

GOAL-ORIENTED OPTIMAL APPROXIMATIONS OF BAYESIAN LINEAR INVERSE PROBLEMS

ALESSIO SPANTINI*, TIANGANG CUI†, KAREN WILLCOX*,
LUIS TENORIO‡, AND YOUSSEF MARZOUK*

Abstract. We propose optimal dimensionality reduction techniques for the solution of goal-oriented linear-Gaussian inverse problems, where the quantity of interest (QoI) is a function of the inversion parameters. These approximations are suitable for large-scale applications. In particular, we study the approximation of the posterior covariance of the QoI as a low-rank negative update of its prior covariance, and prove optimality of this update with respect to the natural geodesic distance on the manifold of symmetric positive definite matrices. Assuming exact knowledge of the posterior mean of the QoI, the optimality results extend to optimality in distribution with respect to the Kullback-Leibler divergence and the Hellinger distance between the associated distributions. We also propose the approximation of the posterior mean of the QoI as a low-rank linear function of the data, and prove optimality of this approximation with respect to a weighted Bayes risk. Both of these optimal approximations avoid the explicit computation of the full posterior distribution of the parameters and instead focus on directions that are well informed by the data and relevant to the QoI. These directions stem from a balance among all the components of the goal-oriented inverse problem: prior information, forward model, measurement noise, and ultimate goals. We illustrate the theory using a high-dimensional inverse problem in heat transfer.

Key words. inverse problems, goal-oriented, Bayesian inference, low-rank approximation, covariance approximation, Riemannian metric, geodesic distance, posterior mean approximation, Bayes risk, optimality

AMS subject classifications. 15A29, 62F15, 68W25

1. Introduction. The Bayesian approach to inverse problems treats the unknown parameters as random variables, endowed with a prior distribution that encodes one’s knowledge before data are collected. The distribution of the data conditioned on any value of the parameters is specified through the likelihood model. Bayes’ rule then combines prior and likelihood information to yield the posterior distribution, i.e., the distribution of the parameters conditioned on the data. The posterior distribution defines the Bayesian solution to the inverse problem. Characterizing this posterior distribution is of primary interest in many engineering and science applications, ranging from computerized tomography and optical imaging to geostatistical modeling. For instance, we might be interested in the posterior marginals, the posterior probability of some functionals of the parameters, or the probability of rare events under the posterior measure. In all these cases we may need to draw samples from the posterior distribution. This sampling task tends to be extremely challenging in large-scale applications, especially when the parameters represent a finite-dimensional approximation to a distributed stochastic process like a permeability or a temperature field. In many applications, however, we are only interested in a particular *function* of the parameters (e.g., the temperature field over a subregion of the entire domain or the probability that the temperature exceeds a critical value). In this paper we exploit such *ultimate goals* to reduce the cost of inversion.

*Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA, {spantini, kwillcox, ymarz}@mit.edu.

†School of Mathematical Sciences, Monash University, Victoria 3800, Australia, tiangang.cui@monash.edu.

‡Applied Mathematics and Statistics, Colorado School of Mines, Golden, CO 80401, USA, ltenorio@mines.edu.

We begin by considering a finite-dimensional linear-Gaussian inverse problem of the form

$$\mathbf{Y} = G \mathbf{X} + \boldsymbol{\mathcal{E}}, \quad (1.1)$$

where $\mathbf{X} \in \mathbb{R}^n$ represents the unknown parameters, $\mathbf{Y} \in \mathbb{R}^d$ denotes the noisy observations, $G \in \mathbb{R}^{d \times n}$ is a deterministic linear forward operator, and $\boldsymbol{\mathcal{E}} \sim \mathcal{N}(0, \Gamma_{\text{obs}})$ is a zero-mean additive Gaussian noise, statistically independent of \mathbf{X} and with covariance matrix $\Gamma_{\text{obs}} \succ 0$. (We use boldface capital letters to denote random vectors.) We prescribe a Gaussian prior distribution, $\mathcal{N}(0, \Gamma_{\text{pr}})$, on \mathbf{X} and assume, without loss of generality, zero prior mean and $\Gamma_{\text{pr}} \succ 0$. One is usually concerned with the posterior distribution of the parameters, $\mathbf{X}|\mathbf{Y} \sim \mathcal{N}(\mu_{\text{pos}}(\mathbf{Y}), \Gamma_{\text{pos}})$,¹ which has mean and covariance given by

$$\mu_{\text{pos}}(\mathbf{Y}) = \Gamma_{\text{pos}} G^\top \Gamma_{\text{obs}}^{-1} \mathbf{Y}, \quad \Gamma_{\text{pos}} = (H + \Gamma_{\text{pr}}^{-1})^{-1}, \quad (1.2)$$

where $H := G^\top \Gamma_{\text{obs}}^{-1} G$ is the Hessian of the negative log-likelihood. In this paper, however, we are not interested in the parameters \mathbf{X} per se, but rather in a quantity of interest (QoI) \mathbf{Z} that is a function of the parameters,

$$\mathbf{Z} = \mathcal{O} \mathbf{X}, \quad (1.3)$$

for some linear and, without loss of generality, full row-rank operator $\mathcal{O} \in \mathbb{R}^{p \times n}$ with $p < n$. Our interests are thus *goal-oriented*, as we wish to characterize only \mathbf{Z} and not the parameters \mathbf{X} . Including such ultimate goals in the inference formulation is an important modeling step in many applications of Bayesian inverse problems. This additional step should reduce the computational complexity of inference by making the ultimate goals explicit. Nevertheless, it is still not clear how to leverage ultimate goals to yield more efficient Bayesian inference algorithms. The present paper will address this issue.

The full Bayesian solution to the goal-oriented inverse problem is the posterior distribution of the QoI, i.e., $\mathbf{Z}|\mathbf{Y}$. It is easy to see that $\mathbf{Z}|\mathbf{Y}$ is once again Gaussian with mean and covariance matrix given by

$$\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) = \mathcal{O} \mu_{\text{pos}}(\mathbf{Y}), \quad \Gamma_{\mathbf{Z}|\mathbf{Y}} = \mathcal{O} \Gamma_{\text{pos}} \mathcal{O}^\top. \quad (1.4)$$

The goal of this paper is to characterize statistically optimal, computationally efficient, and structure-exploiting approximations of the statistics of $\mathbf{Z}|\mathbf{Y}$ whenever the use of direct formulas such as (1.4) is challenging or impractical (perhaps due to high computational complexity or excessive storage requirements). We will approximate $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ as a low-rank negative update of the prior covariance of the QoI. Optimality will be defined with respect to the natural geodesic distance on the manifold of symmetric and positive definite (SPD) matrices [42]. The posterior mean $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ will be approximated as a low-rank function of the data, where optimality is defined by the minimization of the Bayes risk for squared-error loss weighted by $\Gamma_{\mathbf{Z}|\mathbf{Y}}^{-1}$. The essence of these approximations is the restriction of the inference process to directions in the parameter space that are informed by the data relative to the prior *and* that are relevant to the QoI. These directions correspond to the leading generalized eigenpairs of a suitable matrix pencil.

¹ $\mathbf{X}|\mathbf{Y}$ refers to a random variable distributed according to the measure of \mathbf{X} conditioned on \mathbf{Y} .

This paper extends, in several different ways, the work on goal-oriented inference originally presented in [63]. First of all, we introduce the notion of *optimal approximation*, rather than exact computation, for both the posterior covariance matrix and the posterior mean of the QoI. We propose computationally efficient algorithms to determine these optimal approximations. The complexity of our algorithms scales with the intrinsic dimensionality of the goal-oriented problem—which here reflects the dimension of the parameter subspace that is simultaneously relevant to the QoI and informed by the data, as noted above. In particular, the full posterior distribution of the parameters need not be computed at any stage of the algorithms. This is a key contrast with [63]. Moreover, we make it possible to handle high-dimensional QoIs such as those arising from the discretization of a distributed stochastic process. This class of problems is frequently encountered in applications (see, e.g., Section 4).

The ideas and algorithms of this paper are primarily developed in the linear-Gaussian case. On one hand, their application to goal-oriented linear inverse problems is of standalone interest [63]. On the other hand, in the context of high-dimensional nonlinear Bayesian inverse problems, the Gaussian approximation is often the only tractable approximation of the posterior distribution [19, 54]. For example, in [54], a linearization of the parameter-to-observable map and a linearization of the parameter-to-prediction map are performed for a high-dimensional ice sheet model, reducing the nonlinear inverse problem to a goal-oriented linear Gaussian inverse problem. Moreover, the rigorous analysis of dimensionality reduction ideas in linear inverse problems often leads to computationally efficient dimensionality reduction strategies for nonlinear inverse problems² (see, e.g., [97, 31] or [63, 64]).

The remainder of the paper is organized as follows. In Section 2 we introduce statistically optimal approximations of the posterior statistics of the QoI. Section 3 contains a simple proof-of-concept example, while in Section 4 we illustrate the theory using a more realistic inverse problem in heat transfer. Section 5 offers some concluding remarks. Appendix A reviews an important class of metrics between distributions (Rao’s distance), while Appendix B contains the proofs of the main results of this paper.

2. Theory. In this section we introduce optimal approximations of the posterior mean $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ and posterior covariance $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ of the QoI. Section 2.1 reviews a class of metrics between probability distributions and introduces natural loss functions for the approximation of $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ and $\Gamma_{\mathbf{Z}|\mathbf{Y}}$. Section 2.2 then focuses on the approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$, while the posterior mean approximation is addressed in Section 2.3. The main results of this section are Theorem 2.3 and Theorem 2.9.

2.1. Optimality criteria: metrics between distributions. To measure the quality of the approximation of a posterior distribution we employ a metric first introduced by Rao in [84] based on the Fisher information. Rao’s approach to comparing distributions is rooted in differential geometry. The idea is to turn a parametric family of distributions into a Riemannian manifold endowed with a metric based on the geodesic distance [85]. The Riemannian structure is induced by a quadratic form defined by the Fisher information matrix. (See Appendix A for the explicit construction of Rao’s distance.) For a Gaussian family of distributions this metric can be written explicitly at least in two particular cases

²Of course, a comprehensive theory for nonlinear inverse problems will likely require tools beyond those explored in this paper.

[6, 93]. If the family consists of Gaussian distributions with the same covariance matrix, Γ , then the metric reduces to the Mahalanobis distance between the means [67, 85]:

$$d_{\mathcal{R}}(\nu_1, \nu_2) = \|\mu_1 - \mu_2\|_{\Gamma^{-1}}, \quad \nu_1 = \mathcal{N}(\mu_1, \Gamma), \nu_2 = \mathcal{N}(\mu_2, \Gamma), \quad (2.1)$$

where $\|z\|_{\Gamma^{-1}}^2 := z^\top \Gamma^{-1} z$. We will use this metric to define optimal approximations of the posterior mean in Section 2.3. Notice that this metric emphasizes differences in the mean along eigendirections of Γ corresponding to low variance.

If, on the other hand, the family consists of Gaussian distributions with the same mean, μ , then the metric reduces to:

$$d_{\mathcal{R}}(\nu_1, \nu_2) = \sqrt{\frac{1}{2} \sum_i \ln^2(\sigma_i)}, \quad \nu_1 = \mathcal{N}(\mu, \Gamma_1), \nu_2 = \mathcal{N}(\mu, \Gamma_2), \quad (2.2)$$

where (σ_i) are the generalized eigenvalues of the pencil (Γ_1, Γ_2) [55]. That is, (σ_i) are the roots of the characteristic polynomial $\det(\Gamma_1 - \sigma \Gamma_2)$ and satisfy the equation $\Gamma_1 v_i = \Gamma_2 v_i \sigma_i$ for a collection of generalized eigenvectors (v_i) [45]. Since this family of distributions can be identified with the cone of SPD matrices, Sym^+ , (2.2) can also be used as a Riemannian metric on Sym^+ [42, 14]. We will use this Riemannian metric to define optimal approximations of the posterior covariance matrix in Section 2.2. This metric on Sym^+ is also the unique geodesic distance that satisfies the following two important invariance properties:

$$d_{\mathcal{R}}(A, B) = d_{\mathcal{R}}(A^{-1}, B^{-1}) \quad \text{and} \quad d_{\mathcal{R}}(A, B) = d_{\mathcal{R}}(MAM^\top, MBM^\top), \quad (2.3)$$

for any nonsingular matrix M and matrices $A, B \in Sym^+$ (e.g., [16]) making it an ideal candidate to compare covariance matrices. Moreover, it has been used successfully in a variety of applications (e.g., [94, 81, 72, 12, 52, 40]). Notice that the *flat* distance induced by the Frobenius norm does not satisfy the invariance properties (2.3), and has often been shown to be inadequate for comparing covariance matrices [40, 5, 96].

In the most general case of manifolds of Gaussian families parameterized by both the mean and covariance, there seems to be no explicit form for the geodesic distance.

2.2. Approximation of the posterior covariance of the QoI. We first focus on approximating $\Gamma_{\mathbf{Z}|\mathbf{Y}}$. The cost of computing $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ according to (1.4) is dominated by the solution of p linear systems with coefficient matrix Γ_{pos}^{-1} in order to determine $\Gamma_{\text{pos}} \mathcal{O}^\top$. In large-scale inverse problems only the action of the precision matrix Γ_{pos}^{-1} on a vector is usually available; it is not reasonable to expect to have direct factorizations of Γ_{pos}^{-1} , such as Cholesky decomposition. Thus, the solution of linear systems with coefficient matrix Γ_{pos}^{-1} is often necessarily iterative (e.g., via Krylov subspace methods for SPD matrices [50, 45]). Moreover, the storage requirements for $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ scale as $O(p^2)$. If the dimension of the QoI is inherently low, e.g., $p = O(1)$, then the use of direct formulas like (1.4) can be remarkably efficient. For instance, if we are only interested in the average of \mathbf{X} , i.e., $\mathbf{Z} := \frac{1}{n} \sum_i X_i$, then the QoI is only one-dimensional and computing the posterior covariance of the QoI amounts to solving essentially a single linear system. As the dimension of the QoI increases, however, direct formulas like (1.4) become increasingly impractical due to high computational complexity and storage requirements. In many cases of interest, the dimension of the QoI can be arbitrarily large. Consider the following simple example. If \mathbf{X}

represents a finite-dimensional approximation of a spatially distributed stochastic process (e.g., an unknown temperature field), then the QoI could be the restriction of this process to a domain of interest. In this case, the QoI is also a finite-dimensional approximation of a spatially distributed process, and its dimension can be arbitrarily high depending on the chosen level of discretization. (We will revisit this example in Section 4.) There is a clear need for new inference algorithms to efficiently tackle such problems.

2.2.1. Background on optimal low-rank approximations. Even though direct formulas like (1.4) can be intractable when the QoI is high-dimensional, essential features of large-scale Bayesian inverse problems bring additional structure to the Bayesian update: The prior distribution might encode some kind of smoothness or correlation among the inversion parameters. Observations are typically limited in number, indirect, corrupted by noise, and related to the inversion parameters by the action of a forward operator that filters out some information [97, 31]. As a result, data are usually informative, relative to the prior, only about a low-dimensional subspace of the parameter space. That is, the important differences between the prior and posterior distributions are confined to a low-dimensional subspace. This source of low-dimensional structure is key to the development of efficient Bayesian inversion algorithms [39, 31] and also plays a crucial role when dealing with goal-oriented problems.

In [97] we studied the optimal approximation of the posterior covariance of the parameters as a negative definite low-rank update of the prior covariance matrix. Optimality was defined with respect to the Riemannian metric given in Section 2.1, i.e., the geodesic distance on the manifold of SPD matrices [42]. Note that optimality of the covariance approximation in this metric also leads to optimality of the approximation with respect to familiar measures of similarity between probability distributions, such as the Kullback–Leibler divergence and the Hellinger distance [97, 79]. In particular, we focused on the approximation class

$$\mathcal{M}_r = \{\Gamma_{\text{pr}} - KK^\top \succ 0 : \text{rank}(K) \leq r\} \quad (2.4)$$

of positive definite matrices that can be written as a low-rank update of the prior covariance. This class takes advantage of the low-dimensional structure of the prior-to-posterior update.³ The following theorem from [97] characterizes the optimal approximation of Γ_{pos} , and is the launching point for this section. Henceforth whenever we write that (α, v) are eigenpairs of (A, B) we mean that (α, v) are the generalized eigenvalue–eigenvector pairs of the matrix pencil (A, B) .

THEOREM 2.1 (Optimal posterior covariance approximation). *Let (δ_i^2, w_i) be the eigenpairs of $(H, \Gamma_{\text{pr}}^{-1})$ with the ordering $\delta_i^2 \geq \delta_{i+1}^2$, where $H := G^\top \Gamma_{\text{obs}}^{-1} G$ as in (1.2). Then a minimizer $\hat{\Gamma}_{\text{pos}}$ of the Riemannian metric $d_{\mathcal{R}}$ between Γ_{pos} and an element of \mathcal{M}_r is given by*

$$\hat{\Gamma}_{\text{pos}} = \Gamma_{\text{pr}} - KK^\top, \quad KK^\top = \sum_{i=1}^r \frac{\delta_i^2}{1 + \delta_i^2} w_i w_i^\top, \quad (2.5)$$

³Many approximate inference algorithms, especially in the context of Kalman filtering, exploit the class (2.4) to deliver an approximation of Γ_{pos} (e.g., [7, 8, 95]). These algorithms, however, are suboptimal in the sense defined by the forthcoming Theorem 2.1. See [97] for further details and numerical examples.

where the distance between Γ_{pos} and the optimal approximation is

$$d_{\mathcal{R}}^2(\Gamma_{\text{pos}}, \hat{\Gamma}_{\text{pos}}) = \frac{1}{2} \sum_{i>r} \ln^2(1 + \delta_i^2). \quad (2.6)$$

Theorem 2.1 shows that the optimal way to update the prior covariance matrix to obtain an approximation of Γ_{pos} is along the eigenvectors of $(H, \Gamma_{\text{pr}}^{-1})$.⁴ These eigenvectors are the directions most informed by the data, and are obtained from a balance between the forward model, the measurement noise, and prior information. This update is typically low-rank for precisely the reasons discussed above: the data are informative relative to the prior only about a low-dimensional subspace of the parameter space [20, 91]. Notice that (2.5) is not only an optimal and structure-exploiting approximation of Γ_{pos} , but it is also computationally efficient, as the dominant eigenpairs of $(H, \Gamma_{\text{pr}}^{-1})$ can easily be computed using matrix-free algorithms like a Lanczos iteration (including its block version) [60, 76, 32, 19, 57] or a randomized SVD [48, 62]. This approximation of Γ_{pos} , originally introduced in [39] for computational convenience and justified by intuitive arguments, has been employed successfully in many large-scale applications of Bayesian inversion [19, 69, 21, 31, 30, 82]. It is the starting point for our analysis of goal-oriented inverse problems.

2.2.2. A naïve approximation. In this section we introduce an intuitive but suboptimal approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$, which will provide insight and motivation for the structure of the forthcoming optimal approximation. The reader interested exclusively in the optimal approximation may jump directly to Section 2.2.3.

The combination of Theorem 2.1 with the direct formulas (1.4) suggests a first approximation strategy for the posterior covariance of the QoI: just replace Γ_{pos} in (1.4) with the optimal approximation described by Theorem 2.1,

$$\Gamma_{\mathbf{Z}|\mathbf{Y}} \approx \hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}} := \mathcal{O} \hat{\Gamma}_{\text{pos}} \mathcal{O}^\top = \mathcal{O} \Gamma_{\text{pr}} \mathcal{O}^\top - \mathcal{O} K K^\top \mathcal{O}^\top, \quad (2.7)$$

where the low-rank update $K K^\top$ is given by (2.5). Approximation (2.7) is already a major computational improvement over the direct formulas (1.4). There is no need to solve p linear systems; rather, we only need to compute the leading eigenpairs of $(H, \Gamma_{\text{pr}}^{-1})$ with a matrix-free algorithm. The rank of the update depends on the dimension of the parameter subspace that is most informed by the data.

Despite these favorable computational properties, the approximation (2.7) is still not satisfactory as it does not explicitly account for the goal-oriented feature of the problem: $\hat{\Gamma}_{\text{pos}}$ in (2.7) is the optimal approximation of the posterior covariance of the parameters, but is by no means tailored to the QoI. The pencil $(H, \Gamma_{\text{pr}}^{-1})$ used to compute the approximation $\hat{\Gamma}_{\text{pos}}$ does not include the goal-oriented operator \mathcal{O} . In other words, the directions (w_i) defining the optimal prior-to-posterior update in (2.5), though strongly data-informed, need not be relevant to the QoI. For instance, some of the (w_i) could lie in the nullspace of the goal-oriented operator. Computing these eigenvectors would be an unnecessary waste of

⁴The properties of the pencil $(H, \Gamma_{\text{pr}}^{-1})$ have been studied extensively in the literature on classical regularization techniques for linear inverse problems (e.g., [49, 37, 22, 99, 77]). These papers, however, do not adopt a statistical approach to inversion and thus have not considered the optimal approximation of the posterior covariance matrix.

computational resources. Of course, as the rank of the optimal prior-to-posterior update increases, the corresponding approximation $\hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ will continue to improve until eventually $\Gamma_{\mathbf{Z}|\mathbf{Y}} = \hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$. In the worst case scenario, however, $\hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ will be a good approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ only as we start computing eigenpairs of $(H, \Gamma_{\text{pr}}^{-1})$ associated with the smallest nonzero generalized eigenvalues. This is clearly unacceptable as the overall complexity of the approximation algorithm would not depend on the nature of the goal-oriented operator. Therefore, the approximation (2.7) cannot satisfy any reasonable optimality statement in the spirit of Theorem 2.1 and calls for a proper modification.

2.2.3. An optimal approximation. The form of $\hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ in (2.7) shows that the posterior covariance of the QoI can be written as a low-rank update of the prior covariance. (Recall that the prior distribution of the QoI is Gaussian, $\mathbf{Z} \sim \mathcal{N}(0, \Gamma_{\mathbf{Z}})$ with $\Gamma_{\mathbf{Z}} = \mathcal{O} \Gamma_{\text{pr}} \mathcal{O}^\top$.) This is once again consistent with our intuition about the Bayesian update: the data will only inform certain aspects of the QoI. Thus a structure-exploiting approximation class for $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ is given by the set of positive definite matrices that can be written as rank- r negative-definite updates of $\Gamma_{\mathbf{Z}}$:

$$\mathcal{M}_r^{\mathbf{Z}} = \{\Gamma_{\mathbf{Z}} - KK^\top \succ 0 : \text{rank}(K) \leq r\}. \quad (2.8)$$

Before introducing one of the main results of this paper, we observe that \mathbf{Y} and \mathbf{Z} are related by a linear model similar to (1.4). The following lemma clarifies this relationship. (See Appendix B for a proof.)

LEMMA 2.2. *A linear Gaussian model consistent with (1.4) is given by:*

$$\mathbf{Y} = G \mathcal{O}_\dagger \mathbf{Z} + \mathbf{\Delta}, \quad (2.9)$$

where $\mathcal{O}_\dagger := \Gamma_{\text{pr}} \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1}$, $\mathbf{Z} \sim \mathcal{N}(0, \Gamma_{\mathbf{Z}})$ and $\mathbf{\Delta} \sim \mathcal{N}(0, \Gamma_{\mathbf{\Delta}})$ are independent, and $\Gamma_{\mathbf{\Delta}} := \Gamma_{\text{obs}} + G(\Gamma_{\text{pr}} - \Gamma_{\text{pr}} \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1} \mathcal{O} \Gamma_{\text{pr}})G^\top$.

The results of [97, Theorem 2.3] can be extended to the goal-oriented case by applying them to the reduced linear model in (2.9). The following theorem defines the optimal approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ and is one of the main results of this paper. See Appendix B for a proof.

THEOREM 2.3 (Optimal approximation of the posterior covariance of the QoI). *Let (λ_i, q_i) be the eigenpairs of:*

$$(G \Gamma_{\text{pr}} \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1} \mathcal{O} \Gamma_{\text{pr}} G^\top, \Gamma_{\mathbf{Y}}) \quad (2.10)$$

with the ordering $\lambda_i \geq \lambda_{i+1} > 0$ and normalization $q_i^\top G \Gamma_{\text{pr}} \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1} \mathcal{O} \Gamma_{\text{pr}} G^\top q_i = 1$, where $\Gamma_{\mathbf{Y}} := \Gamma_{\text{obs}} + G \Gamma_{\text{pr}} G^\top$ is the covariance matrix of the marginal distribution of \mathbf{Y} . Then, a minimizer $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ of the Riemannian metric $d_{\mathcal{R}}$ between $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ and an element of $\mathcal{M}_r^{\mathbf{Z}}$ is given by:

$$\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \Gamma_{\mathbf{Z}} - KK^\top, \quad KK^\top = \sum_{i=1}^r \lambda_i \hat{q}_i \hat{q}_i^\top, \quad \hat{q}_i := \mathcal{O} \Gamma_{\text{pr}} G^\top q_i, \quad (2.11)$$

where the corresponding minimum distance is:

$$d_{\mathcal{R}}^2(\Gamma_{\mathbf{Z}|\mathbf{Y}}, \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}) = \frac{1}{2} \sum_{i>r} \ln^2(1 - \lambda_i). \quad (2.12)$$

The optimal approximation in Theorem 2.3 yields the best approximation for any given rank of the prior-to-posterior update and, most importantly, never requires the full posterior covariance of the parameters. (This should be contrasted with [63].) The directions (q_i) that define the optimal update are the leading eigenvectors of $(G\Gamma_{\text{pr}}\mathcal{O}^\top\Gamma_{\mathbf{Z}}^{-1}\mathcal{O}\Gamma_{\text{pr}}G^\top, \Gamma_{\mathbf{Y}})$ and stem from a careful balance of all the ingredients of the goal-oriented inverse problem: the forward model, measurement noise, prior information, and ultimate goals. Incorporating ultimate goals reduces the intrinsic dimensionality of the inverse problem: for any fixed approximation error, the rank of the optimal update (2.11) can only be less than or equal to that of the suboptimal approximation introduced in (2.7).

2.2.4. Computational remarks. If square roots of Γ_{pr} and Γ_{obs} are available, such that $\Gamma_{\text{pr}} = S_{\text{pr}}S_{\text{pr}}^\top$ and $\Gamma_{\text{obs}} = S_{\text{obs}}S_{\text{obs}}^\top$, then we can rewrite the pencil (2.10) in a more concise form as follows.

COROLLARY 2.4. *Let $\hat{G} := S_{\text{obs}}^{-1}GS_{\text{pr}}$, and let Π be an orthogonal projector onto the range of $S_{\text{pr}}^\top\mathcal{O}^\top$. Then the eigenvalues of*

$$(\hat{G}\Pi\hat{G}^\top, I + \hat{G}\hat{G}^\top) \quad (2.13)$$

are the same as those of (2.10), and the eigenvectors of (2.13) can be mapped to the eigenvectors of (2.10) with the transformation $w \mapsto S_{\text{obs}}^{-\top}w$.

The proof of the corollary is straightforward once we note that Π can be written as $\Pi = S_{\text{pr}}^\top\mathcal{O}^\top(\mathcal{O}S_{\text{pr}}S_{\text{pr}}^\top\mathcal{O}^\top)^{-1}\mathcal{O}S_{\text{pr}} = S_{\text{pr}}^\top\mathcal{O}^\top\Gamma_{\mathbf{Z}}^{-1}\mathcal{O}S_{\text{pr}}$. Moreover, the action of the projector Π on a vector v can be computed efficiently since $\Pi(v) := S_{\text{pr}}^\top\mathcal{O}^\top x_{\text{ls}}$, where x_{ls} is the least squares solution of the overdetermined linear system $S_{\text{pr}}^\top\mathcal{O}^\top x_{\text{ls}} = v$. There is a variety of techniques for the solution of large-scale matrix-free least squares problems (e.g., [78, 41, 23, 50, 71]).

We now focus our computational remarks on the analysis of the pencil in (2.13), which is well suited for practical implementations of the approximation. To simplify notation, let us rewrite (2.13) as (A, B) , where $A := \hat{G}\Pi\hat{G}^\top$ and $B := I + \hat{G}\hat{G}^\top$.

Finding the leading generalized eigenpairs of (2.13) requires the solution of a Hermitian generalized eigenvalue problem [10]. Unfortunately, it is not easy to reduce (2.13) to a standard eigenvalue problem,⁵ as doing so would require the action of a square root of B or of B^{-1} . Nevertheless, there are a plethora of matrix-free algorithms for large-scale generalized eigenvalue problems: generalized Lanczos iteration [10, Section 5.5], randomized SVD-type methods [88], manifold optimization algorithms [1, 2, 11], the trace minimization algorithm [89, 58], and the inverse-free preconditioned Krylov subspace method [46], to name a few. These algorithms require the iterative solution of linear systems associated with B , in some cases to low accuracy [89, 46]. Applying B to a vector requires the evaluation of the forward model, which may or may not be expensive (e.g., consider PDE based models [39] versus image deblurring problems [56]). In practice, for a *fixed* dimension of the desired eigenspace, algorithms for characterizing the eigenpairs of (A, B) lead to more expensive computations than those for the pencil $(H, \Gamma_{\text{pr}}^{-1})$ used in the naïve approximation of Section 2.2.2. However, the key point is that the optimal approximation of Theorem 2.3 requires the characterization of lower dimensional eigenspaces for a given accuracy.

⁵In contrast, this is often possible in the non-goal-oriented case when dealing with the pencil $(H, \Gamma_{\text{pr}}^{-1})$, as the action of a square root of Γ_{pr}^{-1} , or of its inverse, is available in many cases of interest (e.g., [66]).

Moreover, if we solve the generalized eigenvalue problem using a block Lanczos iteration or a randomized method, then we can also exploit block Krylov methods to solve the associated linear systems—comprised of B and multiple right-hand sides—simultaneously [75, 87]. In particular, the convergence of Krylov methods for solving linear systems of the form $Bx = b$, such as the conjugate gradient algorithm, depends not only on the spectrum of B but also on the right-hand side b [65, 9]. This dependence is especially important when the right-hand side has some structure and is not entirely random: in our case, b lies in the range of the possibly low-rank operator A . For instance, if b is *mostly* contained in a low-dimensional invariant subspace of B (whether associated with small or large eigenvalues), then the Krylov solver will likely converge to an accurate solution in few steps. Conversely, it should be noted that if the range of the operator Π in (2.13)—essentially the subspace of the parameter space that is relevant to the QoI—has non-negligible components along *every* data-informed parameter direction (corresponding to the leading eigenspace of $(H, \Gamma_{\text{pr}}^{-1})$), then it would be difficult to obtain an accurate approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ without exploring the full data-informed subspace.

Even though the naïve approximation is suboptimal, it may be interesting from a practical standpoint to assess its performance, since it is cheaper to compute for a given approximation rank. The following lemma provides useful guidelines in this direction. See Appendix B for a proof of this result.

LEMMA 2.5 (Relationship between approximations). *Let $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}, \hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}} \in \mathcal{M}_r^{\mathbf{Z}}$ be, respectively, the optimal and suboptimal approximations of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ introduced in (2.11) and (2.7). Moreover, let $\hat{\Gamma}_{\text{pos}} \in \mathcal{M}_r$ be the optimal approximation of Γ_{pos} defined in (2.5). Then*

$$d_{\mathcal{R}}(\Gamma_{\mathbf{Z}|\mathbf{Y}}, \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}) \leq d_{\mathcal{R}}(\Gamma_{\mathbf{Z}|\mathbf{Y}}, \hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}) \leq d_{\mathcal{R}}(\Gamma_{\text{pos}}, \hat{\Gamma}_{\text{pos}}). \quad (2.14)$$

Lemma 2.5 has several interesting consequences. First of all, notice that it is possible to bound the accuracy of the naïve approximation, $\hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \mathcal{O} \hat{\Gamma}_{\text{pos}} \mathcal{O}^\top$, using $d_{\mathcal{R}}(\Gamma_{\text{pos}}, \hat{\Gamma}_{\text{pos}})$. The latter distance can easily be bounded as a function of the generalized eigenvalues of $(H, \Gamma_{\text{pr}}^{-1})$, as shown in (2.6). These are precisely the eigenvalues computed by the naïve approximation. Thus, if the eigenvalues of $(H, \Gamma_{\text{pr}}^{-1})$ decay sharply or, equivalently, if the distance $d_{\mathcal{R}}(\Gamma_{\text{pos}}, \hat{\Gamma}_{\text{pos}})$ can be made small with only a low-dimensional (small r) prior-to-posterior update, then Lemma 2.5 says that the naïve approximation, albeit suboptimal, can yield a remarkably efficient approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ —with strong accuracy guarantees in terms of $d_{\mathcal{R}}(\Gamma_{\text{pos}}, \hat{\Gamma}_{\text{pos}})$. Intuitively, if $\hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ already accounts for most of the data-informed directions in the parameter space, then there is no major loss of accuracy in neglecting further directions of the prior-to-posterior update, even if these directions are relevant to the QoI. In this situation, these additional directions would provide very limited information relative to the prior and can be safely neglected.

On the other hand, if the eigenvalues of $(H, \Gamma_{\text{pr}}^{-1})$ do not decay as quickly, then the bound provided in (2.14) becomes useless. This is not to say that the naïve approximation will necessarily perform poorly. (It is possible that the QoI depends *only* on a few of the leading data-informed directions, such that the naïve approximation performs well even for a low rank prior-to-posterior update.) But we cannot quantify the accuracy of the naïve approximation unless we directly compute $d_{\mathcal{R}}(\Gamma_{\mathbf{Z}|\mathbf{Y}}, \hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}})$, which in turn requires the solution of an expensive generalized eigenvalue problem for each rank of the update. This is

not feasible in practice. In such situations, we should resort to the optimal approximation introduced in Theorem 2.3, which offers a useful error bound as well as a concrete possibility for both computational and storage savings.

2.2.5. Properties of the optimal covariance approximation. An important consequence of the optimal approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ with respect to the metric $d_{\mathcal{R}}$ is optimality *in distribution* whenever the posterior mean of the QoI is known. It follows from [97, Lemma 2.2] that the minimizer of the Hellinger distance (or the Kullback–Leibler divergence) between the Gaussian posterior measure of the QoI, $\nu_{\mathbf{Z}|\mathbf{Y}} := \mathcal{N}(\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}), \Gamma_{\mathbf{Z}|\mathbf{Y}})$, and the approximation $\mathcal{N}(\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}), \Gamma)$ for a matrix $\Gamma \in \mathcal{M}_r^{\mathbf{Z}}$, is given by the optimal approximation (2.11) defined in Theorem 2.3. In particular, let $\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}} := \mathcal{N}(\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}), \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}})$ be the measure that optimally approximates $\nu_{\mathbf{Z}|\mathbf{Y}}$, where $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ is defined in (2.11). Then it is easy to show that the Hellinger distance between $\nu_{\mathbf{Z}|\mathbf{Y}}$ and the optimal approximation $\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}$ is given by:

$$d_{\text{Hell}}(\nu_{\mathbf{Z}|\mathbf{Y}}, \tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}) = \sqrt{1 - \prod_{i>r} 2^{1/2} \frac{(1 - \lambda_i)^{1/4}}{(2 - \lambda_i)^{1/2}}} \quad (2.15)$$

where (λ_i) are the generalized eigenvalues defined in Theorem 2.3 (e.g., [97, Appendix A] or [79]).

The Hellinger distance can be used to bound the error of expectations of functions of interest with respect to approximate measures [33]. That is, suppose that we are interested in the posterior expectation $\mathbb{E}_{\nu_{\mathbf{Z}|\mathbf{Y}}}[g]$ of some measurable function $g : \mathbb{R}^p \rightarrow \mathbb{R}$, with certain bounded moments with respect to the prior measure $\nu_{\mathbf{Z}} := \mathcal{N}(0, \Gamma_{\mathbf{Z}})$, and suppose further that we can only evaluate integrals with respect to the approximate measure $\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}$. Then the error resulting from computing $\mathbb{E}_{\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}}[g]$, as opposed to $\mathbb{E}_{\nu_{\mathbf{Z}|\mathbf{Y}}}[g]$, for a fixed realization of the data \mathbf{Y} , can be bounded in terms of the Hellinger distance between the two Gaussian measures using the following lemma, which follows easily from [33, Lemma 7.14]. (See Appendix B for a proof.)

LEMMA 2.6 (Convergence in expectation). *Let $g : \mathbb{R}^p \rightarrow \mathbb{R}$ be a measurable function with $\beta > 2$ bounded moments with respect to the prior measure, i.e., $\mathbb{E}_{\nu_{\mathbf{Z}}} [|g|^\beta] < \infty$. Then:*

$$\left| \mathbb{E}_{\nu_{\mathbf{Z}|\mathbf{Y}}}[g] - \mathbb{E}_{\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}}[g] \right| \leq \mathcal{C}(\mathbf{Y}, g) d_{\text{Hell}}(\nu_{\mathbf{Z}|\mathbf{Y}}, \tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}), \quad (2.16)$$

where $\mathcal{C}(\mathbf{Y}, g) := 2\sqrt{2} \frac{|\Gamma_{\mathbf{Z}}|^{1/4}}{|\Gamma_{\mathbf{Z}|\mathbf{Y}}|^{1/4}} \exp\left(\frac{1}{2(\beta-2)} \|\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})\|_{\Gamma_{\mathbf{Z}}^{-1}}^2\right) \mathbb{E}_{\nu_{\mathbf{Z}}} [|g|^\beta]^{1/\beta}$ and where $|A|$ denotes the determinant of the matrix A .

Notice that the constant $\mathcal{C}(\mathbf{Y}, g)$ in (2.16) is independent of the approximating measure $\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}$. Convergence of the approximation in Hellinger distance thus implies convergence of the expectation $\mathbb{E}_{\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}}[g]$ to $\mathbb{E}_{\nu_{\mathbf{Z}|\mathbf{Y}}}[g]$.

It is interesting to note that the optimal approximation of the posterior covariance matrix of the QoI in Theorem 2.3 is always associated with a corresponding approximation of the posterior covariance of the parameters, Γ_{pos} . The following result clarifies the nature of this approximation. It is proved in Appendix B.

LEMMA 2.7 (Goal-oriented approximation of Γ_{pos}). *Let $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ be the minimizer of the metric $d_{\mathcal{R}}$ between $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ and an element of $\mathcal{M}_r^{\mathbf{Z}}$ as given by (2.11). Then $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ can be*

written as

$$\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \mathcal{O} \hat{\Gamma}_{\text{pos}}^* \mathcal{O}^\top, \quad \hat{\Gamma}_{\text{pos}}^* = \Gamma_{\text{pr}} - \sum_{i=1}^r \lambda_i \tilde{q}_i \tilde{q}_i^\top, \quad \tilde{q}_i := S_{\text{pr}} \Pi S_{\text{pr}}^\top G^\top q_i, \quad (2.17)$$

where the vectors (q_i) are defined in Theorem 2.3, S_{pr} is a square root of the prior covariance matrix such that $\Gamma_{\text{pr}} = S_{\text{pr}} S_{\text{pr}}^\top$, Π is the orthogonal projector onto the range of $S_{\text{pr}}^\top \mathcal{O}^\top$, while $\hat{\Gamma}_{\text{pos}}^*$ satisfies

$$\begin{aligned} \hat{\Gamma}_{\text{pos}}^* &\in \arg \min_{\Gamma} d_{\mathcal{R}}(\Gamma_{\mathbf{Z}|\mathbf{Y}}, \mathcal{O} \Gamma \mathcal{O}^\top) \\ \text{s.t. } \Gamma &\in \mathcal{M}_r := \{\Gamma = \Gamma_{\text{pr}} - K K^\top \succ 0, \text{rank}(K) \leq r\}. \end{aligned} \quad (2.18)$$

The matrix $\hat{\Gamma}_{\text{pos}}^*$ is an optimal *goal-oriented* approximation of Γ_{pos} . This notion of optimality is quite different from that in Theorem 2.1. The prior-to-posterior update directions, (\tilde{q}_i) in (2.17), have the intuitive interpretation of directions, in the parameter space, that are most informed by the data, relative to the prior, *and* that are relevant to the QoI. In particular, it is easy to see that the (\tilde{q}_i) are orthogonal to the nullspace of the goal-oriented operator with respect to the inner product induced by the prior precision, i.e.,

$$h^\top \Gamma_{\text{pr}}^{-1} \tilde{q}_i = (\mathcal{O} h)^\top \Gamma_{\mathbf{Z}|\mathbf{Y}}^{-1} \mathcal{O} \Gamma_{\text{pr}} G^\top q_i = 0, \quad \forall h \in \text{Null}(\mathcal{O}). \quad (2.19)$$

Note that even if $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \mathcal{O} \hat{\Gamma}_{\text{pos}}^* \mathcal{O}^\top$ is a good approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$, $\hat{\Gamma}_{\text{pos}}^*$ need not be a good approximation of Γ_{pos} .

Now we introduce a particularly simple factorization of the optimal approximation $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ from Theorem 2.3, as $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \tilde{S}_{\mathbf{Z}|\mathbf{Y}} \tilde{S}_{\mathbf{Z}|\mathbf{Y}}^\top$ for some matrix $\tilde{S}_{\mathbf{Z}|\mathbf{Y}}$. We can think of $\tilde{S}_{\mathbf{Z}|\mathbf{Y}}$ as a square root of $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$, even though $\tilde{S}_{\mathbf{Z}|\mathbf{Y}}$ need not be a square matrix. Obtaining the action of a square root of $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ on a vector is an essential task if our goal is to sample the distribution $\mathcal{N}(\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}), \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}})$ in truly high-dimensional problems. The key requirement is that $\tilde{S}_{\mathbf{Z}|\mathbf{Y}}$ be easy to compute once we have the optimal approximation $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$. We have deferred the discussion of this topic until now in order to exploit the results of Lemma 2.7 to obtain an explicit characterization of $\tilde{S}_{\mathbf{Z}|\mathbf{Y}}$. The proof of the following lemma can be found in Appendix B.

LEMMA 2.8. *Let (λ_i, q_i) , S_{pr} , and Π be defined as in Lemma 2.7. Then, a non-symmetric square root, $\tilde{S}_{\mathbf{Z}|\mathbf{Y}}$, of $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$, such that $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \tilde{S}_{\mathbf{Z}|\mathbf{Y}} \tilde{S}_{\mathbf{Z}|\mathbf{Y}}^\top$, is given by*

$$\tilde{S}_{\mathbf{Z}|\mathbf{Y}} = \mathcal{O} S_{\text{pr}} \left(\sum_{i=1}^r (\sqrt{1 - \lambda_i} - 1) \bar{q}_i \bar{q}_i^\top + I \right), \quad \bar{q}_i := \Pi S_{\text{pr}}^\top G^\top q_i, \quad (2.20)$$

where I is the identity matrix.

The virtue of this result is that it does not require an invertible square root of $\Gamma_{\mathbf{Z}} = \mathcal{O} \Gamma_{\text{pr}} \mathcal{O}^\top$ as one would expect from [97, Remark 2]. Note that it is easy to apply $\tilde{S}_{\mathbf{Z}|\mathbf{Y}}$ to a vector, which allows efficient sampling from $\mathcal{N}(\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}), \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}})$. An interesting feature of $\tilde{S}_{\mathbf{Z}|\mathbf{Y}} \in \mathbb{R}^{p \times n}$ is that it is a nonsquare matrix with $p < n$. This is certainly not an issue as long as $\tilde{S}_{\mathbf{Z}|\mathbf{Y}} \tilde{S}_{\mathbf{Z}|\mathbf{Y}}^\top = \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$. Notice also that (2.20) contains a square root of

the prior covariance matrix. The action of this matrix is usually available in large-scale applications (e.g., [100, 35, 66, 98, 101]). However, if the action of a square root of Γ_{pr} is truly unavailable, then one can still sample from $\mathcal{N}(\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}), \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}})$ by resorting to the action of the matrix $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ alone (e.g., [24, 43, 80, 92]). It is straightforward to apply $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$ to a vector (see (2.11)).

2.3. Approximation of the posterior mean of the QoI. We conclude this theory section by introducing an optimal approximation of the posterior mean of the QoI. The cost of computing

$$\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) := \mathcal{O} \mu_{\text{pos}}(\mathbf{Y}) = \mathcal{O} \Gamma_{\text{pos}} G^\top \Gamma_{\text{obs}}^{-1} \mathbf{Y} \quad (2.21)$$

for a single realization of the data is usually dominated by the cost of solving a single linear system associated with Γ_{pos}^{-1} to determine $\mu_{\text{pos}}(\mathbf{Y})$. This task can be efficiently tackled with state-of-the-art matrix-free iterative solvers for symmetric linear systems (e.g., [10, 50, 3]) even for million-dimensional parameter spaces [19]. If, however, one is interested in the fast computation of $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ for multiple realizations of the data (e.g., in the context of online inference), then the situation is quite different [26, 44, 68, 53, 29, 25]. Solving a linear system to compute $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ each time a new measurement is available might be infeasible in practical applications. If the dimension of the QoI is small, say $p = \mathcal{O}(1)$, then there is an easy solution to this problem. One can just precompute the matrix $M := \mathcal{O} \Gamma_{\text{pos}} G^\top \Gamma_{\text{obs}}^{-1}$ in an offline stage and then compute the posterior mean of the QoI as $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) = M \mathbf{Y}$ each time a new realization of the data becomes available. Yet the computational efficiency of this procedure breaks down as the dimension of the QoI increases—for instance, if the QoI is a finite-dimensional approximation of some underlying function. In this case, the matrix M would be large and dense, and storing it could be quite inefficient. Moreover, performing a dense matrix-vector product to compute $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) = M \mathbf{Y}$ might become more expensive than solving a single linear system associated with Γ_{pos}^{-1} .

Our goal is thus to characterize computationally efficient and statistically optimal approximations of $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$. In particular, we seek an approximation of $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ as a low-rank linear function of the data, i.e., $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) \approx \tilde{\mu}_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) := A \mathbf{Y}$ for some low-rank matrix A . With such an A , computing $\tilde{\mu}_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ for each new realization of the data would be computationally efficient. We define optimality of the approximation with respect to the Bayes risk for squared-error loss weighted by the posterior precision matrix of the QoI, i.e.,

$$\mathcal{B}(A) := \mathbb{E} \left[\|A \mathbf{Y} - \mathbf{Z}\|_{\Gamma_{\mathbf{Z}|\mathbf{Y}}^{-1}}^2 \right], \quad (2.22)$$

where $\mathcal{B}(A)$ denotes the Bayes risk associated with the matrix A , and where the expectation is taken over the joint distribution of \mathbf{Z} and \mathbf{Y} . (The minimizer of the Bayes risk for squared error loss over all linear functions of the data is precisely the posterior mean of the QoI.) Minimizing the Bayes risk (2.22) is equivalent to minimizing

$$\mathbb{E} \left[\|\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) - \tilde{\mu}_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})\|_{\Gamma_{\mathbf{Z}|\mathbf{Y}}^{-1}}^2 \right] \quad (2.23)$$

over all approximations of the posterior mean of the form $\tilde{\mu}_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) = A \mathbf{Y}$ for some low-rank matrix A . The Mahalanobis distance in (2.23) is precisely a Riemannian metric of the form described in Section 2.1 and thus it is a natural way to assess the quality of a posterior

mean approximation. In particular, the weighted norm in (2.23) penalizes errors in the approximation of $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ more strongly in directions of lower posterior variance. As a result, the approximation of $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ is more likely to fall within the bulk of the posterior density of the QoI. Notice that (2.23) is an average of the squared Riemannian distance between $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ and its approximation $\tilde{\mu}_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ over the distribution of the data \mathbf{Y} .

The following theorem characterizes the optimal approximation of $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$. See Appendix B for a proof.

THEOREM 2.9 (Optimal approximation of $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$). *Let $(\lambda_i, q_i, \hat{q}_i)$ be defined as in Theorem 2.3 and consider the minimization of the following Bayes risk over the set of low-rank matrices:*

$$\min_A \mathbb{E} \left[\|A\mathbf{Y} - \mathbf{Z}\|_{\Gamma_{\mathbf{Z}|\mathbf{Y}}^{-1}}^2 \right], \quad \text{s.t. } \text{rank}(A) \leq r. \quad (2.24)$$

Then a minimizer of (2.24) is given by:

$$A^* = \sum_{i=1}^r \lambda_i \hat{q}_i q_i^\top, \quad (2.25)$$

with minimum Bayes risk:

$$\mathcal{B}(A^*) = \mathbb{E} \left[\|A^*\mathbf{Y} - \mathbf{Z}\|_{\Gamma_{\mathbf{Z}|\mathbf{Y}}^{-1}}^2 \right] = \sum_{i>r} \frac{\lambda_i}{1 - \lambda_i} + n, \quad (2.26)$$

where n is the dimension of the parameter space.

Note that (2.25) can be computed “for free” from the optimal approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ introduced in Theorem 2.3. Also, the optimal approximations of both the posterior mean and the posterior covariance of the QoI become quite accurate as soon as we start including generalized eigenvalues $\lambda \ll 1$ in the corresponding approximations (see minimum loss (2.12) and Bayes risk (2.26)).

3. Proof-of-concept example. Before investigating the numerical performance of our goal-oriented approximations, we illustrate the theory with a simple proof-of-concept example. We consider an identity forward model $G = I$, a diagonal observational noise precision $\Gamma_{\text{obs}}^{-1} = \text{diag}(h_1, \dots, h_n)$, and a diagonal prior covariance $\Gamma_{\text{pr}} = \text{diag}(\mu_1, \dots, \mu_n)$, with $h_i = n - i$ and $\mu_i = i$ for $i = 1, \dots, n$. We may think of this problem as denoising a signal \mathbf{X} [97]. Figure 3.1 shows the normalized eigenvalues of Γ_{obs}^{-1} and Γ_{pr} in blue and red, respectively, for the case $n = 30$. The eigenvectors of both matrices correspond to the canonical vectors in \mathbb{R}^n , i.e., $\mathbf{e}_1, \dots, \mathbf{e}_n$. In this case, the data are most informative—in absolute terms—along directions \mathbf{e}_i with $i \ll n$, since the observational noise precision h_i is a decreasing function of i . On the other hand, the prior variance is large along \mathbf{e}_i when $i \gg 1$, since μ_i is an increasing function of i . Thus the prior is more constraining where the data are more informative. The eigenpairs (δ_i^2, w_i) of the pencil $(H, \Gamma_{\text{pr}}^{-1})$, defined in Theorem 2.1, are given by $\delta_i^2 = h_i \cdot \mu_i = (n - i) \cdot i$ and $w_i \propto \mathbf{e}_i$ for $i = 1, \dots, n$. (These δ_i^2 are not sorted in decreasing order; for simplicity, we retain the same index i as in the problem definition.) From the relative magnitudes of (δ_i^2) —illustrated by the green parabola in Figure 3.1—we can identify the parameter directions that are most informed by the data

relative to the prior: they correspond to \mathbf{e}_i with i around $n/2$ (the middle of the spectrum). These directions define the optimal prior-to-posterior update of Theorem 2.1. Modes \mathbf{e}_i with $i \ll n/2$ are strongly informed by the data in an absolute sense, but not *relative* to the prior; thus their overall importance is limited. In the same way, modes \mathbf{e}_i with $i \gg n/2$ are unimportant to the update since the posterior variance along these directions is roughly equal to the prior variance ($\delta_i^2 \ll 1$), even though both variances are relatively large [97].

Now let the goal-oriented operator $\mathcal{O} : \mathbb{R}^n \rightarrow \mathbb{R}^p$ be defined as follows: $\mathcal{O}\mathbf{x} = (x_1, \dots, x_p)$ for $\mathbf{x} = (x_1, \dots, x_n)$ and $p = n/2$. Simple algebra shows that the goal-oriented eigenpairs (λ_i, q_i) of Theorem 2.3 are given by $q_i \propto \mathbf{e}_i$ for $i = 1, \dots, n$, and $\lambda_i = 1/(1 + 1/(h_i \mu_i))$ for $i \leq p$ and $\lambda_i = 0$ for $i > p$. So that the eigenvalues δ_i^2 and λ_i are comparable in terms of their associated covariance approximation errors—see (2.6) and (2.12)—we plot a nonlinear function of each λ_i in Figure 3.1, namely $\hat{\lambda}_i = f(\lambda_i)$ for $f(x) = 1/(1-x)$. (Since f is strictly increasing on $[0, 1)$, the relative importance of the $(\hat{\lambda}_i)$ is the same as that of the original (λ_i) .) The introduction of a goal-oriented operator reveals directions that can be strongly informed by the data, relative to the prior, but that are irrelevant to the QoI. These modes correspond to (\mathbf{e}_i) for $i > p$, and can be safely neglected when computing the Bayesian update relevant to the QoI.

Of course, in the general case of non-diagonal operators $(G, \Gamma_{\text{obs}}, \Gamma_{\text{pr}})$, the directions (w_i) and (q_i) need not coincide. The following numerical example will illustrate this general situation.




FIG. 3.1. Normalized eigenvalues defined in the proof-of-concept example of Section 3: in blue we show (h_i/h_{\max}) , in red (μ_i/μ_{\max}) , in green $(\delta_i^2/\delta_{\max}^2)$, and in magenta $(\hat{\lambda}_i/\hat{\lambda}_{\max})$, for $i = 1, \dots, n$ and $n = 30$. For any finite collection of eigenvalues (σ_i) , we define σ_{\max} to be the maximum value over that collection. Since for $i \leq p$ and $p = 15$, we have $\hat{\lambda}_i = \delta_i^2$ in this example, we shifted the magenta curve slightly upwards to distinguish it from the green one.

4. Numerical examples. Now we numerically illustrate the performance of our approximations using a goal-oriented inverse problem in heat transfer. In particular, we study the cooling of a CPU by means of a heat sink. Our goal is to infer the (spatially inhomogeneous) temperature of the CPU from noisy pointwise observations of temperature on the heat sink. Figure 4.1 shows the problem setup: the three different layers of material correspond, respectively, to the CPU (\mathcal{D}_1), a thin silicone layer that connects the CPU to the heat sink (\mathcal{D}_2), and an aluminum fin (\mathcal{D}_3). We denote by \mathcal{D} the union of these domains. Each \mathcal{D}_i represents a two-dimensional cross section of material of constant width W along

the horizontal direction and a height L_i . We assume that no heat transfer happens along the third dimension; this is a common engineering approximation [13]. Each material has a constant density ρ_i , a constant specific heat c_i and a constant thermal conductivity k_i . The corresponding thermal diffusivities $\alpha_i = k_i/\rho_i c_i$ are shown in the table at the right of Figure 4.1. The time-dependent temperature field in each domain is $\Theta^{(i)} : \mathcal{D}_i \times T \rightarrow \mathbb{R}$, where $T = (0, t_{\text{end}}]$, for $i = 1, 2, 3$. Jointly, these temperature fields are simply $\Theta : \mathcal{D} \times T \rightarrow \mathbb{R}$.

4.1. Forward, observational and prior models. The time evolution of each temperature field $\Theta^{(i)}$ is described by a linear time-dependent PDE of the form

$$\rho_i c_i \partial_t \Theta^{(i)} = \nabla \cdot (k_i \nabla \Theta^{(i)}), \quad i = 1, \dots, 3, \quad (4.1)$$

where ∂_t denotes partial differentiation with respect to time. We assume no volumetric heat production and use Fourier's law for the heat flux [47]. Equations 4.1 should be complemented with appropriate boundary and initial conditions to define a well-posed forward problem. We use the independent variables s_1 and s_2 to denote, respectively, the horizontal and vertical directions and let $\mathbf{s} = (s_1, s_2)$. The point $\mathbf{s} = (0, 0)$ corresponds to the lower left corner of \mathcal{D} . At the lower boundary of \mathcal{D}_1 we impose a space- and time-dependent heat flux: $k_1 \partial_{\vec{n}} \Theta^{(1)} = q(\mathbf{s}, t)$ for $\mathbf{s} \in \mathcal{D}_{1, \text{bottom}}$, where \vec{n} refers to the outward pointing normal and q is a given nonconstant scalar function in \mathbf{s} . At the interface between domains \mathcal{D}_i and \mathcal{D}_{i+1} we assume heat transfer by conduction with no thermal contact resistance: $k_i \partial_{\vec{n}} \Theta^{(i)} = k_{i+1} \partial_{\vec{n}} \Theta^{(i+1)}$ and $\Theta^{(i)} = \Theta^{(i+1)}$ for $\mathbf{s} \in \text{interface}(\mathcal{D}_i, \mathcal{D}_{i+1})$ and $i = 1, 2$. At the top, left, and right boundaries of \mathcal{D}_3 , we assume heat transfer by convection with a fluid at constant temperature Θ_∞ : $-k_3 \partial_{\vec{n}} \Theta^{(3)} = h_c(\Theta^{(3)} - \Theta_\infty)$ for $\mathbf{s} \in \mathcal{D}_{3, \text{top}} \cup \mathcal{D}_{3, \text{left}} \cup \mathcal{D}_{3, \text{right}}$, where h_c is a convective heat transfer coefficient. Finally, we impose adiabatic conditions (no heat exchange) on the left and right boundaries of \mathcal{D}_1 and \mathcal{D}_2 : $\partial_{\vec{n}} \Theta^{(i)} = 0$ for $\mathbf{s} \in \mathcal{D}_{i, \text{left}} \cup \mathcal{D}_{i, \text{right}}$ and $i = 1, 2$. The initial conditions are not specified here as they are the objective of the forthcoming inverse problem.

We consider a finite element spatial approximation of the weak form of (4.1) by means of linear elements on simplices [83]. We denote by $\Theta_h(t) \in \mathbb{R}^n$ the collection of temperature values at the finite element nodes on \mathcal{D} at time $t \in T$. The function Θ_h satisfies a system of ODEs of the form $M \partial_t \Theta_h(t) + A \Theta_h(t) = \mathbf{f}(t)$, with $t \in T$, for a suitable mass matrix M , stiffness matrix A , known time-dependent forcing term \mathbf{f} and initial conditions $\Theta_{0h} := \Theta_h(t = 0)$.

The initial conditions Θ_{0h} are unknown and must be estimated from local measurements of the temperature field Θ at different locations in space and time. The locations of the sensors s^1, \dots, s^N are shown as black dots in Figure 4.1. Observations are collected every Δt time units for $t \in T$. The first observation happens at time $t = \Delta t$ and we assume that there are M observation times in total. We denote measurements at time $t_i = i\Delta t$ as $\widehat{\mathbf{Y}}_i = [\Theta(s^1, i\Delta t), \dots, \Theta(s^N, i\Delta t)]$. We concatenate the observations into a vector $\widehat{\mathbf{Y}} = (\widehat{\mathbf{Y}}_1, \dots, \widehat{\mathbf{Y}}_M) \in \mathbb{R}^d$. The actual observations are corrupted with additive Gaussian noise: $\mathbf{Y} = \widehat{\mathbf{Y}} + \mathcal{E}$, where $\mathcal{E} \sim \mathcal{N}(0, \sigma_{\text{obs}}^2 I)$ and I is the identity matrix. Notice that $\widehat{\mathbf{Y}}$ is an affine function of Θ_{0h} . This relationship can be made linear by a suitable redefinition of the data vector. Thus, we are lead to a linear Gaussian inverse problem in standard form, $\mathbf{Y} = G \Theta_{0h} + \mathcal{E}$, where G defines the forward operator, $\Theta_{0h} \mapsto \widehat{\mathbf{Y}}$, that can be evaluated implicitly by solving a heat equation with no forcing term and initial conditions Θ_{0h} for a

time interval necessary to collect the corresponding observations $\widehat{\mathbf{Y}}$.

We define a zero-mean Gaussian prior distribution⁶ on Θ_{0h} by modeling Θ_{0h} as a discretized solution of a stochastic PDE of the form

$$\gamma \left(\kappa^2 \mathcal{I} - \Delta \right) \Theta(\mathbf{s}) = \mathcal{W}(\mathbf{s}), \quad \mathbf{s} \in \mathcal{D}, \quad (4.2)$$

where \mathcal{W} is a white noise process, κ is a positive scalar parameter, Δ is the Laplacian operator and \mathcal{I} is the identity operator. In particular, we exploit the explicit link between Gaussian Markov random fields with the Matérn covariance function and solutions to stochastic PDEs as outlined in [66]. In this case, the action of a square root of the prior covariance matrix on a vector is readily available as the solution of an elliptic PDE on \mathcal{D} , and thus it is scalable to very large inverse problems [66].

In this example we use $h_c = 23.8 \text{ W/m}^2 \text{ K}$ for the convective heat transfer coefficient between the aluminum fin and the external fluid (air), which has constant temperature $\Theta_\infty = 283 \text{ K}$. The width of the domain \mathcal{D} in Figure 4.1 is $H = 2 \times 10^{-2} \text{ m}$. The heat flux $q(\mathbf{s}, t)$ is time-independent and nonnegative and can be written as the superposition of two square impulse functions with zero background: one centered at $6 \times 10^{-3} \text{ m}$ with width $8 \times 10^{-3} \text{ m}$ and intensity 0.6 W/m^2 ; and the other centered at $15 \times 10^{-3} \text{ m}$ with width $4 \times 10^{-3} \text{ m}$ and intensity 0.3 W/m^2 . Observations are collected every $\Delta t = 5 \times 10^{-4} \text{ s}$ for a total of $M = 100$ measurements. We use $\sigma_{\text{obs}} = 1/2$ as the standard deviation of the observational noise. The prior parameters in (4.2) are given by $\gamma = 1 \times 10^4$ and $\kappa = \sqrt{8}/\rho_{\text{pr}}$ with $\rho_{\text{pr}} = H/10$. This choice of κ defines a prior with correlation values near $1/10$ at distance ρ_{pr} [66]. The original prior mean is set to $\mu_{\text{pr}} = 318 \text{ K}$. However, we equivalently infer the zero-prior-mean process $\Theta_{0h} - \mu_{\text{pr}}$ as explained in the previous footnote.

4.2. Goal-oriented linear inverse problem. We now introduce the goal-oriented feature of the problem. As stated earlier, we are only interested in the initial temperature distribution over the CPU (i.e., in \mathcal{D}_1). Let \mathbf{Z} be the restriction of Θ_{0h} to the domain of interest \mathcal{D}_1 . Clearly, there is a linear map between \mathbf{Z} and Θ_{0h} , i.e., $\mathbf{Z} = \mathcal{O} \Theta_{0h}$ with $\mathcal{O} \in \mathbb{R}^{p \times n}$ and $p \ll n$. Thus, we have a linear-Gaussian goal-oriented inverse problem as introduced in Section 2:

$$\begin{cases} \mathbf{Y} = G \Theta_{0h} + \mathcal{E} \\ \mathbf{Z} = \mathcal{O} \Theta_{0h}, \end{cases} \quad (4.3)$$

where both the marginal distribution of Θ_{0h} and the likelihood $\mathbf{Y}|\Theta_{0h}$ are specified. (In this example we denote the parameters by Θ_{0h} rather than \mathbf{X} .) We choose a finite element discretization of the temperature field such that $\Theta_{0h} \in \mathbb{R}^{2400}$ and $\mathbf{Z} \in \mathbb{R}^{370}$. Our goal is to characterize optimal approximations of the posterior statistics of the QoI, $\mathbf{Z}|\mathbf{Y}$, for a given set of observations (see Figure 4.2 (left)). In this case, computing the posterior distribution of the QoI using direct formulas like (1.4) is infeasible as the QoI is a finite-dimensional approximation to a distributed stochastic process, $\Theta(0)|_{\mathcal{D}_1}$, and can be arbitrarily high-dimensional depending on the chosen level of discretization.

⁶There is no loss of generality in assuming zero prior mean. If we are given a statistical model of the form $\mathbf{Y} = G \Theta_{0h} + \mathcal{E}$, where $\Theta_{0h} \sim \mathcal{N}(\mu_{\text{pr}}, \Gamma_{\text{pr}})$ has a nonzero prior mean, then we can trivially rewrite this model as $\widehat{\mathbf{Y}} := \mathbf{Y} - G\mu_{\text{pr}} = G(\Theta_{0h} - \mu_{\text{pr}}) + \mathcal{E}$ for a modified data vector $\widehat{\mathbf{Y}}$ and infer, equivalently, a zero-prior-mean process $\Theta_{0h} - \mu_{\text{pr}} \sim \mathcal{N}(0, \Gamma_{\text{pr}})$.

The configuration of this problem highlights a crucial aspect of dimensionality reduction of goal-oriented inverse problems. Ideally we would place the sensors on \mathcal{D}_1 since we are interested in inferring the temperature field on the CPU. However, due to geometrical constraints, we are forced to place our sensors on the heat sink (\mathcal{D}_3). As a result, observations are much more informative about the parameters in \mathcal{D}_3 than in \mathcal{D}_1 . We see a hint of this in Figure 4.2 (*right*), which shows the normalized difference between the prior and posterior variance of the parameters, $(\text{Var}(\Theta_{0h}) - \text{Var}(\Theta_{0h}|\mathbf{Y}))/\text{Var}(\Theta_{0h})$. The prior variance is reduced the most around the sensor locations in \mathcal{D}_3 , which makes intuitive sense as the data are increasingly less informative as we move away from the sensors.

We first focus on the approximation of the posterior covariance of the QoI. If we use the suboptimal approximation introduced in (2.7), then we have to pay a considerable computational price as a result of the data being informative about directions in the parameter space that are not relevant to the QoI. This issue is illustrated by the numerical results in Figure 4.5. To set the stage, we begin with the posterior covariance Γ_{pos} of the parameters Θ_{0h} and construct the optimal approximation $\hat{\Gamma}_{\text{pos}} = \Gamma_{\text{pr}} - KK^\top$ from Theorem 2.1. Though this approximation is optimal for any given rank of the update, its convergence in this problem is rather slow—as shown by the dotted blue line in Figure 4.5—because there are many data-informed directions in the parameter space. (Notice the multitude of sensors on the heat sink in Figure 4.1, each yielding observations at M successive times.) If we now use $\hat{\Gamma}_{\text{pos}}$ to yield an approximation of the actual posterior covariance of interest $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ by means of $\Gamma_{\mathbf{Z}|\mathbf{Y}} \approx \hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \mathcal{O}\hat{\Gamma}_{\text{pos}}\mathcal{O}^\top$ (i.e., the naïve approximation of (2.7)), then the convergence of this approximation is still slow, as seen in green solid line of Figure 4.5. This slow convergence can be easily explained. The optimal approximation $\hat{\Gamma}_{\text{pos}}$ of Γ_{pos} accounts first for those directions that are most informed by the data. These directions correspond to modes with features near the sensors in \mathcal{D}_3 (see Figure 4.3), but they provide little information about the parameters in the region of interest (\mathcal{D}_1).

On the other hand, if we use the optimal approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ defined in Theorem 2.3, then convergence is remarkably fast, as illustrated via the red solid line in Figure 4.5. Now we only need to update $\Gamma_{\mathbf{Z}}$ along a handful of directions—say twenty—to achieve a satisfactory approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$. The key to achieving such fast convergence is to confine the inference to directions in the parameter space that are most informed by the data, relative to the prior, *and* that are relevant to the QoI. Moreover, these fundamental directions can be explicitly extracted from a *goal-oriented* approximation of the posterior covariance of the parameters, as explained in Lemma 2.7; three such directions are shown in Figure 4.4.

We note that $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ is by no means a low-rank matrix. (See its spectrum in Figure 4.7 (*left*)). This situation is fairly typical when dealing with large-scale inverse problems with non-smoothing priors (e.g., Gaussian fields with correlation function of Matérn type) and limited observations. In these situations, seeking an approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ as a low-rank matrix would be inappropriate; that is, classic dimensionality reduction techniques, e.g., Karhunen–Loève reduction [70, 61], are quite inefficient. Instead, low-dimensional structure lies in the change from prior to posterior, due to the data being informative, relative to the prior, only about a low-dimensional subspace of \mathbb{R}^p . This fact justifies the choice of the approximation class $\mathcal{M}_r^{\mathbf{Z}}$ for $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ in (2.8). The efficiency of the approximation class $\mathcal{M}_r^{\mathbf{Z}}$ is also evident from the sharp decay of the red curve in Figure 4.5: only a handful of directions

in the prior-to-posterior *update* are needed for a good approximation of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$.

The optimal approximation of the posterior mean of the QoI as a low-rank linear function of the data, as introduced in Theorem 2.9, also converges very quickly as a function of the rank of the approximation, as shown in Figure 4.6. Once a low-rank approximation of the form (2.25) is available, then one can compute an accurate approximation of $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ for each new realization of the data \mathbf{Y} by simply performing a low-rank ($r = 20$ in this case) matrix-vector product. See [26, 44, 68, 53, 28, 27] for a series of related efforts in a *non-goal-oriented* but possibly non-Gaussian framework.

4.3. A nonlinear QoI. We conclude this section by applying the approximation formulas introduced in this paper to a particular case of nonlinear goal-oriented inference. Suppose that we are only interested in the posterior distribution of the *maximum* temperature over \mathcal{D}_1 (see Figure 4.1). This is a useful QoI because the material properties of a sensitive component (e.g., the CPU) might deteriorate above a certain critical temperature (e.g., [18]). In this case, the QoI $\hat{\mathbf{Z}} := \max_{\mathcal{D}_1} \Theta_{0h}$ is a low-dimensional (in fact scalar-valued) nonlinear function of the parameters. In general, let us write $\hat{\mathbf{Z}} = \mathcal{J}(\Theta_{0h})$ for some nonlinear function $\mathcal{J} : \mathbb{R}^n \rightarrow \mathbb{R}$. Then we can cast the nonlinear goal-oriented Bayesian inverse problem as

$$\begin{cases} \mathbf{Y} = G \Theta_{0h} + \mathcal{E} \\ \hat{\mathbf{Z}} = \mathcal{J}(\Theta_{0h}) \end{cases} \quad (4.4)$$

and try to characterize the posterior $\hat{\mathbf{Z}}|\mathbf{Y}$ for a particular realization of the data. This problem is nontrivial, however, as $\hat{\mathbf{Z}}|\mathbf{Y}$ is non-Gaussian and cannot easily be characterized by just two moments. In the most general case, one needs to resort to sampling techniques such as MCMC [51] to characterize $\hat{\mathbf{Z}}|\mathbf{Y}$, or perhaps some deterministic alternative [73, 90, 34, 74]. Unfortunately, it is still not well understood how to properly adapt these techniques to exploit ultimate goals and bypass full inference of the parameters (see the offline-online strategy of [64] for a related effort in the context of goal-oriented nonlinear Bayesian inference). Though developing computationally efficient techniques to tackle general problems like (4.4) is of fundamental importance, in this particular example we can adopt a much simpler, yet effective, approach. Using the notation of this section, notice that the nonlinear QoI $\hat{\mathbf{Z}}$ can be written as $\hat{\mathbf{Z}} = g(\mathbf{Z})$, where \mathbf{Z} represents the inversion parameters in the region of interest \mathcal{D}_1 , and where $g(\mathbf{x}) = \max_i(x_i)$ for all $\mathbf{x} = (x_1, \dots, x_p)$. Thus we can rewrite (4.4) as:

$$\begin{cases} \mathbf{Y} = G \Theta_{0h} + \mathcal{E} \\ \mathbf{Z} = \mathcal{O} \Theta_{0h} \\ \hat{\mathbf{Z}} = g(\mathbf{Z}). \end{cases} \quad (4.5)$$

Then we can approximate the Gaussian posterior distribution of $\mathbf{Z}|\mathbf{Y}$ using the goal-oriented techniques presented in this paper, and finally we can push forward the latter distribution through the nonlinear operator g to obtain a suitable approximation of $\hat{\mathbf{Z}}|\mathbf{Y}$. The nonlinear operator g is never approximated in this process. That is, we first compute the posterior mean and the optimal goal-oriented approximation of the covariance of $\mathbf{Z}|\mathbf{Y}$ using the results of Theorem 2.3, then we sample from the optimal approximating measure $\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}} := \mathcal{N}(\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}), \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}})$ using the results of Lemma 2.8, and finally we push forward

these samples through $g : \mathbb{R}^p \rightarrow \mathbb{R}$ to obtain approximate samples from the posterior distribution of the nonlinear QoI. We can easily estimate the *quality* of these approximate posterior samples using bounds like (2.16). Note, however, that a bound like (2.16) quantifies only the accuracy of the posterior moments, and does not yield an explicit measure of distance between the non-Gaussian posterior distribution of $\hat{\mathbf{Z}}|\mathbf{Y}$ and its corresponding approximation. Nevertheless, the plot on the right of Figure 4.7 shows that the resulting approximation of the density of $\hat{\mathbf{Z}}|\mathbf{Y}$ is indeed quite good for this particular choice of nonlinear operator g .




FIG. 4.1. (left) CPU cooling problem. Inversion for the initial temperature field on \mathcal{D}_1 given noisy temperature measurements on an aluminum heat sink (\mathcal{D}_3). The figure shows the problem configuration, the locations of the sensors (black dots), and the boundary conditions for the heat equation describing time evolution of the temperature field on the domain $\mathcal{D} := \mathcal{D}_1 \cup \mathcal{D}_2 \cup \mathcal{D}_3$. (right) Material properties of the different layers.




FIG. 4.2. (left) Initial temperature field used to generate synthetic data according to the problem configuration described in Section 4. This temperature field was not drawn from the prior distribution of Θ_{0h} ; instead, it corresponds to a finer discretization of the continuous stochastic process Θ evaluated at the initial time. (right) Normalized difference between the prior and posterior variance of the parameters, i.e., $(\text{Var}(\Theta_{0h}) - \text{Var}(\Theta_{0h}|\mathbf{Y}))/\text{Var}(\Theta_{0h})$. The regions of greatest relative decrease of the variance are localized around the sensor locations (black dots).

5. Conclusions. We have developed statistically optimal and computationally efficient approximations of the posterior statistics of a *quantity of interest* (QoI) in a goal-oriented




FIG. 4.3. Three eigenvectors (w_i) of the matrix pencil $(H, \Gamma_{\text{pr}}^{-1})$ as defined in Theorem 2.1: w_1 (left), w_6 (center), and w_{10} (right). These eigenvectors define the prior-to-posterior update in the optimal approximation (2.5) of the posterior covariance of the parameters Γ_{pos} . Note that these leading eigenvectors have features near the locations of the sensors in \mathcal{D}_3 . This is the region where the data are most informative for the parameters, but not necessarily for the QoI.




FIG. 4.4. Three vectors (\tilde{q}_i) defining the prior-to-posterior update in the optimal goal-oriented approximation of Γ_{pos} introduced in (2.17) (see Lemma 2.7). In particular, we show \tilde{q}_1 (left), \tilde{q}_3 (center), and \tilde{q}_5 (right). One can interpret these vectors as directions in the parameter space that are informed by the data, relative to the prior, and that are relevant to the QoI. The relevant features of the (\tilde{q}_i) are concentrated around the region of interest (\mathcal{D}_1). These directions should be contrasted with the modes in Figure 4.3, which are strongly informed by the data but, at the same time, nearly irrelevant to the QoI.

linear-Gaussian inverse problem. The posterior covariance of the QoI is approximated as a low-rank negative update of the prior covariance of the QoI. Optimality holds with respect to the natural geodesic distance on the manifold of symmetric positive definite matrices. The posterior mean of the QoI is approximated as a low-rank function of the data, and optimality follows from the minimization of the Bayes risk for squared-error loss weighed by the posterior precision matrix of the QoI. The minimization of this Bayes risk is associated with the minimization of a Riemannian metric averaged over the distribution of the data. These optimal approximations avoid computation of the full posterior distribution of the parameters and focus only on directions in the parameter space that are informed by the data *and* that are relevant to the QoI. These directions are obtained as the leading generalized eigenvectors of a suitable matrix pencil, and reflect a balance among all the ingredients of the goal-oriented inverse problem: prior information, the forward model, measurement noise, and the ultimate goals.




FIG. 4.5. (left) Convergence of the covariance approximations in the natural geodesic distance over the manifold of SPD matrices (see Section 2.1). The blue dotted line shows the distance between the covariance of the parameters $\Theta_{0h}|\mathbf{Y}$ (i.e., Γ_{pos}) and its optimal approximation $\hat{\Gamma}_{\text{pos}} = \Gamma_{\text{pr}} - KK^\top$, as a function of the rank of K (see Theorem 2.1). The red line shows the distance between $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ and its optimal approximation introduced in Theorem 2.3, $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \Gamma_{\mathbf{Z}} - KK^\top$, as a function of the rank of K . Finally, the green line shows the distance between $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ and the suboptimal approximation (2.7) obtained as $\mathcal{O}\hat{\Gamma}_{\text{pos}}\mathcal{O}^\top$. (right) Detail of the figure on the left, with both axes rescaled.

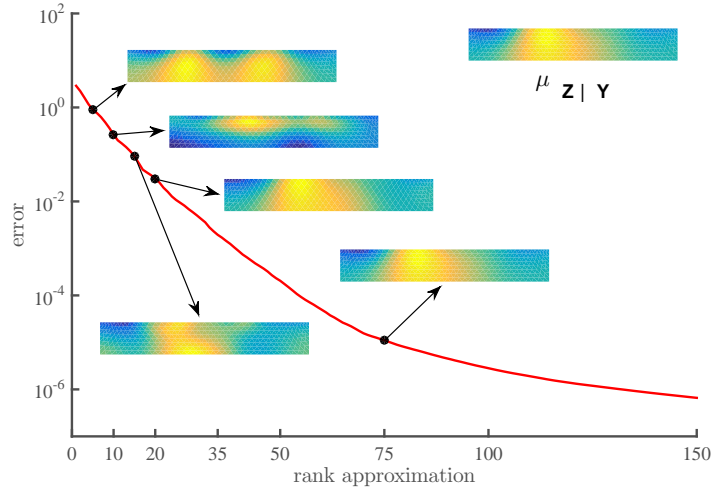


FIG. 4.6. The solid curve shows the error associated with the optimal low-rank approximation of the posterior mean of the QoI, $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$, given in Theorem 2.9). The error is measured as the square root of $\mathbb{E} \left[\left\| \mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) - A^* \mathbf{Y} \right\|_{\Gamma_{\mathbf{Z}|\mathbf{Y}}^{-1}}^2 \right]$ and is a function of $\text{rank}(A^*)$. The top right corner shows $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})$ for a particular realization of \mathbf{Y} (see Figure 4.2 (left)). The snapshots along the solid curve show the corresponding approximation $\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}) \approx A^* \mathbf{Y}$ for various ranks of A^* and for the same realization of \mathbf{Y} . Notice that the approximation of the posterior mean of the QoI is already good with $\text{rank}(A^*) = 20$.

An important avenue for future work is the extension of these optimality results to the case of nonlinear forward operators. Here, we expect that interpreting the QoI posterior approximation as the result of composing the forward model with a carefully chosen projection operator, as in [31], may be quite helpful. Relaxing the Gaussianity assumptions on




FIG. 4.7. (left) *Eigenvalues of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$. For this problem configuration $\mathbf{Z} \in \mathbb{R}^{370}$, so the matrix $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ is not low-rank.* (right) *The blue solid curve shows a kernel density estimate (KDE) of the exact posterior density of the nonlinear QoI $\widehat{\mathbf{Z}} := \max(\mathbf{Z})$, i.e., the density of $\widehat{\mathbf{Z}}|\mathbf{Y}$, constructed from 1×10^6 samples. Notice that this density is non-Gaussian. The red dotted curve shows a KDE constructed with 1×10^6 samples from an approximation of $\widehat{\mathbf{Z}}|\mathbf{Y}$ obtained as follows: First we sample the approximate measure $\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}} := \mathcal{N}(\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y}), \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}})$ obtained from an optimal approximation, $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$, of $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ as a 20-dimensional low rank update of $\Gamma_{\mathbf{Z}}$ (see Theorem 2.3). Then, we push forward these samples through the nonlinear goal-oriented operator $g: \mathbb{R}^p \rightarrow \mathbb{R}$. The quality of the density approximation is already good for a rank-20 update. These results are consistent with the theoretical bounds (2.16). (See, in particular, the error curves in Figure 4.5.)*

both the prior distribution and the measurement noise are also important generalizations of the present work.

Acknowledgments. We would like to thank Olivier Zahm, Jayanth Jagalur, Antoni Musolas, Ahmed Sameh, and Alicia Klinvex for many insightful discussions and for pointing us to key references in the literature. We would also like to thank Pierre-Antoine Absil, Omar Ghattas, and Georg Stadler for many helpful remarks on computational issues. This work was supported by the US Department of Energy, Office of Advanced Scientific Computing (ASCR), under grant number DE-SC0009297.

Appendix A. Rao’s metric between distributions. Let $\mathcal{M} = \{\pi_\theta, \theta \in \Theta\}$ be a parametric family of probability densities indexed by $\theta = (\theta_1, \dots, \theta_n) \in \Theta$ [6]. Rao considered a quadratic differential form given by

$$ds^2 = \sum_{i,j} g_{ij}(\theta) d\theta_i d\theta_j, \quad (\text{A.1})$$

where $g_{ij}(\theta) = \mathbb{E}_{\pi_\theta}[\partial_{\theta_i} \ln \pi_\theta \partial_{\theta_j} \ln \pi_\theta]$ are the entries of the Fisher information matrix, with \mathbb{E}_{π_θ} denoting integration with respect to π_θ [38]. The Fisher information matrix is a central object in mathematical statistics (e.g., the Cramér-Rao inequality [84]). Intuitively, we can interpret (A.1) as the variance of the function that describes the first order relative difference between π_θ and a contiguous density, $\pi_{\theta+d\theta}$, on \mathcal{M} [85]. The definition of a quadratic form like (A.1) allows us to measure curves on \mathcal{M} . Given a smooth curve $\gamma: [0, 1] \rightarrow \Theta \simeq \mathcal{M}$, we can define its length as $\ell(\gamma) := \int_0^1 (\sum_{i,j} g_{ij}(\gamma(t)) d\gamma_i d\gamma_j)^{1/2} dt$ [36]. Thus, Rao’s distance between a pair of distributions on \mathcal{M} is simply their geodesic distance, i.e., the length of the minimum length curve joining these distributions [85]. The quadratic form

defined by the Fisher information matrix is invariant under regular reparameterizations of \mathcal{M} [84]. Thus, this fundamental invariance is also shared by Rao's distance which yields an intrinsic way of comparing distributions on \mathcal{M} . Of course, it is possible to consider more general quadratic differential forms not based on the notion of Fisher information. See [86] for various examples of differential metrics derived from entropy functions or divergence measures between probability distributions. See [4] for a modern treatment of information geometry, the field at the intersection of statistics and differential geometry.

Appendix B. Proofs of the main results. The following two Lemmas, B.1 and 2.2, will be used to prove Theorems 2.3 and 2.9. We start with a result describing the relationships between different eigenpairs of the Schur complements of a particular class of covariance matrices that arises in Bayesian inverse problems.

LEMMA B.1 (Eigenpairs of Schur complements). *Let $\Sigma \succ 0$ be a matrix partitioned as*

$$\Sigma = \begin{pmatrix} A & B \\ B^\top & C \end{pmatrix}, \quad (\text{B.1})$$

where A and C are square matrices and $B \neq 0$. Then, A, C , and the Schur complements, $\mathcal{S}(A) := C - B^\top A^{-1}B$ and $\mathcal{S}(C) := A - BC^{-1}B^\top$, are also SPD matrices. Moreover:

1. If (β, w) is an eigenpair of $(BC^{-1}B^\top, A)$, then $\beta < 1$ and $(1 - \beta, w)$ is an eigenpair of $(\mathcal{S}(C), A)$. Furthermore, if $\beta \neq 0$, then $((1 - \beta)^{-1}, B^\top w)$ is an eigenpair of $(\mathcal{S}(A)^{-1}, C^{-1})$.
2. If $\beta \neq 0$ and (β, w) is an eigenpair of $(BC^{-1}B^\top, A)$, then $(\beta(1 - \beta)^{-1}, B^\top w)$ is an eigenpair of $(C^{-1}B^\top \mathcal{S}(C)^{-1}BC^{-1}, C^{-1})$.
3. If w_1, \dots, w_k are linearly independent eigenvectors of $(BC^{-1}B^\top, A)$ with associated eigenvalues $\beta_1 \geq \beta_2 \geq \dots \geq \beta_k > 0$, then $B^\top w_1, \dots, B^\top w_k$ are linearly independent. Moreover, if $k = \text{rank}(BC^{-1}B^\top)$, then there can be at most k linearly independent eigenvectors of $(C^{-1}B^\top \mathcal{S}(C)^{-1}BC^{-1}, C^{-1})$ associated with strictly positive eigenvalues.

Proof. The fact that $A, C, \mathcal{S}(A)$ and $\mathcal{S}(C)$ are SPD matrices follows from [17]. (1) From $BC^{-1}B^\top w = \beta Aw$ we obtain $\mathcal{S}(C)w = (1 - \beta)Aw$, which also implies that $\beta < 1$ as $\mathcal{S}(C) \succ 0$ and $A \succ 0$. If $\beta \neq 0$, then $B^\top w \neq 0$ and

$$\begin{aligned} \mathcal{S}(A)^{-1}B^\top w &= [C^{-1} + C^{-1}B^\top \mathcal{S}(C)^{-1}BC^{-1}]B^\top w \\ &= [C^{-1}B^\top + C^{-1}B^\top \mathcal{S}(C)^{-1}(A - \mathcal{S}(C))]w \\ &= (1 - \beta)^{-1}C^{-1}B^\top w. \end{aligned}$$

where we used the Woodbury identity to rewrite $\mathcal{S}(A)^{-1}$. (2) It follows from (1) that $((1 - \beta)^{-1}, B^\top w)$ is an eigenpair of $(\mathcal{S}(A)^{-1}, C^{-1})$ and by $\mathcal{S}(A)^{-1} = C^{-1} + C^{-1}B^\top \mathcal{S}(C)^{-1}BC^{-1}$ that $(\beta(1 - \beta)^{-1}, B^\top w)$ is an eigenpair of $(C^{-1}B^\top \mathcal{S}(C)^{-1}BC^{-1}, C^{-1})$. (3) If $\sum_{j=1}^k a_j B^\top w_j = 0$, then $A \sum_{j=1}^k \beta_j a_j w_j = 0$, and therefore $\sum_{j=1}^k \beta_j a_j w_j = 0$ since $A \succ 0$, which leads to $\beta_j a_j = 0$ for $j = 1, \dots, k$ since (w_j) are linearly independent, and thus $a_j = 0$ for $j = 1, \dots, k$ since $\beta_j > 0$. Moreover, notice that $\text{rank}(C^{-1}B^\top \mathcal{S}(C)^{-1}BC^{-1}) =$

$\text{rank}(B^\top \mathcal{S}(C)^{-1}B) = \text{rank}(BC^{-1}B^\top)$. Thus, there can be at most $\text{rank}(BC^{-1}B^\top)$ linearly independent eigenvectors of $(C^{-1}B^\top \mathcal{S}(C)^{-1}BC^{-1}, C^{-1})$ with nonzero eigenvalues. \square

Proof of Lemma 2.2. Consider the identity $\mathbf{Y} = G\mathbf{X} + \mathcal{E} = G\mathcal{O}_\dagger \mathcal{O}\mathbf{X} + G(I - \mathcal{O}_\dagger \mathcal{O})\mathbf{X} + \mathcal{E} = G\mathcal{O}_\dagger \mathbf{Z} + \mathbf{\Delta}$, where $\mathcal{O}_\dagger := \Gamma_{\text{pr}} \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1}$ and $\mathbf{\Delta} := G(I - \mathcal{O}_\dagger \mathcal{O})\mathbf{X} + \mathcal{E}$. A simple computation shows that $\mathbb{E}[(I - \mathcal{O}_\dagger \mathcal{O})\mathbf{X} \mathbf{Z}^\top] = 0$. Hence, $(I - \mathcal{O}_\dagger \mathcal{O})\mathbf{X}$ and \mathbf{Z} are uncorrelated, and, more importantly, independent since they are also jointly Gaussian. It follows that $\mathbf{\Delta}$ and \mathbf{Z} are also independent since \mathcal{E} was independent of \mathbf{X} and $\mathbf{Z} = \mathcal{O}\mathbf{X}$. In the hypothesis of zero prior mean, the mean of $\mathbf{\Delta}$ is also zero. Moreover, $\Gamma_{\mathbf{\Delta}} = \text{Var}[G(I - \mathcal{O}_\dagger \mathcal{O})\mathbf{X}] + \text{Var}[\mathcal{E}]$ since \mathbf{X} and \mathcal{E} are independent. Simple algebra leads to the particular form of $\Gamma_{\mathbf{\Delta}}$. \square

We can now prove the main results of this paper.

Proof of Theorem 2.3. By applying [97, Theorem 2.3] to the linear Gaussian model defined in Lemma 2.2, we know that a minimizer, $\hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$, of the geodesic distance, $d_{\mathcal{R}}$, between $\Gamma_{\mathbf{Z}|\mathbf{Y}}$ and an element of $\mathcal{M}_r^{\mathbf{Z}}$ is given by: $\hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \Gamma_{\mathbf{Z}} - \sum_{i=1}^r \eta_i^2 (1 + \eta_i^2)^{-1} \hat{q}_i \hat{q}_i^\top$, where (η_i^2, \hat{q}_i) are the eigenpairs of $(H_{\mathbf{Z}}, \Gamma_{\mathbf{Z}}^{-1})$, with the ordering $\eta_i^2 \geq \eta_{i+1}^2$, the normalization $\hat{q}_i^\top \Gamma_{\mathbf{Z}}^{-1} \hat{q}_i = 1$ and where $H_{\mathbf{Z}} := \mathcal{O}_\dagger^\top G^\top \Gamma_{\mathbf{\Delta}}^{-1} G \mathcal{O}_\dagger$ is the Hessian of the negative log-likelihood $\mathbf{Y}|\mathbf{Z} \sim \mathcal{N}(G\mathcal{O}_\dagger, \Gamma_{\mathbf{\Delta}})$. Moreover, [97, Theorem 2.3] implies that the distance, at optimality, is given by $d_{\mathcal{R}}^2(\hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}, \Gamma_{\mathbf{Z}|\mathbf{Y}}) = \sum_{i>r} \ln^2(1 + \eta_i^2)$ and that the minimizer is unique if the first r eigenvalues of $(H_{\mathbf{Z}}, \Gamma_{\mathbf{Z}}^{-1})$ are distinct. Now let (λ_i, q_i) be defined as in Theorem 2.3, with $\lambda_i > 0$, and let $\Sigma \succ 0$ be the covariance matrix of the joint distribution of \mathbf{Y} and \mathbf{Z} , i.e.,

$$\Sigma = \begin{pmatrix} \Gamma_{\mathbf{Y}} & G\Gamma_{\text{pr}} \mathcal{O}^\top \\ \mathcal{O} \Gamma_{\text{pr}} G^\top & \Gamma_{\mathbf{Z}} \end{pmatrix}. \quad (\text{B.2})$$

By Lemma B.1[part 2] applied to (B.2), we know that $(\lambda_i(1 - \lambda_i)^{-1}, \mathcal{O} \Gamma_{\text{pr}} G^\top q_i)$ are eigenpairs of $(H_{\mathbf{Z}}, \Gamma_{\mathbf{Z}}^{-1})$. Moreover, by Lemma B.1[part 3] we know that we can always write a maximal set of linearly independent eigenvectors of $(H_{\mathbf{Z}}, \Gamma_{\mathbf{Z}}^{-1})$, associated with nonzero eigenvalues, as $(\mathcal{O} \Gamma_{\text{pr}} G^\top q_i)$. Thus, since $f(\lambda) = \lambda(1 - \lambda)^{-1}$ is a decreasing function of λ as $\lambda \downarrow 0$, we must have $\eta_i^2 = \lambda_i(1 - \lambda_i)^{-1}$ and we can assume, without loss of generality, that $\hat{q}_i = \alpha \mathcal{O} \Gamma_{\text{pr}} G^\top q_i$ for some real $\alpha > 0$. Given the normalizations $\hat{q}_i^\top \Gamma_{\mathbf{Z}}^{-1} \hat{q}_i = 1$ and $q_i^\top (G \Gamma_{\text{pr}} \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1} \mathcal{O} \Gamma_{\text{pr}} G^\top) q_i = 1$, it follows that $\alpha = 1$. Simple algebra then leads to (2.11) and (2.12). Notice, that $\lambda_i > 0$ and $f(\lambda_i) > 0$ imply $\lambda_i < 1$. This property will be useful when proving Lemma 2.8. \square

We now state a standard result that will be used in proving Lemma 2.5.

THEOREM B.2 (Cauchy Interlacing Theorem (e.g., [59, 15])). *Let $A, B \in \mathbb{R}^{n \times n}$ be symmetric matrices with $B \succ 0$ and let $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_n$ be the eigenvalues of (A, B) . For any $P \in \mathbb{R}^{n \times p}$, with $p \leq n$ and full column-rank, let $\mu_1 \geq \mu_2 \geq \dots \geq \mu_p$ be the eigenvalues of $(P^\top A P, P^\top B P)$. Then:*

$$\gamma_k \geq \mu_k \geq \gamma_{n-p+k}, \quad k = 1, \dots, p. \quad (\text{B.3})$$

Proof of Lemma 2.5. The first inequality in (2.14) follows from the optimality statement of Theorem 2.3 since $\hat{\Gamma}_{\mathbf{Z}|\mathbf{Y}} \in \mathcal{M}_r^{\mathbf{Z}}$. The second inequality in (2.14) follows from

the Cauchy interlacing theorem (see Theorem B.2). Let $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_n \geq 1$ be the eigenvalues of $(\widehat{\Gamma}_{\text{pos}}, \Gamma_{\text{pos}})$ and $\mu_1 \geq \mu_2 \geq \dots \geq \mu_p$ be the eigenvalues of $(\widehat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}, \Gamma_{\mathbf{Z}|\mathbf{Y}}) = (\mathcal{O} \widehat{\Gamma}_{\text{pos}} \mathcal{O}^\top, \mathcal{O} \Gamma_{\text{pos}} \mathcal{O}^\top)$, where \mathcal{O} is a full row-rank matrix. Then, by Theorem B.2,

$$\gamma_k \geq \mu_k \geq 1, \quad k = 1, \dots, p. \quad (\text{B.4})$$

In particular, since $\ln^2(x)$ is monotone increasing on $x > 1$, we have:

$$d_{\mathcal{R}}(\Gamma_{\mathbf{Z}|\mathbf{Y}}, \widehat{\Gamma}_{\mathbf{Z}|\mathbf{Y}}) = \frac{1}{2} \sum_k \ln^2(\mu_k) \leq \frac{1}{2} \sum_k \ln^2(\gamma_k) \leq d_{\mathcal{R}}(\Gamma_{\text{pos}}, \widehat{\Gamma}_{\text{pos}}), \quad (\text{B.5})$$

where clearly $\sum_{k>p} \ln^2(\gamma_k) \geq 0$. \square

The following two lemmas will be used in proving Lemma 2.6.

LEMMA B.3. *If $\Gamma_1 \succeq \Gamma_2 \succ 0$, then $|\Gamma_1| \geq |\Gamma_2|$*

Proof. If $\Gamma_1 \succeq \Gamma_2$, then there exists a $S \succeq 0$ such that $\Gamma_1 = \Gamma_2 + S$. Thus, $|\Gamma_1| |\Gamma_2|^{-1} = |I + \Gamma_2^{-1/2} S \Gamma_2^{-1/2}| \geq 1$. \square

LEMMA B.4. *Let $X \sim \mathcal{N}(\mu, \Sigma)$ and $Y \sim \mathcal{N}(0, \Gamma)$ with $\Gamma \succeq \Sigma \succ 0$. Let g be a measurable real-valued function such that*

$$\mathbb{E}[|g|^{2+\alpha}(Y)] < \infty \quad (\text{B.6})$$

for some $\alpha > 0$. Then,

$$\mathbb{E}[g^2(X)] \leq \frac{|\Gamma|^{1/2}}{|\Sigma|^{1/2}} \exp\left(\frac{\mu^\top \Gamma^{-1} \mu}{\alpha}\right) \mathbb{E}[|g|^{2+\alpha}(Y)]^{1/(1+\alpha/2)} \quad (\text{B.7})$$

Proof. Let f_X and f_Y be the densities of X and Y , respectively, and M_Y the moment generating function of Y . Since we have $\Sigma^{-1} = \Gamma^{-1} + S$ for some $S \succeq 0$, it follows that for all x ,

$$f_X(x) \leq K \exp(\mu^\top \Gamma^{-1} x) f_Y(x), \quad (\text{B.8})$$

where $K := |\Gamma|^{1/2} |\Sigma|^{-1/2} \exp(-\mu^\top \Gamma^{-1} \mu/2)$. Now we use Hölder's inequality, with $p = 1 + \alpha/2$ and $q = p/(p-1)$ so that $1/p + 1/q = 1$, to obtain:

$$\begin{aligned} \mathbb{E}[g^2(X)] &\leq K \mathbb{E}[g^2(Y) \exp(\mu^\top \Gamma^{-1} Y)] \\ &\leq K (\mathbb{E}[|g|^{2p}(Y)])^{1/p} (\mathbb{E}[\exp(q \mu^\top \Gamma^{-1} Y)])^{1/q} \\ &= K (\mathbb{E}[|g|^{2p}(Y)])^{1/p} M_Y^{1/q}(q \Gamma^{-1} \mu) \\ &= K (\mathbb{E}[|g|^{2p}(Y)])^{1/p} \exp(q \mu^\top \Gamma^{-1} \mu/2) \\ &= |\Gamma|^{1/2} |\Sigma|^{-1/2} (\mathbb{E}[|g|^{2p}(Y)])^{1/p} \exp((q-1) \mu^\top \Gamma^{-1} \mu/2) \\ &= |\Gamma|^{1/2} |\Sigma|^{-1/2} \exp(\mu^\top \Gamma^{-1} \mu/\alpha) (\mathbb{E}[|g|^{2+\alpha}(Y)])^{1/(1+\alpha/2)}, \end{aligned}$$

where we used the fact that $M_Y(t) = \exp(t^\top \Gamma t/2)$ since $Y \sim \mathcal{N}(0, \Gamma)$. \square

Proof of Lemma 2.6. By [33, Lemma 7.14] we have:

$$\left| \mathbb{E}_{\nu_{\mathbf{Z}|\mathbf{Y}}}[g] - \mathbb{E}_{\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}}[g] \right| \leq 2 \sqrt{\int |g|^2 (\pi_{\mathbf{Z}|\mathbf{Y}} + \tilde{\pi}_{\mathbf{Z}|\mathbf{Y}}) d_{\text{Hell}}(\nu_{\mathbf{Z}|\mathbf{Y}}, \tilde{\nu}_{\mathbf{Z}|\mathbf{Y}})} \quad (\text{B.9})$$

where $\pi_{\mathbf{Z}|\mathbf{Y}}$ and $\tilde{\pi}_{\mathbf{Z}|\mathbf{Y}}$ are, respectively, the densities of $\nu_{\mathbf{Z}|\mathbf{Y}}$ and $\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}$ with respect to the Lebesgue measure. Now notice that $\Gamma_{\mathbf{Z}|\mathbf{Y}} \preceq \Gamma_{\mathbf{Z}}$ as well as $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} \preceq \Gamma_{\mathbf{Z}}$. Thus, by Lemma B.4, we have:

$$\mathbb{E}_{\nu_{\mathbf{Z}|\mathbf{Y}}} [|g|^2] + \mathbb{E}_{\tilde{\nu}_{\mathbf{Z}|\mathbf{Y}}} [|g|^2] \leq 2 \frac{|\Gamma_{\mathbf{Z}}|^{1/2}}{|\Gamma_{\mathbf{Z}|\mathbf{Y}}|^{1/2}} \exp\left(\frac{1}{\beta-2} \|\mu_{\mathbf{Z}|\mathbf{Y}}(\mathbf{Y})\|_{\Gamma_{\mathbf{Z}}^{-1}}^2\right) \mathbb{E}_{\nu_{\mathbf{Z}}} [|g|^\beta]^{2/\beta} \quad (\text{B.10})$$

where we used the fact that $|\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}| \geq |\Gamma_{\mathbf{Z}|\mathbf{Y}}|$ since $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} \succeq \Gamma_{\mathbf{Z}|\mathbf{Y}}$ (see Lemma B.3). Thus, (2.16) follows from simple algebra. \square

LEMMA B.5. *Let $M := A(I - BB^\top)A^\top \succ 0$ for a pair of compatible matrices A, B , and let P be the orthogonal projector onto the range of A^\top . Then $C := I - PBB^\top P \succ 0$ and $M = ACA^\top$.*

Proof. Since $M \succ 0$, A^\top must be full column rank. Thus, by definition, $P = A^\top(AA^\top)^{-1}A$ and $PA^\top = I = AP$. Hence $M = ACA^\top$. Now let $Q := I - P$ and notice that $PQ = 0$ and $CQ = Q$. Thus, for $z \neq 0$, $\langle Cz, z \rangle = \langle CPz, Pz \rangle + \langle Qz, Qz \rangle = \langle CPz, Pz \rangle + \|Qz\|^2$. In particular, $\langle CPz, Pz \rangle = \langle PCPz, z \rangle = \langle M(AA^\top)^{-1}Az, (AA^\top)^{-1}Az \rangle \geq 0$ and it is zero only if $Pz = 0$, in which case $Qz \neq 0$ and $\|Qz\| > 0$. Thus $C \succ 0$. \square

Proof of Lemma 2.7. We first need to show the equivalence $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} \equiv \mathcal{O}\hat{\Gamma}_{\text{pos}}^* \mathcal{O}^\top$, where $\hat{\Gamma}_{\text{pos}}^*$ is defined in (2.17). Notice that $\mathcal{O}\hat{\Gamma}_{\text{pos}}^* \mathcal{O}^\top = \Gamma_{\mathbf{Z}} - \sum_i \lambda_i v_i v_i^\top$ with $v_i := \mathcal{O}S_{\text{pr}} \Pi S_{\text{pr}}^\top G^\top q_i = \mathcal{O}\Gamma_{\text{pr}} G^\top q_i$ since Π is a projector onto the rowspace of $\mathcal{O}S_{\text{pr}}$. The desired equivalence follows by comparison with (2.11). In particular, it follows that $\mathcal{O}\hat{\Gamma}_{\text{pos}}^* \mathcal{O}^\top \succ 0$. Thus, in order to show that $\hat{\Gamma}_{\text{pos}}^*$ is in the feasible set of (2.18) it remains to prove that $\hat{\Gamma}_{\text{pos}}^* \in \mathcal{M}_r$. Clearly, it just suffices to show that $\hat{\Gamma}_{\text{pos}}^* \succ 0$. Notice that $\tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}} = \mathcal{O}S_{\text{pr}}(I - \Pi BB^\top \Pi)S_{\text{pr}}^\top \mathcal{O}^\top \succ 0$ where $BB^\top := S_{\text{pr}}^\top G^\top \sum_{i=1}^r q_i q_i^\top G S_{\text{pr}}$. Thus, we can apply Lemma B.5 with $M := \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$, $A := \mathcal{O}S_{\text{pr}}$, $C := I - \Pi BB^\top \Pi$, and get $C \succ 0$. In particular, this shows that $\hat{\Gamma}_{\text{pos}}^* = S_{\text{pr}} C S_{\text{pr}}^\top \succ 0$ and thus $\hat{\Gamma}_{\text{pos}}^*$ is in the feasible set of (2.17). Optimality of $\hat{\Gamma}_{\text{pos}}^*$ then follows almost immediately. By Theorem 2.3:

$$d_{\mathcal{R}}(\Gamma_{\mathbf{Z}|\mathbf{Y}}, \mathcal{O}\hat{\Gamma}_{\text{pos}}^* \mathcal{O}^\top) = d_{\mathcal{R}}(\Gamma_{\mathbf{Z}|\mathbf{Y}}, \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}) \leq d_{\mathcal{R}}(\Gamma_{\mathbf{Z}|\mathbf{Y}}, \tilde{\Gamma}) \quad \forall \tilde{\Gamma} \in \mathcal{M}_r^{\mathbf{Z}}. \quad (\text{B.11})$$

In particular, we can consider $\tilde{\Gamma}$ of the form $\tilde{\Gamma} = \mathcal{O}\Gamma\mathcal{O}^\top$ for $\Gamma \in \mathcal{M}_r$. Notice that $\tilde{\Gamma} \succ 0$ since \mathcal{O} is assumed to be full row-rank. This shows optimality of $\hat{\Gamma}_{\text{pos}}^*$ according to (2.18). \square

Proof of Lemma 2.8. We first provide an explicit square root factorization of $\hat{\Gamma}_{\text{pos}}^*$, defined in (2.17) (Lemma 2.7), as $\hat{\Gamma}_{\text{pos}}^* = \hat{S}_{\text{pos}}^* (\hat{S}_{\text{pos}}^*)^\top$ for some matrix \hat{S}_{pos}^* . We claim that

$$\hat{S}_{\text{pos}}^* = S_{\text{pr}} \left(\sum_{i=1}^r (\sqrt{1 - \lambda_i} - 1) \bar{q}_i \bar{q}_i^\top + I \right) \quad (\text{B.12})$$

where I is the identity matrix. and the (\bar{q}_i) are defined in (2.20). First of all, notice that (B.12) is well defined since $1 > \lambda_i > 0$ for all $i = 1, \dots, r$ (see the proof of Theorem 2.3). One can verify that (B.12) is indeed a valid square root of $\hat{\Gamma}_{\text{pos}}^*$. Their key observation is

that $\bar{q}_i = S_{\text{pr}}^{-1} \tilde{q}_i$ and that the vectors (\tilde{q}_i) are Γ_{pr}^{-1} -orthogonal, i.e., $\tilde{q}_i^\top \Gamma_{\text{pr}}^{-1} \tilde{q}_j = \delta_{ij}$. To see this, consider the following identities:

$$\tilde{q}_i^\top \Gamma_{\text{pr}}^{-1} \tilde{q}_j = q_i^\top G S_{\text{pr}} \Pi S_{\text{pr}}^\top \Gamma_{\text{pr}}^{-1} S_{\text{pr}} \Pi S_{\text{pr}}^\top G^\top q_j = q_i^\top G S_{\text{pr}} \Pi S_{\text{pr}}^\top G^\top q_j. \quad (\text{B.13})$$

Since $\Pi = S_{\text{pr}}^\top \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1} \mathcal{O} S_{\text{pr}}$, it must be that $\tilde{q}_i^\top \Gamma_{\text{pr}}^{-1} \tilde{q}_j = q_i^\top G \Gamma_{\text{pr}} \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1} \mathcal{O} \Gamma_{\text{pr}} G^\top q_j = \delta_{ij}$, for the (q_i) are the generalized eigenvectors of the pencil (2.10), properly normalized. Now notice that $\tilde{S}_{\mathbf{Z}|\mathbf{Y}} = \mathcal{O} \hat{S}_{\text{pos}}^*$ and thus (2.17) implies that $\tilde{S}_{\mathbf{Z}|\mathbf{Y}} \tilde{S}_{\mathbf{Z}|\mathbf{Y}}^\top = \tilde{\Gamma}_{\mathbf{Z}|\mathbf{Y}}$. \square

Proof of Theorem 2.9. By applying [97, Theorem 4.1] to the linear Gaussian model defined in Lemma 2.2, we know that a minimizer of (2.24) is given by: $A^* = \sum_{i=1}^r \eta_i (1 + \eta_i^2)^{-1} \hat{q}_i \hat{v}_i^\top$, where (η_i^2, \hat{q}_i) are eigenpairs of $(H_{\mathbf{Z}}, \Gamma_{\mathbf{Z}}^{-1})$ with normalization $\hat{q}_i^\top \Gamma_{\mathbf{Z}}^{-1} \hat{q}_i = 1$, whereas (\hat{v}_i) are eigenvectors of $(G \mathcal{O}_\dagger \Gamma_{\mathbf{Z}} \mathcal{O}_\dagger^\top G^\top, \Gamma_{\mathbf{\Delta}})$ with normalization $\hat{v}_i^\top \Gamma_{\mathbf{\Delta}} \hat{v}_i = 1$. Moreover, [97, Theorem 4.1] tells us that the Bayes risk associated with the minimizer A^* can be written as: $\mathbb{E}[\|A^* \mathbf{Y} - \mathbf{Z}\|_{\Gamma_{\mathbf{Z}|\mathbf{Y}}}^2] = \sum_{i>r} \eta_i^2 + n$, where n is the dimension of the parameter space. The fact that the vectors (\hat{q}_i) can be written as $\hat{q}_i = \mathcal{O} \Gamma_{\text{pr}} G^\top q_i$ for $\eta_i^2 > 0$ was proved in Theorem 2.3. Furthermore, in the proof of Theorem 2.3 we showed that $\eta_i^2 = \lambda_i (1 - \lambda_i)^{-1}$. Using the latter expression we can rewrite the minimizer as $A^* = \sum_{i=1}^r \sqrt{\lambda_i (1 - \lambda_i)} \hat{q}_i \hat{v}_i^\top$. If (\hat{v}_i) are eigenvectors of $(G \mathcal{O}_\dagger \Gamma_{\mathbf{Z}} \mathcal{O}_\dagger^\top G^\top, \Gamma_{\mathbf{\Delta}})$, then they must also be eigenvectors of $(G \Gamma_{\text{pr}} \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1} \mathcal{O} \Gamma_{\text{pr}} G^\top, \Gamma_{\mathbf{Y}})$. In particular, we can set $\hat{v}_i = \alpha q_i$ for some real $\alpha > 0$. Given the normalizations $q_i^\top G \Gamma_{\text{pr}} \mathcal{O}^\top \Gamma_{\mathbf{Z}}^{-1} \mathcal{O} \Gamma_{\text{pr}} G^\top q_i = 1$ and $\hat{v}_i^\top \Gamma_{\mathbf{\Delta}} \hat{v}_i = 1$, it must be $\alpha = \lambda_i^{1/2} (1 - \lambda_i)^{-1/2}$. Simple algebra then leads to (2.25). \square

REFERENCES

- [1] P. A. ABSIL, C. G. BAKER, AND K. A. GALLIVAN, *A truncated-CG style method for symmetric generalized eigenvalue problems*, J. Comput. Appl. Math, 189 (2006), pp. 274–285.
- [2] P. A. ABSIL, C. G. BAKER, K. A. GALLIVAN, AND A. SAMEH, *Adaptive model trust region methods for generalized eigenvalue problems*, in International Conference on Computational Science, Springer, 2005, pp. 33–41.
- [3] V. AKÇELIK, G. BIROS, O. GHATTAS, J. HILL, D. KEYES, AND B. VAN BLOEMEN WAANDERS, *Parallel algorithms for PDE-constrained optimization*, Parallel Process. Sci. Comput, 20 (2006), p. 291.
- [4] S. AMARI AND H. NAGAOKA, *Methods of Information Geometry*, vol. 191, American Mathematical Soc., 2007.
- [5] V. ARSIGNY, P. FILLARD, X. PENNEC, AND N. AYACHE, *Geometric means in a novel vector space structure on symmetric positive-definite matrices*, SIAM Journal on Matrix Analysis and Applications, 29 (2007), pp. 328–347.
- [6] C. ATKINSON AND A. F. S. MITCHELL, *Rao's distance measure*, Sankhyā: The Indian Journal of Statistics, Series A, 43 (1981), pp. 345–365.
- [7] H. AUVINEN, J. M. BARDSLEY, H. HAARIO, AND T. KAURANNE, *Large-scale Kalman filtering using the limited memory BFGS method*, Electronic Transactions on Numerical Analysis, 35 (2009), pp. 217–233.
- [8] ———, *The variational Kalman filter and an efficient implementation using limited memory BFGS*, International Journal for Numerical Methods in Fluids, 64 (2010), pp. 314–335.
- [9] O. AXELSSON AND I. KAPORIN, *On the sublinear and superlinear rate of convergence of conjugate gradient methods*, Numerical Algorithms, 25 (2000), pp. 1–22.
- [10] Z. BAI, J. DEMMEL, J. DONGARRA, A. RUHE, AND H. VAN DER VORST, *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, vol. 11, SIAM, 2000.

- [11] C. G. BAKER, P. A. ABSIL, AND K. A. GALLIVAN, *An implicit Riemannian trust-region method for the symmetric generalized eigenproblem*, in International Conference on Computational Science, Springer, 2006, pp. 210–217.
- [12] A. BARACHANT, S. BONNET, M. CONGEDO, AND C. JUTTEN, *Classification of covariance matrices using a Riemannian-based kernel for BCI applications*, Neurocomputing, 112 (2013), pp. 172–178.
- [13] T. L. BERGMAN, F. P. INCROPERA, AND A. S. LAVINE, *Fundamentals of Heat and Mass Transfer*, John Wiley & Sons, 2011.
- [14] R. BHATIA, *Positive Definite Matrices*, Princeton University Press, 2009.
- [15] ———, *Matrix Analysis*, vol. 169, Springer Science & Business Media, 2013.
- [16] S. BONNABEL AND R. SEPULCHRE, *Riemannian metric and geometric mean for positive semidefinite matrices of fixed rank*, SIAM Journal on Matrix Analysis and Applications, 31 (2009), pp. 1055–1070.
- [17] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, 2004.
- [18] C. M. BRANCO, R. RITCHIE, AND V. SKLENICKA, *Mechanical behaviour of materials at high temperature*, vol. 15, Springer Science & Business Media, 1996.
- [19] T. BUI-THANH, C. BURSTEDDE, O. GHATTAS, J. MARTIN, G. STADLER, AND L. WILCOX, *Extreme-scale UQ for Bayesian inverse problems governed by PDEs*, in Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, IEEE Computer Society Press, 2012, p. 3.
- [20] T. BUI-THANH AND O. GHATTAS, *Analysis of the Hessian for inverse scattering problems: I. Inverse shape scattering of acoustic waves*, Inverse Problems, 28 (2012), p. 055001.
- [21] T. BUI-THANH, O. GHATTAS, AND D. HIGDON, *Adaptive Hessian-based nonstationary Gaussian process response surface method for probability density approximation with application to Bayesian solution of large-scale inverse problems*, SIAM Journal on Scientific Computing, 34 (2012), pp. A2837–A2871.
- [22] D. CALVETTI AND E. SOMERSALO, *Priorconditioners for linear systems*, Inverse Problems, 21 (2005), p. 1397.
- [23] S. T. CHOI, *Iterative Methods for Singular Linear Equations and Least-Squares Problems*, PhD thesis, Stanford University, 2006.
- [24] E. CHOW AND Y. SAAD, *Preconditioned Krylov subspace methods for sampling multivariate Gaussian distributions*, SIAM Journal on Scientific Computing, 36 (2014), pp. A588–A608.
- [25] J. CHUNG AND M. CHUNG, *Computing optimal low-rank matrix approximations for image processing*, in Signals, Systems and Computers, 2013 Asilomar Conference on, IEEE, 2013, pp. 670–674.
- [26] ———, *An efficient approach for computing optimal low-rank regularized inverse matrices*, Inverse Problems, 30 (2014), p. 114009.
- [27] J. CHUNG, M. CHUNG, AND D.P. O’LEARY, *Designing optimal spectral filters for inverse problems*, SIAM Journal on Scientific Computing, 33 (2011), pp. 3132–3152.
- [28] ———, *Optimal filters from calibration data for image deconvolution with data acquisition error*, Journal of Mathematical Imaging and Vision, 44 (2012), pp. 366–374.
- [29] J. CHUNG, M. CHUNG, AND D. P. O’LEARY, *Optimal regularized low rank inverse approximation*, Linear Algebra and its Applications, 468 (2015), pp. 260 – 269.
- [30] T. CUI, K. J. H. LAW, AND Y. M. MARZOUK, *Dimension-independent likelihood-informed MCMC*, Journal of Computational Physics, 304 (2016), pp. 109–137.
- [31] T. CUI, J. MARTIN, Y. M. MARZOUK, A. SOLONEN, AND A. SPANTINI, *Likelihood-informed dimension reduction for nonlinear inverse problems*, Inverse Problems, 30 (2014), p. 114015.
- [32] J. CULLUM AND W. DONATH, *A block Lanczos algorithm for computing the q algebraically largest eigenvalues and a corresponding eigenspace of large, sparse, real symmetric matrices*, in Decision and Control including the 13th Symposium on Adaptive Processes, 1974 IEEE Conference on, IEEE, 1974, pp. 505–509.
- [33] M. DASHTI AND A. M. STUART, *The Bayesian approach to inverse problems*, arXiv preprint arXiv:1302.6989, (2013).
- [34] J. DICK, R. N. GANTNER, Q. T. L. GIA, AND C. SCHWAB, *Higher order quasi-Monte Carlo integration for Bayesian estimation*, arXiv preprint arXiv:1602.07363, (2016).
- [35] C. R. DIETRICH AND G. N. NEWSAM, *Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix*, SIAM Journal on Scientific Computing, 18

- (1997), pp. 1088–1107.
- [36] M. P. DO CARMO, *Riemannian Geometry*, Birkhäuser, 1992.
 - [37] L. DYKES AND L. REICHEL, *Simplified GSVD computations for the solution of linear discrete ill-posed problems*, Journal of Computational and Applied Mathematics, 255 (2014), pp. 15–27.
 - [38] R. A. FISHER, *On the mathematical foundations of theoretical statistics*, Philosophical Transactions of the Royal Society of London. Series A, 222 (1922), pp. 309–368.
 - [39] H. P. FLATH, L. WILCOX, V. AKÇELİK, J. HILL, B. VAN BLOEMEN WAANDERS, AND O. GHATTAS, *Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations*, SIAM Journal on Scientific Computing, 33 (2011), pp. 407–432.
 - [40] P. T. FLETCHER, C. LU, S. M. PIZER, AND S. JOSHI, *Principal geodesic analysis for the study of nonlinear statistics of shape*, Medical Imaging, IEEE Transactions on, 23 (2004), pp. 995–1005.
 - [41] D. C. FONG AND M. SAUNDERS, *LSMR: An iterative algorithm for sparse least-squares problems*, SIAM Journal on Scientific Computing, 33 (2011), pp. 2950–2971.
 - [42] W. FÖRSTNER AND B. MOONEN, *A metric for covariance matrices*, in Geodesy-The Challenge of the 3rd Millennium, Springer, 2003, pp. 299–309.
 - [43] C. FOX AND A. PARKER, *Convergence in variance of Chebyshev accelerated Gibbs samplers*, SIAM Journal on Scientific Computing, 36 (2014), pp. A124–A147.
 - [44] S. FRIEDLAND AND A. TOROKHTI, *Generalized rank-constrained matrix approximations*, SIAM Journal on Matrix Analysis and Applications, 29 (2007), pp. 656–659.
 - [45] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, vol. 3, JHU Press, 2012.
 - [46] G. H. GOLUB AND Q. YE, *An inverse free preconditioned Krylov subspace method for symmetric generalized eigenvalue problems*, SIAM Journal on Scientific Computing, 24 (2002), pp. 312–334.
 - [47] G. GUGLIELMINI AND C. PISONI, *Elementi di trasmissione del calore*, Veschi, 1990.
 - [48] N. HALKO, P. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 53 (2011), pp. 217–288.
 - [49] P. C. HANSEN, *Regularization, GSVD and truncated GSVD*, BIT Numerical Mathematics, 29 (1989), pp. 491–504.
 - [50] M. R. HESTENES AND E. STIEFEL, *Methods of conjugate gradients for solving linear systems*, vol. 49, National Bureau of Standards Washington, DC, 1952.
 - [51] D. HIGDON, C. S. REESE, J. D. MOULTON, J. A. VRUGT, AND C. FOX, *Posterior exploration for computationally intensive forward models*, Handbook of Markov chain Monte Carlo, CHAPMAN & HALL/CRC, (2011), pp. 401–418.
 - [52] I. HOREV, F. YGER, AND M. SUGIYAMA, *Geometry-aware principal component analysis for symmetric positive definite matrices*, in JMLR: Workshop and Conference Proceedings, 2015.
 - [53] Y. HUA AND W. LIU, *Generalized Karhunen-Loeve transform*, IEEE Signal Processing Letters, 5 (1998), pp. 141–142.
 - [54] T. ISAAC, N. PETRA, G. STADLER, AND O. GHATTAS, *Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the antarctic ice sheet*, Journal of Computational Physics, 296 (2015), pp. 348–368.
 - [55] S. T. JENSEN. Private communication, 1976. (As cited in Atkinson, Mitchell, “Rao’s Distance Measure,” Sankhyā: The Indian Journal of Statistics, Series A, 43 (1981), pp. 345–365.)
 - [56] JARI KAIPIO AND ERKKI SOMERSALO, *Statistical inverse problems: Discretization, model reduction and inverse crimes*, Journal of Computational and Applied Mathematics, 198 (2007), pp. 493–504.
 - [57] A. G. KALMIKOV AND P. HEIMBACH, *A Hessian-based method for uncertainty quantification in global ocean state estimation*, SIAM Journal on Scientific Computing, 36 (2014), pp. S267–S295.
 - [58] A. KLINVEX, F. SAIED, AND A. SAMEH, *Parallel implementations of the trace minimization scheme tracemin for the sparse symmetric eigenvalue problem*, Computers & Mathematics with Applications, 65 (2013), pp. 460–468.
 - [59] D. KRESSNER, M. M. PANDUR, AND M. SHAO, *An indefinite variant of LOBPCG for definite matrix pencils*, Numerical Algorithms, 66 (2014), pp. 681–703.
 - [60] C. LANCZOS, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, United States Governm. Press Office, 1950.

- [61] W. LI AND O. A. CIRPKA, *Efficient geostatistical inverse methods for structured and unstructured grids*, Water Resources Research, 42 (2006).
- [62] E. LIBERTY, F. WOOLFE, P. MARTINSSON, V. ROKHLIN, AND M. TYGERT, *Randomized algorithms for the low-rank approximation of matrices*, Proceedings of the National Academy of Sciences, 104 (2007), pp. 20167–20172.
- [63] C. LIEBERMAN AND K. WILLCOX, *Goal-oriented inference: approach, linear theory, and application to advection diffusion*, SIAM Journal on Scientific Computing, 34 (2012), pp. A1880–A1904.
- [64] ———, *Nonlinear goal-oriented Bayesian inference: application to carbon capture and storage*, SIAM Journal on Scientific Computing, 36 (2014), pp. B427–B449.
- [65] J. LIESEN AND P. TICHÝ, *Convergence analysis of Krylov subspace methods*, GAMM-Mitteilungen, 27 (2004), pp. 153–173.
- [66] F. LINDGREN, H. RUE, AND J. LINDSTRÖM, *An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach*, Journal of the Royal Statistical Society: Series B, 73 (2011), pp. 423–498.
- [67] P. C. MAHALANOBIS, *On the generalized distance in statistics*, Proceedings of the National Institute of Sciences (Calcutta), 2 (1936), pp. 49–55.
- [68] I. MARKOVSKY, *Structured low-rank approximation and its applications*, Automatica, 44 (2008), pp. 891–909.
- [69] J. MARTIN, L. WILCOX, C. BURSTEDDE, AND O. GHATTAS, *A stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion*, SIAM Journal on Scientific Computing, 34 (2012), pp. A1460–A1487.
- [70] Y. M. MARZOUK AND H. N. NAJM, *Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems*, Journal of Computational Physics, 228 (2009), pp. 1862–1902.
- [71] X. MENG, M. A. SAUNDERS, AND M. W. MAHONEY, *LSRN: A parallel iterative solver for strongly over- or underdetermined systems*, SIAM Journal on Scientific Computing, 36 (2014), pp. C95–C118.
- [72] M. MOAKHER AND M. ZÉRAÏ, *The Riemannian geometry of the space of positive-definite matrices and its application to the regularization of positive-definite matrix-valued data*, Journal of Mathematical Imaging and Vision, 40 (2011), pp. 171–187.
- [73] T. MOSELHY AND Y. MARZOUK, *Bayesian inference with optimal maps*, Journal of Computational Physics, 231 (2012), pp. 7815–7850.
- [74] J. B. NAGEL AND B. SUDRET, *Spectral likelihood expansions for Bayesian inference*, Journal of Computational Physics, 309 (2016), pp. 267–294.
- [75] D. P. O’LEARY, *The block conjugate gradient algorithm and related methods*, Linear Algebra and its Applications, 29 (1980), pp. 293–322.
- [76] C. C. PAIGE, *Computational variants of the Lanczos method for the eigenproblem*, IMA Journal of Applied Mathematics, 10 (1972), pp. 373–381.
- [77] C. C. PAIGE AND M. A. SAUNDERS, *Towards a generalized singular value decomposition*, SIAM Journal on Numerical Analysis, 18 (1981), pp. 398–405.
- [78] ———, *LSQR: An algorithm for sparse linear equations and sparse least squares*, ACM Transactions on Mathematical Software (TOMS), 8 (1982), pp. 43–71.
- [79] L. PARDO, *Statistical Inference Based on Divergence Measures*, CRC Press, 2005.
- [80] A. PARKER AND C. FOX, *Sampling Gaussian distributions in Krylov spaces with conjugate gradients*, SIAM Journal on Scientific Computing, 34 (2012), pp. B312–B334.
- [81] X. PENNEC, P. FILLARD, AND N. AYACHE, *A Riemannian framework for tensor computing*, International Journal of Computer Vision, 66 (2006), pp. 41–66.
- [82] N. PETRA, J. MARTIN, G. STADLER, AND O. GHATTAS, *A computational framework for infinite-dimensional Bayesian inverse problems, Part II: stochastic Newton MCMC with application to ice sheet flow inverse problems*, SIAM Journal on Scientific Computing, 36 (2014), pp. A1525–A1555.
- [83] A. QUARTERONI AND A. VALLI, *Numerical Approximation of Partial Differential Equations*, vol. 23, Springer Science & Business Media, 2008.
- [84] C. R. RAO, *Information and the accuracy attainable in the estimation of statistical parameters*, Bulletin of the Calcutta Mathematical Society, 37 (1945), pp. 81–89.

- [85] ———, *On the distance between two populations*, *Sankhya*, 9 (1949), pp. 246–248.
- [86] ———, *Differential metrics in probability spaces*, *Differential geometry in statistical inference*, 10 (1987), pp. 217–240.
- [87] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, SIAM, 2003.
- [88] A. K. SAIBABA, J. LEE, AND P. K. KITANIDIS, *Randomized algorithms for generalized Hermitian eigenvalue problems with application to computing Karhunen–Loève expansion*, *Numerical Linear Algebra with Applications*, (2015).
- [89] A. H. SAMEH AND J. A. WISNIEWSKI, *A trace minimization algorithm for the generalized eigenvalue problem*, *SIAM Journal on Numerical Analysis*, 19 (1982), pp. 1243–1259.
- [90] C. SCHILLINGS AND C. SCHWAB, *Sparse, adaptive Smolyak quadratures for Bayesian inverse problems*, *Inverse Problems*, 29 (2013), p. 065011.
- [91] ———, *Scaling limits in computational Bayesian inversion*, Research Report No. 2014-26, ETH-Zürich, (2014).
- [92] M. K. SCHNEIDER AND A. S. WILLSKY, *A Krylov subspace method for covariance approximation and simulation of random processes and fields*, *Multidimensional Systems and Signal Processing*, 14 (2003), pp. 295–318.
- [93] L. T. SKOVGAARD, *A Riemannian geometry of the multivariate normal model*, *Scandinavian Journal of Statistics*, (1984), pp. 211–223.
- [94] S. T. SMITH, *Covariance, subspace, and intrinsic Cramér-Rao bounds*, *Signal Processing, IEEE Transactions on*, 53 (2005), pp. 1610–1630.
- [95] A. SOLONEN, H. HAARIO, J. HAKKARAINEN, H. AUVINEN, I. AMOUR, AND T. KAURANNE, *Variational ensemble Kalman filtering using limited memory BFGS*, *Electronic Transactions on Numerical Analysis*, 39 (2012), pp. 271–285.
- [96] S. SOMMER, F. LAUZE, S. HAUBERG, AND M. NIELSEN, *Manifold valued statistics, exact principal geodesic analysis and the effect of linear approximations*, in *Computer Vision–ECCV 2010*, Springer, 2010, pp. 43–56.
- [97] A. SPANTINI, A. SOLONEN, T. CUI, J. MARTIN, L. TENORIO, AND Y. MARZOUK, *Optimal low-rank approximations of Bayesian linear inverse problems*, *SIAM Journal on Scientific Computing*, 37 (2015), pp. A2451–A2487.
- [98] A. M. STUART, *Inverse problems: a Bayesian perspective*, *Acta Numerica*, 19 (2010), pp. 451–559.
- [99] C. F. VAN LOAN, *Generalizing the singular value decomposition*, *SIAM Journal on Numerical Analysis*, 13 (1976), pp. 76–83.
- [100] A. T. A. WOOD AND G. CHAN, *Simulation of stationary Gaussian processes in $[0, 1]^d$* , *Journal of Computational and Graphical Statistics*, 3 (1994), pp. 409–432.
- [101] Y. YUE AND P. L. SPECKMAN, *Nonstationary spatial Gaussian Markov random fields*, *Journal of Computational and Graphical Statistics*, 19 (2010), pp. 96–116.