

Optimal projection of observations in a Bayesian setting

Loïc Giraldi^{a,*}, Olivier P. Le Maître^b, Ibrahim Hoteit^c, Omar M. Knio^a

^a*Division of Computer, Electrical and Mathematical Sciences and Engineering, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia*

^b*LIMSI, CNRS, Université Paris-Saclay, France*

^c*Division of Physical Sciences and Engineering, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia*

Abstract

Optimal dimensionality reduction methods are proposed for the Bayesian inference of a Gaussian linear model with additive noise in presence of overabundant data. Three different optimal projections of the observations are proposed based on information theory: the projection that minimizes the Kullback-Leibler divergence between the posterior distributions of the original and the projected models, the one that minimizes the expected Kullback-Leibler divergence between the same distributions, and the one that maximizes the mutual information between the parameter of interest and the projected observations. The first two optimization problems are formulated as the determination of an optimal subspace and therefore the solution is computed using Riemannian optimization algorithms on the Grassmann manifold. Regarding the maximization of the mutual information, it is shown that there exists an optimal subspace that minimizes the entropy of the posterior distribution of the reduced model; a basis of the subspace can be computed as the solution to a generalized eigenvalue problem; an a priori error estimate on the mutual information is available for this particular solution; and that the dimensionality of the subspace to exactly conserve the mutual information between the input and the output of the models is less than the number of parameters to be inferred. Numerical applications to linear and nonlinear models are used to assess the efficiency of the proposed approaches, and to highlight their advantages compared to standard approaches based on the principal component analysis of the observations.

Keywords: Optimal dimensionality reduction, Optimal data reduction, Gaussian linear model, Information theory

2000 MSC: 65C60, 62F15, 62B10

1. Introduction

We consider the problem of Bayesian inference in the case of overabundant data. The goal is to compute an optimal approximation of the posterior distribution by projection of the observations. These projections are computed solving the following optimization problem

$$\min_V \mathcal{J} (P(X | Y = y), P(X | W = V^T y)),$$

where X is the inferred parameter, Y the observations of the Bayesian model with values in \mathbb{R}^n , $y \in \mathbb{R}^n$ the data, $W = V^T Y$ the reduced observations with values in \mathbb{R}^r , $V \in \mathbb{R}^{n \times r}$ the deterministic matrix defining the projection from the full to the reduced observations, and \mathcal{J} a functional defining the optimality criterion.

*Correspondence to: 4700 King Abdullah University of Science and Technology (KAUST), Thuwal, 23955–6900, Kingdom of Saudi Arabia. Tel.: +966 12 808 0214.

Email address: loic.giraldi@kaust.edu.sa (Loïc Giraldi)

In the literature, the most popular dimensionality reduction techniques result from the optimal approximation of the observations Y with respect to the L^2 norm defined by

$$\|Y\|_{L^2}^2 = \mathbb{E}(Y^T Y),$$

with \mathbb{E} the expectation operator. A best low-rank approximation of Y with respect to the L^2 norm is computed by a singular value decomposition. Then, with V being the matrix composed of the dominant left eigenvectors of Y , the approximation is given by $Y \approx VW$ where $W = V^T Y$. The reader could refer to [15, Section 4.4.3] for the presentation of the singular value decomposition in a general case. When applied to the centered random vector $Y - \mathbb{E}(Y)$, this decomposition is also called truncated Karhunen-Loève expansion [20, 21, 22], or principal component analysis [17, 19, 26].

Let Ω denote a sample space. When considering the random vector as a map $\Omega \rightarrow \mathbb{R}^n$, an approximation of the observations Y can be used to define the matrix V based on the L^∞ norm defined by

$$\|Y\|_{L^\infty} = \sup_{\omega \in \Omega} \max_{1 \leq i \leq n} |Y_i(\omega)|.$$

The empirical interpolation method [4] provides an approximation of the form $\omega \mapsto VW(\omega)$ of a parametric vector $\omega \mapsto Y(\omega)$ based on this supremum norm. Given that computing the supremum is not tractable, the sample space is restricted to a finite sample $\Omega_N = \{\omega_i\}_{i=1}^N \subset \Omega$ and the norm is approximated by

$$\|Y\|_{L^\infty} \approx \max_{\omega \in \Omega_N} \max_{1 \leq i \leq n} |Y_i(\omega)|.$$

The interpolation is then defined to be exact on a subsample $\{\omega_j^*\}_{j=1}^r \subset \Omega_N$, and on a restricted number of indices $\{i_j\}_{j=1}^r \subset \{1, \dots, n\}$. The interpolation points $((\omega_j^*, i_j))_{j=1}^r$ are selected in a greedy fashion using the supremum norm of the error, such that

$$(\omega_r^*, i_r) = \arg \max_{\omega \in \Omega_N} \max_{1 \leq i \leq n} \left| Y_i - \sum_{j=1}^{r-1} V_{ij} W_j \right|.$$

The approach reduces to the interpolation of the matrix $M_{ij} = Y_i(\omega_j)$ and is therefore very close to the cross approximation method [5] for the low-rank interpolation of a matrix. A detailed comparison between the singular value decomposition, the empirical interpolation method and the cross approximation is provided in [6]. A weighted variant of the empirical interpolation method was introduced in [8] in order to take into account probability measures. Given a positive weight $w : \Omega \rightarrow \mathbb{R}_+$, the supremum norm is modified such that

$$\|Y\|_{L^\infty}^w = \sup_{\omega \in \Omega} \max_{1 \leq i \leq n} w(\omega) |Y_i(\omega)|,$$

yielding different interpolation points. While the approximations based on the L^2 and L^∞ norms are widely used, they are optimal with respect to the output of the model and are not directly related to the distribution of the parameter to infer. An exception to this rule is [13] where a weighted singular value decomposition was used to accommodate uniform priors.

Geppert et al. [12] propose to reduce the number of observations in a Bayesian regression framework using random projections. The methodology relies on ε -subspace embeddings that are particular maps of the form $\Pi \in \mathbb{R}^{r \times n}$ satisfying

$$(1 - \varepsilon) \|Bx\|^2 \leq \|\Pi Bx\|^2 \leq (1 + \varepsilon) \|Bx\|^2,$$

for a particular $B \in \mathbb{R}^{n \times q}$, and for any $x \in \mathbb{R}^q$ with a probability $1 - \alpha$. In order to obtain an error on the posterior distribution of order ε in terms of the Wasserstein metric, it is shown that the number of observations required is $\mathcal{O}((q + \log(1/\alpha))/\varepsilon^2)$, $\mathcal{O}(q \log(q/\alpha)/\varepsilon^2)$ or $\mathcal{O}(q^2/(\alpha\varepsilon^2))$, depending on the embedding. Even though the dimension r can be drastically smaller than the original number of observations, n , it can still be relatively large for a small risk α and a small error ε .

Another related technique is introduced in a series of papers [9, 28, 29]. Given a Gaussian linear model, the goal is to directly compute an approximate posterior covariance matrix as a low-rank update of the prior covariance. Given a particular loss function depending only on the covariance, it is shown that an optimal low-rank update can be derived from a generalized eigenvalue problem. The resulting distribution is then optimal in terms of the Hellinger distance and Kullback-Leibler divergence under the assumption that the mean is exactly recovered. An optimal mean is also derived as a linear projection of the data by minimizing the Bayes risk defined as the expected Mahalanobis distance between the parameter of interest X and the approximate mean. Regarding this methodology, two disadvantages are notable: there is an inconsistency between the optimality criteria of the mean and the covariance, and the computation of the different matrices requires the inversion of the covariance matrix of the noise, which is of large dimension $n \times n$.

In this work, the approximate mean and covariance are defined similarly, namely as an affine function of the data and as a low-rank update of the prior covariance, respectively. However, they result from the optimal projections of the statistical model using criteria from information theory, namely the Kullback-Leibler divergence, the expected Kullback-Leibler divergence, and the mutual information, which is the first contribution of this paper. The second contribution concerns the choice of the practical numerical algorithm for the minimization of the expected Kullback-Leibler divergence between the posterior distributions of the full and reduced model. It is moreover shown that a solution to the optimization problem defined as the maximization of the mutual information $\mathcal{I}(X, W)$ between the parameter X and the reduced observations W is given by the solution of a generalized eigenvalue problem that does not require the inversion of a large matrix. We can moreover estimate the loss $\mathcal{I}(X, Y) - \mathcal{I}(X, W)$ and show that no more than q projections are required to recover the full mutual information, where q is the size of X . The last contribution of this work concerns the illustration of the method on linear and nonlinear examples.

This paper is organized as follows. In Section 2, the full and reduced linear models are presented in the Gaussian case, as well as other required definitions. The posterior distributions are then provided explicitly in closed form. Section 3 introduces the three different optimization problems that are used to define the alternative optimal projections of the observations. The analysis of the corresponding optimal subspace and the numerical algorithms for their computation are then provided. The methodologies are finally applied and illustrated to a Bayesian linear regression problem in Section 4 and to a nonlinear problem in Section 5. Major conclusions are summarized in Section 5.4.

2. Linear Gaussian model

2.1. Models

We consider an abstract probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is the sample space, \mathcal{F} is a σ -algebra and \mathbb{P} a probability measure. Given an \mathbb{R}^n -valued random vector Z , we denote by $P(Z)$ the pushforward probability measure such that $P(A) = \mathbb{P}(Z^{-1}(A))$ for any set A in the Borel algebra of \mathbb{R}^n , and f_Z the probability density function defined with respect to the Lebesgue measure.

We consider the following linear model

$$Y = BX + E, \tag{1}$$

where $B \in \mathbb{R}^{n \times q}$ is the design matrix, X is the random parameter we want to infer and E is the random noise. The random vector X (resp. E) is supposed to follow the multivariate normal distribution $\mathcal{N}(m_X, C_X)$ (resp. $\mathcal{N}(m_E, C_E)$) with mean $m_X \in \mathbb{R}^q$ (resp. $m_E \in \mathbb{R}^n$) and covariance $C_X \in \mathbb{R}^{q \times q}$ (resp. $C_E \in \mathbb{R}^{n \times n}$). The input parameter X and the noise E are assumed to be independent.

In order to compress the amount of data used for the inference, we introduce $V = (v_i)_{i=1}^r \in \mathbb{R}^{n \times r}$, a reduced basis of observations. In the following, the term reduced space may be used for V , as we look for the projection of the observations on the space spanned by the columns of V . The linear model expressed in the reduced coordinates is therefore

$$W = V^T BX + V^T E, \tag{2}$$

and the reduction is efficient if, for $r \ll n$, the posterior distribution $P(X | W)$ is close to $P(X | Y)$ in some sense defined in Section 3. Our main goal is to compute a suitable matrix V which satisfies this condition.

In order to subsequently apply the different methodologies to nonlinear models of the form

$$Y = A(X) + E,$$

it is beneficial to consider the random vector $A(X) = BX$. In the linear case, this random vector follows the distribution $\mathcal{N}(m_A, C_A)$ where

$$m_A = Bm_X \quad \text{and} \quad C_A = BC_XB^T.$$

We also denote by C_{AX} the covariance between A and X , that is

$$C_{AX} = \mathbb{E} \left((A(X) - m_A)(X - m_X)^T \right) = BC_X.$$

Given the structure of the problem, the random vectors Y and W are also distributed according to the multivariate normal distribution $\mathcal{N}(m_Y, C_Y)$ and $\mathcal{N}(m_W, C_W)$, respectively, with $m_Y = m_A + m_E$, $C_Y = C_A + C_E$, $m_W = V^T m_Y$, and $C_W = V^T C_Y V$. The extension of the approach to nonlinear problems is based on the three quantities m_A , C_A and C_{AX} .

2.2. Posterior distributions

Given the linear Gaussian structure of Equation (1), an observation y of Y , and a reduced basis V , the posterior distributions $P(X | Y = y)$ and $P(X | W = V^T y)$ can be analytically derived. The result is summarized in Proposition 2.1.

Proposition 2.1. *The posterior distribution $P(X | Y = y)$ follows the multivariate normal distribution $\mathcal{N}(m_\star, C_\star)$, where*

$$C_\star = C_X(C_X + C_{AX}^T C_E^{-1} C_{AX})^{-1} C_X = C_X - C_{AX}^T C_Y^{-1} C_{AX}, \quad (3)$$

and

$$m_\star = G_\star(y - m_Y) + h_\star, \quad (4)$$

with $G_\star = C_{AX}^T C_Y^{-1}$ and $h_\star = C_\star C_X^{-1} m_X + G_\star m_A$.

Regarding the posterior distribution of the reduced model, if the matrix $V \in \mathbb{R}^{n \times r}$ is full-rank, the distribution $P(X | W = V^T y)$ follows the multivariate normal distribution $\mathcal{N}(m_V, C_V)$, where

$$\begin{aligned} C_V &= C_X \left(C_X + C_{AX}^T V (V^T C_E V)^{-1} V^T C_{AX} \right)^{-1} C_X \\ &= C_X - C_{AX}^T V (V^T C_Y V)^{-1} V^T C_{AX}, \end{aligned} \quad (5)$$

and

$$m_V = G_V V^T (y - m_Y) + h_V, \quad (6)$$

with

$$G_V = C_{AX}^T V (V^T C_Y V)^{-1},$$

and

$$h_V = C_V C_X^{-1} m_X + G_V V^T m_A.$$

Proof. See Appendix [Appendix A](#). □

Regarding Proposition 2.1, we can first notice that the two expressions

$$C_\star = C_X(C_X + C_{AX}^T C_E^{-1} C_{AX})^{-1} C_X$$

$$\text{and } C_V = C_X \left(C_X + C_{AX}^T V (V^T C_E V)^{-1} V^T C_{AX} \right)^{-1} C_X$$

show that the matrices C_\star and C_V are always symmetric positive definite, even for a nonlinear model. In the following, we denote by GL_r the set of invertible matrices in $\mathbb{R}^{r \times r}$, and obtain an invariance property expressed in Proposition 2.2.

Proposition 2.2. *For all matrices $M \in \text{GL}_r$, we have*

$$m_{VM} = m_V \quad \text{and} \quad C_{VM} = C_V.$$

Therefore, the posterior distribution $P(X | W = V^T y) \sim \mathcal{N}(m_V, C_V)$ is invariant under invertible linear transformation of the matrix V on the right.

Proof. See Appendix AppendixB. □

In practice, this proposition means that V is less important than $\text{range}(V)$ in the determination of the posterior distribution. Indeed, rescaling, rotating or permuting the observations in Equation (2) does not affect the posterior distribution $P(X | W = V^T y)$.

Formally, the Grassmann manifold $\text{Gr}(r, n)$ defined as the set of r dimensional subspace of \mathbb{R}^n is therefore the set of interest to determine the optimal reduced observations. In this work, we identify $\text{Gr}(r, n)$ with the quotient manifold $\mathbb{R}_*^{n \times r} / \text{GL}_r$ following [2], where $\mathbb{R}_*^{n \times r}$ is the set of full rank matrices of $\mathbb{R}^{n \times r}$ and the quotient space is defined by

$$\text{Gr}(r, n) = \mathbb{R}_*^{n \times r} / \text{GL}_r = \{[V]; V \in \mathbb{R}_*^{n \times r}\}, \text{ where } [V] = \{VM; M \in \text{GL}_r\}.$$

Finally, Proposition (2.2) means that it is more important to identify the equivalence class $[V] \in \mathbb{R}_*^{n \times r}$ than a particular matrix $V \in \mathbb{R}_*^{n \times r}$.

The next section presents the different proposed optimization problems, where the Grassmann manifold $\text{Gr}(r, n)$ has an important role.

3. Optimality criteria for the definition of the reduced basis

3.1. Kullback-Leibler divergence minimization

Given two distributions $P(Z_0)$ and $P(Z_1)$, the Kullback-Leibler divergence between them is defined by

$$D_{\text{KL}}(P(Z_0) \parallel P(Z_1)) = \mathbb{E}_{Z_0} \left(\log \frac{f_{Z_0}}{f_{Z_1}} \right). \quad (7)$$

This divergence quantifies the ‘‘information lost when $[P(Z_1)]$ is used to approximate $[P(Z_0)]$ ’’ according to [7, Section 2.1]. The Kullback-Leibler divergence is always positive and null if and only if the two distributions are identical, therefore defining a generalized distance between distributions.

This interpretation of the Kullback-Leibler divergence leads us to consider the following functional $\mathcal{J}_0 : \mathbb{R}_*^{n \times r} \rightarrow \mathbb{R}$ defined by

$$\mathcal{J}_0(V) = D_{\text{KL}}(P(X | Y = y) \parallel P(X | W = V^T y)).$$

The domain definition of \mathcal{J}_0 must be restricted to the set of full rank matrices $\mathbb{R}_*^{n \times r}$ in order to comply with Proposition 2.1 characterizing the posterior distributions. Given that we are working with Gaussian distributions, the computation of the Kullback-Leibler divergence is always well-posed (i.e. f_{Z_1} is always strictly positive in Equation (7)). The general expression of the Kullback-Leibler divergence between two Gaussian distribution is given in Proposition 3.1.

Proposition 3.1. Assuming that $Z_0 \sim \mathcal{N}(m_0, C_0)$ and $Z_1 \sim \mathcal{N}(m_1, C_1)$ are \mathbb{R}^q -valued random variables, the Kullback-Leibler divergence between $P(Z_0)$ and $P(Z_1)$ is expressed by

$$D_{\text{KL}}(P(Z_0) \parallel P(Z_1)) = \frac{1}{2} (D_{\ell d}(C_0, C_1) + D_{C_1}(m_0, m_1)),$$

where $D_{\ell d}(C_0, C_1)$ is the Bregman log det divergence between C_0 and C_1 defined by

$$D_{\ell d}(C_0, C_1) = \text{tr}(C_0 C_1^{-1}) - \log \det(C_0 C_1^{-1}) - q,$$

and $D_{C_1}(m_0, m_1)$ is the Mahalanobis divergence defined by

$$D_{C_1}(m_0, m_1) = (m_0 - m_1)^T C_1^{-1} (m_0 - m_1).$$

Proof. See Appendix [Appendix C](#). □

As a consequence of Proposition 3.1, the functional \mathcal{J}_0 has a closed form depending on m_* , C_* , m_V and C_V :

$$\begin{aligned} \mathcal{J}_0(V) &= D_{\text{KL}}(P(X | Y = y) \parallel P(X | W = V^T y)) \\ &= \frac{1}{2} (D_{\ell d}(C_*, C_V) + D_{C_V}(m_*, m_V)) \\ &= \frac{1}{2} (\text{tr}(C_* C_V^{-1}) - \log \det(C_* C_V^{-1}) - q \\ &\quad + (m_* - m_V)^T C_V^{-1} (m_* - m_V)). \end{aligned} \tag{8}$$

Given Proposition 2.2, for all $M \in \text{GL}_r$, we have $\mathcal{J}_0(VM) = \mathcal{J}_0(V)$. It means that we are in fact interested in the map defined on $\text{Gr}(r, n)$ by $[V] \mapsto \mathcal{J}_0(V)$. The minimization problem of interest is therefore

$$\min_{[V] \in \text{Gr}(r, n)} D_{\text{KL}}(P(X | Y = y) \parallel P(X | W = V^T y)). \tag{9}$$

We can show that the following result holds.

Theorem 3.2. *There exists a solution to Problem (9).*

Proof. See Appendix [Appendix D](#). □

Note that the minimization of the Kullback-Leibler divergence in Problem (9) results in an *a posteriori* reduction in the sense that a realization y of Y is required to evaluate the cost function. In the following, other functionals are proposed that circumvent this issue.

3.2. Expected Kullback-Leibler divergence minimization

The first possibility to remove the dependence on the data is to work on the expected Kullback-Leibler divergence with respect to the observation, where the measurement Y is treated as a random variable. Similarly to Section 3.1, let $\mathcal{J}_1 : \mathbb{R}_*^{n \times r} \rightarrow \mathbb{R}$ be defined by

$$\mathcal{J}_1(V) = \mathbb{E}_Y (D_{\text{KL}}(P(X | Y) \parallel P(X | W = V^T Y))).$$

The expected Kullback-Leibler divergence admits a closed form as well, presented in the next proposition.

Proposition 3.3. *We have the following equality*

$$\mathcal{J}_1(V) = \frac{1}{2} (D_{\ell d}(C_*, C_V) + \mathbb{E}_Y (D_{C_V}(m_*, m_V))), \tag{10}$$

where

$$\begin{aligned} \mathbb{E}_Y (D_{C_V}(m_*, m_V)) &= \\ &\text{tr} \left(C_V^{-1} (G_* - G_V V^T) C_Y (G_* - G_V V^T)^T \right) + (h_* - h_V) C_V^{-1} (h_* - h_V). \end{aligned}$$

Proof. See Appendix [AppendixE](#). □

Using Proposition [2.2](#) and Equation [\(10\)](#), we can show that $\mathcal{J}_1(V) = \mathcal{J}_1(VM)$ for all matrices $M \in \text{GL}_r$. We are therefore interested in finding the optimal equivalence class $[V]$ and solving the minimization problem

$$\min_{[V] \in \text{Gr}(r,n)} \mathbb{E}_Y \left(\text{D}_{\text{KL}} \left(P(X | Y) \parallel P(X | W = V^T Y) \right) \right). \quad (11)$$

As in Section [3.1](#), we can prove the following result.

Theorem 3.4. *There exists a solution to Problem [\(11\)](#).*

Proof. The proof is similar to the one in Appendix [AppendixD](#), replacing \mathcal{J}_0 by \mathcal{J}_1 . □

Remark 3.5. *The minimization of the log det divergence $\text{D}_{\text{ld}}(C_\star, C_V)$ has also been considered, being the data-free part of the Kullback-Leibler divergence. It has been ignored in the paper as it did not bring additional insights on the optimal construction of the reduced observations.*

3.3. Mutual information maximization and entropy minimization

In this section the Shannon entropy and the mutual information are introduced. The entropy $H(Z)$ (sometimes denoted $H(P(Z))$) quantifies the uncertainty or the amount of information contained in a random variable $Z \sim P(Z)$ and is defined by

$$H(Z) = \mathbb{E}_Z(-\log(f_Z(Z))).$$

The mutual information $\mathcal{I}(Z_0, Z_1)$ between the two random variables $Z_0 \sim P_0(Z_0)$ and $Z_1 \sim P_1(Z_1)$ is a measure of the information that Z_0 contains about Z_1 , and is defined by

$$\mathcal{I}(Z_0, Z_1) = H(Z_0) + H(Z_1) - H(Z_0, Z_1),$$

where $H(Z_0, Z_1)$ is the entropy of the *joint distribution* of $Z = (Z_0, Z_1)$. From this definition, it is clear that the mutual information is symmetric.

The new definition of the reduced basis, introduced in this section, is related to the definition of the mutual information. We would like the reduced observations W to contain as much information as possible about X . We therefore consider the following maximization problem

$$\max_{V \in \mathbb{R}_*^{n \times r}} \mathcal{I}(W, X). \quad (12)$$

Note that another expression of the mutual information is

$$\mathcal{I}(W, X) = \mathbb{E}_W(\text{D}_{\text{KL}}(P(X | W) \parallel P(X))),$$

showing that this strategy aims at maximizing the expected information gain between the prior and the posterior distributions of X .

The optimization problem in Equation [\(12\)](#) admits a simple solution presented in Theorem [3.6](#). Moreover, we shall show that the maximization of the mutual information is equivalent to the minimization of the entropy of the posterior distribution $P(X | W = V^T y)$.

Theorem 3.6. *The following equalities hold*

$$\begin{aligned} \mathcal{I}(W, X) &= \frac{1}{2} \log \det \left((V^T C_Y V) (V^T C_E V)^{-1} \right) \\ \text{and } H(P(X | W = V^T y)) &= -\mathcal{I}(W, X) + \frac{1}{2} \log \det C_X + \frac{q}{2} \log(2\pi e). \end{aligned}$$

As a consequence, the maximization of the mutual information $\mathcal{I}(W, X)$ and the minimization of the entropy of the posterior distribution $H(P(X | W = V^T y))$ with respect to V admit the same solutions for any realization y of Y . We have the equality

$$\max_{V \in \mathbb{R}_*^{n \times r}} \mathcal{I}(W, X) = \frac{1}{2} \sum_{i=1}^r \log \lambda_i, \quad (13)$$

where $(\lambda_i)_{i=1}^r$ are the r dominant eigenvalues of the following generalized eigenvector problem

$$C_Y v = \lambda C_E v, \quad \lambda \in \mathbb{R}, \quad v \in \mathbb{R}^n. \quad (14)$$

A solution to the optimization Problem (13) is given by the matrix V with columns being eigenvectors $(v_i)_{i=1}^r$ associated to the dominant eigenvalues of Problem (14).

Proof. See Appendix [AppendixF](#). □

Several remarks follow this result. First, the map $\mathcal{J}_2 : V \mapsto \mathcal{I}(W, X)$ is also invariant under the transformation $\mathcal{J}_2(VM) = \mathcal{J}_2(V)$ for any invertible matrix $M \in \text{GL}_r$, and therefore the solution should be searched in the Grassmann manifold. However in the present case, a particular solution admits a simple characterization.

The generalized eigenvalue problem in Equation (14) is used to define the optimal mean in [29] to minimize the Bayes risk. It is however unclear how this optimal mean is related to the mean defined in Equation (6). Note moreover that the computation of the optimal mean from [29] requires the inversion of the matrix C_E , which is not needed in the presently developed approach.

Another interesting feature of Theorem 3.6 is that it provides an a priori estimate for the reduction error, based on the mutual information, summarized in the following Corollary 3.7.

Corollary 3.7. *Let $V \in \mathbb{R}^{n \times r}$ be a particular solution to Problem (13) and W be the reduced model associated to V , and $(\lambda_i)_{i=1}^n$ be the eigenvalues associated to Problem (14) sorted in a decreasing order. Then, the relative error on the mutual information is given by*

$$\frac{\mathcal{I}(Y, X) - \mathcal{I}(W, X)}{\mathcal{I}(Y, X)} = 1 - \frac{\sum_{i=1}^r \log \lambda_i}{\sum_{i=1}^n \log \lambda_i}.$$

In fact, the entire spectrum of C_Y is not required to estimate the error. In practice, we only need to determine the eigenvalues ν_i associated to the following problem

$$C_A v = \nu C_E v, \quad \nu \geq 0, \quad v \in \mathbb{R}^n. \quad (15)$$

If λ is an eigenvalue associated to Problem (14), then $\nu = \lambda - 1$ is an eigenvalue associated to Problem (15). Considering Problem (15) is beneficial in practice because C_A is at most a rank- q matrix. This remark leads to the following important result on the number of required projections to get the same mutual information between the observations and the parameter of interest, for the full and the reduced model.

Corollary 3.8. *Let $V \in \mathbb{R}^{n \times r}$ be a solution to Problem (13) and W be the reduced model associated to V , and $(\lambda_i)_{i=1}^n$ be the eigenvalues associated to Problem (14), sorted in a decreasing order. Let $(\nu_i)_{i=1}^n$ be the eigenvalues associated to Problem (15) (i.e. $\lambda_i = 1 + \nu_i$), and let $m \leq q \ll n$ be the rank of $B \in \mathbb{R}^{n \times q}$ (see Equation (1)). Then $C_A = BC_X B^T$ is a rank- m matrix, and the relative error on the mutual information is given by*

$$\frac{\mathcal{I}(Y, X) - \mathcal{I}(W, X)}{\mathcal{I}(Y, X)} = 1 - \frac{\sum_{i=1}^r \log(1 + \nu_i)}{\sum_{i=1}^m \log(1 + \nu_i)}. \quad (16)$$

The condition $r \geq m$ implies that $\mathcal{I}(W, X) = \mathcal{I}(Y, X)$ and the mutual informations between the observations and the parameter of interest are the same for the full and the reduced model. In particular, the condition is satisfied for $m = q$.

As a side note, the principal component analysis of the observations Y yields a reduced basis defined as the dominant eigenvectors of C_Y . Therefore, the resulting reduced space is optimal with respect to the mutual information in the case of a white noise, i.e. $C_E = \sigma^2 \mathbf{I}_n$. However, denoting $(\chi_i)_{i=1}^n$ the eigenvalues of C_Y sorted in a decreasing order, the corresponding estimate of the relative reduction error on the mutual information is given by

$$\frac{\mathcal{I}(Y, X) - \mathcal{I}(W, X)}{\mathcal{I}(Y, X)} = 1 - \frac{\sum_{i=1}^r \log\left(\frac{\chi_i}{\sigma^2}\right)}{\sum_{i=1}^n \log\left(\frac{\chi_i}{\sigma^2}\right)}.$$

Note that the usual error criteria used in the principal component analysis between the random variable Y and its rank- r truncated version Y_r controls the L_2 norm and is given by (see e.g. [6, Proposition 2.1])

$$\frac{\mathbb{E}\left(\|Y - Y_r\|_2^2\right)}{\mathbb{E}\left(\|Y\|_2^2\right)} = 1 - \frac{\sum_{i=1}^r \chi_i^2}{\sum_{i=1}^n \chi_i^2}.$$

3.4. Numerical solution to the optimization problems

For any functional $\mathcal{J} \in \{\mathcal{J}_0, \mathcal{J}_1, \mathcal{J}_2\}$ involved in the optimization problems presented in Section 3, the following property holds

$$\mathcal{J}(V) = \mathcal{J}(VM), \quad \forall M \in \text{GL}_r.$$

As a consequence of this invariance, there exists an infinite number of solutions to the optimization problems and the Hessian of the functional $\nabla^2 \mathcal{J}$ is ill-conditioned in a neighbourhood of a solution. The main consequence is that we cannot use a standard Newton algorithm to solve these nonlinear problems without regularizing the optimization problem first.

In order to circumvent this issue, we consider here the restriction of the optimization problem to the Grassmann manifold $\text{Gr}(r, n)$, replacing the search for a $n \times r$ matrix by the search of a r -dimensional linear subspace of $\mathbb{R}^{n \times n}$. In order to solve Problems (9) and (11) we choose to use a specific algorithm exploiting the smooth manifold structure of $\text{Gr}(r, n)$, that is the Riemannian trust-region algorithm [1] implemented in the *Pymanopt* library [32]. The derivatives of the cost functions are computed by automatic differentiation with the *autograd* library [24].

Given a finite dimensional vector space \mathcal{V} equipped with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{V}}$ and the associated norm $\|\cdot\|_{\mathcal{V}}$, the trust-region algorithm consists in correcting the current iterate $V \in \mathcal{V}$ with $W \in \mathcal{V}$ using a quadratic approximation of the functional \mathcal{J} . W is defined as the solution to

$$\min_{W \in \mathcal{V}} m(W) = \mathcal{J}(V) + \langle \nabla \mathcal{J}(V), W \rangle_{\mathcal{V}} + \frac{1}{2} \langle \nabla^2 \mathcal{J}(V) W, W \rangle_{\mathcal{V}},$$

such that $\|W\|_{\mathcal{V}}^2 \leq \Delta^2$, where $\nabla \mathcal{J}$ (resp. $\nabla^2 \mathcal{J}$) is the gradient (resp. Hessian) of \mathcal{J} . The trust-region radius Δ is adapted at each iteration of the algorithm based on the quantity

$$\rho = \frac{\mathcal{J}(V) - \mathcal{J}(V + W)}{m(0) - m(W)}.$$

If ρ is close to 1, the quadratic approximation is good and the radius Δ can be expanded. Otherwise, Δ is shrunk.

The Riemannian version of the algorithm consists in considering the tangent space to the manifold for the search space, which is locally mapped to the smooth manifold. Formally, let \mathcal{M} be a smooth manifold equipped with the Riemannian metric $\langle \cdot, \cdot \rangle_{\mathcal{M}, V}$ and the associated norm $\|\cdot\|_{\mathcal{M}, V}$ defined on the tangent space $T_V \mathcal{M}$ to \mathcal{M} at V . We denote by $R_V : T_V \mathcal{M} \rightarrow \mathcal{M}$ a retraction which is a first-order approximation of the exponential map that maps locally the tangent space to the manifold. The retraction is such that $R_V(0) = V$. The correction is now defined as

$$\min_{W \in T_V \mathcal{M}} m(W) = \mathcal{J}(V) + \langle \nabla \mathcal{J}(V), W \rangle_{\mathcal{M}, V} + \frac{1}{2} \langle \nabla^2 \mathcal{J}(V) W, W \rangle_{\mathcal{M}, V}, \quad (17)$$

such that $\|W\|_{\mathcal{M}}^2 \leq \Delta^2$, where $\nabla \mathcal{J}$ (resp. $\nabla^2 \mathcal{J}$) is the Riemannian gradient (resp. Riemannian Hessian) of \mathcal{J} . The correction that belongs to the tangent space is mapped to the manifold using the retraction, such that the new iterate is defined by $R_V(W)$. The trust-region radius is now adapted according to the ratio

$$\rho = \frac{\mathcal{J}(V) - \mathcal{J}(R_V(W))}{m(0) - m(W)}.$$

The quadratic subproblem presented in Equation (17) is solved with a truncated conjugate gradient method. We refer the reader to [1] for an exhaustive description and analysis of the algorithm as well as its application on the Grassmann manifold.

4. Application to Bayesian linear regression

4.1. Inference problem

The goal of this section is to illustrate the results of Section 3, and numerically assess the methods in the case of a Bayesian polynomial regression. Given a uniformly distributed sample $(s_i)_{i=1}^n$ in $(-1, 1)$, we want to infer the random variable, X , from the following linear model

$$Y_i = \sum_{j=0}^{q-1} T_j(s_i) X_j + E(s_i), \quad \forall i \in \{1, \dots, n\}, \quad (18)$$

where T_j is the Chebyshev polynomial of the first kind [14] of order j and $q = 30$. The two moments of the prior distribution of $X \sim \mathcal{N}(m_X, C_X)$ are defined by

$$(m_X)_i = -1 + 2 \frac{i-1}{q-1}$$

and

$$(C_X)_{ij} = \sigma_X^2 \left(1 + \sqrt{1200} \frac{|i-j|}{(q-1)} \right) \exp \left(-\sqrt{1200} \frac{|i-j|}{(q-1)} \right),$$

with $\sigma_X = 1$. Note that the covariance C_X is a Matérn 3/2 covariance matrix, prescribing that polynomial coefficients associated to Chebyshev polynomials with distant degrees are less correlated than close ones. The noise E is a stationary Gaussian process with mean and covariance functions defined respectively by $\mu_E(s) = \cos(4\pi s)$ and

$$k_E(s, s') = \sigma_{E,1}^2 \exp \left(-\frac{|s-s'|}{\ell_E} \right) + \sigma_{E,2}^2 \delta(s-s'),$$

with $\sigma_{E,1} = 0.6$, $\ell_E = 0.05$, and $\sigma_{E,2} = 10^{-3}$. The model presented in Equation (18) is equivalent to the linear model from Equation (1) with $B_{ij} = T_{j-1}(s_i)$.

Figure 1 illustrates the data y used for the observations and the maximum a posteriori fit $Bm_* + m_E$. The rest of Section 4 is dedicated to the optimal estimation of the posterior distribution $P(X | Y = y)$ using $P(X | W = V^T y)$, where V has been computed according to the criteria introduced in Section 3.

4.2. Numerical results

We first consider three types of approaches for the computation of the reduced space V based on the principal component analysis. They are denoted PCA-A, PCA-Y and PCA-YN and are respectively computed as the dominant eigenvectors of the three following eigenvalue problems

$$C_A v = \lambda v, \quad C_Y v = \lambda v \quad \text{and} \quad C_Y C_E^{-1} v = \lambda v. \quad (19)$$

PCA-A corresponds to the principal component analysis of $A(X) = BX$, PCA-Y to the analysis of Y , and PCA-YN to the analysis of Y using the Mahalanobis distance, the metric induced by the inverse of the noise




Figure 1: Comparison between the observations y and the MAP estimate $Bm_* + m_E$.

covariance C_E^{-1} . The latter has been successfully used in a Bayesian inference context in [13], where the metric is directly involved in the posterior distribution due to uniform priors.

We denote by KLD (resp. EKLD, MI) the solutions obtained using the minimization of the Kullback-Leibler divergence (resp. minimization of the expected Kullback-Leibler divergence, maximization of the mutual information).

For a particular realization y , we compute the Kullback-Leibler divergence between the posterior distribution $P(X | W = V^T y)$ and $P(X | Y = y)$, and analyze its dependence on the dimension of the reduced space, r . The results are plotted in Figure 2 for the different dimensionality reduction methods. We conclude that the information theoretic based methods (KLD, EKLD, MI) with $r = q$ dimensions yield the exact posterior distribution within machine accuracy, and outperform the PCA-based approaches. Given that we are measuring the error using the Kullback-Leibler divergence, the KLD method performs better than the others. We can however note that the EKLD and MI techniques are robust to the realization y . We will observe below that the PCA methods require a dimension of the order of the total number of observations, n , to achieve a similar accuracy.

Figure 3 depicts the dependence the *expected* Kullback-Leibler divergence between the posterior distributions of the reduced and the full models on the dimension of the reduced spaces; plotted are results obtained using the different projection techniques. Similar to Figure 2, the information theoretic approaches converge to the posterior distribution with subspaces of dimension $r = q$, which is not the case for the PCA methods. We also note that even when the expected Kullback-Leibler divergence is used as error criterion, the EKLD method does not really improve the speed of convergence of the distributions compared to the other information theoretic approaches.

Figure 4 illustrates the relative error with respect to the dimension of the reduced space between the mutual information of:

- the observations and the parameter of interest, $\mathcal{I}(Y, X)$; and,
- the projected observations and the parameter of interest.

We are in fact looking at the criterion introduced in Corollaries 3.7 and 3.8. We note that for $r \geq q = 30$, all the information theoretic methods converge to the minimal value of the relative error. This behavior is predicted by Corollary 3.8 for the MI approach, as illustrated in the figure by the fact that the error




Figure 2: Kullback-Leibler divergence versus the dimension of the reduced space for the different numerical methods.

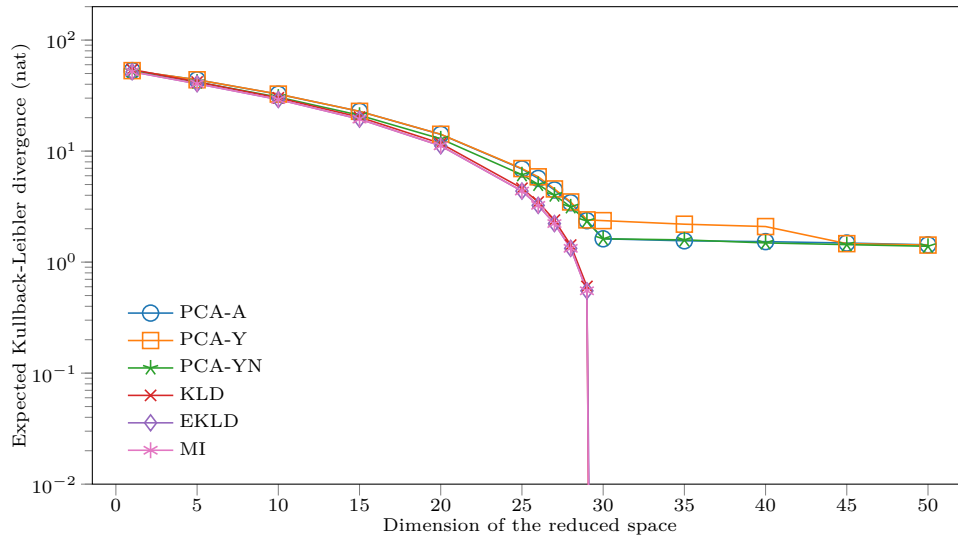


Figure 3: Expected Kullback-Leibler divergence versus the dimension of the reduced space for the different numerical methods.

estimator (16) overlaps with the error of the MI approach. Again, the PCA based methods perform poorly when compared to the information theoretic approaches.

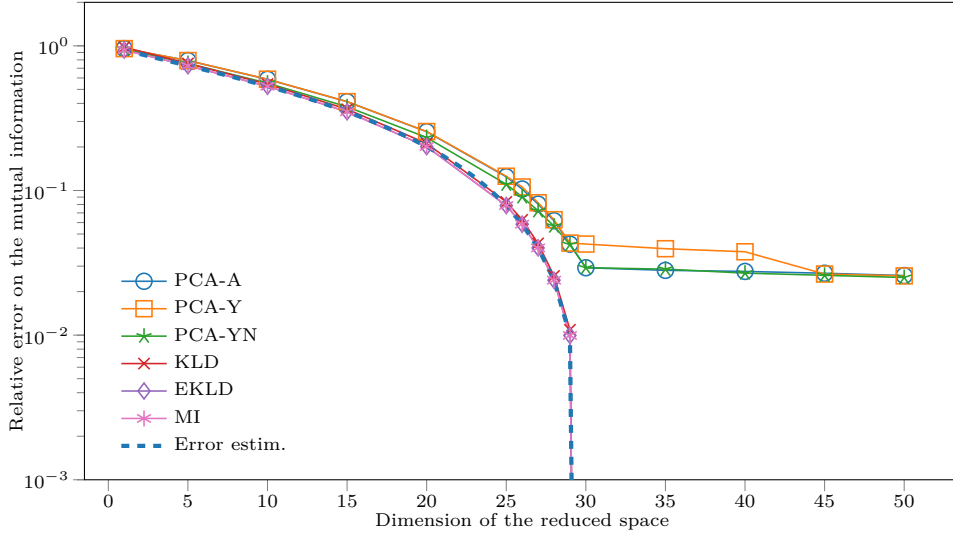


Figure 4: Relative error between the mutual information of the observations and the parameter of interest of the full model and the one of the reduced model and error estimator (16) versus the dimension of the reduced space for the different numerical methods.

In Figure 5, the different divergences and the absolute error on the entropy with respect to the dimension of the reduced space are illustrated for the PCA-based methods for larger values of the dimension and compared to the MI approach. The absolute error on the entropy is equivalent to the error on the mutual information up to a constant according to Theorem 3.6. One can see that the dimension of the reduced space must be an order of magnitude larger compared to the MI technique to reach the same accuracy for all three convergence criteria.

In Figure 6, the normalized singular values $(\sigma_i/\sigma_1)_{i=1}^n$ computed for the PCA methods are illustrated. The singular values are defined by

$$\sigma_i = \sqrt{\lambda_i}, \quad \text{such that } \sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n,$$

where $(\lambda_i)_{i=1}^n$ are the eigenvalues involved in Equation (19). It is shown that the spectrum resulting from the PCA-YN method decays faster than the other approaches. Moreover note that the eigenvalues involved in the MI approach (i.e. eigenvalues of Problem (14)) are strictly equal to the eigenvalues of the PCA-YN technique, see [13] for more details.

For the last experiment, we only consider the PCA-Y and MI approaches. The convergence of the Kullback-Leibler divergence, the expected Kullback-Leibler divergence, and the entropy with respect to the dimension of the reduced space is plotted in Figure 7 for the PCA-Y and MI methods using an even larger number of observations ($n = 2000$). For a dimension $r = 30$, the three quantities of interest are null within machine precision for the MI method, whereas the PCA-Y approaches needs a dimension $r = 700$ to get a value of 10^{-2} nat. This highlights that the accuracy of the MI method is more related to the number of parameters ($q = 30$) than the number of observations ($n = 500$ in Figures 2, 3, and 4, and $n = 2000$ in Figure 7) as predicted by Corollary 3.8 for the relative error on the mutual information.

4.3. Summary

We have seen that regarding the accuracy on the posterior, the KLD, EKLD, and MI approaches perform much better in terms of Kullback-Leibler divergence than the PCA approaches. On the other hand, in terms of computational efforts to evaluate the basis, the PCA methods and the MI technique only require the




Figure 5: Kullback-Leibler divergence (top), expected Kullback-Leibler divergence (middle), and entropy (bottom) versus the dimension of the reduced space for PCA-A, PCA-Y, PCA-YN and MI methods.




Figure 6: Normalized singular values for the PCA-A, PCA-Y, PCA-YN methods.

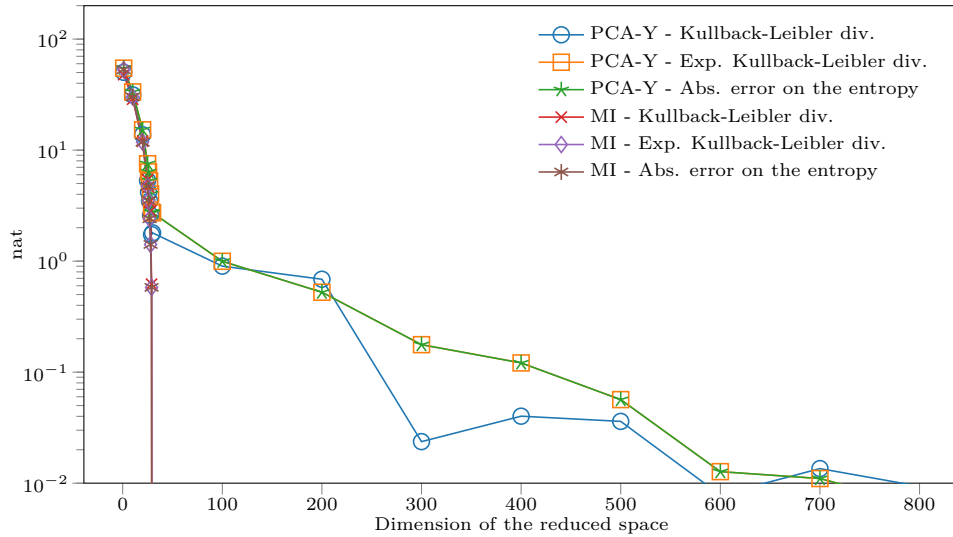


Figure 7: Kullback-Leibler divergence, expected Kullback-Leibler divergence, and the entropy for the PCA-Y and MI methods versus the dimension of the reduced space for a larger number of observations ($n = 2000$).

solution to an eigenvalue problem, whereas the others need more advanced strategies like the Riemannian optimization algorithm presented in Section 3.4. The maximization of the mutual information therefore exhibits a good balance between posterior distribution accuracy and computational difficulty, further providing an a priori error estimate as well as an upper bound on the number of required projections.

5. Application to nonlinear problems

We focused in the previous sections on the case of a linear problem where X , $A(X)$ and E follow the multivariate normal distribution. In [13], we showed that when A is nonlinear and X is drawn according to a uniform distribution, the PCA-YN is an appropriate dimensionality reduction method. The reason is that the probability density function of the posterior distribution of X is

$$\log f_X(X | Y = y) = -\frac{1}{2} \|y - A(X)\|_{C_E^{-1}}^2 + \text{constant},$$

and that the PCA-YN aims to approximate the random variable Y with respect to the Mahalanobis norm $\|\cdot\|_{C_E^{-1}}$. This illustrates that the appropriate dimensionality reduction method depends on the statistical model used for the inference.

In this section, we assess the benefits of the proposed approaches when the normality assumption is violated for $A(X)$, specifically using a log-normal model. In particular, we numerically evaluate the robustness of the approach using two values of the variance of the underlying Gaussian random vector, i.e. a small variance yielding a model that could be well approximated with a Gaussian process, and a large variance where the Gaussian assumption no longer holds. The last example also involves dimensionality reduction in a large-scale data setting, where the number of observations is drastically reduced.

5.1. Inference problem

For $\mathbf{s} \in (-1, 1)^2$, let F be a centered stationary Gaussian process with covariance function given by

$$k_F(\mathbf{s}, \mathbf{s}') = \sigma_{F,1}^2 \exp\left(-\frac{\|\mathbf{s} - \mathbf{s}'\|^2}{2\ell_F^2}\right) + \sigma_{F,2}^2 \delta(\mathbf{s} - \mathbf{s}').$$

The nonlinear regression model of interest is based on the PCA of the random vector $(F(\mathbf{s}_i))_{i=1}^n$, where $(\mathbf{s}_i)_{i=1}^n$ is a uniformly distributed sample in $(-1, 1)^2$. Let C_F be the covariance matrix of the random vector, and its eigenpairs $(\lambda_i, w_i)_{i=1}^n \in (\mathbb{R}^+ \times \mathbb{R}^n)^n$ ordered such that $\lambda_1 \geq \dots \geq \lambda_n$. The nonlinear model of interest is then given by

$$Y_i = \exp\left(\sum_{j=1}^q B_{ij} X_j\right) + E(\mathbf{s}_i), \quad \text{where } B_{ij} = (w_j)_i.$$

The prior distribution of X is deduced from the PCA of F . Given that F is a Gaussian process, X is chosen to follow the multivariate normal distribution $\mathcal{N}(0, C_X)$ such that $(C_X)_{ij} = \lambda_i \delta_{ij}$. The noise E is a centered Gaussian process independent of X and F with covariance function

$$k_E(\mathbf{s}, \mathbf{s}') = \sigma_{E,1}^2 \exp\left(-\frac{\|\mathbf{s} - \mathbf{s}'\|^2}{\ell_E}\right) + \sigma_{E,2}^2 \delta(\mathbf{s} - \mathbf{s}').$$

The synthetic data y are generated using the nonlinear model

$$\tilde{Y}_i = \exp(F(\mathbf{s}_i)) + E(\mathbf{s}_i), \quad \forall i \in \{1, \dots, n\}.$$

We are therefore introducing a model error accounting for the truncation to q terms of the PCA-based expansion of F .

In the applications below, the number of parameters is set to $q = 20$. Two different sets of values for the standard deviation parameters ($\sigma_{F,1}$, $\sigma_{F,2}$, $\sigma_{E,1}$ and $\sigma_{E,2}$) will be tested to control the nonlinearity of the mapping between the predictions and the observations. The correlation lengths are set to $\ell_F = 0.2$ and $\ell_E = 0.05$. Finally, we shall use $n = 2,000$ observation points.

5.2. Computation of the bases and error estimation

To compute the reduced bases, we rely on the expressions of the linear case which need the determination of the second moments of the nonlinear model $A(X) = (A_i(X))_{i=1}^n$, where $A_i(X) = \exp((BX)_i)$. The analytical expressions of the mean $m_A = \mathbb{E}(A(X))$ and the covariances $C_A = \mathbb{E}((A(X) - m_A)(A(X) - m_A)^T)$ and $C_{AX} = \mathbb{E}((A(X) - m_A)(X - m_X)^T)$ are given by

$$(m_A)_i = \exp\left(\frac{1}{2}D_{ii}\right), \quad (C_A)_{ij} = \exp\left(\frac{1}{2}(D_{ii} + D_{jj})\right) (\exp(D_{ij}) - 1),$$

$$\text{and } (C_{AX})_{ij} = (m_A)_i (BC_X)_{ij},$$

where $D = BC_X B^T$.

To assess the reduction error, the Kullback-Leibler divergence and the mutual information are not available in closed form. Their accurate numerical estimation is challenging and would require prohibitive sampling of the posterior distributions, for instance using a Markov-Chain Monte Carlo method, and an estimation of the probability density function with inherent source of error. The situation is even more complicated for the expected Kullback-Leibler divergence, requiring a repetitive sampling of the posterior distribution for the estimation of only one value of this quantity. Therefore, we choose to characterize the reduction error by its impact on the MAP value of the parameter. The MAP is computed by solving

$$\max_x \log f_X(x | Y = y) \Leftrightarrow \max_x \log f_Y(y | X = x) + \log f_X(x), \quad (20)$$

for the full (unreduced) approach and, in the case of the reduced models,

$$\max_x \log f_X(x | W = V^T y) \Leftrightarrow \max_x \log f_W(V^T y | X = x) + \log f_X(x). \quad (21)$$

These optimization problems are solved with a trust-region Newton method, using automatic differentiation for the evaluation of the gradient and the Hessian of the log density function.

We denote by x^{MAP} (resp. x_V^{MAP}) the MAP estimate of the full (resp. reduced) model. Since x^{MAP} is a stationary point of the log density function of the posterior distribution, the second-order Taylor expansion of $f_X(\cdot | Y = y)$ is given by

$$\log f_X(x | Y = y) \approx \log f_X(x^{\text{MAP}} | Y = y) + \frac{1}{2}(x - x^{\text{MAP}})^T \nabla^2 \log f_X(x^{\text{MAP}} | Y = y)(x - x^{\text{MAP}}).$$

Approximating locally the distribution by the multivariate normal distribution $\mathcal{N}(x^{\text{MAP}}, C^{\text{MAP}})$, where

$$C^{\text{MAP}} = -(\nabla^2 \log f_X(x^{\text{MAP}} | Y = y))^{-1},$$

gives the so-called Laplace approximation of the distribution [31]. Similarly, the posterior distribution of the reduced model will be approximated by the multivariate normal distribution $\mathcal{N}(x_V^{\text{MAP}}, C_V^{\text{MAP}})$ where

$$C_V^{\text{MAP}} = -(\nabla^2 \log f_X(x_V^{\text{MAP}} | W = V^T y))^{-1}.$$

In the following, we monitor the convergence of x_V^{MAP} to x^{MAP} with the dimension of the reduced space, as well as the convergence of the Hessian $(C_V^{\text{MAP}})^{-1}$ to $(C^{\text{MAP}})^{-1}$ in Frobenius norm. Note that it is empirically checked that the posterior distribution is unimodal by solving 200 times the Problems (20) and (21) with random initial guesses drawn according to the prior distribution. We denote by ϵ and ϵ_H the (\tilde{Y} -averaged) relative errors on the MAP and Hessian, respectively defined by

$$\epsilon = \frac{\mathbb{E}_{\tilde{Y}} (\|x_V^{\text{MAP}} - x^{\text{MAP}}\|)}{\mathbb{E}_{\tilde{Y}} (\|x^{\text{MAP}}\|)} \quad \text{and} \quad \epsilon_H = \frac{\mathbb{E}_{\tilde{Y}} \left(\left\| (C_V^{\text{MAP}})^{-1} - (C^{\text{MAP}})^{-1} \right\|_{\text{Fro}} \right)}{\mathbb{E}_{\tilde{Y}} \left(\left\| (C^{\text{MAP}})^{-1} \right\|_{\text{Fro}} \right)}. \quad (22)$$

The expectations appearing in the errors ϵ and ϵ_H are estimated by a crude Monte-Carlo method with a sample of size 70. This low sample size was found enough to obtain sufficiently correct error estimates, reflecting the robustness of all approaches which exhibit moderate dependences of the reduction error with the particular realization of \tilde{Y} .

5.3. Weak nonlinearity

In this section, the case of a weak nonlinearity is considered, setting the standard deviations to

$$\sigma_{F,1} = 0.3, \quad \sigma_{F,2} = 0.001, \quad \sigma_{E,1} = 0.1, \quad \text{and} \quad \sigma_{E,2} = 0.001.$$

The error estimates ϵ and ϵ_H introduced in Equation (22) are plotted in Figure 8 against the dimension of the reduced space. First, we observe that all the methods converge in terms of ϵ or ϵ_H . All the principal




Figure 8: Error ϵ and ϵ_H versus the dimension of the reduced space for the case of a weak nonlinearity.

component analysis based approaches perform poorly compared to the information theoretic techniques introduced here, with more than one order of magnitude difference when considering a reduced space of dimension 100. As a consequence, the normality assumption for the computation of the reduced basis is shown to improve the quality of the posterior distribution even when the statistical model does not have a Gaussian structure anymore.

It is interesting to note that the maximization of the mutual information (MI method) yields a basis that performs slightly better than the KLD or the EKLD approaches regarding the error on the MAP parameter ϵ .

The difference is less significant when considering the error on the Hessian ϵ_H but the information theoretic methods converges faster than the principal component analysis based approaches which tend to stagnate.

A comparison between the MAP estimates of the field, $A(x^{\text{MAP}})$ and $A(x_V^{\text{MAP}})$, for the same sample y of \tilde{Y} and the PCA-Y and MI methods is provided in Figure 9 for the reduction with $r = 60$. The plots highlight the better approximation for the MI method.




Figure 9: Contour plot of $A(x^{\text{MAP}})$ (dashed lines) and $A(x_V^{\text{MAP}})$ (solid lines) for the PCA-Y (left) and the MI (right) methods with a dimension of the reduced space $r = 60$.

5.4. Strong nonlinearity

A strong nonlinearity is considered by considerably increasing $\sigma_{F,1}$ and $\sigma_{E,1}$ compared to Section 5.3. The standard deviations are now set to

$$\sigma_{F,1} = 1.5 \quad \text{and} \quad \sigma_{E,1} = 0.6,$$

while $\sigma_{F,2}$ and $\sigma_{E,2}$ are identically set to 0.001. We expect now that the Gaussian assumption to be less useful than in Section 5.3.

Figure 10 depicts the convergence of the error estimators ϵ and ϵ_H with respect to the dimension of the reduced space for the different methods. In contrast to Section 5.3, all the approaches exhibit a similar convergence in terms of the error criteria ϵ and ϵ_H . Note that the PCA-Y method performs slightly better, especially for the error on the Hessian matrix. One major difference with the previous convergence curves reported previously in Figure 8 is the larger dimension of the reduced space needed to achieve a given relative error. Indeed, the dimension of the reduced space varies from 1 to 100 in Figure 8 and from 1 to 1000 in Figure 10. It indicates that a larger amount of observations is required to identify the posterior distribution of the model parameters, with similar relative accuracy, because of the non-linearities. Even if the normality assumption is violated, the information theoretic approaches are shown to be robust and converge to the original posterior distribution at the same rate as the PCA based methods.

The estimates $A(x^{\text{MAP}})$ and $A(x_V^{\text{MAP}})$ of the field are compared in Figure 11 for the PCA-Y and MI methods and dimension $r = 400$ and the same sample of \tilde{Y} . It confirms that for this highly non-linear case and this dimension of the reduced space, the two reduction approaches yield similar accuracy.

5.5. Large-scale problem

The objective of this section is to demonstrate the feasibility, robustness and efficiency of the proposed information-based reduction method in the context of large-scale simulations and large-dimensional observations. To this end, we consider the problem of identifying three values $\kappa_{\Omega_{1,2,3}}$ associated with the three




Figure 10: Error estimates versus the dimension of the reduced space for the case of a strong nonlinearity.




Figure 11: Contour plot of $A(x^{\text{MAP}})$ (dashed lines) and $A(x_V^{\text{MAP}})$ (solid lines) for the PCA-Y (left) and the MI (right) methods with a dimension of the reduced space $r = 400$.

subdomains, $\Omega_{1,2,3}$, of the two-dimensional domain Ω depicted in the left plot of Figure 12. These κ_{Ω_j} are independent and follow a log-normal distribution with parameters μ_{κ} , σ_{κ} . They are therefore expressed as

$$\kappa_{\Omega_j} = \exp[\mu_{\kappa} + \sigma_{\kappa} X_j], \quad X_j \sim \mathcal{N}(0, 1).$$

Thus, the vector of parameters to be inferred is $X \in \mathbb{R}^q$, $q = 3$. For simplicity, but without loss of generality, we shall use hereafter $\mu_{\kappa} = 0$ and $\sigma_{\kappa} = 1$. The inference uses a large set of $n \approx 32,000$ observations Y_i modeled as

$$Y_i = A_i(X) + E_i, \quad (23)$$

where $A_i(X) := U(x_i)$ is the solution at the observation point $x_i \in \Omega$ of the elliptic partial differential equation with uncertain parameters κ_{Ω_j} :

$$\nabla \cdot (\kappa(x) \nabla U(x)) = -1, \quad \kappa(x \in \Omega_j) = \kappa_{\Omega_j}.$$

The model equation is equipped with homogeneous Dirichlet (resp. Neumann) boundary conditions on the vertical and horizontal (resp. oblique) boundaries of Ω . The model for the E_i is again the independent centered Gaussian model with variance σ_{ϵ}^2 .




Figure 12: Left plot: Schematic of the problem domain (contained in a 3×3 square) and its three subdomains, Ω_j , over which $\kappa = \kappa_j$ is constant. Centre plots: centroids location for 20 and 100 clusters. Right plot: observation points x_i .

For the reduction, we consider the maximization of the mutual information (MI), requiring the solution of (15). Since C_E is diagonal, the reduced basis is given by the dominant eigenspace of C_A . Different

approaches can be used to estimate C_A . Here, we rely on a Polynomial Chaos (PC) method [21], exploiting the low dimensionality of X , and a standard, second-order finite element method for the spatial discretization of the elliptic problem on a very fine mesh. As expected from the low dimensionality of X , the decay of the spectrum of C_A is very fast. In Figure 13 we plot the first five dominant modes of C_A using the observation points shown in the right plot of Figure 12. Note that these observation points cover well the entire domain Ω .




Figure 13: The five leading reduced modes (from left to right) of the MI method plotted against the $n \approx 32,000$ observation points shown in the right plot of Figure 12.

For comparison purposes, we also consider more reduction approaches based on observation clustering. Indeed, the amount of observations ($n \approx 32,000$) appears an overkill to learn just $q = 3$ parameters. It is consequently tempting to disregard some observations and retain only $k > 0$ of them to carry the inference. However, we want to maintain a sufficient coverage of the domain, and so we rely on a clustering method (k-means [16, 23]) to partition the observations set into $k > 0$ distinct subsets, minimizing the Euclidean distances between the x_i and their respective cluster’s centroids. The k-means procedure randomly generates clusters with a roughly equal number of observations. In each cluster, the position x_i of the selected observation is the one closest to the corresponding cluster centroid. Two examples of selected observation points are depicted in the two center plots of Figure 12, for $k = 20$ and 100 clusters respectively. We shall refer to this reduction approach as “Centroids.” Disregarding all observations but the k -th closest to the centroids is clearly a brutal reduction approach, which is more susceptible to be affected by the noise compared to an approach involving the projection of all observations. Consequently, one may prefer to average (with equal weight) all the observations belonging to a cluster to define the corresponding reduced observation. This approach is referred to Cluster Averages (CAv) in the following.

The MI, Centroids and CAv reduction approaches are compared for three noise levels. The measurements y_i are randomly generated from (23) and plotted in Figure 14 to appreciate the noise to signal ratio.




Figure 14: Measurements y_i for noise level $\sigma_\epsilon = 0.01, 0.1$ and 0.5 from left to right.

To quantify the reduction errors, we consider as before the distance to the unreduced MAP point and Hessian:

$$\hat{\epsilon}(y) = \frac{\|x_V^{\text{MAP}} - x^{\text{MAP}}\|}{\|x^{\text{MAP}}\|} \quad \text{and} \quad \hat{\epsilon}_H(y) = \frac{\|(C_V^{\text{MAP}})^{-1} - (C^{\text{MAP}})^{-1}\|_{\text{Fro}}}{\|(C^{\text{MAP}})^{-1}\|_{\text{Fro}}}.$$

Note that we do not average over random observations Y , and restrict the analysis to a unique measurement y , because of the involved computational times. The convergence of the errors $\hat{\epsilon}(y)$ and $\hat{\epsilon}_H(y)$ with the dimension of the reduced spaces is reported in Figure 15, for the three approaches and the highest noise level (σ_ϵ). It is seen that the MI reduction converges for roughly 10 reduced modes, and outperforms the cluster-based reduction methods that converges at a much lower rate. As one may have expected, the convergence of the errors in the cluster-based methods is also noisier than in MI, with Centroids exhibiting higher sensitivity to noise than CAv.




Figure 15: Convergence with the reduction dimension of the MI, Centroids and Cluster Averages errors on MAP ($\hat{\epsilon}(y)$, left) and Hessian ($\hat{\epsilon}_H(y)$, right). Case of high noise level $\sigma_\epsilon = 0.5$.

However, the slow convergence of the cluster-based methods is due to the large noise in the previous example. This can be appreciated from the results reported in Figure 16, which show that $\hat{\epsilon}$ and $\hat{\epsilon}_H$ decrease with the noise level in the CAv method, but that the convergence rate remains the same. Also note that the convergence rate of the MI method appears to be insensitive to the noise level.

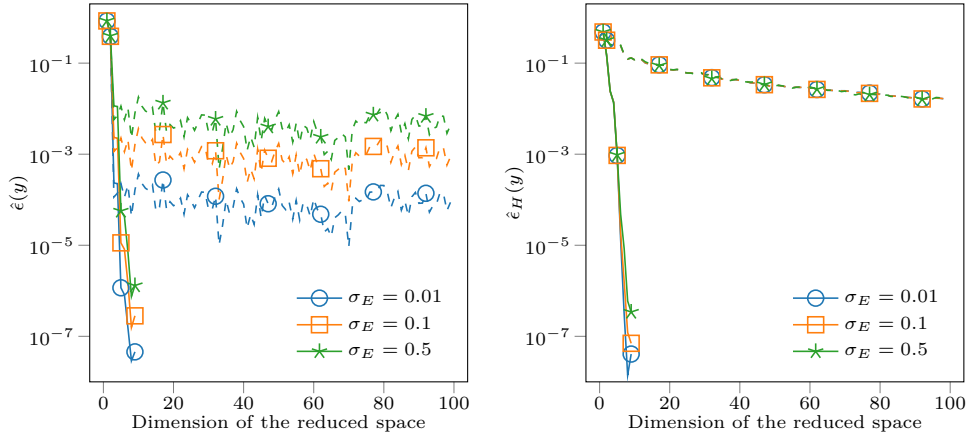


Figure 16: Convergence with the reduction dimension of the MI (solid lines) and CAv (dashed lines) MAP errors ($\hat{\epsilon}(y)$, left) and Hessian errors ($\hat{\epsilon}_H(y)$, right). Plotted are the errors for different noise intensities as indicated.

5.6. Summary

The numerical experiments of Sections 5.3–5.5 suggest that the information theoretic approaches yield robust reductions even though they were developed for linear Gaussian models. We have shown in particular

that they perform better than the PCA-based approaches, except in the strongly nonlinear case where all approaches behave similarly. Moreover, the solution to the maximization of the mutual information is significantly simpler to compute than in the KLD and EKLD techniques. Indeed, it only requires the solution of an eigenvalue problem and has therefore a computational complexity similar to the computation of the principal component analysis.

Moreover, the proposed approaches are robust to model errors as illustrated in Section 5.5. Indeed, even if we truncate the PCA-based expansion of the Gaussian process, F , the information theoretic reduction methods provide the lowest errors on the posterior distribution approximation.

6. Conclusions and perspectives

6.1. Conclusions

Different optimal reductions of observations by projection in a Bayesian framework are investigated in this work. The proposed methods are optimal in an information theoretic sense and aim at conserving the information about the posterior distribution of interest for Gaussian linear models with correlated additive noise.

Three optimization problems are proposed. First, the Kullback-Leibler divergence between the posterior distribution of the full and the reduced models is minimized. This corresponds to an a posteriori approach in the sense that a realization of the observations (a measurement) is required to compute the optimal projection. Second, we consider the minimization of the expected value of the previous Kullback-Leibler divergence, where the expectation is taken with respect to the observations. As a consequence, no measurement is required to compute the optimal reduced space and this strategy yields an a priori technique. The last proposed approach is the maximization of the mutual information between the projected observations and the parameters of interest. This last approach is equivalent to the minimization of the entropy of the posterior distribution.

Solving the first two optimization problems requires specific numerical algorithms. We use in this work the Riemannian trust-region algorithm on a manifold that take into account the invariance of the problems. In contrast, the mutual information maximization only requires the solution to a generalized eigenvalue problem. The computational cost and efficiency of the Riemannian algorithms will be addressed in a future work when large scale model will be considered.

Regarding the resulting posterior distributions, the three approaches are similar in terms of (possibly expected) Kullback-Leibler divergence and mutual information, and perform much better, on the considered examples, than the methods based on the principal component analysis of the observations. We advocate therefore that the mutual information maximization is the most appropriate approach for the determination of the optimal observation projection, given the balance between accuracy and computational complexity. For this particular approach, an a priori error estimate on the mutual information loss is readily available as well as a bound on the number of required projections. It is shown that no more projections than the rank of the linear model are required, which is in particular lower than the number of parameters to be inferred.

Moreover, we addressed the linear Gaussian case in this work. However, the proposed approaches only require the first two moments of the distributions and have been successfully applied to nonlinear non-Gaussian examples, in which optimality is no longer ensured.

6.2. Perspectives

In future works, the method will be applied to extreme hydrological flow problems (e.g. [13, 30]). In particular, we plan to apply the approach to the framework of Ensemble Kalman filters (EnKF) [11] for large datasets. The EnKF is a recursive Bayesian estimation technique for dynamical models of the form

$$\begin{aligned} X^{(k+1)} &= HX^{(k)} + L^{(k)}, \\ Y^{(k+1)} &= BX^{(k+1)} + E^{(k)}, \end{aligned}$$

where $X^{(0)}$, $L^{(k)}$, and $E^{(k)}$ are independent Gaussian vectors. Note that the equation above is the same as (1). To estimate the posterior distribution of $X^{(k+1)}$, the Kalman filter requires the inversion of the

covariance matrix $C_{Y^{(k+1)}}$ at each iteration of the discrete dynamical system. However, in the EnKF, $C_{X^{(k+1)}}$ is estimated using a Monte-Carlo estimator with a sample size that can be much lower than the total number of observations n . As a consequence, the covariance of the forward state $X^{(k+1)}$ is low-rank and we showed in this paper that a low number of projections of the observations are enough to recover the mutual information between the estimated distribution of the state $X^{(k+1)}$ and the observations $Y^{(k+1)}$.

Additional challenges arise when the datasets are high dimensional. Considering the mutual information based technique, the problem could be first tackled using high performance computing. Given that we a priori know an upper bound on the number of projections that is already low, we only need an efficient matrix product computation (e.g. [10]) to implement the algorithm from [1]. Further developments are required to appropriately use these approaches in a streaming environment.

Acknowledgments

This work is supported by King Abdullah University of Science and Technology Awards CRG3-2156 and OSR-2016-RPP-3268.

Appendix A. Proof of Proposition 2.1

According to Bayes' theorem, the posterior distribution is such that

$$f_X(x | Y = y) \propto f_Y(y | X = x)f_X(x),$$

or equivalently,

$$\begin{aligned} \log f_X(x | Y = y) &= \log f_Y(y | X = x) + \log f_X(x) + k_0 \\ &= -\frac{1}{2} \left((y - Bx - m_E)^T C_E^{-1} (y - Bx - m_E) \right. \\ &\quad \left. + (x - m_X)^T C_X^{-1} (x - m_X) \right) + k_1, \end{aligned}$$

where k_0 and k_1 are constants. Because the log probability density function is quadratic with respect to x , we conclude that the posterior distribution is also a multivariate normal distribution, i.e. $P(X | Y = y) \sim \mathcal{N}(m_\star, C_\star)$. This implies that, up to a constant k_2 , the following equality holds

$$\log f_X(x | Y = y) = -\frac{1}{2} \left((x - m_\star)^T C_\star^{-1} (x - m_\star) \right) + k_2.$$

Identifying the quadratic terms in x and using the Woodbury matrix identity [18, Equation (29)] gives

$$\begin{aligned} C_\star^{-1} &= C_X^{-1} + B^T C_E^{-1} B = C_X^{-1} (C_X + C_{AX}^T C_E^{-1} C_{AX}) C_X^{-1} \\ \text{and } C_\star &= C_X - C_X B^T (C_E + B C_X B^T)^{-1} B C_X = C_X - C_{AX}^T C_Y^{-1} C_{AX}. \end{aligned}$$

Identifying the linear term w.r.t. x yields

$$\begin{aligned} C_\star^{-1} m_\star &= C_X^{-1} m_X + B^T C_E^{-1} (y - m_E) \\ &= B^T C_E^{-1} (y - m_Y) + B^T C_E^{-1} m_A + C_X^{-1} m_X. \end{aligned}$$

We finally have

$$m_\star = G_\star (y - m_Y) + h_\star,$$

with

$$\begin{aligned} G_\star &= C_\star B^T C_E^{-1} = C_{AX}^T (I - C_Y^{-1} C_A) C_E^{-1} = C_{AX}^T C_Y^{-1}, \\ h_\star &= C_\star C_X^{-1} m_X + G_\star m_A. \end{aligned}$$

For the posterior distribution of the reduced model, we substitute $(y - m_Y)$, C_Y , m_E , C_E , m_A , C_{AX} , and C_A respectively by $V^T(y - m_Y)$, $V^T C_Y V$, $V^T m_E$, $V^T C_E V$, $V^T m_A$, $V^T C_{AX}$, and $V^T C_A V$ in the full model. The fact that V is full-rank ensures that $V^T C_Y V$, $V^T C_E V$ and

$$C_X + C_{AX}^T V (V^T C_E V)^{-1} V^T C_{AX}$$

are symmetric positive definite matrices and hence are invertible.

Appendix B. Proof of Proposition 2.2

For $M \in \text{GL}_r$ we have

$$C_{AX} V M (M^T V^T C_E V M)^{-1} M^T V^T C_{AX} = C_{AX} V (V^T C_E V)^{-1} V^T C_{AX},$$

so we deduce that $C_{VM} = C_V$. Moreover, given that

$$G_{VM} = C_{AX}^T V M (M^T V^T C_Y V M)^{-1} = C_{AX}^T V (V^T C_Y V)^{-1} M^{-T} = G_V M^{-T},$$

we conclude that $G_{VM} (VM)^T = G_V V^T$, $h_{VM} = h_V$ and finally $m_{VM} = m_V$.

Appendix C. Proof of Proposition 3.1

Using the definition of the Kullback-Leibler divergence (7), we have

$$\begin{aligned} \text{D}_{\text{KL}}(P(Z_0) \parallel P(Z_1)) &= \mathbb{E}_Z \left(\log \left(\frac{\det(C_1)^{\frac{1}{2}}}{\det(C_0)^{\frac{1}{2}}} \right) \right. \\ &\quad \left. - \frac{1}{2} (Z - m_0)^T C_0^{-1} (Z - m_0) + \frac{1}{2} (Z - m_1)^T C_1^{-1} (Z - m_1) \right). \end{aligned}$$

Given that $Z \sim P(Z_0)$, we deduce that

$$\begin{aligned} \mathbb{E}_Z \left((Z - m_0)^T C_0^{-1} (Z - m_0) \right) &= \mathbb{E}_Z \left(\text{tr} \left((Z - m_0)^T C_0^{-1} (Z - m_0) \right) \right) \\ &= \text{tr}(\mathbb{E}_Z((Z - m_0)(Z - m_0)^T) C_0^{-1}) \\ &= \text{tr}(C_0 C_0^{-1}) = q. \end{aligned} \tag{C.1}$$

Moreover we have

$$\begin{aligned} &\mathbb{E}_Z \left((Z - m_1)^T C_1^{-1} (Z - m_1) \right) \\ &= \mathbb{E}_Z \left((Z - m_0 + m_0 - m_1)^T C_1^{-1} (Z - m_0 + m_0 - m_1) \right) \\ &= \mathbb{E}_Z \left((Z - m_0)^T C_1^{-1} (Z - m_0) + (m_0 - m_1)^T C_1^{-1} (m_0 - m_1) \right. \\ &\quad \left. + 2 (Z - m_0)^T C_1^{-1} (m_0 - m_1) \right). \end{aligned}$$

Using the same trace technique as in Equation C.1, and using the fact that $\mathbb{E}_Z(Z) = m_0$, the term $\mathbb{E}_Z((Z - m_1)^T C_1^{-1} (Z - m_1))$ is equal to

$$\mathbb{E}_Z \left((Z - m_1)^T C_1^{-1} (Z - m_1) \right) = \text{tr}(C_0 C_1^{-1}) + (m_0 - m_1)^T C_1^{-1} (m_0 - m_1),$$

yielding the final result

$$\begin{aligned} \text{D}_{\text{KL}}(P(Z_0) \parallel P(Z_1)) &= \\ &\frac{1}{2} \left(\text{tr}(C_0 C_1^{-1}) - \log \det(C_0 C_1^{-1}) - q + (m_0 - m_1)^T C_1^{-1} (m_0 - m_1) \right). \end{aligned}$$

AppendixD. Proof of Theorem 3.2

First, the map \mathcal{J}_0 is smooth ($\in \mathcal{C}^\infty$) as the sum and composition of smooth functions, noting that the determinant is always strictly positive.

Let $\pi : \mathbb{R}_*^{n \times r} \rightarrow \text{Gr}(r, n)$ denotes the canonical projection defined by $\pi(V) = [V]$. Let $\mathcal{K}_0 : \text{Gr}(r, n) \rightarrow \mathbb{R}$ be the map defined by $\mathcal{J}_0(V) = \mathcal{K}_0 \circ \pi(V)$. \mathcal{K}_0 is in fact the functional we are minimizing in Problem (9).

According to [3, Proposition 3.4.5], the smoothness of \mathcal{J}_0 implies that \mathcal{K}_0 is smooth and in particular continuous. According to [25, Lemma 5.1] $\text{Gr}(r, n)$ is compact, the extreme value theorem concludes the proof.

AppendixE. Proof of Proposition 3.3

Since only m_\star and m_V depend on Y in Equation (8), the expected Kullback-Leibler divergence admits the form

$$\mathbb{E}_Y (\text{D}_{\text{KL}} (P(X | Y) \parallel P(X | W = V^T Y))) = \frac{1}{2} (\text{D}_{\text{el}}(C_\star, C_V) + \mathbb{E}_Y (\text{D}_{C_V}(m_\star, m_V))).$$

Note that $m_\star = G_\star(Y - m_Y) + h_\star$ and $m_V = G_V V^T(Y - m_Y) + h_V$, hence

$$m_\star - m_V = (G_\star - G_V V^T)(Y - m_Y) + (h_\star - h_V),$$

and

$$\mathbb{E}_Y (\text{D}_{C_V}(m_\star, m_V)) = \text{tr} \left(C_V^{-1} (G_\star - G_V V^T) C_Y (G_\star - G_V V^T)^T \right) + (h_\star - h_V) C_V^{-1} (h_\star - h_V),$$

which yields the final result.

AppendixF. Proof of Theorem 3.6

For a normally distributed \mathbb{R}^n -valued random variable $Z \sim \mathcal{N}(m_Z, C_Z)$, the entropy $H(Z)$ is given by

$$H(Z) = \frac{1}{2} \log(\det(C_Z)) + \frac{n}{2} \log(2\pi e).$$

Given that X and W are normally distributed, we immediatly deduce

$$H(X) = \frac{1}{2} \log(\det(C_X)) + \frac{q}{2} \log(2\pi e),$$

and $H(W) = \frac{1}{2} \log(\det(V^T C_Y V)) + \frac{r}{2} \log(2\pi e).$

In order to compute the joint-entropy $H(W, X)$, we need to characterize the covariance of (W, X) . Note that we already know that (W, X) is drawn according to a Gaussian distribution. In order to obtain the covariance $C_{(W, X)}$, we identify the quadratic terms in the following equality between the probability density functions:

$$\log f_{(W, X)}(W, X) = \log f_W(W | X) + \log f_X(X),$$

where the likelihood $f_W(W | X)$ is directly deduced from Equation (2). The conditional random distribution $P(W | X)$ follows the Gaussian distribution $\mathcal{N}(V^T(A X + m_E), V^T C_E V)$. Identifying the quadratic terms yields

$$C_{(W, X)}^{-1} = \begin{pmatrix} (V^T C_E V)^{-1} & - (V^T C_E V)^{-1} V^T A \\ -A^T V (V^T C_E V)^{-1} & C_X^{-1} + A^T V (V^T C_E V)^{-1} V^T A \end{pmatrix}.$$

According to [27, Section 9.1.2], the determinant of the precision matrix $C_{(W,X)}^{-1}$ is given by

$$\det\left(C_{(W,X)}^{-1}\right) = \det\left(\left(V^T C_E V\right)^{-1}\right) \det\left(C_X^{-1}\right).$$

We immediatly have

$$\begin{aligned} H(W, X) &= \frac{1}{2} \log \det(C_{(W,X)}) + \frac{r+q}{2} \log(2\pi e) \\ &= -\frac{1}{2} \log \det(C_{(W,X)}^{-1}) + \frac{r+q}{2} \log(2\pi e) \\ &= -\frac{1}{2} \left(\log \det(C_X^{-1}) + \log \det\left(\left(V^T C_E V\right)^{-1}\right) \right) + \frac{r+q}{2} \log(2\pi e), \end{aligned}$$

and the mutual information reduces to

$$\begin{aligned} \mathcal{I}(W, X) &= \frac{1}{2} \left(\log \det(V^T C_Y V) + \log \det\left(\left(V^T C_E V\right)^{-1}\right) \right) \\ &= \frac{1}{2} \log \det\left(\left(V^T C_Y V\right) \left(V^T C_E V\right)^{-1}\right), \end{aligned}$$

which proves the first equality.

Regarding the entropy of the posterior distribution, we know that $P(X | W = V^T y) \sim \mathcal{N}(m_V, C_V)$, yielding

$$H(P(X | W = V^T y)) = \frac{1}{2} \log \det(C_V) + \frac{q}{2} \log(2\pi e),$$

and the entropy does not depend on the realization y of Y . Using Equation (5), we have

$$\begin{aligned} \log \det(C_V) &= \log \det\left(C_X^{\frac{1}{2}} \left(\mathbf{I} - C_X^{\frac{1}{2}} B^T V \left(V^T C_Y V\right)^{-1} V^T B C_X^{\frac{1}{2}}\right) C_X^{\frac{1}{2}}\right) \\ &= \log \det\left(\mathbf{I} - C_X^{\frac{1}{2}} B^T V \left(V^T C_Y V\right)^{-1} V^T B C_X^{\frac{1}{2}}\right) + \log \det C_X \\ &= \log \det\left(V^T C_Y V - V^T B C_X B^T V\right) - \log \det\left(V^T C_Y V\right) + \log \det C_X. \end{aligned}$$

The last equality is obtained using the identity $\det(\mathbf{I} + MN) = \det(\mathbf{I} + NM)$, with $M = (V^T C_Y V)^{-1/2} V^T B C_X^{1/2}$ and $N = M^T$ and factorizing the resulting expression by $(V^T C_Y V)^{-1/2}$ on the left and right in the determinant. We finally find that

$$\begin{aligned} \log \det(C_V) &= \log \det\left(V^T C_Y V - V^T C_A V\right) - \log \det\left(V^T C_Y V\right) + \log \det C_X \\ &= \log \det\left(V^T C_E V\right) - \log \det\left(V^T C_Y V\right) + \log \det C_X \\ &= \log \det C_X - \log \det\left(\left(V^T C_Y V\right) \left(V^T C_E V\right)^{-1}\right), \end{aligned}$$

and the entropy is

$$\begin{aligned} H(P(X | W = V^T y)) &= \\ &= -\frac{1}{2} \log \det\left(\left(V^T C_Y V\right) \left(V^T C_E V\right)^{-1}\right) + \frac{1}{2} \log \det(C_X) + \frac{q}{2} \log(2\pi e), \end{aligned}$$

that proves the second equality.

For the last part of the proof, we consider the maximization problem

$$\max_{V \in \mathbb{R}_+^{n \times r}} \log \det\left(\left(V^T C_Y V\right) \left(V^T C_E V\right)^{-1}\right) = \max_{V \in \mathbb{R}_+^{n \times r}} 2 \mathcal{I}(W, X). \quad (\text{F.1})$$

First let us introduce a change of variable, setting $U = C_E^{\frac{1}{2}}V$. The optimization problem becomes

$$\max_{U \in \mathbb{R}_*^{n \times r}} \det \left(\left(U^T C_E^{-\frac{1}{2}} C_Y C_E^{-\frac{1}{2}} U \right) (U^T U)^{-1} \right).$$

Then, the quantity $\mathcal{K}(U) = \det((U^T C_Y U)(U^T U)^{-1})$ is invariant under any invertible linear transformation on the right, meaning that $\mathcal{K}(U) = \mathcal{K}(UQ)$ for any $Q \in \mathbb{R}^{r \times r}$ invertible. With $\text{St}(r, n)$ denoting the Stiefel manifold defined by

$$\text{St}(r, n) = \{M \in \mathbb{R}^{n \times r}; M^T M = \mathbf{I}_r\},$$

there exists a matrix $\hat{U} \in \text{St}(r, n)$ such that $\mathcal{K}(U) = \mathcal{K}(\hat{U})$. Such a matrix \hat{U} can be computed using, for instance, a thin QR factorization. We can therefore consider the following equivalent optimization problem

$$\max_{U \in \text{St}(r, n)} \det \left(U^T C_E^{-\frac{1}{2}} C_Y C_E^{-\frac{1}{2}} U \right). \quad (\text{F.2})$$

In order to conclude the proof, we need the following result.

Lemma Appendix F.1. *Let $K \in \mathbb{R}^{n \times n}$ be a symmetric positive definite matrix with eigenvalues $(\lambda_i)_{i=1}^n$ in a decreasing order. Then we have*

$$\max_{U \in \text{St}(r, n)} \log \det (U^T K U) = \sum_{i=1}^r \log \lambda_i. \quad (\text{F.3})$$

Moreover, any solution to the optimization Problem (F.3) is an invariant subspace of K and a particular solution is given by the matrix U whose columns are the eigenvectors of K associated to the eigenvalues $(\lambda_i)_{i=1}^r$.

Proof. First, a solution to Problem (F.3) exists using the fact that $\mathcal{F} : U \mapsto \log \det(U^T K U)$ is continuous and $\text{St}(r, n)$ is compact. It is closed as the inverse image of $\{0\}$ by the continuous function $U \mapsto U^T U - \mathbf{I}$, and bounded because $\|U\|_{\text{Fro}}^2 = r$ for all $U \in \text{St}(r, n)$. The extreme value theorem implies the existence of a maximizer.

Let us introduce the map $\mathcal{H} : \mathbb{R}^{n \times r} \times \mathbb{R}^{r \times r} \rightarrow \mathbb{R}$ be defined by

$$\mathcal{H}(U, \Psi) = 2 \text{tr} \left((U^T U - \mathbf{I}) \Psi \right),$$

and consider the Lagrangian function $\mathcal{L}(U, \Psi) = \mathcal{F}(U) + \mathcal{H}(U, \Psi)$ associated to the constrained optimization Problem (F.3). An optimal solution (U_*, Ψ_*) satisfies the equation

$$D_U \mathcal{L}(U_*, \Psi_*)[\delta U] = 0, \quad \forall \delta U \in \mathbb{R}^{n \times r},$$

where $D_U \mathcal{L}(U_*, \Psi_*)[\delta U]$ denotes the Gâteaux derivative of the Lagrangian \mathcal{L} at U_* in the direction δU with respect to the first parameter. Given the formula

$$\frac{d}{dt} \log \det (U + t\delta U) = 2 \text{tr} \left((U + t\delta U)^{-1} \delta U \right),$$

we conclude that the Gâteaux derivative $D_U \mathcal{F}(U)[\delta U]$ is

$$D_U \mathcal{F}(U)[\delta U] = 2 \text{tr} \left((U^T K U)^{-1} U^T K \delta U \right),$$

and similarly we have

$$D_U \mathcal{H}(U)[\delta U] = 2 \text{tr} \left((\Psi_* + \Psi_*^T) U^T \delta U \right).$$

Hence, for all $\delta U \in \mathbb{R}^{n \times r}$, a solution U_* to Problem (F.3) satisfies

$$D_U \mathcal{L}(U_*, \Psi_*)[\delta U] = 2 \text{tr} \left(\left((U_*^T K U_*)^{-1} U_*^T K + (\Psi_* + \Psi_*^T) U_*^T \right) \delta U \right) = 0.$$

The result holding for all δU , we conclude that U_\star satisfies

$$\begin{aligned} (U_\star^T K U_\star)^{-1} U_\star^T K + (\Psi_\star + \Psi_\star^T) U_\star^T &= 0 \\ \Leftrightarrow K U_\star &= -U_\star (\Psi_\star + \Psi_\star^T) U_\star^T K U_\star. \end{aligned}$$

Finally, multiplying this last equation on the left by U_\star^T and on the right by $(U_\star^T K U_\star)^{-1}$ gives that $\Psi_\star + \Psi_\star^T = -I_r$ and

$$K U_\star = U_\star U_\star^T K U_\star,$$

meaning that U_\star spans an r -dimensional invariant subspace of K .

To conclude the proof, let \mathcal{U} be the r -dimensional subspace spanned by the columns of U_\star , i.e. $\mathcal{U} = \text{range } U_\star$, and consider K as a linear map on \mathbb{R}^n .

K being diagonalizable, the restriction $K|_{\mathcal{U}}$ of K to its invariant subspace \mathcal{U} is also diagonalizable. Hence there exists an orthonormal basis of \mathcal{U} formed of eigenvectors of $K|_{\mathcal{U}}$ and therefore of eigenvectors of K . Given the invariance $\mathcal{F}(UQ) = \mathcal{F}(U)$ for every orthogonal matrix $Q \in \mathbb{R}^{r \times r}$, we can arbitrary set the columns of U_\star to be eigenvectors of K . As a consequence, the determinant is

$$\log \det(U_\star^T K U_\star) = \sum_{i \in \mathcal{I}} \log \lambda_i,$$

where \mathcal{I} is a subset of $\{1, \dots, n\}$ such that $\#\mathcal{I} = r$. The sum is maximized by picking the r largest eigenvalues $(\lambda_i)_{i=1}^r$, and therefore a solution U_\star is given by a matrix whose columns corresponds to r eigenvectors associated to the dominant eigenvalues. \square

Since $C_E^{-\frac{1}{2}} C_Y C_E^{-\frac{1}{2}}$ is symmetric positive definite, Lemma [AppendixF.1](#) gives first that a solution to Problem [\(F.2\)](#) is given by the matrix U whose columns are the dominant eigenvectors of $C_E^{-\frac{1}{2}} C_Y C_E^{-\frac{1}{2}}$. Using the equality $V = C_E^{-\frac{1}{2}} U$, we finally find that a solution to Problem [\(F.3\)](#) is given by the matrix V whose columns are r dominant eigenvectors associated to the generalized eigenvalue problem

$$C_Y v = \lambda C_E v, \quad \lambda \in \mathbb{R}, v \in \mathbb{R}^n.$$

- [1] P.-A. Absil, C.G. Baker, and K.A. Gallivan. Trust-region methods on Riemannian manifolds. *Foundations of Computational Mathematics*, 7(3):303–330, 2007.
- [2] P.-A. Absil, R. Mahony, and R. Sepulchre. Riemannian geometry of Grassmann manifolds with a view on algorithmic computation. *Acta Applicandae Mathematica*, 80(2):199–220, Jan 2004.
- [3] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.
- [4] M. Barrault, Y. Maday, N. Cuong Nguyen, and A. T. Patera. An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Mathematique*, 339(9):667 – 672, 2004.
- [5] M. Bebendorf. Approximation of boundary element matrices. *Numerische Mathematik*, 86(4):565–589, Oct 2000.
- [6] M. Bebendorf, Y. Maday, and B. Stamm. *Comparison of Some Reduced Representation Approximations*, chapter 3, pages 67–100. Springer International Publishing, Cham, 2014.
- [7] K. P. Burnham and D. R. Anderson. *Model selection and multimodel inference: a practical information-theoretic approach*. Springer Science & Business Media, 2003.
- [8] P. Chen, A. Quarteroni, and G. Rozza. A weighted empirical interpolation method: a priori convergence analysis and applications. *ESAIM: Mathematical Modelling and Numerical Analysis*, 48(4):943–953, 2014.
- [9] T. Cui, J. Martin, Y. M. Marzouk, A. Solonen, and A. Spantini. Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems*, 29:114015, 2014.
- [10] P. Drineas, R. Kannan, and M. W. Mahoney. Fast Monte Carlo Algorithms for Matrices I: Approximating Matrix Multiplication. *SIAM Journal on Computing*, 36(1):132–157, jan 2006.
- [11] G. Evensen. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99(C5):10143, may 1994.
- [12] L. N. Geppert, K. Ickstadt, A. Munteanu, J. Quedenfeld, and C. Sohler. Random projections for bayesian regression. *Statistics and Computing*, 27(1):79–101, Jan 2017.
- [13] L. Giraldi, O. P. Le Maître, K. T. Mandli, C. N. Dawson, I. Hoteit, and O. M. Knio. Bayesian inference of earthquake parameters from buoy data using a polynomial chaos-based surrogate. *Computational Geosciences*, 21(4):683–699, Aug 2017.

- [14] I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals, Series, and Products*. Elsevier, eight edition edition, 2014.
- [15] W. Hackbusch. *Tensor Spaces and Numerical Tensor Calculus*, volume 42 of *Springer Series in Computational Mathematics*. Springer-Verlag Berlin Heidelberg, 2012.
- [16] J. A. Hartigan and M. A. Wong. Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1):100–108, 1979.
- [17] H. Hotelling. Analysis of a complex of statistical variables with principal components. *Journal of Educational Psychology*, 24:417–441, 1933.
- [18] A. S. Householder. A survey of some closed methods for inverting matrices. *Journal of the Society for Industrial and Applied Mathematics*, 5(3):155–169, 1957.
- [19] I. T. Jolliffe. *Principal Component Analysis*. Springer, 2002.
- [20] K. Karhunen. Über lineare methoden in der wahrcheinlichkeitsrechnung. *Annales Academiae Scientiarum Fennicae Series A1, Mathematical Physics* 37, 37:1–79, 1947.
- [21] O. P. Le Maître and O. M. Knio. *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*. Springer, 1st editio edition, 2010.
- [22] M. Loève. Fonctions aléatoires du second ordre. *Processus stochastiques et mouvement Brownien*, 1948.
- [23] J. Mac Queen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pages 281–297, Berkeley, Calif., 1967. University of California Press.
- [24] D. Maclaurin, D. Duvenaud, M. Johnson, and R. P. Adams. Autograd: Reverse-mode differentiation of native Python, 2017.
- [25] J. W. Milnor and J. D. Stasheff. *Characteristic Classes*, volume 76 of *Ann. of Math. Stud.* Princeton Univ. Press, Princeton, NJ, 1974.
- [26] K. Pearson. Liii. on lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(11):559–572, 1901.
- [27] K. B. Petersen and M. S. Pedersen. The matrix cookbook, nov 2012. Version 20121115.
- [28] A. Spantini, T. Cui, K. Willcox, L. Tenorio, and Y. M. Marzouk. Goal-oriented optimal approximations of bayesian linear inverse problems. *SIAM Journal on Scientific Computing*, in press, 2017.
- [29] 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(6):A2451–A2487, 2015.
- [30] I. Sraj, K. T. Mandli, O. M. Knio, C. N. Dawson, and I. Hoteit. Uncertainty quantification and inference of manning’s friction coefficients using dart buoy data during the tōhoku tsunami. *Ocean Modelling*, 83:82 – 97, 2014.
- [31] L. Tierney and J. B. Kadane. Accurate approximations for posterior moments and marginal densities. *Journal of the american statistical association*, 81(393):82–86, 1986.
- [32] J. Townsend, N. Koep, and S. Weichwald. Pymanopt: A Python toolbox for optimization on manifolds using automatic differentiation. *Journal of Machine Learning Research*, 17(137):1–5, 2016.