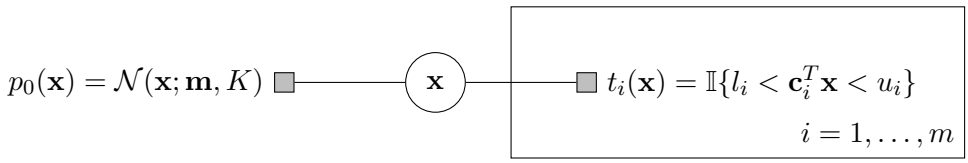

Figure 2: Factor graph for the polyhedral Gaussian integration problem. Note the onstensibly very similar setup to the rectangular problem in Figure 1. The only difference is a replacement of axis-aligned box functions with box functions on the axis defined by $\mathbf{c}_i$, and a change to the number of restrictions, from n to some different m.

21. In fact the only difference is the cavity parameters $\{\mu_{\backslash i}, \sigma^2_{\backslash i}\}$, which have the general form of Equation 17:

$$\mu_{\backslash i} \;=\; \sigma^2_{\backslash i}\Big(\frac{\mathbf{c}_i^T\boldsymbol{\mu}}{\mathbf{c}_i^T\Sigma\mathbf{c}_i} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\Big), \quad \text{and} \quad \sigma^2_{\backslash i} = \big((\mathbf{c}_i^T\Sigma\mathbf{c}_i)^{-1} - \tilde{\sigma}_i^{-2}\big)^{-1}. \tag{25}$$

From this point of view, it might seem like there is very little difference between the axis-aligned hyperrectangular setup of the previous section and the more general polyhedral setup of this section. However, our experiments in Section 6 indicate that EP can have very different performance between these two cases, a finding discussed at length in subsequent sections.

3.3 Implementation

To summarise, we have an EP algorithm that calculates $q(\mathbf{x}) = Z\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \Sigma)$ such that $D_{KL}(p\|q)$ is hopefully small. Though there is no proof for when EP methods converge to the global KL minimiser, our results (in the following sections) suggest that EPMGP does converge to a fixed point reasonably close to this optimum in the hyperrectangular integration case. In the general polyhedral case, performance is often excellent but can also be considerably worse, such that results in this case should be treated with caution. The EPMGP algorithm typically converges in fewer than 10 EP iterations (of all m factors) and is insensitive to initialisation and factor ordering. EPMGP pseudocode is given in Algorithm 1, and our MATLAB implementation is available at `<url to come with publication>`.

3.4 Features of the EPMGP algorithm

Before evaluating the numerical results of EPMGP, it is worth noting a few attractive features of the EP approach, which may motivate its use in many applied settings.

3.4.1 Derivatives with respect to the distribution's parameters

A common desire in machine learning and statistics is model selection, which involves optimising (or perhaps integrating) model parameters over an objective function. In the case of Gaussian probabilities, researchers may want to maximise the likelihood $F(\mathcal{A})$ by varying

Algorithm 1 EPMGP: Calculate $F(\mathcal{A})$, the probability of $p_0(\mathbf{x})$ on the region $\mathcal{A}$.

Initialise with any $q(\mathbf{x})$ defined by $Z, \boldsymbol{\mu}, \Sigma$ (typically the parameters of $p_0(\mathbf{x})$).
Initialise messages $\tilde{t}_i$ with zero precision.
while $q(\mathbf{x})$ has not converged **do**
  for $i \leftarrow 1 : m$ **do**
    form cavity local $q_{\setminus i}(\mathbf{x})$ by Equation 17 (stably calculated by 51).
    calculate moments of $q_{\setminus i}(\mathbf{x})t_i(\mathbf{x})$ by Equations 21-23.
    choose $\tilde{t}_i(\mathbf{x})$ so $q_{\setminus i}(\mathbf{x})\tilde{t}_i(\mathbf{x})$ matches above moments by Equation 16.
  end for
  update $\boldsymbol{\mu}, \Sigma$ with new $\tilde{t}_i(\mathbf{x})$ (stably using Equations 53 and 58).
end while
calculate Z, the total mass of $q(\mathbf{x})$ using Equation 19 (stably using Equation 60).
return Z, the approximation of $F(\mathcal{A})$.

$\{\mathbf{m}, K\}$ subject to some other problem constraints. Numerical integration methods do not offer this feature currently. Because the EP approach yields the analytic approximation shown in Equation 19, we can take derivatives of $\log Z$ so as to approximate $\partial F(\mathcal{A})/\partial \mathbf{m}$ and $\partial F(\mathcal{A})/\partial K$ for use in a gradient solver. Those derivatives are detailed in Appendix C.

As a caveat, it is worth noting that numerical integration methods like Genz were likely not designed with these derivatives in mind. Thus, though it has not been shown in the literature, perhaps an implementation could be derived to give derivative approximations via a clever reuse of integration points (though it is not immediately clear how to do so, as it would be in a simple Monte Carlo scheme). It is an advantage of the EPMGP algorithm that it can calculate these derivatives naturally and analytically, but we do not claim that it is necessarily a fundamental characteristic of the EP approach unshared by any other approaches.

3.4.2 Fast runtime and low computational overhead

By careful consideration of rank one updates and a convenient reparameterisation, we can derive a quadratic run time factor update as opposed to a naive cubic runtime (in the dimensionality of the Gaussian n). This quadratic factor update is done at each iteration, in addition to a determinant calculation in the $\log Z$ calculation, leading to a complexity of $\mathcal{O}(n^3)$ in theory (or $\mathcal{O}(mn^2 + n^3)$ in the case of polyhedral regions). Numerical integration methods such as Genz also require a Cholesky factorisation of the covariance matrix K and are thus $\mathcal{O}(n^3)$ in theory, and we have not found a more detailed complexity analysis in the literature. Further, the Genz method uses potentially large numbers of sample points (we use 5×10^5 to derive our high accuracy estimates), which can present a serious runtime and memory burden. As such, we find in practice that the EP-based algorithm always runs as fast as the Genz methods, and typically runs one to three orders of magnitude faster (the hyperrectangular EPMGP runs significantly faster than the general EPMGP, as the factor updates can be batched for efficiency in languages such as MATLAB). Furthermore, we find that EP always converges to the same fixed point in a small number of steps (typically 10 or fewer iterations through all factors). Thus, we claim there are computational advantages to running the EPMGP algorithm over other numerical methods.

There are three caveats to the above claims. First, since Genz numerical integration methods have a user-selected number of sample points, there is a clear runtime-accuracy tradeoff that does not exist with EP. EP is a deterministic approximation, and it has no guarantees about convergence to the correct $F(\mathcal{A})$ in the limit of infinite computation. Claims of computational efficiency must be considered in light of these two different characteristics. Second, since both EP and Genz algorithms are theoretically $\mathcal{O}(n^3)$, calculated runtime differences are highly platform dependent. Our implementations use MATLAB, the runtime results of which should not be trusted as reflective of a fundamental feature of the algorithm. Thirdly, since it is unclear if Genz's method (available from author) was written with computational efficiency in mind, a more rigorous comparison seems unfair. Nonetheless, Section 6 presents runtime/accuracy comparisons for a handful of representative cases so that the applied researcher may be further informed about these methods.

Despite these caveats, the EPMGP algorithm runs very fast and stably up to problems of high dimension, and it always converges quickly to what appears to be a unique mode of the EP energy landscape (empirical observation). Thus, we claim that runtime is an attractive feature of EP for calculating Gaussian probabilities.

3.4.3 Tail probabilities

In addition to fast runtime, careful reparameterisation and consideration of the EP parameters leads to a highly numerically stable implementation that is computed in log space. The lengthy but simple details of these steps are presented in Appendix B. Importantly, this stability allows arbitrarily small "tail" probabilities to be calculated, which are often important in applied contexts. In our testing the EPMGP algorithm can calculate tail probabilities with $\log Z < -10^5$, that is, $F(\mathcal{A}) \approx 10^{-50000}$. This is vastly below the range of Genz and other methods, which snap to $\log Z = -\infty$ anywhere below roughly $\log Z \approx -150$. These numbers reflect our own testing of the two methods. To be able to present results from both methods, we present the results of Section 6 over a much smaller range of probabilities, in the range of $\log Z \approx -20$ to $\log Z = 0$ ($F(\mathcal{A})$ from 10^{-9} to 1).

Again, it is worth noting that the Genz method was likely not designed with tail probabilities in mind, and perhaps an implementation could be created to handle tail probabilities (though it is not obvious how to do so). Thus, it is an advantage of the EPMGP algorithm that it can calculate tail probabilities naturally, but we do not claim that it is necessarily a unique characteristic of the EP approach.

3.5 The fundamental similarity between hyperrectangular and polyhedral integration regions

We here note important similarities between the hyperrectangular and polyhedral cases. Though brief, this section forms the basis for much of the discussion and exploration of the strengths and weaknesses of EP. Since we have developed EP for arbitrary polyhedral Gaussian probabilities, we can somewhat simplify our understanding of the problem. We can always whiten the space of the problem via a substitution $\mathbf{y} = L^{-1}(\mathbf{x} - \mathbf{m})$ where $K = LL^T$. In this transformed space, we have a standard $\mathcal{N}(0, I)$ Gaussian, and we have a transformed set of polyhedral faces $L^T\mathbf{c}_i$ (which must also be normalised as the $\mathbf{c}_i$ are). This whitening is demonstrated in Figure 3.

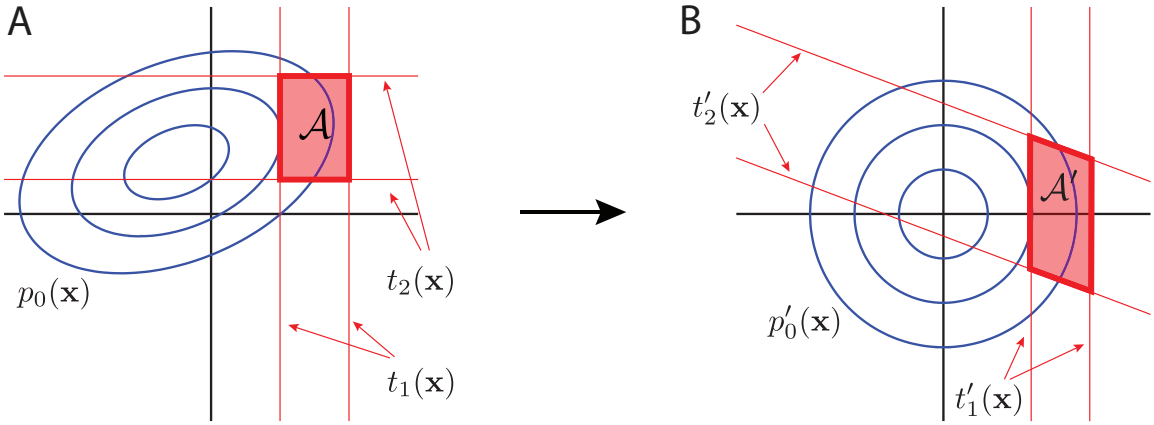

Figure 3: The effect of whitening the Gaussian probability based on the covariance of the Gaussian. Panel A shows a Gaussian probability problem in original form, with the integration region in red. Panel B shows the same problem when we have whitened (and mean-centered) space with respect to the Gaussian. We can then think of these problems only in terms of the geometry of the "coloured" integration region in the whitened space.

Thus, in trying to understand sources of error, we can then attend to properties of this transformed region. Intuition tells us that, the more hyperrectangular this transformed region is, the better performance, as hyperrectangular cases (in a white space) decompose into a product of solvable univariate problems. On the other hand, the more redundant/colinear/coloured those faces are, the higher the error that we should expect. We will thus consider metrics on the geometry of transformed integration region that exists in the whitened space. Further, we then see that hyperrectangular cases are in fact not a fundamental distinction. Indeed they are simply polyhedra where $m = n$, that is, the number of polyhedral constraints equals the number of dimensions of the problem. This seemingly small constraint, however, often makes the difference between an algorithm that performs well and one that does not, as the results will show.

Finally, it is important to connect this change to our EP approach. Because each step of EP minimises a KL-divergence on distributions, and because KL is invariant to invertible change of variables, each step of EP is invariant to invertible change of variables. This invariance requires that the factors, approximate factors, etc. are all transformed accordingly, so that the overall distribution is unchanged (as noted in Seeger (2008), Section 4). For Gaussian EP, the approximate factors are Gaussian, so the family is invariant to invertible *linear* transformations. Thus we can equivalently analyse our Gaussian EP algorithm for multivariate Gaussian probabilities in the white space, where $p_0(\mathbf{x})$ is a spherical Gaussian as in Panel B of Figure 3. Considering the whitened space and transformed integration region is conceptually useful and theoretically sound. This consideration allow us to maintain two views which have different focus: covariance structure in coloured space and a rectangular region (in the hyperectangular case), and geometry of the transformed polyhedron in whitened space. Much more will be made of this notion in the Results and Discussion.

4. Other Multivariate Gaussian Probability algorithms

This section details other methods to which we compare EPMGP. First, we describe the algorithms due to Genz (Section 4.1), which represent the state of the art for calculating Gaussian probabilities. Second, we briefly describe Monte Carlo-based methods (Section 4.2). Third, in some special cases there exist exact (*i.e.*, exact to machine precision) methods; we describe those in brief in Section 4.3. EPMGP and the algorithms discussed here will be the algorithms studied in numerical simulation.

4.1 Genz Numerical Integration Method

The Genz method (first described in Genz (1992), but the book Genz and Bretz (2009) is an excellent current reference) is a sophisticated numerical integration method for calculating Gaussian probabilities. The core idea is to make a series of transformations of the region $\mathcal{A}$ and the Gaussian $p_0(\mathbf{x})$ with the hope that this transformed region can be accurately integrated numerically. This approach performs three transformations to $F(\mathcal{A})$. It first whitens the Gaussian integrand $p_0(\mathbf{x})$ (via a Cholesky factorisation of K, which changes the integration bounds), and secondly it removes the integrand altogether by transforming the integration bounds with a function using the cdf and inverse cdf of a univariate Gaussian. Finally, a further transformation is done (using points uniformly drawn from the $[0, 1]$ interval) to set the integration region to the unit box (in $\mathbb{R}^n$). Once this is done, Genz (1992) makes intelligent choices on dimension ordering to improve accuracy. With all orderings and transformations completed, numerical integration is carried out over the unit hypercube, using either quasi-random integration rules (Niederreiter, 1972; Cranley and Patterson, 1976) or lattice-point rules (Cranley and Patterson, 1976; Nuyens and Cools, 2004). The original algorithm (Genz, 1992) reported, "it is possible to reliably compute moderately accurate multivariate normal probabilities for practical problems with as many as ten variables." Further developments including (Genz and Kwong, 2000; Genz and Brentz, 1999, 2002; Genz, 2004) have improved the algorithm considerably and have generalised it to general polyhedra. Our use of this algorithm was enabled by the author's MATLAB code `QSILATMVNV`, which is directly analogous to EPMGP for axis-aligned hyperrectangles. The author's more general algorithm for arbitrary polyhedra is available from the same code base as `QSCLATMVNV`: this is directly comparable to EPMGP. These functions run a vectorised, lattice-point version of the algorithm from (Genz, 1992), which we found always produced lower error estimates (this algorithm approximates the accuracy of its result to the true probability) than the vectorised, quasi-random `QSIMVNV` (with similar run-times). We also found that these vectorised versions had lower error estimates and significantly faster run-times than the non-vectorised `QSIMVN`. Thus, all results shown as "Genz" were gathered using `QSCLATMVNV`, which was available at the time of this report at: `http://www.math.wsu.edu/faculty/genz/software/software.html`. The software allows the user to define the number of lattice points used in the evaluation. For the Genz methods to which EPMGP was compared, we used 5×10^5 points. We found that many fewer points increased the error considerably, and many more points increased run-time with little improvement to the estimates. We note also that several statistical software packages have bundled some version of the Genz method into their product (*e.g.*, `mvncdf` in MATLAB, `pmvnorm` in R). Particular choices and calculations outside the literature are

typically made in those implementations (for example, MATLAB does not allow $n > 25$), so we have focused on just the code available from the author himself.

4.2 Sampling methods

It is worth briefly noting that Monte Carlo methods are not particularly well suited to this problem. In the simplest case, one could use a rejection sampler, drawing samples from $p_0(\mathbf{x})$ and rejecting them if outside the region $\mathcal{A}$ (and calculating the fraction of kept samples). This method requires huge numbers of samples to be accurate, and it scales very poorly with dimension n. The problem with this simple Monte Carlo approach is that the draws from $p_0(\mathbf{x})$ neglect the location of the integration region $\mathcal{A}$, so the rejection sampler can be terribly inefficient. We also considered importance sampling from the region $\mathcal{A}$ (uniformly with a "pinball" or "hit-and-run" MCMC algorithm (Herbrich, 2002)). However, this approach has the complementary problem of wasting samples in parts of $\mathcal{A}$ where $p_0(\mathbf{x})$ is small. Further, this scheme is inappropriate for open polyhedra. In an attempt to consider both $\mathcal{A}$ and $p_0(\mathbf{x})$, we used elliptical slice sampling (Murray et al., 2010) to draw samples from $p(\mathbf{x}) = p_0(\mathbf{x}) \prod_i t_i(\mathbf{x})$ (the familiar Bayesian posterior). We then calculated the empirical mean and covariance of these samples and used that posterior Gaussian as the basis of an importance sampler with the true distribution. Though better than naive Monte Carlo, this and other methods are deeply suboptimal in terms of runtime and accuracy, particularly in light of the Genz algorithms, which are highly accurate for calculating the "true" value, which allows a baseline for comparison to EPGMP. Thus we will not consider samplers further in this paper.

4.3 Special case methods

It is also useful to consider special cases where we can analytically find the true probability $F(\mathcal{A})$. The most obvious is to consider hyperrectangular integration regions over a diagonal covariance K. Indeed both EPMGP and Genz methods return the correct $F(\mathcal{A})$ to machine precision, since these problems trivially decompose to univariate probabilities. While a good sanity check, this finding is not instructive. Instead, here we consider a few cases where the integration region yields a geometric solution in closed form.

4.3.1 Orthant Probabilities

Orthants generalise quadrants to high dimension, for example the nonnegative orthant $\{\mathbf{x} \in \mathbb{R}^n : x_i > 0 \, \forall i\}$. For zero-mean Gaussians of low dimension, calculating the probability $F(\mathcal{A})$ when $\mathcal{A}$ is an orthant can be solved geometrically. In $\mathbb{R}^2$, any orthant (quadrant) can be whitened by the Gaussian covariance K, which transforms the orthant into a wedge. The angle subtended by that wedge, divided by 2π, is the desired probability $F(\mathcal{A})$. With that intuition in $\mathbb{R}^2$, researchers have also derived exact geometric solutions in $\mathbb{R}^3, \mathbb{R}^4$, and $\mathbb{R}^5$ (Genz and Bretz, 2009; Sinn and Keller, 2010). We can thus compare EP and Genz methods to the true exact answers in the case of these orthant probabilities. Orthant probabilities also have particularly important application in applied settings such as probit analysis (Ashford and Sowden, 1970; Sickles and Taubman, 1986; Bock and Gibbons, 1996), so knowing the performance of various methods here is valuable.

4.3.2 Trapezoid Probabilities

Trapezoid probabilities are another convenient construction where we can know the true answer. If we choose our Gaussian to be zero-mean with a diagonal covariance K, and we choose $\mathcal{A}$ to be a hyperrectangle that is symmetric about zero (and thus the centroid is at zero), then this probability decomposes simply into a product of univariate probabilities. We can then take a non-axis cut through the center of this region. If we restrict this cut to be in only two dimensions, then the region in those two dimensions is now a trapezoid (or a triangle if the cut happens to pass through the corner, but this happens with zero probability). By the symmetry of the Gaussian and this region, we have also cut the probability in half. Thus, we can calculate the simple rectangular probability and half it to calculate the probability of this generalised trapezoid.

The virtue of this construction is that this trapezoidal region appears as a nontrivial, general polyhedron to the EPMGP and Genz methods, and thus we can use this simple example to benchmark those methods. Unlike orthant probabilities, trapezoids are a contrived example so that we can test these methods. We do not know of any particular application of these probabilities, so these results should be seen solely for vetting of the EP and Genz algorithms.

5. Cases for Comparing Gaussian Probability Algorithms

The following section details our procedure for drawing example integration regions $\mathcal{A}$ and example Gaussians $p_0(\mathbf{x})$ for the numerical experiments. Certainly understanding the behavior of these algorithms across the infinite space of parameter choices is not possible, but our testing in several different scenarios indicates that the choices below are representative.

5.1 Drawing Gaussian distributions $p_0(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{m}, K)$

To define a Gaussian $p_0(\mathbf{x})$, we first randomly drew one positive-valued n-vector from an exponential distribution with mean 10 ($\lambda = 0.1$, n independent draws from this distribution), and we call the corresponding diagonal matrix the matrix of eigenvalues S. We then randomly draw an $n \times n$ orthogonal matrix U (orthogonalising a matrix of iid $\mathcal{N}(0, 1)$ random variables with the singular value decomposition suffices), and we form the Gaussian covariance matrix $K = USU^T$. This procedure produces a good spread of differing covariances K with quite different eigenvalue spectra and condition numbers. We note that this procedure produced a more interesting range of K - in particular, a better spread of condition numbers - than using draws from a Wishart distribution. We then set the mean of the Gaussian $\mathbf{m} = 0$ without loss of generality (note that, were $\mathbf{m} \neq 0$, we could equivalently pick the region $\mathcal{A}$ to be shifted by $\mathbf{m}$, and then the problem would be unchanged; instead, we leave $\mathbf{m} = 0$, and we allow the randomness of $\mathcal{A}$ to suffice). For orthant probabilities, no change to this procedure is required. For trapezoid probabilities, a diagonal covariance is required, and thus we set $U = I$ in the above procedure.

5.2 Drawing integration regions $\mathcal{A}$

Now that we have the Gaussian $p_0(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{m}, K)$, we must define the region $\mathcal{A}$ for the Gaussian probability $F(\mathcal{A})$.

5.2.1 Axis-aligned hyperrectangular regions

The hyperrectangle $\mathcal{A}$ can be defined by two n-vectors: the upper and lower bounds $\mathbf{u}$ and $\mathbf{l}$, where each entry corresponds to a dimension of the problem. To make these vectors, we first randomly drew a point from the Gaussian $p_0(\mathbf{x})$ and defined this point as an interior point of $\mathcal{A}$. We then added and subtracted randomly chosen lengths to each dimension of this point to form the bounds u_i and l_i. Each length from the interior point was chosen uniformly between 0.01 and the dimensionality n. Having the region size scale with the dimensionality n helped to prevent the probabilities $F(\mathcal{A})$ from becoming vanishingly small with increasing dimension (which, as noted previously, is handled fine by EP but problematic for other methods).

5.2.2 Arbitrary polyhedral regions

The same procedure is repeated for drawing the interior point and the bounds l_i and u_i for all m constraints. The only difference is that we also choose axes of orientation, which are unit-norm vectors $\mathbf{c}_i \in \mathbb{R}^n$ for each of the m factors. These vectors are drawn from a uniform density on the unit hypersphere (which is achieved by normalising n iid $\mathcal{N}(0,1)$ draws). The upper and lower bounds are defined with respect to these axes.

With these parameters defined, we now have a randomly chosen $\mathcal{A}$ and the Gaussian $p_0(\mathbf{x})$, and we can test all methods as previously described. This procedure yields a variety of Gaussians and regions, so we believe our results are strongly representative.

It is worth noting here that arbitrary polyhedra are not uniquely defined by the "intersection of boxes" or more conventional "intersection of halfspaces" definition. Indeed, there can be inactive, redundant constraints which change the representation but not the polyhedron itself. We can define a minimal representation of any polyhedron $\mathcal{A}$ to be the polyhedron where all constraints are active and unique. This can be achieved by pulling in all inactive constraints until they support the polyhedron $\mathcal{A}$ and are thus active, and then we can prune all repeated/redundant constraints. These steps are solvable in polynomial time as a series of linear programs, and they bear important connection to problems considered in centering (analytical centering, Chebyshev centering, etc.) for cutting-plane methods and other techniques in convex optimisation (*e.g.*, see Boyd and Vandenberghe (2004)). These steps are detailed in Appendix D. To avoid adding unnecessary complexity, we do not minimalise polyhedra in our results, as that can mask other features of EP that we are interested in exploring. Anecdotally, we report that while minimalising polyhedra *will* change EP's answer (a different set of factors), the effect is not particularly large in these randomly sampled polyhedral cases. In our empirical and theoretical discussion of the shortcomings of EP, we will heavily revisit this notion of constraint redundancy.

6. Results

We have thus far claimed that EP provides a nice framework for Gaussian probability calculations, and it offers a number of computational and analytical benefits (runtime, tail probabilities, derivatives). We have further claimed that existing numerical integration methods, of which the Genz method is state of the art, can reliably compute Gaussian probabilities to high accuracy, such that they can be used as "ground truth" when evaluating

EP. Thirdly, we claim that this Gaussian probability calculation elucidates a pathology in EP that is not often discussed in the literature. Giving evidence to those claims is the purpose of this Results section, which will then be discussed from a theoretical perspective in Section 7.

The remainder of this results section is as follows: we first demonstrate the accuracy of the EP method in the hyperrectangular case across a range of dimensionalities (Section 6.1). Using the same setup, we then demonstrate a higher error rate for the general polyhedral case in Section 6.2, testing the sensitivity of this error to both dimensionality n and the number of polyhedral constraints m. In Section 6.3, we evaluate the empirical results for the special cases where the true Gaussian probability can be calculated by geometric means, as introduced in the previous Section 4.3, which allows us to compare the EP and Genz approaches. We next test the empirical results for the commonly tested case of equicorrelated Gaussian probabiliites (Section 6.4). The equicorrelation and special case results point out the shortcomings of the EP approach. We can then proceed to construct pathological cases that clarify where and why EP gives poor approximations of Gaussian probabilities, which is done in Section 6.5. This final section gives the reader guidelines about when EP can be effectively used for Gaussian probabilities, and it provides a unique perspective and leads to a theoretical discussion about negative aspects of EP rarely discussed in the literature.

6.1 EP results for hyperrectangular integration regions

Figure 4 shows the empirical results for the EPMGP algorithm in the case of hyperrectangular integration regions. For each dimension $n = \{2, 3, 4, 5, 10, 20, 50, 100\}$, we used the above procedure to choose a random Gaussian $p_0(\mathbf{x})$ and a random region $\mathcal{A}$. This procedure was repeated 1000 times at each n. Each choice of $p_0(\mathbf{x})$ and $\mathcal{A}$ defines a single Gaussian probability problem, and we can use EPMGP and Genz methods to approximate the answer. We calculate the relative error between these two estimates. Each of these errors is plotted as a point in Figure 4. We plot only 250 of the 1000 points at each dimensionality, for purposes of visual clarity in the figures. The four panels of Figure 4 plot the same data on different x-axes, in an effort to demonstrate what aspects of the problem imply different errors. As described in the figure caption, Panel A plots the errors of EPMGP by dimensionality n. Horizontal offset at each dimension is added jitter to aid the reader in seeing the distribution of points (panel A only). The black dashed line indicates the $1\% \log Z$ error. Finally, the black line below the EP results is the median error estimate given by Genz. Since this line is almost always a few orders of magnitude below the distribution of EP errors, it is safe to use the high accuracy Genz method as a proxy to ground truth.

Panel A demonstrates that the EP method gives a reliable estimate, typically with errors on the order of 10^{-4}, of Gaussian probabilities with hyperrectangular integration regions. The shape of the blue curve in Panel A seems to indicate that error peaks around $n = 10$. We caution the reader not to over-interpret this result, as this feature may be just as much about the random cases chosen (as described in Section 5) as a fundamental feature of Gaussians. Nonetheless, we do notice across many test cases that error decreases at higher dimension, perhaps in part due to the concentration of mass of high dimensional Gaussians on the 1σ covariance ellipsoid.

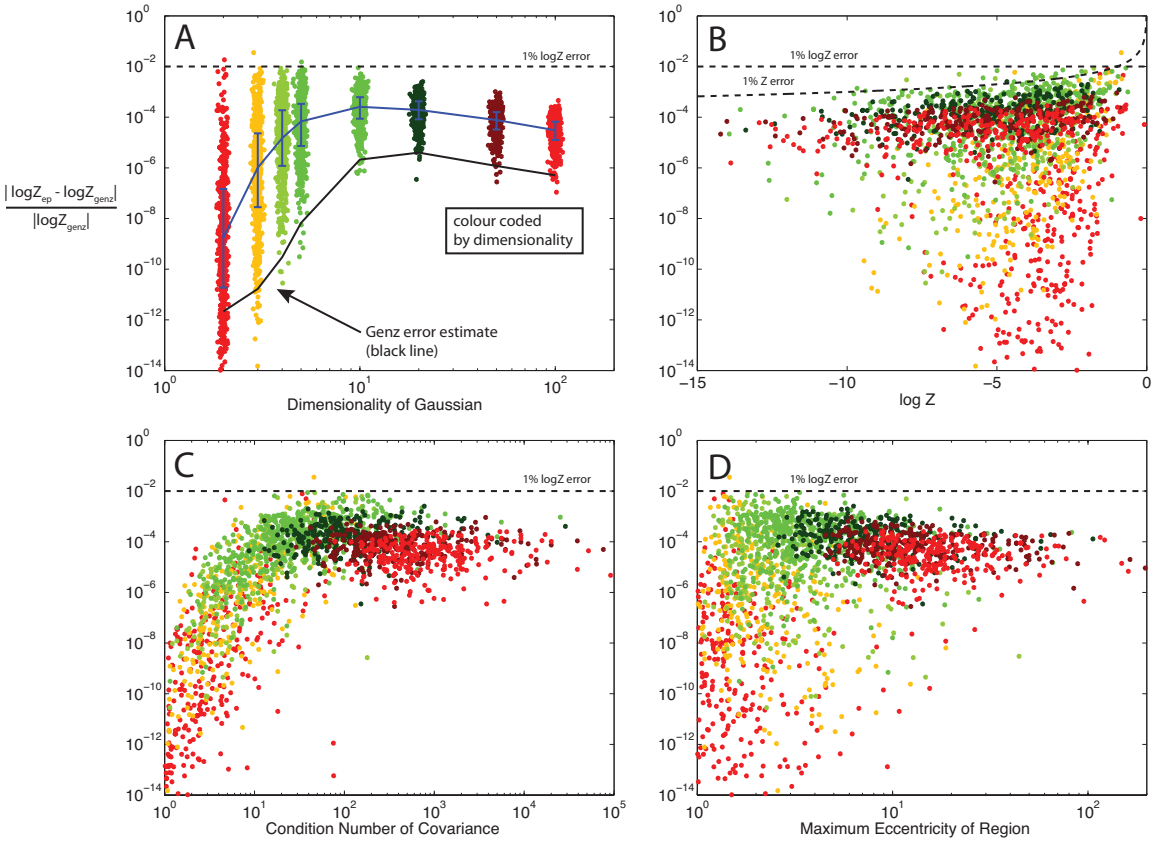

Figure 4: Empirical results for EPMGP with hyperrectangular integration regions. Each point represents a random Gaussian $p_0(\mathbf{x})$ and a random region $\mathcal{A}$, as described in Section 5. In all panels the colour coding corresponds to the dimensionality of the Gaussian integral (chosen to be $n = \{2, 3, 4, 5, 10, 20, 50, 100\}$. All panels show the relative log error of the EP method with respect to the high accuracy Genz method. Panel A plots the errors of EPMGP by dimensionality (note log-log scale). The horizontal offset at each dimension is added jitter to aid the reader in seeing the distribution of points (panel A only). The blue line and error bars plot the median and $\{25\%, 75\%\}$ quartiles at each dimension n. The black dashed line indicates the 1% $\log Z$ error. Finally, the black line below the EP results is the median error estimate given by Genz. Since this line is almost always a few orders of magnitude below the EP distribution, it is safe to use the high accuracy Genz method as a proxy to ground truth. Panel B plots the errors by the value of $\log Z$ itself. On this panel we also plot a 1% error for $Z \approx F(\mathcal{A})$. Panel C repeats the same errors but plotted by the condition number of the covariance K. Panel D plots the errors vs the maximum eccentricity of the hyperrectangle, namely the largest width divided by the smallest width.

Panel B plots the errors by the value of $\log Z$ itself, where we use the $\log Z$ calculated from Genz. On this panel we can also plot a 1% error for $Z \approx F(\mathcal{A})$, as the log scale can suggest misleadingly large or small errors. A priori we would not expect any significant trend to be caused by the actual value of the probability, and indeed there seems to be a weak or nonexistent correlation of error with $\log Z$. Panel C repeats the same errors but plotted by the condition number of the covariance K. In this case perhaps we would expect to see trend effects. Certainly the EP algorithm should be correct to machine precision in the case of a hyperrectangular region and a white covariance, since this problem decomposes into n one-dimensional problems. Indeed we do see this accuracy in our tests. Increasing the eccentricity of the covariance, which can be roughly measured by the condition number, should and does increase the error, as shown in Panel C. This effect can be seen more clearly in later figures (when we test equicorrelation matrices). One might also wonder about the condition number of the correlation, as this value would neglect axis-scaling effects. We have also tested this and do not show this panel as it looks nearly identical to Panel C.

Panel D plots the errors vs the maximum eccentricity of the hyperrectangle, namely the largest width divided by the smallest width. As with Panel B, it is unclear why the eccentricity of a hyperrectangular region should have bearing on the accuracy of the computation, and indeed there appears to be little effect. The colour coding helps to clarify this lack of effect, because each dimensionality seems to have insensitivity to error based on changing eccentricity. Thus any overall trend appears more due to the underlying explanatory factor of dimensionality than due to the region eccentricity itself.

Finally, as an implementation detail, we discuss the Genz error estimate (as seen in the black line of Panel A). It is indeed true that this error estimate can sometimes be significant, but that happens rarely enough ($< 10\%$) that it has little bearing on the trends shown here. To be strict, one can say that the lowest 10% of EP errors are conservative, noisy errors: they could be significantly smaller, but not significantly larger. Due to the log scale of these plots, the larger EP errors (the ones we care about) are reliable. Furthermore, we note that the Genz errors increase significantly at higher dimension $n > 100$. The EP method can compute larger dimensions without trouble, but the Genz method achieves a high enough error estimate that many more, often a majority, of EP errors are to be considered conservative, noisy estimates. Empirically, it seems the errors in higher dimensions continue their downward trend, but they are noisier measurements, so we do not report them to avoid over penalising the EP method (as they could be much smaller).

Overall, the empirical results shown here indicate that EPMGP is a successful candidate algorithm for Gaussian probabilities over hyperrectangular integration regions. The error rate is non-zero but generally quite low, with median errors less 10^{-4} and individual errors rarely in excess of 1% across a range of dimensions, which may be acceptable in many applied settings.

6.2 EP results for general polyhedral cases

We have previously shown how the EP framework readily can be extended to arbitrary polyhedral regions with only a minor change to the calculation of cavity parameters (Equation 25). Figure 5 shows the empirical results of this change, which are perhaps surprisingly large. In the same way we drew 1000 test cases at each dimensionality n. The only differ-

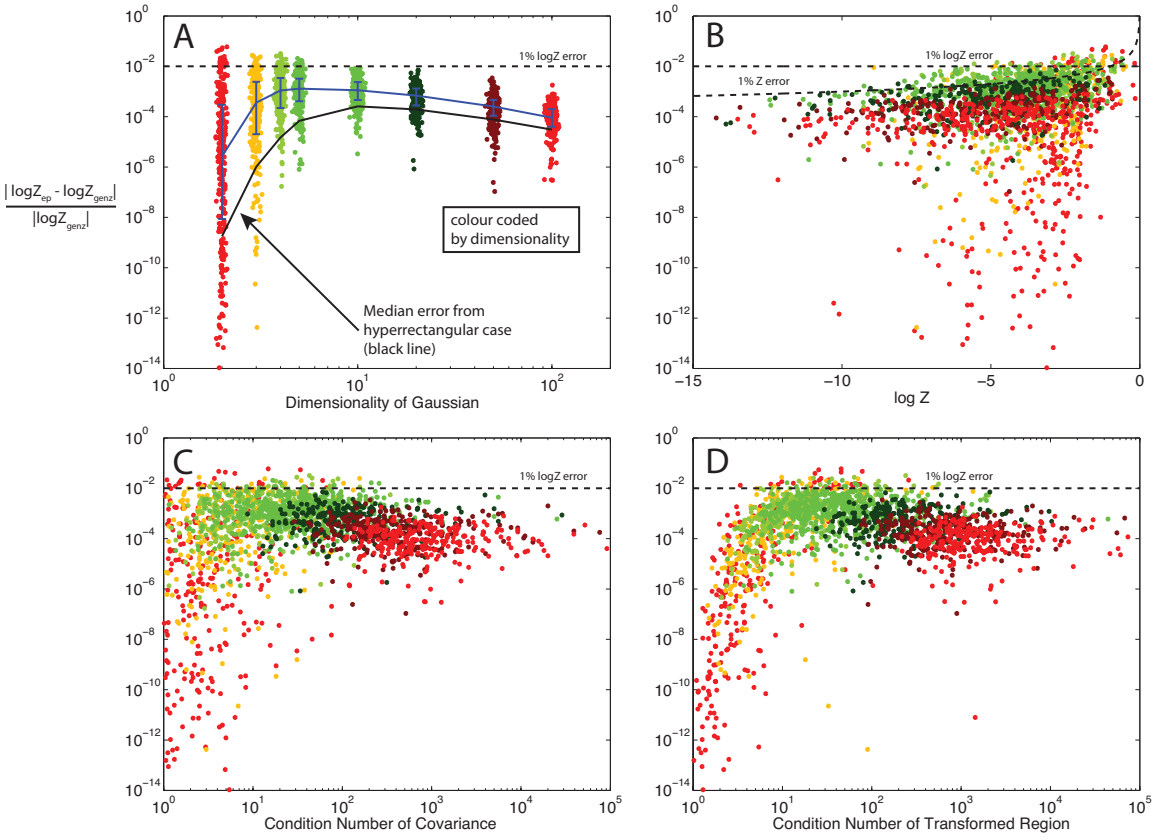

Figure 5: Empirical results for EPMGP with polyhedral integration regions. This figure has largely the same description as Figure 4. The only difference from that previous figure is that panel A depicts as a black line the median errors from EPMGP in the hyperrectangular case. This demonstrates the consistently higher errors of general polyhedral EPMGP. Panels B and C have the same description as Figure 4. Panel D plots the condition number of the transformed region. As discussed in the text, any Gaussian probability problem can be whitened to consider a Gaussian with identity covariance. This step merely linearly transforms the region $\mathcal{A}$. Panel D shows the errors plotted vs. the condition number of the integration region transformed appropriately by the Gaussian covariance K.

ence is that we also drew $m = n$ random polyhedral faces as described in Section 5, instead of the fixed hyperrectangular faces in the previous case.

Figure 5 has the same setup as 4, with only a few exceptions. First, the black line in Panel A shows the median EP errors in the axis aligned case (*i.e.*, the blue line in Figure 4, Panel A). The purpose of this line is to show that indeed for the same Gaussian cases, the error rate of polyhedral EPMGP is typically an order of magnitude or two higher than in the hyperrectangular case. Genz errors are not shown (as they were in Figure 4), but we note that they tell the same story as in the hyperrectangular case. The only other change is Panel D. In Section 3.5, we introduced the concept of whitening the space to consider just the region integrated over a standard Gaussian. We consider this region as a matrix with columns equal to $L^T \mathbf{c}_i$ (normalised), the transformed polyhedral faces. If this region were hyerrectangular (in the transformed space), then this matrix would be I, and the method would be perfectly accurate (decomposed problem). Thus, to test the transformed region's proximity to I, we can use the condition number. We do so in Panel D, where we plot errors against that value. As in Panel C, there appears to be some legitimate trend based on the condition number of the transformed region. This finding confirms our intuition that less hyperrectangular transformed integration regions will have higher error. There are a number of other metrics on the transformed region that one might consider that attempt to account for colinearity (non-orthogonality) in the transformed region constraints. We have produced similar plots with Frobenius (l_2) and l_1 norms of the elements of the scaled Gram matrix (instead of condition number), and we report that the trends found there are similar to these figures: there is some upward trend in error with a metric of colinearity.

Figure 5 only shows cases where the number of polyhedral constraints is the same as the dimensionality of the Gaussian ($m = n$). The sensitivity of error to the number of constraints m is also an important question. Figure 6 gives these empirical results. We use the same setup as Figure 5 with only one change: Panel A plots errors by number of polyhedral face constraints m instead of by Gaussian dimension n. In this figure all cases use $n = 10$ dimensions, and we show polyhedral sizes of $m = \{2, 4, 8, 10, 12, 16, 32, 64\}$. It seems from Panel A that larger numbers of constraints/polyhedral faces do imply larger error, trending up by roughly an order of magnitude or two. This is an interesting finding that will be revisited in the theoretical discussion: for a given dimension n, more polyhedral constraints (more factors) implies higher EP error. The errors shown in B and C seem largely invariant to the condition number of the covariance and the value $\log Z$ itself. However, Panel D still shows the error in the whitened space and indicates again some upward trend, as we would expect.

Figure 5 and 6 show error rates that may in many cases be unacceptable. Though EP still performs well often, the reliability of the algorithm is not what it is in the hyperrectangular case of Figure 4. We now test special cases where true answers are known, which will further our understanding of the strengths and weaknesses of the EP approach.

6.3 EP and Genz results for special cases with known answers

Thus far we have relied on the Genz method to provide the baseline of comparisons. Though we have argued that this step is valid, it is still useful to consider the handful of special

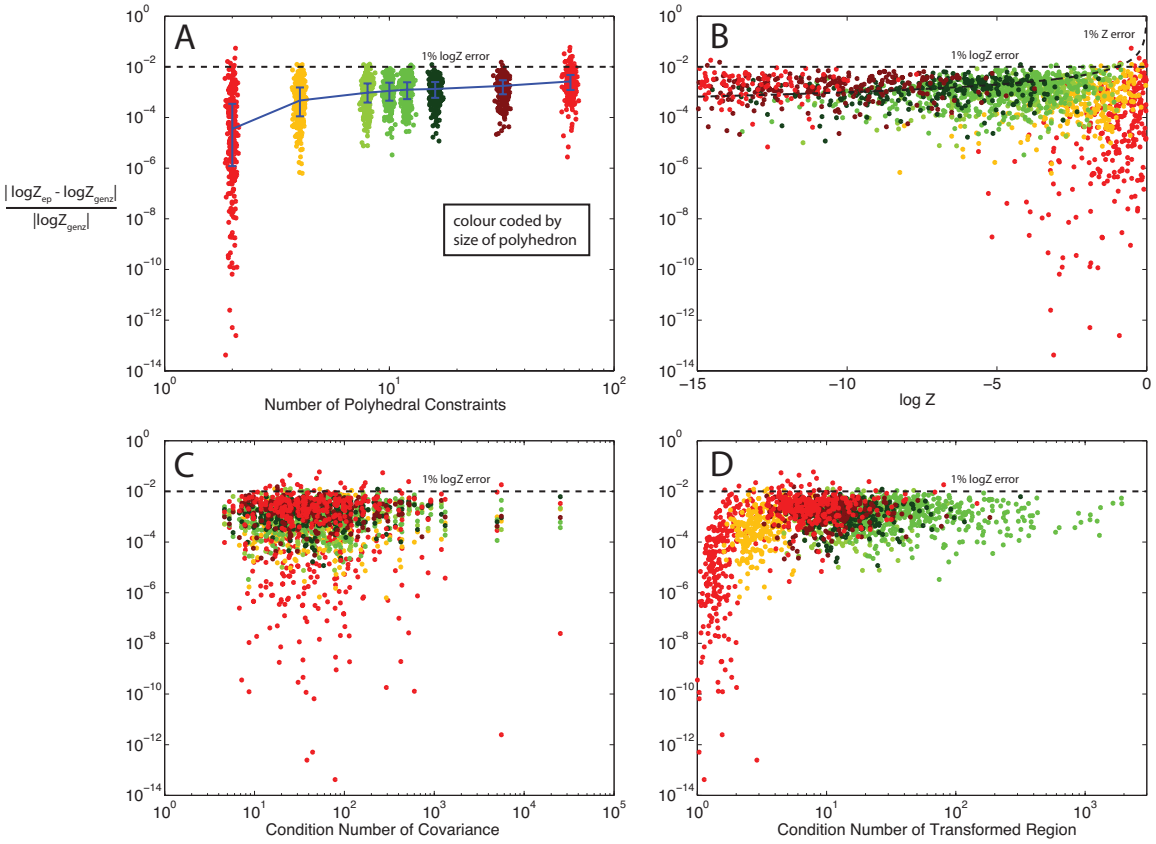

Figure 6: Empirical results for EPMGP with polyhedral integration regions. Panel A calculates the same error metric as in the previous figures for $n = 10$ dimensional Gaussians with varying numbers of polyhedral constraints $m = \{2, 4, 8, 10, 12, 16, 32, 64\}$. Note the difference with Figure 5 where dimensionality n is swept (and $m = n$). Other panels are as previously described in Figure 5.

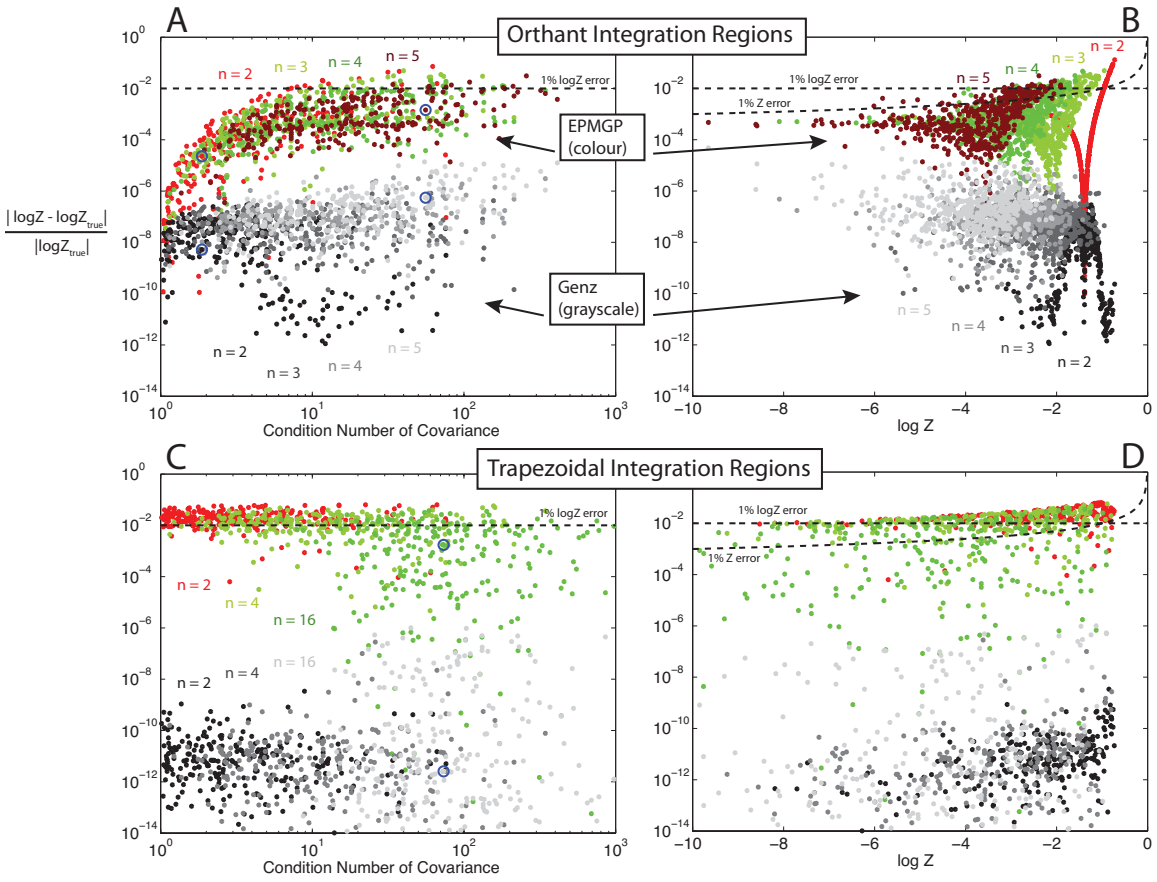

Figure 7: Empirical results for the special cases where an analytical true answer can be calculated. Panels A and B show results for orthant probabilities of dimension $n = \{2, 3, 4, 5\}$ (See Section 4.3). Panel A shows relative $\log Z$ error, and Panel B shows relative Z error. The coloured points indicate EP errors, and the grayscale points indicate the errors of the Genz method. Though not shown in this figure, the Genz estimates of error only slightly and occasionally overlap the coloured distribution of EP errors, which again clarifies that Genz can be effectively used as a proxy to ground truth. Panels C and D plot errors for trapezoid probabilities (see Section 4.3), again with EP and Genz errors. Note also the three pairs of points in Panel A and C which are circled in blue. These are three representative examples: one each of an orthant probability in $n = 2$, an orthant probability in $n = 5$, and a trapezoid probability in $n = 16$. We use these representative examples in the following Figure 8 to evaluate the runtime/accuracy tradeoff of the EP and Genz methods.

cases where the true probability can be calculated from geometric arguments. We do so here, and those results are shown in Figure 7.

Panels A and B of Figure 7 show errors for both the EP method (colour) and the Genz method (grayscale). The two panels are the usual errors plotted as a function of condition number of covariance (Panel A) and the true probability $\log Z$ (Panel B). The four cases shown are orthant probabilities at $n = \{2, 3, 4, 5\}$. There are a few interesting features

worthy of mention. First, we note that there is a clear separation between the Genz errors and the EP errors (in log-scale no less). This finding helps to solidify the earlier claim that the Genz numerical answer can be used as a reasonable proxy to ground truth. We can also view this figure with the Genz error estimates as error bars on the grayscale points. We have done so and report that the majority of the coloured points remain above those error bars, and thus the error estimate reinforces this difference between Genz and EP. We do not show those error bars for clarity of the figure.

Second, it is interesting to note that there is significant structure in the EP errors with orthant probabilities when plotted by $\log Z$ in Panel B. Each dimension has a "V" shape of errors. This can be readily explained by reviewing what an orthant probability is. For a zero-mean Gaussian, an orthant probability in $\mathbb{R}^2$ is simply the mass of one of the quadrants. If that quadrant has a mass of $\log Z = \log 0.25$, then the correlation must be white, and hence EP will produce a highly accurate answer (the problem decomposes). Moving away from white correlation, EP will produce error. This describes the shape of the red curve for $n = 2$, which indeed is minimised at $\log 0.25 \approx -1.39$. The same argument can be extended for why there should be a minimum in $\mathbb{R}^3$ at $\log Z = \log 0.125$, and so on. It is interesting to note that the Genz errors have structure of their own in this view, which is presumably related to the same fact about orthant probabilities. Investigating the pattern of Genz errors further is beyond the scope of this work, however.

Panels C and D plot the same figures as Panels A and B, but they do so with trapezoid probabilities. As described above, trapezoid probabilities involve mean-centered hyper-rectangular integration regions and Gaussians with diagonal covariance. There is then an additional polyhedral constraint that cuts two dimensions of the region through the origin such that the probability is halved. This appears to EP and to Genz as an arbitrary polyhedron in any dimensionality. We show here results for $n = \{2, 4, 16\}$. Panels C and D clarify that EP works particularly poorly in this case, often with significant errors. This finding highlights again that EP with general polyhedral regions can yield unreliable results.

In Figure 7, Panels A and C also show 3 pairs of points circled in blue. Those points are three representative test cases - an orthant probability in $n = 2$, an orthant probability in $n = 5$, and a trapezoid probability in $n = 16$. We chose these as representative cases to demonstrate runtime/accuracy tradeoffs, which is shown in Figure 8. This figure is offered for the applied researcher who may have runtime/accuracy constraints that need to be balanced. Panel A shows the familiar $\log Z$ errors, and Panel B shows the same cases with Z error. Each colour gives the error plotted as a function of runtime for one of the pair of points circled in Figure 7. For the Genz method, the lines are formed by sweeping the number of integration points used from 50 to 10^6. For the EP method (results in bold), we vary the number of EP iterations allowed, from 2 to 20. Since the EP method rarely takes more than 10 iterations to converge numerically, the bold lines are rather flat. Each point in these panels is the median runtime of 100 repeated runs (the results and errors were identical for each case, as both methods are deterministic).

Figure 8 shows that EPMGP runs one to three orders of magnitude faster than the Genz method. However, since EP is a Gaussian approximation and not a numerical computation of the correct answer, one can not simply run EPMGP longer to produce a better answer. The Genz method on the other hand has this clear runtime/accuracy tradeoff. Our results throughout the other figures use the high accuracy Genz result (5×10^5 integration points

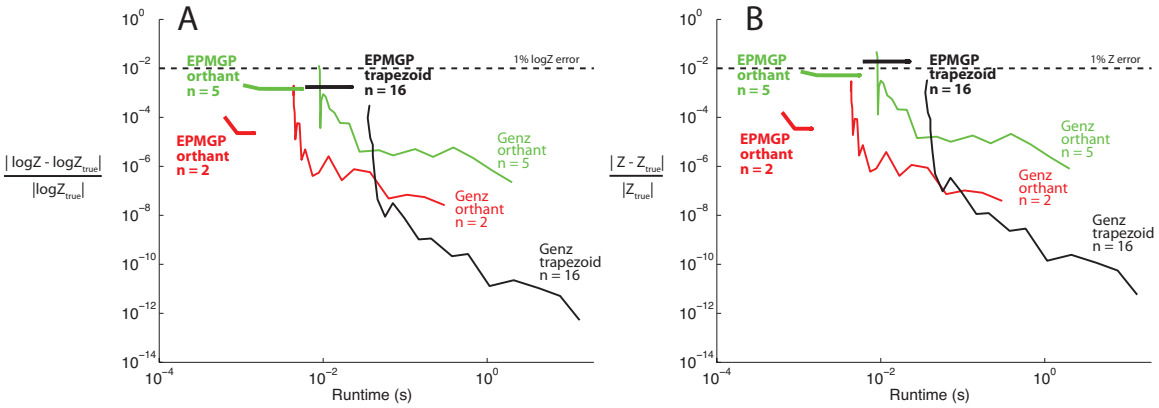

Figure 8: Runtime analysis of the EPMGP and Genz algorithms. The three pairs of lines are three representative examples, which were circled in blue in Figure 7. They include orthant probabilities at dimension $n = 2$ and $n = 5$, and a trapezoid probability at $n = 16$. Each point in this Figure is the median run time of 100 repeated trials (both methods give the same error on each repetition). The Genz method has a clear runtime/accuracy tradeoff, as more integration points significantly increases accuracy and computational burden. For Genz we swept the number of integration points from 50 to 10^6. The EP method was run with a fixed maximum number of iterations (from 1 to 20). Because the method converges quickly, rarely in more than 10 steps, these lines typically appear flat. Note that EP runs one to three orders of magnitude more quickly than Genz, but more computation will not make EP more accurate.

- the rightmost point on the Genz lines) and the fully converged EP method (again the rightmost points on the EP lines). Finally, we note that we chose three representative examples, but these are highly reflective of the typical runtime/accuracy tradeoff we have seen across many more test cases.

6.4 EP results for equicorrelated Gaussian probabilities

Another interesting case for study is the case of Gaussians with equicorrelated covariance matrices. These Gaussians have covariances with all diagonal entries equal to 1, and all off-diagonal entries equal to ρ, where $-1/(n-1) \leq \rho \leq 1$ (these bounds ensure a positive semidefinite matrix). This case is interesting primarily because it allows us to look at a parameterised family of covariance matrices that include the identity matrix ($\rho = 0$, where we believe EP should produce a highly accurate answer) and highly elongated matrices ($\rho \to 1$ or $\rho \to -1/(n-1)$, where Figures 5 and 7 suggest that EP should do poorly). This case is also important as it connects to previous literature on the subject of Gaussian probabilities, where equicorrelated matrices are a common test case (such as Gassman et al. (2002); Shervish (1984)).

Each panel tests the three cases of $n = \{2, 4, 8\}$, with the lines colour coded as in previous figures. The x-axis of these figures sweeps the valid range of ρ values. At each ρ value, which fully specifies a Gaussian covariance, we draw the region $\mathcal{A}$ and Gaussian mean $\mathbf{m}$ in the usual way, as previously described. We do this 100 times at each value of ρ, calculating the EP and Genz results, from which we can compute errors in $\log Z$ and Z in the usual way. Again the line and error bars at each point show the median and $\{25\%, 75\%\}$ quartiles. Panels A and B show these results with hyperrectangular integration regions, and Panels C and D show the results with general polyhedral regions.

Panels A and B behave as intuition suggests: when the covariance K is nearly white, the error is minimal. As the correlation becomes larger and the matrix more eccentric, error increases. Also in the hyperrectangular case we see that errors become meaningful - around 1% - only in the most eccentric cases where the matrix is nearly singular.

Panels C and D appear to have a much flatter error, which perhaps disagrees with intuition. The error seems relatively insensitive to the eccentricity of the correlation structure. Errors do have a slight downward trend towards 0, but it is by no means significant. This finding can be readily explained by returning to the notion of the whitened space (Figure 3 and Panel D of Figures 5 and 6). This result also helps us transition into understanding pathologies of EP. Recall that the problem can be whitened by the covariance of the Gaussian, such that the Gaussian is standardised to identity covariance and the region carries all notions of eccentricity. If we revisit Panels A and B in this light, we see that the EP method is doing well when the transformed region is close to hyperrectangular ($\rho \approx 0$), but then error grows as the region becomes coloured/non-rectangular ($\rho \to 1$ or $\rho \to -1/(n-1)$). The intuition and the results are unchanged in Panels A and B. Turning to Panels C and D, we can imagine the same effect as ρ is changed. However, because we are in the case of polyhedral integration regions, the region we draw is *already* coloured before we whiten the space. Thus, when we whiten the space, we will equally often make a region less hyperrectangular and more hyperrectangular. In this view, sweeping this correlation matrix should

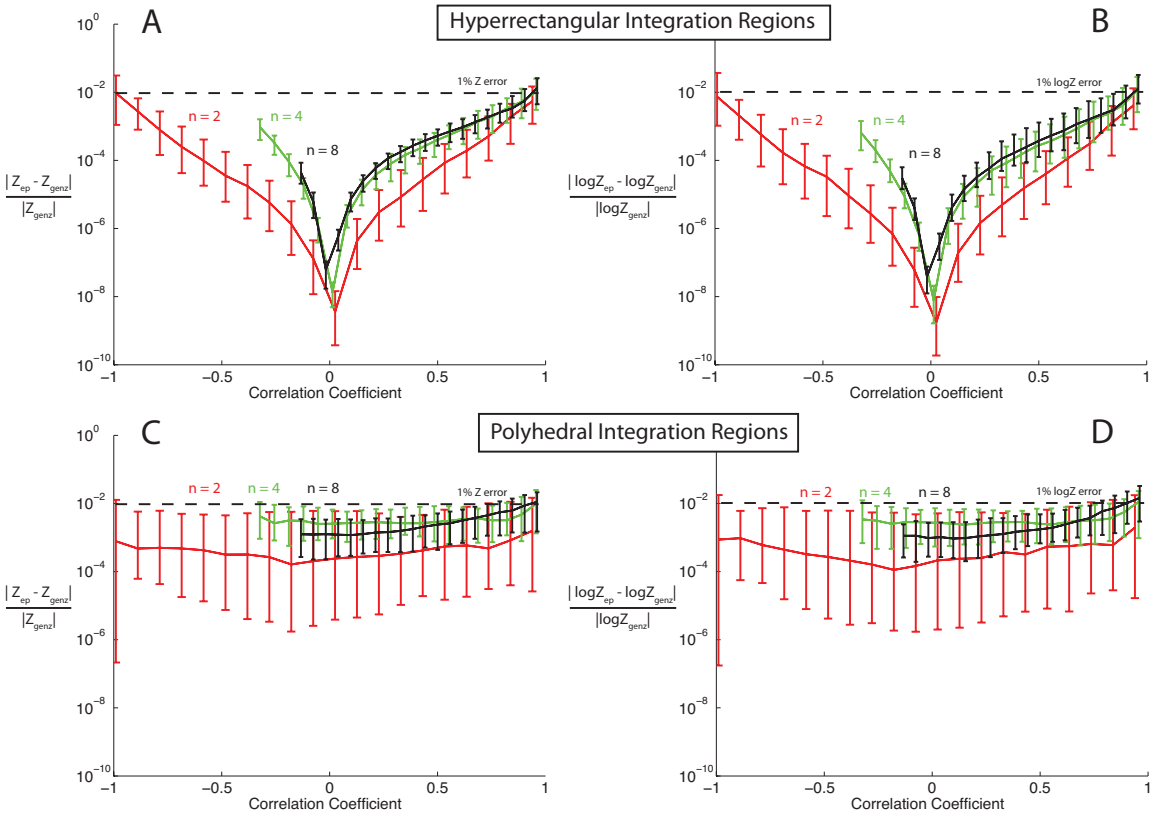

Figure 9: Empirical results for EPMGP over equicorrelated Gaussians. Historically, a common test case for Gaussian probability algorithms has been to test error on correlation matrices of equal correlation (see *e.g.*, Gassman et al. (2002); Shervish (1984)). All other quantities - the region $\mathcal{A}$ and Gaussian mean $\mathbf{m}$ - are still drawn randomly as previously described. Here we test error for both the hyperrectangular case (panels A and B) and the polyhedral case (panels C and D), showing relative error in Z (panels A and C) and relative error in $\log Z$ (panels B and D). We show three different dimensionalities $n = \{2, 4, 8\}$ coloured as red, green, and black, respectively. 100 random cases were run at each correlation value. Each error bar shows the median and $\{25\%, 75\%\}$ quartiles. The x-axis sweeps the correlation coefficient ρ throughout its range for each dimension (which is the $[-1/(n-1), 1]$ interval; matrices outside this range are not valid, positive semidefinite correlation matrices).

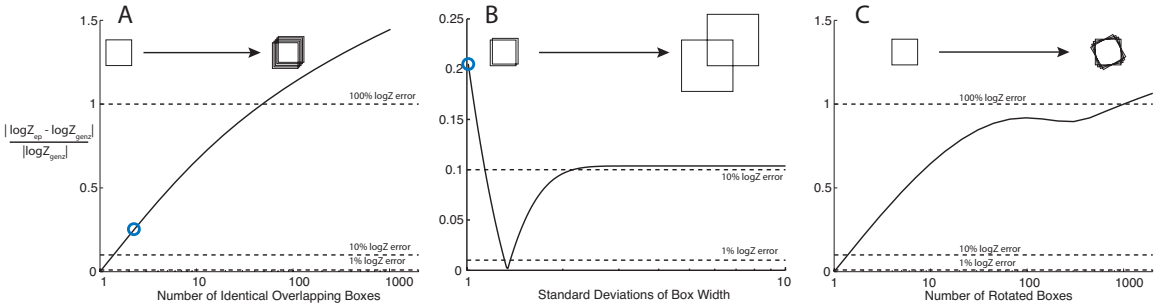

Figure 10: Empirical results from intentionally pathological cases that explore the shortcomings of the EP approach. All panels show the problem of integrating $\mathcal{N}(0, I)$ Gaussian in $n = 2$ dimensions over the $[-1, 1] \times [-1, 1]$ box. Panel depicts the error as repeated box factors are added to the problem. The true integral $F(\mathcal{A})$ is unchanged, but the EP error grows arbitrarily bad. Here EP consistently *underestimates* the true $F(\mathcal{A})$. Panel B depicts the error as extra factor mass is added. Though the integration region is unchanged (still the $[-1, 1] \times [-1, 1]$ box that is the intersection of larger boxes), the factors have extra mass truncated by other factors. EP can not ignore this extra mass and categorically *overestimates* the true probability. Here it is important to note that the error in Panel B does not start at 0 (when there is no extra mass), because that case corresponds to 2 repeated boxes. This case is shown as blue dots in Panels A and B to make that connection (the blue dots are the same problem). The error starts as an underestimate, and it crosses 0 error at the point at which the redundancy effect (underestimator from Panel A) and extra mass (overestimator from Panel B) compensate for each other. Both Panels A and B can be "fixed" by removing redundant or inactive constraints: minimalising the polyhedron as described in Appendix D. However, Panel C shows that these redundancy and extra mass problems can not always be corrected. Here we show rotated boxes (each rotated by $\pi/1000$ with respect to the previous box), which define the region uniquely and with no inactive constraints. There are still significant errors due to these effects. Panel C underestimates the probability, which is sensible because there is more redundancy than there is extra mass in this pathological case. Though all these examples are contrived, they illustrate serious empirical issues with EP that are not widely discussed.

have minimal average effect on the error. This lack of effect is indeed what we see in Panels C and D of Figure 9.

6.5 Contrived, pathological cases to illustrate shortcomings of EP

The above results already suggest that certain Gaussian probabilities may be ill-suited to EP. In particular, we gained intuition from viewing the problem in the whitened space, where the colinearity of the integration region can be seen to imply greater or lesser error when using EP. Now we want to use that intuition to intentionally create pathological cases for the EP method. Here we will introduce our terminology for the issues that cause EP to give poor results, which will provide a transition to a more theoretical discussion of the shortcomings of EP in the subsequent Discussion. Figure 10 introduces these pathological cases. In all cases, we want to integrate the $\mathcal{N}(0, I)$ Gaussian in $n = 2$ dimensions over

the $[-1, 1] \times [-1, 1]$ box, which is simply the product of two axis-aligned univariate box functions ($t_1(x_1) = \mathbb{I}\{-1 < x_1 < 1\}$ and $t_2(x_2) = \mathbb{I}\{-1 < x_2 < 1\}$). We note that this choice of Gaussian and integration region is convenient, but that the effects shown here exist in all integration regions that we have tested, to varying extents. This case is a common and representative case in terms of error magnitudes. The y-axis shows the errors as usual, and the x-axis is a feature chosen to create large errors in the EP method.

First, when looking at a transformed integration region, it seems plausible that some of the error from EP derives from the fact that approximate factors $\tilde{t}_i(\mathbf{x})$ are not orthogonal, and thus they may double count the mass in some way when forming the EP approximation $q(\mathbf{x})$. An extreme, contrived example of this would be adding two repeats of identical factors $t_i(\mathbf{x})$. Though the integration region is unchanged, EP is not invariant to this change, as it still must make an approximate factor for each true factor. This fact has been noted previously in Power-EP (Minka, 2004), a connection which we explore deeply in the Discussion. Panel A describes this pathological case, which we call the "redundancy" issue. As the cartoon in Panel A depicts, along the x-axis we use EP to solve the same probability with increasing numbers of redundant factors. We begin with a single box at left, where EP gets the exact answer as expected, since this problem decomposes to a product of univariate computations. Moving right, we add up to 1000 copies of the same integration region, and we see that the error grows to an answer that is arbitrarily bad. In mathematical terms, we go from using EP to solve $\int p_0(\mathbf{x})t(x_1)t(x_2)d\mathbf{x}$, to using EP to solve $\int p_0(\mathbf{x})t(x_1)^{1000}t(x_2)^{1000}d\mathbf{x}$. Though the true $F(\mathcal{A})$ is unchanged, EP returns a wildly different answer. More specifically, EP in this redundancy case is *underestimating* the true probability, which is expected and fully explained in the Discussion.

Second, by returning again to the transformed integration region, it also seems plausible that some of the error from EP derives from the approximate factors accounting for mass inappropriately. For example, when two $t_i(\mathbf{x})$ are not orthogonal, the moment matching of the cavity and the factor to the cavity and the approximate factor will include mass that is outside the integration region. Panel B of Figure 10 elucidates this concept. The integration region of interest here is still the $[-1, 1] \times [-1, 1]$ box, but that can be described by the intersection two boxes of any size, as the cartoon in Panel B shows. However, during the EP iterations, EP must consider the mass that exists in each true box factor *individually*. Because the Gaussian family has infinite support, there will always be nonzero mass outside the desired intersection that is the $[-1, 1] \times [-1, 1]$ box, and EP approximate factors will incorporate this mass. Hence we expect that EP will not be invariant to this change, and further that EP should overestimate the true probability when there is "extra mass," which is the term we use to describe this pathology of EP. Indeed Panel B shows precisely this. As extra mass is increased along the x-axis, error increases and overestimates the true probability. There are two interesting features here: first, note that the case where there is no extra mass - at left - is still a case where there are two redundant boxes, and hence we expect the underestimation error from Panel A (these corresponding points are circled in blue in Panels A and B). As extra mass is added, the EP result increases, the error crosses zero, and the EP result becomes a significant overestimate of the true probability. The second interesting feature is that the error levels off more quickly than in the redundancy case. This is not surprising, as the exponential decay of the Gaussian implies that the extra

mass issue can not arbitrarily punish EP - adding larger and larger boxes gives diminishing additional probability mass that EP must incorporate.

Finally, one might suggest that the problems of Panels A and B are entirely avoidable with appropriate preprocessing. Specifically, we can pull in all inactive polyhedral constraints, which would remove the extra mass issue of Panel B. Further, we can prune all redundant constraints, which would remove the redundancy issue of Panel A. We have discussed this previously in terms of a minimal representation of the polyhedron $\mathcal{A}$, and these steps are detailed in Appendix D. However, these trimming and pruning operations will not remove these issues in anything other than these specific cases, as can be seen in Panel C. Here we essentially replicate the result of Panel A. However, instead of precisely repeating the box, we add a sequence of slightly rotated boxes. Thus, this polyhedron is a minimal description of the integration region, and yet still we have identical issues as Panel A and B. One can imagine a series of heuristic choices to circumvent this issue, but the fact remains that EP will produce poor results for many simpler cases (see Figures 5 and 7, for example the trapezoid case), and thus we feel an investigation of these error modes is warranted and an interesting EP contribution.

From this figure the reader might correctly suggest that both the redundancy issue of Panel A and the extra mass issue of Panel B are two sides of the same coin: EP enforces a Gaussian approximation to a hard box-function factor, which is a highly non-Gaussian form. Thus redundancy and extra mass generally go hand in hand: if there are constraints that are not orthogonal, they will each have to consider both mass already considered by another factor (redundancy) and mass that lies outside the actual polyhedron (extra mass). Nonetheless, keeping these two ideas separate in these results will help us consider the theoretical underpinnings and point towards possible ways to improve the EP estimate of Gaussian probability and EP more generally, which is the subject of the Discussion section.

6.6 Final note: other moments

A novelty of our work here is that we have brought an approximate inference framework to bear on the problem of approximate integration. While the problem of calculating Gaussian probabilities - the zeroth moment of $p(\mathbf{x})$ - is interesting in its own right, many researchers will ask how well other moments, namely the "posterior" mean and covariance of the approximation $q(\mathbf{x})$ match the true moments. Throughout our results we calculated these moments with EP also. While demonstrating this claim rigorously is beyond the scope of this work (which focuses on the zeroth moment), we report anecdotally that we see similar accuracy and error regimes for the mean and covariance estimates from EP. As Genz methods do not return these moments, we calculated posterior means and covariances with elliptical slice sampling (Murray et al., 2010), which further complicates this claim due to the nontrivial error rate and computational burden of this sampling method. Nonetheless, to answer this common question and broaden the implications of these results beyond simply the zeroth moment, we note our finding that other EP moments follow similar trends.

7. Discussion

This paper has introduced a number of concepts and results around both Gaussian probabilities and Expectation Propagation. Here we discuss these findings.

7.1 EP allows approximate integration and can be used to calculate Gaussian probabilities

At the highest level, this paper introduces a novel approach to calculating multivariate Gaussian probabilities, which are a difficult and important problem throughout science and applied statistics. We used an approximate inference framework - Expectation Propagation - in an unexpected way to perform this Gaussian integration. From the perspective of Gaussian probabilities, this work introduces EPMGP as a method that can be used to give good, often highly accurate, numerical approximations. Our results show that under a variety of circumstances, relative EP errors rarely exceed 1%, and median errors are typically two to four orders of magnitude smaller. While we spent much of the results exploring the regions of poor performance (which we continue to explore below), the general strength of the method should not be neglected. The EPMGP algorithm also has some nice features such as analytical derivatives, attractive computational properties, and a natural ability to compute tail probabilities. On the other hand, existing numerical methods, in particular the Genz method, typically demonstrate high numerical accuracy and a clear runtime/accuracy tradeoff, which we have also demonstrated here. Furthermore, as we noted, one can readily construct cases wherein the EP method will perform poorly. Thus we again note that the purpose here is not to compare EPMGP to the Genz method, as those approaches share different properties, but rather to analyse the methods for Gaussian probability calculations across a range of scenarios so that applied researchers can make an informed and practical choice about which methods suit particular needs.

For machine learning researchers, another value of this work is the use of a popular approximate inference framework for an application other than approximate inference. Exploring the generality of the machine learning toolset for other purposes is important for understanding broader implications of our field's technologies.

Perhaps the most interesting aspect of this work is that the Gaussian probability problem, by its simple geometric interpretation, allows clear investigation into the strengths and weaknesses of EP. This work helps us explore empirically and theoretically the circumstances of and reasons behind EP's sometimes excellent and sometimes poor performance, which bears relevance to applications well beyond just Gaussian probabilities.

7.2 The strengths and weaknesses of EP: empirical evidence

Section 6 presents extensive empirical results of the performance of EP for the Gaussian probability problem. Figures 4 and 5 indicate that the EP algorithm works well for hyperrectangular regions but has considerably lower accuracy (one to two orders of magnitude) for general polyhedral regions. While this discrepancy may seem fundamental, we showed in Figure 3 that in fact hyperrectangular and polyhedral regions are closely related in a whitened space, where all but trivial cases are non-rectangular polyhedra. We have introduced the intuitive concepts of extra mass and redundancy, which suggest why EP should be error prone as the integration region (in the whitened space) gets further from rectangular. Figure 6 provided a first example of this feature. For a fixed dimensionality n, increasing the number of polyhedral faces m increases the redundancy, since each subsequently added constraint will have increasing colinearity with the existing constraint set. Panel A of Figure 6 shows this increasing error as a function of m. We also investigated

cases where the Gaussian problem has special structure, such as orthant and trapezoid integration regions in Figure 7 and equicorrelated Gaussians in Figure 9. Focusing now on the redundancy and extra mass issues that we have introduced, Panel B of Figure 7 and all of Figure 9 are particularly useful. They both demonstrate considerable EP error trend as the Gaussian moves away from white. Considered from the perspective of the space whitened with respect to the Gaussian, these figures demonstrate that EP error trends reliably with how non-rectangular or coloured the integration region is. We can concretely see the extra mass and redundancy issues causing errors in this view. Finally, Figure 10 presents pointed (albeit contrived) examples to clarify the effect of these issues.

In total, our empirical evidence suggests that EP can often be used to calculate Gaussian probabilities quickly and to high accuracy, and further that consideration of the Gaussian and integration region at hand can help predict when this computation will not be reliable. We introduced the intuitive concepts of redundancy and extra mass, which are helpful in understanding why problems come about. We attempt to condense these intuitions into a more rigorous theoretical understanding in the following section.

7.3 The strengths and weaknesses of EP: theoretical underpinnings

We now delve into a more theoretical investigation of EP, first by considering a factor graph representation, and second by considering the divergences that EP is trying to minimise in various problem settings.

7.3.1 Factorisation perspective

As we explained in section 3.5, we can equivalently represent any hyperrectangular or polyhedral integration region as a polyhedral region in the whitened space where the Gaussian has identity covariance. This standardisation enables us to focus on the geometry of the region. If this transformed region is actually axis-aligned, we obtain a fully factorised distribution. The factor graph representing this situation is shown in Panel A of Figure 11, where we use one node for each dimension to explicitly show the factorisation. Because the distribution is fully factorised, the integration also factorises as a product of univariate integrals, and so local moment matching in this case also yields global moment matching. Mathematically, the global KL-divergence can now be directly related to the local divergences such that $D_{\mathrm{KL}}(p\|q) = \sum_i D_{\mathrm{KL}}(t_i p'_{0i} \| \tilde{t}_i p'_{0i})$, where p'_{0i} is the ith dimension of the whitened Gaussian $p'_0(\mathbf{x})$ (standard univariate normals as in Figure 11). Accordingly, the KL minimisation problem separates, and EP will yield the exact zeroth moment (to within the machine precision of the error function). To be clear, EP is still making an approximation (the Gaussian approximation q), but we know that this q will have the exact zeroth, first, and second moments of p. And thus the Gaussian probability will be correct.

On the other hand, non axis-aligned constraints in the whitened space can be seen as linking multiple nodes together, thereby inducing dependencies in the factor graph and breaking the equivalence between local moment matching and global moment matching. These non axis-aligned constraints could either arise from the original integration region where a constraint links variables together, or from a non-diagonal Gaussian covariance which colours the integration region when space is whitened with respect to that covariance. We show an example of these dependencies in Panel B of Figure 11. This factor graph is

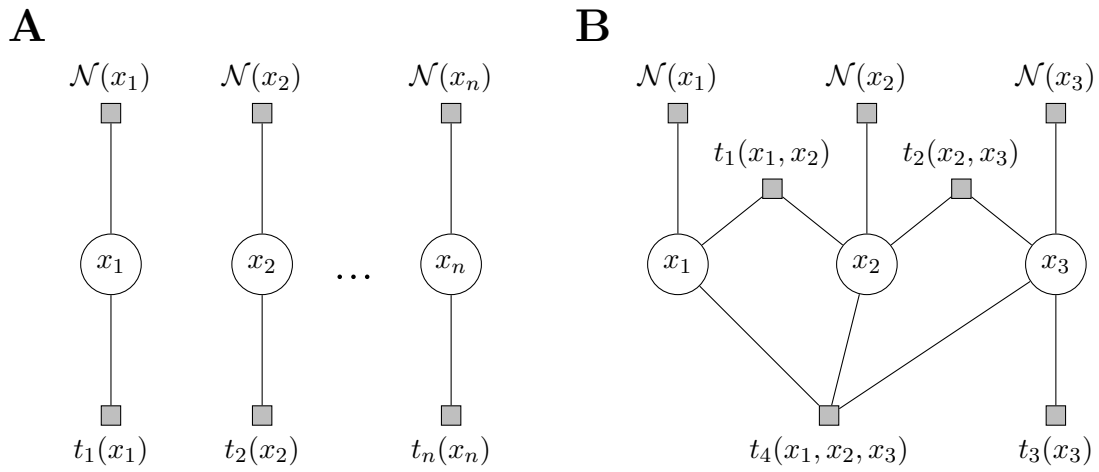

Figure 11: Factorisation perspective. Panel A represents the factor graph for an axis-aligned region in the whitened space ($\mathcal{N}(x_i)$ is a standard Gaussian on x_i, which defines the whitened prior $p'_0(\mathbf{x})$). Panel B represents an example of a region which is not axis-aligned in the whitened space: we show a polyhedral region with $n = 3$ which had a bidiagonal prior covariance structure, yielding dependent factors in the whitened space.

obtained by taking a bidiagonal prior covariance matrix on a polyhedral region for $n = 3$ and transforming it to the whitened space as in Figure 3. The polyhedral region is an axis-aligned hyperrectangle plus an additional factor constraint that links all dimensions. The covariance matrix induces pairwise factors in the whitened space, and the polyhedral constraint induces a factor on all variables. These non axis-aligned factors give rise to loops[3] in the factor graph which we expect to reduce the accuracy of EP, as seen in Figure 5 for polyhedral regions.

Little is formally known about the accuracy of EP on loopy graphs. In the case of Gaussian probabilities, fortunately we can use the geometry of the integration region to interpret the effect of additional graph edges. In particular, factors having several variables in common correspond to constraints which have a non-zero dot product (somewhat co-linear). Thus, these factors create approximate redundancy (Panel C of Figure 10) and extra mass (Panel B of the same figure). While these EP issues exist in all EP problem setups (not only Gaussian probabilities), the geometrical concepts available here provide a unique opportunity to understand the accuracy of EP which would not be apparent in standard factor graphs. We formalise these concepts in the next section.

3. We note in passing that the EP factor formulation that we are using is different than the standard one used on factor graphs which replaces each factor with a product of univariate approximate factors – see for example Section 10.7.2 of Bishop (2006) – and which reduces to loopy belief propagation in the case of discrete variables. Our formulation keeps the full approximate factor $\tilde{t}_i(\mathbf{x})$, much as when EP is applied on the Bayesian posterior of an i.i.d. model.

7.3.2 Alpha-divergence perspective

The redundancy and extra mass problems can be viewed in a unifying way as changing the effective *alpha-divergence* that EP is minimising, as we explain in this section. The alpha-divergence $D_\alpha(p \parallel q)$ is a generalisation of the KL-divergence. The relationship of alpha-divergences to EP and other variational schemes has been studied: Minka (2005) shows that several message passing algorithms on factor graphs can be interpreted as the local minimisation of different alpha-divergences. Following the notation from Minka (2005), $\mathrm{KL}(q \parallel p)$, the Hellinger distance, and $\mathrm{KL}(p \parallel q)$ corresponds to $\alpha = 0$, $\alpha = 0.5$, and $\alpha = 1$ respectively. Hence EP can be interpreted as the minimisation of local alpha-divergences with $\alpha = 1$.

The Power-EP algorithm (Minka, 2004)[4] is an extension of EP where the factors are raised to a specific power α_i before doing moment matching. One of the motivations for this algorithm is that raising the factor t_i to a power can make the projection more tractable. That paper also notes that running EP on a model where each factor t_i is repeated n_i times is equivalent to running Power-EP (with the factors unrepeated) with $\alpha_i = 1/n_i$. Importantly, the Power-EP algorithm can also be derived as the minimisation of local α_i-divergences.

This repetition of factors is precisely the "redundancy" issue that we created by repeating constraints n_i times (as in Panel A of Figure 10). Our EP algorithm in the pure redundant case can thus be interpreted as running Power-EP with $\alpha_i = 1/n_i$. This relationship is useful to us for four reasons.

First, the zeroth moment Z is always *underestimated* when doing global minimisation of *exclusive* alpha-divergence with $\alpha < 1$ (Minka, 2005). Also, Z is only correctly estimated for $\alpha = 1$ (KL-divergence), and it is *overestimated* for *inclusive* alpha-divergences which have $\alpha > 1$. This provides a theoretical explanation for the systematic underestimation of Z in the redundancy case of Panel A, Figure 10: EP in this case is the same as Power-EP with $\alpha_i = 1/n_i < 1$. As in regular EP, a subtle aspect here is that Power-EP is doing *local* alpha-divergence minimisation, whereas the overestimation/underestimation of the zeroth moment Z holds for *global* alpha-divergence minimisation, and the relationship between the two is not yet fully understood. However, our construction in Panel A of Figure 10 was a fully factorised case, so there is direct correspondence between global and local divergences, and as such we did systematically observe underestimation in these experiments. We also add that our experiments with Panel C of this figure and other similar cases indicate similar underestimation even when there is not direct global-local correspondence.

Second, this relationship between EP and alpha-divergence gives us a notion of *effective* alpha-divergence α_{eff} for each factor. In particular, running EP with the factor t_i repeated n_i times is the same as the local alpha-divergence minimisation by Power-EP with $\alpha_{\text{eff}} = 1/n_i$. When the constraints are almost-repeated (they have a large dot product, but are not fully colinear, such as in Panel C of Figure 10), we lose the exact correspondence between EP and Power-EP, but we could still imagine that EP would correspond to Power-EP with a suitable $\frac{1}{n_i} < \alpha_{\text{eff}} < 1$ in this case as well (using a continuity argument).

4. We warn the reader that the notation for alpha-divergence varies between papers (and in particular between Minka (2004) and Minka (2005)). We use the notation from Minka (2005) – the corresponding α' used in Minka (2004) is obtained as $\alpha' = 2\alpha - 1$.

Third, we can now think of *correcting* for this α_{eff} by using Power-EP with $\alpha' = 1/\alpha_{\text{eff}}$ instead of standard EP. In the case of exactly repeated constraints, this means using $\alpha'_i = n_i$. Indeed, our additional experiments (not shown) demonstrate that Power-EP does yield back the correct answer in this simple case. In other words, using this notion of correcting for α_{eff}, the errors of Panel A in Figure 10 disappear. In the case of almost-repeated constraints, we can still imagine correcting the effect by choosing a suitable α_{eff} correction term and using Power-EP. Our early investigation of this correction has given some promising results. An important and interesting direction of future work is characterising α_{eff} from geometric quantities such as dot products of constraints.

Fourth and finally, the extra mass issue can also be viewed in terms of alpha-divergences, though the connection is not as rigorous as the redundancy issue. The extra mass problem is one of inclusivity: the Gaussian approximate factor is including mass that it should not (as that mass lies outside the true polyhedron), which is a property of *inclusive* alpha-divergences as mentioned previously, *i.e.*, divergences with $\alpha > 1$. These we would expect to overestimate the true probability, and indeed this is what we see in Panel B of Figure 10. We can also consider correcting for these α_{eff} with Power EP. The situation is slightly more complicated here as the extra mass problem involves understanding how α_{eff} changes with the decay of the Gaussian distribution, but our initial attempts at these corrections indicate that the overestimation error can be significantly reduced.

To summarise, both the redundancy and extra mass problems can be cast in the language of effective alpha-divergence. We know that minimising D_α only returns moment matching when $\alpha = 1$ (KL-divergence), and thus any efforts to drive α_{eff} closer to 1 should improve all the performance results seen here. We have noted our success in removing errors entirely in the simplest cases such as Panel A of Figure 10 and in removing errors significantly in other slightly more complicated scenarios. Thus we feel that a judicious choice of α, to correct for the effective α in a Power-EP framework, is a fruitful direction of future work for both the theoretical understanding of and the empirical performance of EP across many applications.

7.4 Broader implications for EP

Having delved deeply into EP for Gaussian probabilities, it is worth stepping back to consider the breadth of these findings.

First we draw connections to other common machine learning applications. In the noise-free Bayes Point Machine model (Herbrich, 2002; Minka, 2001b), the model evidence is the integral of a white Gaussian over a polyhedral cone (intersection of halfspaces going through the origin), and so is a special case of our Gaussian probability framework. EP was analysed for the Bayes Point Machine in Minka (2001a,b) and was found to have good performance. On the other hand, its accuracy was only evaluated on a synthetic example with 3 points in 2 dimensions, which in our terminology would be $\{m = 3, n = 2\}$, which forms but a small part of our empirical analyses. Furthermore, prediction accuracy on a test set was used to evaluate EP on real world examples, but as ground truth was unavailable, the estimation accuracy remains unclear. Minka (2001a) observed our redundancy problem in a synthetic example: the author notes repetitions of the same datapoints (and so repeating the halfspace constraint in our formulation) yielded much poorer approximations of the

evidence or the mean. By comparing to numerical integration methods, we have provided a more extensive analysis of EP's accuracy for this problem. Further, we anticipate that a method for estimating α_{eff} (and using Power-EP) could improve the performance of Bayes Point Machines when the data is highly redundant.

Another important existing connection point is Gaussian Process classification (Rasmussen and Williams, 2006; Kuss and Rasmussen, 2005). When a probit likelihood (namely a univariate Gaussian cdf) is used, the data likelihood becomes a Gaussian probability problem. The same is true of probit regression (Ochi and Prentice, 1984), where changing the representation of the more standard prior and likelihood reveals a high-dimensional Gaussian probability. Though these problems are usually not considered as Gaussian probabilities, it is another point of broad importance for EP and for Gaussian probabilities. Furthermore, as EP is often used in a different way in these problems, exploring the similarities between EP for Gaussian probabilities and EP for problems that can be cast as Gaussian probabilities is an interesting subject for future investigation.

We have revealed some interesting empirical and theoretical properties for EP as applied to approximate Gaussian integration, but understanding how much of these issues stems from EP itself and how much from the application is an important clarification. Gaussian probabilities are a particularly instructive application, as they encourage the geometric intepretation that we have maintained throughout this paper. However, the issues of redundancy and extra mass are not specific to strict box functions, and the interpretation in terms of α_{eff} is no less applicable in the case of smooth functions. To test this question, we could repeat the setup of Figure 10 but using smooth functions such as probit curves. This problem would then bear similarity to Gaussian Process classification, and we could stress test the sensitivity of EP (in terms of extra mass and redundancy) to changing amounts of smoothness in the probit, where one end of this continuum is our Gaussian probability results. Though this remains a subject of interesting further work, we anticipate a result consistent with our findings here.

7.5 Future Directions

This work has explored both Gaussian probabilities and Expectation Propagation, and there are opportunities for valuable future work in both. We feel the most interesting direction for further investigation regards EP, specifically the theoretical discussion of alpha-divergences as in the section above. Any insights into understanding α_{eff} based on the geometry of the distributions involved, and further methods by which to correct for that α_{eff} via Power-EP approaches, could have broad implications for the continued success of EP as a powerful approximate inference framework. Here we have explored the geometrically and functionally simple case of box functions (which give Gaussian probablities), but future work should explore this for a variety of other factors (likelihood terms) as well.

The other focus for future work could be improvements to EPMGP specifically. For example, EP can often be improved by updating multiple approximate factors together. Here we considered a rank-one factor update, but also highly redundant factors could be considered in a rank-two update together. Because some of these low dimensional Gaussian probabilities can be calculated (such as orthants), these updates will be tractable in certain cases.

Finally, it is worth noting that our EP approach to approximate integration highlights the connections between approximate inference and integration, so some interesting connections here could be explored. For example, we noted that the Genz methods do not currently offer the ability to compute tail probabilities, calculate higher moments than the zeroth, or offer analytical derivatives. The absence of these features should not in principle be fundamental to the numerical integration approach, so reconsidering the Genz methods in the light of approximate inference could yield novel techniques for machine learning. Furthermore, hybrid approaches could be considered, drawing strengths from EP and numerical integration approaches.

8. Conclusion

We have presented an algorithm constructing analytic approximations to polyhedral integrals on multivariate Gaussian distributions. Such approximations are of value for a number of applications in many fields, particularly in hyperrectangular integration cases. We presented extensive numerical results for this case and the more general polyhedral case. We showed an interesting result that a small generalisation of the same algorithm is considerably less reliable. We explored the empirical and theoretical reasons behind these performance differences, which elucidated some features of EP not often discussed. These results also reveal Gaussian Probabilities as a useful testbed for analysing EP. As Expectation Propagation is becoming a more popular approximate inference method, this cautionary result suggests that EP should not be trusted blindly in new applications. Nonetheless, the uses of this EP framework will hopefully enable future work that uses the important object that is Gaussian probability.

Acknowledgments

We thank Tom Minka, Thore Graepel, Ralf Herbrich, Carl Rasmussen, Zoubin Ghahramani, Ed Snelson, and David Knowles for helpful discussions. JPC was supported by the UK Engineering and Physical Sciences Research Council (EPSRC EP/H019472/1); PH and SLJ were supported by a grant from Microsoft Research Ltd.

Appendix A. Gaussian EP cavity distributions for rank one factors $t_i(\mathbf{x})$

The purpose of this appendix is to derive the forms of the EP cavity distributions given in Equation 17. The only specific requirement is that the factors are rank one, that is: $t_i(\mathbf{x}) = t_i(\mathbf{c}_i^T\mathbf{x})$. We begin with the high dimensional Gaussian cavity distribution:

$$q^{\backslash i}(\mathbf{x}) \quad = \quad Z^{\backslash i}\mathcal{N}(\mathbf{x}; \mathbf{u}^{\backslash i}, V^{\backslash i}) \quad \propto \quad p_0(\mathbf{x})\prod_{j\neq i}\tilde{t}_j(\mathbf{x}) \tag{26}$$

where

$$V^{\backslash i} = \left(K^{-1} + \sum_{j\neq i}\frac{1}{\tilde{\sigma}_j^2}\mathbf{c}_j\mathbf{c}_j^T\right)^{-1} \tag{27}$$

and

$$\mathbf{u}^{\backslash i} = V^{\backslash i}\left(K^{-1}\mathbf{m} + \sum_{j\neq i}\frac{\tilde{\mu}_j}{\tilde{\sigma}_j^2}\mathbf{c}_j\right) \tag{28}$$

using standard properties of the normal distribution (*e.g.*, Rasmussen and Williams (2006)). Alternatively, as is often done for EP, we can consider subtracting the ith approximate factor from the overall posterior approximation $q(\mathbf{x})$. That is, since

$$q(\mathbf{x}) = Z\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \Sigma) \quad \propto \quad p_0(\mathbf{x})\prod_{j=1}^{m}\tilde{t}_j(\mathbf{x}), \tag{29}$$

we can say that $q^{\backslash i}(\mathbf{x})$ has parameters:

$$V^{\backslash i} = \left(\Sigma^{-1} - \frac{1}{\tilde{\sigma}_i^2}\mathbf{c}_i\mathbf{c}_i^T\right)^{-1} \tag{30}$$

and

$$\mathbf{u}^{\backslash i} = V^{\backslash i}\left(\Sigma^{-1}\boldsymbol{\mu} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\mathbf{c}_i\right). \tag{31}$$

Note that these steps are just like standard EP, save that we are subtracting rank one factors $\tilde{t}_i(\mathbf{x})$ instead of the more common axis-aligned factors. We will use this cavity to moment match $q^{\backslash i}(\mathbf{x})\tilde{t}_i(\mathbf{x})$ to $q^{\backslash i}(\mathbf{x})t_i(\mathbf{x})$. Since our $t_i(\mathbf{x})$ are rank one, we can equivalently marginalise out all other dimensions and moment match in a single dimension, which simplifies computation and presentation. This equivalence of marginal moment matching with the high-dimensional moment matching for EP requires that the exponential family used for the approximate factors is closed under marginalisation (which is the case for Gaussians), as noted in Section 3 of Seeger (2008). Accordingly, we want to form a univariate cavity distribution by marginalising over all dimensions orthogonal to the univariate factor of interest:

$$q_{\backslash i}(\mathbf{c}_i^T\mathbf{x}) = \int_{\backslash \mathbf{c}_i;\mathbf{x}} q^{\backslash i}(\mathbf{x}')d\mathbf{x}' \tag{32}$$

where the integration region $\backslash\mathbf{c}_i;\mathbf{x}$ indicates marginalisation over all dimensions orthogonal to $\mathbf{c}_i$ (an $n-1$ dimensional integral) and we move the cavity-defining "$\backslash i$" superscript down

to the subscript to indicate going from high dimension to single dimension. In standard EP, this notation would be $\backslash \mathbf{e}_i$ for the axis defined by the unit vector $\mathbf{e}_i$; in other words, we marginalise all variables but the ith. Mathematically, $\backslash \mathbf{c}_i; \mathbf{x}$ is the $n-1$ dimensional affine space perpendicular to $\mathbf{c}_i$ and defined by a particular choice of $\mathbf{c}_i^T \mathbf{x}$, that is $\backslash \mathbf{c}_i; \mathbf{x} = U_{\mathbf{c}_i} + (\mathbf{c}_i^T \mathbf{x})\mathbf{c}_i$ and $U_{\mathbf{c}_i} = \left\{ \mathbf{x}' \in \mathbb{R}^n \mid \mathbf{c}_i^T \mathbf{x}' = 0 \right\}$ (note that $\mathbf{c}_i$ is assumed to be normalised).

To do this marginalisation in closed form, we do a simple change of variable. We exploit the fact that, for any unit norm vector $\mathbf{c}_i$, there exists an orthonormal basis $\{\mathbf{c}_i, \mathbf{a}_2, ..., \mathbf{a}_n\}$ of $\mathbb{R}^n$. We define the orthogonal matrix

$$A = \begin{bmatrix} \mathbf{c}_i & \mathbf{a}_2 & \dots & \mathbf{a}_n \end{bmatrix}, \quad \text{with} \quad A^T A = A A^T = I, \tag{33}$$

and change the coordinates for our integration of $q^{\backslash i}(\mathbf{x})$ with the change of variable $\mathbf{y} = A^T \mathbf{x}'$. Since A is orthogonal, the Jacobian determinant is one. The Gaussian $q^{\backslash i}(\mathbf{x}') = Z^{\backslash i} \mathcal{N}(\mathbf{x}'; \mathbf{u}^{\backslash i}, V^{\backslash i})$ simply becomes $Z^{\backslash i} \mathcal{N}(\mathbf{y}; A^T \mathbf{u}^{\backslash i}, A^T V^{\backslash i} A)$, and so:

$$\begin{aligned} q_{\backslash i}(\mathbf{c}_i^T \mathbf{x}) &= \int_{\backslash \mathbf{c}_i; \mathbf{x}} q^{\backslash i}(\mathbf{x}') d\mathbf{x}' && (34) \\ &= \int_{\backslash \mathbf{e}_1; (\mathbf{c}^T \mathbf{x}) \mathbf{e}_1} Z^{\backslash i} \mathcal{N}(\mathbf{y}; A^T \mathbf{u}^{\backslash i}, A^T V^{\backslash i} A) d\mathbf{y} && (35) \\ &= Z^{\backslash i} \mathcal{N}(\mathbf{c}_i^T \mathbf{x}; \mathbf{c}_i^T \mathbf{u}^{\backslash i}, \mathbf{c}_i^T V^{\backslash i} \mathbf{c}_i) && (36) \end{aligned}$$

where the second to last equality follows from the fact that integrating out all dimensions orthogonal to $\mathbf{c}_i$ is equivalent to marginalising all but the first axis ($\mathbf{e}_1$ of $A^T \mathbf{x}$), by our definition of the matrix A. Since those orthogonal dimensions all normalise to 1, the high dimensional integral becomes univariate. The last equality follows again from the definition of A, and thus we are left with a simple univariate cavity distribution. By using the matrix inversion lemma and the definitions from Equations 30 and 31, we can derive simple forms for these cavity parameters:

$$\begin{aligned} \sigma_{\backslash i}^2 &= \mathbf{c}_i^T V^{\backslash i} \mathbf{c}_i && (37) \\ &= \mathbf{c}_i^T (\Sigma^{-1} - \frac{1}{\tilde{\sigma}_i^2} \mathbf{c}_i \mathbf{c}_i^T)^{-1} \mathbf{c}_i && (38) \\ &= \mathbf{c}_i^T (\Sigma + \Sigma \mathbf{c}_i (\tilde{\sigma}_i^2 - \mathbf{c}_i^T \Sigma \mathbf{c}_i)^{-1} \mathbf{c}_i^T \Sigma) \mathbf{c}_i && (39) \\ &= \mathbf{c}_i^T \Sigma \mathbf{c}_i \Big(1 + \frac{\mathbf{c}_i^T \Sigma \mathbf{c}_i}{\tilde{\sigma}_i^2 - \mathbf{c}_i^T \Sigma \mathbf{c}_i} \Big) && (40) \\ &= \frac{\tilde{\sigma}_i^2 \mathbf{c}_i^T \Sigma \mathbf{c}_i}{\tilde{\sigma}_i^2 - \mathbf{c}_i^T \Sigma \mathbf{c}_i} && (41) \\ &= \left((\mathbf{c}_i^T \Sigma \mathbf{c}_i)^{-1} - \tilde{\sigma}_i^{-2} \right)^{-1} && (42) \end{aligned}$$

which matches the form of Equation 17. So too, for the mean parameter of the univariate cavity distribution, we have:

$$
\begin{aligned}
\mu_{\backslash i} &= \mathbf{c}_i^T \mathbf{u}^{\backslash i} && (43)\\
&= \mathbf{c}_i^T V^{\backslash i}(\Sigma^{-1}\boldsymbol{\mu} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\mathbf{c}_i) && (44)\\
&= \mathbf{c}_i^T V^{\backslash i}\Sigma^{-1}\boldsymbol{\mu} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\sigma^2_{\backslash i} && (45)\\
&= \mathbf{c}_i^T(\Sigma + \Sigma\mathbf{c}_i(\tilde{\sigma}_i^2 - \mathbf{c}_i^T\Sigma\mathbf{c}_i)^{-1}\mathbf{c}_i^T\Sigma)\Sigma^{-1}\boldsymbol{\mu} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\sigma^2_{\backslash i} && (46)\\
&= \mathbf{c}_i^T\boldsymbol{\mu} + \mathbf{c}_i^T\Sigma\mathbf{c}_i(\tilde{\sigma}_i^2 - \mathbf{c}_i^T\Sigma\mathbf{c}_i)^{-1}\mathbf{c}_i^T\boldsymbol{\mu} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\sigma^2_{\backslash i} && (47)\\
&= \mathbf{c}_i^T\boldsymbol{\mu}\Big(1 + \frac{\mathbf{c}_i^T\Sigma\mathbf{c}_i}{\tilde{\sigma}_i^2 - \mathbf{c}_i^T\Sigma\mathbf{c}_i}\Big) - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\sigma^2_{\backslash i} && (48)\\
&= \frac{\mathbf{c}_i^T\boldsymbol{\mu}}{\mathbf{c}_i^T\Sigma\mathbf{c}_i}\sigma^2_{\backslash i} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\sigma^2_{\backslash i} && (49)\\
&= \sigma^2_{\backslash i}\Big(\frac{\mathbf{c}_i^T\boldsymbol{\mu}}{\mathbf{c}_i^T\Sigma\mathbf{c}_i} - \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\Big), && (50)
\end{aligned}
$$

which again matches the result of Equation 17. These cavity parameters $\{\mu_{\backslash i}, \sigma^2_{\backslash i}\}$ are intuitively sensible. Comparing to the standard EP cavity form (*e.g.*, Rasmussen and Williams (2006)), we see that instead of using the axis aligned univariate marginal defined by $\{\mu_i, \sigma_i^2\}$ (the ith entry and diagonal entry of μ and Σ, the parameters of the approximation $q(\mathbf{x})$), the cavities over arbitrary polyhedra require the univariate projection onto the axis $\mathbf{c}_i$.

Appendix B. Efficient and stable EPMGP implementation

EP gives us a simple form for calculating the parameters of the approximation $q(\mathbf{x})$ (from which we use the normalising constant Z as the Gaussian probability approximation). However, there are two potential problems that deserve investigation. First, each update of $\{\boldsymbol{\mu}, \Sigma\}$ naively requires $\mathcal{O}(n^3)$ computation, which may be needlessly burdensome in terms of run time and may introduce numerical instabilities. Second, as factor parameters $\tilde{\sigma}_i^2$ can tend to infinity (*i.e.*, that factor $t_i(\mathbf{x})$ has no effect on the MGP calculation), we want to reparametrise the calculation for numerical stability. This reparameterisation is also critical for calculating tail probabilities.

We introduce the natural parameters $\tilde{\tau}_i = \frac{1}{\tilde{\sigma}_i^2}$ and $\tilde{\nu}_i = \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}$, and we rewrite Equation 17 in terms of these parameters, as is standard for EP. In particular the cavity parameters are calculated as:

$$\mu_{\backslash i} \;=\; \sigma_{\backslash i}^2\Big(\frac{\mathbf{c}_i^T\boldsymbol{\mu}}{\mathbf{c}_i^T\Sigma\mathbf{c}_i} - \tilde{\nu}_i\Big) \quad \text{and} \quad \sigma_{\backslash i}^2 = \big((\mathbf{c}_i^T\Sigma\mathbf{c}_i)^{-1} - \tilde{\tau}_i\big)^{-1}. \tag{51}$$

We note that we do not reparameterise the cavity parameters (just the site parameters within that calculation), as the cavity parameters remain well behaved and numerically stable. Next, we rewrite the parameters of $q(\mathbf{x})$ from Equation 18 as

$$\boldsymbol{\mu} = \Sigma\Big(K^{-1}\mathbf{m} + \sum_{i=1}^m \tilde{\nu}_i\mathbf{c}_i\Big) \qquad \text{and} \qquad \Sigma = \Big(K^{-1} + \sum_{i=1}^m \tilde{\tau}_i\mathbf{c}_i\mathbf{c}_i^T\Big)^{-1} \tag{52}$$

Again these updates have cubic run time complexity. Let $\Delta\tilde{\tau}_i = \tilde{\tau}_i^{\text{new}} - \tilde{\tau}_i^{\text{old}}$. Now we consider updating Σ_{old} to Σ_{new} with an update $\Delta\tilde{\tau}_i$:

$$\begin{aligned}
\Sigma_{\text{new}} &= \Big(K^{-1} + \sum_{j\neq i}^m \tilde{\tau}_j^{\text{old}}\mathbf{c}_j\mathbf{c}_j^T + \tilde{\tau}_i^{\text{new}}\mathbf{c}_i\mathbf{c}_i^T\Big)^{-1} && (53)\\
&= \Big(K^{-1} + \sum_{i=1}^m \tilde{\tau}_i^{\text{old}}\mathbf{c}_i\mathbf{c}_i^T + \Delta\tilde{\tau}_i\mathbf{c}_i\mathbf{c}_i^T\Big)^{-1} && (54)\\
&= (\Sigma_{\text{old}}^{-1} + \Delta\tilde{\tau}_i\mathbf{c}_i\mathbf{c}_i^T)^{-1} && (55)\\
&= \Sigma_{\text{old}} - \Sigma_{\text{old}}\mathbf{c}_i\Big(\frac{1}{\Delta\tilde{\tau}_i} + \mathbf{c}_i^T\Sigma_{\text{old}}\mathbf{c}_i\Big)^{-1}\mathbf{c}_i^T\Sigma_{\text{old}} && (56)\\
&= \Sigma_{\text{old}} - \left(\frac{\Delta\tilde{\tau}_i}{1 + \Delta\tilde{\tau}_i\mathbf{c}_i^T\Sigma_{\text{old}}\mathbf{c}_i}\right)(\Sigma_{\text{old}}\mathbf{c}_i)(\Sigma_{\text{old}}\mathbf{c}_i)^T, && (57)
\end{aligned}$$

which is conveniently just a rank one (quadratic run time) update of the previous covariance. By a similar set of steps, we derive an update the mean of $q(\mathbf{x})$, which gives:

$$\boldsymbol{\mu}_{\text{new}} = \boldsymbol{\mu}_{\text{old}} + \left(\frac{\Delta\tilde{\nu}_i - \Delta\tilde{\tau}_i\mathbf{c}_i^T\boldsymbol{\mu}_{\text{old}}}{1 + \Delta\tilde{\tau}_i\mathbf{c}_i^T\Sigma_{\text{old}}\mathbf{c}_i}\right)\Sigma_{\text{old}}\mathbf{c}_i, \tag{58}$$

which again involves only quadratic computations and uses numerically stable natural parameters for the approximate factor $\tilde{t}_i(\mathbf{x})$ terms.

The final term that warrants consideration is the normalisation constant Z, which in fact is our approximation to the high dimensional Gaussian probability and the entire point of the EPMGP algorithm. Starting from Equation 19 and including the definition of $\tilde{Z}_i$ from Equation 16, that term is:

$$\begin{aligned}
\log Z &= -\frac{1}{2}\log|K| - \frac{1}{2}\sum_{i=1}^{m}\log\tilde{\sigma}_i^2 + \frac{1}{2}\log|\Sigma| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} + \frac{1}{2}\boldsymbol{\mu}^T\Sigma^{-1}\boldsymbol{\mu} - \frac{1}{2}\sum_{i=1}^{m}\frac{\tilde{\mu}_i^2}{\tilde{\sigma}_i^2} \\
&\quad + \sum_{i=1}^{m}\log\tilde{Z}_i - \frac{m}{2}\log(2\pi) \\
&= -\frac{1}{2}\log|K| - \frac{1}{2}\sum_{i=1}^{m}\log\tilde{\sigma}_i^2 + \frac{1}{2}\log|\Sigma| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} + \frac{1}{2}\boldsymbol{\mu}^T\Sigma^{-1}\boldsymbol{\mu} - \frac{1}{2}\sum_{i=1}^{m}\frac{\tilde{\mu}_i^2}{\tilde{\sigma}_i^2} \\
&\quad + \sum_{i=1}^{m}\log\hat{Z}_i + \frac{m}{2}\log(2\pi) + \frac{1}{2}\sum_{i=1}^{m}\log(\sigma_{\setminus i}^2 + \tilde{\sigma}_i^2) + \frac{1}{2}\sum_{i=1}^{m}\frac{(\mu_{\setminus i} - \tilde{\mu}_i)^2}{(\sigma_{\setminus i}^2 + \tilde{\sigma}_i^2)} - \frac{m}{2}\log(2\pi) \\
&= -\frac{1}{2}\log|K| - \frac{1}{2}\sum_{i=1}^{m}\log\tilde{\sigma}_i^2 + \frac{1}{2}\log|\Sigma| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} + \frac{1}{2}\boldsymbol{\mu}^T\Sigma^{-1}\boldsymbol{\mu} - \frac{1}{2}\sum_{i=1}^{m}\frac{\tilde{\mu}_i^2}{\tilde{\sigma}_i^2} \\
&\quad + \sum_{i=1}^{m}\log\hat{Z}_i + \frac{1}{2}\sum_{i=1}^{m}\log(\sigma_{\setminus i}^2 + \tilde{\sigma}_i^2) + \frac{1}{2}\sum_{i=1}^{m}\frac{(\mu_{\setminus i} - \tilde{\mu}_i)^2}{(\sigma_{\setminus i}^2 + \tilde{\sigma}_i^2)} \\
&= -\frac{1}{2}\log|K| + \frac{1}{2}\log|\Sigma| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} + \frac{1}{2}\boldsymbol{\mu}^T\Sigma^{-1}\boldsymbol{\mu} - \frac{1}{2}\sum_{i=1}^{m}\frac{\tilde{\mu}_i^2}{\tilde{\sigma}_i^2} \\
&\quad + \sum_{i=1}^{m}\log\hat{Z}_i + \frac{1}{2}\sum_{i=1}^{m}\log(1 + \tilde{\sigma}_i^{-2}\sigma_{\setminus i}^2) + \frac{1}{2}\sum_{i=1}^{m}\frac{(\mu_{\setminus i} - \tilde{\mu}_i)^2}{(\sigma_{\setminus i}^2 + \tilde{\sigma}_i^2)} \\
&= -\frac{1}{2}\log|K| + \frac{1}{2}\log|\Sigma| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} + \frac{1}{2}\boldsymbol{\mu}^T\Sigma^{-1}\boldsymbol{\mu} \\
&\quad + \sum_{i=1}^{m}\log\hat{Z}_i + \frac{1}{2}\sum_{i=1}^{m}\log(1 + \tilde{\sigma}_i^{-2}\sigma_{\setminus i}^2) + \frac{1}{2}\sum_{i=1}^{m}\frac{(\mu_{-i}^2 - 2\mu_{\setminus i}\tilde{\mu}_i - \tilde{\mu}_i^2\tilde{\sigma}_i^{-2}\sigma_{\setminus i}^2)\tilde{\sigma}_i^{-2}}{(1 + \tilde{\sigma}_i^{-2}\sigma_{\setminus i}^2)} \\
&= -\frac{1}{2}\log|K| + \frac{1}{2}\log|\Sigma| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} + \frac{1}{2}\boldsymbol{\mu}^T\Sigma^{-1}\boldsymbol{\mu} \\
&\quad + \sum_{i=1}^{m}\log\hat{Z}_i + \frac{1}{2}\sum_{i=1}^{m}\log(1 + \tilde{\tau}_i\sigma_{\setminus i}^2) + \frac{1}{2}\sum_{i=1}^{m}\frac{\mu_{\setminus i}^2\tilde{\tau}_i - 2\mu_{\setminus i}\tilde{\nu}_i - \tilde{\nu}_i^2\sigma_{\setminus i}^2}{(1 + \tilde{\tau}_i\sigma_{\setminus i}^2)}, \qquad (59)
\end{aligned}$$

where the second equation is just expanding the $\log\tilde{Z}_i$ terms, the fourth equation combines the second and eighth terms of the previous equation, the fifth equation combines the fifth and the eighth terms of the previous equation, and the final equation rewrites those in terms of natural parameters. These tedious but straightforward steps enable calculation of $\log Z$ using only the natural parameters of the approximate factors. Doing so yields a numerically stable implementation.

Finally, we can go a few steps further to simplify the run time and the numerical stability of calculating $\log Z$. Since we often want to calculate small probabilities with EPMGP, and

do so quickly, these extra steps are important. Noting the definition of Σ (as in Equation 18), we write:

$$\begin{aligned}
\log Z &= -\frac{1}{2}\log|K| + \frac{1}{2}\log|\Sigma| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} + \frac{1}{2}\boldsymbol{\mu}^T\Sigma^{-1}\boldsymbol{\mu} \\
&\quad + \sum_{i=1}^{m}\log\hat{Z}_i + \frac{1}{2}\sum_{i=1}^{m}\log(1+\tilde{\tau}_i\sigma^2_{\backslash i}) + \frac{1}{2}\sum_{i=1}^{m}\frac{\mu^2_{\backslash i}\tilde{\tau}_i - 2\mu_{\backslash i}\tilde{\nu}_i - \tilde{\nu}_i^2\sigma^2_{\backslash i}}{(1+\tilde{\tau}_i\sigma^2_{\backslash i})}, \\
&= -\frac{1}{2}\log|K| - \frac{1}{2}\log\Big|\Big(K^{-1} + \sum_{i=1}^{m}\tilde{\tau}_i\mathbf{c}_i\mathbf{c}_i^T\Big)\Big| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} + \frac{1}{2}\boldsymbol{\mu}^T\Big(K^{-1} + \sum_{i=1}^{m}\tilde{\tau}_i\mathbf{c}_i\mathbf{c}_i^T\Big)\boldsymbol{\mu} \\
&\quad + \sum_{i=1}^{m}\log\hat{Z}_i + \frac{1}{2}\sum_{i=1}^{m}\log(1+\tilde{\tau}_i\sigma^2_{\backslash i}) + \frac{1}{2}\sum_{i=1}^{m}\frac{\mu^2_{\backslash i}\tilde{\tau}_i - 2\mu_{\backslash i}\tilde{\nu}_i - \tilde{\nu}_i^2\sigma^2_{\backslash i}}{(1+\tilde{\tau}_i\sigma^2_{\backslash i})}, \\
&= -\frac{1}{2}\log\Big|\Big(I + \sum_{i=1}^{m}\tilde{\tau}_i(L^T\mathbf{c}_i)(L^T\mathbf{c}_i)^T\Big)\Big| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} + \frac{1}{2}\boldsymbol{\mu}^T K^{-1}\boldsymbol{\mu} + \frac{1}{2}\sum_{i=1}^{m}\tilde{\tau}_i(\mathbf{c}_i^T\boldsymbol{\mu})^2 \\
&\quad + \sum_{i=1}^{m}\log\hat{Z}_i + \frac{1}{2}\sum_{i=1}^{m}\log(1+\tilde{\tau}_i\sigma^2_{\backslash i}) + \frac{1}{2}\sum_{i=1}^{m}\frac{\mu^2_{\backslash i}\tilde{\tau}_i - 2\mu_{\backslash i}\tilde{\nu}_i - \tilde{\nu}_i^2\sigma^2_{\backslash i}}{(1+\tilde{\tau}_i\sigma^2_{\backslash i})}, \qquad (60)
\end{aligned}$$

where $K = LL^T$ is the Cholesky factorisation of the prior covariance K. We can precompute L at the beginning of the algorithm, and hence the terms $\mathbf{m}^T K^{-1}\mathbf{m}$ and $\boldsymbol{\mu}^T K^{-1}\boldsymbol{\mu}$ are also easy to solve. Then, all terms are highly numerically stable, and only that first log determinant has cubic run time. Importantly, we never have to invert or solve Σ, and we are able to operate solely in the space of natural parameters for the factor approximations.

To sum up, these steps give us fast and numerically stable implementations for calculating the updates to $q(\mathbf{x})$ and for calculating the final normalising constant, which is the result of our EPMGP algorithm. Our simple MATLAB implementation can be downloaded from `<url to come with publication>`.

Appendix C. Derivatives of the EP result $\log Z$

The main text has claimed that an advantage of EP methods is that they provide an analytical approximation to the Gaussian (log) probability $\log Z$ such that derivatives can be taken with respect to parameters of the Gaussian $p_0(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{m}, K)$. We produce those derivatives here as reference. Throughout applied uses of EP, derivatives with respect to $\log Z$ are often interesting, as $\log Z$ in a standard approximate inference model corresponds to the marginal likelihood of the observed data. One often wants to optimise (and thus take derivatives of) this quantity with respect to parameters of the prior $p_0(\mathbf{x})$. Two useful references for discussion and application of these derivatives are Seeger (2008); Rasmussen and Williams (2006). Most importantly, we leverage the result from Seeger (2008) (Section 2.1), where it is shown that the converged messages - the approximate factor parameters $\{\tilde{Z}_i, \tilde{\mu}_i, \tilde{\sigma}_i^2\}$ - have 0 derivative with respect to parameters of the prior $p_0(\mathbf{x})$. This consideration is sometimes described as a distinction between *explicit* dependencies (where $\log Z$ actually has the prior terms $\{\mathbf{m}, K\}$) and *implicit* dependencies where a message parameter such as $\tilde{Z}_i$ might depend implicitly on $\{\mathbf{m}, K\}$ (as in Rasmussen and Williams (2006)). The fact that EP converges to a fixed point of its objective function, and thus makes the latter implicit derivatives equal 0, simplifies our task considerably.

The simplest place to begin is with the original form of the normaliser in Equation 19. We note that the second line of that equation is entirely due to the approximate factors, so by Seeger (2008), we need only consider the first and third lines, yielding:

$$\log Z \quad = \quad \frac{1}{2}\boldsymbol{\mu}^T\Sigma^{-1}\boldsymbol{\mu} + \frac{1}{2}\log|\Sigma| - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m} - \frac{1}{2}\log|K| + \text{constants}. \tag{61}$$

Typically, these terms can be simplified into the familiar form of a normaliser of a product of Gaussians (*e.g.* A.7 in Rasmussen and Williams (2006)). However, because the factor contributions need not be full rank, that standard treatment requires a bit more detail. First, we revisit the definition of $\{\boldsymbol{\mu}, \Sigma\}$:

$$\boldsymbol{\mu} = \Sigma\Big(K^{-1}\mathbf{m} + \mathbf{r}\Big), \qquad \Sigma = \Big(K^{-1} + RR^T\Big)^{-1}, \tag{62}$$

where we have hidden several of the approximate factor parameters with:

$$\mathbf{r} = \sum_{i=1}^{m} \frac{\tilde{\mu}_i}{\tilde{\sigma}_i^2}\mathbf{c}_i, \qquad RR^T = \sum_{i=1}^{m} \frac{1}{\tilde{\sigma}_i^2}\mathbf{c}_i\mathbf{c}_i^T. \tag{63}$$

RR^T is a singular value decomposition of the matrix above (in general R will not be square). This simplification allows us to focus only on the parts of the equation that are dependent on $\{\mathbf{m}, K\}$. Thus we can rewrite $\log Z$ as:

$$\begin{aligned}
\log Z &= -\frac{1}{2}\log|K| + \frac{1}{2}\log|(K^{-1}+RR^T)^{-1}| + \text{constants} \qquad (64)\\
&\quad + \frac{1}{2}(K^{-1}\mathbf{m}+\mathbf{r})^T(K^{-1}+RR^T)^{-1}(K^{-1}\mathbf{m}+\mathbf{r}) - \frac{1}{2}\mathbf{m}^T K^{-1}\mathbf{m}\\
&= -\frac{1}{2}\log|(I+R^TKR)| + \text{constants} \qquad (65)\\
&\quad - \frac{1}{2}(\mathbf{m}-K\mathbf{r})^T\Big(R(I+R^TKR)^{-1}R^T\Big)(\mathbf{m}-K\mathbf{r}) + \mathbf{r}^T(\mathbf{m}+\frac{1}{2}K\mathbf{r})
\end{aligned}$$

where the second line comes from combining the two determinants and using Sylvester's determinant rule, and the third line comes from the typical algebra of using the matrix inversion lemma to rearrange these terms. Some readers may prefer the second or the third line; we use the third as it prevents any explicit inversions of K and allows significant computational savings if the number of columns of R is smaller than that of K.

With that setup, we consider two possible and common derivative quantities of interest. First and most simply, we may want the gradient with respect to the Gaussian mean $\mathbf{m}$:

$$\nabla_{\mathbf{m}}\log Z = -\Big(R(I+R^TKR)^{-1}R^T\Big)(\mathbf{m}-K\mathbf{r}) + \mathbf{r} \qquad (66)$$

This form has a closed form solution and is instructive for two reasons. First, we note that if RR^T is not full rank, then $\mathbf{m}$ is underdetermined. This is intuitively reasonable, as a subrank RR^T means that the polyhedron is open in some dimension (*e.g.* a slab), and thus there should be a linear subspace of optimal positions for the Gaussian. Second, this solution is interesting in its connections to centering algorithms as previously mentioned (see for example Appendix D). This optimised mean provides a center for a K shaped Gaussian within the polyhedron $\mathcal{A}$. Understanding the geometric implications of this centering approach is an interesting question, but well beyond the scope of this work.

Second, we can consider parameterised K. We presume there is some parameter θ (such as a kernel hyperparameter) on the covariance $K = K(\theta)$. Then we see:

$$\begin{aligned}
\frac{\partial \log Z}{\partial \theta} &= \frac{\partial}{\partial\theta}\Bigg[-\frac{1}{2}\log|(I+R^TKR)|\\
&\qquad - \frac{1}{2}(\mathbf{m}-K\mathbf{r})^T\Big(R(I+R^TKR)^{-1}R^T\Big)(\mathbf{m}-K\mathbf{r}) + \mathbf{r}^T(\mathbf{m}+\frac{1}{2}K\mathbf{r})\Bigg]\\
&= -\frac{1}{2}\text{trace}\left(R(I+R^TKR)^{-1}R^T\frac{\partial K}{\partial\theta}\right)\\
&\quad + \mathbf{r}^T\frac{\partial K}{\partial\theta}\Big(R(I+R^TKR)^{-1}R^T\Big)(\mathbf{m}-K\mathbf{r})\\
&\quad + \frac{1}{2}(\mathbf{m}-K\mathbf{r})^T\Big(R(I+R^TKR)^{-1}R^T\frac{\partial K}{\partial\theta}R(I+R^TKR)^{-1}R^T\Big)(\mathbf{m}-K\mathbf{r})\\
&\quad + \frac{1}{2}\mathbf{r}^T\frac{\partial K}{\partial\theta}\mathbf{r}. \qquad (67)
\end{aligned}$$

Within this form, the matrix $R(I + R^T K R)^{-1} R^T$ reverts to the familiar inverse sum of covariances when R is full rank. Thus, though this form may look different, it is simply a subrank version of the familiar product of Gaussians marginalisation constant as often used in marginal likelihood derivatives for EP.

Finally, we note that we have only focused on derivatives with respect to parameters of the Gaussian $p_0(\mathbf{x})$. One might also want to take parameters with respect to the region $\mathcal{A}$. The convenience of ignoring implicit dependencies would no longer hold in such a case, so taking these derivatives would involve taking more complicated derivatives of the approximate factors. As this is an aside to the current report, we leave that investigation for future work.

Appendix D. Minimalising polyhedral regions $\mathcal{A}$

Our box-function definition of polyhedra, or similarly the more conventional halfspace definition, does not in general uniquely specify a polyhedral region $\mathcal{A}$. For any polyhedron defined by $\mathcal{A} = \prod_i t_i(\mathbf{x})$ (our conventional definition throughout the paper), we can always add another constraint $t_j(\mathbf{x})$ such that $\mathcal{A} = \prod_i t_i(\mathbf{x}) = t_j(\mathbf{x}) \prod_i t_i(\mathbf{x})$. This redundant constraint could be colinear with an existing $t_i(\mathbf{x})$ factor that includes at least the $[l_i, u_i]$ interval, but it can also be any constraint such that the entire projection of $\mathcal{A}$ onto $\mathbf{c}_j$, the rank one dimension of $t_j(\mathbf{x})$, is contained inside $[l_j, u_j]$. Such factors are typically called inactive constraints.

Importantly, these inactive constraints do not change the polyhedron and therefore do not change the probability $F(\mathcal{A})$ of any Gaussian across this region. On the other hand, inactive constraints do change the representation and thus the EP approximation. Given our discussion of the extra mass and redundancy issues (Sections 6 and 7), removing these inactive constraints can potentially improve the accuracy of our EP algorithms. To do so, we define a minimal representation of any polyhedron $\mathcal{A}$ to be the polyhedron where all constraints are active and unique. This can be achieved by pulling in all inactive constraints until they support the polyhedron $\mathcal{A}$ and are thus active, and then we can prune all repeated/redundant constraints. These steps are solvable in polynomial time as a series of linear programs, and they bear important connection to problems considered in centering (analytical centering, Chebyshev centering, etc.) for cutting-plane methods and other techniques in convex optimisation (*e.g.*, see Boyd and Vandenberghe (2004)). The first problem one might consider is to find a point on the interior of the polyhedron, which can be solved by introducing a slack variable t into the linear program:

$$\begin{aligned} \underset{\{\mathbf{x},t\}}{\text{minimise}} \quad & t \\ \text{subject to} \quad & -\mathbf{c}_i^T\mathbf{x} + l_i < t \;\; \forall\, i = 1 \ldots m \\ & \mathbf{c}_i^T\mathbf{x} - u_i < t \;\; \forall\, i = 1 \ldots m. \end{aligned} \tag{68}$$

This program will either return an interior point $\mathbf{x}$ and a negative slack t or will return a positive t, which is a certificate that the polyhedron is in fact empty. Standard LP solvers can be used to find this result. Next, we can consider each factor $t_j(\mathbf{x})$ to test if it is inactive. First to consider the lower bound, for all j we solve:

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimise}} \quad & \mathbf{c}_j^T\mathbf{x} \\ \text{subject to} \quad & -\mathbf{c}_i^T\mathbf{x} + l_i < t \;\; \forall\, i = 1 \ldots m \\ & \mathbf{c}_i^T\mathbf{x} - u_i < t \;\; \forall\, i = 1 \ldots m. \end{aligned} \tag{69}$$

Having solved for the lower bounds, we then solve the upper bound problems for all j:

$$\begin{aligned} \underset{\mathbf{x}}{\text{maximise}} \quad & \mathbf{c}_j^T\mathbf{x} \\ \text{subject to} \quad & -\mathbf{c}_i^T\mathbf{x} + l_i < t \;\; \forall\, i = 1 \ldots m \\ & \mathbf{c}_i^T\mathbf{x} - u_i < t \;\; \forall\, i = 1 \ldots m. \end{aligned} \tag{70}$$

With the answers for these $2m$ LPs, we can now make all constraints active. If the solution to Equation 69 is larger than l_j, then we know that the lower bound l_j is inactive. In words, this means that the other constraints have cut off $\mathcal{A}$ inside of l_j in direction $\mathbf{c}_j$. Equivalently, if the solution to Equation 70 is smaller than the upper bound u_j, we know u_j is inactive. If both lower and upper bounds are found to be inactive, the factor $t_j(\mathbf{x})$ is entirely redundant and can be pruned from the representation without changing $\mathcal{A}$. Since we know redundancy can lead to EP underestimates, pruning can be useful (see for example Panel A of Figure 10). If only one of the lower or upper bounds is found to be inactive by Equation 69 or 70, then this bound can be tightened to *activate* it, at which point it will support the polyhedron $\mathcal{A}$. The value it should be tightened to is the result of the corresponding LP that found that bound to be inactive. This tightening operation would prevent errors such as those seen in Panel B of Figure 10. One small nuisance is simple repeated active factors, as each will technically appear active. However, this case can be dealt with by checking for such repetitions in preprocessing, and thus is not a difficulty in practice.

With these steps done, we have a minimal representation of the polyhedron $\mathcal{A}$. This set of operations is convenient because using it as a preprocessing step will ensure that each polyhedron is uniquely represented, and thus results from EP and other such approaches (as with centering algorithms) are consistent. However, in our results and our presentation we do not minimalise polyhedra to avoid adding additional complexity. Furthermore, as in Figure 10, doing so can mask other features of EP that we are interested in exploring. Anecdotally, we report that while minimalising polyhedra *will* change EP's answer, the effect is not particularly large in these randomly sampled polyhedral cases.

Appendix E. Minimising KL-divergence does moment matching

In this section, we review the equivalence between moment matching and minimising the unnormalised KL-divergence in an exponential family. This fact is often mentioned without proof in the literature (as for example in Minka (2005)) – we provide here the derivation for completeness. In particular, this implies that minimising the KL-divergence corresponds to matching the zeroth, first, and second moments of $q(\mathbf{x})$ to $p(\mathbf{x})$, when the approximating family is Gaussian. As approximately minimising the KL-divergence is the goal of EP, this result will show that EPMGP is an appropriate choice to solve the Gaussian probability calculation (zeroth moment). We provide the general derivation here.

First, as defined in the main text, $p(\mathbf{x})$ and $q(\mathbf{x})$ do not normalise to 1. Thus, we use the general definition of the KL-divergence for non-negative unnormalised distributions $p(\mathbf{x})$ and $q(\mathbf{x})$:

$$D_{KL}\big(p(\mathbf{x}) \parallel q(\mathbf{x})\big) = \int p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} d\mathbf{x} \quad + \quad \int q(\mathbf{x}) d\mathbf{x} \quad - \quad \int p(\mathbf{x}) d\mathbf{x}, \tag{71}$$

as in (Zhu and Rohwer, 1995; Minka, 2005). Note that the typically-seen normalised KL-divergence (*e.g.*, Cover and Thomas (1991)) is recovered when both $p(\mathbf{x})$ and $q(\mathbf{x})$ normalise to 1.

We are interested in the distribution $q(\mathbf{x})$ that minimises $D_{KL}\big(p(\mathbf{x}) \parallel q(\mathbf{x})\big)$ when $q(\mathbf{x})$ is restricted to be in an (unnormalised) exponential family, and where $p(\mathbf{x})$ could be arbitrary. The exponential families are flexible parametric families of distributions which contain many familiar families such as the Gaussian, Dirichlet, Poisson, etc. (Bishop, 2006; Seeger, 2003). Reusing the notation from Seeger (2003), an exponential family is defined by specifying the sufficient statistics vector function $\phi(\mathbf{x})$[5], and so the unnormalised $q(\mathbf{x})$ takes the form:

$$q(\mathbf{x}) = Z_q \exp\Big\{\boldsymbol{\theta}^T \phi(\mathbf{x}) - \Phi(\boldsymbol{\theta})\Big\}, \tag{72}$$

where $\Phi(\boldsymbol{\theta}) = \log \int \exp\Big\{\boldsymbol{\theta}^T \phi(\mathbf{x})\Big\} d\mathbf{x}$ is the log-partition function (the usual normaliser) and Z_q is an extra parameter accounting for the unnormalised $q(\mathbf{x})$. Substituting this form of $q(\mathbf{x})$ in (71), we obtain:

$$D_{KL}\big(p(\mathbf{x}) \parallel q(\mathbf{x})\big) = \int p(\mathbf{x}) \left(\Phi(\boldsymbol{\theta}) - \boldsymbol{\theta}^T \phi(\mathbf{x}) - \log Z_q\right) d\mathbf{x} + Z_q - \int p(\mathbf{x}) d\mathbf{x} - H[p] \tag{73}$$

$$= Z_p \Phi(\boldsymbol{\theta}) - \boldsymbol{\theta}^T \int \phi(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} - Z_p \log Z_q + (Z_q - Z_p) - H[p], \tag{74}$$

where $H[p] = -\int p(\mathbf{x}) \log p(\mathbf{x}) d\mathbf{x}$ is the entropy of p (a constant with respect to q), and we defined the normalisation constant of p as $Z_p = \int p(\mathbf{x}) d\mathbf{x}$. To minimise this KL, we first

5. Strictly speaking, the base measure μ for the density also needs to be specified (e.g. the counting measure for a Poisson distribution). For simplicity of presentation, we will show our derivations with the Lebesgue measure (appropriate for the Gaussian distribution e.g.), though the result holds for any exponential family in the natural parameterisation.

take the derivative with respect to the normaliser of $q(\mathbf{x})$, namely Z_q:

$$\frac{d}{dZ_q} D_{KL}\big(p(\mathbf{x}) \parallel q(\mathbf{x})\big) = -\frac{Z_p}{Z_q} + 1 = 0 \tag{75}$$

$$\implies Z_q^* = Z_p = \int p(\mathbf{x}) d\mathbf{x}. \tag{76}$$

The second derivative is $1/Z_q^2 > 0$, so this is the unique global minimum.

From this, we see that minimising the KL-divergence with respect to the normalising constant Z_q matches the zeroth moments, as we expected. Indeed, it is this equation that motivates the entire EPMGP approach. We obtain the optimal natural parameter $\boldsymbol{\theta}^*$ of $q(\mathbf{x})$ by setting its derivative of the KL to zero:

$$\frac{d}{d\boldsymbol{\theta}} D_{KL}\big(p(\mathbf{x}) \parallel q(\mathbf{x})\big) = Z_p \nabla_{\boldsymbol{\theta}} \Phi(\boldsymbol{\theta}) - \int p(\mathbf{x}) \phi(\mathbf{x}) d\mathbf{x} = \mathbf{0} \tag{77}$$

$$\implies \nabla_{\boldsymbol{\theta}} \Phi(\boldsymbol{\theta}) = \int \frac{p(\mathbf{x})}{Z_p} \phi(\mathbf{x}) d\mathbf{x}. \tag{78}$$

It is a standard fact of the exponential family that the gradient of the log-partition function is equal to the expected sufficient statistics; to see this, from the definition of $\Phi(\cdot)$ (after Equation 72 above), we have:

$$\nabla_{\boldsymbol{\theta}} \Phi(\boldsymbol{\theta}) = \frac{\int \phi(\mathbf{x}) \exp\Big\{\boldsymbol{\theta}^T \phi(\mathbf{x})\Big\} d\mathbf{x}}{\int \exp\Big\{\boldsymbol{\theta}^T \phi(\mathbf{x})\Big\} d\mathbf{x}} \tag{79}$$

$$= \int \phi(\mathbf{x}) \exp\Big\{\boldsymbol{\theta}^T \phi(\mathbf{x}) - \Phi(\boldsymbol{\theta})\Big\} d\mathbf{x} \tag{80}$$

$$= \int \frac{q(\mathbf{x})}{Z_q} \phi(\mathbf{x}) d\mathbf{x}. \tag{81}$$

Combining (81) with (78), we get that the expected sufficient statistics of q need to match those of p, namely, the higher-order moment matching condition, as we wanted to show. The second derivative with respect to $\boldsymbol{\theta}$ of the KL-divergence is the Hessian of the log-partition function, which is strictly positive definite if the exponential family is *minimal* (Wainwright and Jordan, 2008) and in this case gives a unique $\boldsymbol{\theta}^*$; otherwise, the moment matching condition still gives a global minimum, but there can be multiple parameters with this property. For the Gaussian family, the sufficient statistics are a vector of all elements x_i and all pairs of elements $x_i x_j$. This means that the moment matching condition is to match the zeroth, first and second order statistics.

To summarise, these final equations (76), (81) and (78) tell us that, to uniquely minimise the global KL-divergence between the truncated distribution $p(\mathbf{x})$ (or any target distribution) and the Gaussian $q(\mathbf{x})$, we do two things: first, we set Z to be the total mass (zeroth moment) of $p(\mathbf{x})$; and second, we set $\boldsymbol{\theta}$, the natural parameters of $q(\mathbf{x})$ such that the mean (first moment) and covariance (second moment) of $q(\mathbf{x})$ equal exactly the first and second moments of $p(\mathbf{x})$.

As it pertains to calculating Gaussian probabilities, minimising the KL-divergence calculates the zeroth moment of $p(\mathbf{x})$, which is exactly $F(\mathcal{A})$, the probability of interest. As such, EPMGP is a sensible choice for multivariate Gaussian probability calculations.

References

J. R. Ashford and R. R. Sowden. Multi-variate probit analysis. *Biometrics*, 26(3):535–546, 1970.

C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, New York, 2006.

R. D. Bock and R. D. Gibbons. High-dimensional multivariate probit analysis. *Biometrics*, 52(4):1183–1194, 1996.

S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004.

P. P. Boyle, Y. Lai, and K. S. Tan. Pricing options using lattice rules. *N. Amer. Actuarial Journal*, 9(3), 2005.

A. Buccianti, G. Nardi, R. Potenza (eds, M. L. Serre, P. Bogaert, and G. Christakos. Computational investigations of Bayesian maximum entropy spatiotemporal mapping. *Proc. of the Intl. Assoc. Math. Geology*, 1998.

W. J. Cody. Rational Chebyshev approximations for the error function. *Math. Comp.*, pages 631–638, 1969.

T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, New York, 1991.

R. Cranley and T. N. L. Patterson. Randomization of number theoretic methods for multiple integration. *SIAM J. Numer. Anal.*, 13:904–914, 1976.

R. N. Curnow and C. W. Dunnett. The numerical evaluation of certain multivariate normal integrals. *Annals of Mathematical Statistics*, 33(2):571–579, June 1962.

E. de Klerk, D. V. Pasechnik, and J. P. Warners. On approximate graph colouring and MAX-k-CUT algorithms based on the ϑ-function, 2000.

Z. Drezner. Computation of the trivariate normal integral. *Math. of Comp.*, 63:289–294, 1994.

Z. Drezner and G. O. Wesolowsky. On the computation of the bivariate normal integral. *J. Stat. Comp. Simul.*, 35:101–107, 1989.

H. I. Gassman, I. Deak, and T. Szantai. Computing multivariate normal probabilities: A new look. *J Comp Graph Stats*, 11(4):920–949, 2002.

A. Genz. Numerical computation of multivariate normal probabilities. *J. Comp. Graph. Stat.*, 1:141–149, 1992.

A. Genz. Numerical computation of rectangular bivariate and trivariate normal and t probabilities. *Stat. and Comput.*, 14:251–260, 2004.

A. Genz and F. Brentz. Numerical computation of multivariate t probabilities with application to power calculation of multiple contrasts. *J. Stat. Comp. Simul.*, 63:361–378, 1999.

A. Genz and F. Brentz. Comparison of methods for the computation of multivariate t probabilities. *J. Comp. Graph. Stats.*, 11:950–971, 2002.

A. Genz and K. S. Kwong. Numerical evaluation of singular multivariate normal distributions. *J. Stat. Comp. Simul.*, 68:1–21, 2000.

Alan Genz and Frank Bretz. *Information Theory, Inference, and Learning Algorithms.* Springer, 2009.

R. Herbrich. *Learning Kernel Classifiers.* MIT Press, 2002.

R. Herbrich, T. Minka, and T. Graepel. TrueSkillTM: A Bayesian skill rating system. *Advances in Neural Information Processing Systems*, 20:569–576, 2007.

F. J. Hickernell and H. S. Hong. The asymptotic efficiency of randomized nets for quadrature. *Math. Comp*, 68:767–791, 1999.

T. Hothorn, K. Hornik, A. Zeileis, and W. Wien. Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical Statistics*, 2006:651–674, 2005.

J. Huguenin, F. Pelgrin, and A. Holly. Estimation of multivariate probit models by exact maximum likelihood. Technical report, University of Lausanne, Institute of Health and Economics Management (IEMS), 2009.

J. W. Jawitz. Moments of truncated continuous univariate distributions. *Advances in Water Resources*, 27:269–281, 2004.

H. Joe. Approximations to multivariate normal rectangle probabilities based on conditional expectations. *J. Am. Stat. Assoc.*, 90(431):957–964, 1995.

D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques.* MIT Press, 2009.

M. Kuss and C.E. Rasmussen. Assessing approximate inference for binary Gaussian process classification. *Journal of Machine Learning Res.*, 6:1679–1704, 2005.

H. Lazard-Holly and A. Holly. Computation of the probability that a d-dimensional normal variable belongs to a polyhedral cone with arbitrary vertex. Technical report, Mimeo, 2003.

S. Lerman and C. Manski. On the use of simulated frequencies to approximate choice probabilities. *Structural analysis of discrete data with econometric applications*, 10:305–319, 1981.

E. Lesaffre and G. Molenberghs. Multivariate probit analysis: a neglected procedure in medical statistics. *Statistics in Medicine*, 10(9):1391–1403, 1991.

X. Liao, H. Li, and L. Carin. Quadratically gated mixture of experts for incomplete data classification. In *Proceedings of the 24th International Conference on Machine Learning*, 2007.

D. McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026, 1989.

T. P. Minka. *A family of algorithms for approximate Bayesian inference*. MIT Ph.D. Thesis, 2001a.

T. P. Minka. Expectation Propagation for approximate Bayesian inference. *Uncertainty in AI*, pages 362–369, 2001b.

T. P. Minka. Power EP. Technical report, Microsoft Research, 2004.

T. P. Minka. Divergence measures and message passing. *Microsoft Research Tech. Report MSR-TR-2005-173*, 2005.

T. P. Minka. EP: a quick reference. *Techincal Report*, 2008.

I Murray, R. P. Adams, and D. J. C. MacKay. Elliptical slice sampling. *Proc. 13th Intl Conference on AISTATS (AISTATS)*, 2010.

H. Niederreiter. On a number-theoretical integration method. *Aequationes Mathematicae*, 8:304–311, 1972.

D. Nuyens and R. Cools. Fast component-by-component construction, a reprise for different kernels. In H. Niederreiter and D. Talay, editors, *Monte Carlo and quasi-Monte Carlo methods*, pages 371–385. Springer-Verlag, 2004.

Y. Ochi and R. L. Prentice. Likelihood inference in a correlated probit regression model. *Biometrika*, 71(3):531–543, 1984.

M. Opper and O. Winther. From naive mean field theory to the TAP equations. In M. Opper and D. Saad, editors, *Advanced Mean Field Methods: Theory and Practice*. MIT Press, Cambridge, MA, 2001.

A. Pakes and D. Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057, 1989.

J. W. Pillow, L. Paninski, and E. P. Simoncelli. Maximum likelihood estimation of a stochastic integrate-and-fire neural model. *Neural Computation*, 16:2533–2561, 2004.

R. L. Plackett. A reduction formula for normal multivariate integrals. *Biometrika*, 41(3–4): 351–360, 1954.

C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, 2006.

M. Seeger. Bayesian Gaussian process models: PAC-Bayesian generalisation error bounds and sparse approximations. *Ph.D. Thesis, U. Edinburgh*, 2003.

M. Seeger. Expectation progapation for exponential families. *Technical Report, U.C. Berkeley*, 2008.

M. J. Shervish. Algorithm as 195: Multivariate normal probabilities with error bound. *Journal of the Royal Statistical Society Series C*, pages 81–94, 1984.

R.C. Sickles and P. Taubman. An analysis of the health and retirement status of the elderly. *Econometrica*, 54(6):1339–1356, 1986.

M. Sinn and K. Keller. Covariances of zero crossings in Gaussian processes. *Theory Prob and Appl*, 2010. To appear.

D. H. Stern, R. Herbrich, and T. Graepel. Matchbox: large scale online Bayesian recommendations. In *Proceedings of the 18th international conference on World Wide Web*, pages 111–120. ACM, 2009.

R. Thiebaut and H. Jacqmin-Gadda. Mixed models for longitudinal left-censored repeated measures. *Comput Methods Programs Biomed*, 74(3):255–260, 2004.

M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.

Y. Zhao, P. M. Grambsch, and J. D. Neaton. Comparison of numerical algorithms for bivariate sequential tests based on marginal criteria. *Comp. Stat. and Data Anal.*, 49: 631–641, 2005.

H. Zhu and R. Rohwer. Information geometric measurements of generalisation. *Aston Univ. Tech. Report NCRG/4350*, 1995.