

Geometric Variational Inference

Philipp Frank ^{1,2,*} , Reimar Leike ¹ and Torsten A. Enßlin ^{1,2} 

¹ Max-Planck Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85748 Garching, Germany; reimar@mpa-garching.mpg.de (R.L.); enssln@mpa-garching.mpg.de (T.A.E.)

² Faculty of Physics, Ludwig-Maximilians-Universität München, Geschwister-Scholl-Platz 1, 80539 München, Germany

* Correspondence: philipp@mpa-garching.mpg.de

Abstract: Efficiently accessing the information contained in non-linear and high dimensional probability distributions remains a core challenge in modern statistics. Traditionally, estimators that go beyond point estimates are either categorized as Variational Inference (VI) or Markov-Chain Monte-Carlo (MCMC) techniques. While MCMC methods that utilize the geometric properties of continuous probability distributions to increase their efficiency have been proposed, VI methods rarely use the geometry. This work aims to fill this gap and proposes geometric Variational Inference (geoVI), a method based on Riemannian geometry and the Fisher information metric. It is used to construct a coordinate transformation that relates the Riemannian manifold associated with the metric to Euclidean space. The distribution, expressed in the coordinate system induced by the transformation, takes a particularly simple form that allows for an accurate variational approximation by a normal distribution. Furthermore, the algorithmic structure allows for an efficient implementation of geoVI which is demonstrated on multiple examples, ranging from low-dimensional illustrative ones to non-linear, hierarchical Bayesian inverse problems in thousands of dimensions.

Keywords: variational methods; Bayesian inference; Fisher information metric; Riemann manifolds



Citation: Frank, P.; Leike, R.; Enßlin, T.A. Geometric Variational Inference. *Entropy* **2021**, *23*, 853. <https://doi.org/10.3390/e23070853>

Academic Editor:
Carlos Alberto De Bragança Pereira

Received: 21 May 2021
Accepted: 30 June 2021
Published: 2 July 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

In modern statistical inference and machine learning it is of utmost importance to access the information contained in complex and high dimensional probability distributions. In particular in Bayesian inference, it remains one of the key challenges to approximate samples from the posterior distribution, or the distribution itself, in a computationally fast and accurate way. Traditionally, there have been two distinct approaches towards this problem: the direct construction of posterior samples based on Markov Chain Monte-Carlo (MCMC) methods [1–3], and the attempt to approximate the probability distribution with a different one, chosen from a family of simpler distributions, known as variational inference (VI) [4–7] or variational Bayes' (VB) methods [8–10]. While MCMC methods are attractive due to their theoretical guarantees to reproduce the true distribution in the limit, they tend to be more expensive compared to variational alternatives. On the other hand, the family of distributions used in VI is typically chosen ad hoc. While VI aims to provide an appropriate approximation within the chosen family, the entire family may be a poor approximation to the true distribution.

In recent years, MCMC methods have been improved by incorporating geometric information of the posterior, especially by means of Riemannian manifold Hamilton Monte-Carlo (RMHMC) [11], a particular hybrid Monte-Carlo (HMC) [12,13] technique that constructs a Hamiltonian system on a Riemannian manifold with a metric tensor related to the Fisher information metric of the likelihood distribution and the curvature of the prior. For VI methods, however, the geometric structure of the true distribution has rarely been utilized to motivate and enhance the family of distributions used during optimization. One of the few examples being [14] where the Fisher metric has been used to reformulate the task of VI by means of α -divergencies in the mean-field setting.

In addition, a powerful variational approximation technique for the family of normal distributions utilizing infinitesimal geometric properties of the posterior is Metric Gaussian Variational Inference (MGVI) [15]. In MGVI the family is parameterized in terms of the mean m , and the covariance matrix is set to the inverse of the metric tensor evaluated at m . This choice ensures that the true distribution and the approximation obtain the same geometric properties infinitesimally, i.e., at the location of the mean m . In this work we extend the geometric correspondence used by MGVI to be valid not only at m , but also in a local neighborhood of m . We achieve this extension by means of an invertible coordinate transformation from the coordinate system used within MGVI, in which the curvature of the prior is the identity, to a new coordinate system in which the metric of the posterior becomes (approximately) the Euclidean metric. We use a normal distribution in these coordinates as the approximation to the true distribution and thereby establish a non-Gaussian posterior in the MGVI coordinate system. The resulting algorithm, called geometric Variational Inference (geoVI) can be computed efficiently and is inherently similar to the implementation of MGVI. This is not by mere coincidence: To linear order, geoVI reproduces MGVI. In this sense, the geoVI algorithm is a non-linear generalization of MGVI that captures the geometric properties encoded in the posterior metric not only infinitesimally, but also in a local neighborhood of this point. We include an implementation of the proposed geoVI algorithm into the software package Numerical Information Field Theory (NIFTy [16]), a versatile library for signal inference algorithms.

Mathematical Setup

Throughout this work, we consider the joint distribution $P(d, s)$ of observational data $d \in \Omega$ and the unknown, to be inferred signal s . This distribution is factorized into the likelihood of observing the data, given the signal $P(d|s)$, and the prior distribution $P(s)$. In general, only a subset of the signal, denoted as s' , may be directly constrained by the likelihood, such that $P(d|s) = P(d|s')$, and therefore there may be additional hidden variables in s , that are unobserved by the data, but part of the prior model. Thus, the prior distribution $P(s)$ may possess a hierarchical structure that summarizes our knowledge about the system prior to the measurement, and s represents everything in the system that is of interest to us, but about which our knowledge is uncertain a priori. We do not put any constraints on the functional form of $P(s)$, and assume that the signal s solely consists of continuous real valued variables, i.e., $s \in X \subset \mathbb{R}^M$. This enables us to regard s as coordinates of the space on which $P(s)$ is defined and to use geometric concepts such as coordinate transformations to represent probability distributions in different coordinate systems. Probability densities transform in a probability mass preserving fashion. Specifically let $f : \mathbb{R}^M \rightarrow X$ be an invertible function, and let $s = f(\xi)$. Then the distributions $P(s)$ and $P(\xi)$ relate via

$$\int P(s) ds = \int P(\xi) d\xi . \quad (1)$$

This allows us to express $P(s)$ by means of the pushforward of $P(\xi)$ by f . We denote the pushforward as

$$P(s) = (f \star P(\xi))(s) = \int \delta(s - f(\xi)) P(\xi) d\xi = \left(P(\xi) \left\| \frac{df}{d\xi} \right\|^{-1} \right) \Big|_{\xi=f^{-1}(s)} . \quad (2)$$

Under mild regularity conditions on the prior distribution, there always exists an f that relates the complex hierarchical form of $P(s)$ to a simple distribution $P(\xi)$ [17]. We choose f such that $P(\xi)$ takes the form of a normal distribution with zero mean and unit covariance and call such a distribution a *standard distribution*:

$$P(\xi) = \mathcal{N}(\xi; 0, 1) , \quad (3)$$

where $\mathcal{N}(\xi; m, D)$ denotes a multivariate normal distribution in the random variables ξ with mean m and covariance D .

We may express the likelihood in terms of ξ as

$$P(d|\xi) \equiv P(d|s' = f'(\xi)), \quad (4)$$

where f' is the part of f that maps onto the observed quantities s' . In general, f' is a non-invertible function and is commonly referred to as *generative model* or *generative process*, as it encodes all the information necessary to transform a standard distribution into the observed quantities, subject to our prior beliefs. Using Equation (4) we get by means of Bayes' theorem, that the posterior takes the form

$$P(\xi|d) = \frac{P(\xi, d)}{P(d)} = \frac{P(d|\xi) \mathcal{N}(\xi; 0, \mathbb{1})}{P(d)}. \quad (5)$$

Using the push-forward of the posterior, we can recover the posterior statistics of s via

$$P(s|d) = (f \star P(\xi|d))(s), \quad (6)$$

which means that we can fully recover the posterior properties of s , which typically has a physical interpretation as opposed to ξ . In particular Equation (6) implies that if we are able to draw samples from $P(\xi|d)$ we can simply generate posterior samples for s since $s = f(\xi)$.

2. Geometric Properties of Posterior Distributions

In order to access the information contained in the posterior distribution $P(\xi|d)$, in this work, we wish to exploit the geometric properties of the posterior, in particular with the help of Riemannian geometry. Specifically, we define a Riemannian manifold using a metric tensor, related to the Fisher Information metric of the likelihood and a metric for the prior, and establish a (local) isometry of this manifold to Euclidean space. The associated coordinate transformation gives rise to a coordinate system in which, hopefully, the posterior takes a simplified form despite the fact that probabilities do not transform in the same way as metric spaces do. As we will see, in cases where the isometry is global, and in addition the transformation is (almost) volume-preserving, the complexity of the posterior distribution can be absorbed (almost) entirely into this transformation.

To begin our discussion, we need to define an appropriate metric for posterior distributions. To this end, consider the negative logarithm of the posterior, sometimes also referred to as information Hamiltonian, which takes the form

$$\mathcal{H}(\xi|d) \equiv -\log(P(\xi|d)) = \mathcal{H}(d|\xi) + \mathcal{H}(\xi) - \mathcal{H}(d). \quad (7)$$

A common choice to extract geometric information from this Hamiltonian is the Hessian \mathcal{C} of \mathcal{H} . Specifically

$$\mathcal{C}(\xi) \equiv \frac{\partial^2 \mathcal{H}(\xi|d)}{\partial \xi \partial \xi'} = \frac{\partial^2 \mathcal{H}(d|\xi)}{\partial \xi \partial \xi'} + \mathbb{1} \equiv \mathcal{C}_{d|\xi}(\xi) + \mathbb{1}, \quad (8)$$

where the identity matrix arises from the curvature of the prior (information Hamiltonian). While \mathcal{C} provides information about the local geometry, it turns out to be unsuited for our approach to construct a coordinate transformation, as it is not guaranteed to be positive definite for all ξ . An alternative, positive definite, measure for the curvature can be obtained by replacing the Hessian of the likelihood with its Fisher information metric [18], defined as

$$\mathcal{M}_{d|\xi}(\xi) = \left\langle \frac{\partial \mathcal{H}}{\partial \xi} \frac{\partial \mathcal{H}}{\partial \xi'} \right\rangle_{P(d|\xi)} = \left\langle \frac{\partial^2 \mathcal{H}(d|\xi)}{\partial \xi \partial \xi'} \right\rangle_{P(d|\xi)} = \left\langle \mathcal{C}_{d|\xi}(\xi) \right\rangle_{P(d|\xi)}. \quad (9)$$

The Fisher information metric can be understood as a Riemannian metric defined over the statistical manifold associated with the likelihood [19], and is a core element in the field of information geometry [20] as it provides a distance measure between probability distributions [21]. Replacing $\mathcal{C}_{d|\xi}$ with $\mathcal{M}_{d|\xi}$ in Equation (8) we find

$$\mathcal{M}(\xi) \equiv \mathcal{M}_{d|\xi}(\xi) + \mathbb{1} = \left\langle \mathcal{C}_{d|\xi}(\xi) \right\rangle_{P(d|\xi)} + \mathbb{1} = \langle \mathcal{C}(\xi) \rangle_{P(d|\xi)}, \quad (10)$$

which, from now on, we refer to as the metric \mathcal{M} . As the Fisher metric of the likelihood is a symmetric, positive-semidefinite matrix, we get that \mathcal{M} is a symmetric, positive-definite matrix for all ξ . It is noteworthy that upon insertion, we find that the metric \mathcal{M} is defined as the expectation value of the Hessian of the posterior Hamiltonian \mathcal{C} w.r.t. the likelihood $P(d|\xi)$. Therefore, in some way, we may regard \mathcal{M} as the measure for the curvature in case the observed data d is unknown, and the only information given is the structure of the model itself, as encoded in $P(d|\xi)$. This connection is only of qualitative nature, but it highlights a key limitation of \mathcal{M} when used as the defining property of the posterior geometry. From a Bayesian perspective, only the data d that is actually observed is of relevance as the posterior is conditioned on d . Therefore a curvature measure that arises from marginalization over the data must be sub-optimal compared to a measure conditional to the data, as it ignores the local information that we gain from observing d . Nevertheless, we find that in many practical applications \mathcal{M} encodes enough relevant information about the posterior geometry that it provides us with a valuable metric to construct a coordinate transformation. It is noteworthy that attempts have been provided to resolve this issue via a more direct approach to recover a positive definite matrix from the Hessian of the posterior while retaining the local information of the data. e.g., in [22], the SoftAbs non-linearity is applied to the Hessian and the resulting positive definite matrix is used as a curvature measure. In our practical applications, however, we are particularly interested in solving very high dimensional problems, and applying a matrix non-linearity is currently too expensive to give rise to a scalable algorithm for our purposes. Therefore we rely on the metric \mathcal{M} as a measure for the curvature of the posterior, and leave possible extensions to future research.

2.1. Coordinate Transformation

Our goal is to construct a coordinate system y and an associated transformation g , that maps from ξ to y , in which the posterior metric \mathcal{M} takes the form of the identity matrix $\mathbb{1}$. The motivation is that if \mathcal{M} captures the geometric properties of the posterior, a coordinate system in which this metric becomes trivial should also be a coordinate system in which the posterior distribution takes a particularly simple form. For an illustrative example see Figure 1. To do so, we require the Fisher metric of the likelihood $\mathcal{M}_{d|\xi}$ to be the pullback of the Euclidean metric. Specifically we propose a function $x(\xi)$ such that

$$\mathcal{M}_{d|\xi} \stackrel{!}{=} \left(\frac{\partial x}{\partial \xi} \right)^T \frac{\partial x}{\partial \xi}, \quad (11)$$

where T denotes the adjoint of a matrix. As outlined in Appendix A, for many practically relevant likelihoods such a decomposition is possible by means of an inexpensive to evaluate function x (Here with “inexpensive” we mean that applying the function $x(\xi)$ has a similar computational cost compared to applying the likelihood function $P(d|\xi)$ to a specific ξ). Given x , we can rewrite the posterior metric \mathcal{M} as

$$\mathcal{M} = \left(\frac{\partial x}{\partial \xi} \right)^T \frac{\partial x}{\partial \xi} + \mathbb{1}. \quad (12)$$

In order to relate this metric to Euclidean space, we aim to find the isometry g that relates the Riemannian manifold associated with the metric \mathcal{M} to Euclidean space. Specifically we seek to find an invertible function g satisfying

$$\mathcal{M}(\xi) = \left(\frac{\partial x}{\partial \xi}\right)^T \frac{\partial x}{\partial \xi} + \mathbb{1} \stackrel{!}{=} \left(\frac{\partial g}{\partial \xi}\right)^T \frac{\partial g}{\partial \xi}. \tag{13}$$

In general, i.e., for a general function $x(\xi)$, however, this decomposition does not exist globally. Nevertheless, there exists a transformation $g(\xi; \bar{\xi})$ based on an approximative Taylor series around an expansion point $\bar{\xi}$, that results in a metric $\tilde{\mathcal{M}}(\xi)$ such that

$$\mathcal{M}(\xi) \approx \tilde{\mathcal{M}}(\xi) \equiv \left(\frac{\partial g(\xi; \bar{\xi})}{\partial \xi}\right)^T \frac{\partial g(\xi; \bar{\xi})}{\partial \xi}, \tag{14}$$

in the vicinity of $\bar{\xi}$. This transformation g can be obtained up to an integration constant by Taylor expanding Equation (14) around $\bar{\xi}$ and solving for the Taylor coefficients of g in increasing order. We express g in terms of its Taylor series using the Einstein sum convention as

$$g(\xi; \bar{\xi})^i = \bar{g}^i + \bar{g}^i_{,j} (\xi - \bar{\xi})^j + \bar{g}^i_{,jk} (\xi - \bar{\xi})^j (\xi - \bar{\xi})^k + \dots, \tag{15}$$

where repeated indices get summed over, $a_{,i}$ denotes the partial derivative of a w.r.t. the i th component of ξ , and \bar{s} denotes a (tensor) field $s(\xi)$, evaluated at the expansion point $\bar{\xi}$. We begin to expand Equation (14) around $\bar{\xi}$ and obtain for the zeroth order

$$\bar{\mathcal{M}}_{ij} \equiv \bar{x}^\alpha_{,i} \bar{x}^\alpha_{,j} + \delta_{ij} \stackrel{!}{=} \bar{g}^\alpha_{,i} \bar{g}^\alpha_{,j}. \tag{16}$$

Expanding Equation (14) to first order yields

$$\bar{x}^\alpha_{,ik} \bar{x}^\alpha_{,j} + \bar{x}^\alpha_{,i} \bar{x}^\alpha_{,jk} \stackrel{!}{=} \bar{g}^\alpha_{,ik} \bar{g}^\alpha_{,j} + \bar{g}^\alpha_{,i} \bar{g}^\alpha_{,jk}, \tag{17}$$

and therefore

$$\bar{g}^i_{,jk} = \tilde{\mathcal{M}}^{i\gamma} \bar{g}^\beta_{,\gamma} \bar{x}^\alpha_{,\beta} \bar{x}^\alpha_{,jk}, \tag{18}$$

where $\tilde{\mathcal{M}}^{ij} = (\bar{\mathcal{M}}^{-1})_{ij}$ denotes the components of the inverse of $\bar{\mathcal{M}}$.

Thus, to first order in the metric (meaning to second order in the transformation) the expansion remains solvable for a general x . Proceeding with the second order, however, we get that

$$\bar{x}^\alpha_{,ikl} \bar{x}^\alpha_{,j} + \bar{x}^\alpha_{,ik} \bar{x}^\alpha_{,jl} + \bar{x}^\alpha_{,il} \bar{x}^\alpha_{,jk} + \bar{x}^\alpha_{,i} \bar{x}^\alpha_{,jkl} \stackrel{!}{=} \bar{g}^\alpha_{,ikl} \bar{g}^\alpha_{,j} + \bar{g}^\alpha_{,ik} \bar{g}^\alpha_{,jl} + \bar{g}^\alpha_{,il} \bar{g}^\alpha_{,jk} + \bar{g}^\alpha_{,i} \bar{g}^\alpha_{,jkl}, \tag{19}$$

which does not exhibit a general solution for $\bar{g}^i_{,jkl}$ in higher dimensions due to the fact that the third derivative has to be invariant under arbitrary permutation of the latter three indices jkl . However, in analogy to Equation (18), we may set

$$\bar{g}^i_{,jkl} = \tilde{\mathcal{M}}^{i\gamma} \bar{g}^\beta_{,\gamma} \bar{x}^\alpha_{,\beta} \bar{x}^\alpha_{,jkl}, \tag{20}$$

which cancels the first and the last term of Equation (19), and study the remaining error which takes the form

$$\begin{aligned} \epsilon_{ijkl} &= \bar{x}^\alpha_{,ik} \bar{x}^\alpha_{,jl} + \bar{x}^\alpha_{,il} \bar{x}^\alpha_{,jk} - \bar{g}^\alpha_{,ik} \bar{g}^\alpha_{,jl} - \bar{g}^\alpha_{,il} \bar{g}^\alpha_{,jk} \\ &= \bar{x}^\alpha_{,ik} \bar{x}^\alpha_{,jl} + \bar{x}^\alpha_{,il} \bar{x}^\alpha_{,jk} - \bar{x}^\alpha_{,ik} \bar{x}^\gamma_{,\alpha} \tilde{\mathcal{M}}^{\gamma\delta} \bar{x}^\delta_{,\beta} \bar{x}^\beta_{,jl} - \bar{x}^\alpha_{,il} \bar{x}^\gamma_{,\alpha} \tilde{\mathcal{M}}^{\gamma\delta} \bar{x}^\delta_{,\beta} \bar{x}^\beta_{,jk} \\ &= \bar{x}^\alpha_{,ik} \left(\delta_{\alpha\beta} - \bar{x}^\gamma_{,\alpha} \tilde{\mathcal{M}}^{\gamma\delta} \bar{x}^\delta_{,\beta} \right) \bar{x}^\beta_{,jl} + \bar{x}^\alpha_{,il} \left(\delta_{\alpha\beta} - \bar{x}^\gamma_{,\alpha} \tilde{\mathcal{M}}^{\gamma\delta} \bar{x}^\delta_{,\beta} \right) \bar{x}^\beta_{,jk}. \end{aligned} \tag{21}$$

Let $X^i_j \equiv \bar{x}^i_{,j}$, the expression in the parentheses takes the form

$$\mathbb{1} - X \tilde{\mathcal{M}}^{-1} X^T = \mathbb{1} - X \left(\mathbb{1} + X^T X \right)^{-1} X^T = \left(\mathbb{1} + X X^T \right)^{-1} \equiv M, \tag{22}$$

and thus Equation (21) reduces to

$$\epsilon_{ijkl} = \bar{x}_{,ik}^\alpha M_{\alpha\beta} \bar{x}_{,jl}^\beta + \bar{x}_{,il}^\alpha M_{\alpha\beta} \bar{x}_{,jk}^\beta. \tag{23}$$

The impact of this error contribution can be qualitatively studied using the spectrum $\lambda(M)$ of the matrix M . This spectrum may exhibit two extreme cases, a so-called likelihood dominated regime, where the spectrum $\lambda(XX^T) \gg 1$, and a prior dominated regime where $\lambda(XX^T) \ll 1$. In the likelihood dominated regime, we get that $\lambda(M) \ll 1$ and thus the contribution of the error is small, whereas in the prior dominated regime $\lambda(M) \approx 1$ which yields an $\mathcal{O}(1)$ error. However, in the prior dominated regime, the entire metric \mathcal{M} is close to the identity as we are in the standard coordinate system of the prior ζ and therefore higher order derivatives of x are small. As a consequence, the error is of the order $\mathcal{O}(1)$ only in regimes where the third (and higher) order of the expansion is negligible compared to the first and second order. An exception occurs when the expansion point is close to a saddle point of x , i.e., in cases where the first derivative of x becomes small (and therefore the metric is close to the identity), but higher order derivatives of x may be large. For the moment, we proceed under the assumption that the change of x , as a function of ζ , is sufficiently monotonic throughout the expansion regime. We discuss the implications of violating this assumption in Section 5.3.




Figure 1. Non-linear posterior distribution $P(\zeta|d)$ in the standard coordinate system of the prior distribution ζ (left) and the transformed distribution $P(y|d)$ (right) in the coordinate system y where the posterior metric becomes (approximately) the identity matrix. $P(y|d)$ is obtained from $P(\zeta|d)$ via the push-forward through the transformation g which relates the two coordinate systems. The functional form of g is derived in Section 2.1 and depends on an expansion point $\bar{\zeta}$ (orange dot in the left image), and g is set up such that $\bar{\zeta}$ coincides with the origin in y . To visualize the transformation, the coordinate lines of y (black mesh grid on the right) are transformed back into ζ -coordinates using the inverse coordinate transformation g^{-1} and are displayed as a black mesh in the original space on the left. In addition, note that while the transformed posterior $P(y|d)$ arguably takes a simpler form compared to $P(\zeta|d)$, it does not become trivial (e.g., identical to a standard distribution) as there remain small asymmetries in the posterior density. There are multiple reasons for these deviations which are discussed in more detail in Section 2.2 once we established how the transformation g is constructed.

If we proceed to higher order expansions of Equation (14), we notice that a repetitive picture emerges: The leading order derivative tensor $\bar{g}^i_{,j\dots}$ that appears in the expansion may be set in direct analogy to Equations (18) and (20) as

$$\bar{g}^i_{,j\dots} = \bar{\mathcal{M}}^{i\gamma} \bar{g}^{\beta}_{,\gamma} \bar{x}^{\alpha}_{,\beta} \bar{x}^{\alpha}_{,j\dots}, \tag{24}$$

where \dots denotes the higher order derivatives. The remaining error contributions at each order take a similar form as in Equation (23), where the matrix M reappears in between all possible combinations of the remaining derivatives of x that appear using the Leibniz rule. Note that for increasing order, the number of terms that contribute to the error also increases. Specifically for the n th order expansion of Equation (14) we get $m = \sum_{k=1}^{n-1} \binom{n}{k}$ contributions to the error. Therefore, even if each individual contribution by means of M is small, the expansion error eventually becomes large once high order expansions become relevant. Therefore the proposed approximation only remains locally valid around $\bar{\zeta}$.

Nevertheless, we may proceed to combine the derivative tensors of g determined above in order to get the Jacobian of the transformation g as

$$\begin{aligned} g^i_{,j}(\zeta) &\equiv \bar{g}^i_{,j} + \bar{g}^i_{,jk}(\zeta - \bar{\zeta})^k + \frac{1}{2} \bar{g}^i_{,jkl}(\zeta - \bar{\zeta})^k(\zeta - \bar{\zeta})^l + \dots \\ &= \bar{g}^i_{,j} + \bar{\mathcal{M}}^{i\alpha} \bar{g}^{\beta}_{,\alpha} \bar{x}^{\gamma}_{,\beta} \left(\bar{x}^{\gamma}_{,jk}(\zeta - \bar{\zeta})^k + \frac{1}{2} \bar{x}^{\gamma}_{,jkl}(\zeta - \bar{\zeta})^k(\zeta - \bar{\zeta})^l + \dots \right), \end{aligned} \tag{25}$$

or equivalently

$$\bar{g}^{\alpha}_{,i} \bar{g}^{\alpha}_{,j}(\zeta) = \delta_{ij} + \bar{x}^{\alpha}_{,i} \left(\bar{x}^{\alpha}_{,j} + \bar{x}^{\alpha}_{,jk}(\zeta - \bar{\zeta})^k + \frac{1}{2} \bar{x}^{\alpha}_{,jkl}(\zeta - \bar{\zeta})^k(\zeta - \bar{\zeta})^l + \dots \right). \tag{26}$$

From the zeroth order, Equation (16), we get that $\bar{g}^i_{,j} = \left(\sqrt{\bar{\mathcal{M}}}\right)^i_j$ up to a unitary transformation, and we can sum up the Taylor series in x of Equation (26) to arrive at an index free representation of the Jacobian as

$$\frac{\partial g}{\partial \zeta} = \sqrt{\bar{\mathcal{M}}}^{-1} \left(\mathbb{1} + \left(\frac{\partial x}{\partial \zeta} \Big|_{\bar{\zeta}} \right)^T \frac{\partial x}{\partial \zeta} \right). \tag{27}$$

Upon integration, this yields a transformation

$$g(\zeta) - g(\bar{\zeta}) = \sqrt{\bar{\mathcal{M}}}^{-1} \left(\zeta - \bar{\zeta} + \left(\frac{\partial x}{\partial \zeta} \Big|_{\bar{\zeta}} \right)^T (x(\zeta) - x(\bar{\zeta})) \right). \tag{28}$$

The resulting transformation takes an intuitive form: The approximation to the distance between a point $g(\zeta)$ and the transformed expansion point $g(\bar{\zeta})$ consists of the distance w.r.t. the prior measure $(\zeta - \bar{\zeta})$ and the distance w.r.t. the likelihood measure $(x(\zeta) - x(\bar{\zeta}))$, back-projected into the prior domain using the local transformation at $\bar{\zeta}$. Finally, the metric at $\bar{\zeta}$ is used as a measure for the local curvature. Equation (28) is only defined up to an integration constant, and therefore, without loss of generality, we may set $g(\bar{\zeta}) = 0$ to obtain the final approximative coordinate transformation as

$$g(\zeta; \bar{\zeta}) = \sqrt{\bar{\mathcal{M}}}^{-1} \left(\zeta - \bar{\zeta} + \left(\frac{\partial x}{\partial \zeta} \Big|_{\bar{\zeta}} \right)^T (x(\zeta) - x(\bar{\zeta})) \right) \equiv \sqrt{\bar{\mathcal{M}}}^{-1} \tilde{g}(\zeta; \bar{\zeta}). \tag{29}$$

2.2. Basic Properties

In order to study a few basic properties of this transformation, for simplicity, we first consider a posterior distribution with a metric that allows for an exact isometry g_{iso} to

Euclidean space. Specifically let g_{iso} be a coordinate transformation satisfying Equation (13). The posterior distribution in coordinates $y = g_{\text{iso}}(\xi)$ is given via the push-forward of $P(\xi|d)$ through g_{iso} as

$$P(y|d) \propto (g_{\text{iso}} \star P(\xi|d))(y) = \left(P(\xi|d) \left\| \frac{\partial g_{\text{iso}}}{\partial \xi} \right\|^{-1} \right) \Big|_{\xi=g_{\text{iso}}^{-1}(y)} = \frac{P(\xi|d)}{\sqrt{|\mathcal{M}(\xi)|}} \Big|_{\xi=g_{\text{iso}}^{-1}(y)}, \tag{30}$$

and the information Hamiltonian takes the form

$$\mathcal{H}(y|d) = \left(\mathcal{H}(\xi|d) + \frac{1}{2} \log(|\mathcal{M}(\xi)|) \right) \Big|_{\xi=g_{\text{iso}}^{-1}(y)} + \mathcal{H}_0 \equiv \tilde{\mathcal{H}}(\xi = g_{\text{iso}}^{-1}(y)) + \mathcal{H}_0, \tag{31}$$

where \mathcal{H}_0 denotes y independent contributions. We may study the curvature of the posterior in coordinates y given as:

$$\mathcal{C}(y) = \frac{\partial \xi}{\partial y} \left(\frac{\partial^2 \tilde{\mathcal{H}}(\xi)}{\partial \xi \partial \xi'} \right) \left(\frac{\partial \xi'}{\partial y'} \right)^T + \frac{\partial \tilde{\mathcal{H}}(\xi)}{\partial \xi} \frac{\partial^2 \xi}{\partial y \partial y'} \quad \text{with} \tag{32}$$

$$\xi = g_{\text{iso}}^{-1}(y) \quad \text{and} \quad \xi' = g_{\text{iso}}^{-1}(y'),$$

which we can use to construct a metric $\mathcal{M}(y)$ in analogy to Equation (10) by taking the expectation value of the curvature w.r.t. the likelihood. This yields

$$\begin{aligned} \mathcal{M}(y) &= \frac{\partial \xi}{\partial y} \mathcal{M}(\xi) \left(\frac{\partial \xi'}{\partial y'} \right)^T + \frac{1}{2} \frac{\partial \xi}{\partial y} \left(\frac{\partial^2 \log(|\mathcal{M}(\xi)|)}{\partial \xi \partial \xi'} \right) \left(\frac{\partial \xi'}{\partial y'} \right)^T + \left\langle \frac{\partial \tilde{\mathcal{H}}(\xi)}{\partial \xi} \right\rangle_{P(d|\xi)} \frac{\partial^2 \xi}{\partial y \partial y'} \\ &\equiv \frac{\partial \xi}{\partial y} \mathcal{M}(\xi) \left(\frac{\partial \xi'}{\partial y'} \right)^T + \mathcal{R}(y) = \mathbb{1} + \mathcal{R}(y). \end{aligned} \tag{33}$$

The first terms yields the identity, as it is the defining property of g_{iso} . Furthermore, in case we are able to say that $\mathcal{R}(y)$ is small compared to the identity, we notice that the quantity $\mathcal{M}(\xi)$ (Equation (10)), that we referred to as the posterior metric, approximately transforms like a proper metric under g_{iso} . In this case we find that the isometry g_{iso} between the Riemannian manifold associated with $\mathcal{M}(\xi)$ and the Euclidean space is also a transformation that removes the complex geometry of the posterior. To further study $\mathcal{R}(y)$, we consider its two contributions separately, where for the first part, the log-determinant (or logarithmic volume), we get that it becomes small compared to the identity if

$$\mathcal{M}(\xi) \gg \frac{1}{2} \left(\frac{\partial^2 \log(|\mathcal{M}(\xi)|)}{\partial \xi \partial \xi'} \right). \tag{34}$$

Therefore, the curvature of the log-determinant of the metric has to be much smaller than the metric itself. To study the second term of $\mathcal{R}(y)$, we may again split the discussion into a prior and a likelihood dominated regime, depending on the ξ at which we evaluate the expression. In a prior dominated regime we get that

$$\frac{\partial^2 \xi}{\partial y \partial y'} \approx 0, \tag{35}$$

as the metric is close to the identity in this regime (and therefore $\xi \approx y$). In a likelihood dominated regime we get that $\tilde{\mathcal{H}} \approx \mathcal{H}(d|\xi)$ and therefore

$$\left\langle \frac{\partial \tilde{\mathcal{H}}(\xi)}{\partial \xi} \right\rangle_{P(d|\xi)} \approx \left\langle \frac{\partial \mathcal{H}(d|\xi)}{\partial \xi} \right\rangle_{P(d|\xi)} = - \left\langle \frac{1}{P(d|\xi)} \frac{\partial P(d|\xi)}{\partial \xi} \right\rangle_{P(d|\xi)} = 0. \tag{36}$$

So at least in a prior dominated regime, as well as a likelihood dominated regime, the posterior Hamiltonian transforms in an analogous way as the manifold, under the transformation g_{iso} , if Equation (34) also holds true.

For a practical application, however, in all but the simplest cases the isometry g_{iso} is not accessible, or might not even exist. Therefore, in general, we have to use the approximation $g(\xi; \bar{\xi})$, as defined in Equation (29), instead. We may express the transformation of the metric using $g(\xi; \bar{\xi})$, and find that

$$\begin{aligned} \mathcal{M}(y) = & \sqrt{\mathcal{M}} \left(\mathbb{1} + \left(\frac{\partial x}{\partial \bar{\xi}} \right)^T \frac{\partial x}{\partial \bar{\xi}} \Big|_{\bar{\xi}} \right)^{-1} \left(\mathbb{1} + \left(\frac{\partial x}{\partial \bar{\xi}} \right)^T \frac{\partial x'}{\partial \bar{\xi}'} \right) \left(\mathbb{1} + \left(\frac{\partial x}{\partial \bar{\xi}} \right)^T \frac{\partial x'}{\partial \bar{\xi}'} \right)^{-1} \sqrt{\mathcal{M}} \\ & + \tilde{\mathcal{R}}(y), \end{aligned} \tag{37}$$

now with $\xi = g^{-1}(y; \bar{\xi})$ and analogous for ξ' . $\tilde{\mathcal{R}}$ is defined by replacing g_{iso} with g for the entire expression of \mathcal{R} . We notice that this transformation does not yield the identity, except when evaluated at the expansion point $\xi = \bar{\xi}$. Therefore, in addition to the error $\tilde{\mathcal{R}}$ there is a deviation from the identity related to the expansion error as one moves away from $\bar{\xi}$.

At this point we would like to emphasize that the posterior Hamiltonian \mathcal{H} and the Riemannian manifold constructed from \mathcal{M} are only loosely connected due to the errors described by $\tilde{\mathcal{R}}$ and the additional expansion error. They are arguably small in many cases and in the vicinity of $\bar{\xi}$, but we do not want to claim that this correspondence is valid in general (see Section 5.3). Nevertheless, we find that in many cases this correspondence works well in practice. Some illustrative examples are given in Section 3.1.2.

3. Posterior Approximation

Utilizing the derived coordinate transformation for posterior approximation is mainly based on the idea that in the transformed coordinate system, the posterior takes a simpler form. In particular we aim to remove parts (if not most) of the complex geometry of the posterior, such that a simple probability distribution, e.g., a Gaussian distribution, yields a good approximation.

3.1. Direct Approximation

Assuming that all the errors discussed in the previous section are small enough, we may attempt to directly approximate the posterior distribution via a unit Gaussian in the coordinates y as in this case the transformed metric $\mathcal{M}(y)$ is close to the identity. As the coordinate transformation g , defined via Equation (29), is only known up to an integration constant by construction, the posterior approximation is achieved by a shifted unit Gaussian in y . This shift needs to be determined, which we can do by maximizing the transformed posterior distribution

$$P(y|d) \propto \left(P(\xi, d) \left| \left| \frac{\partial g}{\partial \bar{\xi}} \right| \right|^{-1} \right) \Big|_{\xi=g^{-1}(y; \bar{\xi})}, \tag{38}$$

w.r.t. y . Here $g^{-1}(y; \bar{\xi})$ denotes the inverse of $g(\xi; \bar{\xi})$ w.r.t. its first argument. Equivalently we can minimize the information Hamiltonian $\mathcal{H}(y|d)$, defined as

$$\mathcal{H}(y|d) \equiv -\log(P(y|d)) = \left(\mathcal{H}(\xi, d) + \frac{1}{2} \log(|\mathcal{M}|) \right) \Big|_{\xi=g^{-1}(y; \bar{\xi})} \equiv \tilde{\mathcal{H}}(\xi = g^{-1}(y; \bar{\xi})). \tag{39}$$

Minimizing $\mathcal{H}(y|d)$ yields the maximum a posterior solution y^* which, in case the posterior is close to a unit Gaussian in the coordinates y , can be identified with the shift in y . As g is an invertible function, we may instead minimize $\tilde{\mathcal{H}}$ w.r.t. ξ and apply g to the result in order to obtain y^* . Specifically

$$y^* \equiv \underset{y}{\operatorname{argmin}}(\mathcal{H}(y|d)) = g \left(\underset{\xi}{\operatorname{argmin}}(\tilde{\mathcal{H}}(\xi)) \right). \tag{40}$$

Therefore we can circumvent the inversion of g at any point during optimization. Now suppose that we use any gradient based optimization scheme to minimize for ζ , starting from some initial position ζ^0 . If we set the expansion point $\bar{\zeta}$, used to construct g , to be equal to ζ^0 , we notice that

$$\tilde{\mathcal{M}}(\bar{\zeta}) = \mathcal{M}(\bar{\zeta}) \quad (41)$$

$$\left. \frac{\partial \tilde{\mathcal{M}}}{\partial \bar{\zeta}} \right|_{\bar{\zeta}=\bar{\zeta}} = \left. \frac{\partial \mathcal{M}}{\partial \zeta} \right|_{\zeta=\bar{\zeta}}, \quad (42)$$

as the expansion of the metric is valid to first order by construction. Therefore if we set the expansion point $\bar{\zeta}$ to the current estimate of ζ after every step, we can replace the approximated metric $\tilde{\mathcal{M}}$ with the true metric \mathcal{M} and arrive at an optimization objective of the form

$$\bar{\zeta} = \underset{\zeta}{\operatorname{argmin}} \left(\mathcal{H}(\zeta, d) + \frac{1}{2} \log(|\mathcal{M}(\zeta)|) \right). \quad (43)$$

Note that $g(\bar{\zeta}; \bar{\zeta}) = 0$ by construction, and therefore $y^* = 0$, as there is a degeneracy between a shift in y and a change of the expansion point $\bar{\zeta}$. Once the optimal expansion point $\bar{\zeta}$ is found, we directly retrieve a generative process to sample from our approximation to the posterior distribution. Specifically

$$P(y|d) \approx \mathcal{N}(y; 0, \mathbf{1}) \quad (44)$$

$$\zeta = g^{-1}(y; \bar{\zeta}), \quad (45)$$

where g^{-1} is only implicitly defined using Equation (29) and therefore its inverse application has to be approximated numerically in general.

3.1.1. Numerical Approximation to Sampling

Recall that

$$y = g(\zeta; \bar{\zeta}) = \sqrt{\bar{\mathcal{M}}}^{-1} \tilde{g}(\zeta; \bar{\zeta}). \quad (46)$$

To generate a posterior sample for ζ we have to draw a random realization for y from a unit Gaussian, and then solve Equation (46) for ζ . To avoid the matrix square root of $\bar{\mathcal{M}}$, we may instead define

$$z \equiv \sqrt{\bar{\mathcal{M}}} y = \tilde{g}(\zeta; \bar{\zeta}) \quad \text{with} \quad P(z) = \mathcal{N}(z; 0, \bar{\mathcal{M}}). \quad (47)$$

Sampling from $P(z)$ is much more convenient than constructing the matrix square root, since

$$\bar{\mathcal{M}} = \mathbf{1} + \left(\left(\frac{\partial x}{\partial \bar{\zeta}} \right)^T \frac{\partial x}{\partial \bar{\zeta}} \right) \Bigg|_{\bar{\zeta}}, \quad (48)$$

and therefore a random realization may be generated using

$$z = \eta_1 + \left(\frac{\partial x}{\partial \bar{\zeta}} \Bigg|_{\bar{\zeta}} \right)^T \eta_2 \quad \text{with} \quad \eta_i \sim \mathcal{N}(\eta_i; 0, \mathbf{1}), \quad i \in \{1, 2\}. \quad (49)$$

Finally, a posterior sample ζ is retrieved by inversion of Equation (47). We numerically approximate the inversion by minimizing the squared difference between z and $\tilde{g}(\zeta)$. Specifically,

$$\bar{\zeta} = \underset{\zeta}{\operatorname{argmin}} \left(\frac{1}{2} (z - \tilde{g}(\zeta))^T (z - \tilde{g}(\zeta)) \right). \quad (50)$$

Note that if g is invertible then also \tilde{g} is invertible as $\bar{\mathcal{M}}$ is a symmetric positive definite matrix. Therefore the quadratic form of Equation (50) has a unique global optimum at zero which corresponds to the inverse of Equation (47).

In practice, this optimum is typically only reached approximately. For an efficient numerical approximation, throughout this work, we employ a second order quasi-Newton method, named NewtonCG [23], as implemented in the NIFTy framework. Within the NewtonCG algorithm, we utilize the metric $\tilde{\mathcal{M}}(\xi)$ as a positive-definite approximation to the curvature of the quadratic form in Equation (50). Furthermore, its inverse application, required for the second order optimization step of NewtonCG, is approximated with the conjugate gradient (CG) [24] method, which requires the metric to be only implicitly available via matrix-vector products. In addition, in practice we find that the initial position ξ^0 of the minimization procedure can be set to be equal to the prior realization η_1 used to construct z (Equation (49)) in order to improve convergence as $\xi = \eta_1$ is the solution of Equation (47) for all degrees of freedom unconstrained by the likelihood. Alternatively, for weakly non-linear problems, initializing ξ^0 as the solution of the linearized problem

$$z = \tilde{g}(\xi; \bar{\xi})|_{\xi=\bar{\xi}} + \left. \frac{\partial \tilde{g}}{\partial \xi} \right|_{\xi=\bar{\xi}} (\xi - \bar{\xi}) = \left(\mathbb{1} + \left(\frac{\partial x}{\partial \xi} \right)^T \frac{\partial x}{\partial \xi} \right) \Big|_{\xi=\bar{\xi}} (\xi - \bar{\xi}), \quad (51)$$

can significantly improve the convergence. The full realization of the sampling procedure is summarized in Algorithm 1.

Algorithm 1: Approximate posterior samples using inverse transformation

```

1 Function drawSample(Location  $\bar{\xi}$ , Transformation  $x(\xi)$ , Jacobian  $\frac{\partial x}{\partial \xi}$ ):
2    $A \leftarrow \left. \frac{\partial x}{\partial \xi} \right|_{\xi=\bar{\xi}}$ 
3    $\eta_1 \sim \mathcal{N}(\eta_1; 0, \mathbb{1})$ 
4    $\eta_2 \sim \mathcal{N}(\eta_2; 0, \mathbb{1})$ 
5    $z \leftarrow \eta_1 + A^T \eta_2$ 
6    $\xi^0 \leftarrow \eta_1$  or  $\xi^0 \leftarrow \text{Solve}(z = (\mathbb{1} + A^T A)(\xi^0 - \bar{\xi}))$  for  $\xi^0$  (see Equation (51))
7   Function Energy( $\xi$ ):
8      $\tilde{g} \leftarrow \xi - \bar{\xi} + A^T (x(\xi) - x(\bar{\xi}))$ 
9     return  $\frac{1}{2} (z - \tilde{g})^T (z - \tilde{g})$ 
10   $\xi^* \leftarrow \text{NewtonCG}(\text{Energy}, \xi^0)$ 
11  return  $\xi^*$ 

```

3.1.2. Properties

We may qualitatively study some basic properties of the coordinate transformation and the associated approximation using illustrative one and two dimensional examples. To this end, consider a one dimensional log-normal prior model with zero mean and standard deviation σ_p of the form

$$s(\xi) = e^{\sigma_p \xi} \quad \text{with} \quad P(\xi) = \mathcal{N}(\xi; 0, 1), \quad (52)$$

from which we obtain a measurement d subject to independent, additive Gaussian noise with standard deviation σ_n such that the likelihood takes the form

$$P(d|\xi) = \mathcal{N}(d; s(\xi), \sigma_n^2). \quad (53)$$

The posterior distribution is given as

$$P(\xi|d) \propto P(d|\xi) P(\xi) = \mathcal{N}(d; s(\xi), \sigma_n^2) \mathcal{N}(\xi; 0, 1), \quad (54)$$

and its metric takes the form

$$\mathcal{M}(\xi) = \left(\frac{1}{\sigma_n} \frac{\partial s(\xi)}{\partial \xi} \right)^2 + 1 = \left(\frac{\sigma_p}{\sigma_n} \right)^2 e^{2\sigma_p \xi} + 1. \quad (55)$$

In this one dimensional example we can construct the exact transformation g_{iso} that maps from ζ to the transformed coordinates y , by integrating the square root of Equation (55) over ζ . The resulting transformation can be seen in the central panel of Figure 2, for an example with $\sigma_p = 3$ and $\sigma_n = 0.3$ and measured data $d = 0.5$. In addition, we depict the approximated transformation $g(\zeta; \bar{\zeta})$ for multiple expansion points $\bar{\zeta} \in \{-1, -0.6, -0.2\}$. We see that the function approximation quality depends on the choice of the expansion point $\bar{\zeta}$ as the approximation error is smallest in the vicinity of $\bar{\zeta}$. In order to transform the posterior distribution P (Equation (54)) into the new coordinated system, not all parts of the transformation are equally relevant and therefore different expansion points result in more/less complex transformed distributions (see top panel of Figure 2). Finally, if we use a standard distribution in the transformed coordinates y and transform it back using the inverse transformations $g^{-1}(y; \bar{\zeta})$, we find that the approximation quality of the resulting distributions $Q_{\bar{\zeta}}$ depends on $\bar{\zeta}$. The distributions are illustrated in the left panel of Figure 2 together with the Kullback–Leibler divergence KL between the true posterior distribution P and the approximations $Q_{\bar{\zeta}}$. We also illustrate the “geometrically optimal” approximation using a standard distribution in y and the optimal transformation g_{iso} and find that while the approximation error becomes minimal in this case, it remains non-zero. Considering the discussion in Section 2.2, this result is to be expected due to the error contribution from the change in volume associated with the transformation g . As a comparison we also depict the optimal linear approximation of P , that is a normal distribution in the coordinates ζ with optimally chosen mean and standard deviation. We see that even the worst expansion point $\bar{\zeta} = -0.2$, that is far away from the optimum, still yields a better approximation of the posterior.

As a second example we consider the task of inferring the mean m and variance v of a single, real valued Gaussian random variable d . In terms of $s = (m, v)$, the likelihood takes the form

$$P(d|s) = \mathcal{N}(d; m, v). \quad (56)$$

Furthermore we assume a prior model for s by means of a generative model of the form

$$m = \zeta_1 \quad \text{and} \quad vs. = \exp[3(\zeta_2 + 2\zeta_1)], \quad (57)$$

where ζ_1 and ζ_2 follow standard distributions a priori. This artificial model results in a linear prior correlation between the mean and the log-variance and thus introduces a non-linear coupling between m and v . The resulting two dimensional posterior distribution $P(\zeta_1, \zeta_2)$ can be seen in the left panel of Figure 3, together with the two marginals $P(\zeta_1)$ and $P(\zeta_2)$ for a given measurement $d = 0$. We approximate this posterior distribution following the direct approach described in Section 3.1, where the expansion point $\bar{\zeta}$ is obtained from minimizing the sum of the posterior Hamiltonian and the log-determinant of the metric (see Equation (43)). The resulting approximative distribution Q_D is shown in the right panel of Figure 3, where the location of $\bar{\zeta}$ is indicated as a blue cross. In comparison to the true distribution, we see that both, the joint distribution as well as the marginals are in a good agreement qualitatively, which is also supported quantitatively by a small difference of the KL between P and Q_D (see Figure 3). The difference between P and Q_D appears to increase in regions further away from the expansion point, which is to be expected due to the local nature of the approximation. However, non-linear features such as the sharp peak at the “bottom” of P (Figure 3), are also present in Q_D , although slightly less prominent. This demonstrates that relevant non-linear structure can, to some degree, be captured by the coordinate transformation g derived from the metric \mathcal{M} of the posterior.

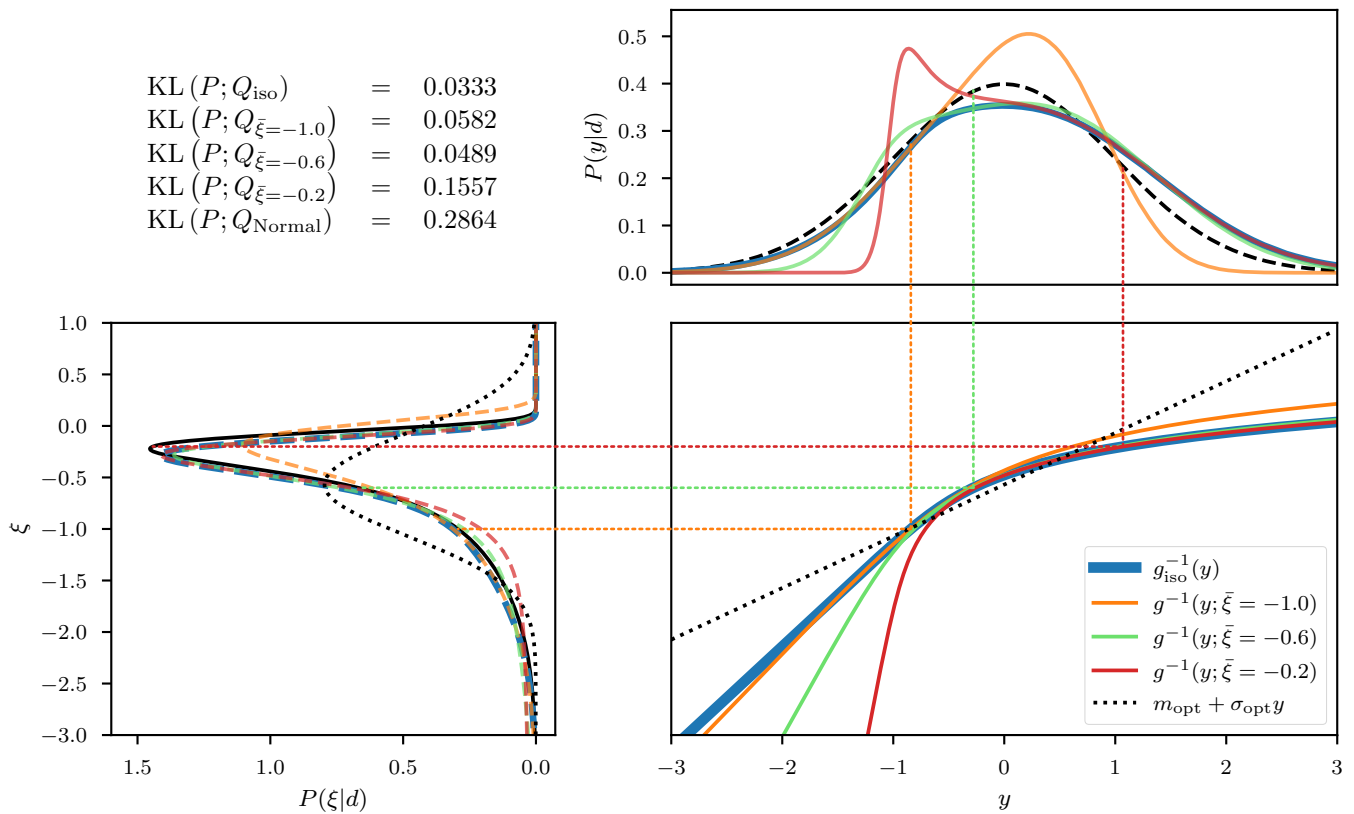


Figure 2. Illustration of the coordinate transformation for the one-dimensional log-normal model (Equation (54)). The true posterior $P(\xi|d)$, displayed as the black solid line in the left panel, is transformed into the coordinate system y using the optimal transformation g_{iso} (blue), as well as three approximations g thereof with expansion points $\bar{\xi} \in \{-1, -0.6, -0.2\}$ (orange, green, red). The resulting distributions $P(y|d)$ are displayed in the top panel of the figure as solid lines, color coded according to the used transformation g (or g_{iso} in case of blue). The black, dashed line in the top panel displays a standard distribution in y . The location of the expansion point $\bar{\xi}$, and its associated point in y , is highlighted via the color coded, dotted lines. Finally, the direct approximations to the posterior associated with the transformations, meaning the push-forwards of the standard distribution in y using the inverse of the various transformations g^{-1} , are displayed in the left panel as dashed lines, color coded according to their used transformation. As a comparison, the “optimal linear approximation” (black dotted line in the central panel), which corresponds to the optimal approximation of the posterior with a normal distribution in ξ (black dotted line in left panel), is displayed as a comparison. To numerically quantify the information distance between the true distribution P and its approximations Q_{\bullet} , the Kullback–Leibler (KL) divergences between P and Q_{\bullet} are displayed in the top left of the image. The numerical values of the KL are given in nats (meaning the KL is evaluated in the basis of the natural logarithm).

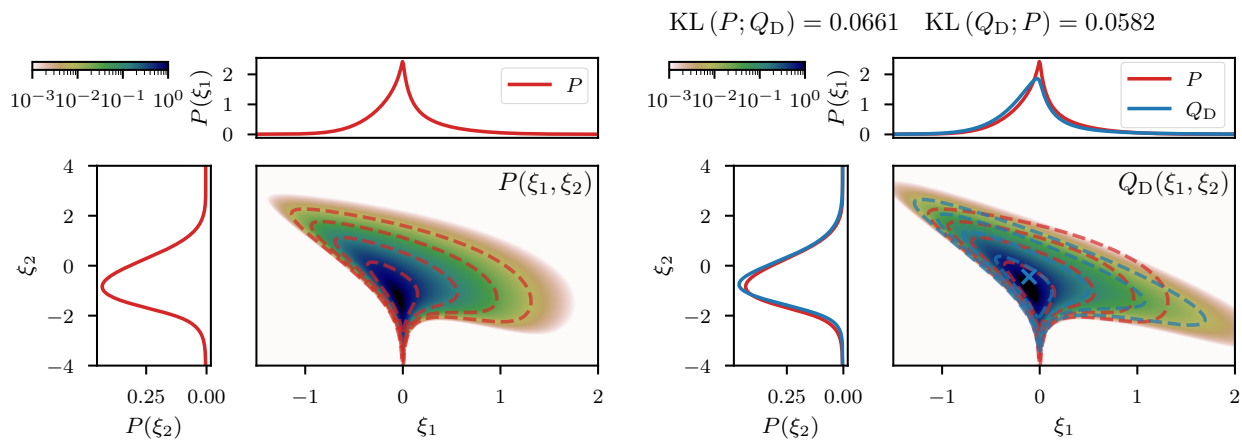


Figure 3. Left: posterior distribution P in the standard coordinates $\xi_{1/2}$ for the inference of the mean and variance of a normal distribution (Equations (56) and (57)). The central panel shows the two dimensional density and the red dashed lines are logarithmically spaced contours. The top and left sub-panels display the marginal posterior distributions for ξ_1 and ξ_2 , respectively. Right: Approximation Q_D to the posterior distribution using the direct method (Section 3.1). As a comparison, the contours (red dashed) and the marginal distributions (red solid) of the true posterior distribution P are displayed in addition to the approximation. The blue cross in the central panel denotes the location of the expansion point used to construct Q_D . Above the panel we display the optimal ($KL(P; Q_D)$) and variational ($KL(Q_D; P)$) Kullback–Leibler divergences between P and Q_D .

Although these low-dimensional, illustrative examples appear promising, there remains one central issue left to be addressed before the approach can be applied to high-dimensional problems. In particular, the direct approach possesses a substantial additional computational burden compared to, e.g., a maximum a posteriori (MAP) estimate in ξ which is obtained by minimizing the posterior Hamiltonian \mathcal{H} . For the direct approach, the optimization objective Equation (43) consists not only of \mathcal{H} , but also of the log-determinant of the metric \mathcal{M} . In all but the simplest examples this term cannot be computed directly but has to be approximated numerically as in high dimensions an explicit representation of the matrix becomes infeasible and \mathcal{M} is only implicitly accessible through matrix vector products (MVPs). There are a variety of stochastic log-determinant (more specifically trace-log) estimators based on combining Hutchinsons’ trace-estimation [25] with approximations to the matrix logarithm using, e.g., Chebychev polynomials [26], Krylov subspace methods [27], or moment constrained estimation based on Maximum Entropy [28]. While all these methods provide a significant improvement in performance compared to directly computing the determinant, they nevertheless typically require many MVPs in order to yield an accurate estimate. For large and complex problems, evaluating an MVP of \mathcal{M} is dominated by applying the Jacobian of x , more precisely of the generative process $s'(\xi)$, and its adjoint to a vector. Similarly, evaluating the gradient of \mathcal{H} is also dominated by an MVP that invokes applying the adjoint Jacobian of $s'(\xi)$. Therefore the computational overhead compared to a MAP estimate in ξ is, roughly, multiplicative in the number of MVPs. For large, non-linear problems, this quickly becomes infeasible as nonlinear optimization typically requires many steps to reach a sensible approximation to the optimum.

Nevertheless there remain some important exceptions, in which a fast and scalable algorithm emerges. In particular recall that

$$\log(|\mathcal{M}|) = \text{tr} \left(\log \left(\mathbb{1} + \left(\frac{\partial x}{\partial \xi} \right)^T \frac{\partial x}{\partial \xi} \right) \right) = \text{tr} \left(\log \left(\mathbb{1} + \frac{\partial x}{\partial \xi} \left(\frac{\partial x}{\partial \xi} \right)^T \right) \right), \quad (58)$$

where the last equality arises from applying the matrix determinant lemma. Therefore in cases where the dimensionality of the so-called data-space (i.e., the target space of x) is much smaller than the dimensionality of the signal space (the domain of ξ), the latter

representation of the metric is of much smaller dimension. Thus, in cases where either the signal- or the data-space is small, or in weakly non-linear cases (i.e., if \mathcal{M} is close to the identity), the log-determinant may be approximated efficiently enough to give rise to a fast and scalable algorithm. For the (arguably most interesting) class of problems where neither of these assumptions is valid, however, the direct approach to obtain the optimal expansion point becomes too expensive for practical purposes as none of the log-determinant estimators scale linearly with the size of the problem in general.

3.2. Geometric Variational Inference (geoVI)

As we shall see, it is possible to circumvent the need to compute the log-determinant of the metric at any point, if we employ a specific variant of a variational approximation to obtain the optimal expansion point. To this end, we start with a variational approximation to the posterior P , assuming that the approximative distribution \tilde{Q} is given as the unit Gaussian in y transformed via g . To this end let

$$\tilde{Q}(\xi|\bar{\xi}) = \mathcal{N}(g(\xi; \bar{\xi}); 0, \mathbb{1}) \left\| \frac{\partial g(\xi; \bar{\xi})}{\partial \xi} \right\|, \tag{59}$$

denote the approximation to the posterior conditional to the expansion point $\bar{\xi}$. The variationally optimal $\bar{\xi}$ can be found by optimization of the forward Kullback–Leibler divergence between \tilde{Q} and P , as given via

$$\begin{aligned} \text{KL}(\tilde{Q}|P) &\equiv \int \log\left(\frac{\tilde{Q}(\xi|\bar{\xi})}{P(\xi|d)}\right) \tilde{Q}(\xi|\bar{\xi}) d\xi \\ &= \langle \mathcal{H}(\xi|d) \rangle_{\tilde{Q}(\xi|\bar{\xi})} - \langle \mathcal{H}_{\tilde{Q}}(\xi|\bar{\xi}) \rangle_{\tilde{Q}(\xi|\bar{\xi})} \\ &= \langle \mathcal{H}(\xi|d) \rangle_{\tilde{Q}(\xi|\bar{\xi})} + \frac{1}{2} \langle \log(|\tilde{\mathcal{M}}(\xi)|) \rangle_{\tilde{Q}(\xi|\bar{\xi})} + \text{KL}_0, \end{aligned} \tag{60}$$

where KL_0 denotes contributions independent of $\bar{\xi}$, and $\mathcal{H}(\xi|d)$ and $\mathcal{H}_{\tilde{Q}}$ denote the Hamiltonians of the posterior and the approximation, respectively. We notice that in this form, a minimization of the KL w.r.t. $\bar{\xi}$ does not circumvent a computation of the log-determinant of the metric. Within the KL, this term arises from the entropy of the approximation \tilde{Q} , and can be understood as a measure of the volume associated with the distribution. In order to avoid this term, our idea is to propose an alternative family of distributions $Q_m(\xi|\bar{\xi})$, defined as a shifted version of \tilde{Q} . Specifically we let $\xi \rightarrow m + \xi - \bar{\xi}$ such that the distribution may be written as

$$Q_m(\xi|\bar{\xi}) = \tilde{Q}(\xi|\bar{\xi})|_{\xi=\xi+\bar{\xi}-m} \equiv Q(r|\bar{\xi})|_{r=\xi-\bar{\xi}} \quad \text{with } r = \xi - \bar{\xi}, \tag{61}$$

where we also introduced the residual r , which measures the deviations from $\bar{\xi}$, and the associated distribution $Q(r|\bar{\xi})$. In words, $Q_m(\xi|\bar{\xi})$ is the distribution using the residual statistics r , around an expansion point $\bar{\xi}$, but shifted to m . One can easily verify that the entropy related to Q_m becomes independent of m , as shifts are volume-preserving transformations. Therefore we may use some fixed expansion point $\bar{\xi}$, and find the optimal shift m using the KL which now may be written as

$$\begin{aligned} \text{KL}(Q_m|P) &= \langle \mathcal{H}(\xi = m + r, d) \rangle_{Q(r|\bar{\xi})} + \frac{1}{2} \langle \log(|\tilde{\mathcal{M}}(\xi)|)|_{\xi=\bar{\xi}+r} \rangle_{Q(r|\bar{\xi})} + \text{KL}_0 \\ \widehat{\text{KL}} &= \langle \mathcal{H}(\xi = m + r, d) \rangle_{Q(r|\bar{\xi})}, \end{aligned} \tag{62}$$

where $\widehat{\text{KL}}$ denotes the KL up to m independent contributions. After optimization for m , we can update to a new expansion point, and use it to define a new family of distributions Q_m which are a more appropriate class of approximations. In general, the expectation value in

$\widehat{\text{KL}}$ cannot be computed analytically, but it can be approximated using a set of N samples $\{r_i^*\}_{i \in \{1, \dots, N\}}$, drawn from $Q(r|\bar{\xi})$, which yields

$$\widehat{\text{KL}}(Q_m|P) \approx \frac{1}{N} \sum_{i=1}^N \mathcal{H}(\xi = m + r_i^*, d) \quad \text{with } r_i^* \sim Q(r|\bar{\xi}). \quad (63)$$

Sampling from $Q(r|\bar{\xi})$ is defined as in Section 3.1.1, where the sampling procedure for $\bar{Q}(\xi|\bar{\xi})$ is described, with the addition that a sample r^* is obtained from a sample for ξ^* as $r^* = \xi^* - \bar{\xi}$.

Optimizing $\widehat{\text{KL}}$ w.r.t. m yields the variational optimum for the distribution $Q_m(\xi|\bar{\xi})$, given a fixed, predetermined expansion point $\bar{\xi}$. In order to move the expansion point $\bar{\xi}$ towards the optimal point, its location is updated subsequently and the KL is re-estimated using novel samples from $Q(r|\bar{\xi})$ with an updated $\bar{\xi}$. Specifically, we initialize the optimization algorithm at some position m^0 , set $\bar{\xi} = m^0$ to obtain a set of samples $\{r_i^*\}_{i \in \{1, \dots, N\}}^{(0)}$, and use this set to approximate the KL. This approximation is then used to obtain an optimal shift m^1 . Given this optimal shift, a new expansion point $\bar{\xi} = m^1$ is defined and used to obtain a novel set of samples $\{r_i^*\}_{i \in \{1, \dots, N\}}^{(1)}$ which defines a new estimate for the KL. This estimate is furthermore used to obtain a novel optimal m , and so on. An illustrative view of this procedure is given in Figure 4. Finally, the entire procedure of optimizing the KL for m and re-estimation of the KL via a novel expansion point is repeated until the algorithm converges to an optimal point $m^* = \bar{\xi}^*$. To optimize $\widehat{\text{KL}}$ for m , we again employ the NewtonCG algorithm, and use the average of the metric \mathcal{M} as a proxy for the curvature of $\widehat{\text{KL}}$ to perform the optimization step. Specifically we use

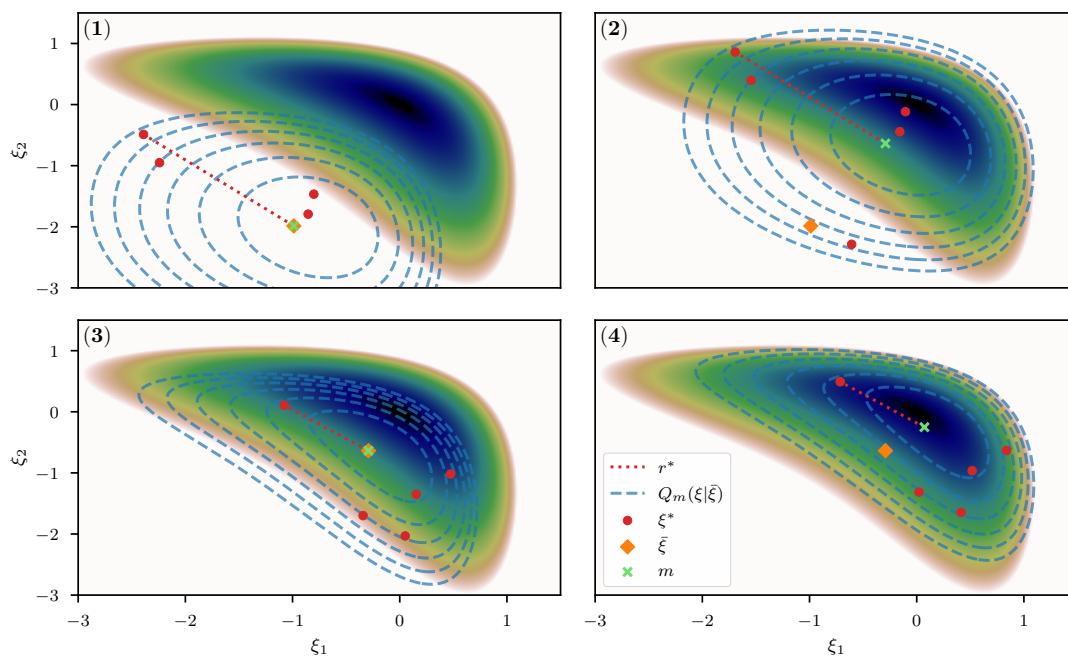


Figure 4. (1)–(4): Visualization of the geoVI steps. (1): A randomly initialized shift m (green cross) is used to set the initial expansion point $\bar{\xi}$ (orange dot) which in turn defines the initial approximation $Q_m(\xi|\bar{\xi})$ (blue dashed contours) used to generate a set of samples ξ^* (red dots). (2): The KL (Equation (63)), estimated from the samples, is used to optimize for m , which results in a shift of $Q_m(\xi|\bar{\xi})$ away from the expansion point $\bar{\xi}$. The residual statistics r^* derived from the geometry around $\bar{\xi}$, however, remains unchanged during this shift and therefore, at the new location m , becomes a bad representation of the local geometry. Thus, in (3), the expansion point is set to the current estimate of m , which yields an update to the approximation $Q_m(\xi|\bar{\xi})$. Finally, we generate samples from this update and use them to optimize the re-estimated KL for m which again results in a shift as seen in (4). Within the full geoVI algorithm this procedure is iterated until convergence.

$$\widehat{\mathcal{M}}(m) = \frac{1}{N} \sum_{i=1}^N \mathcal{M}(\zeta = m + r_i^*), \quad (64)$$

as the metric of $\widehat{\text{KL}}$. We call this algorithm the *geometric Variational Inference* (geoVI) method. A pseudo-code summary of geoVI is given in Algorithm 2.

Algorithm 2: Geometric Variational Inference (geoVI)

Input: Likelihood $\mathcal{H}(d|\zeta)$, Transformation $x(\zeta)$, Jacobian $\frac{\partial x}{\partial \zeta}$

```

1 Function Energy( $\zeta$ ):
2   | return  $\mathcal{H}(d|\zeta) + \frac{1}{2}\zeta^T\zeta$ 
3  $m \sim \mathcal{N}(m, 0, \mathbf{1})$ 
4 while  $m$  not converged do
5   |  $\bar{\zeta} \leftarrow m$ 
6   | samples  $\leftarrow$  empty list
7   | for  $i = 1$  to  $N$  do
8     |  $\zeta^* \leftarrow \text{drawSample}(\bar{\zeta}, x, \frac{\partial x}{\partial \zeta})$  (see Algorithm 1)
9     |  $r^* \leftarrow \zeta^* - \bar{\zeta}$ 
10    | Insert  $r^*$  into samples
11  | end
12  | Function geoKL( $\zeta$ ):
13    |  $kl \leftarrow 0$ 
14    | for  $r^*$  in samples do
15      |  $kl \leftarrow kl + \text{Energy}(\zeta + r^*)$ 
16    | end
17    | return  $\frac{1}{N} kl$ 
18  |  $m^* \leftarrow \text{NewtonCG}(\text{geoKL}, m)$ 
19  |  $m \leftarrow m^*$ 
20 end
21 posteriorSamples  $\leftarrow$  empty list
22 for  $r^*$  in samples do
23   |  $\zeta^* \leftarrow m + r^*$ 
24   | Insert  $\zeta^*$  into posteriorSamples
25 end
Output: posteriorSamples

```

3.2.1. Numerical Sampling within geoVI

It is noteworthy that, as described in Section 3.1.1, an implementation of the proposed sampling procedure for the residual r , and as a result also of the geoVI method itself, inevitably relies on numerical approximations to realize a sample for r . To better understand the impact of such approximations, we have to consider its impact on the distribution $Q(r|\bar{\zeta})$. To this end, we denote with f the function that, given the expansion point $\bar{\zeta}$, turns two standard distributed random vectors η_1 and η_2 into a random realization of r . Specifically

$$r = f(\eta_1, \eta_2; \bar{\zeta}) \quad \text{with} \quad \eta_{1/2} \sim \mathcal{N}(\eta_{1/2}; 0, \mathbf{1}), \quad (65)$$

where the functional form of f is defined by combination of Equations (49) and (50). using f we may write the geoVI distribution Q as

$$Q(r|\bar{\zeta}) = \int \int \delta(r - f(\eta_1, \eta_2; \bar{\zeta})) \mathcal{N}(\eta_1; 0, \mathbf{1}) \mathcal{N}(\eta_2; 0, \mathbf{1}) d\eta_1 d\eta_2. \quad (66)$$

Any numerical algorithm used to approximate the sampling, irrespective of its exact form, may be described by replacing the function f , leading to exact sampling from Q , with some approximation \hat{f} which leads to an approximation of the distribution for r , which we

denote as $\widehat{Q}(r|\bar{\xi})$. Therefore, in a way, the geoVI result using a numerical approximation for sampling can be understood as the variational optimum chosen from the family of distributions \widehat{Q} , rather than Q . Therefore, even for a non-zero approximation error in \widehat{f} , the result remains a valid optimum of a variational approximation, it is simply the family of distributions used for approximation that has changed. This finding is of great relevance in practice, as there is typically a trade off between numerical accuracy of the generated samples and computational efforts. Thus, we may achieve faster convergence at a cost of accuracy in the approximation, but without completely detaching from the theoretical optimum, so long as \widehat{f} remains sufficiently close to f . Nevertheless, as motivated in the introduction, it is important for the chosen family to contain distributions close to the true posterior, and therefore it remains important that the family \widehat{Q} remains close to the family of Q as only for Q the geometric correspondence to the posterior has been established. A detailed study to further quantify this result, is left to future work.

3.2.2. MGVI as a First Order Approximation

We can compare the geoVI algorithm to the aforementioned variational approximation technique called Metric Gaussian variational inference (MGVI), and notice some key similarities. In particular the optimization heuristics with repeated alternation between sampling of r^* and optimization for m is entirely equivalent. The difference occurs in the distribution $Q(r|\bar{\xi})$ used for approximation. In MGVI, Q is assumed to be a Gaussian distribution in r , as opposed to the Gaussian distribution in the transformed space y used in geoVI. Specifically

$$Q_{\text{MGVI}}(r|\bar{\xi}) \equiv \mathcal{N}(r; 0, \bar{\mathcal{M}}^{-1}), \quad (67)$$

where the inverse of the posterior metric \mathcal{M} , evaluated at the expansion point $\bar{\xi}$, is used as the covariance. As it turns out, the distribution Q_{MGVI} arises naturally as a first order approximation to the coordinate transformation used in the geoVI approach. Specifically if we consider the geoVI distribution of r given in terms of a generative process

$$r = g^{-1}(y; \bar{\xi}) - \bar{\xi} \quad \text{with} \quad y \sim \mathcal{N}(y; 0, \mathbf{1}), \quad (68)$$

and expand it around $y = 0$ to first order, we get that

$$\begin{aligned} r &= g^{-1}(0, \bar{\xi}) - \left(\left. \frac{\partial g(\bar{\xi}, \bar{\xi})}{\partial \bar{\xi}} \right|_{\bar{\xi}=\bar{\xi}} \right)^{-1} y + \mathcal{O}(y^2) - \bar{\xi} \\ &= \bar{\xi} - \bar{\xi} - \sqrt{\bar{\mathcal{M}}} \left(\mathbf{1} + \left(\left. \left(\frac{\partial x}{\partial \bar{\xi}} \right)^T \frac{\partial x}{\partial \bar{\xi}} \right) \right|_{\bar{\xi}=\bar{\xi}} \right)^{-1} y + \mathcal{O}(y^2) \\ &= -(\sqrt{\bar{\mathcal{M}}})^{-1} y + \mathcal{O}(y^2). \end{aligned} \quad (69)$$

Therefore, to first order in y , we get that

$$Q(r|\bar{\xi}) = \int \delta\left(r + (\sqrt{\bar{\mathcal{M}}})^{-1} y\right) \mathcal{N}(y; 0, \mathbf{1}) \, dy = \mathcal{N}(r; 0, \bar{\mathcal{M}}^{-1}) = Q_{\text{MGVI}}(r|\bar{\xi}). \quad (70)$$

This correspondence shows that geoVI is a generalization of MGVI in non-linear cases. This is a welcome result, as numerous practical applications [29–31] have shown that already MGVI provides a sensible approximation to the posterior distribution. On the other hand, it provides further insight in which cases the MGVI approximation remains valid, and when it reaches its limitations. In particular if

$$\mathcal{M}(m+r) \approx \bar{\mathcal{M}}, \quad \forall r = g^{-1}(y, \bar{\xi}) - \bar{\xi} \quad \text{with} \quad y \sim \mathcal{N}(y; 0, \mathbf{1}), \quad (71)$$

we get that the first order approximation of Equation (69) yields a close approximation of the inverse and geoVI reduces to the MGVI algorithm. In contrast, geoVI with its non-linear inversion requires the log determinant of the metric \mathcal{M} to be approximately constant throughout the sampling regime. This is a much less restrictive requirement than Equation (71), as the variation of eigenvalues of \mathcal{M} is considered on a logarithmic scale whereas it is considered on linear scale in Equation (71). Furthermore, the log-determinant is invariant under unitary transformations which means that local rotations of the metric, and therefore changes in orientation as we move along the manifold, can be captured by the non-linear approach, whereas Equation (71) does not hold any more if the orientation varies as a function of r . Therefore we expect the proposed approach to be applicable in a more general context, while still retaining the MGVI properties, as it reproduces MGVI in the linear limit.

3.3. Examples

We can visually compare the geoVI and the MGVI algorithm using the two-dimensional example previously mentioned in Section 3.1.2. In analogy to Figure 3 we depict the approximation to the posterior density together with its two marginals in Figure 5. We see that geoVI yields a similar result compared to the direct approach here, while it provides a significant improvement compared to the approximation capacity of MGVI.

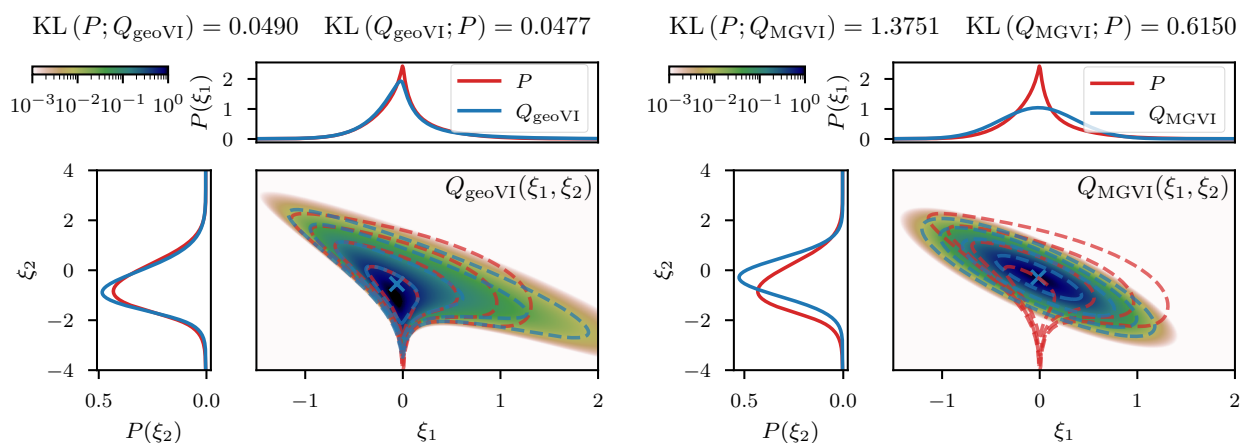


Figure 5. The geoVI and MGVI approximations of the two-dimensional example described in Section 3.1.2. We display the same quantities as for the direct approximation shown in Figure 3.

To conclude the illustrative examples, we consider a single observation of the product of a normal and a log-normal distributed quantity subject to independent, additive Gaussian noise. The full model consists of a likelihood and a prior of the form

$$P(d|\xi_1, \xi_2) = \mathcal{N}(d; \xi_1 e^{\xi_2}, \sigma_n^2) \quad \text{with} \quad \xi_{1/2} \sim \mathcal{N}(\xi_{1/2}; 0, 1). \tag{72}$$

This example should serve as an illustration of the challenges that arise when attempting a separation of non-linearly coupled quantities from a single observation. Such separation problems reappear in Section 4 in much more intertwined and high dimensional examples, but much of the structural challenges can already be seen in this simple two-dimensional problem. Figure 6 displays the results of the direct approach as well as the geoVI and MGVI methods for a measurement setting of $d = -0.3$ and $\sigma_n = 0.1$. As a comparison, we also depict the results from performing a variational approximation using a normal distribution with a diagonal covariance, also known as a mean-field approximation (MFVI), as well as an approximation with a normal distribution using a full-rank matrix as its covariance (FCVI). Both, the diagonal as well as the full-rank covariance are considered parameters of the distribution, and have to be optimized for in addition to the mean of the normal distribution. An efficient implementation thereof is described in [7]. We notice that both,

the direct and the geoVI approach manage to approximate the true posterior distribution well, although the KL values indicate that the approximation by geoVI is worse by ≈ 0.016 nats compared to the direct approach. Here the passive update of the expansion point used in this approach reaches its limitations as in cases where the posterior distribution becomes increasingly narrow towards the optimal expansion point, the static sample statistics of r can get stuck during optimization and increasingly repeated re-sampling becomes necessary as one moves closer to the optimum. Nevertheless, the geoVI approximation remains a good approximation to the true distribution, especially when compared to the approaches using a normal distribution such as MGVI, MFVI, and FCVI.

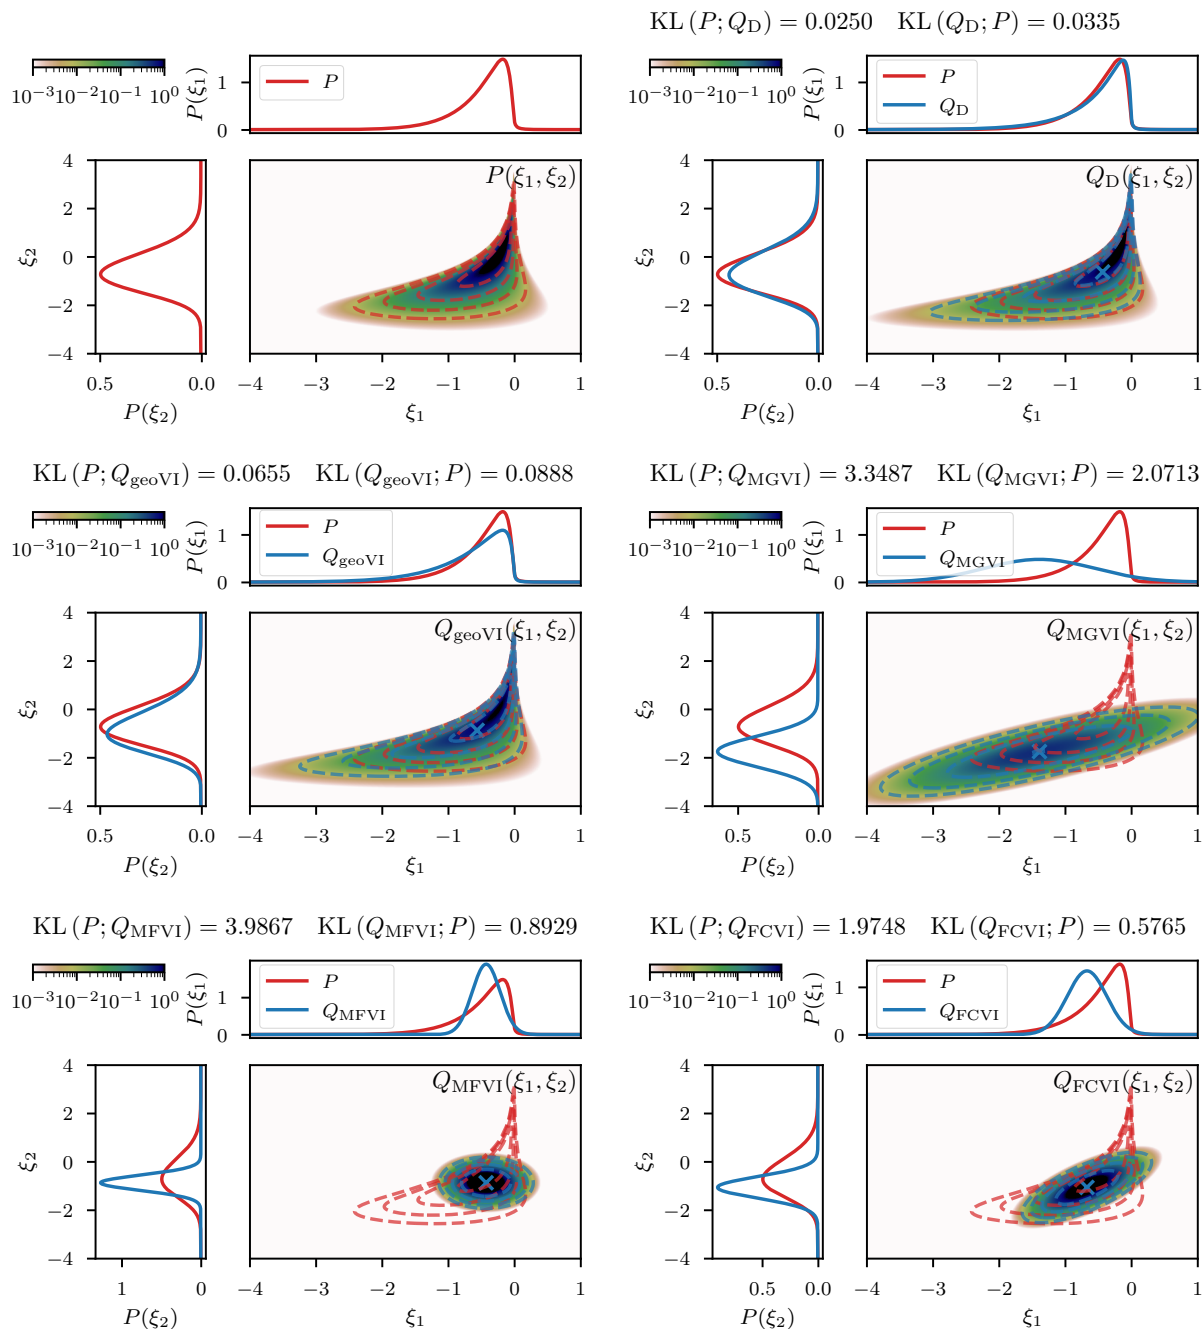


Figure 6. Same setup as in Figures 3 and 5 but for a Gaussian measurement of the product of a normal distributed quantity ξ_1 and a log-normal distributed one ξ_2 as described in the second example of Section 3.3. From top to bottom and from left to right: ground truth P , direct approximation Q_D , geoVI approximation Q_{geoVI} , MGVI approximation Q_{MGVI} , mean-field approximation Q_{MFVI} , and the normal approximation with a full-rank covariance Q_{FCVI} .

4. Applications

To investigate the performance of the geoVI algorithm in high dimensional imaging problems, we apply it to two mock data examples and compare it to the results using MGVI. In the first example, which serves as an illustration, the geoVI results are additionally compared to the results obtained from applying a Hamiltonian Monte-Carlo (HMC) sampler [12] to the mock example (see Section 5.2 for further information on HMC). The second example is an illustration of a typical problem encountered in astrophysical imaging. Both examples consist of hierarchical Bayesian models with multiple layers which are represented as a generative process. One particularly important process for the class of problems at hand are statistically homogeneous and isotropic Gaussian processes with unknown power spectral density, for which a flexible generative model has been presented in [32]. This process is at the core of a variety of astrophysical imaging applications [32–35], and therefore an accurate posterior approximation of problems involving this model is crucial. To better understand the inference challenges that arise in problems using this particular model, we briefly summarize some of its key properties.

4.1. Gaussian Processes with Unknown Power Spectra

Consider a zero mean, square integrable random process s_x defined on a L -dimensional space subject to periodic boundary conditions which, for simplicity, we assume to have size one. Specifically let $x \in \Lambda = [0, 1]^L$ and thus $s \in \mathcal{L}^2(\Lambda)$. A Gaussian process

$$P(s) = \mathcal{N}(s; 0, S), \quad (73)$$

with mean zero and covariance function S_{xy} is said to be statistically homogeneous and isotropic, if S is a function of the Euclidean distance between two points i.e.,

$$S_{xy} = S(|x - y|). \quad (74)$$

Furthermore, as implied by the Wiener–Khinchin theorem [36], the linear operator associated with S becomes diagonal in the Fourier space associated with Λ , and therefore s may be represented in terms of a Fourier series with coefficients \tilde{s}_k , where k labels the Fourier coefficients. These coefficients are independent, zero mean Gaussian random variables with variance

$$\langle |\tilde{s}_k|^2 \rangle_{P(s)} \equiv P_s(|k|), \quad (75)$$

which is also known as the power spectrum P_s of s . As P_s encodes the correlation structure of s , its functional form is crucial to determine the prior statistical properties of s . In [32], a flexible, non-parametric prior process for the power spectrum has been proposed by means of a Gauss–Markov process on log-log-scale. This process models the spectrum as a straight line on log-log-scale (resulting in a power law in $|k|$ on linear scale) with possible continuous deviations thereof. These deviations are itself defined as a Gauss–Markov process (specifically an integrated Wiener process) and their respective variance is, among others, an additional scalar parameter steering the properties of this prior process that are also considered to be random variables that have to be inferred. These parameters are summarized in Table 1. A more formal derivation of this model in terms of a generative process relating standard distributed random variables ζ_p to a random realization $P_s(\zeta_p)$ of this prior model, is given in Appendix B.

Table 1. Table of additional parameters.

Name	Description	Prior Distribution
offset std.	Prior standard deviation of the overall offset of s from zero	Log-normal
fluctuations	Prior amplitude of the variation of s around its offset	Log-normal
slope	Exponent of the power law related to P_s	Normal
flexibility	Amplitude of deviations from the power-law on log-log-scale	Log-normal
asperity	Smoothness of the deviations as a function of $\log(k)$	Log-normal

In order to use this prior within a larger inference model, the underlying space has to be discretized such that the solution of the resulting discrete problem remains consistent with the continuum. We achieve this discretization by means of a truncated Fourier series for s such that s may be written as

$$s = \mathcal{F}^+ \left(\sqrt{P_k(\xi_p)} \xi \right) \quad \text{with} \quad P(\xi) = \mathcal{N}(\xi; 0, \mathbf{1}), \quad (76)$$

where \mathcal{F} denotes a discrete Fourier transformation (DFT) and \mathcal{F}^+ its back-transformation. If we additionally evaluate s on a regular grid on Λ , we can replace the DFT with a fast Fourier transformation (FFT) which is numerically more efficient. For a detailed description on how the spatial discretization is constructed please refer to [37,38]. In this work, however, we are primarily interested in evaluating the approximation quality of the proposed algorithm geoVI, and therefore, from now on, we regard all inference problems involving this random process to be high, but finite, dimensional Bayesian inference problems and ignore the fact that it was constructed from a corresponding continuous, infinite dimensional, inference problem.

4.2. Log-normal Process with Noise Estimation

As a first example we consider a log-normal process e^s , defined over a one-dimensional space, with s being a priori distributed according to the aforementioned Gaussian process prior with unknown power spectrum. The observed data d (see top panel of Figure 7) consists of a partially masked realization of this process subject to additive Gaussian noise with standard deviation σ_n . In addition to s and its power spectrum P_s , we also assume σ_n to be unknown prior to the observation and place a log normal prior on it. Therefore the corresponding likelihood takes the form

$$P(d|s, \sigma_n) = \mathcal{N}(d; Re^s, \sigma_n^2). \quad (77)$$

We apply the geoVI algorithm (Figure 7), the MGVI algorithm (Figure 8), and an HMC sampler (Figure 9) to this problem and construct a set of 3000 approximate posterior samples for all methods. The HMC results serve as the true reference here, as the true posterior distribution is too high dimensional to be accessible directly and HMC is known to reproduce the true posterior in the limit of infinite samples. Considering solely the reconstruction of e^s , we see that both methods, geoVI and MGVI, agree with the true signal largely within their corresponding uncertainties. Overall we find that the geoVI solution is slightly closer to the ground truth compared to MGVI and the posterior uncertainty is smaller for geoVI in most regions, with the exception of the unobserved region, where it is larger compared to MGVI (see residual plot of Figures 7 and 8). In this region MGVI appears to slightly underestimate the posterior uncertainty. In addition, in the bottom panels of Figures 7 and 8, we depict the posterior distribution of the noise standard deviation σ_n as well as the posterior mean of the power spectrum P_s , together with corresponding posterior samples. Here the difference between geoVI and MGVI becomes evident more visibly, which, to some degree, is to be expected due to the more non-linear coupling of σ_n and P_s to the data compared to e^s . Indeed we find that the posterior distribution of σ_n recovered using MGVI is overestimating the noise level of the reconstruction. The geoVI algorithm

is also slightly overestimating σ_n , however we find that for geoVI the posterior yields $\sigma_n^{\text{geoVI}} = 0.220 \pm 0.026$ which places the true value of $\sigma_n = 0.2$ approximately 0.8-sigma away from the posterior mean. In contrast, for MGVI, we get that $\sigma_n^{\text{MGVI}} = 0.233 \pm 0.011$ for with the ground truth is almost a 3-sigma event. Considering the HMC results (bottom panel of Figure 9), the geoVI results appear to be closer to the HMC distribution compared to MGVI, although the HMC distribution for σ_n is broader and even closer to the true value than geoVI. In addition we find that the overall reconstruction quality of the power spectrum P_s is significantly increased when moving from MGVI to geoVI. While MGVI manages to recover the overall slope of the power-law, it fails to reconstruct the deviations from this power-law as well as the overall statistical properties of P_s as encoded in the parameters of Table 1. In contrast, the geoVI algorithm is able to pick up some of these features and recovers posterior statistical properties of the power spectrum similar to the ground truth. In addition the posterior uncertainty appears to be on a reasonable scale, as opposed to the MGVI reconstruction which significantly underestimates the posterior uncertainty. The structures on the smallest scales (largest values for $|k|$), however, appear to be underestimated by the geoVI mean, although the posterior uncertainty increases significantly in these regimes. In comparison to HMC we find that the results are in agreement for the large scales, although the geoVI uncertainties appear to be slightly larger. On small scales, the methods deviate stronger, and the under-estimation of the spectrum seen by geoVI is absent in the HMC results.

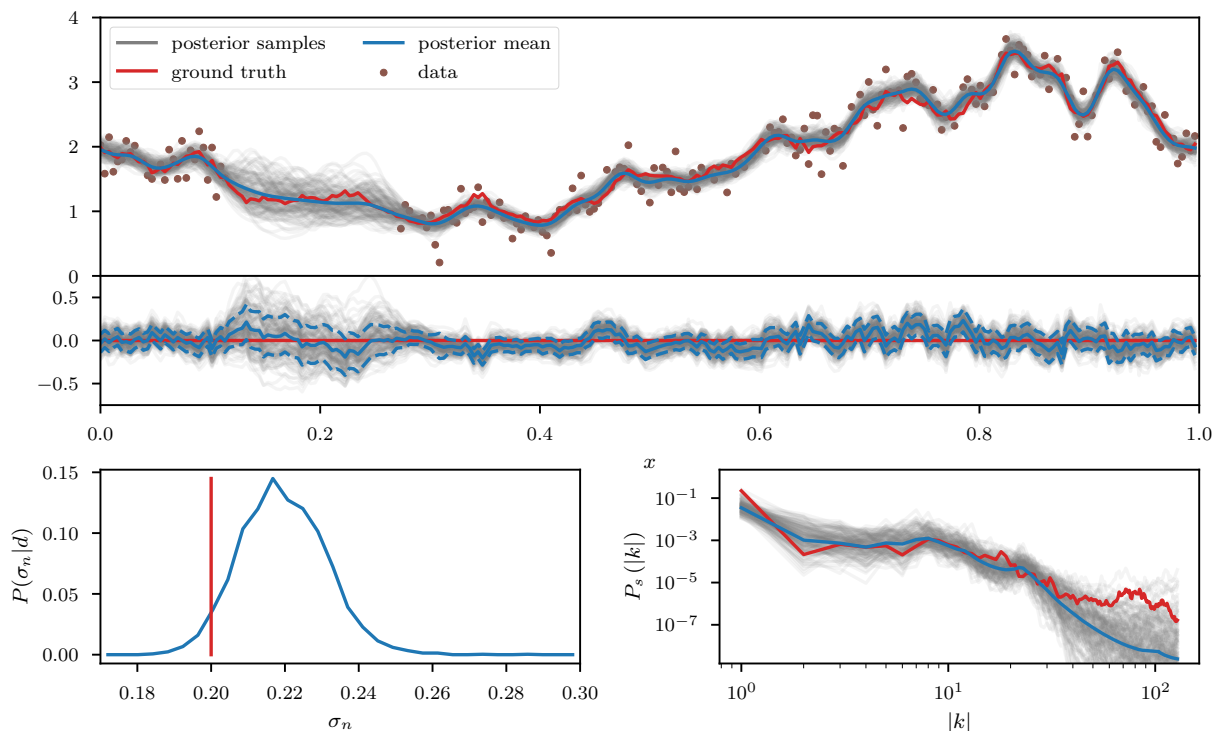


Figure 7. Posterior approximation using the geoVI algorithm for the log-normal process described in Section 4.2. Top: The ground truth realization of the log-normal process e^s (red line) and the corresponding data (brown dots) used for reconstruction. The blue line is the posterior mean, and the gray lines are a subset of the posterior samples obtained from the geoVI approximation. Below we depict the residual between the ground truth and reconstruction, including the residuals for the posterior samples. The blue dashed line corresponds to the one-sigma uncertainty of the reconstruction. Bottom left: Approximation to the marginal posterior distribution (blue) of the noise standard deviation σ_n . The red vertical line indicates the true value of $\sigma_n = 0.2$ used to construct the data. Bottom right: Power spectrum P_s of the logarithmic quantity s . Red displays the ground truth, blue the posterior mean, and the gray lines are posterior samples of the power spectrum.

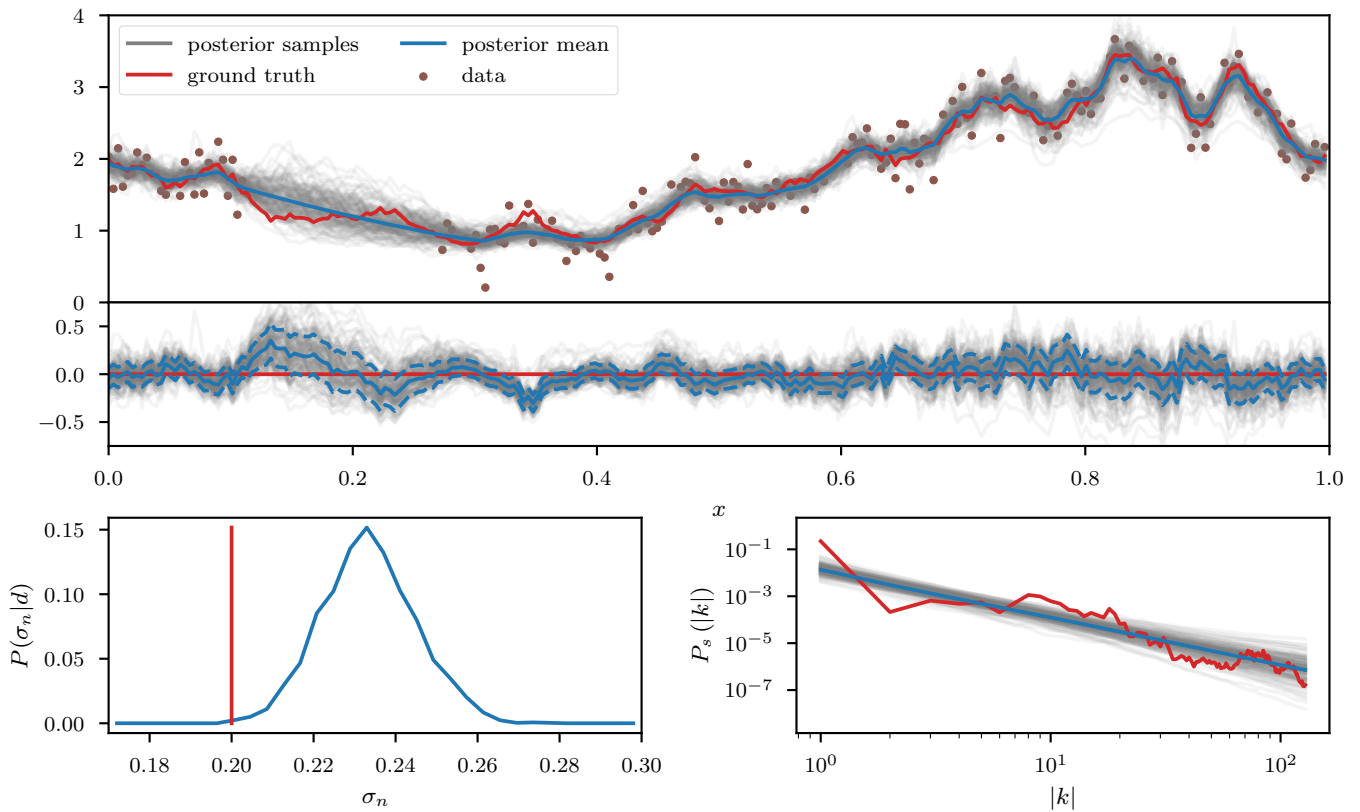


Figure 8. Same setup as in Figure 7, but for the approximation using the MGVI algorithm.

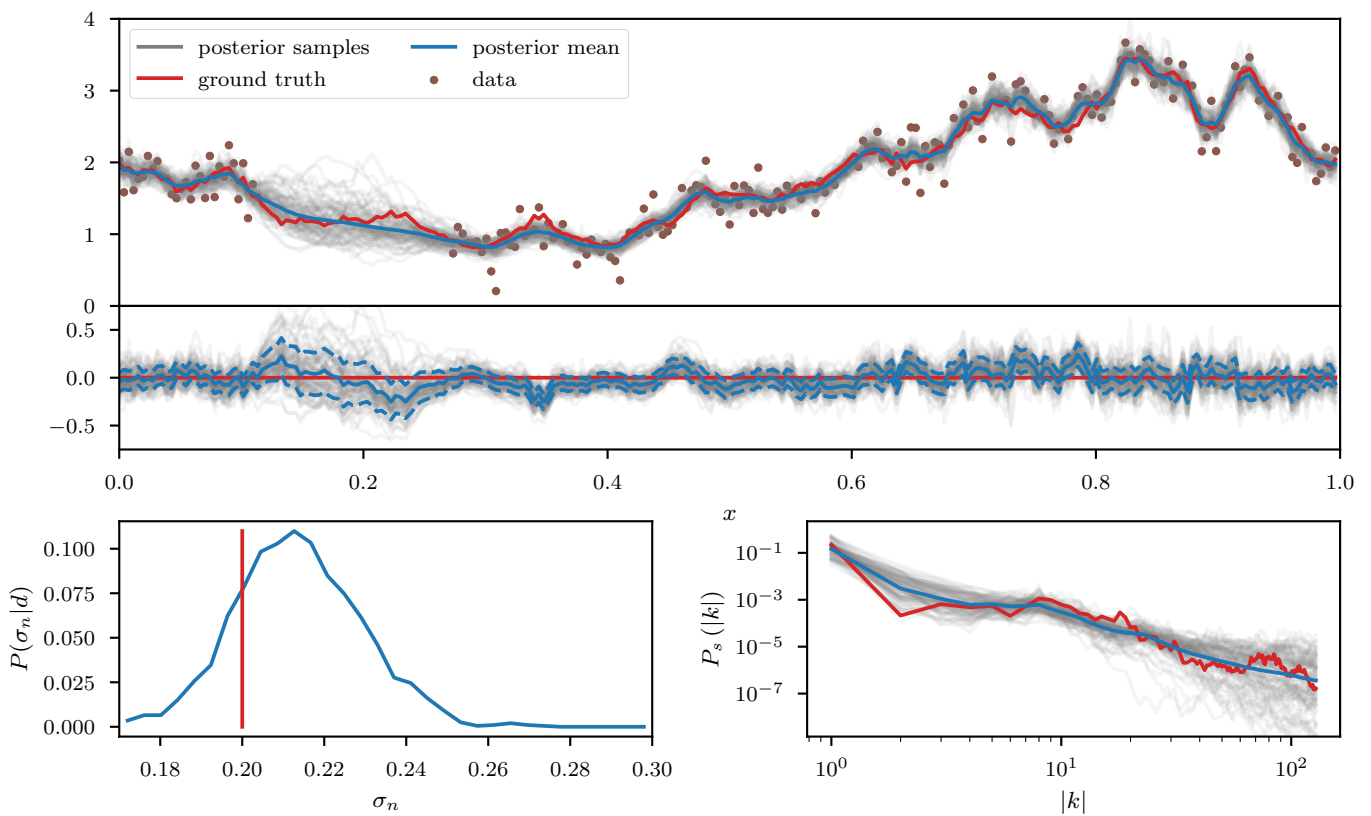


Figure 9. Same setup as in Figure 7, but for the approximation using the HMC sampling.

To further study the posterior distribution of the various scalar parameters that enter the power spectrum model (see Table 1), as well as the noise standard deviation σ_n , we depict the reconstructed marginal posterior distributions for all pairs of inferred scalar parameters. Figures 10–12 show the posterior distributions recovered using geoVI, MGVI, and HMC, respectively. All parameters are displayed in their corresponding standard coordinate system, i.e., they all follow a zero-mean unit variance normal distribution prior to the measurement. From an inference perspective, some of these parameters are very challenging to reconstruct, as their coupling to the observed data is highly non-linear and influenced by many other parameters of the model. In turn, their values are highly influential to the statistical properties of more interpretable variables such as the observed signal e^s and its spectrum P_s . We see that despite these challenges the geoVI posterior appears to give reasonable results, that are largely in agreement with the ground truth, within uncertainties. Thus, the algorithm appears to be able to pick up parts of the non-linear structure of the posterior, which is validated when compared to the MGVI algorithm, as for these parameters the MGVI reconstruction (Figure 11) does not yield reliable results any more. This is to be expected in case of significant non-linearity as MGVI is the first order approximation of geoVI. In comparison to HMC (Figure 12), however, we find that there remain some differences in the recovered posterior distributions. The HMC results regarding the “fluctuations” and “noise std.” parameters are more centered on the ground truth and in particular the posterior distribution of the “slope” parameter is significantly different and more constrained, compared to the geoVI results. These differences indicate that there remain some limitations to the recovered geoVI results in the regime of highly non-linear parameters of the model which we may associate to the theoretical limitations discussed in Section 2.2.

4.3. Separation of Diffuse Emission from Point Sources

In a second inference problem we consider the imaging task of separating diffuse, spatially extended and correlated emission e^s from, bright, but uncorrelated point sources p in an image. This problem is often encountered within certain astrophysical imaging problems [39] where the goal is to recover the emission of spatially extended structures such as gas or galactic dust. This emission usually gets superimposed by the bright emission of stars (point sources) in the image plane, and only their joint emission can be observed. In this example we assume that the emission is observed through a detection device that convolves the incoming emission with a spherical symmetric point spread function R and ultimately measures photon counts on a pixelated grid. Specifically we may define a Poisson process with count rate

$$\lambda = R(p + e^s) \quad \text{with} \quad P(d|\lambda) = \prod_i \frac{(\lambda_i)^{d_i} e^{-\lambda_i}}{(d_i)!}, \quad (78)$$

where i labels the pixels of the detector. We assume the diffuse emission to follow a statistically homogeneous and isotropic log-normal distribution with unknown prior power spectrum. Thus, s is again distributed according to the prior process previously given in Section 4.1. The point sources follow an inverse-gamma distribution at every point (x, y) of the image plane, given as

$$P(p_{xy}) = \frac{q^\alpha}{\Gamma(\alpha)} (p_{xy})^{-\alpha-1} \exp\left(-\frac{q}{p_{xy}}\right), \quad (79)$$

where in the particular example we used $(\alpha, q) = (2, 3)$. A visualization of the problem setup with the various stages of the observation process is given in Figure 13.

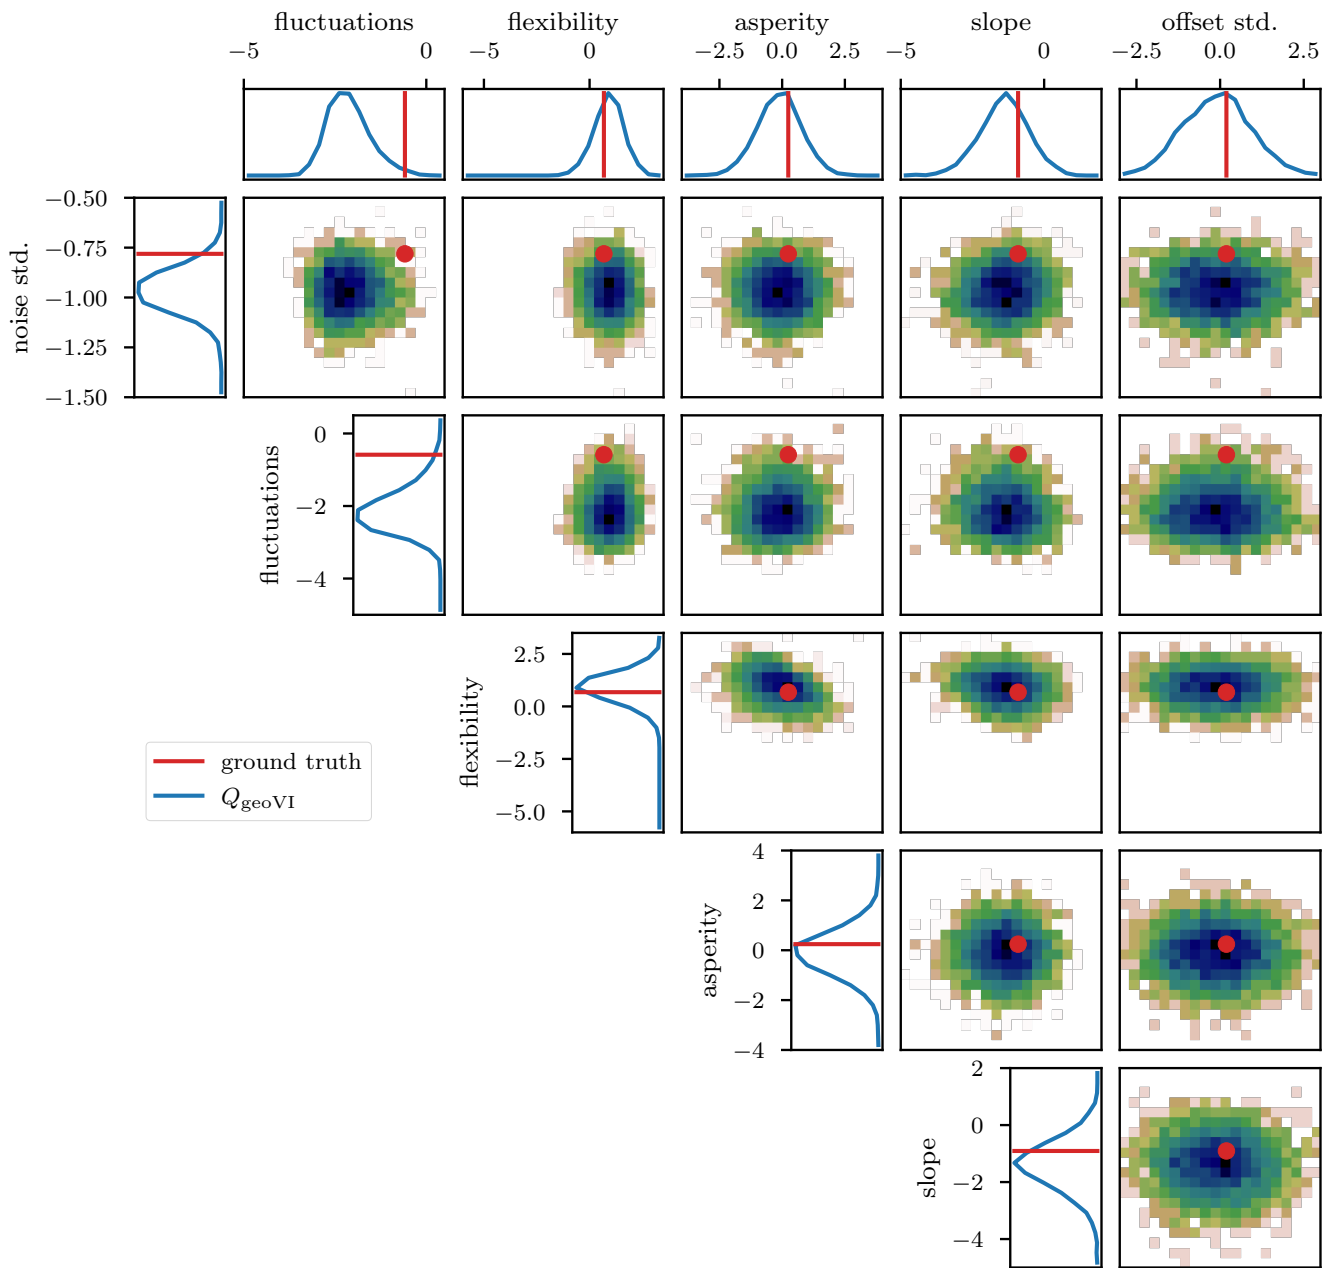


Figure 10. Posterior distributions of the scalar parameters that enter the forward model of the power spectrum (Table 1), and the noise standard deviation. All parameters, including the noise parameter, are given in their corresponding prior standard coordinate system, i.e., have a normal distribution with zero mean and variance one as a prior distribution. Each square panel corresponds to the joint posterior of the parameter in the respective row and column. In addition, for each row and each column the one-dimensional marginal posteriors of the corresponding parameter are displayed as blue lines. The red lines in the 1-D, and the red dots in the 2-D plots denote the values of the ground truth used to realize the ground truth values of the spectrum P_s , the signal e^s , and finally the observed data d .

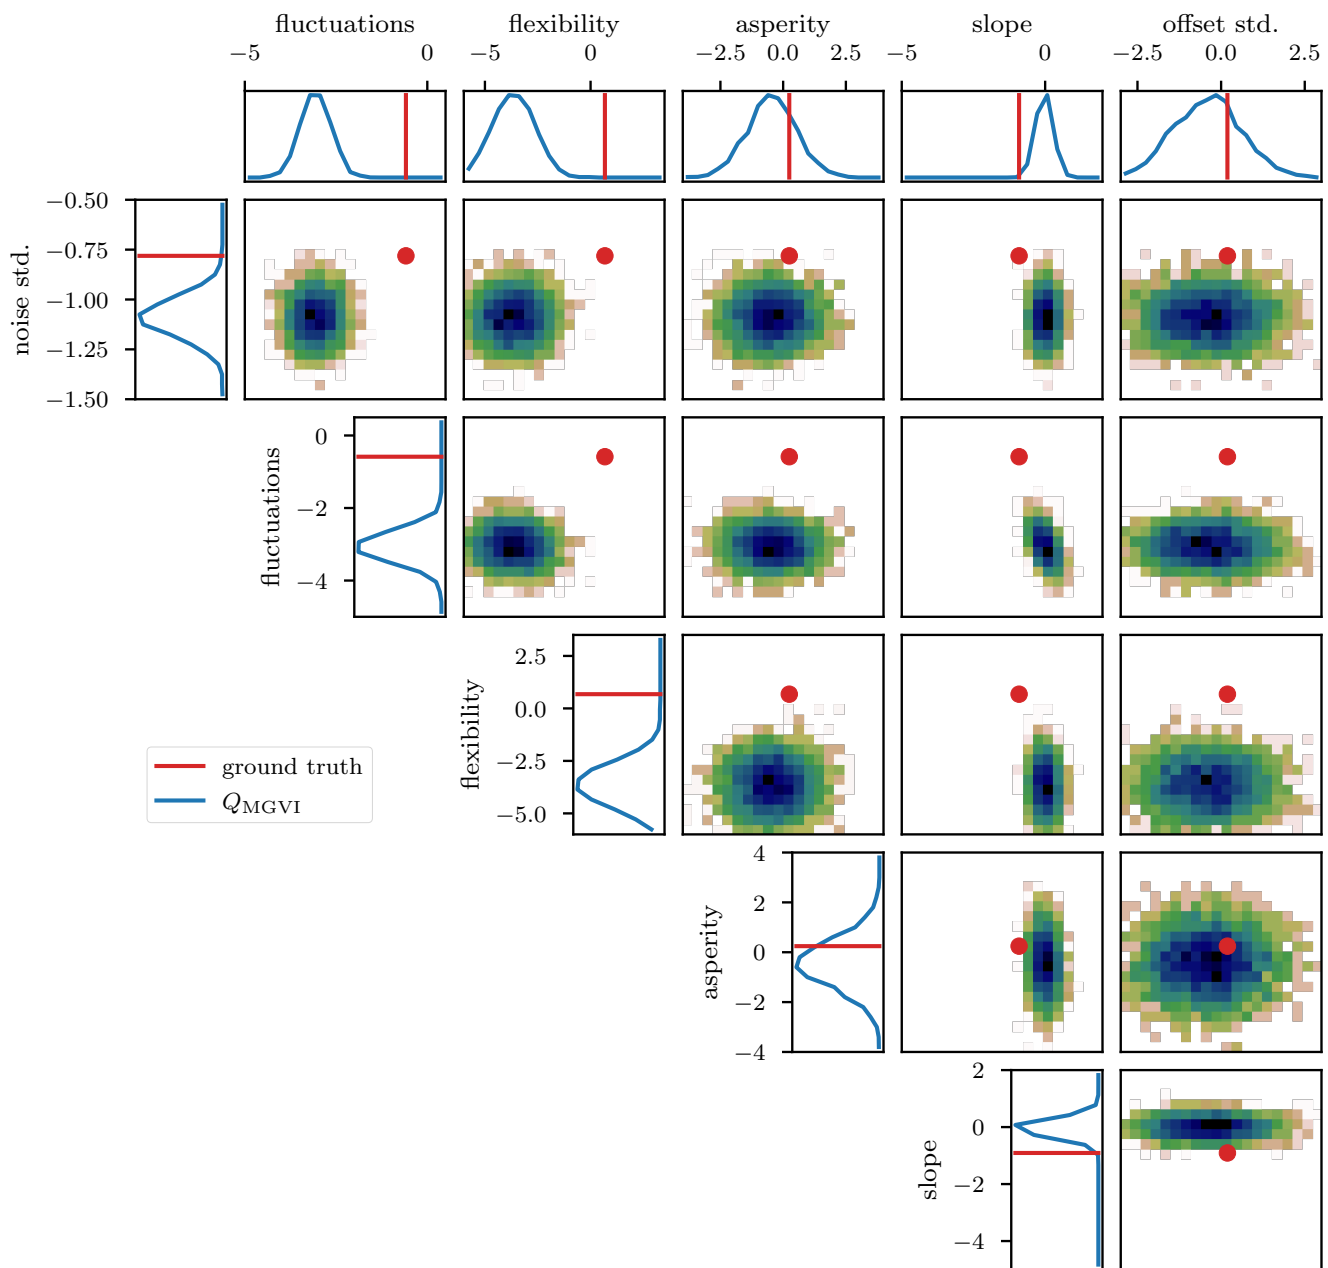


Figure 11. Same setup as in Figure 10, but for the approximation using the MGVI algorithm.

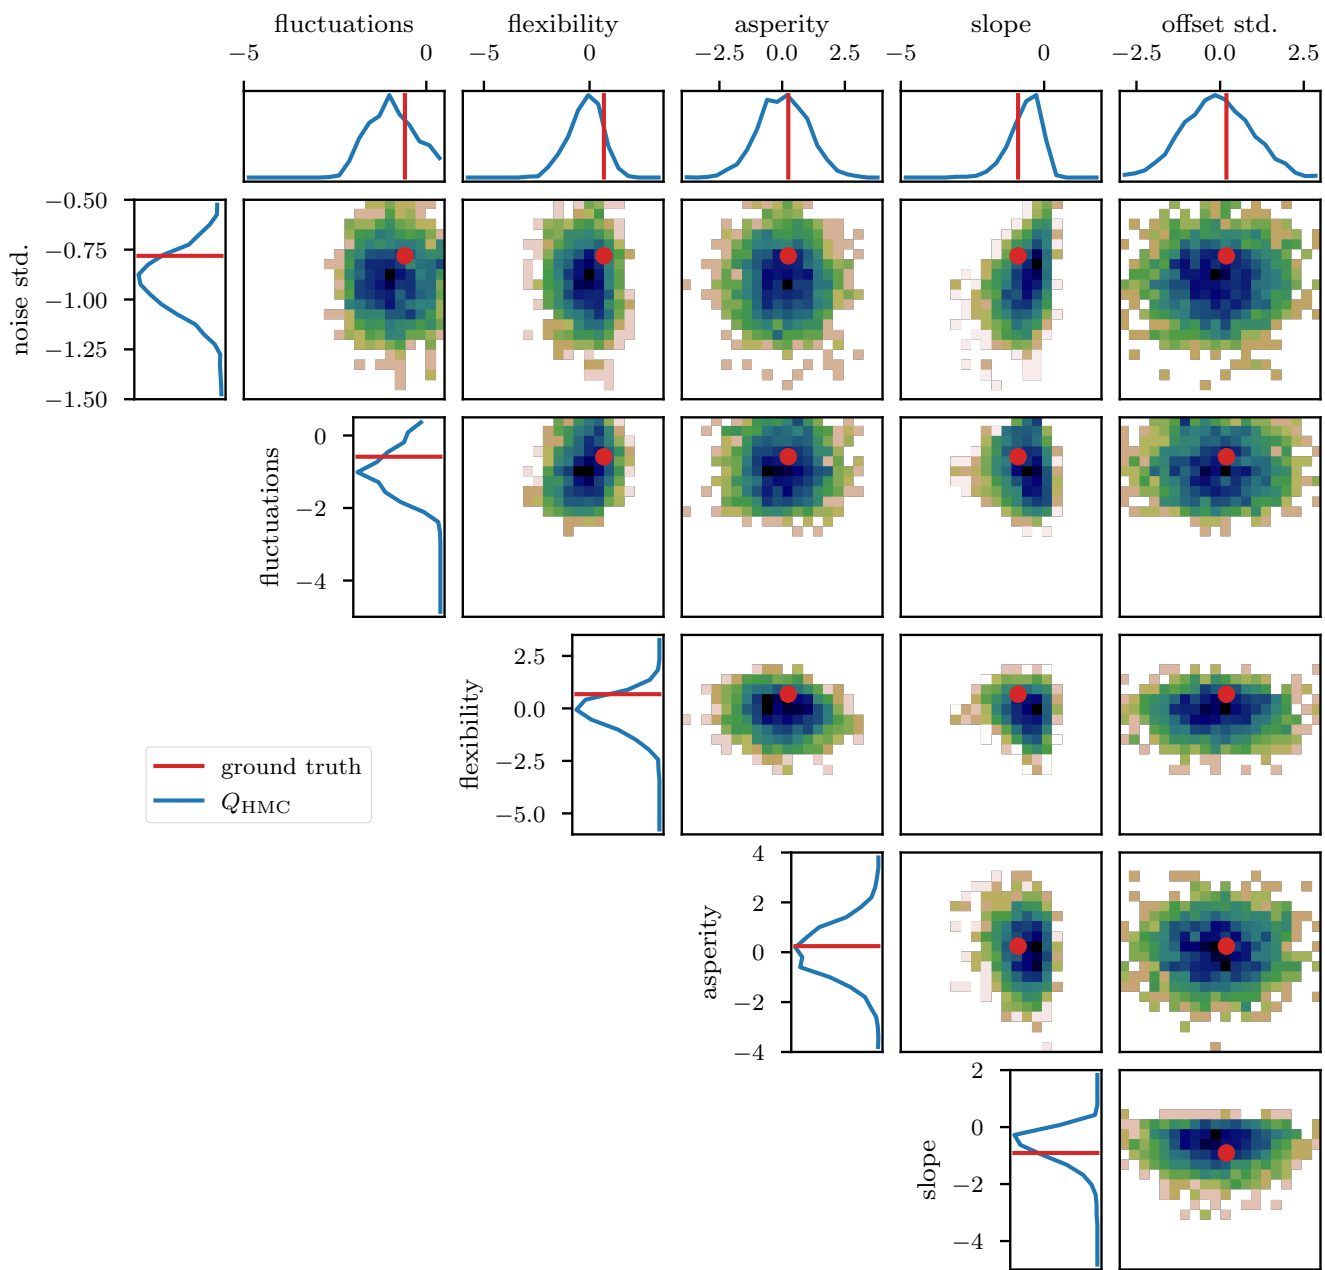


Figure 12. Same setup as in Figure 10, but for the approximation using HMC sampling.

We employ the geoVI and MGVI algorithms to infer all, a priori standard distributed, degrees of freedom of the model and recover the power spectrum P_s for the diffuse emission together with its hyper parameters, the realized emission e^s and the point sources p , from the Poisson count data d . The reconstructed two dimensional images of p and e^s are displayed in Figure 14 together with the recovered count rate λ , and compared to their respective ground truth. We find that in this example there is barely a visible difference between the reconstructed diffuse emission of MGVI and geoVI. Both reconstructions are in good agreement with the ground truth. For the point sources, we find that the brightest sources are well recovered by both algorithms, while geoVI manages to infer a few more of the fainter sources as opposed to MGVI. Nevertheless, for both algorithms, the posterior mean does not recover very faint sources present in the true source distribution. This can also be seen in Figure 15, where we depict the per-pixel flux values for all locations in the image against their reconstructed values, for both, the diffuse emission and the point

sources. We find that the MGVI and the geoVI are, on average, in very good agreement. It is noteworthy that the deviations between the true and the reconstructed flux values increases towards smaller values, which is to be expected due to the larger impact of the Poisson noise. For the spatially independent point sources, there appears to be a transition regime around a flux of ≈ 50 , below which point sources become barely detectable. All in all, for the diffuse emission as well as the point sources, both reconstruction methods apparently yield similar results, consistent with the ground truth. In addition, in Figure 16, we depict the inferred power spectra. We find that the overall shape is reconstructed well by both algorithms, but smaller, more detailed features can only be recovered using geoVI. In addition we find that the statistical properties of the spectrum, as indicated by, e.g., the roughness of the spectrum as a function of the Fourier modes $|k|$, are well reconstructed by geoVI and in agreement with the true spectrum, whereas the MGVI reconstruction, including the posterior samples, appear to be systematically too smooth compared to the ground truth. As discussed in the previous example in more detail, the parameters that enter the model to determine these properties of the spectrum are highly non-linearly coupled and influenced by the observed data and therefore the linear approximation as used in MGVI becomes, at some point, invalid.

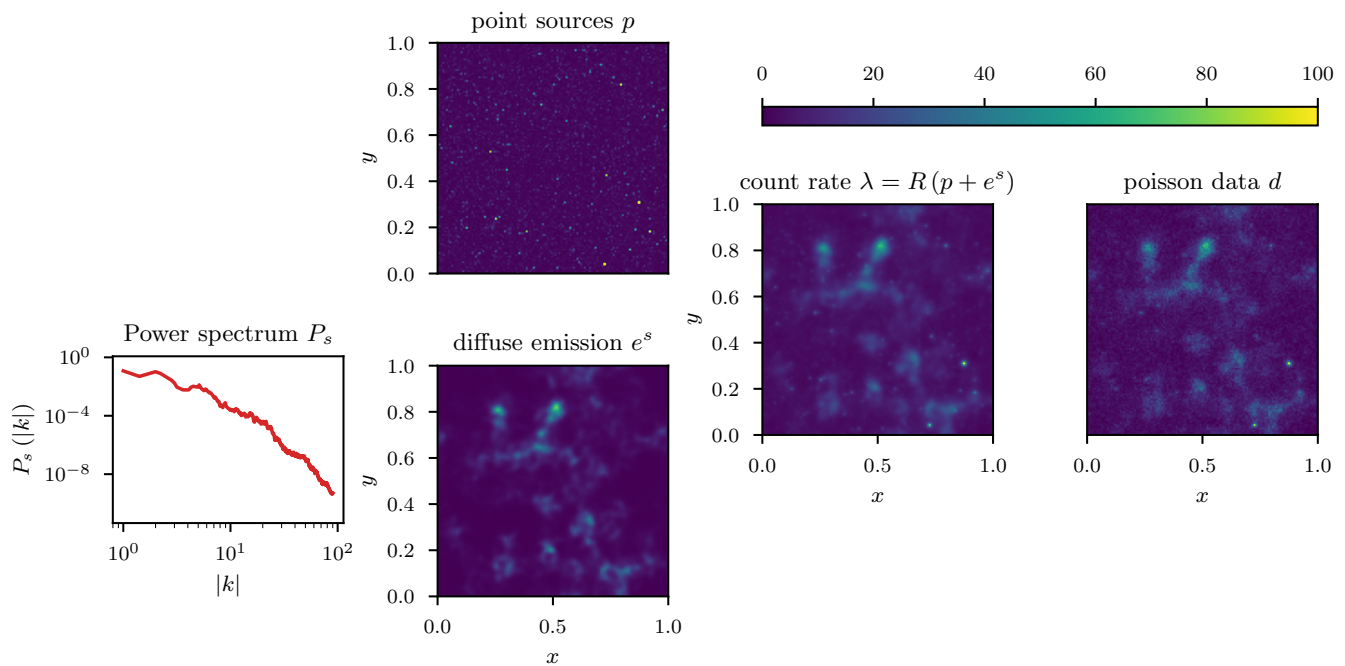


Figure 13. Graphical setup of the separation problem discussed in Section 4.3. Random realization of the power spectrum P_s (left) which is used to generate the log-signal s , which, after exponentiation, models the diffuse emission on the sky e^s . The point sources p (top panel in the middle), which are a realization of the position-independent inverse-gamma process, get combined with the diffuse emission and the result is convolved with a spherical symmetric point spread function R to yield the per-pixel count rate λ which is ultimately used as the rate in a Poisson distribution used to realize the count data d .




Figure 14. Comparison of the ground truth (top row) to the geoVI (middle row) and the MGVI (bottom row) algorithms. The middle and bottom rows show the posterior means for (from left to right) the point sources p , the diffuse emission e^s , and the count rate λ .




Figure 15. Comparison of the per-pixel flux between the ground truth (y-axis) and the reconstruction (x-axis) for the diffuse emission e^s (top row), and the point sources p (bottom row). The left column shows the geoVI result where the density of pixels is color-coded ranging from blue, where the density is highest, to green towards lower densities. The red lines indicate contours of equal density. The right column displays the same for the MGVI reconstruction, with the corresponding density contours now displayed in light blue. The red dashed contours are the density contours of the geoVI case, shown for comparison.

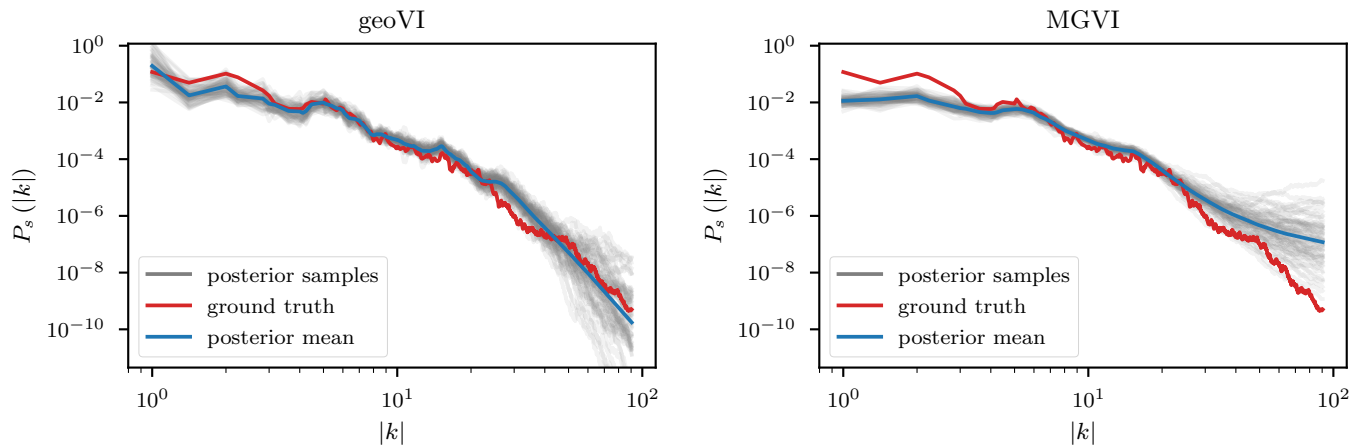


Figure 16. Power spectrum P_s of the logarithm of the diffuse emission s . The red line is the ground truth, the blue line the posterior mean, and the gray lines a subset of posterior samples for the geoVI (left) and MGVI (right) approximations.

5. Further Properties and Challenges

Aside from the apparent capacity to approximate non-linear and high-dimensional posterior distributions, there are some further properties that can be derived from geoVI and the associated coordinate transformation. In the following, we show how to obtain a lower bound to the evidence using the geoVI results. Furthermore, we outline a way to utilize the coordinate transformation in the context of Hamilton Monte-Carlo (HMC) sampling. Finally, some limitations remain to the approximation capacity of geoVI in its current form, which are discussed in Section 5.3.

5.1. Evidence Lower Bound (ELBO)

With the results of the variational approximation at hand, we can provide an Evidence lower bound (ELBO). To this end consider the Hamiltonian $\mathcal{H}(\xi|d)$ of the posterior which takes the form

$$\mathcal{H}(\xi|d) = \mathcal{H}(\xi, d) - \mathcal{H}(d) = \mathcal{H}(d|\xi) + \frac{1}{2}\xi^T \xi + \frac{1}{2} \log(|2\pi\mathbb{1}|) - \mathcal{H}(d), \quad (80)$$

and the Hamiltonian of the approximation \mathcal{H}_Q as a function of r , given as

$$\mathcal{H}_Q(r|\bar{\xi}) = \frac{1}{2}g(\bar{\xi} + r; \bar{\xi})^T g(\bar{\xi} + r; \bar{\xi}) + \frac{1}{2} \log(|2\pi\mathbb{1}|) - \frac{1}{2} \log(|\tilde{\mathcal{M}}(\bar{\xi} + r)|). \quad (81)$$

Using these Hamiltonians, we may write the variational approximation as

$$\begin{aligned} \text{KL}(Q; P) &= \langle \mathcal{H}(\xi = \bar{\xi} + r|d) \rangle_{Q(r|\bar{\xi})} - \langle \mathcal{H}_Q(r|\bar{\xi}) \rangle_{Q(r|\bar{\xi})} \\ &= \langle \mathcal{H}(\xi = \bar{\xi} + r, d) \rangle_{Q(r|\bar{\xi})} - \mathcal{H}(d) - \langle \mathcal{H}_Q(r|\bar{\xi}) \rangle_{Q(r|\bar{\xi})}. \end{aligned} \quad (82)$$

As $\mathcal{H}(d) = -\log(P(d))$, we can derive a lower bound for the logarithmic evidence $P(d)$ using the KL as

$$\log(P(d)) \geq \log(P(d)) - \text{KL}(Q; P), \quad (83)$$

where the lower bound becomes maximal if the KL becomes minimal. Thus, we may use our final expansion point $\bar{\xi}$ obtained from minimizing the KL together with the Hamiltonians (Equations (80) and (81)) to arrive at

$$\begin{aligned} \log(P(d)) - \text{KL}(Q; P) &= \\ &= \frac{1}{2} \text{tr}(\mathbb{1}) - \left\langle \mathcal{H}(d|\xi = \bar{\xi} + r) + \frac{1}{2} (\bar{\xi} + r)^T (\bar{\xi} + r) + \frac{1}{2} \log(|\tilde{\mathcal{M}}(\bar{\xi} + r)|) \right\rangle_{Q(r|\bar{\xi})} \\ &\approx \frac{1}{2} \text{tr}(\mathbb{1}) - \frac{1}{N} \sum_{i=1}^N \left(\mathcal{H}(d|\xi = \bar{\xi} + r_i^*) + \frac{1}{2} (\bar{\xi} + r_i^*)^T (\bar{\xi} + r_i^*) + \frac{1}{2} \log(|\tilde{\mathcal{M}}(\bar{\xi} + r_i^*)|) \right), \quad (84) \end{aligned}$$

where $\{r_i^*\}_{i \in \{1, \dots, N\}}$ are a set of samples drawn from the approximation $Q(r|\bar{\xi})$. Under the assumption that the log determinant of $\tilde{\mathcal{M}}$ is approximately constant throughout the typical set reached by Q , we may replace its sample average with the value obtained at $\bar{\xi}$ to arrive at

$$\begin{aligned} \log(P(d)) - \text{KL}(Q; P) &\approx \frac{1}{2} \text{tr}(\mathbb{1}) - \frac{1}{2} \log(|\tilde{\mathcal{M}}|) \\ &\quad - \frac{1}{N} \sum_{i=1}^N \left(\mathcal{H}(d|\xi = \bar{\xi} + r_i^*) + \frac{1}{2} (\bar{\xi} + r_i^*)^T (\bar{\xi} + r_i^*) \right), \quad (85) \end{aligned}$$

where we also replaced the metric of the expansion $\tilde{\mathcal{M}}$ with the metric of the posterior \mathcal{M} as they are identical when evaluated at the expansion point $\bar{\xi}$ (see Equation (41)). The assumption that the log determinant does not vary significantly within the typical set is also a requirement for the approximation Q to be a close match for the posterior P and in turn it is a necessary condition for the ELBO to be a tight lower bound to the evidence as only in this case the KL may become small. Therefore Equation (85) is a justified simplification in case the entire variational approximation itself is justified. Nevertheless it may be useful to compute the log determinant also for the posterior samples if feasible, as it provides a valuable consistency check for the method itself.

5.2. RMHMC with Metric Approximation

As initially discussed in the introduction, aside from Variational inference methods there exist Markov chain Monte-Carlo (MCMC) methods that utilize the geometry of posterior to increase sampling efficiency. A recently introduced Hybrid Monte-Carlo (HMC) method called Riemannian manifold HMC (or RMHMC) utilizes the same posterior metric as discussed in this work in order to define a Riemannian manifold on which the HMC integration is performed. As one of the key results presented here yields an approximate isometry for this manifold, we like to study the impact of the proposed coordinate transformation on RMHMC. To do so, recall that in HMC the random variable $\xi \in \mathbb{R}^M$, which is distributed according to a posterior distribution $P(\xi)$, is accompanied by another random variable $p \in \mathbb{R}^M$, called momentum, and their joint distribution $P(\xi, p)$ is factorized by means of the posterior $P(\xi)$, and the conditional distribution $P(p|\xi)$. The main idea of HMC is to regard the joint Hamiltonian $\mathcal{H}(\xi, p) = -\log(P(\xi, p))$ as an artificial Hamiltonian system that can be used to construct a new posterior sample from a previous one by following trajectories of the Hamiltonian dynamics. In particular suppose that we are given some random realization ξ^0 of $P(\xi)$, we may use the conditional distribution $P(p|\xi^0)$ to generate a random realization p^0 . Given a pair (ξ^0, p^0) , HMC solves the dynamical system associated with the Hamiltonian $\mathcal{H}(\xi, p)$ for some integration time t^* , to obtain a new pair (ξ^*, p^*) . As Hamiltonian dynamics is both energy and volume preserving by construction, one can show that if (ξ^0, p^0) is a random realization of $P(\xi, p)$, then also (ξ^*, p^*) is. This procedure may be repeated until a desired number of posterior samples is collected. In practice, the performance of an HMC implementation for a specific distribution $P(\xi)$ strongly depends on the choice of conditional distribution $P(p|\xi)$. To simplify the Hamiltonian trajectories and enable a fast traversal of the posterior,

RMHMC has been proposed which utilizes a position dependent metric for the conditional distribution of the momentum which takes the form

$$P(p|\xi) = \mathcal{N}(p; 0, \mathcal{M}(\xi)), \quad (86)$$

where $\mathcal{M}(\xi)$ denotes the metric associated with the posterior $P(\xi)$ as introduced in Section 2 in Equation (10). The associated Hamiltonian takes the form

$$\mathcal{H}(p, \xi) = \frac{1}{2} p^T \mathcal{M}(\xi)^{-1} p + \frac{1}{2} \log(|\mathcal{M}(\xi)|) + \mathcal{H}(\xi). \quad (87)$$

In direct analogy of the discussion in Section 2, the motivation of utilizing the metric is that the resulting Hamiltonian system can be understood as being defined on the Riemannian manifold associated with \mathcal{M} . Therefore the geometric complexity is absorbed into the shape of the manifold, and the trajectories become particularly simple. In practice, however, numerical integration of the system related to Equation (87) is challenging, as in general $\mathcal{H}(p, \xi)$ is non-separable. Here, our coordinate transformation may come in handy, as a Hamiltonian using the approximated metric $\tilde{\mathcal{M}}$ (Equation (14)) instead of \mathcal{M} becomes separable. Specifically replacing \mathcal{M} in Equation (87) yields

$$\begin{aligned} \mathcal{H}(p, \xi) &= \frac{1}{2} p^T \tilde{\mathcal{M}}(\xi; \bar{\xi})^{-1} p + \frac{1}{2} \log(|\tilde{\mathcal{M}}(\xi; \bar{\xi})|) + \mathcal{H}(\xi) \\ &= \frac{1}{2} p^T \left(\left(\frac{\partial g(\xi; \bar{\xi})}{\partial \xi} \right)^T \frac{\partial g(\xi; \bar{\xi})}{\partial \xi} \right)^{-1} p + \tilde{\mathcal{H}}(\xi; \bar{\xi}). \end{aligned} \quad (88)$$

This modified system allows for a canonical transformation of the form

$$\begin{pmatrix} y \\ v \end{pmatrix} \leftarrow \begin{pmatrix} g(\xi; \bar{\xi}) \\ \left(\frac{\partial g(\xi; \bar{\xi})}{\partial \xi} \right)^T p \end{pmatrix}, \quad (89)$$

in which the Hamiltonian (Equation (88)) takes the form

$$\mathcal{H}(v, y) = \frac{1}{2} v^T v s. + \tilde{\mathcal{H}}(\xi; \bar{\xi})|_{\xi=g^{-1}(y; \bar{\xi})} \equiv T(v) + V(y), \quad (90)$$

and therefore \mathcal{H} is separable in the momenta v and the position y . This separability is an interesting property as it has the potential to simplify the integration step used within RMHMC. However, in its current form, we find that there are multiple issues with this approach that prevent an efficient implementation in practice. For one, the transformation g depends on an expansion point $\bar{\xi}$, which becomes a hyper-parameter of the method that has to be determined (possibly in the warm-up phase). In addition, unlike the direct approach discussed in Section 3.1, we cannot circumvent the inversion of g , which is only implicitly available in general, as it has to be computed for every integration step related to $\mathcal{H}(v, y)$ (Equation (90)). Therefore, numerical integration of the system may be simpler, but evaluation of $\mathcal{H}(v, y)$ becomes more expensive. Finally, the approximation of the metric may become invalid as we move far away from the expansion point $\bar{\xi}$, and therefore the applicability compared to an RMHMC implementation using the full metric \mathcal{M} is limited. Nevertheless we find the existence of a separable approximation to the Hamiltonian system very interesting, and think that the (or a similar) transformation g and its associated coordinate system (y, v) might be of relevance in the future development of RMHMC algorithms.

5.3. Pathological Cases

As discussed in Section 2.1, one property that can violate our assumptions are non-monotonic changes in the metric. To this end, consider a sigmoid-normal distributed random variable, and a measurement subject to additive, independent noise of the form

$$P(d|\xi) = \mathcal{N}(d; \sigma(\sigma_p \xi), \sigma_n^2) \quad \text{with} \quad P(\xi) = \mathcal{N}(\xi; 0, 1), \quad (91)$$

where $\sigma(\bullet)$ denotes the sigmoid function. The resulting posterior, its associated coordinate transformation, as well as its geoVI approximation, is displayed in Figure 17 for a case with $(\sigma_p, \sigma_n, d) = (3, 0.2, 0.2)$. We find that similar to the one-dimensional log-normal example of Section 3.1.2, the approximation quality depends on the chosen expansion point. However, the changes in approximation quality are much more drastic as in the log-normal example. In particular, due to the sigmoid non-linearity, there exists a turning point in the coordinate transformation g , and if we choose an expansion point close to this point, we see that the approximation to the transformation strongly deviates from the optimal transformation as we move away from this point. As a result, in this case the approximation to the posterior (left panel of Figure 17) obtains a heavy tail that is neither present in the true posterior nor the approximation using the optimal transformation. Nevertheless there may very well also exist a case where such a heavy tail is present in the optimal approximation to the transformation. Even in the depicted case, where the tail is only present for sub-optimal choices of the expansion point, an optimization algorithm might have to traverse this sub-optimal region to reach the optimum. Thus, the heavy tail can lead to extreme samples for some intermediate approximation, and therefore the geoVI algorithm could become unstable.

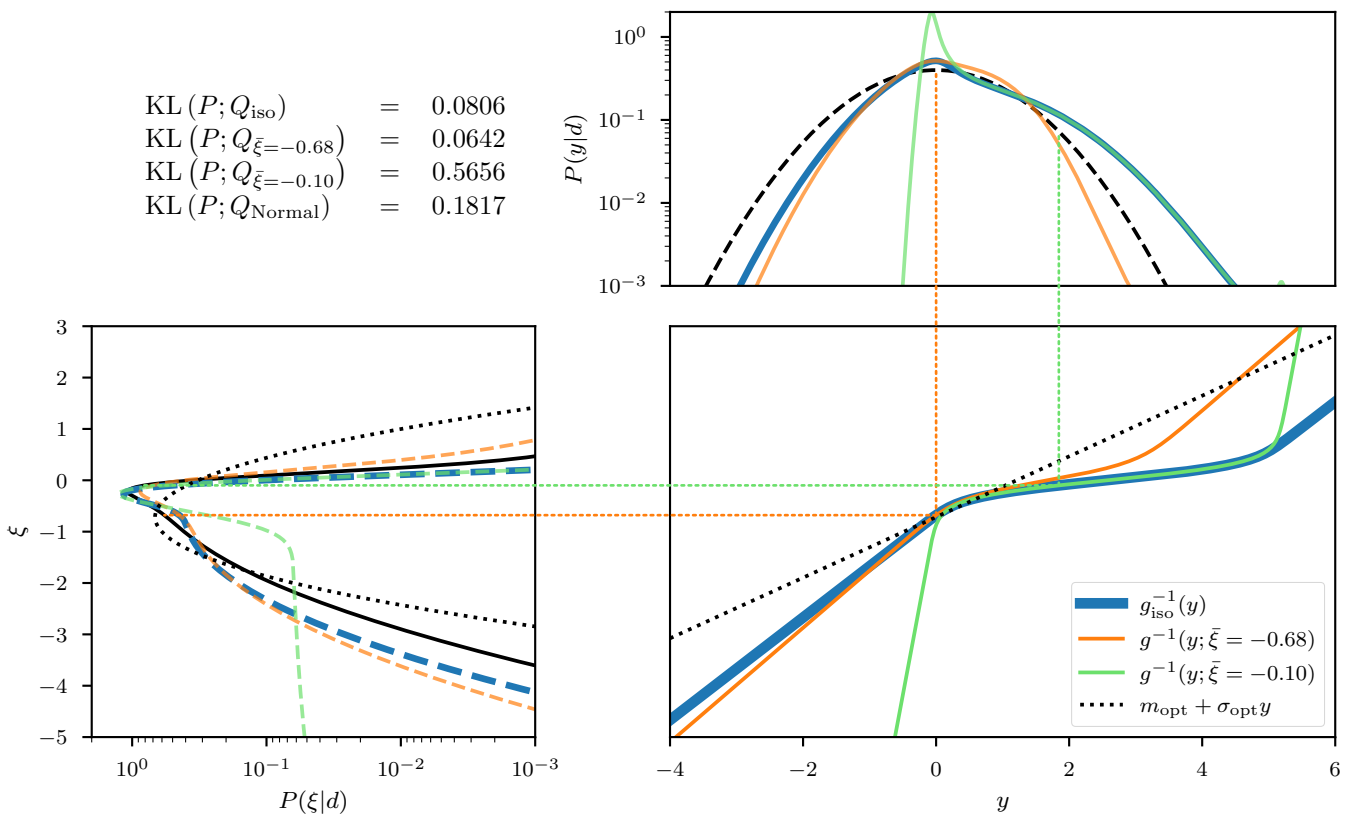


Figure 17. Same setup as in Figure 2, but for the sigmoid-normal distributed case. In addition to the exact isometry g_{iso} , the approximation using the optimal expansion point $\bar{\xi} = -0.68$ and a pathological heavy-tail example using $\bar{\xi} = -0.1$ is displayed.

In a second example we consider a bi-modal posterior distribution, generated from a Gaussian measurement of a polynomial. Specifically we consider a likelihood of the form

$$P(d|\xi) = \mathcal{N}(d; \xi^4 + \xi, 1), \tag{92}$$

with ξ being a priori standard distributed. As can be seen in Figure 18, this scenario leads to a bi-modal posterior distribution with two well separated, asymmetric modes. We find that the geometrically optimal transformation g_{iso} also leads to a bi-modal distribution in the transformed coordinates, however the local asymmetry and curvature of each mode has approximately been removed. Thus, while an approximation of the posterior by means of a single unit Gaussian distribution is apparently not possible, each mode may be approximated individually, at least in case the modes are well separated. If we consider the approximation of the coordinate transformation used within geoVI, and choose as an expansion point the optimal point associated with one of the modes, we get that for the chosen mode the approximation remains valid and the transformation is close to the optimal transformation. However, if we move away from the mode towards the other mode, the approximation quickly deviates from g_{iso} and eventually becomes non-invertible. Therefore only the approximation of one of the modes is possible. Here, care must be taken, as in practical applications the inversion of g is performed numerically and one has to ensure that the inversion does not end up on the second branch of the transformation.




Figure 18. Second pathological example, given as a bi-modal posterior distribution. The setup is similar to Figures 2 and 17, where in this example only the (locally) optimal expansion point $\bar{\xi} = 1.08$ is used.

This summarizes the two main issues that may render a geoVI approximation of a posterior distribution invalid. The challenges and issues related to multi modality appear to be quite fundamental, as in its current form, the geoVI method falls into the category of methods that utilize local information of the posterior which all suffer from the inability to deal with more than a single mode. The problems related to turning points are more

specific for geoVI, and its implications need to be further studied in order to generalize its range of applicability in the future. One promising finding is that this issue appears to be solely related to the local approximation of the transformation with a “bad” expansion point, as the geometrically optimal transformation g_{iso} apparently does not show such behavior. Therefore an extension of the current approximation technique using, e.g., multiple expansion points, or identifying and excluding these “bad” expansion points during optimization, may provide a solution to this problem. At the current stage of the development, however, it is unclear how to incorporate such ideas into the algorithm without loss of the functional form of g that allows for the numerically efficient implementation at hand.

6. Summary and Outlook

In this work we introduced a coordinate transformation for posterior distributions that yields a coordinate system in which the distributions take a geometrically simple form. In particular we construct a metric as the sum of the Fisher metric of the likelihood and the identity matrix for a standard prior distribution, and construct the transformation that relates this metric to the Euclidean metric. Using this transformed coordinate system, we introduce geometric Variational Inference (geoVI), where we perform a variational approximation of the transformed posterior distribution with a normal distribution with unit covariance. As the coordinate transformation is only approximately available and utilizes an expansion point around which it is most accurate, the VI task reduces to finding the optimal expansion point such that the variational KL between the true posterior and the approximation becomes minimal. There exists a numerically efficient realization that enables high-dimensional applications of geoVI because even though the transformation is non-volume preserving, geoVI avoids a computation of the related log-determinant of the Jacobian of the transformation at any point. The expansion point used to generate intermediate samples is only passively updated. Furthermore, the application of the constructed coordinate transformation is similar to the cost of computing the gradient of the posterior Hamiltonian. In addition, to generate random realizations, computing the appearing matrix square root of the metric can be entirely avoided, and the inverse transformation is achieved implicitly by second order numerical inversion.

Despite being an approximation method, we find that geoVI is successfully applicable in non-linear, but uni-modal settings, which we demonstrated with multiple examples. We see that non-linear features of the posterior distribution can accurately be captured by the coordinate transformation in low-dimensional examples. This property may translate into high dimensions, as it increases the overall reconstruction quality there when compared to its linearized version MGVI. Nevertheless we also find remaining pathological cases in which further development is necessary to achieve a good approximation quality.

In addition to posterior approximation, geoVI results can be used in order to provide an evidence lower bound (ELBO) which is used for model comparison. Finally we demonstrate the overlap to another posterior sampling technique based on Hamilton Monte-Carlo (HMC), that utilizes the same metric used in geoVI, called Riemannian manifold HMC.

All in all, the geoVI algorithm, and more generally the constructed approximative coordinate transformation, are a fast and accurate way to approximate non-linear and high-dimensional posterior distributions.

Author Contributions: Conceptualization, P.F. and R.L.; methodology, P.F. and R.L.; software, P.F.; validation, R.L. and T.A.E.; writing—Original draft preparation, P.F.; writing—Review and editing, R.L. and T.A.E.; visualization, P.F.; supervision, T.A.E.; project administration, T.A.E. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Not applicable.

Acknowledgments: The authors would like to thank Philipp Arras for their detailed feedback on the manuscript, Sebastian Hutschenreuther for their hands on feedback to the early versions of

geoVI, Jakob Knollmüller for the development of the MGVI algorithm, and Martin Reinecke for his contributions to NIFTy.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

VI	Variational Inference
MCMC	Markov-Chain Monte-Carlo
geoVI	geometric Variational Inference
VB	Variational Bayes'
RMHMC	Riemannian manifold Hamilton Monte-Carlo
HMC	Hamilton (Hybrid) Monte-Carlo
MGVI	Metric Gaussian Variational Inference
NIFTy	Numerical Information Field Theory
KL	Kullback–Leibler divergence
MVP	matrix vector product
MAP	Maximum a posterior
DFT	Discrete Fourier transform
FFT	Fast Fourier transform
ELBO	Evidence lower bound

Appendix A. Likelihood Transformations

In order to construct the coordinate transformation $x(\xi)$ introduced in Section 2.1, we require that the Fisher metric of the likelihood $\mathcal{M}_{d|\xi}$ may be written as the pullback of the Euclidean metric. Recall that the likelihood expressed in coordinates ξ is obtained from the likelihood $P(d|s')$ with $s' = f'(\xi)$ (see Equation (4)). Therefore we may express $\mathcal{M}_{d|\xi}$ as

$$\mathcal{M}_{d|\xi}(\xi) = \left(\frac{\partial s'}{\partial \xi} \right)^T \mathcal{M}_{d|s'} \frac{\partial s'}{\partial \xi}. \quad (\text{A1})$$

Thus, the task reduces to construct a transformation $x(s')$ that recovers $\mathcal{M}_{d|s'}$ from the Euclidean metric if we set the full transformation to be $x(\xi) \equiv x(s' = f'(\xi))$. Specifically we require for $x(s')$

$$\mathcal{M}_{d|s'} \stackrel{!}{=} \left(\frac{\partial x}{\partial s'} \right)^T \frac{\partial x}{\partial s'}. \quad (\text{A2})$$

Below, in Table A1, we give a summary of multiple commonly used likelihoods, their respective Fisher metric, and the associated transformation $x(s')$.

Table A1. List of common likelihood distributions with their respective Hamiltonian $\mathcal{H}(d|s')$, their Fisher Metric $\mathcal{M}(d|s')$, and the associated coordinate transformation $x(s')$ satisfying Equation (A2).

Name	$\mathcal{H}(d s')$	Metric \mathcal{M}	Trafo. $x(s')$
Normal	$\frac{1}{2}(d - s')^T N^{-1}(d - s') + \text{cst.}$	N^{-1}	$\sqrt{N^{-1}}s'$
Poisson	$1^T s' - d^T \log(s') + \text{cst.}$	$1/s'$	$\frac{1}{2}\sqrt{s'}$
Inv. Gamma	$(\alpha + 1)^T \log(s') + \beta^T \left(\frac{1}{s'}\right) + \text{cst.}$	$\frac{\alpha+1}{s'^2}$	$\sqrt{\alpha + 1} \log(s')$
Student-T	$\frac{\theta+1}{2} \log\left(1 + \frac{s'^2}{\theta}\right) + \text{cst.}$	$\frac{\theta+1}{\theta+3}$	$\sqrt{\frac{\theta+1}{\theta+3}}s'$
Bernoulli	$-d^T \log(s') - (1 - d)^T \log(1 - s') + \text{cst.}$	$\frac{1}{s'(1-s')}$	$-2 \tan^{-1}(\sqrt{s'})$

For some likelihoods, however, such a decomposition is not accessible in a simple form. One example that is being used in this work is a normal distribution with unknown mean m and variance v . The Hamiltonian of a one dimensional example takes the form

$$\mathcal{H}(d|m, v) = \frac{1}{2} \frac{(d-m)^2}{v} + \frac{1}{2} \log(v) + \text{cst.}, \quad (\text{A3})$$

and the corresponding Fisher metric for $s' = (m, v)$ is

$$\mathcal{M}_{d|s'} = \begin{pmatrix} \frac{1}{v} & 0 \\ 0 & \frac{1}{2v^2} \end{pmatrix}. \quad (\text{A4})$$

While there is no simple decomposition by means of the Jacobian of some function x , there is an approximation available for which x takes the form

$$x(s') = \begin{pmatrix} \frac{d-m}{\sqrt{v}} \\ \frac{1}{2} \log(v) \end{pmatrix} \quad \text{with} \quad \frac{\partial x}{\partial s'} = \begin{pmatrix} -\frac{1}{\sqrt{v}} & -\frac{d-m}{2v^{3/2}} \\ 0 & \frac{1}{2v} \end{pmatrix}. \quad (\text{A5})$$

We can compute the approximation to the metric and find

$$\left(\frac{\partial x}{\partial s'} \right)^T \frac{\partial x}{\partial s'} = \begin{pmatrix} \frac{1}{v} & \frac{d-m}{2v^2} \\ \frac{d-m}{2v^2} & \frac{1}{4v^2} + \frac{(d-m)^2}{4v^2} \end{pmatrix}. \quad (\text{A6})$$

Note that as opposed to the Fisher metric, this approximation depends on the observed data d . In fact we can recover the Fisher metric from this approximation by taking the expectation value w.r.t. the likelihood. Specifically

$$\left\langle \left(\frac{\partial x}{\partial s'} \right)^T \frac{\partial x}{\partial s'} \right\rangle_{\mathcal{N}(d;m,v)} = \begin{pmatrix} \frac{1}{v} & 0 \\ 0 & \frac{1}{2v^2} \end{pmatrix} = \mathcal{M}_{d|s'}, \quad (\text{A7})$$

and therefore it may be regarded as a local approximation using the observed data. All examples of this work that use a normal distribution where in addition to the mean also the variance is inferred, use this approximation.

Multiple Likelihoods

In general, we may encounter measurement situations where multiple likelihoods are involved, e.g., if we aim to constrain s' with multiple data-sets simultaneously. Specifically consider a set of D data-sets $\{d_i\}_{i \in \{1, \dots, D\}}$, and an associated mutually independent set of likelihoods, such that the joint likelihood takes the form

$$P(d_1, \dots, d_D | s') = \prod_{i=1}^D P(d_i | s'), \quad (\text{A8})$$

we get that the corresponding Fisher metric takes the form

$$\mathcal{M}_{d_1, \dots, d_D | s'}(s') = \sum_{i=1}^D \mathcal{M}_{d_i | s'}. \quad (\text{A9})$$

If we assume that we have, for every individual metric $\mathcal{M}_{d_i | s'}$, an associated transformation $x_i(s')$ available that satisfies Equation (A2), we see that we can stack them together to form a combined transformation

$$x(s') \equiv (x_1(s'), \dots, x_D(s'))^T, \quad (\text{A10})$$

that automatically satisfies (A2) for the joint metric $\mathcal{M}_{d_1, \dots, d_D | s'}$.

Appendix B. Correlated Field Model

Here, we give a brief description of the generative model for power spectra and resulting Gaussian processes used in Section 4. For a detailed and extended derivation please refer to [32].

A random realization $s \in \mathcal{L}(\Lambda)$ of a statistically homogeneous and isotropic Gaussian process $P(s)$, defined over an L dimensional domain $\Lambda = [0, 1]^L$, subject to periodic boundary conditions along each dimension, may be represented as a Fourier series via

$$s_x = \left(\mathcal{F}^\dagger A \tilde{\zeta}\right)_x \equiv \sum_k e^{-2\pi i k x} A(|k|) \tilde{\zeta}_k \quad \text{with} \quad \tilde{\zeta}_k \sim \mathcal{N}(\tilde{\zeta}; 0, 1) \quad \forall k, \quad (\text{A11})$$

where $k = (k_1, \dots, k_L) \in \mathcal{Z}^L$ is a multi-index labeling the individual Fourier components, and $|k|$ denotes its Euclidean norm. In order to discretize s on a computer, we may truncate this Fourier series, i.e., by replacing the infinite index k with a finite index that truncates at some maximal k_{\max} . The operator \mathcal{F} denotes the Fourier transformation and \mathcal{F}^\dagger its corresponding back-transformation (or their discrete versions in case of truncation). The so-called amplitude spectrum A may be identified with the square root of the power-spectrum P_s of the process (specifically P_s being the eigen-spectrum of the linear operator associated with the covariance of the prior probability $P(s)$). Therefore we proceed to construct a model for A rather than P_s as it is more convenient for a generative model. The non-parametric prior model for A is largely built on the assumption that power spectra (and therefore also amplitude spectra) do not vary arbitrarily for similar $|k|$, which in turn allows us to assume that the values of A are, to some degree, correlated. A prominent example of a physically plausible spectrum is a power-law $P_s = |k|^\alpha$ and therefore it turns out to be more convenient to represent A on a log-log-scale, specifically

$$\tau_l \equiv \log(A(|k|))|_{|k|=e^l}, \quad (\text{A12})$$

since power-laws become straight lines on these scales. As k is a regularly spaced index, the new index $l = \log(|k|)$ is an irregularly spaced index starting from the smallest non-zero mode labeled as l_0 (the origin with $|k| = 0$ is treated separately). To exploit correlations in the prior of τ_l , we define a random process $\tilde{\tau}(l)$ over a continuous domain $O = [l_0, \infty)$ ($O = [l_0, l_{\max} = \log(|k_{\max}|)]$ in the truncated case), and evaluate this process on the irregularly spaced locations on which τ_l is defined. The prior process used for $\tilde{\tau}$ is a Gauss–Markov process given in terms of a linear stochastic differential equation of the form

$$\frac{\partial}{\partial l} \begin{pmatrix} \tilde{\tau}(l) \\ y(l) \end{pmatrix} + \begin{pmatrix} 0 & -1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \tilde{\tau}(l) \\ y(l) \end{pmatrix} = \sigma \begin{pmatrix} \epsilon \eta_l \\ \xi_l \end{pmatrix}, \quad \text{with} \quad \eta_l/\xi_l \sim \mathcal{N}(\eta_l/\xi_l; 0, 1) \quad \forall l \in O. \quad (\text{A13})$$

A Markov process can easily be realized on an irregular grid utilizing its transition probability which in this case takes the form

$$P \left(\begin{pmatrix} \tau_l \\ y_l \end{pmatrix} \middle| \begin{pmatrix} \tau_{l_0} \\ y_{l_0} \end{pmatrix} \right) = \mathcal{N} \left(\begin{pmatrix} \tau_l \\ y_l \end{pmatrix}; \begin{pmatrix} 1 & \Delta_l \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \tau_{l_0} \\ y_{l_0} \end{pmatrix}, \sigma^2 \begin{pmatrix} \Delta_l^3/3 + \epsilon^2 \Delta_l & \Delta_l^2/2 \\ \Delta_l^2/2 & \Delta_l \end{pmatrix} \right) \quad (\text{A14})$$

with $\Delta_l = l - l_0$. We notice that in absence of stochastic deviations (e.g., if $\sigma = 0$), the solution is a straight line with slope y_{l_0} that determines the exponent of the power-law, and therefore becomes a variable of the model on which we place a Gaussian prior with a negative prior mean (to a priori favor falling power laws). The offset τ_{l_0} becomes, after exponentiation, an overall scaling factor that sets the variance of the stochastic process s . Thus, τ_{l_0} (specifically its exponential) is also a variable of the model which we refer to as “fluctuations” in Table 1. Similarly, the zero-mode (i.e., $A(|k| = 0)$), which is not included in τ , is set to be a log-normal distributed random variable which we refer to as “offset std.”. Finally σ (named flexibility) and ϵ (named asperity) become both log-normal distributed

variables (again see Table 1) that determine the variance and shape of the deviations of τ from a straight line (i.e., the deviations of A from a power-law).

References

- Geyer, C.J. Practical markov chain monte carlo. *Stat. Sci.* **1992**, *7*, 473–483. [[CrossRef](#)]
- Brooks, S.; Gelman, A.; Jones, G.; Meng, X.L. *Handbook of Markov Chain Monte Carlo*; CRC Press: Boca Raton, FL, USA, 2011.
- Marjoram, P.; Molitor, J.; Plagnol, V.; Tavaré, S. Markov chain Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci. USA* **2003**, *100*, 15324–15328. [[CrossRef](#)]
- Blei, D.M.; Kucukelbir, A.; McAuliffe, J.D. Variational inference: A review for statisticians. *J. Am. Stat. Assoc.* **2017**, *112*, 859–877. [[CrossRef](#)]
- Hoffman, M.D.; Blei, D.M.; Wang, C.; Paisley, J. Stochastic variational inference. *J. Mach. Learn. Res.* **2013**, *14*, 1303–1347.
- Rezende, D.; Mohamed, S. Variational inference with normalizing flows. In Proceedings of the International Conference on Machine Learning, PMLR, Lille, France, 6–11 July 2015; pp. 1530–1538.
- Kucukelbir, A.; Tran, D.; Ranganath, R.; Gelman, A.; Blei, D.M. Automatic differentiation variational inference. *J. Mach. Learn. Res.* **2017**, *18*, 430–474.
- Kingma, D.P.; Welling, M. Auto-encoding variational bayes. *arXiv* **2013**, arXiv:1312.6114.
- Fox, C.W.; Roberts, S.J. A tutorial on variational Bayesian inference. *Artif. Intell. Rev.* **2012**, *38*, 85–95. [[CrossRef](#)]
- Šmídl, V.; Quinn, A. *The Variational Bayes Method in Signal Processing*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2006.
- Girolami, M.; Calderhead, B. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2011**, *73*, 123–214. [[CrossRef](#)]
- Duane, S.; Kennedy, A.D.; Pendleton, B.J.; Roweth, D. Hybrid monte carlo. *Phys. Lett. B* **1987**, *195*, 216–222. [[CrossRef](#)]
- Betancourt, M. A conceptual introduction to Hamiltonian Monte Carlo. *arXiv* **2017**, arXiv:1701.02434.
- Saha, A.; Bharath, K.; Kurtek, S. A Geometric Variational Approach to Bayesian Inference. *J. Am. Stat. Assoc.* **2020**, *115*, 822–835. [[CrossRef](#)] [[PubMed](#)]
- Knollmüller, J.; Enßlin, T.A. Metric gaussian variational inference. *arXiv* **2019**, arXiv:1901.11033.
- Arras, P.; Baltac, M.; Ensslin, T.A.; Frank, P.; Hutschenreuter, S.; Knollmueller, J.; Leike, R.; Newrzella, M.N.; Platz, L.; Reinecke, M.; et al. Nifty5: Numerical information field theory v5. *Astrophys. Source Code Libr.* **2019**, ascl-1903.008.
- Bogachev, V.I.; Kolesnikov, A.V.; Medvedev, K.V. Triangular transformations of measures. *Sb. Math.* **2005**, *196*, 309. [[CrossRef](#)]
- Fisher, R.A. *Theory of Statistical Estimation*; Cambridge University Press: Cambridge, UK, 1925; Volume 22, pp. 700–725.
- Rao, C.R. Information and the accuracy attainable in the estimation of statistical parameters. In *Breakthroughs in Statistics*; Springer: Berlin, Germany, 1992; pp. 235–247.
- Amari, S.; Nagaoka, H. *Methods of Information Geometry*; Translations of Mathematical Monographs; American Mathematical Society: Providence, RI, USA, 2000.
- Cencov, N.N. *Statistical Decision Rules and Optimal Inference*; Number 53; American Mathematical Society: Providence, RI, USA, 2000.
- Betancourt, M. A general metric for Riemannian manifold Hamiltonian Monte Carlo. In Proceedings of the International Conference on Geometric Science of Information, Paris, France, 28–30 August 2013; Springer: Berlin, Germany, 2013; pp. 327–334.
- Nocedal, J.; Wright, S.J. *Numerical Optimization*, 2nd ed.; Springer: New York, NY, USA, 2006; p. 168.
- Hestenes, M.R.; Stiefel, E. *Methods of Conjugate Gradients for Solving Linear Systems*; NBS: Washington, DC, USA, 1952; Volume 49.
- Hutchinson, M.F. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Commun. Stat. Simul. Comput.* **1989**, *18*, 1059–1076. [[CrossRef](#)]
- Han, I.; Malioutov, D.; Shin, J. Large-scale log-determinant computation through stochastic Chebyshev expansions. In Proceedings of the International Conference on Machine Learning, PMLR, Lille, France, 6–11 July 2015; pp. 908–917.
- Ubaru, S.; Chen, J.; Saad, Y. Fast estimation of $\text{tr}(f(a))$ via stochastic lanczos quadrature. *SIAM J. Matrix Anal. Appl.* **2017**, *38*, 1075–1099. [[CrossRef](#)]
- Fitzsimons, J.; Granzio, D.; Cutajar, K.; Osborne, M.; Filippone, M.; Roberts, S. Entropic trace estimates for log determinants. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*; Springer: Berlin, Germany, 2017; pp. 323–338.
- Hutschenreuter, S.; Anderson, C.S.; Betti, S.; Bower, G.C.; Brown, J.A.; Brüggem, M.; Carretti, E.; Clarke, T.; Clegg, A.; Costa, A.; et al. The Galactic Faraday rotation sky 2020. *arXiv* **2021**, arXiv:2102.01709.
- Welling, C.; Frank, P.; Enßlin, T.; Nelles, A. Reconstructing non-repeating radio pulses with Information Field Theory. *J. Cosmol. Astropart. Phys.* **2021**, *2021*, 071. [[CrossRef](#)]
- Arras, P.; Bester, H.L.; Perley, R.A.; Leike, R.; Smirnov, O.; Westermann, R.; Enßlin, T.A. Comparison of classical and Bayesian imaging in radio interferometry—Cygnus A with CLEAN and resolve. *Astron. Astrophys.* **2021**, *646*, A84. [[CrossRef](#)]
- Arras, P.; Frank, P.; Haim, P.; Knollmüller, J.; Leike, R.; Reinecke, M.; Enßlin, T. The variable shadow of M87. *arXiv* **2020**, arXiv:2002.05218.
- Leike, R.H.; Glatzle, M.; Enßlin, T.A. Resolving nearby dust clouds. *Astron. Astrophys.* **2020**, *639*, A138. [[CrossRef](#)]
- Arras, P.; Frank, P.; Leike, R.; Westermann, R.; Enßlin, T.A. Unified radio interferometric calibration and imaging with joint uncertainty quantification. *Astron. Astrophys.* **2019**, *627*, A134. [[CrossRef](#)]

-
35. Hutschenreuter, S.; Enßlin, T.A. The Galactic Faraday depth sky revisited. *Astron. Astrophys.* **2020**, *633*, A150. [[CrossRef](#)]
 36. Wiener, N. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series, with Engineering Applications*; Stationary Time Series; Technology Press of the Massachusetts Institute of Technology: Cambridge, MA, USA, 1950; Number ix, 163p.
 37. Lassas, M.; Saksman, E.; Siltanen, S. Discretization-invariant Bayesian inversion and Besov space priors. *Inverse Probl. Imaging* **2009**, *3*, 87. [[CrossRef](#)]
 38. Frank, P.; Leike, R.; Enßlin, T.A. Field Dynamics Inference for Local and Causal Interactions. *Ann. Phys.* **2021**, *533*, 2000486. [[CrossRef](#)]
 39. Bertin, E.; Arnouts, S. SExtractor: Software for source extraction. *Astron. Astrophys. Suppl. Ser.* **1996**, *117*, 393–404. [[CrossRef](#)]