

RESEARCH ARTICLE

Variational inference as iterative projection in a Bayesian Hilbert space with application to robotic state estimation

Timothy D. Barfoot*  and Gabriele M. T. D’Eleuterio

University of Toronto Institute for Aerospace Studies, 4925 Dufferin Street, Ontario, M3H 5T6, Canada

*Corresponding author. E-mail: tim.barfoot@utoronto.ca

Received: 18 May 2021; **Revised:** 12 September 2022; **Accepted:** 20 September 2022;

First published online: 24 October 2022

Keywords: Aitchison geometry, Bayesian Hilbert spaces, variational inference, Bayesian inference, compositional data, stochastic algebra

Abstract

Variational Bayesian inference is an important machine learning tool that finds application from statistics to robotics. The goal is to find an approximate probability density function (PDF) from a chosen family that is in some sense “closest” to the full Bayesian posterior. Closeness is typically defined through the selection of an appropriate loss functional such as the Kullback-Leibler (KL) divergence. In this paper, we explore a new formulation of variational inference by exploiting the fact that (most) PDFs are members of a Bayesian Hilbert space under careful definitions of vector addition, scalar multiplication, and an inner product. We show that, under the right conditions, variational inference based on KL divergence can amount to iterative projection, in the Euclidean sense, of the Bayesian posterior onto a subspace corresponding to the selected approximation family. We work through the details of this general framework for the specific case of the Gaussian approximation family and show the equivalence to another Gaussian variational inference approach. We furthermore discuss the implications for systems that exhibit sparsity, which is handled naturally in Bayesian space, and give an example of a high-dimensional robotic state estimation problem that can be handled as a result. We provide some preliminary examples of how the approach could be applied to non-Gaussian inference and discuss the limitations of the approach in detail to encourage follow-on work along these lines.

1. Introduction

In 1763, Richard Price published on behalf of his recently deceased friend, the Reverend Thomas Bayes, a paper that introduced what would become the atomic element of probabilistic inference: Bayes’ rule [16]. The paper though was widely ignored. About a decade later, the same rule was discovered by Pierre-Simon Laplace and, while Laplace laid its foundations, the theory of inference based on this rule became known as Bayesian inference. So confident was Laplace in the theory that he famously calculated the odds at 11,000 to 1 that the mass of Saturn as determined by a former student was correct to within 1%, 1,000,000-to-1 odds on the mass of Jupiter [[29], translated from 1825 French edition, pp. 46–47]. (Based on the most recent available data, he would have collected on the bet on Saturn.) Bayesian inference has been used in a great variety of applications from Henri Poincaré’s defense of Captain Dreyfus to Alan Turing’s breaking of the Enigma code [34]. In modern day, it provides the crucial framework for inference in such fields as statistics, decision theory, computational neuroscience, machine learning, computer vision, state estimation, and robotics.

The objective common to all these applications is the determination of a *posterior* probability to test some hypothesis or to calculate some estimate based on *prior* information and observed measurements. However, it is not always possible to find the posterior exactly. Indeed, we must often resort to approximate techniques. One such technique, which will occupy us here, is that of *variational inference* or *variational Bayes* [17]. In this variational approach, the goal is to find the probability density function

(PDF) that comes closest to the posterior as determined by minimizing a loss functional subject to the constraint that the distribution sought be drawn from a tractable class of densities, for example, where the posterior has to take the form of a Gaussian distribution. A common choice for the loss functional is the Kullback-Leibler (KL) divergence [5, 8, 17, 18, 20, 25, 27] although others such as Bregman [1, 36], Wasserstein [7], and Rényi divergences [30] have been used.

The field of variational inference based on the KL divergence is already well trodden although the research is hardly exhausted. The chosen class of densities from which the approximate posterior is to be shaped is key to variational inference. In the mean-field approximation, for example, the solution to the minimization of the divergence is constructed as a product of densities from a chosen family of admissible functions such as a Bayesian mixture of Gaussians [17]. Another possibility is using Bayesian mixtures of exponential families [5, 44]. A number of algorithms by which to execute the minimization exist including the variational EM algorithm, natural gradient descent, and Gaussian variational inference.

In [27], it is observed that “there is not as yet a systematic algebra that allows particular variational transformations to be matched optimally to particular graphical models.” While this was written two decades ago and specifically about graphical models, the remark finds resonance in the present work.

In previous work [14], we developed a practical robotic state estimation tool based on variational inference and compared it to maximum a posteriori (MAP), showing some advantages in certain situations. For example, the method we developed, dubbed Exactly Sparse Gaussian Variational Inference, can be used to solve the famous simultaneous localization and mapping (SLAM) problem. The current paper shows this existing method can be viewed through a different lens, that of iterative projections in a special space known as a *Bayesian Hilbert space* or *Bayes space* for short [43]. The primary contribution of this paper is therefore to make this connection between two quite different fields and hopefully to open the door to future extensions.

Our aim is to introduce a kind of information algebra to variational inference that not only provides a convenient and effective framework for analysis but also reveals key relationships to past work. This algebra has its origins in the work of [2] on *compositional data* in statistics. Compositional data can be represented on a simplex as with probability distributions for a finite set of discrete events. The resulting *Aitchison geometry* or *Aitchison simplex* establishes a vector space, in which vector addition is a normalized multiplication (perturbation) and scalar multiplication is a normalized exponentiation (powering). With an appropriate inner product, the set of PDFs over a finite discrete space was formalized as a Hilbert space by [37] and independently investigated by [9] and [12] in their *stochastic algebra*. The extension to continuous variables was first published by [23] and also studied by [13] for the case of finite domains. The generalization to include probabilities and measures on the whole real line was made by [42, 43], who introduced the term *Bayesian Hilbert space*.

In such a space, Bayes’ venerable rule becomes

$$p(\mathbf{x}|\mathbf{z}) = p(\mathbf{z}|\mathbf{x}) \oplus p(\mathbf{x}) \quad (1)$$

where \oplus indicates vector addition. (The normalization inherent in the operations accounts for the marginal $p(\mathbf{z})$ automatically.) Each new measurement made to refine the posterior becomes one more term added to the sum. It is this linear feature of a Bayesian Hilbert space that makes the structure ideally suited to variational inference.

The set of Gaussians, in an appropriately extended sense, constitutes a subspace of Bayes space as do exponential families. An arbitrary PDF in one of these subspaces can be expressed in the simple and usual manner as a linear combination of a basis for the subspace. The problem of variational inference can thus be expressed as the minimization of a divergence over a set of Fourier coefficients.

The linear-algebraic structure of these spaces affords us a new perspective and provides new insight. We show, for example, that the solution to variational inference based on the KL divergence can be viewed as an iterative projection, in the Euclidean sense, onto a given subspace. Indeed, this algorithm is essentially a Newton-like iteration scheme to solve for the minimum of the divergence, having a form identical to the natural-gradient-descent technique of [4]. Moreover, using a subspace of Gaussians reproduces the recent results of [14].

We also employ an information measure using a norm for Bayes space. This allows for a metric to be defined on the space, which can be interpreted as the distance between two PDFs. A (symmetric and quadratic) divergence between PDFs can be based on the distance metric. It is notable that each step in our iterative projection scheme is a local minimization of this divergence.

While this connection between variational inference and projection in Bayes space is exciting, there are still some open challenges around the approach. In the current formulation, there is no guarantee that our projection scheme will result in a valid PDF, although in practice we find the case to be so quite often, particularly with a good initial guess for the posterior estimate. Throughout we attempt to pinpoint the specific challenges and limitations of our approach so that improvements may follow in future work.

We shall begin with an overview of Bayesian Hilbert spaces in the next section. In Section 3, we discuss subspaces and bases, including exponentiated Hermite polynomials and Gaussian distributions. The variational inference problem for the KL divergence as viewed from the purchase of a Bayesian Hilbert space is considered in Section 4. The specific case of using a Gaussian subspace, that is, Gaussian variational inference, is treated in Section 5. Discussion is provided in Section 6 and we end with a few concluding remarks.

2. Bayesian Hilbert spaces

Let us consider some domain \mathcal{X} for our PDFs, for example, \mathbb{R}^N ; we shall refer to $\mathbf{x} \in \mathcal{X}$ as the *state*. A PDF $p(\mathbf{x})$ assigns a nonnegative, finite value to each element of \mathcal{X} such that

$$\int_{\mathcal{X}} p(\mathbf{x}) d\mathbf{x} = 1. \quad (2)$$

It turns out that this condition provides challenges when it comes to defining *Bayes space* on an infinite domain. As we will see, not all members of Bayes space (as we define it) will be PDFs and not all PDFs will be members of Bayes space; however, there is a large enough intersection between the two sets that Bayes space will be of practical use. Notationally, we will use $p(\mathbf{x})$ to mean a member of Bayes space throughout, indicating when the member is a valid PDF.

We provide a lightweight explanation of Bayes space, referring to [43] for more detail. We define the following set of functions:

$$\mathcal{B}^2 = \left\{ p(\mathbf{x}) = c \exp(-\phi(\mathbf{x})) \mid 0 < c < \infty, \int_{\mathcal{X}} \phi(\mathbf{x})^2 \nu(\mathbf{x}) d\mathbf{x} < \infty \right\}, \quad (3)$$

where $\nu(\mathbf{x})$ is an appropriate measure for \mathcal{X} (loosely, a weighting function); we will assume that $\nu(\mathbf{x})$ is in fact a PDF (and from \mathcal{B}^2) throughout although this is not necessary. Essentially, each member of \mathcal{B}^2 is an exponentiated function from \mathcal{L}^2 , the set of square-integrable functions under our chosen measure. Importantly, there is no requirement for $p(\mathbf{x}) \in \mathcal{B}^2$ to be a valid PDF; however, if we have that $c^{-1} = \int_{\mathcal{X}} \exp(-\phi(\mathbf{x})) d\mathbf{x}$, it will be so. Moreover, not all PDFs are members of \mathcal{B}^2 as we do not allow members to take on the value of zero anywhere in the domain¹, meaning only those PDFs that are strictly positive are contained (e.g., Gaussians and other exponential families).

We say that two members, $p_1(\mathbf{x}) = c_1 \exp(-\phi_1(\mathbf{x}))$, $p_2(\mathbf{x}) = c_2 \exp(-\phi_2(\mathbf{x})) \in \mathcal{B}^2$, are equivalent (equal) if and only if $\phi_1(\mathbf{x}) = \phi_2(\mathbf{x})$; in other words, the normalization constants, c_1 and c_2 , need not be the same. Under these conditions, we have that \mathcal{B}^2 is isomorphic to \mathcal{L}^2 .

We define *vector addition* [42], $\oplus : \mathcal{B}^2 \times \mathcal{B}^2 \rightarrow \mathcal{B}^2$, between two elements $p_1, p_2 \in \mathcal{B}^2$ to be $p_1 \oplus p_2$:

$$(\forall \mathbf{x} \in \mathcal{X}) \quad (p_1 \oplus p_2)(\mathbf{x}) = p_1(\mathbf{x})p_2(\mathbf{x}) = c_1 c_2 \exp(-(\phi_1(\mathbf{x}) + \phi_2(\mathbf{x}))) \in \mathcal{B}^2, \quad (4)$$

and likewise *scalar multiplication* [42], $\cdot : \mathbb{R} \times \mathcal{B}^2 \rightarrow \mathcal{B}^2$, of $p \in \mathcal{B}^2$ by $\alpha \in \mathbb{R}$ to be $\alpha \cdot p$:

$$(\forall \mathbf{x} \in \mathcal{X}) \quad (\alpha \cdot p)(\mathbf{x}) = (p(\mathbf{x}))^\alpha = c^\alpha \exp(-\alpha\phi(\mathbf{x})) \in \mathcal{B}^2. \quad (5)$$

¹[43] explain the details around letting members of Bayes space take the value zero; our more restrictive definition sidesteps some complications.

With these operations, \mathcal{B}^2 is established as a vector space, termed a *Bayesian linear space*, over the field \mathbb{R} [42]. Notably, the *zero vector*² is simply any constant function, $c \exp(0)$. Vector subtraction [42] is defined in the usual way, $p_1 \ominus p_2 = p_1 \oplus (-1) \cdot p_2$:

$$(\forall \mathbf{x} \in \mathcal{X}) \quad (p_1 \ominus p_2)(\mathbf{x}) = \frac{c_1}{c_2} \exp(-(\phi_1(\mathbf{x}) - \phi_2(\mathbf{x}))) \in \mathcal{B}^2. \tag{6}$$

We note that subtraction, or the inverse additive operation, is equivalent to the Radon-Nikodym derivative [42].

To turn a member of \mathcal{B}^2 into a valid PDF we define the *normalization operator*, $\downarrow p$:

$$(\forall \mathbf{x} \in \mathcal{X}) \quad (\downarrow p)(\mathbf{x}) = \frac{p(\mathbf{x})}{\int_{\mathcal{X}} p(\mathbf{z}) d\mathbf{z}} \in \mathcal{B}^2. \tag{7}$$

This operation can only be applied to those members of \mathcal{B}^2 that are equivalent to a valid PDF; in other words, it must be that $\int_{\mathcal{X}} p(\mathbf{x}) d\mathbf{x} < \infty$. We will refer to the *subset* of \mathcal{B}^2 whose members are equivalent to a valid PDF as $\downarrow \mathcal{B}^2 \subset \mathcal{B}^2$; note that this subset is not a subspace under our chosen addition and scalar multiplication operators. As a point of order, the normalization operator is not strictly required in the establishment of \mathcal{B}^2 , only when we want to make the connection to a valid PDF.

As mentioned above, Bayes' rule can be rendered as $p(\mathbf{x}|\mathbf{z}) = p(\mathbf{z}|\mathbf{x}) \oplus p(\mathbf{x})$. The normalizing marginal $p(\mathbf{z})$ is accounted for in the implied equivalence of the “=” operator. We could also write $p(\mathbf{x}|\mathbf{z}) = \downarrow(p(\mathbf{z}|\mathbf{x}) \oplus p(\mathbf{x}))$, which then makes the right-hand side a valid PDF through normalization.

2.1. Inner product

We endow the vector space with an *inner product* [43] defined as

$$\langle p_1, p_2 \rangle = \frac{1}{2} \int_{\mathcal{X}} \int_{\mathcal{X}} \ln\left(\frac{p_1(\mathbf{x})}{p_1(\mathbf{y})}\right) \ln\left(\frac{p_2(\mathbf{x})}{p_2(\mathbf{y})}\right) \nu(\mathbf{x}) \nu(\mathbf{y}) d\mathbf{x} d\mathbf{y}, \tag{8}$$

where $\nu(\cdot)$ is again a density function corresponding to an appropriate measure for \mathcal{X} . Notably, we see that because of the way the inner product is defined the normalization constants, c_1 and c_2 , associated with p_1 and p_2 play no role.

Because ν is a valid PDF, we can also write the inner product in (8) as

$$\langle p_1, p_2 \rangle = \mathbb{E}_{\nu}[\ln p_1 \ln p_2] - \mathbb{E}_{\nu}[\ln p_1] \mathbb{E}_{\nu}[\ln p_2], \tag{9}$$

where $\mathbb{E}_{\nu}[\cdot]$ is the expectation with respect to ν . To be clear, when we use expectations the argument, $f(\mathbf{x})$, and the measure (a PDF), $\nu(\mathbf{x})$, are defined over the same space, \mathcal{X} :

$$\mathbb{E}_{\nu(\mathbf{x})}[f(\mathbf{x})] = \int_{\mathcal{X}} f(\mathbf{x}) \nu(\mathbf{x}) d\mathbf{x}, \tag{10}$$

although sometimes we will abbreviate this as $\mathbb{E}_{\nu}[f]$. In this work, we shall always take the measure to be a PDF (and from \mathcal{B}^2); however, we shall refer to it as the measure to distinguish it from the other densities involved. Following [43], then, we can claim that \mathcal{B}^2 with inner product (8) forms a separable Hilbert space, which is referred to as a *Bayesian Hilbert space*. We shall sometimes briefly refer to it as a *Bayesian space* or *Bayes space*.

2.2. Information and divergence

The *norm* [43] of $p \in \mathcal{B}^2$ can be taken as $\|p\| = \langle p, p \rangle^{1/2}$. Accordingly, we can define the *distance* between two members of \mathcal{B}^2 , p and q , simply as $d(p, q) = \|p \ominus q\|$, which induces a metric on Bayes space.

²[43] make the point that the origin of Bayes space (i.e., the zero vector) can be shifted to be any valid member of \mathcal{B}^2 including a PDF, although we do not find this necessary here.

The norm of p can be used to express the information content of the PDF (if it is in \mathcal{B}^2). In fact, we shall define

$$I(p) = \frac{1}{2} \|p\|^2 = \frac{1}{2} \langle p, p \rangle \tag{11}$$

as the *information*³ in p . (The reason for the factor of $\frac{1}{2}$ will become evident.) As an example, consider $p = \mathcal{N}(\mu, \sigma^2)$ (over the domain \mathbb{R}) and measure $\nu = \mathcal{N}(0, 1)$. The information is $I(p) = (1 + 2\mu^2)/4\sigma^4$. The smaller the variance, the larger the information indicating that the PDF concentrates its probability mass more tightly about its mean; that is, we know better where to expect the state so we may say that we have more information about it.

We shall furthermore find it useful to define a *divergence* between two members of \mathcal{B}^2 , p and q , as

$$I(p \ominus q) = \frac{1}{2} \langle p \ominus q, p \ominus q \rangle. \tag{12}$$

This is the information contained in the difference of p and q . Unlike the KL divergence, this divergence is symmetric in p and q and quadratic in Bayesian space. Clearly, $p = q$ if and only if $I(p \ominus q) = 0$. Geometrically, the divergence is (half) the squared Euclidean distance between p and q in Bayes space.

2.3. Stochastic derivative

Accompanying this algebra is a functional calculus. Consider $p(\mathbf{x}|\theta) \in \mathcal{B}^2$ depending continuously on some parameter θ . We define the *stochastic partial derivative* of p with respect to θ as [13, 22]

$$\frac{\partial p}{\partial \theta} = \lim_{\lambda \rightarrow 0} \frac{1}{\lambda} \cdot (p(\mathbf{x}|\theta + \lambda) \ominus p(\mathbf{x}|\theta)). \tag{13}$$

Note that the result of this operation remains an element in \mathcal{B}^2 . We can also define directional derivatives and a gradient operator, but these will not be required here.

3. Subspaces and bases

While \mathcal{B}^2 is an infinite-dimensional space, it contains infinitely many finite-dimensional subspaces. We can in fact build a subspace \mathcal{Q} by taking the span of a set of M vectors $B = \{b_1, \dots, b_M\}$, namely,

$$\mathcal{Q} = \text{span} \{b_1, \dots, b_M\}. \tag{14}$$

If we choose B to be linearly independent, it will form a basis for \mathcal{Q} . We can accordingly write every vector q in \mathcal{Q} as a linear combination of B , that is,

$$q = \bigoplus_{m=1}^M \alpha_m \cdot b_m, \tag{15}$$

where $\alpha_m \in \mathbb{R}$ are unique. We use the notation $\bigoplus_{m=1}^M \alpha_m \cdot b_m$ to mean $\alpha_1 \cdot b_1 \oplus \dots \oplus \alpha_M \cdot b_M$, paralleling $\sum_{m=1}^M$ for normal addition.

As a shorthand, we will denote $\mathbf{b} = [b_1 \ b_2 \ \dots \ b_M]^T$ as the basis. The inner products between all pairs of basis vectors form the *Gram matrix*,

$$\langle \mathbf{b}, \mathbf{b} \rangle = [\langle b_m, b_n \rangle], \tag{16}$$

where (m, n) are the indices of the matrix entries. We furthermore have an orthonormal basis if $\langle b_m, b_n \rangle = \delta_{mn}$, the Kronecker delta, in which case $\langle \mathbf{b}, \mathbf{b} \rangle = \mathbf{1}$, the identity matrix.

³This is different than the information of [39].




Figure 1. Projection onto a subspace, \mathcal{Q} , of the Bayesian Hilbert space, \mathcal{B}^2 .

3.1. Projections

Given a subspace \mathcal{Q} of \mathcal{B}^2 and $p \in \mathcal{B}^2$, the $q^* \in \mathcal{Q}$ that minimizes the distance to, as well as the divergence (12) from, p is the projection of p onto \mathcal{Q} , that is,

$$q^* = \text{proj}_{\mathcal{Q}} p. \tag{17}$$

As in Euclidean geometry, we can view p as being decomposed into a component p_{\parallel} lying in \mathcal{Q} and a component p_{\perp} perpendicular to it; therefore $q^* = p_{\parallel}$ (see Fig. 1).

The coordinates of q^* can be calculated as

$$\alpha^* = \langle \mathbf{b}, \mathbf{b} \rangle^{-1} \langle \mathbf{b}, p \rangle. \tag{18}$$

We may also write the projection as an *outer product* operation on p , as detailed in Appendix B.

3.2. Example: one-dimensional Gaussian

To make the concept of Bayes space more tangible, consider the canonical one-dimensional Gaussian PDF defined over $x \in \mathbb{R}$:

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2} \frac{(x - \mu)^2}{\sigma^2}\right), \tag{19}$$

where μ is the mean and σ^2 the variance. In the language of \mathcal{B}^2 , we can write this as

$$p(x) = \underbrace{\left(-\frac{\mu}{\sigma^2}\right)}_{\alpha_1} \cdot \underbrace{\exp(-x)}_{b_1} \oplus \underbrace{\left(\frac{1}{\sqrt{2\sigma^2}}\right)}_{\alpha_2} \cdot \underbrace{\exp\left(-\frac{(x^2 - 1)}{\sqrt{2}}\right)}_{b_2} = \alpha_1 \cdot b_1 \oplus \alpha_2 \cdot b_2. \tag{20}$$

In other words, every Gaussian can be written as a linear combination of the two vectors, b_1 and b_2 , where the coefficients, α_1 and α_2 , depend on the mean and variance. Note, we can neglect the normalizing constant as equivalence is implied in the “=” operator.

The choice of b_1 and b_2 is not arbitrary in this example. They constitute the first two basis vectors in an orthonormal basis for \mathcal{B}^2 , which can be established using the *probabilist’s Hermite polynomials*; Appendix C.1 provides the details of this *Hermite basis*. In fact, we can define a new space \mathcal{G} as the span of these two basis vectors:

$$\mathcal{G} = \text{span} \{b_1, b_2\}, \tag{21}$$

which is a subspace of \mathcal{B}^2 . Importantly, every Gaussian PDF of the form in (19) is a member of \mathcal{G} , but not every member of \mathcal{G} is a valid Gaussian PDF. Only those members of \mathcal{G} that have $\sigma^2 > 0$ are valid Gaussian PDFs. We shall refer to \mathcal{G} as the *indefinite-Gaussian subspace* of \mathcal{B}^2 while $\downarrow\mathcal{G} \subset \mathcal{G}$ will denote the *positive-definite-Gaussian subset*. Figure 2 shows the relationships between the various spaces and how we can view a Gaussian as a point in \mathcal{G} .




Figure 2. On the left is a depiction of the relationships between Bayes space, \mathcal{B}^2 , and the set of all PDFs. We see the subset of strictly positive PDFs, $\downarrow\mathcal{B}^2$, the indefinite-Gaussian subspace, \mathcal{G} , and the positive-definite-Gaussian subset, $\downarrow\mathcal{G}$. The right image shows how a valid Gaussian PDF can be viewed as a point in a plane with coordinates that depend on its mean μ and variance σ^2 ; only the open upper-half plane admits valid Gaussian PDFs since we must have $\sigma^2 > 0$.

3.3. Example: projecting to a Gaussian

Let us consider a simple one-dimensional, nonlinear estimation problem as a numerical example motivated by the type of inverse-distance nonlinearity found in a stereo camera model. This same experiment (with the same parameter settings) was used as a running example by [[10], Section 4]. We assume that our true state is drawn from a Gaussian prior:

$$x \sim \mathcal{N}(\mu_p, \sigma_p^2). \tag{22}$$

We then generate a measurement according to

$$z = \frac{fb}{x} + n, \quad n \sim \mathcal{N}(0, \sigma_r^2), \tag{23}$$

where n is the measurement noise. The numerical values of the parameters used were

$$\begin{aligned} \mu_p &= 20 \text{ [m]}, & \sigma_p^2 &= 9 \text{ [m}^2\text{]}, \\ f &= 400 \text{ [pixel]}, & b &= 0.1 \text{ [m]}, & \sigma_r^2 &= 0.09 \text{ [pixel}^2\text{]}. \end{aligned} \tag{24}$$

The true posterior is given by

$$p(x|z) = \downarrow \exp(-\phi(x)), \quad \phi(x) = \underbrace{\frac{1}{2} \frac{(x - \mu_p)^2}{\sigma_p^2}}_{\text{prior}} + \underbrace{\frac{1}{2} \left(z - \frac{fb}{x}\right)^2}_{\text{measurement}}. \tag{25}$$

This problem can also be viewed as the correction step of the Bayes filter [26]: Start from a prior and correct it based on the latest (nonlinear) measurement.

We seek to find $q(x) \in \downarrow\mathcal{G}$ that is a good approximation to the true posterior $p(x|z)$. To do this, we will simply project the posterior onto the indefinite-Gaussian subspace, using the method described in Section 3.1, and then normalize the result. Figure 3 shows the results of doing this for two cases that differ only in the measure ν that we associate with Bayes space. The expectations used in the projections were computed with generic numerical integration although, as discussed by [14], there are several other options including Gaussian quadrature. In the top case, the measure is chosen as the prior estimate, while in the bottom case it is chosen to be closer to the posterior. In both cases, we see that our projection method produces a valid PDF, but in the bottom case the result is much closer to the true posterior. This simple example provides motivation for the main point of this paper, which is that to use the tools of Bayes space effectively, we will seek to iteratively update the measure used to carry out our projections such that we can best approximate a posterior. Intuitively, this makes sense since the measure is providing a weighting to different parts of \mathbb{R} so we would like to choose it to pay close attention where the posterior ends up.

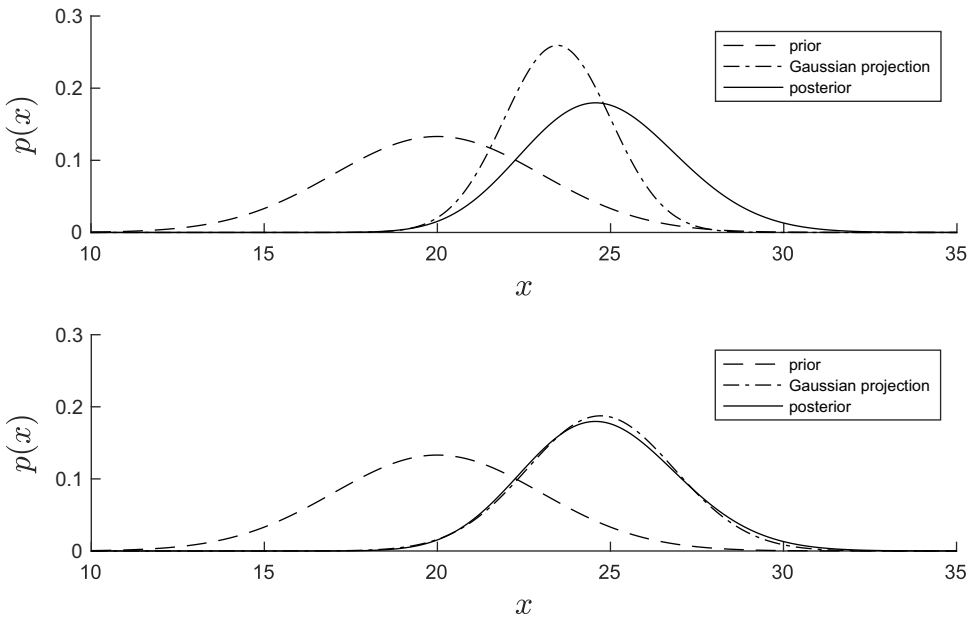


Figure 3. An example of projecting a non-Gaussian posterior onto the indefinite-Gaussian subspace. The top panel shows the case where the measure associated with \mathcal{B}^2 was chosen to be the same as the (Gaussian) prior, $v(x) = \mathcal{N}(20, 9)$. The bottom panel does the same with a Gaussian measure selected to be closer to the posterior, $v(x) = \mathcal{N}(24, 4)$. We see that the Gaussian projection of the posterior is much closer to the true posterior in the bottom case.

4. Variational Bayesian inference

Motivated by the example in Section 3.3, we shall now address the problem of variational Bayesian inference using the algebraic tools of Bayes space.

4.1. Variation on the KL divergence

In variational Bayesian inference, we seek to find an approximation, q , from some family of distributions constituting a subspace \mathcal{Q} to the true Bayesian posterior $p \in \mathcal{B}^2$. In general,

$$\mathcal{Q} \subseteq \mathcal{B}^2, \tag{26}$$

where equality will always ensure that $q = p$ will match the posterior exactly. But \mathcal{B}^2 is infinite-dimensional and, in practice, $\mathcal{Q} \subset \mathcal{B}^2$ is a finite-dimensional subspace.

There are many possible *divergences* that can be defined to characterize the “closeness” of q to p including the KL divergence [28], Bregman divergence [19], Wasserstein divergence/Earth mover’s distance [35], and Rényi divergence [38]. We shall focus on the KL divergence, which is defined as

$$\text{KL}(q||p) = - \int_{\mathcal{X}} q(\mathbf{x}) \ln \left(\frac{p(\mathbf{x}|\mathbf{z})}{q(\mathbf{x})} \right) d\mathbf{x} = -\mathbb{E}_q[\ln p - \ln q]. \tag{27}$$

Sometimes the reverse of this is used: $\text{KL}(p||q)$. Note, we show the divergence with respect to the posterior, $p(\mathbf{x}|\mathbf{z})$, but in practice during the calculations we use that $p(\mathbf{x}|\mathbf{z}) = p(\mathbf{x}, \mathbf{z})/p(\mathbf{z}) = \downarrow p(\mathbf{x}, \mathbf{z})$ since the joint likelihood is easy to construct and then the $p(\mathbf{z})$ can be dropped for it does not result in a KL term that depends on \mathbf{x} . We will generically use p in what follows to keep the notation clean.

4.2. KL gradient

We assume a basis $B = \{b_1, b_2 \dots b_M\}$ for \mathcal{Q} and we write q as

$$q = \downarrow \bigoplus_{m=1}^M \alpha_m \cdot b_m. \tag{28}$$

We desire to minimize the KL divergence with respect to the coordinates α_m . The gradient of $\text{KL}(q||p)$ can be computed as follows:

$$\frac{\partial \text{KL}}{\partial \alpha_n} = - \int_{\mathcal{X}} \left(\frac{\partial q}{\partial \alpha_n} (\ln p - \ln q) - q \frac{\partial \ln q}{\partial \alpha_n} \right) dx. \tag{29}$$

Exploiting (E4) and (E6), this reduces to

$$\frac{\partial \text{KL}}{\partial \alpha_n} = -\mathbb{E}_q[\ln b_n (\ln p - \ln q)] + \mathbb{E}_q[\ln b_n] \mathbb{E}_q[\ln p - \ln q] = - \langle b_n, p \ominus q \rangle_q \tag{30}$$

or, collecting these in matrix form,

$$\frac{\partial \text{KL}}{\partial \alpha^T} = - \langle \mathbf{b}, p \ominus q \rangle_q, \tag{31}$$

where $\mathbf{b} = [b_1 \ b_2 \ \dots \ b_M]^T$. Implicit in this statement is that when employed as the measure, we have normalized the current approximation, $\downarrow q^{(i)}$, since we always take the measure to be a valid PDF. The necessary condition for a minimum of the KL divergence is that the gradient is zero. Newton’s method suggests the manner in which we might iteratively solve for the optimal distribution. Following the established procedure, the iteration for the coordinates is given by

$$\alpha^{(i+1)} = \alpha^{(i)} + \mathbf{H}^{(i-1)} \langle \mathbf{b}, p \ominus q^{(i)} \rangle_{q^{(i)}}, \tag{32}$$

where \mathbf{H} is the Hessian of the KL divergence.

4.3. KL Hessian

The (m, n) entry of the Hessian is

$$\frac{\partial^2 \text{KL}}{\partial \alpha_m \partial \alpha_n} = - \frac{\partial}{\partial \alpha_m} \langle b_n, p \ominus q \rangle_q. \tag{33}$$

This differentiation must take into account the effect of the “measure” q . The product rule applies here and we can break down the differentiation as

$$\frac{\partial^2 \text{KL}}{\partial \alpha_m \partial \alpha_n} = - \left(\frac{\partial}{\partial \alpha_m} \langle b_n, p \ominus q \rangle \right)_q - \langle b_n, p \ominus q \rangle_{\partial q / \partial \alpha_m}, \tag{34}$$

the first term of which is to be read as the derivative of the inner product holding the measure fixed and the second of which deals with the derivative of the measure while holding the arguments of inner product fixed. The first term is

$$\left(\frac{\partial}{\partial \alpha_m} \langle b_n, p \ominus q \rangle \right)_q = \left(\frac{\partial}{\partial \alpha_m} \langle b_n, p \rangle - \frac{\partial}{\partial \alpha_m} \sum_k \alpha_k \langle b_n, b_k \rangle \right)_q = - \langle b_n, b_m \rangle_q = - \langle b_m, b_n \rangle_q. \tag{35}$$

As shown in Appendix E.2, the second becomes

$$\langle b_n, p \ominus q \rangle_{\partial q / \partial \alpha_m} = \left\langle b_n, \frac{\partial \ln q}{\partial \alpha_m} \cdot (p \ominus q) \right\rangle_q - \mathbb{E}_q[\ln p - \ln q] \langle b_m, b_n \rangle_q. \tag{36}$$

We advise that the coefficient $\partial \ln q / \partial \alpha_m$ of $p \ominus q$ is in fact a function of the state and as such cannot be transferred to the other argument of the inner product as would be possible for a scalar in the field \mathbb{R} .

We also recognize the factor of the last term as $\text{KL}(q||p)$. Therefore, substituting (35) and (36) into (34) yields

$$\frac{\partial^2 \text{KL}}{\partial \alpha_m \partial \alpha_n} = (1 - \text{KL}(q||p)) \langle b_m, b_n \rangle_q - \left\langle b_n, \frac{\partial \ln q}{\partial \alpha_m} \cdot (p \ominus q) \right\rangle_q. \tag{37}$$

We observe that the second term on the right-hand side is symmetric in the indices as the substitution of (E6) will attest. In matrix form, the Hessian is

$$\mathbf{H} = \frac{\partial^2 \text{KL}}{\partial \boldsymbol{\alpha}^T \partial \boldsymbol{\alpha}} = (1 - \text{KL}(q||p)) \mathbf{I}_\alpha - \left\langle \mathbf{b}, \frac{\partial \ln q}{\partial \boldsymbol{\alpha}^T} \cdot (p \ominus q) \right\rangle_q, \tag{38}$$

where \mathbf{I}_α is the *Fisher information matrix* (FIM) or *Gram matrix* and is described in detail in Appendix E.1. Newton’s method (32) can now be implemented. But the Hessian bears a closer look.

The Hessian can also be explicitly written as

$$\begin{aligned} \frac{\partial^2 \text{KL}}{\partial \alpha_m \partial \alpha_n} &= \langle b_m, b_n \rangle_q - \mathbb{E}_q[\ln b_m \ln b_n (\ln p - \ln q)] \\ &+ \mathbb{E}_q[\ln b_m \ln b_n] \mathbb{E}_q[\ln p - \ln q] + \mathbb{E}_q[\ln b_n] \mathbb{E}_q[\ln b_m (\ln p - \ln q)] \\ &+ \mathbb{E}_q[\ln b_m] \mathbb{E}_q[\ln b_n (\ln p - \ln q)] - 2 \mathbb{E}_q[\ln b_m] \mathbb{E}_q[\ln b_n] \mathbb{E}_q[\ln p - \ln q], \end{aligned} \tag{39}$$

the terms of which can be collected as

$$\frac{\partial^2 \text{KL}}{\partial \alpha_m \partial \alpha_n} = \langle b_m, b_n \rangle_q + \langle -b_{mn} \oplus \mathbb{E}_q[\ln b_n] \cdot b_m \oplus \mathbb{E}_q[\ln b_m] \cdot b_n, p \ominus q \rangle_q, \tag{40}$$

where $b_{mn} = \exp(\ln b_m \ln b_n)$. The symmetry in the Hessian is plainly evident in this version.

4.4. Iterative projection

In the vicinity of the optimal distribution, with a sufficiently large subspace \mathcal{Q} , we may expect $p \ominus q$ to be small almost everywhere. This makes all the terms in the Hessian of first order except \mathbf{I}_α , which is of zeroth order. The gradient (68) is also of first order. Thus to keep Newton’s descent to this order, we may approximate the Hessian as $\mathbf{H} \simeq \mathbf{I}_\alpha$ and the iterative procedure (32) becomes simply

$$\boldsymbol{\alpha}^{(i+1)} = \boldsymbol{\alpha}^{(i)} + \mathbf{I}_\alpha^{(i)-1} \langle \mathbf{b}, p \ominus q^{(i)} \rangle_{q^{(i)}}. \tag{41}$$

However, as $q^{(i)} = \downarrow \bigoplus_m \alpha_m^{(i)} \cdot b_m$,

$$\langle \mathbf{b}, p \ominus q^{(i)} \rangle_{q^{(i)}} = \langle \mathbf{b}, p \rangle_{q^{(i)}} - \langle \mathbf{b}, \mathbf{b} \rangle_{q^{(i)}} \boldsymbol{\alpha}^{(i)} = \langle \mathbf{b}, p \rangle_{q^{(i)}} - \mathbf{I}_\alpha^{(i)} \boldsymbol{\alpha}^{(i)}. \tag{42}$$

Hence, (41) becomes

$$\boldsymbol{\alpha}^{(i+1)} = \mathbf{I}_\alpha^{(i)-1} \langle \mathbf{b}, p \rangle_{q^{(i)}}. \tag{43}$$

The iterative update to q is $q^{(i+1)} = \downarrow \bigoplus_m \alpha_m^{(i+1)} \cdot b_m$. That is, the procedure can be viewed as an *iterative projection*,

$$q^{(i+1)} = \downarrow \underset{(\mathcal{Q}, q^{(i)})}{\text{proj}} p, \tag{44}$$

where we explicitly indicate that we normalize the result as the output of our algorithm should be a PDF. Figure 4 depicts the scheme. The procedure is essentially the application of Newton’s method on the KL divergence with the Hessian approximated as the FIM. This is precisely the approximation made in natural gradient descent [4]. In our Bayesian space, the operating point of the Newton step becomes the measure for the inner product. This highlights a key aspect of using the algebra associated with a Bayesian space. It recognizes the dual role of q : On the one hand, it is the approximating PDF, and on the other it serves as a measure that weights the difference between the approximation and the approximated.




Figure 4. Iterative projection onto a sequence of Bayesian Hilbert spaces, $(\mathcal{Q}, q^{(i)})$.

Convergence of iterative projection is guaranteed if the Hessian is positive definite. Provided that the subspace is large enough, we can expect convergence when we begin in a neighborhood of optimal q where the first-order terms in the Hessian are sufficiently small.

It is notable that each step of the iterative projection is equivalent to the local minimization of the divergence $I(p \ominus q)$ with the measure fixed at $q^{(i)}$ because

$$I(p \ominus (q^{(i)} \oplus \delta q)) = I(p \ominus q^{(i)}) + \delta \alpha^T \left(\frac{\partial I}{\partial \alpha^T} \right)_{q^{(i)}} + \frac{1}{2} \delta \alpha^T \left(\frac{\partial^2 I}{\partial \alpha^T \partial \alpha} \right)_{q^{(i)}} \delta \alpha, \tag{45}$$

where $\delta \alpha = \alpha^{(i+1)} - \alpha^{(i)}$ and

$$\left(\frac{\partial I}{\partial \alpha^T} \right)_{q^{(i)}} = -\langle \mathbf{b}, p \ominus q^{(i)} \rangle_{q^{(i)}}, \quad \left(\frac{\partial^2 I}{\partial \alpha^T \partial \alpha} \right)_{q^{(i)}} = \langle \mathbf{b}, \mathbf{b} \rangle_{q^{(i)}} \equiv \mathbf{I}_\alpha^{(i)}, \tag{46}$$

which are identical to the linearized forms for the KL divergence.

Throughout this section, we have assumed that the basis B remains constant across iterations, but this need not be the case. We may also choose to update the basis along with the measure to maintain, for example, orthonormality. This is explored in the next example and further in Section 5 on Gaussian variational inference.

4.5. Example: iteratively projecting to a Gaussian

In the example of Section 3.3, we saw that selecting a measure that was closer to the posterior resulted in a projection that was also closer to the posterior. We now redo this example using the iterative projection concepts from this section. We will still project onto the indefinite-Gaussian subspace and employ a Gaussian measure, only now with each iteration the measure will be taken to be the (normalized) projection from the previous iteration.

We initialized the estimate to the prior, which corresponds to the first panel in Fig. 5. The next three panels show subsequent iterations of the estimate. The last panel shows the KL divergence between the estimate and the true posterior for 10 iterations. We see the estimate converged in a few iterations.

Note that as the measure changes from one iteration to the next, we then have to update the basis to retain the desired orthogonality. This can be accomplished by using the reparameterization “trick” (see Appendix C.1) to adjust the basis to be orthogonal with respect to the current Gaussian measure.

4.6. Exploiting sparsity

One of the major advantages of thinking of B^2 as a vector space with the definition of vector addition \oplus is that Bayesian inference in general can be viewed as the addition of vectors. Consider the posterior $p(\mathbf{x}|\mathbf{z})$ where \mathbf{z} are some measurements. Bayes’ rule states that

$$p(\mathbf{x}|\mathbf{z}) = \frac{p(\mathbf{z}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{z})} = p(\mathbf{z}|\mathbf{x}) \oplus p(\mathbf{x}), \tag{47}$$

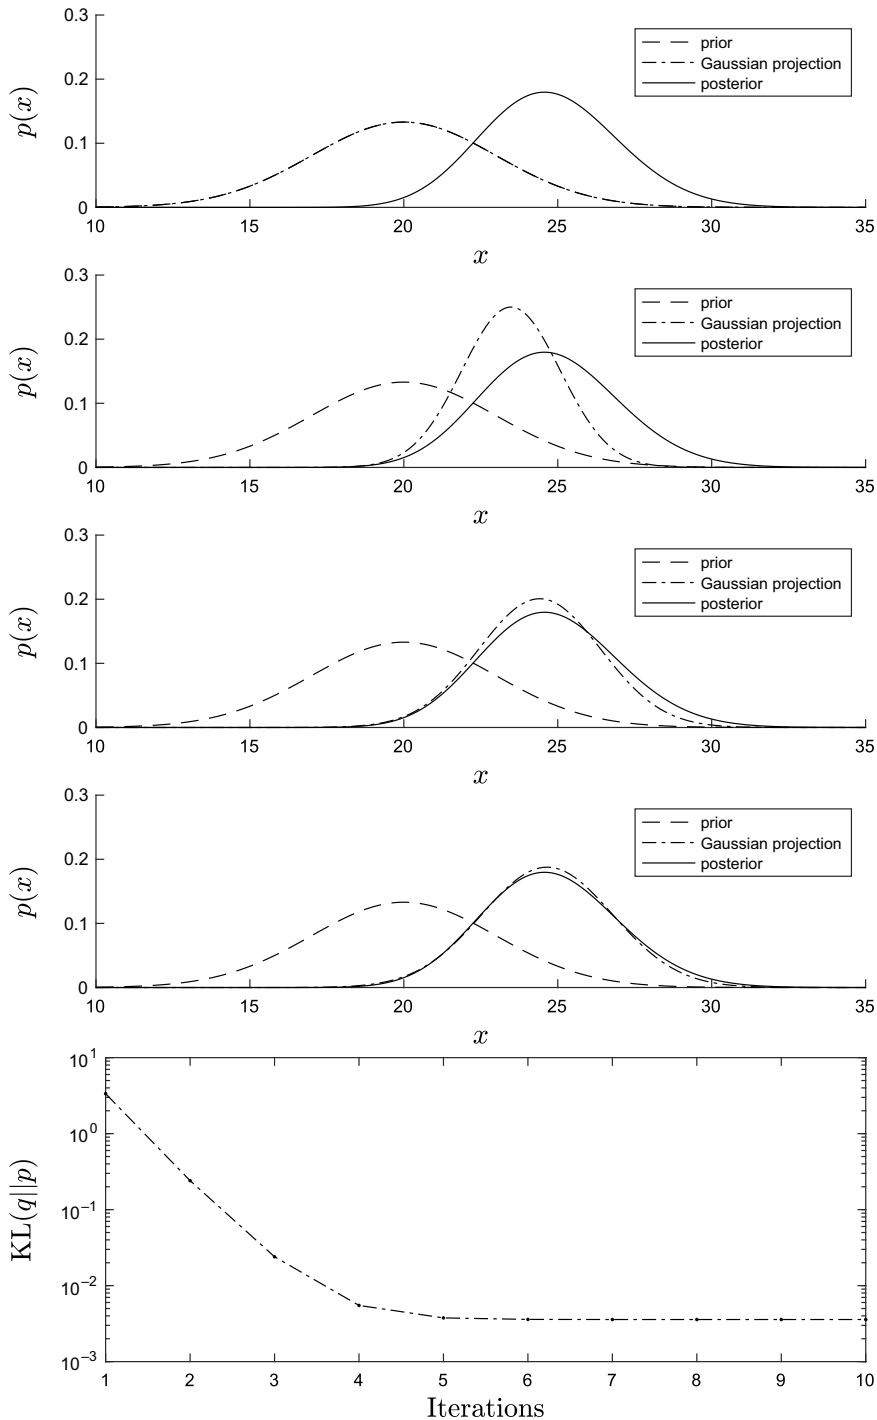


Figure 5. Example of iterative projection onto the indefinite-Gaussian subspace spanned by two Hermite basis functions, where the measure is taken to be the estimate $q^{(i)}$ at the previous iteration and the basis reorthogonalized at each iteration as described in Section 5. The estimate was initialized to the prior (first panel) and then iteratively updated (next three panels). The last panel shows the KL divergence between the estimate and the true posterior for 10 iterations, with convergence occurring at approximately 5 iterations.

where $p(\mathbf{x})$ is a prior, $p(\mathbf{z}|\mathbf{x})$ is a measurement factor and, as mentioned earlier, we need not introduce the normalization constant $p(\mathbf{z})$ explicitly when writing the posterior as a vector addition in Bayesian space. To be clear, addition is defined for two members of the same Bayes space and here we interpret $p(\mathbf{z}|\mathbf{x})$ as a function of \mathbf{x} since \mathbf{z} is a constant (e.g., a known measurement).

If we have several measurements that are statistically independent, then this can be factored as

$$p(\mathbf{x}|\mathbf{z}) = p(\mathbf{x}) \oplus \bigoplus_{k=1}^K p(\mathbf{z}_k|\mathbf{x}_k), \tag{48}$$

where $\mathbf{x}_k = \mathbf{P}_k \mathbf{x}$ is a subset of the variables in \mathbf{x} , \mathbf{P}_k is a projection matrix, and \mathbf{z}_k is the k th measurement. This expresses sparsity in the state description and in the measurements. To keep the notation economical, we shall simply write

$$p = \bigoplus_{k=0}^K p_k, \tag{49}$$

where p is the posterior and the p_k comprise the prior and the measurements, corresponding to statistically independent data. In other words, the factorization becomes a summation in the Bayesian space \mathcal{B}^2 .

Now consider our projective approach to inference. As usual, given a subspace $\mathcal{Q} \subset \mathcal{B}^2$, the optimal estimate to (49) is given by

$$q^* = \text{proj}_{\mathcal{Q}} p = \text{proj}_{\mathcal{Q}} \bigoplus_{k=0}^K p_k = \bigoplus_{k=0}^K \text{proj}_{\mathcal{Q}} p_k. \tag{50}$$

That is, the projection of the sum is the sum of the projections. Each individual projection can be done separately because we are in a linear space. This is of enormous practical advantage because it means that we do not need all of \mathcal{Q} to represent each projection.

We can see this more clearly by defining $\mathcal{B}_k^2 \subset \mathcal{B}^2$ as the subspace corresponding to the variables \mathbf{x}_k . Then

$$p = \bigoplus_{k=0}^K p_k \in \mathcal{B}_0^2 \oplus \mathcal{B}_1^2 \oplus \dots \oplus \mathcal{B}_K^2 \subseteq \mathcal{B}^2. \tag{51}$$

In other words, p is contained in the direct sum of the subspaces \mathcal{B}_k^2 . Each constituent part p_k may be confined to a smaller subspace of \mathcal{B}^2 , depending on the variable dependencies in each term.

If we wish to project $p_k \in \mathcal{B}_k^2$ onto \mathcal{Q} it will suffice to consider the projection on just $\mathcal{Q}_k = \mathcal{B}_k^2 \cap \mathcal{Q}$, that is,

$$\text{proj}_{\mathcal{Q}} p_k = \text{proj}_{\mathcal{Q}_k} p_k. \tag{52}$$

The subspace \mathcal{Q}_k may, and ideally would, be smaller than \mathcal{Q} . We may refer to \mathcal{Q}_k as the *marginal subspace* of \mathcal{Q} with respect to the subset of variables \mathbf{x}_k .

Therefore, the optimal estimate will be given by

$$q^* = \text{proj}_{\mathcal{Q}} p = \bigoplus_{k=0}^K \text{proj}_{\mathcal{Q}_k} p_k. \tag{53}$$

This means that we can project the PDF associated with each measurement onto a smaller subspace and simply add up the estimates, lifting the overall estimate up into a potentially much larger space. Naturally, when employed in practice we will normalize q^* to ensure our algorithm outputs a valid PDF. The decomposition and reconstitution are illustrated in Fig. 6. Just as with the total posterior, we may describe q^* as being an element of a direct sum of the individual subspaces of \mathcal{Q} , that is,

$$q^* \in \mathcal{Q}_0 \oplus \mathcal{Q}_1 \oplus \dots \oplus \mathcal{Q}_K \subseteq \mathcal{Q}. \tag{54}$$




Figure 6. Exploiting sparsity by projecting individual measurements onto marginal subspaces, \mathcal{Q}_k , and then recombining the results.

The subspace sum may be substantially smaller than \mathcal{Q} but again it will depend on the variable dependencies of each term.

This is the key result that allows most practical inference frameworks to function in a tractable way. Depending on the chosen basis for \mathcal{Q} , many of the coordinates can potentially be zero and thus it will not be necessary to waste effort computing them or space storing them.

5. Application: Iterative projection for multivariate Gaussians

Let us investigate a little more closely iterative projection to multivariate Gaussian PDFs, given their importance in statistics and estimation theory.

5.1. Projections

As mentioned at the end of the last section, we do not have to maintain the same basis from step to step as long as each basis spans the same subspace. This is a particularly useful maneuver when using the subspace \mathcal{G} of indefinite Gaussians, which are discussed in detail for the multivariate case in Appendix D.1. Denote the mean and variance of $q^{(i)} \in \mathcal{G}$ as $\mu^{(i)}$ and $\Sigma^{(i)}$ and let the basis $\mathbf{g}^{(i)}$ be defined as in (D3) and (D4). Note that this basis is orthonormal with respect to $q^{(i)}$. As such, $\mathbf{I}_\alpha^{(i)} = \langle \mathbf{g}^{(i)}, \mathbf{g}^{(i)} \rangle_{q^{(i)}} = \mathbf{1}$. Imagine the PDF to be approximated is expressed as $p = \downarrow \exp(-\phi(\mathbf{x})) \in \downarrow \mathcal{B}^2$. The coordinates resulting from the next projection are given by (D13), namely,

$$\alpha_1^{(i+1)} = \langle \mathbf{g}_1^{(i)}, p \rangle = \mathbf{L}^{(i)T} \mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right], \tag{55}$$

$$\alpha_2^{(i+1)} = \langle \mathbf{g}_2^{(i)}, p \rangle = \sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D}} \text{vech} \left(\mathbf{L}^{(i)T} \mathbb{E}_{q^{(i)}} \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \mathbf{L}^{(i)} \right),$$

where $\mathbf{L}^{(i)}$ issues from the Cholesky decomposition of $\Sigma^{(i)}$.

The new iteration is

$$q^{(i+1)} = \downarrow \text{proj}_{(\mathcal{G}, q^{(i)})} p = \downarrow \exp \left(-\alpha_1^{(i+1)T} \boldsymbol{\gamma}_1^{(i)} - \alpha_1^{(i+1)T} \boldsymbol{\gamma}_1^{(i)} \right). \tag{56}$$

Using (D14), this becomes

$$q^{(i+1)} = \downarrow \exp \left(-(\mathbf{x} - \boldsymbol{\mu}^{(i)})^T \mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right] - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}^{(i)})^T \mathbb{E}_{q^{(i)}} \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] (\mathbf{x} - \boldsymbol{\mu}^{(i)}) \right), \tag{57}$$

which we may cast into the form,

$$q^{(i+1)} = \downarrow \exp \left(-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}^{(i+1)})^T \boldsymbol{\Sigma}^{(i+1)-1} (\mathbf{x} - \boldsymbol{\mu}^{(i+1)}) \right). \tag{58}$$

Herein

$$\boldsymbol{\Sigma}^{(i+1)-1} = \mathbb{E}_{q^{(i)}} \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right], \tag{59a}$$

$$\boldsymbol{\Sigma}^{(i+1)-1} \delta \boldsymbol{\mu} = -\mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right], \tag{59b}$$

$$\boldsymbol{\mu}^{(i+1)} = \boldsymbol{\mu}^{(i)} + \delta \boldsymbol{\mu} \tag{59c}$$

give the updates from $q^{(i)} = \mathcal{N}(\boldsymbol{\mu}^{(i)}, \boldsymbol{\Sigma}^{(i)})$ to $q^{(i+1)} = \mathcal{N}(\boldsymbol{\mu}^{(i+1)}, \boldsymbol{\Sigma}^{(i+1)})$ and these are exactly the same as those used in the iterative Gaussian variational inference approach presented by [14]. We have arrived at the same variational updates but have done so from the framework of a Bayesian Hilbert space, where it becomes abundantly clear that the minimization algorithm is in fact a slightly simplified version of Newton’s method. This also provides the connection back to the classic Gaussian filtering and smoothing algorithms as discussed by [14].

5.2. Sparsity in Gaussian inference

The effect of sparsity as it applies to iterative Gaussian inference is of particular interest. Let us consider the decomposition of a posterior p in accordance to the general sparsity discussion in Section 4.6, that is,

$$p = \downarrow \exp(-\phi(\mathbf{x})) = \downarrow \exp \left(-\sum_{k=0}^K \phi_k(\mathbf{x}_k) \right) = \downarrow \bigoplus_{k=0}^K \exp(-\phi_k(\mathbf{x}_k)) = \downarrow \bigoplus_{k=0}^K p_k, \tag{60}$$

where $\phi_k(\mathbf{x}_k)$ is the k th (negative log) factor expression and $\mathbf{x}_k = \mathbf{P}_k \mathbf{x}$.

As in (D14), we may express the variational estimate as

$$q^{(i+1)} = \downarrow \text{proj}_{(\mathcal{G}, q^{(i)})} p = \downarrow \exp \left(-(\mathbf{x} - \boldsymbol{\mu}^{(i)})^T \mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right] - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}^{(i)})^T \mathbb{E}_{q^{(i)}} \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] (\mathbf{x} - \boldsymbol{\mu}^{(i)}) \right), \tag{61}$$

using the measure $q^{(i)} = \mathcal{N}(\boldsymbol{\mu}^{(i)}, \boldsymbol{\Sigma}^{(i)})$. To take advantage of sparsity, we need to have it reflected in the expectations herein. The first one leads to

$$\begin{aligned} \mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right] &= \sum_{k=0}^K \mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi_k(\mathbf{x}_k)}{\partial \mathbf{x}^T} \right] = \sum_{k=0}^K \left(\frac{\partial \mathbf{x}_k}{\partial \mathbf{x}} \right)^T \mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi_k(\mathbf{x}_k)}{\partial \mathbf{x}_k^T} \right] \\ &= \sum_{k=0}^K \mathbf{P}_k^T \mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi_k(\mathbf{x}_k)}{\partial \mathbf{x}_k^T} \right] = \sum_{k=0}^K \mathbf{P}_k^T \mathbb{E}_{q_k^{(i)}} \left[\frac{\partial \phi_k(\mathbf{x}_k)}{\partial \mathbf{x}_k^T} \right], \end{aligned} \tag{62}$$

given that $\mathbf{x}_k = \mathbf{P}_k \mathbf{x}$. For each factor k , then, we are able to shift the differentiation from \mathbf{x} to \mathbf{x}_k . We draw attention to the last equality, where the expectation simplifies to using $q_k^{(i)} = q_k^{(i)}(\mathbf{x}_k)$, the marginal of the

measure for just the variables in factor k . In a similar fashion,

$$\mathbb{E}_{q^{(i)}} \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] = \sum_{k=0}^K \mathbf{P}_k^T \mathbb{E}_{q_k^{(i)}} \left[\frac{\partial^2 \phi_k(\mathbf{x}_k)}{\partial \mathbf{x}_k^T \partial \mathbf{x}_k} \right] \mathbf{P}_k \tag{63}$$

accounts for the second expectation in (61).

The implication of the factorization is that each factor, identified by $\phi_k(\mathbf{x}_k)$, is projected onto \mathcal{G}_k , the marginal subspace associated with variables \mathbf{x}_k . The results can then be recombined for the full variational estimate as

$$q^{(i+1)} = \downarrow \text{proj}_{(\mathcal{G}, q^{(i)})} p = \downarrow \bigoplus_{k=0}^K \text{proj}_{(\mathcal{G}_k, q_k^{(i)})} p_k = \downarrow \bigoplus_{k=0}^K q_k^{(i+1)}. \tag{64}$$

The individual projections of $p_k = \downarrow \exp(-\phi_k(\mathbf{x}_k))$ onto $(\mathcal{G}_k, q^{(i)})$ are

$$q_k^{(i+1)} = \downarrow \text{proj}_{(\mathcal{G}_k, q_k^{(i)})} p_k = \downarrow \exp \left(-\frac{1}{2} (\mathbf{x}_k - \boldsymbol{\mu}_k^{(i+1)})^T \boldsymbol{\Sigma}_k^{(i+1)-1} (\mathbf{x}_k - \boldsymbol{\mu}_k^{(i+1)}) \right) \tag{65}$$

$$= \downarrow \exp \left(-\frac{1}{2} (\mathbf{x} - \mathbf{P}_k^T \boldsymbol{\mu}_k^{(i+1)})^T (\mathbf{P}_k^T \boldsymbol{\Sigma}_k^{(i+1)-1} \mathbf{P}_k) (\mathbf{x} - \mathbf{P}_k^T \boldsymbol{\mu}_k^{(i+1)}) \right), \tag{65}$$

where

$$\boldsymbol{\mu}_k^{(i+1)} = \boldsymbol{\mu}_k^{(i)} - \boldsymbol{\Sigma}_k^{(i+1)} \mathbb{E}_{q_k^{(i)}} \left[\frac{\partial \phi_k(\mathbf{x}_k)}{\partial \mathbf{x}_k^T} \right], \quad \boldsymbol{\Sigma}_k^{(i+1)-1} = \mathbb{E}_{q_k^{(i)}} \left[\frac{\partial^2 \phi_k(\mathbf{x}_k)}{\partial \mathbf{x}_k^T \partial \mathbf{x}_k} \right]. \tag{66}$$

It is straightforward to show that the vector sum of q_k from (65) reproduces (61). (Note that $\mathbf{P}_k \mathbf{P}_k^T = \mathbf{I}$ as \mathbf{P}_k is a projection matrix and \mathbf{P}_k^T the corresponding dilation.)

As explained in detail by [14], it would be too expensive for practical problems to construct first $\boldsymbol{\Sigma}^{(i)}$ and then extract the required blocks for the marginals, $q_k^{(i)} = \mathcal{N}(\boldsymbol{\mu}_k^{(i)}, \boldsymbol{\Sigma}_k^{(i)}) = \mathcal{N}(\mathbf{P}_k \boldsymbol{\mu}^{(i)}, \mathbf{P}_k \boldsymbol{\Sigma}^{(i)} \mathbf{P}_k^T)$. We see from the above development that we actually only require the blocks of $\boldsymbol{\Sigma}^{(i)}$ corresponding to the nonzero blocks of its inverse and the method of [41] can be used to extract the required blocks efficiently. In Ref. [14], the authors provide numerical experiments showing the efficacy of this approach.

5.3. Example: SLAM

A main purpose in the current paper was to show the connection between Bayes space [43] and Gaussian variational inference [14]. We see that minimizing the KL divergence between a true Bayesian posterior and an approximation can be viewed as iterative projection in Bayes space. Moreover, by exploiting the sparsity that Bayes space makes clear, the method has the potential to be applied to quite large inference problems.

Following [14], we consider a batch SLAM problem with a robot driving around and building a map of landmarks as depicted in Fig. 7. The robot is equipped with a laser rangefinder and wheel odometers and must estimate its own trajectory and the locations of a number of tubular landmarks. This dataset has been used previously by [15] to test SLAM algorithms. Groundtruth for both the robot trajectory and landmark positions (this is a unique aspect of this dataset) is provided by a Vicon motion capture system. The whole dataset is 12,000 timesteps long (approximately 20 min of real time). It was assumed that the data association (i.e., which measurement corresponds to which landmark) is known in this experiment to restrict testing to the state estimation part of the problem.

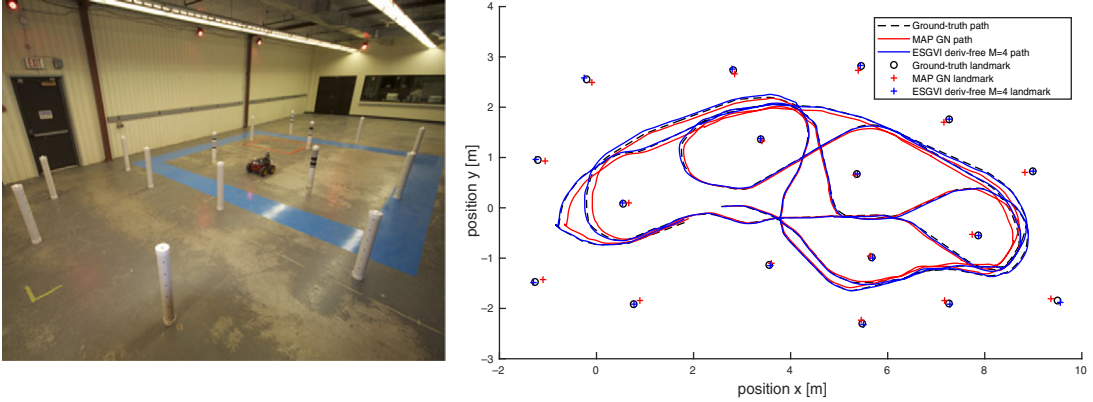


Figure 7. A special case of the approach presented in this paper was previously demonstrated [14] to be a useful and practical tool for robotic state estimation. The method, called exactly sparse Gaussian variational inference (ESGVI), was used to solve the simultaneous localization and mapping (SLAM) problem, outperforming the standard maximum a posteriori (MAP) estimation in certain cases. The current paper reinterprets this earlier work in a new mathematical formalism. Figure reproduced from [14].

The state to be estimated is

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_K \\ \mathbf{m}_1 \\ \vdots \\ \mathbf{m}_L \end{bmatrix}, \quad \mathbf{x}_k = \begin{bmatrix} x_k \\ y_k \\ \theta_k \\ \dot{x}_k \\ \dot{y}_k \\ \dot{\theta}_k \end{bmatrix}, \quad \mathbf{m}_\ell = \begin{bmatrix} x_\ell \\ y_\ell \end{bmatrix}, \tag{67}$$

where \mathbf{x}_k is a robot state (position, orientation, velocity, angular velocity) and \mathbf{m}_ℓ a landmark position.

For the (linear) prior factor on the robot states, we have

$$\varphi_k = \begin{cases} \frac{1}{2}(\mathbf{x}_0 - \check{\mathbf{x}}_0)^T \check{\mathbf{P}}^{-1} (\mathbf{x}_0 - \check{\mathbf{x}}_0) & k = 0 \\ \frac{1}{2}(\mathbf{x}_k - \mathbf{A}\mathbf{x}_{k-1})^T \mathbf{Q}^{-1} (\mathbf{x}_k - \mathbf{A}\mathbf{x}_{k-1}) & k > 0 \end{cases}, \tag{68}$$

with

$$\check{\mathbf{P}} = \text{diag}(\sigma_x^2, \sigma_y^2, \sigma_\theta^2, \sigma_{\dot{x}}^2, \sigma_{\dot{y}}^2, \sigma_{\dot{\theta}}^2), \quad \mathbf{A} = \begin{bmatrix} \mathbf{1} & T\mathbf{1} \\ \mathbf{0} & \mathbf{1} \end{bmatrix},$$

$$\mathbf{Q} = \begin{bmatrix} \frac{1}{3}T^3\mathbf{Q}_C & \frac{1}{2}T^2\mathbf{Q}_C \\ \frac{1}{2}T^2\mathbf{Q}_C & T\mathbf{Q}_C \end{bmatrix}, \quad \mathbf{Q}_C = \text{diag}(Q_{C,1}, Q_{C,2}, Q_{C,3}), \tag{69}$$

where T is the discrete-time sampling period, $Q_{C,i}$ are power spectral densities, and $\sigma_x^2, \sigma_y^2, \sigma_\theta^2, \sigma_{\dot{x}}^2, \sigma_{\dot{y}}^2, \sigma_{\dot{\theta}}^2$ are variances. The robot state prior encourages constant velocity [[10], §3, p.85].

The (nonlinear) odometry factors, derived from wheel encoder measurements, are

$$\psi_k = \frac{1}{2} (\mathbf{v}_k - \mathbf{C}_k \mathbf{x}_k)^T \mathbf{S}^{-1} (\mathbf{v}_k - \mathbf{C}_k \mathbf{x}_k), \tag{70}$$

where

$$\mathbf{v}_k = \begin{bmatrix} u_k \\ v_k \\ \omega_k \end{bmatrix}, \quad \mathbf{C}_k = \begin{bmatrix} 0 & 0 & 0 & \cos \theta_k & \sin \theta_k & 0 \\ 0 & 0 & 0 & -\sin \theta_k & \cos \theta_k & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{S} = \text{diag}(\sigma_u^2, \sigma_v^2, \sigma_\omega^2). \tag{71}$$

The \mathbf{v}_k consists of measured forward, lateral, and rotational speeds in the robot frame, derived from wheel encoders; we set $v_k = 0$, which enforces the nonholonomy of the wheels as a soft constraint. The σ_u^2, σ_v^2 , and σ_ω^2 are measurement noise variances.

The (nonlinear) bearing measurement factors, derived from a laser rangefinder, are

$$\psi_{\ell,k} = \frac{1}{2} \frac{(\beta_{\ell,k} - g(\mathbf{m}_\ell, \mathbf{x}_k))^2}{\sigma_r^2}, \tag{72}$$

with

$$g(\mathbf{m}_\ell, \mathbf{x}_k) = \tan^{-1}(y_\ell - y_k - d \sin \theta_k, x_\ell - x_k - d \cos \theta_k) - \theta_k, \tag{73}$$

where $\beta_{\ell,k}$ is a bearing measurement from the k th robot pose to the ℓ th landmark, d is the offset of the laser rangefinder from the robot center in the longitudinal direction, and σ_r^2 is the measurement noise variance. Although the dataset provides range to the landmarks as well, we chose to neglect these measurements to make the problem difficult and therefore show the benefit of taking a variational inference approach. Our setup is similar to a monocular camera situation, which is known to be a challenging SLAM problem.

Putting all the factors together, the posterior that we would like to project to a Gaussian is

$$p(\mathbf{x}) = \downarrow \exp(-\phi(\mathbf{x})), \quad \phi(\mathbf{x}) = \sum_{k=0}^K \varphi_k + \sum_{k=0}^K \psi_k + \sum_{k=1}^K \sum_{\ell=1}^L \psi_{\ell,k} + \text{constant}, \tag{74}$$

where it is understood that not all $L = 17$ landmarks are actually seen at each timestep and thus we must remove the factors for unseen landmarks.

We then carry out iterative projection to a multivariate Gaussian estimate, $q^{(i)} = \mathcal{N}(\boldsymbol{\mu}^{(i)}, \boldsymbol{\Sigma}^{(i)})$, according to

$$q^{(i+1)} = \downarrow \underset{(\mathcal{G}, q^{(i)})}{\text{proj}} p, \tag{75}$$

where we use (59) at implementation.

In a sub-sequence of the full dataset comprising 2000 timestamps with $(x, y, \theta, \dot{x}, \dot{y}, \dot{\theta})$ for the robot state at each time and 17 landmarks with (x, y) position, the dimension of the state estimation problem is $N = 12034$. Spanning the indefinite-Gaussian subspace requires N basis functions for the mean and $\frac{1}{2}N(N + 3)$ basis functions for the covariance, or 72, 438, 663 basis functions total. To naively apply the idea of iterative projection to Bayes space would be intractable. However, by exploiting the sparsity that Bayes space affords (see Sections 4.6 and 5.2) we are able to do this in a very computationally efficient way. In Ref. [14], the authors provide further implementation details of this experiment.

Figure 8 shows estimation error plots for the section of the robot’s trajectory in Fig. 7. Not only does this show the concepts of Bayes space can be applied to a large problem (N in the range of thousands) but also that there are some situations where it performs slightly better than the standard MAP Gauss-Newton (GN) algorithm.

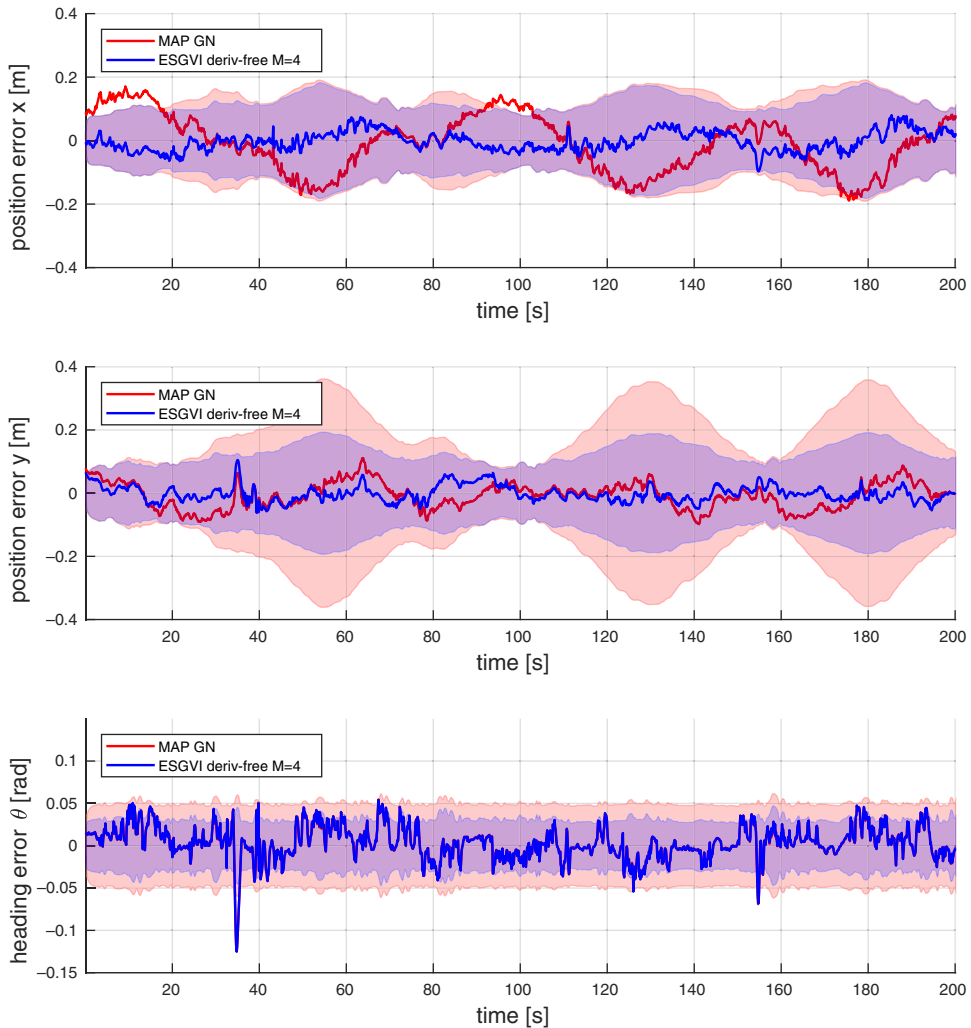


Figure 8. Error plots for a portion of the trajectory in the SLAM problem conducted by [14] and discussed in Section 5.3. The exactly sparse Gaussian variational inference (ESGVI) algorithm (red) is equivalent to the iterative projection approach described herein. The maximum a posteriori (MAP) Gauss-Newton (GN) algorithm (blue) is the more standard approach to solving this type of problem. Here, we see ESGVI performing slightly better than MAP GN in terms of smaller errors and more consistency (i.e., errors staying within covariance envelope). Note, in the heading error plot, the red mean line is hidden behind the blue one.

6. Discussion

6.1. Beyond Gaussians

Much of our discussion has centered on projection to the indefinite-Gaussian subspace and also the use of Gaussian measures in our definition of Bayes space. This is primarily because we wanted to show the connection between Bayes space and the Gaussian variational inference framework of [14]. However, we have attempted to lay out the framework to be as general as possible.

As a teaser of applying the methods beyond Gaussians, we can use $M \geq 2$ Hermite basis functions to see if we can better approximate a PDF. Figure 9 shows that indeed as we project to a higher-dimensional subspace, we are able to better approximate the stereo camera posterior introduced in Section 3.3. Here

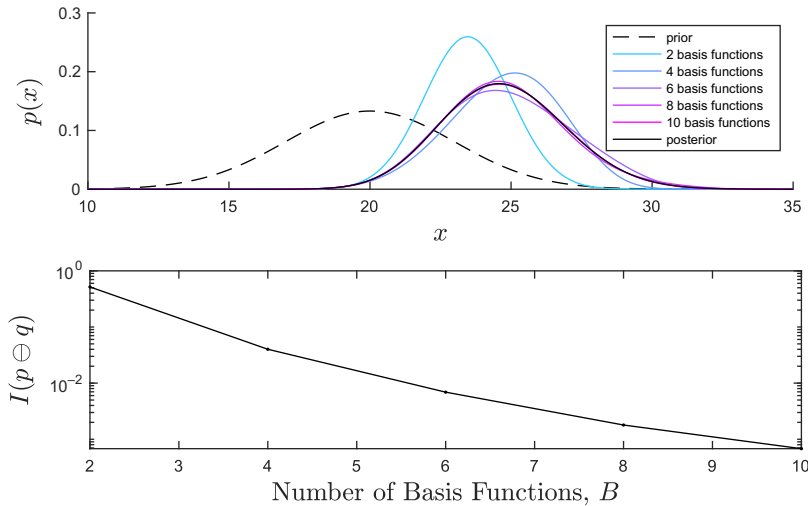


Figure 9. An example of projection of a posterior onto a finite basis with increasing number of basis functions. The top panel qualitatively shows that adding more basis functions brings the approximation closer to the posterior. The bottom shows the same quantitatively where $I(p \ominus q)$ decreases exponentially fast with more basis functions. The measure was taken to be the prior in this example.

we took the measure ν to be equal to the prior for the problem. This shows that even without iteratively updating the measure, we can better approximate the posterior by using more basis functions.

Moreover, we can repeat the iterative projection experiment from Section 4.5, this time with both 2 and 4 basis functions for the approximation. Figure 10 shows the results. We see that the 4-basis-function estimate requires a few more iterations to converge than the 2-basis-function one, but it arrives at a better final approximation as demonstrated by the lower final KL divergence.

6.2. Limitations and future work

While the results of the previous section make the use of high-dimensional subspaces look promising, there are some limitations still to overcome, which we discuss here.

First, while the establishment of \mathcal{B}^2 is mathematically sound, it is actually $\downarrow \mathcal{B}^2$ that we are primarily interested in, since we want to approximate valid PDFs by other valid (simpler) PDFs. It seems through our experiments that we have been lucky in the sense that the results of our projections to Bayesian subspaces are valid PDFs, but there is nothing that actually guarantees this for some of our approximation problems. For example, consider our one-dimensional Gaussian again:

$$p(x) = \underbrace{\left(-\frac{\mu}{\sigma^2}\right)}_{\alpha_1} \cdot \underbrace{\exp(-x)}_{b_1} \oplus \underbrace{\left(\frac{1}{\sqrt{2}\sigma^2}\right)}_{\alpha_2} \cdot \underbrace{\exp\left(-\frac{(x^2-1)}{\sqrt{2}}\right)}_{b_2} = \alpha_1 \cdot b_1 \oplus \alpha_2 \cdot b_2, \tag{76}$$

which is a member of $\downarrow \mathcal{G}$ when $\sigma^2 > 0$. If we project this vector onto $\text{span}\{b_1\}$, just the first Hermite basis vector, the result is

$$\text{proj}_{\text{span}\{b_1\}} p = \alpha_1 \cdot b_1, \tag{77}$$

which is no longer a member of $\downarrow \mathcal{G}$ since it cannot be normalized to become a valid PDF. Extrapolating from this simple example, it means that truncating a Fourier series at some arbitrary number of terms does not guarantee that the result will be a valid PDF. If we want to extend the Gaussian results to higher-dimensional subspaces, we need to better understand this issue.

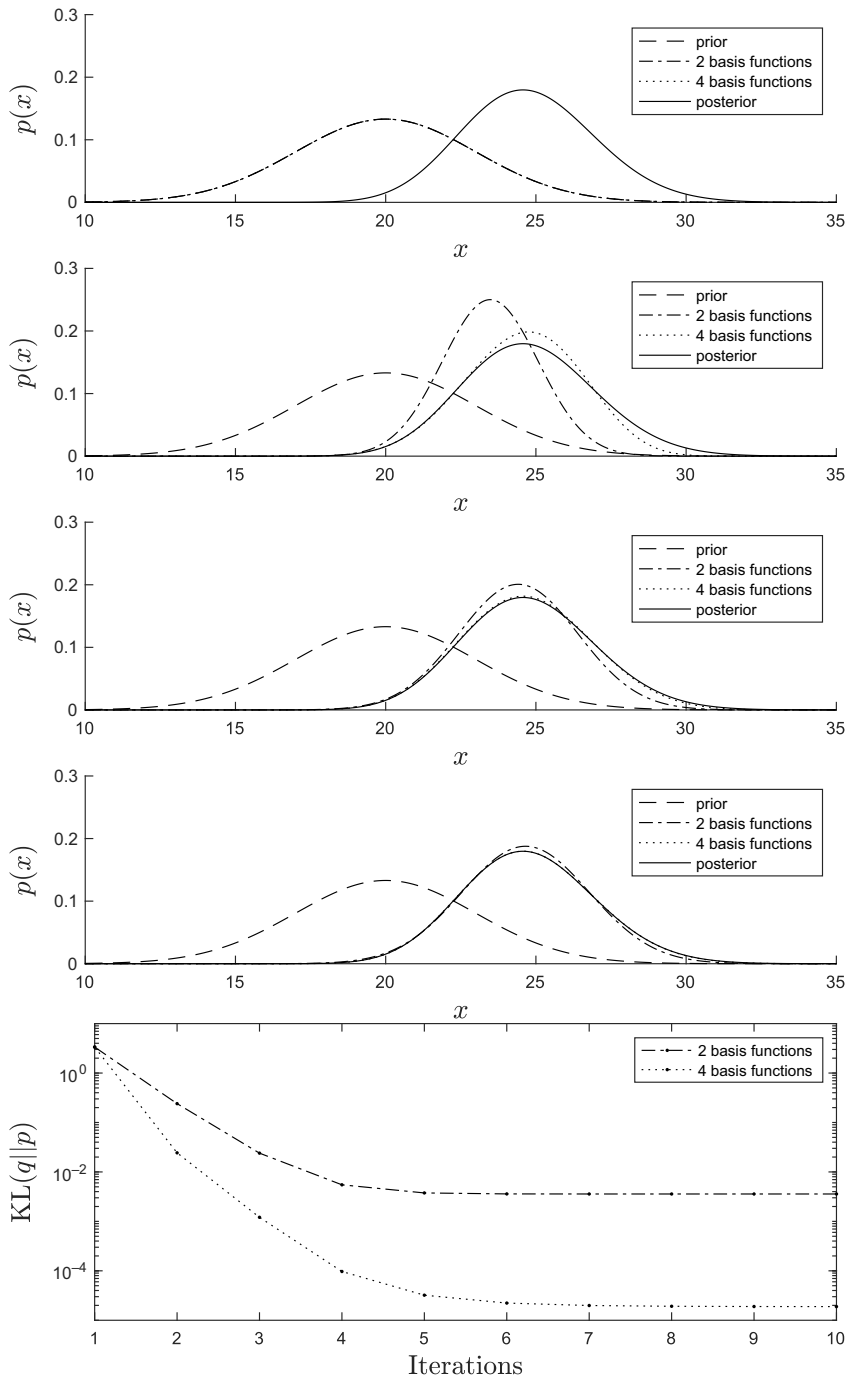


Figure 10. Example of iterative projection onto subspaces spanned by 2 and 4 Hermite basis functions, where the measure is taken to be the estimate $q^{(i)}$ at the previous iteration (projected to the indefinite-Gaussian subspace) and the basis reorthogonalized at each iteration as described in Section 5. The estimates were initialized to the prior (first panel) and then iteratively updated (next three panels). The last panel shows the KL divergence between the estimates and the true posterior for 10 iterations. We see that the estimate using 4 basis functions took slightly longer to converge but in the end produced a better approximation of the posterior.

Second, even in the case of projecting to the indefinite-Gaussian subspace, guaranteeing that the result is in $\downarrow\mathcal{G}$ is quite restrictive. If the PDF to be projected is

$$p(\mathbf{x}) = c \exp(-\phi(\mathbf{x})), \tag{78}$$

we saw that the projection to the indefinite-Gaussian subspace (see Appendix D.2) has the form

$$\text{proj}_{(\mathcal{G}, \nu)} p = \exp\left(-(\mathbf{x} - \boldsymbol{\mu})^T \mathbb{E}_\nu \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right] - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \underbrace{\mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right]}_{\Sigma^{-1}} (\mathbf{x} - \boldsymbol{\mu})\right), \tag{79}$$

where we have indicated the resulting inverse covariance is $\Sigma^{-1} = \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right]$. To guarantee $\Sigma^{-1} > 0$ which would make this a valid PDF for any choice of the measure ν , we require that $\phi(\mathbf{x})$ is a *convex* function of \mathbf{x} . This is clearly too restrictive for most real estimation problems involving nonlinear measurement models. If $\phi(\mathbf{x})$ is locally convex, it suggests the measure ν must be chosen so that its probability mass coincides with this region of local convexity. This perhaps emphasizes the need to iteratively update the measure in our proposed projection scheme. However, when the Bayesian posterior and prior are far apart, there is work to be done to understand how best to initialize the measure to ensure the projections wind up in $\downarrow\mathcal{B}^2$ in the general setup.

Finally, in the general case of projecting to a high-dimensional subspace, the measure itself could also be something other than a Gaussian, depending on the basis that is established. How to carry out the expectations in a computationally efficient and stable way in this case is again future work. The Hermite basis is cooperative in that the basis functions are orthonormal with respect to a Gaussian measure. In the high-dimensional SLAM problem that we discussed in Section 5.3, we (i) exploited sparsity inherent in the problem to require only taking expectations over marginals for each measurement factor, and (ii) were also able to exploit the fact that the measure was Gaussian in order to use Gaussian cubature to carry out the expectations somewhat efficiently [14]. Perhaps there are other bases that could be used for certain problems that admit similar computational conveniences.

It is worth noting that many of the challenges in working with Bayes space stem from the fact that we are attempting to work on infinite domains. If our interest lies with practical robotic state estimation, Bayes space defined over a finite domain [23] may be both mathematically simpler and more realistic from a practical point of view. This would of course mean giving up on using Gaussians, for example, which could be replaced by a truncated alternative.

7. Concluding remarks

Our principal goal in this work has been to provide a new perspective on the problem of variational inference. This new vantage point is afforded by considering PDFs as elements in a Bayesian Hilbert space, where vector addition is a multiplication (perturbation) that accounts for Bayes’ rule and scalar multiplication is an exponentiation (powering). Gaussians and, more generally, exponential families, which are often used in variational inference, are associated with subspaces. We thus have at our disposal all the familiar instruments of linear algebra.

The use of the KL divergence $\text{KL}(q\|p)$ in variational inference to find the best approximation q to a given posterior p is widespread. In most approaches, the canvas on which the minimization is carried out is a set, usually convex, or a manifold of admissible functions [1, 3, 5, 20, 21]. “Projections” of p onto the set or manifold are *ipso facto* the PDF q that minimizes the divergence. However, in Bayesian space, we may interpret projections as standard linear-algebraic projections, reminding us of a Euclidean world.

We take particular note of the information geometry of Csiszár and Amari. They along with their colleagues [6, 21] separately developed the *em* algorithm—not to be confused with the EM (expectation-maximization) algorithm although the two are in many cases equivalent—to solve the generalized variational problem, which involves a dual minimization of q over its manifold and p over its own.

(The minimum is therefore the minimum “distance” between manifolds.) The e -step of the algorithm is performed by making the manifold “flat,” that is, linear, as a result of using an exponential family of densities. This flattening is equivalent to thinking in terms of a Bayesian space as we have done here. Indeed, as we have shown, the natural-gradient-descent algorithm of [4] can be explained using this framework as a Newton-like iterative projection.

Based on the inner product of our Bayesian space, we have employed an information measure. It is proportional to the squared norm of a probability distribution, which can be used to establish a (symmetric and quadratic) divergence between two PDFs. The connection to the KL divergence is worthwhile mentioning. Each step in the iterative projection algorithm presented here for variational inference based on the KL divergence amounts to a local minimization of our Bayesian-space divergence. Admittedly, our iterative projection approach has some limitations and open issues to resolve. Particularly, we as yet cannot guarantee that the result of a projection will be a valid PDF and hence starting far from the final posterior estimate can pose challenges. We hope others may pick up where we have left off to further develop these ideas and overcome current limitations.

The linear structure of Bayes space furthermore allows us to treat sparsity in measurement data very neatly as the vector sum of the measurements, each of which can be expressed as an element in a subspace restricted to the local variables dictated by the sparsity of the problem, for example, as in the SLAM problem in robotics [14]. The mean-field approximation in variational inference can be handled in much the same way in this framework. The factorization of a distribution with respect to a desired family of distributions would again be rendered as a vector sum of PDFs.

In his fictional autobiography, *Zen and the Art of Motorcycle Maintenance*, Robert M. Pirsig notes that “One geometry cannot be more true than another; it can only be more convenient.” The same can be said of algebra. Whether one takes a geometric or algebraic tack in analyzing a problem, it can be agreed that different perspectives offer different views and given a particular problem or even a particular class of problem one tack may sometimes be more convenient than others. We hope the perspective presented here on variational inference using a Bayesian Hilbert space offers not only convenience in some respects but insight and a degree of elegance as well.

Authors’ contributions. Barfoot and D’Eleuterio both contributed to the concepts and writing of this work.

Financial support. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

Conflicts of interest. There are no conflicts of interest, financial or other, for this work.

Ethical considerations. There are no ethical considerations for this work.

References

- [1] M. Adamčík, “The information geometry of Bregman divergences and some applications in multi-expert reasoning,” *Entropy* **16**(12), 6338–6381 (2014).
- [2] J. Aitchison, “The statistical analysis of compositional data (with discussion),” *J. R. Stat. Soc. Ser. B* **44**, 139–177 (1982).
- [3] S.-I. Amari, “The EM algorithm and information geometry in neural network learning,” *Neural Comput.* **7**(1), 13–18 (1995).
- [4] S.-I. Amari, “Natural gradient works efficiently in learning,” *Neural Comput.* **10**(2), 251–276 (1998).
- [5] S.-I. Amari. *Information Geometry and Its Applications* (Springer, Japan, 2016).
- [6] S.-I. Amari, K. Kurata and H. Nagaoka, “Information geometry of Boltzmann machines,” *IEEE T rans. Neural Netw.* **3**(2), 260–271 (1992).
- [7] L. Ambrogioni, U. Güçlü, Y. Güçlütürk, M. Hinne, E. Maris and M. A. J. van Gerven, “Wasserstein Variational Inference,” **In: 32nd Conference on Neural Information Processing Systems (NIPS 2018)**, Montréal, Canada (2018).
- [8] D. Barber. *Bayesian Reasoning and Machine Learning* (Cambridge University Press, Cambridge, UK, 2012).
- [9] T. D. Barfoot, Stochastic Decentralized Systems, Ph.D. Thesis (University of Toronto, 2002).
- [10] T. D. Barfoot. *State Estimation for Robotics* (Cambridge University Press, Cambridge, UK, 2017).
- [11] T. D. Barfoot. *Multivariate Gaussian Variational Inference by Natural Gradient Descent*, “Technical report (Autonomous Space Robotics Lab, University of Toronto, (2020), arXiv:2001.10025 [stat.ML].

- [12] T. D. Barfoot and G. M. T. D'Eleuterio, "An Algebra for the Control of Stochastic Systems: Exercises in Linear Algebra," **In: Proceedings of the 5th International Conference on Dynamics and Control of Structures in Space (DCSS)**, Cambridge, England (2002).
- [13] T. D. Barfoot and G. M. T. D'Eleuterio, "Stochastic Algebra for Continuous Variables, Technical report," University of Toronto Institute for Aerospace Studies (2003).
- [14] T. D. Barfoot, J. R. Forbes and D. J. Yoon, "Exactly sparse Gaussian variational inference with application to derivative-free batch nonlinear state estimation," *Int. J. Robot. Res. (IJRR)* **39**(13), 1473–1502 (2020), 1911, (arXiv:1911.08333 [cs.RO]).
- [15] T. D. Barfoot, C. H. Tong and S. Sarkka, "Batch Continuous-Time Trajectory Estimation as Exactly Sparse Gaussian Process Regression," **In: Proceedings of Robotics: Science and Systems (RSS)**, Berkeley, USA (2014).
- [16] T. Bayes, "Essay towards solving a problem in the doctrine of chances," *Philos. Trans. R. Soc. Lond.* **53**, 370–418 (1763).
- [17] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, New York, 2006).
- [18] D. M. Blei, A. Kucukelbir and J. D. McAuliffe, "Variational inference: A review for statisticians," *J. Am. Stat. Assoc.* **112**(518), 859–877 (2017).
- [19] L. M. Bregman, "The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming," *USSR Comput. Math. Math. Phys.* **7**(3), 200–217 (1967).
- [20] I. Csiszár, "I-Divergence geometry of probability distributions and minimization problems," *Ann. Probab.* **3**(1), 146–158 (1975).
- [21] I. Csiszár and G. Tusnády, "Information Geometry and Alternating Minimization Procedures," **In: Statistics and Decisions, Supplement Issue No. 1, R. Oldenberg**, (1984).
- [22] J. Egozcue, V. Pawlowsky-Glahn, R. Tolosana-Delgado, M. I. Ortego and K. G. van den Boogaart, "Bayes spaces: use of improper distributions and exponential families," *RACSAM: Rev. Real Acad. Ciencias Exactas, Físicas Naturales. Ser. A. Mat.* **107**(2), 475–486 (2013).
- [23] J. J. Egozcue, J. L. Diaz-Barrero and V. Pawlowsky-Glahn, "Hilbert space of probability density functions based on Aitchison geometry," *Acta Math. Sin.* **22**(4), 1175–1182 (2006).
- [24] R. A. Fisher, "On the mathematical foundations of theoretical statistics," *Philos. Trans. R. Soc. Lond. Ser. A, Containing Papers of a Mathematical or Physical Character* **222**(594-604), 309–368 (1922).
- [25] G. E. Hinton and D. van Camp, "Keeping Neural Networks Simple by Minimizing the Description Length of the Weights," **In: Sixth ACM Conference on Computational Learning Theory**, Santa Cruz, California (1993).
- [26] A. H. Jazwinski, *Stochastic Processes and Filtering Theory* (Academic, New York, 1970).
- [27] M. I. Jordan, Z. Ghahramani, T. Jaakkola and L. K. Saul, "An introduction to variational methods for graphical models," *Mach. Learn.* **37**(2), 183–233 (1999).
- [28] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Stat.* **22**(1), 79–86 (1951).
- [29] P.-S. Laplace, , *Philosophical Essay on Probabilities*, Springer, (1995). translated by Andrew I. Dale from Fifth French Edition, 1825.
- [30] Y. Li and R. E. Turner, "Rényi Divergence Variational Inference," **In: 30th Conference on Neural Information Processing Systems (NIPS 2016)**, Barcelona, Spain (2016).
- [31] J. R. Magnus and H. Neudecker, "The elimination matrix: some lemmas and applications," *SIAM J. Algebraic Discret. Methods* **1**(4), 422–449 (1980).
- [32] J. R. Magnus and H. Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics* (John Wiley & Sons, Hoboken, NJ and Chichester, West Sussex, 2019).
- [33] J. H. Manton and P.-O. Amblard, "A Primer on Reproducing Kernel Hilbert Spaces, Technical report," The University of Melbourne and CNRS (2015), arXiv:1408.0952v2 [math.HO].
- [34] S. B. McGrayne, *The Theory That Would Not Die: How Bayes' Rule Cracked the Enigma Code, Hunted Down Russian Submarines, and Emerged Triumphant from Two Centuries of Controversy* (Yale University Press, New Haven, Connecticut, 2011).
- [35] G. Monge, "Mémoire sur la théorie des déblais et des remblais," **In: Histoire de l'Académie Royale des Sciences de Paris**, (1781).
- [36] A. Painsky and G. G. Wornell, Bregman divergence bounds and universality properties of the logarithmic loss, Department of Industrial Engineering, Tel Aviv University (2020). Technical report, arXiv: 1810.07014v2 [cs.IT].
- [37] V. Pawlowsky-Glahn and J. J. Egozcue, "Geometric approach to statistical analysis on the simplex," *Stoch. Environ. Res. Risk Assess.* **15**(5), 384–398 (2001).
- [38] A. Rényi, "On Measure of Entropy and Information," **In: Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability**, vol. 1, (1961) pp. 547–561.
- [39] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.* **27**(3), 379–423, 623–656 (1948).
- [40] C. M. Stein, "Estimation of the mean of a multivariate normal distribution," *Ann. Stat.* **9**(6), 1135–1151 (1981).
- [41] K. Takahashi, J. Fagan and M.-S. Chen, "A Sparse Bus Impedance Matrix and its Application to Short Circuit Study," **In: Proceedings of the PICA Conference** (1973).
- [42] K. G. van den Boogaart, J. J. Egozcue and V. Pawlowsky-Glahn, "Bayes linear spaces," *Stat. Oper. Res. Trans.* **34**(2), 201–222 (2010).
- [43] K. G. van den Boogaart, J. J. Egozcue and V. Pawlowsky-Glahn, "Bayes Hilbert spaces," *Aust. N. Z. Stat.* **56**(2), 171–194 (2014).
- [44] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Mach. Learn.* **1**(1-2), 1–305 (2008).

Appendix A. Kronecker Product, vec and vech Operators, and Duplication Matrices

For the benefit of the reader, we summarize several identities, which will be used in subsequent appendices, involving the *Kronecker product* \otimes and the *vectorization* operator $\text{vec}(\cdot)$ that stacks the columns of a matrix:

$$\begin{aligned}
 \text{vec}(\mathbf{a}) &\equiv \mathbf{a} \\
 \text{vec}(\mathbf{a}\mathbf{b}^T) &\equiv \mathbf{b} \otimes \mathbf{a} \\
 \text{vec}(\mathbf{A}\mathbf{B}\mathbf{C}) &\equiv (\mathbf{C}^T \otimes \mathbf{A}) \text{vec}(\mathbf{B}) \\
 \text{vec}(\mathbf{A})^T \text{vec}(\mathbf{B}) &\equiv \text{tr}(\mathbf{A}^T \mathbf{B}) \\
 (\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) &\equiv (\mathbf{A}\mathbf{C}) \otimes (\mathbf{B}\mathbf{D}) \\
 (\mathbf{A} \otimes \mathbf{B})^{-1} &\equiv \mathbf{A}^{-1} \otimes \mathbf{B}^{-1} \\
 (\mathbf{A} \otimes \mathbf{B})^T &\equiv \mathbf{A}^T \otimes \mathbf{B}^T.
 \end{aligned} \tag{A1}$$

It is worth noting that \otimes and $\text{vec}(\cdot)$ are linear operators.

As we will be working with (symmetric) covariance matrices when discussing Gaussians, we would like to be able to represent them parsimoniously in terms of only their unique variables. Following [32, §18], we introduce the *half-vectorization* operator $\text{vech}(\cdot)$ that stacks up the elements in a column matrix, excluding all the elements above the main diagonal. The *duplication matrix* \mathbf{D} allows us to recover a full symmetric matrix from its unique parts:

$$\text{vec}(\mathbf{A}) = \mathbf{D} \text{vech}(\mathbf{A}) \quad (\text{symmetric } \mathbf{A}). \tag{A2}$$

It is helpful to consider a simple 2×2 example:

$$\mathbf{A} = \begin{bmatrix} a & b \\ b & c \end{bmatrix}, \quad \text{vec}(\mathbf{A}) = \begin{bmatrix} a \\ b \\ b \\ c \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \text{vech}(\mathbf{A}) = \begin{bmatrix} a \\ b \\ c \end{bmatrix}. \tag{A3}$$

The *Moore-Penrose pseudoinverse* of \mathbf{D} will be denoted \mathbf{D}^\dagger and is given by

$$\mathbf{D}^\dagger = (\mathbf{D}^T \mathbf{D})^{-1} \mathbf{D}^T. \tag{A4}$$

We can then use \mathbf{D}^\dagger to convert the vectorization of a matrix into its half-vectorization:

$$\text{vech}(\mathbf{A}) = \mathbf{D}^\dagger \text{vec}(\mathbf{A}) \quad (\text{symmetric } \mathbf{A}). \tag{A5}$$

For our 2×2 example, we have

$$\mathbf{D}^\dagger = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{A6}$$

Useful identities involving \mathbf{D} are then

$$\begin{aligned}
 \mathbf{D}^\dagger \mathbf{D} &\equiv \mathbf{1} \\
 \mathbf{D}^\dagger{}^T \mathbf{D}^T &\equiv \mathbf{D}\mathbf{D}^\dagger \\
 \mathbf{D}\mathbf{D}^\dagger \text{vec}(\mathbf{A}) &\equiv \text{vec}(\mathbf{A}) \quad (\text{symmetric } \mathbf{A}) \\
 \mathbf{D}\mathbf{D}^\dagger (\mathbf{A} \otimes \mathbf{A}) \mathbf{D} &\equiv (\mathbf{A} \otimes \mathbf{A}) \mathbf{D} \quad (\text{any } \mathbf{A}),
 \end{aligned} \tag{A7}$$

which can be found in ref. [31].

Appendix B. Outer Products

The *outer product* $\Phi: \mathcal{B}^2 \rightarrow \mathcal{B}^2$ of two vectors $b = b(\mathbf{x}), c = c(\mathbf{x}') \in \mathcal{B}^2$, denoted $\Phi(\mathbf{x}, \mathbf{x}') = b(\mathbf{x}) \langle c(\mathbf{x}') \rangle$ or briefly $\Phi = b \langle c$, is defined by its operation on arbitrary $d = d(\mathbf{x}') \in \mathcal{B}^2$ as

$$\Phi(\mathbf{x}, \mathbf{x}') \otimes d(\mathbf{x}') = b(\mathbf{x}) \langle c(\mathbf{x}') \rangle \otimes d(\mathbf{x}') = b(\mathbf{x}) \cdot \langle c, d \rangle = \langle c, d \rangle \cdot b(\mathbf{x}). \tag{B1}$$

Thus, dropping the functional dependence,

$$\langle a, \Phi \otimes d \rangle = \langle a, b \rangle \langle c, d \rangle \tag{B2}$$

for arbitrary $a \in \mathcal{B}^2$. More generally,

$$\Phi = \bigoplus_{i=1}^M \bigoplus_{j=1}^N \phi_{ij} \cdot b_i \langle c_j, \tag{B3}$$

where $b_i, c_j \in \mathcal{B}^2$ and $\phi_{ij} \in \mathbb{R}$, so that

$$\Phi \otimes d = \bigoplus_{i=1}^M \sum_{j=1}^N \phi_{ij} \langle c_j, d \rangle \cdot b_i \tag{B4}$$

and

$$\langle a, \Phi \otimes d \rangle = \sum_{i=1}^M \sum_{j=1}^N \phi_{ij} \langle a, b_i \rangle \langle c_j, d \rangle. \tag{B5}$$

Defining the matrix $\Phi = [\phi_{ij}] \in \mathbb{R}^{M \times N}$ and

$$\mathbf{b}(\mathbf{x}) = \begin{bmatrix} b_1(\mathbf{x}) \\ b_2(\mathbf{x}) \\ \vdots \\ b_M(\mathbf{x}) \end{bmatrix}, \quad \mathbf{c}(\mathbf{x}) = \begin{bmatrix} c_1(\mathbf{x}) \\ c_2(\mathbf{x}) \\ \vdots \\ c_N(\mathbf{x}) \end{bmatrix}, \tag{B6}$$

we may abbreviate (B3) to

$$\Phi(\mathbf{x}, \mathbf{x}') = \mathbf{b}(\mathbf{x}) \Phi \langle \mathbf{c}(\mathbf{x}') \rangle \tag{B7}$$

and hence $\langle a, \Phi \otimes d \rangle = \Phi \langle \mathbf{c}, d \rangle$, where $\langle a, \mathbf{b} \rangle$ is interpreted as a row and $\langle \mathbf{c}, d \rangle$ as a column.

Given an orthonormal basis $\{b_1, b_2 \dots b_M\}$ for a subspace $\mathcal{S} \subset \mathcal{B}^2$, we define

$$Q = \bigoplus_{m=1}^M b_m \langle b_m \equiv \mathbf{b} \rangle \langle \mathbf{b}, \tag{B8}$$

and thus, for any $s \in \mathcal{S}$, $Q \otimes s = s$ [33]. For a nonorthonormal basis,

$$Q = \bigoplus_{m=1}^M \bigoplus_{n=1}^M \kappa_{mn} \cdot b_m \langle b_n \equiv \mathbf{b} \rangle \langle \mathbf{b}, \mathbf{b} \rangle^{-1} \langle \mathbf{b}, \tag{B9}$$

where κ_{mn} is the (m, n) entry in $\langle \mathbf{b}, \mathbf{b} \rangle^{-1}$. Notationally, $\cdot \rangle \mathbf{A} \langle \cdot$ indicates an outer product weighted in the middle by an appropriately sized matrix \mathbf{A} , which in the above example serves to normalize the basis. In normal matrix algebra, it would be equivalent, for example, to writing $\mathbf{a}(\mathbf{a}^T \mathbf{a})^{-1} \mathbf{a}^T$, for some column \mathbf{a} .

Using the outer product, we can write a projection as

$$q^*(\mathbf{x}) = Q(\mathbf{x}, \mathbf{x}') \otimes p(\mathbf{x}'), \tag{B10}$$

where

$$Q(\mathbf{x}, \mathbf{x}') = \mathbf{b}(\mathbf{x}) \langle \mathbf{b}, \mathbf{b} \rangle^{-1} \langle \mathbf{b}(\mathbf{x}'), \tag{B11}$$

which plays a similar role to projection matrix.

Appendix C. Hermite Basis

C.1. Basis for \mathbb{R}

Consider the domain over which members of \mathcal{B}^2 are defined to be \mathbb{R} . We can use the exponentiated Hermite polynomials as a basis for our infinite-dimensional \mathcal{B}^2 ; in fact, they prove to be a natural choice [43]. In one dimension, the first few *probabilist's Hermite polynomials* are

$$H_1(\xi) = \xi, \quad H_2(\xi) = \xi^2 - 1, \quad H_3(\xi) = \xi^3 - 3\xi, \quad H_4(\xi) = \xi^4 - 6\xi^2 + 3. \tag{C1}$$

(We exclude $H_0(\xi) = 1$ as the resulting vector is the zero vector; however, it will need to be introduced when considering the domain \mathbb{R}^N as explained in Appendix C.2.) Owing to the properties of the Hermite polynomials, namely, that

$$\int_{-\infty}^{\infty} H_n(\xi)v(\xi) d\xi = 0, \quad \int_{-\infty}^{\infty} H_m(\xi)H_n(\xi)v(\xi) d\xi = n! \delta_{mn}, \quad m, n = 1, 2, 3 \dots, \tag{C2}$$

where $v(\xi) = \mathcal{N}(0, 1)$ is the standard normal density, we can construct an orthonormal basis for \mathcal{B}^2 following [23]. Accordingly,

$$\mathbb{E}_v[H_n] = 0, \quad \mathbb{E}_v[H_m H_n] = n! \delta_{mn}, \quad m, n = 1, 2, 3 \dots \tag{C3}$$

Our basis functions are

$$h_n(\xi) = \exp(-\eta_n(\xi)), \quad \eta_n(\xi) = \frac{1}{\sqrt{n!}} H_n(\xi). \tag{C4}$$

Orthogonality follows as

$$\begin{aligned} \langle h_m, h_n \rangle &= \mathbb{E}_v[\eta_m \eta_n] - \mathbb{E}_v[\eta_m] \mathbb{E}_v[\eta_n] \\ &= \frac{1}{\sqrt{m!n!}} \int_{-\infty}^{\infty} H_m(\xi)H_n(\xi) \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\xi^2\right) d\xi = \delta_{mn}. \end{aligned} \tag{C5}$$

An arbitrary member p of \mathcal{B}^2 can be expanded in terms of this Hermite basis. However, we first need two lemmata, resting on the recursive definition of Hermite polynomials; these are

Lemma C.1. *For the standard normal measure, $v \sim \mathcal{N}(0, 1)$,*

$$\mathbb{E}_v[H_{n+1}(\xi)f(\xi)] = \mathbb{E}_v\left[H_n(\xi) \frac{\partial f(\xi)}{\partial \xi}\right]. \tag{C6}$$

where $f(\xi)$ is a differentiable function and is such that the expectations exist.

Proof. The $n = 0$ case,

$$\mathbb{E}_v[\xi f] = \mathbb{E}_v\left[\frac{\partial f}{\partial \xi}\right], \tag{C7}$$

is immediately true by Stein's lemma [40]. For general case n ,

$$\begin{aligned} \mathbb{E}_v[H_{n+1}f] &= \mathbb{E}_v\left[\left(\xi H_n - \frac{\partial H_n}{\partial \xi}\right) f\right] = \mathbb{E}_v\left[\frac{\partial}{\partial \xi}(H_n f) - \frac{\partial H_n}{\partial \xi} f\right] = \mathbb{E}_v\left[\frac{\partial H_n}{\partial \xi} f + H_n \frac{\partial f}{\partial \xi} - \frac{\partial H_n}{\partial \xi} f\right] \\ &= \mathbb{E}_v\left[H_n \frac{\partial f}{\partial \xi}\right], \end{aligned} \tag{C8}$$

where we have used the recurrence relation,

$$H_{n+1} = \xi H_n - \frac{\partial H_n}{\partial \xi}, \tag{C9}$$

for the Hermite polynomials. □

Lemma C.2. For the standard normal measure, $\nu \sim \mathcal{N}(0, 1)$,

$$\mathbb{E}_\nu [H_n(\xi) f(\xi)] = \mathbb{E}_\nu \left[\frac{\partial^n f(\xi)}{\partial \xi^n} \right], \tag{C10}$$

where $f(\xi)$ is an n -fold differentiable function and is such that the expectations exist.

Proof. Repeatedly applying Lemma C.1,

$$\mathbb{E}_\nu [H_n f] = \mathbb{E}_\nu \left[H_{n-1} \frac{\partial f}{\partial \xi} \right] = \mathbb{E}_\nu \left[H_{n-2} \frac{\partial^2 f}{\partial \xi^2} \right] = \dots = \mathbb{E}_\nu \left[H_1 \frac{\partial^{n-1} f}{\partial \xi^{n-1}} \right] = \mathbb{E}_\nu \left[H_0 \frac{\partial^n f}{\partial \xi^n} \right] = \mathbb{E}_\nu \left[\frac{\partial^n f}{\partial \xi^n} \right], \tag{C11}$$

yields the desired result. □

Now consider any $p \in \mathcal{B}^2$ expressed as $p(\xi) = c \exp(-\phi(\xi))$. The coordinates are given by

$$\alpha_n = \langle h_n, p \rangle = \frac{1}{\sqrt{n!}} \mathbb{E}_\nu \left[\frac{\partial^n \phi(\xi)}{\partial \xi^n} \right] \tag{C12}$$

and hence

$$p(\xi) = \bigoplus_{n=1}^{\infty} \alpha_n \cdot h_n(\xi) = \exp \left(- \sum_{n=1}^{\infty} \frac{1}{n!} \mathbb{E}_\nu \left[\frac{\partial^n \phi(\xi)}{\partial \xi^n} \right] H_n(\xi) \right). \tag{C13}$$

We can account for measures other than the standard normal density, say $\nu \sim \mathcal{N}(\mu, \sigma^2)$, by the well-known reparameterization “trick,”

$$x = \mu + \sigma \xi, \tag{C14}$$

which leads to

$$p(x) = \bigoplus_{n=1}^{\infty} \alpha_n \cdot h_n \left(\frac{x - \mu}{\sigma} \right) = \exp \left(- \sum_{n=1}^{\infty} \frac{\sigma^n}{n!} \mathbb{E}_\nu \left[\frac{\partial^n \phi(x)}{\partial x^n} \right] H_n \left(\frac{x - \mu}{\sigma} \right) \right). \tag{C15}$$

It is instructive to rewrite this expression by replacing $-\phi$ with $\ln p$ giving

$$p(x) = \exp \left(\sum_{n=1}^{\infty} \frac{\sigma^n}{n!} \mathbb{E}_\nu \left[\frac{\partial^n \ln p(x)}{\partial x^n} \right] H_n \left(\frac{x - \mu}{\sigma} \right) \right). \tag{C16}$$

This is a Taylor-like expansion of p pivoting on a given mean μ and standard deviation σ .

Any subset of the basis functions $\{h_1, h_2, \dots\}$ establishes a subspace of \mathcal{B}^2 ; however, as far as such subspaces are concerned, it would be natural to choose an M -dimensional subspace \mathcal{H} spanned by the first M basis functions. As the basis is orthonormal, the Gram matrix is $(\mathbf{h}, \mathbf{h}) = \mathbf{1}$.

The Hermite functions can also be used to generate a basis for \mathcal{B}^2 on the domain \mathbb{R}^N , which we detail in the next subsection.

C.2. Basis for \mathbb{R}^N

We can extend the results of the previous subsection to create a Hermite basis for \mathcal{B}^2 on \mathbb{R}^N . Let

$$\boldsymbol{\eta}(\xi) = \frac{1}{\sqrt{n!}} \begin{bmatrix} H_0(\xi) \\ H_1(\xi) \\ \vdots \\ H_M(\xi) \end{bmatrix}. \tag{C17}$$

Note that we have reintroduced $H_0(\xi)$ because the basis will be created by all possible combinatorial N -products of these functions, one for each variable in $\xi \in \mathbb{R}^N$. However, we will have to exclude the combination made up of only H_0 because once again this function gives the zero vector of \mathcal{B}^2 . We may express this operation as a Kronecker product, that is,

$$\eta(\xi) = \mathbf{C} (\eta(\xi_1) \otimes \eta(\xi_2) \otimes \cdots \otimes \eta(\xi_N)), \tag{C18}$$

where $\mathbf{C} = [\mathbf{0} \ \mathbf{1}]$ contains zero in the first column followed by the identity matrix; this removes the offending function. Observe that $\mathbf{C}\mathbf{C}^T = \mathbf{1}$. The basis is then

$$\mathbf{h}(\xi) = \exp(-\eta(\xi)). \tag{C19}$$

The total number of basis functions is $(M + 1)^N - 1$.

This set of basis functions retains its orthonormality because

$$\begin{aligned} \langle \mathbf{h}(\xi), \mathbf{h}(\xi) \rangle &= \mathbb{E}_v [(\mathbf{C}\eta(\xi_1) \otimes \cdots \otimes \eta(\xi_N)) (\mathbf{C}\eta(\xi_1) \otimes \cdots \otimes \eta(\xi_N))^T] \\ &= \mathbf{C} \mathbb{E}_v [\eta(\xi_1)\eta(\xi_1)^T \otimes \cdots \otimes \eta(\xi_N)\eta(\xi_N)^T] \mathbf{C}^T \end{aligned} \tag{C20}$$

by a property of the Kronecker product (Appendix A). Now

$$\begin{aligned} &\mathbb{E}_v [\eta(\xi_1)\eta(\xi_1)^T \otimes \cdots \otimes \eta(\xi_N)\eta(\xi_N)^T] \\ &= \int_{-\infty}^{\infty} \eta(\xi_1)\eta(\xi_1)^T v(\xi_1) d\xi_1 \otimes \cdots \otimes \int_{-\infty}^{\infty} \eta(\xi_N)\eta(\xi_N)^T v(\xi_N) d\xi_N = \mathbf{1}_{(M+1)^N \times (M+1)^N}, \end{aligned} \tag{C21}$$

wherein each of the integrals expresses the orthonormality of the Hermite functions and results in an $(M + 1) \times (M + 1)$ identity matrix. Hence

$$\langle \mathbf{h}(\xi), \mathbf{h}(\xi) \rangle = \mathbf{C} \mathbf{1}_{(M+1)^N \times (M+1)^N} \mathbf{C}^T = \mathbf{1}_{[(M+1)^N - 1] \times [(M+1)^N - 1]}. \tag{C22}$$

To determine the coordinates of an arbitrary $p \in \mathcal{B}^2$, we shall require the multivariate version of Lemma 2:

Lemma C.3. *For the standard normal measure, $v \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$,*

$$\mathbb{E}_v [H_{n_1}(\xi_1)H_{n_2}(\xi_2) \cdots H_{n_N}(\xi_N)f(\xi)] = \mathbb{E}_v \left[\frac{\partial^{n_1+n_2+\cdots+n_N} f(\xi)}{\partial \xi_1^{n_1} \partial \xi_2^{n_2} \cdots \partial \xi_N^{n_N}} \right], \quad n_k = 1, 2 \cdots M. \tag{C23}$$

where $f : \mathbb{R}^N \rightarrow \mathbb{R}$ is n_k -fold differentiable in ξ_k and is such that the expectations exist.

The proof relies on the use of Lemma 2 for each individual partial derivative; for example, with respect to the variable ξ_1 ,

$$\mathbb{E}_v [H_{n_1}(\xi_1)H_{n_2}(\xi_2) \cdots H_{n_N}(\xi_N)f(\xi)] = \mathbb{E}_v \left[H_{n_1-1}(\xi_1)H_{n_2}(\xi_2) \cdots H_{n_N}(\xi_N) \frac{\partial f(\xi)}{\partial \xi_1} \right]. \tag{C24}$$

The product $H_{n_1}(\xi_2) \cdots H_{n_N}(\xi_N)$ has no dependence on ξ_1 and can therefore be treated as a constant. Doing the same for all the other variables and for the indicated number of times leads to the stated result.

We can streamline the notation by defining

$$\partial_\xi = \left[1 \ \frac{\partial}{\partial \xi} \ \frac{\partial^2}{\partial \xi^2} \ \cdots \ \frac{\partial^M}{\partial \xi^M} \right] \tag{C25}$$

and, as above,

$$\partial_\xi = \mathbf{C}(\partial_{\xi_1} \otimes \partial_{\xi_2} \otimes \cdots \otimes \partial_{\xi_N}). \tag{C26}$$

Using the measure $v = \mathcal{N}(0, 1)$, then,

$$\alpha = \langle \mathbf{h}(\xi), p(\xi) \rangle = \mathbb{E}_v [\partial_\xi p(\xi)] \tag{C27}$$

are the coordinates of $p(\xi) \in \mathcal{B}^2$, truncated to however many basis functions we decide to keep.

Appendix D. Multivariate Gaussians

D.1. Basis for multivariate Gaussians

Multivariate Gaussians are quintessentially important to statistics and estimation theory. Gaussians, as traditionally defined with a positive-definite covariance matrix, do not in themselves form a subspace of \mathcal{B}^2 . We need to expand the set to include covariance matrices that are sign-indefinite. Let us accordingly define an N -dimensional *indefinite-Gaussian PDF* as

$$p(\mathbf{x}) = c \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right), \tag{D1}$$

which has mean, $\boldsymbol{\mu}$, and symmetric covariance, $\boldsymbol{\Sigma}$. The set of all N -dimensional, indefinite Gaussians is

$$\mathcal{G} = \left\{ p(\mathbf{x}) = c \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right) \mid \boldsymbol{\mu} \in \mathbb{R}^N, \boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}, 0 < c < \infty \right\}. \tag{D2}$$

It is easy to show that \mathcal{G} is in fact a subspace of \mathcal{B}^2 as the zero vector is contained therein ($\boldsymbol{\Sigma}^{-1} = \mathbf{0}$, allowing that $\boldsymbol{\Sigma} \rightarrow \infty$) and the set is closed under vector addition and scalar multiplication.

To establish \mathcal{G} as a Bayesian Hilbert space, we must have an appropriate measure, ν . In our case, we choose the measure to also be a Gaussian, $\nu = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \in \mathcal{G}$. We may thus declare \mathcal{G} to be a Bayesian Hilbert space for a measure $\nu \in \mathcal{G}$. We will refer to the set of Gaussian PDFs with positive-definite covariance, $\boldsymbol{\Sigma} > 0$, as $\downarrow \mathcal{G} \subset \mathcal{G}$.

Several possibilities exist to parameterize Gaussians [11]. There are $\frac{1}{2}N(N + 3)$ unique elements contained in the mean and the symmetric covariance matrix on \mathbb{R}^N ; hence the dimension of \mathcal{G} is $\frac{1}{2}N(N + 3)$. We shall construct our basis on a positive-definite choice of covariance $\boldsymbol{\Sigma}$ that we can decompose in Cholesky fashion, that is, $\boldsymbol{\Sigma} = \mathbf{L}\mathbf{L}^T$. Now consider

$$\boldsymbol{\gamma}_1(\mathbf{x}) = \mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu}), \quad \boldsymbol{\gamma}_2(\mathbf{x}) = \sqrt{\frac{1}{2}\mathbf{D}^T \mathbf{D}} \text{vech}\left(\mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{L}^{-T}\right), \tag{D3}$$

wherein $\text{vech}(\cdot)$ is the *half-vectorization* of its matrix argument and \mathbf{D} is the associated duplication matrix (see Appendix A). Note that $\boldsymbol{\gamma}_1$ is an $N \times 1$ column and $\boldsymbol{\gamma}_2$ is an $\frac{1}{2}N(N + 1) \times 1$ column. With a little abuse of notation, we set the basis functions as

$$\mathbf{g}(\mathbf{x}) = \begin{bmatrix} \mathbf{g}_1(\mathbf{x}) \\ \mathbf{g}_2(\mathbf{x}) \end{bmatrix} = \exp\left(-\begin{bmatrix} \boldsymbol{\gamma}_1(\mathbf{x}) \\ \boldsymbol{\gamma}_2(\mathbf{x}) \end{bmatrix}\right); \tag{D4}$$

that is, the exponential is applied elementwise. We claim that $\mathbf{g}(\mathbf{x})$ is a basis for \mathcal{G} .

It is instructive to show that $\mathbf{g}(\mathbf{x})$ spans \mathcal{G} as well as serving as the proof that it is a basis. Consider again the reparameterization “trick” given by

$$\mathbf{x} = \boldsymbol{\mu} + \mathbf{L}\boldsymbol{\xi} \tag{D5}$$

with $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$. This renders (D4) as

$$\mathbf{g}(\boldsymbol{\xi}) = \begin{bmatrix} \mathbf{g}_1(\boldsymbol{\xi}) \\ \mathbf{g}_2(\boldsymbol{\xi}) \end{bmatrix} = \exp \left(- \begin{bmatrix} \boldsymbol{\gamma}_1(\boldsymbol{\xi}) \\ \boldsymbol{\gamma}_2(\boldsymbol{\xi}) \end{bmatrix} \right) = \exp \left(- \begin{bmatrix} \boldsymbol{\xi} \\ \sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D}} \text{vech } \boldsymbol{\xi} \boldsymbol{\xi}^T \end{bmatrix} \right). \tag{D6}$$

A (normalized) linear combination of the basis functions can be written as

$$p(\boldsymbol{\xi}) = \downarrow \exp \left(-\boldsymbol{\alpha}_1^T \boldsymbol{\gamma}_1(\boldsymbol{\xi}) - \boldsymbol{\alpha}_2^T \boldsymbol{\gamma}_2(\boldsymbol{\xi}) \right). \tag{D7}$$

Now

$$\boldsymbol{\alpha}_1^T \boldsymbol{\gamma}_1 = \boldsymbol{\alpha}_1^T \boldsymbol{\xi}. \tag{D8}$$

Also, we can in general express the second set of coordinates as

$$\boldsymbol{\alpha}_2 = \sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D}} \text{vech } \mathbf{S} \tag{D9}$$

for some symmetric \mathbf{S} that can easily be reconstructed from $\boldsymbol{\alpha}_2$. Hence

$$\boldsymbol{\alpha}_2^T \boldsymbol{\gamma}_2 = \frac{1}{2} (\text{vech } \mathbf{S})^T \mathbf{D}^T \mathbf{D} \text{vech } \boldsymbol{\xi} \boldsymbol{\xi}^T = \frac{1}{2} (\text{vec } \mathbf{S})^T \text{vec } \boldsymbol{\xi} \boldsymbol{\xi}^T \tag{D10}$$

given the identities $\text{vech } \mathbf{A} = \mathbf{D}^\dagger \text{vec } \mathbf{A}$ and $\mathbf{D} \mathbf{D}^\dagger \text{vec } \mathbf{A} = \text{vec } \mathbf{A}$, where \mathbf{D}^\dagger is the Moore-Penrose inverse of \mathbf{D} (Appendix A). Moreover, the identity $(\text{vec } \mathbf{A})^T \text{vec } \mathbf{b} = \text{tr } \mathbf{A} \mathbf{b}$ leads to

$$\boldsymbol{\alpha}_2^T \boldsymbol{\gamma}_2 = \frac{1}{2} \text{tr} (\mathbf{S} \boldsymbol{\xi} \boldsymbol{\xi}^T) = \frac{1}{2} \boldsymbol{\xi}^T \mathbf{S} \boldsymbol{\xi}. \tag{D11}$$

Then

$$\begin{aligned} p(\mathbf{x}) &= \downarrow \exp \left(-\boldsymbol{\alpha}_1^T \boldsymbol{\xi} - \frac{1}{2} \boldsymbol{\xi}^T \mathbf{S} \boldsymbol{\xi} \right) = \downarrow \exp \left(-\frac{1}{2} (\boldsymbol{\xi} + \mathbf{S}^{-1} \boldsymbol{\alpha}_1)^T \mathbf{S} (\boldsymbol{\xi} + \mathbf{S}^{-1} \boldsymbol{\alpha}_1) \right) \\ &= \downarrow \exp \left(-\frac{1}{2} (\mathbf{x} - (\boldsymbol{\mu} - \mathbf{L} \mathbf{S}^{-1} \boldsymbol{\alpha}_1))^T \mathbf{L}^{-T} \mathbf{S} \mathbf{L}^{-1} (\mathbf{x} - (\boldsymbol{\mu} - \mathbf{L} \mathbf{S}^{-1} \boldsymbol{\alpha}_1)) \right). \end{aligned} \tag{D12}$$

This can represent any Gaussian distribution, where the mean is $\boldsymbol{\mu} - \mathbf{L} \mathbf{S}^{-1} \boldsymbol{\alpha}_1$ and the covariance $\mathbf{L} \mathbf{S}^{-1} \mathbf{L}^T$. Thus, \mathbf{g} spans \mathcal{G} . Furthermore, as the dimension of \mathcal{G} is $\frac{1}{2}N(N + 3)$, the number of functions in \mathbf{g} , \mathbf{g} is a basis for \mathcal{G} .

This basis is, in addition, orthonormal as can be proven in a straightforward fashion by using the reparameterized form $\mathbf{g}(\boldsymbol{\xi})$ and recognizing that the entries in $\boldsymbol{\gamma}_1(\boldsymbol{\xi})$ are ξ_i and those in $\boldsymbol{\gamma}_2(\boldsymbol{\xi})$ are either $\xi_i \xi_j$ ($i \neq j$) or $\xi_i^2 / \sqrt{2}$. Hence, $\langle \mathbf{g}, \mathbf{g} \rangle = \mathbf{1}$.

It can be shown that

$$\begin{aligned} \boldsymbol{\alpha}_1 &= \langle \mathbf{g}_1, p \rangle = \mathbf{L}^T \mathbb{E}_v \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right], \\ \boldsymbol{\alpha}_2 &= \langle \mathbf{g}_2, p \rangle = \sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D}} \text{vech} \left(\mathbf{L}^T \mathbb{E}_v \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \mathbf{L} \right) \end{aligned} \tag{D13}$$

are the coordinates for $p(\mathbf{x}) = \downarrow \exp \left(-\phi(\mathbf{x}) \right) \in \mathcal{G}$. Another rendering of (D12) is

$$p(\mathbf{x}) = \downarrow \exp \left(-(\mathbf{x} - \boldsymbol{\mu})^T \mathbb{E}_v \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right] - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \mathbb{E}_v \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] (\mathbf{x} - \boldsymbol{\mu}) \right), \tag{D14}$$

which also expresses the projection of a PDF in \mathcal{B}^2 onto \mathcal{G} .

D.2. Coordinates of multivariate Gaussian projection

Let $p(\mathbf{x}) = c \exp(-\phi(\mathbf{x})) \in \mathcal{B}^2$. Projecting onto \mathcal{G} , the coordinates associated with basis functions \mathbf{g}_1 are

$$\begin{aligned} \boldsymbol{\alpha}_1 &= \langle \mathbf{g}_1, p \rangle \\ &= \mathbb{E}_v[\boldsymbol{\gamma}_1(\mathbf{x})\phi(\mathbf{x})] - \mathbb{E}_v[\boldsymbol{\gamma}_1(\mathbf{x})] \mathbb{E}_v[\phi(\mathbf{x})] \\ &= \mathbb{E}_v[\mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})\phi(\mathbf{x})] - \underbrace{\mathbb{E}_v[\mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})]}_0 \mathbb{E}_v[\phi(\mathbf{x})] \\ &= \mathbf{L}^{-1} \boldsymbol{\Sigma} \mathbb{E}_v \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right] \\ &= \mathbf{L}^T \mathbb{E}_v \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right], \end{aligned} \tag{D15}$$

where we have employed Stein’s lemma [40] to go from the third line to the fourth. Taking the inner product of these coefficients with the associated basis functions, we have

$$\boldsymbol{\alpha}_1^T \boldsymbol{\gamma}_1(\mathbf{x}) = \mathbb{E}_v \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right]^T \mathbf{L} \mathbf{L}^{-1} (\mathbf{x} - \boldsymbol{\mu}) = \mathbb{E}_v \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right]^T (\mathbf{x} - \boldsymbol{\mu}). \tag{D16}$$

The coordinates associated with basis functions \mathbf{g}_2 are

$$\begin{aligned} \boldsymbol{\alpha}_2 &= \langle \mathbf{g}_2, p \rangle \\ &= \mathbb{E}_v[\boldsymbol{\gamma}_2(\mathbf{x})\phi(\mathbf{x})] - \mathbb{E}_v[\boldsymbol{\gamma}_2(\mathbf{x})] \mathbb{E}_v[\phi(\mathbf{x})] \\ &= \mathbb{E}_v \left[\sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D} \text{vech}(\mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{L}^{-T})} \phi(\mathbf{x}) \right] \\ &\quad - \mathbb{E}_v \left[\sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D} \text{vech}(\mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{L}^{-T})} \right] \mathbb{E}_v[\phi(\mathbf{x})] \\ &= \sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D} \text{vech}(\mathbf{L}^{-1}(\mathbb{E}_v[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \phi(\mathbf{x})] - \boldsymbol{\Sigma} \mathbb{E}_v[\phi(\mathbf{x})]) \mathbf{L}^{-T})} \\ &= \sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D} \text{vech}(\mathbf{L}^{-1} \boldsymbol{\Sigma} \mathbb{E}_v \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \boldsymbol{\Sigma} \mathbf{L}^{-T})} \\ &= \sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D} \text{vech}(\mathbf{L}^T \mathbb{E}_v \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \mathbf{L})}, \end{aligned} \tag{D17}$$

where we have again used Stein’s lemma to go from the fourth line to the fifth, this time a double application. Taking the inner product of these coefficients with the associated basis functions, we have

$$\begin{aligned}
 \alpha_2^T \gamma_2(\mathbf{x}) &= \frac{1}{2} \text{vech} \left(\mathbf{L}^T \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \mathbf{L} \right)^T \mathbf{D}^T \mathbf{D} \text{vech} \left(\mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{L}^{-T} \right) \\
 &= \frac{1}{2} \text{vec} \left(\mathbf{L}^T \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \mathbf{L} \right)^T \underbrace{\mathbf{D}^{\dagger T} \mathbf{D}^T \mathbf{D} \mathbf{D}^\dagger}_{\mathbf{D}} \text{vec} \left(\mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{L}^{-T} \right) \\
 &= \frac{1}{2} \text{vec} \left(\mathbf{L}^T \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \mathbf{L} \right)^T \underbrace{\mathbf{D} \mathbf{D}^\dagger}_{\text{use (A7) third line}} \text{vec} \left(\mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{L}^{-T} \right) \\
 &= \frac{1}{2} \text{vec} \left(\mathbf{L}^T \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \mathbf{L} \right)^T \text{vec} \left(\mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{L}^{-T} \right) \tag{D18} \\
 &= \frac{1}{2} \text{tr} \left(\mathbf{L}^T \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \mathbf{L} \mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{L}^{-T} \right) \\
 &= \frac{1}{2} \text{tr} \left((\mathbf{x} - \boldsymbol{\mu})^T \mathbf{L}^{-T} \mathbf{L}^T \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] \mathbf{L} \mathbf{L}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \right) \\
 &= \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] (\mathbf{x} - \boldsymbol{\mu}).
 \end{aligned}$$

Combining these, we have

$$\downarrow_{(\mathcal{G}, \nu)} \text{proj } p = \downarrow \exp \left(-(\mathbf{x} - \boldsymbol{\mu})^T \mathbb{E}_\nu \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right] - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] (\mathbf{x} - \boldsymbol{\mu}) \right) \tag{D19}$$

for the projection in terms of its Gaussian basis.

D.3. Gaussian information

We calculate here the information I contained in a multivariate Gaussian distribution, $g(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}', \boldsymbol{\Sigma}') \in \downarrow \mathcal{G}$. We have

$$g(\mathbf{x}) = \downarrow \exp(-\phi(\mathbf{x})) \tag{D20}$$

with

$$\phi(\mathbf{x}) = \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}')^T \boldsymbol{\Sigma}'^{-1} (\mathbf{x} - \boldsymbol{\mu}'). \tag{D21}$$

The measure is taken as $\nu = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

Using our orthonormal basis for \mathcal{G} , the information in g is

$$I(g) = \frac{1}{2} \|g\|^2 = \frac{1}{2} \langle g, g \rangle = \frac{1}{2} (\boldsymbol{\alpha}_1^T \boldsymbol{\alpha}_1 + \boldsymbol{\alpha}_2^T \boldsymbol{\alpha}_2), \tag{D22}$$

where $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$ are the coordinates. As

$$\mathbb{E}_\nu \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^T} \right] = \boldsymbol{\Sigma}'^{-1} (\boldsymbol{\mu} - \boldsymbol{\mu}'), \quad \mathbb{E}_\nu \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^T \partial \mathbf{x}} \right] = \boldsymbol{\Sigma}'^{-1}, \tag{D23}$$

these coordinates are, by (D13),

$$\boldsymbol{\alpha}_1 = \mathbf{L}^T \boldsymbol{\Sigma}'^{-1} (\boldsymbol{\mu} - \boldsymbol{\mu}'), \quad \boldsymbol{\alpha}_2 = \sqrt{\frac{1}{2} \mathbf{D}^T \mathbf{D} \mathbf{D}^\dagger} \text{vec} \left(\mathbf{L}^T \boldsymbol{\Sigma}'^{-1} \mathbf{L} \right). \tag{D24}$$

Hence, from (D22),

$$I(g) = \frac{1}{2} \left((\boldsymbol{\mu} - \boldsymbol{\mu}')^T \boldsymbol{\Sigma}'^{-1} \boldsymbol{\Sigma} \boldsymbol{\Sigma}'^{-1} (\boldsymbol{\mu} - \boldsymbol{\mu}') + \frac{1}{2} \text{tr} \boldsymbol{\Sigma}'^{-1} \boldsymbol{\Sigma} \boldsymbol{\Sigma}'^{-1} \boldsymbol{\Sigma} \right), \tag{D25}$$

where the second term is a result of the fourth identity in (A1) and the third in (A7). It will, however, be instructive to rewrite the terms as

$$\begin{aligned}
 (\boldsymbol{\mu} - \boldsymbol{\mu}')^T \boldsymbol{\Sigma}'^{-1} \boldsymbol{\Sigma} \boldsymbol{\Sigma}'^{-1} (\boldsymbol{\mu} - \boldsymbol{\mu}') &= \boldsymbol{\mu}'^T \boldsymbol{\Sigma}'^{-1} \boldsymbol{\Sigma} \boldsymbol{\Sigma}'^{-1} \boldsymbol{\mu}' - 2\boldsymbol{\mu}'^T \boldsymbol{\Sigma}'^{-1} \boldsymbol{\Sigma} (\boldsymbol{\mu}^T \otimes \mathbf{1}) \text{vec } \boldsymbol{\Sigma}'^{-1} \\
 &+ (\text{vec } \boldsymbol{\Sigma}'^{-1})^T (\boldsymbol{\mu} \otimes \mathbf{1}) \boldsymbol{\Sigma} (\boldsymbol{\mu}^T \otimes \mathbf{1}) \text{vec } \boldsymbol{\Sigma}'^{-1} \tag{D26}
 \end{aligned}$$

$$\text{tr } \boldsymbol{\Sigma}'^{-1} \boldsymbol{\Sigma} \boldsymbol{\Sigma}'^{-1} \boldsymbol{\Sigma} = (\text{vec } \boldsymbol{\Sigma}'^{-1})^T (\boldsymbol{\Sigma} \otimes \boldsymbol{\Sigma}) \text{vec } \boldsymbol{\Sigma}'^{-1}, \tag{D27}$$

with the help of the third and fourth identities in (A1). Now (D25) can be neatly expressed as

$$I(g) = \frac{1}{2} \begin{bmatrix} \boldsymbol{\Sigma}'^{-1} \boldsymbol{\mu}' \\ \text{vec } \boldsymbol{\Sigma}'^{-1} \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\Sigma} & -\boldsymbol{\Sigma} (\boldsymbol{\mu}^T \otimes \mathbf{1}) \\ -(\boldsymbol{\mu} \otimes \mathbf{1}) \boldsymbol{\Sigma} & \frac{1}{2} (\boldsymbol{\Sigma} \otimes \boldsymbol{\Sigma}) + (\boldsymbol{\mu} \otimes \mathbf{1}) \boldsymbol{\Sigma} (\boldsymbol{\mu}^T \otimes \mathbf{1}) \end{bmatrix} \begin{bmatrix} \boldsymbol{\Sigma}'^{-1} \boldsymbol{\mu}' \\ \text{vec } \boldsymbol{\Sigma}'^{-1} \end{bmatrix}. \tag{D28}$$

This is the information contained in the Gaussian $\mathcal{N}(\boldsymbol{\mu}', \boldsymbol{\Sigma}')$ although it is conditioned by the choice of measure $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ used to define the inner product. Note that as $\boldsymbol{\Sigma}'^{-1}$ tends to zero, indicating a broadening of the distribution, the information also goes to zero. The expression (D28) can also be interpreted as simply writing the information using a different basis associated with the so-called *natural parameters* of a Gaussian [11].

Appendix E. Variational Inference Details

E.1. Fisher information matrix

This section reviews the *Fisher information matrix* (FIM) and shows that with respect to the coordinates used in a given subspace it is simply the Gram matrix of the chosen basis.

Let $q(\mathbf{x}|\theta) \in \mathcal{Q}$, a finite-dimensional subspace of \mathcal{B}^2 with basis B , depending on some parameter θ . The Fisher information on θ with respect to the measure ν is defined to be the covariance of the *score* [24], that is,

$$I_\theta = \mathbb{E}_\nu \left[\left(\frac{\partial \ln q}{\partial \theta} - \mathbb{E}_\nu \left[\frac{\partial \ln q}{\partial \theta} \right] \right)^2 \right] = \mathbb{E}_\nu \left[\left(\frac{\partial \ln q}{\partial \theta} \right)^2 \right] - \left(\mathbb{E}_\nu \left[\frac{\partial \ln q}{\partial \theta} \right] \right)^2. \tag{E1}$$

While our Fisher information may appear slightly unfamiliar, by taking the measure to be the density $\nu = \downarrow q$ then $\mathbb{E}_q[\partial \ln q / \partial \theta] = 0$ and we have the traditional version. We purposely delay setting $\nu = \downarrow q$ to show the connection to Bayes space.

Take q to be expressed as a normalized linear combination of the basis functions b_n , that is,

$$q(\mathbf{x}|\theta) = \downarrow \bigoplus_n \alpha_n(\theta) \cdot b_n. \tag{E2}$$

The score is

$$\frac{\partial \ln q}{\partial \theta} = \frac{1}{q} \frac{\partial q}{\partial \theta}. \tag{E3}$$

As $q = \prod_n b_n^{\alpha_n} / \int \prod_n b_n^{\alpha_n} d\mathbf{x}$,

$$\begin{aligned}
 \frac{\partial q}{\partial \theta} &= \frac{\partial}{\partial \theta} \left(\frac{\prod_n b_n^{\alpha_n}}{\int \prod_n b_n^{\alpha_n} d\mathbf{x}} \right) \\
 &= \sum_m \left(\frac{\partial \alpha_m}{\partial \theta} \right) \ln b_m \frac{\prod_n b_n^{\alpha_n}}{\int \prod_n b_n^{\alpha_n} d\mathbf{x}} - \frac{\prod_n b_n^{\alpha_n}}{\int \prod_n b_n^{\alpha_n} d\mathbf{x}} \frac{\sum_m (\partial \alpha_m / \partial \theta) \int \ln b_m \prod_n b_n^{\alpha_n}}{\int \prod_n b_n^{\alpha_n} d\mathbf{x}} \\
 &= q \sum_m \left(\frac{\partial \alpha_m}{\partial \theta} \right) (\ln b_m - \mathbb{E}_q[\ln b_m]). \tag{E4}
 \end{aligned}$$

Hence

$$\frac{\partial \ln q}{\partial \theta} = \sum_m \left(\frac{\partial \alpha_m}{\partial \theta} \right) (\ln b_m - \mathbb{E}_q[\ln b_m]). \tag{E5}$$

Substituting (E5) into (E1) produces

$$I_\theta = \sum_m \sum_n \left(\frac{\partial \alpha_m}{\partial \theta} \right) \left(\frac{\partial \alpha_n}{\partial \theta} \right) (\mathbb{E}_v[\ln b_m \ln b_n] - \mathbb{E}_v[\ln b_m] \mathbb{E}_v[\ln b_n]) = \left(\frac{\partial \alpha}{\partial \theta} \right)^T \langle \mathbf{b}, \mathbf{b} \rangle \left(\frac{\partial \alpha}{\partial \theta} \right). \tag{E6}$$

The traditional Fisher information uses $\downarrow q$ as the measure and we will indicate that explicitly with a subscript on the inner product, for example, $\langle \mathbf{b}, \mathbf{b} \rangle_q$. We also note that (E6) still holds in the event that q is not normalized, owing to the nature of the inner product.

We mention for interest that the stochastic derivative of $q(\mathbf{x}|\theta)$ with respect to θ is

$$\frac{\partial q}{\partial \theta} = \bigoplus_m \left(\frac{\partial \alpha_m}{\partial \theta} \right) \cdot b_m(\mathbf{x}) \tag{E7}$$

and so

$$I_\theta = \left(\frac{\partial \alpha}{\partial \theta} \right)^T \langle \mathbf{b}, \mathbf{b} \rangle \left(\frac{\partial \alpha}{\partial \theta} \right) = \left\langle \bigoplus_m \left(\frac{\partial \alpha_m}{\partial \theta} \right) \cdot b_m, \bigoplus_n \left(\frac{\partial \alpha_n}{\partial \theta} \right) \cdot b_n \right\rangle = \left\langle \frac{\partial q}{\partial \theta}, \frac{\partial q}{\partial \theta} \right\rangle, \tag{E8}$$

which makes the inner product expression of the Fisher information coordinate free.

For multiple parameters, $\theta_1, \theta_2 \dots \theta_K$, the (m, n) entry in the Fisher information matrix (FIM) is

$$\begin{aligned} I_{\theta, mn} &= \mathbb{E}_v \left[\left(\frac{\partial \ln q}{\partial \theta_m} - \mathbb{E}_v \left[\frac{\partial \ln q}{\partial \theta_m} \right] \right) \left(\frac{\partial \ln q}{\partial \theta_n} - \mathbb{E}_v \left[\frac{\partial \ln q}{\partial \theta_n} \right] \right) \right] \\ &= \mathbb{E}_v \left[\frac{\partial \ln q}{\partial \theta_m} \frac{\partial \ln q}{\partial \theta_n} \right] - \mathbb{E}_v \left[\frac{\partial \ln q}{\partial \theta_m} \right] \mathbb{E}_v \left[\frac{\partial \ln q}{\partial \theta_n} \right] \end{aligned} \tag{E9}$$

leading to

$$\mathbf{I}_\theta = \left(\frac{\partial \alpha}{\partial \theta} \right)^T \langle \mathbf{b}, \mathbf{b} \rangle \left(\frac{\partial \alpha}{\partial \theta} \right). \tag{E10}$$

We shall be particularly interested in the FIM with respect to the coordinates for a given basis, that is, when $\theta = \alpha$. In this case, the FIM is simply the Gram matrix,

$$\mathbf{I}_\alpha = \langle \mathbf{b}, \mathbf{b} \rangle. \tag{E11}$$

When q is used as the measure, we shall write $\mathbf{I}_\alpha = \langle \mathbf{b}, \mathbf{b} \rangle_q$.

E.2. Derivation of Eq. (36): Derivative of the measure in the inner product

We consider the inner product

$$\langle p, q \rangle_v = \mathbb{E}_v[\ln p \ln q] - \mathbb{E}_v[\ln p] \mathbb{E}_v[\ln q] \tag{E12}$$

with

$$v = \bigoplus_m \alpha_m \cdot b_m. \tag{E13}$$

We emphasize that here p and q are held fixed. The partial derivative with respect to α_n is

$$\frac{\partial}{\partial \alpha_n} \langle p, q \rangle_v = \frac{\partial}{\partial \alpha_n} \mathbb{E}_v[\ln p \ln q] - \left(\frac{\partial}{\partial \alpha_n} \mathbb{E}_v[\ln p] \right) \mathbb{E}_v[\ln q] - \mathbb{E}_v[\ln p] \left(\frac{\partial}{\partial \alpha_n} \mathbb{E}_v[\ln q] \right). \tag{E14}$$

In general,

$$\frac{\partial}{\partial \alpha_n} \mathbb{E}_v[\ln r] = \int \frac{\partial v}{\partial \alpha_n} \ln r d\mathbf{x} = \int v \frac{\partial \ln v}{\partial \alpha_n} \ln r d\mathbf{x} = \mathbb{E}_v \left[\frac{\partial \ln v}{\partial \alpha_n} \ln r \right]. \tag{E15}$$

(This quantity may in fact alternatively be written as $\langle b_n, r \rangle_v$.) The last two derivatives in (E14) are accounted for; as for the first, replacing $\ln r$ with $\ln p \ln q$ above, gives

$$\frac{\partial}{\partial \alpha_n} \mathbb{E}_v[\ln p \ln q] = \mathbb{E}_v \left[\frac{\partial \ln v}{\partial \alpha_n} \ln p \ln q \right]. \tag{E16}$$

Thus

$$\frac{\partial}{\partial \alpha_n} \langle p, q \rangle_v = \mathbb{E}_v \left[\frac{\partial \ln v}{\partial \alpha_n} \ln p \ln q \right] - \mathbb{E}_v \left[\frac{\partial \ln v}{\partial \alpha_n} \ln p \right] \mathbb{E}_v[\ln q] - \mathbb{E}_v[\ln p] \mathbb{E}_v \left[\frac{\partial \ln v}{\partial \alpha_n} \ln q \right]. \tag{E17}$$

Now we may rewrite this as

$$\frac{\partial}{\partial \alpha_n} \langle p, q \rangle_v = \mathbb{E}_v \left[\ln p \ln q^{\partial \ln v / \partial \alpha_n} \right] - \mathbb{E}_v[\ln p] \mathbb{E}_v \left[\ln q^{\partial \ln v / \partial \alpha_n} \right] - \mathbb{E}_v \left[\frac{\partial \ln v}{\partial \alpha_n} \ln p \right] \mathbb{E}_v[\ln q]. \tag{E18}$$

We recognize that $q^{\partial \ln v / \partial \alpha_n}$ is not a PDF; however, the self-normalizing feature of the inner product allows us to write

$$\mathbb{E}_v \left[\ln p \ln q^{\partial \ln v / \partial \alpha_n} \right] - \mathbb{E}_v[\ln p] \mathbb{E}_v \left[\ln q^{\partial \ln v / \partial \alpha_n} \right] = \left\langle p, q^{\partial \ln v / \partial \alpha_n} \right\rangle_v = \left\langle p, \frac{\partial \ln v}{\partial \alpha_n} \cdot q \right\rangle_v. \tag{E19}$$

For the last term in (E18), we use (E5) yielding

$$\mathbb{E}_v \left[\frac{\partial \ln v}{\partial \alpha_n} \ln p \right] \mathbb{E}_v[\ln q] = (\mathbb{E}_v[\ln b_n \ln p] - \mathbb{E}_v[\ln b_n] \mathbb{E}_v[\ln p]) \mathbb{E}_v[\ln q] = \mathbb{E}_v[\ln q] \langle b_n, p \rangle_v. \tag{E20}$$

Finally then

$$\frac{\partial}{\partial \alpha_n} \langle p, q \rangle_v = \left\langle p, \frac{\partial \ln v}{\partial \alpha_n} \cdot q \right\rangle_v - \mathbb{E}_v[\ln q] \langle b_n, p \rangle_v. \tag{E21}$$

As the inner product is symmetric in its arguments, this is also

$$\frac{\partial}{\partial \alpha_n} \langle p, q \rangle_v = \frac{\partial}{\partial \alpha_n} \langle q, p \rangle_v = \left\langle q, \frac{\partial \ln v}{\partial \alpha_n} \cdot p \right\rangle_v - \mathbb{E}_v[\ln p] \langle b_n, q \rangle_v. \tag{E22}$$

There is a caveat, however, in that we cannot transfer $\partial v / \partial \alpha_n$ as the coefficient of p to that of q ; this is because the coefficient is a function of the domain variables of the PDFs. That transformation, though, may be expressed as

$$\left\langle q, \frac{\partial \ln v}{\partial \alpha_n} \cdot p \right\rangle_v = \left\langle p, \frac{\partial \ln v}{\partial \alpha_n} \cdot q \right\rangle_v + \mathbb{E}_v[\ln p] \langle b_n, q \rangle_v - \mathbb{E}_v[\ln q] \langle b_n, p \rangle_v. \tag{E23}$$

We have used the shorthand $\langle p, q \rangle_{\partial v / \partial \alpha_n}$ to denote the derivative in (E21) as in (36).