



**HAL**  
open science

# Observable Dictionary Learning for High-Dimensional Statistical Inference

Lionel Mathelin, Kévin Kasper, Hisham Abou-Kandil

► **To cite this version:**

Lionel Mathelin, Kévin Kasper, Hisham Abou-Kandil. Observable Dictionary Learning for High-Dimensional Statistical Inference. Archives of Computational Methods in Engineering, 2018, 25 (1), pp.103 - 120. 10.1007/s11831-017-9219-2 . hal-01726821

**HAL Id: hal-01726821**

**<https://hal.science/hal-01726821>**

Submitted on 19 Jan 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Observable dictionary learning for high-dimensional statistical inference

Lionel Mathelin<sup>a,\*</sup>, Kévin Kasper<sup>b,c</sup>, Hisham Abou-Kandil<sup>b</sup>

<sup>a</sup>LIMSI-CNRS (UPR 3251), rue John von Neumann, Campus Universitaire, Bât. 508, 91405 Orsay cedex, France

<sup>b</sup>Ecole Normale Supérieure de Cachan, SATIE-CNRS (UMR 8029), 61 av. Président Wilson, 94230 Cachan, France

<sup>c</sup>Université Grenoble Alpes, IUT1, 151 rue de la Papeterie, 38400 Saint-Martin-d'Hères, France

---

## Abstract

This paper introduces a method for efficiently inferring a high-dimensional distributed quantity from a few observations. The quantity of interest (QoI) is approximated in a basis (dictionary) learned from a training set. The coefficients associated with the approximation of the QoI in the basis are determined by minimizing the misfit with the observations. To obtain a probabilistic estimate of the quantity of interest, a Bayesian approach is employed. The QoI is treated as a random field endowed with a hierarchical prior distribution so that closed-form expressions can be obtained for the posterior distribution. The main contribution of the present work lies in the derivation of *a representation basis consistent with the observation chain* used to infer the associated coefficients. The resulting dictionary is then tailored to be both observable by the sensors and accurate in approximating the posterior mean. An algorithm for deriving such an observable dictionary is presented. The method is illustrated with the estimation of the velocity field of an open cavity flow from a handful of wall-mounted point sensors. Comparison with standard estimation approaches relying on Principal Component Analysis and K-SVD dictionaries is provided and illustrates the superior performance of the present approach.

*Keywords:* Bayesian inference, inverse problem, dictionary learning, sparsity promotion, K-SVD, observability, hierarchical prior.

---

## 1. Introduction

State estimation and parameter inference are very common problems occurring in many situations. As an example, when controlling a physical system, state feedback control achieves the best results but the state vector is rarely fully measurable and it is hence necessary to estimate it from the scarce observations (sensors) available. In another context, numerical simulations in general have now reached a degree of maturity and reliability such that the bottleneck of the accuracy of the output of the simulations typically is the quality of the input parameters. These parameters can be initial conditions, physical properties, geometry, *etc.* Accurately estimating these input parameters is key to reliable simulations. In both the situations discussed above, the estimation of the hidden variables typically has to be done from scarce, indirect, observations of the system at hand while the information of interest is often a high-, or even infinite-, dimensional quantity. This leads to mathematically severely ill-posed problems.

A uniform grid-based description of the quantity of interest (QoI) is often employed and yields a representation under the form of a (very) high-dimensional vector, expressing a high degree of redundancy and suffering from over-parameterization. Identification of such a grid-based description through inverse techniques is a challenge. Several approaches can be followed to address this issue. A popular one is to reduce the number of degrees of freedom to describe the quantity of interest, at the price of an approximation. If

---

\*Corresponding author

*Email addresses:* mathelin@limsi.fr (Lionel Mathelin), kevin.kasper@univ-grenoble-alpes.fr (Kévin Kasper), hisham.abou-kandil@satie.ens-cachan.fr (Hisham Abou-Kandil)

expertise knowledge is available, it can be exploited to guide the choice of the approximation space. For instance, physical quantities have finite-energy, motivating an approximation in, say, a  $L^2$ -space. Other typical examples include smoothness, known class of probability distribution, *etc.* This naive approach can advantageously be improved by an approach exploiting correlations of the QoI among the entries of its grid-based description. For instance, in the case the QoI is a spatial field of physical properties, as encountered, for example, in subsurface reconstruction problems, this amounts to exploiting spatial correlations and the fact that the field to be inferred typically enjoys some degree of smoothness.<sup>1</sup> A representation basis is chosen and characterizing the QoI reduces to estimating the coefficients associated with the elements of the basis. This certainly is the most popular approach to inverse problems where the estimation typically reformulates as finding the small set of scalar-valued coefficients associated with the dominant modes of the approximation basis. Among many efforts from the literature, this strategy has been adopted in Buffoni *et al.* (2008); Yang *et al.* (2010) or Leroux *et al.* (2015), where a Principal Component Analysis (PCA), also termed Proper Orthogonal Decomposition (POD), Holmes *et al.* (2012); Sirovich (1987), was used as an approximation basis. Once the approximation basis is determined, estimating the associated coefficients is done by minimizing the data misfit with the available observations. With very few observations, this data misfit minimization can itself be ill-posed and the solution method must involve some kind of regularization. A typical choice, often justified from physics-related arguments, is to look for the least energy solution, *i.e.*, that with the lowest  $L^2$ -norm: the ubiquitous Tikhonov regularization. Another approach is to determine an approximation basis in which the QoI is likely to be sparse. The data misfit minimization then includes a penalization of the cardinality of the solution, in the form of a  $L^0$ -(quasi-)norm, and is efficiently solved using sparse coding algorithms such as Orthogonal Matching Pursuit (OMP), Mallat & Zhang (1993), CoSaMP, Needell & Tropp (2008), or Iterative Hard Thresholding (IHT), Blumensath & Davies (2009), to cite only a few. For instance, this approach is followed by Khaninezhad *et al.* (2012*a,b*) who have used a K-SVD algorithm, Aharon *et al.* (2006); Rubinstein *et al.* (2008), to determine such a sparsity-promoting approximation basis in a subsurface reconstruction context.

Another strategy to address the ill-posedness of inverse problems is to explicitly introduce prior knowledge. For instance, a deterministic QoI can be treated as a random field and endowed with a prior distribution. A Bayesian approach can then be employed to “assimilate” the QoI from the prior to the posterior distribution using the available observations. This is a rigorous approach and results in a complete probabilistic estimation of the parameters at hand, Evans & Stark (2002); Kaipio & Somersalo (2005); Marzouk & Najm (2009); Stuart (2010). The QoI is then described in a probabilistic sense, with closed-form expressions in some specific situations (linear forward model from the input to the observations, Gaussian distribution of the likelihood, *etc.*) or, more generally, by sampling the input space according to the posterior measure density with a Markov-chain Monte-Carlo (MCMC) technique, Brooks *et al.* (2013).

Note that the two strategies can be combined and most efficient methods use both an approximation basis for the input space and a Bayesian approach. As an example, a Karhunen-Loève-based approximation of the input space, subsequently sampled with a MCMC scheme, can be used, Efendiev *et al.* (2006); Marzouk & Najm (2009).

Besides the quality of the approximation basis or the relevance of the expert knowledge introduced as a prior in the estimation method, the performance of the recovery procedure critically depends on observations, both quantity and quality of which are key factors for an accurate inference. Among other things, including the type of acquired information, the sensors location is a critical point. In an operational situation, sensors are limited in number and are subjected to a series of constraints. For instance, in a fluid mechanics context, they are usually mounted on a solid wall, as opposed to measuring in the bulk flow. Optimal location of the sensors requires an exhaustive search and constitutes an NP-hard problem. Many approaches have been proposed to determine a reasonably good solution in a computationally acceptable way. These include Effective Independence, Kammer (1991), Optimal Driving Point, Sensor Set Expansion, Kinetic Energy Method and Variance Method, Bakir (2011); Meo & Zumpano (2005), the maximization of the smallest singular value of the Fisher Information Matrix (FIM), Alonso *et al.* (2004), compressed-sensing-inspired

---

<sup>1</sup>Note that for physical quantities exhibiting sharp discontinuities, dedicated approaches have been proposed relying on zonation or multi-resolution, see for instance Ellam *et al.* (2016), or multi-point statistics, Strebelle (2002).

approaches, Bright *et al.* (2013); Sargsyan *et al.* (2015); Brunton *et al.* (2016), FRAMESENSE, Ranieri *et al.* (2014) or the SENSORS-SPACE method introduced by the authors of the present work, Kasper *et al.* (2015).

The quality of the reconstruction from the sensor information largely depends on the *observability* of the QoI from the sensors. Clearly, if the trace of the QoI on the support of the point sensors vanishes, the observations do not bear information from the QoI and the estimation is not possible. In the case of linear dynamical systems, observability of the state vector from the sensors is quantified by studying the spectrum of the observability Gramian matrix. A similar situation occurs when the QoI is estimated within an approximation basis. In this case, the quality of the estimation of the approximation coefficients largely depends on the observability of the elements of the basis from the sensors. It may happen that the sensors are well informed by some approximation modes while being blind to others. As will be seen below in our examples, this situation is quite common and strongly impacts the recovery performance.

The motivation for the present work is the following: we advocate that an efficient estimation procedure should make use of prior expertise allowing an educated choice of a suitable approximation space *whose basis functions should be consistent with the way the associated coefficients will be estimated from the sensor measurements*. As will be shown below, deriving an approximation basis disregarding the way the coefficients are actually determined, as is common practice with the use of “generic” bases such as PCA/POD or K-SVD, may significantly affect the performance of the recovery method.

The paper is structured as follows. The general framework of inverse problems is introduced in Sec. 2, together with the main notations of the paper. Approximation of both an infinite- and finite-dimensional quantity is introduced and solution methods to estimate the coefficients associated with the approximation basis from observations are discussed. The need for an approximation basis aware of the sensors subsequently used to inform the coefficients is advocated and our observability-constrained basis learning approach is presented in Sec. 3. Both a deterministic and a stochastic estimation methods are presented, leading to a deterministic estimate or to a full probabilistic description of the QoI when considered as a random field. The proposed basis learning method allows both cases. Computational aspects are briefly discussed and an algorithm to learn the approximation basis is presented, together with the full solution method. Our approach is illustrated with the flow field estimation from a few wall-mounted point sensors in a configuration of numerically simulated two-dimensional flow over an open cavity, Sec. 4. Robustness of the recovery is studied with respect to noise in the measurements. Comparisons with standard recovery techniques relying on both PCA/POD and K-SVD show the superior performance of our approach. The paper concludes with closing remarks in Sec. 5.

## 2. Inverse problem

### 2.1. Problem statement

Let us introduce the problem and associated notations. In inverse problems, one typically wants to estimate a real-valued quantity of interest  $y \in \mathcal{V}_y$  from a set of information  $\mathcal{S}_s$ . Few pieces of information are typically available so that the  $n_s$  real-valued elements of  $\mathcal{S}_s$  are collected in a vector  $\mathbf{s} \in \mathcal{V}_s \subset \mathbb{R}^{n_s}$  of *observations*. The QoI can be a set of variables or a distributed quantity such as a field. The dimension of  $\mathcal{V}_y$  can be finite or infinite and we note  $n_y := \dim(\mathcal{V}_y)$  where it is understood that  $n_y$  can  $\rightarrow +\infty$  in which case  $\mathcal{V}_y$  is assumed to be a Hilbert space  $\mathcal{H}$  of real-valued functions defined over a domain  $\Omega$ ,  $\mathcal{H} \equiv L^2(\Omega; \mathbb{R})$ , and endowed with an inner product  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ .

Observations are typically scarce while the representation of the field of interest involves a large number of degrees of freedom (DOFs), so that  $n_s \ll n_y$ .

In many situations, a good approximation of  $y$  is sufficient for the application of interest. Further, the structure of the QoI along the coordinates  $\tau \in \Omega$ , the field  $y(\tau)$  is indexed upon, can often be exploited to lower the number of DOFs necessary for a good description. The field of interest can be approximated in a low dimensional linear subspace  $\mathcal{V}_D$  on the Grassmann manifold  $\text{Gr}(n_D, \mathcal{V}_y)$ :  $y \approx y^{(n_D)} \in \mathcal{V}_D \subset \mathbb{R}^{n_y}$ . Let  $\{d_l\}_l$  be a set of *modes* defining a basis of  $\mathcal{V}_D$ , the field of interest is approximated as

$$y(\tau) \approx y^{(n_D)}(\tau) = \sum_{l=1}^{n_D} x^l d_l(\tau). \quad (1)$$

The set  $D := \{d_1, d_2, \dots, d_{n_D}\}$  gathers the approximation modes and is termed the *dictionary*. Under this approximation, the field of interest is completely characterized in  $\mathcal{V}_D$  by the set of coefficients  $\mathbf{x} \in \mathbb{R}^{n_D}$ . The inverse problem then reformulates as, given a dictionary  $D$ , inferring  $\mathbf{x}$  from  $\mathbf{s}$ .

For the inference problem to admit a solution, observations  $\mathbf{s}$  must be informed by the field of interest. Introducing the observation form  $F^\infty : y \in \mathcal{V}_y \mapsto \mathbf{s} \in \mathcal{V}_s$ , observations are related to  $y$  via  $\mathbf{s} = F^\infty(y)$ . In practice, as discussed above, the QoI is approximated in a low dimensional space  $\mathcal{V}_D$ . Let  $F$  be a finite-dimensional continuous operator,  $F : y^{(n_D)} \in \mathcal{V}_D \mapsto \mathbf{s} \in \mathcal{V}_s$ , the observations obey

$$\mathbf{s} = F \left( y^{(n_D)} \right) + \boldsymbol{\eta}, \quad (2)$$

where  $\boldsymbol{\eta} \in \mathcal{V}_s$  denotes the contribution of  $y \setminus \mathcal{V}_D \in \mathcal{V}_y \setminus \mathcal{V}_D$  through  $F^\infty$  and measurement noise, if any.

## 2.2. Approximation basis

To allow for a reduction of the dimensionality of the quantity to infer, an approximation basis must be chosen, see Eq. (1). It should allow for an accurate approximation of  $y$  while being low dimensional. Prior or expertise knowledge on the QoI is typically used to guide the derivation of a suitable basis. In the context of an unknown spatially-distributed quantity over a spatial domain  $\Omega$ , the QoI is typically considered as a statistical realization of a random field. Consider the probability space  $(\Theta, \sigma, \mu)$ , where  $\Theta$  is the set of random elementary events  $\theta$ ,  $\sigma$  the  $\sigma$ -algebra of the events and  $\mu$  a probability measure. The space  $\mathcal{W} \equiv L^2(\Theta, \mu)$  of real-valued, second-order, random variables with inner product  $\langle \cdot, \cdot \rangle_{\mathcal{W}}$  is such that

$$\|u\|_{\mathcal{W}}^2 := \langle u, u \rangle_{\mathcal{W}} < +\infty, \quad \langle u, v \rangle_{\mathcal{W}} := \int_{\Theta} u(\theta) v(\theta) d\mu(\theta), \quad \forall (u, v) \in L^2(\Theta, \mu). \quad (3)$$

If the QoI is a realization  $y(\tau; \theta)$  of a second-order random field, it can be efficiently approximated in  $\mathcal{H} \otimes \mathcal{W}$  using the Hilbert-Karhunen-Loève decomposition, Levy & Rubinstein (1999):

$$y(\tau, \theta) = \bar{y}(\tau) + \sum_{l=1}^{\infty} \lambda_l^{1/2} \varphi_l(\tau) \psi_l(\theta), \quad \bar{y}(\tau) \equiv \mathcal{E}_\mu [y(\tau, \theta)] := \int_{\Theta} y(\tau, \theta) d\mu(\theta). \quad (4)$$

The random variables  $\{\psi_l\}_l$  are uncorrelated in the sense of  $\langle \cdot, \cdot \rangle_{\mathcal{W}}$ ,  $\langle \psi_l, \psi_{l'} \rangle_{\mathcal{W}} = 0$  if  $l \neq l'$ , and  $(\lambda_l, \varphi_l) \in \mathbb{R}^+ \times \mathcal{H}$  is the ordered (decreasing  $\lambda_l$ ) sequence of eigenvalues and eigenfunctions of the linear operator  $\mathcal{T}_y(u) : \mathcal{H} \rightarrow \mathcal{H} := \langle \kappa_y, u \rangle_{\mathcal{H}} = \mathcal{E}_\mu [(y - \bar{y}) \langle y - \bar{y}, u \rangle_{\mathcal{H}}]$ . The correlation kernel  $\kappa_y \in \mathcal{H} \otimes \mathcal{H}$  is defined as

$$\kappa_y(\tau, \tau') := \mathcal{E}_\mu [(y(\tau, \theta) - \bar{y}(\tau)) (y(\tau', \theta) - \bar{y}(\tau'))]. \quad (5)$$

For a known kernel  $\kappa$ , a popular choice of approximation basis of a realization in  $\mathcal{H}$  is the set of eigenfunctions  $\{\varphi_l\}_l$ . This is justified by the classic optimality result: letting  $y^{(n_D)}(\tau, \theta) := \bar{y}(\tau) + \sum_{l=1}^{n_D} \lambda_l^{1/2} \varphi_l(\tau) \psi_l(\theta)$  be the  $n_D$ -term truncation of the series in Eq. (4) and  $\|\cdot\|_{\mathcal{H} \otimes \mathcal{W}}$  the norm induced by the scalar product  $\langle \cdot, \cdot \rangle_{\mathcal{H} \otimes \mathcal{W}} := \mathcal{E}_\mu [\langle \cdot, \cdot \rangle_{\mathcal{H}}]$ , one infers

$$\left\| y - y^{(n_D)} \right\|_{\mathcal{H} \otimes \mathcal{W}} = \min_{\varphi_l \in \mathcal{H}, \psi_l \in \mathcal{W}} \left\| y - \bar{y} - \sum_{l=1}^{n_D} \lambda_l^{1/2} \varphi_l \psi_l \right\|_{\mathcal{H} \otimes \mathcal{W}}. \quad (6)$$

For a given number of terms retained, the approximation error is then of minimal norm, in the sense of Eq. (6). Approximating a given instance  $y(\tau; \theta)$  of the random field in the linear span  $\{\varphi_l\}_l$  hence minimizes the *average* spatial error norm  $\|y - y^{(n_D)}\|_{\mathcal{H}}$  in the sense of  $\mu$ .

In practice, the correlation kernel is rarely known. However, if realizations of the random field are available, the correlation kernel can be approximated by the empirical correlation. These available realizations are usually finite-dimensional, as produced by numerical simulations or from past inference problems. A realization  $y(\tau; \theta) \in \mathcal{V}_y$  can then be identified with  $\mathbf{y} \in \mathbb{R}^{n_y}$  and the empirical covariance is determined from

a collection of  $n_{\text{train}}$  known instances of the field  $\{\mathbf{y}_i\}_{i=1}^{n_{\text{train}}}$ . Letting  $Y := (\mathbf{y}_1^\circ \mathbf{y}_2^\circ \dots \mathbf{y}_{n_{\text{train}}}^\circ) \in \mathbb{R}^{n_y \times n_{\text{train}}}$  be the *training* matrix of centered known realizations  $\mathbf{y}_i^\circ := \mathbf{y}_i - \bar{\mathbf{y}}$ ,  $\bar{\mathbf{y}} := (n_{\text{train}})^{-1} \sum_{i=1}^{n_{\text{train}}} \mathbf{y}_i$ , the empirical covariance  $C_y \in \mathbb{R}^{n_y \times n_y}$  is given by

$$C_y := (n_{\text{train}} - 1)^{-1} Y Y^\top \succeq 0. \quad (7)$$

From the eigenvalues decomposition of the symmetric positive definite matrix  $C_y \stackrel{\text{eigen}}{=} \Phi \Lambda \Phi^\top$ , the approximation basis can be defined by the dictionary of the eigenvectors of  $C_y$ :  $D = \Phi = (\boldsymbol{\varphi}_1 \boldsymbol{\varphi}_2 \dots \boldsymbol{\varphi}_{n_y})$  and the training matrix can then express as  $Y = \Phi \Lambda^{1/2} \Psi^\top$  with a suitable matrix  $\Psi \in \mathbb{R}^{n_{\text{train}} \times n_y}$ . However, one typically has a limited training set while the realizations are finely discretized so that  $n_{\text{train}} < n_y$ . The matrix  $C_y$  is hence not full-rank and leads to a poor approximation of the true covariance. Further, it involves the dimension  $n_y$  of the discretized field, which typically is large, leading to a significant computational burden.

From now on, and without loss of generality, the data are assumed centered. With a collection of finite-dimensional (centered) training instances  $\{\mathbf{y}_i\}_i$ , and in the context of a second-order QoI,  $y \in L^2(\mathbb{R}^{n_y}) \otimes L^2(\mathbb{R}^{n_{\text{train}}}) \equiv \mathcal{H}_h \otimes \mathcal{W}_h$ ,  $\mathcal{H}_h \subset \mathcal{V}_y$ , the expression in Eq. (6) is reformulated in

$$\left\| Y - Y^{(n_D)} \right\|_F = \min_{\boldsymbol{\varphi}_l \in \mathcal{H}_h, \boldsymbol{\psi}_l \in \mathcal{W}_h} \left\| Y - \sum_{l=1}^{n_D} \lambda_l^{1/2} \boldsymbol{\varphi}_l \otimes \boldsymbol{\psi}_l \right\|_F. \quad (8)$$

The Eckart-Young theorem, Eckart & Young (1936), allows to directly identify  $\{\boldsymbol{\varphi}_l\}_l$  with the  $n_D$  dominant left singular vectors of  $Y$  and  $\{\boldsymbol{\psi}_l\}_l$  and  $\{\lambda_l\}_l$  with the corresponding right singular vectors and singular values. Letting  $\Sigma = \Lambda^{1/2}$  be the diagonal  $n_D \times n_D$  matrix of the positive square root of (positive) eigenvalues  $\{\lambda_l\}_l$ , the training matrix is best approximated with a rank- $n_D$  matrix as  $Y \stackrel{\text{rank-}n_D}{\approx} \Phi \Sigma \Psi^\top$ . The dictionary  $D$  is then given by the dominant left singular vectors and is efficiently determined from the solution  $\Psi$  of the following  $n_{\text{train}}$ -dimensional eigen problem:

$$(Y^\top Y) \Psi \stackrel{\text{eigen}}{=} \Psi \Lambda, \quad D = Y \Psi \Lambda^{-1/2} = (\mathbf{d}_1 \mathbf{d}_2 \dots \mathbf{d}_{n_D}). \quad (9)$$

The resulting approximation modes  $\{\mathbf{d}_l\}_l$  are sometimes termed *principal components* (PCA) or *empirical modes* in the context of the Proper Orthogonal Decomposition (POD), Holmes *et al.* (2012).

### 2.3. Direct approach

Once an approximation basis for the QoI is determined, the problem becomes to identify the coefficients associated with the approximation,  $\mathbf{y} \approx D \mathbf{x}$ , from the only pieces of available information  $\mathbf{s}$ . The link between measurements  $\mathbf{s}$  and the QoI  $\mathbf{y}$  is similar as in Eq. (2) and the estimation  $\hat{\mathbf{y}}$  can formulate as minimizing the  $L^2$ -norm of the data misfit:

$$\mathbf{y} \approx \hat{\mathbf{y}} = D \mathbf{x}, \quad \text{with} \quad \mathbf{x} \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^{n_D}} \|\mathbf{s} - F(D \tilde{\mathbf{x}})\|_2^2. \quad (10)$$

In the present study,  $F : \mathcal{H}_h \rightarrow L^2(\mathbb{R}^{n_s})$  is a bounded finite-dimensional linear operator and the measurements have finite energy,  $\mathbf{s} \in L^2(\mathbb{R}^{n_s}) \equiv \mathcal{S}_h \subset \mathcal{V}_s$ . Minimizing the data misfit then reduces to

$$\mathbf{x} \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^{n_D}} \|\mathbf{s} - F D \tilde{\mathbf{x}}\|_2^2. \quad (11)$$

Since one has limited measurements compared to the size of the dictionary,  $n_s \leq n_D$ , the linear system of normal equations associated with Eq. (11) is underdetermined. Additional knowledge or constraints must hence be considered to select among the infinite number of solutions. A popular choice is to retain the minimum  $L^2$ -norm solution. This is partly justified by the heuristic that one looks for the QoI of minimal energy compatible with the observations. The inferred field is then given by

$$\mathbf{y} \approx \hat{\mathbf{y}} = D \mathbf{x}, \quad \text{with} \quad \mathbf{x} = (F D)^+ \mathbf{s}, \quad (12)$$

where the superscript  $+$  denotes the Moore-Penrose pseudo-inverse. The lift-up operator  $L : \mathcal{S}_h \rightarrow \mathcal{H}_h$  from measurements to the estimated QoI is then  $L = D (F D)^+$  and  $\hat{\mathbf{y}} = L \mathbf{s}$ .



#### 2.4. Parsimonious approach

An alternative choice among the infinite number of solutions to problem (11) is to select a solution involving at most  $n_s$  distinct non-zero coefficients. Upon knowledge of the corresponding set of indices  $\mathcal{I} \subseteq \{1, 2, \dots, n_D\} \in \mathbb{N}^*$ ,  $|\mathcal{I}| \leq n_s$ , one can define the coefficient vector restricted to these non-zero coefficients:  $\mathbf{x}^{\mathcal{I}} \in \mathbb{R}^{|\mathcal{I}|}$ . Let  $D_{\mathcal{I}} \in \mathbb{R}^{n_y \times |\mathcal{I}|}$  be the restriction of the dictionary to its columns indexed by  $\mathcal{I}$ , the inferred field is given by the least squares solution

$$\mathbf{y} \approx \hat{\mathbf{y}} = D_{\mathcal{I}} \mathbf{x}^{\mathcal{I}}, \quad \text{with} \quad \mathbf{x}^{\mathcal{I}} = \left( (F D_{\mathcal{I}})^{\top} F D_{\mathcal{I}} \right)^{-1} (F D_{\mathcal{I}})^{\top} \mathbf{s}. \quad (13)$$

This approach results in a well-posed problem for  $\mathbf{x}$  *once*  $\mathcal{I}$  is determined.

Determining a sparse solution to a linear system of equations has attracted significant efforts in the recent past. A decisive step was made with the Compressed Sensing theory, Candès & Tao (2004); Candès *et al.* (2005); Donoho (2006); Candès *et al.* (2006), which has rigorously justified the formulation of an underdetermined linear problem  $A \mathbf{x} \approx \mathbf{b}$ ,  $A$  being a tall  $n \times p$  matrix,  $n < p$ , with a sparse solution as a constrained convex optimization problem (basis pursuit denoising problem, BPDN):

$$\mathbf{x} \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^p} \|A \tilde{\mathbf{x}} - \mathbf{b}\|_2^2 + \rho \|\tilde{\mathbf{x}}\|_1, \quad (14)$$

where  $\rho > 0$  is a regularization parameter.

In case the cardinality of  $\mathbf{x}$  is known or imposed, dedicated algorithms relying on a greedy approach, such as Orthogonal Matching Pursuit (OMP), Mallat & Zhang (1993), CoSaMP, Needell & Tropp (2008) or Iterative Hard Thresholding (IHT), Blumensath & Davies (2009), can be employed.

#### 2.5. Dictionary learning

As discussed in the previous section, if the solution to a linear system is sparse, efficient solution methods exist. However, the approximation of the field of interest  $\mathbf{y}$  with a sparse coefficient vector  $\mathbf{x}$  associated with the dictionary  $D$  is poor, in the general case. As an example, the field of interest  $\mathbf{y}$  is unlikely to admit a sparse representation in the PCA basis. The benefit of a well-posed problem for estimating the set of (sparse) coefficients is then lost by the poor resulting approximation. This situation asks for a dictionary specifically tailored so that the associated coefficients are more likely to be sparse. That is, instead of using a generic approximation basis such as PCA, one should determine the basis<sup>2</sup> with the specific aim of favoring sparse solutions. Many algorithms exist to this aim and we here briefly describe the K-SVD method, Aharon *et al.* (2006); Rubinstein *et al.* (2008), as it will be useful hereafter in the paper. For a given training set  $Y \in \mathbb{R}^{n_y \times n_{\text{train}}}$ , the dictionary learning problem consists in finding  $D \in \mathbb{R}^{n_y \times n_D}$  and  $X \in \mathbb{R}^{n_D \times n_{\text{train}}}$  such that

$$\{D, X\} \in \arg \min_{\tilde{D}, \tilde{X}} \left\| Y - \tilde{D} \tilde{X} \right\|_F^2 \quad \text{s.t.} \quad \|\mathbf{x}_i\|_0 \leq K, \quad \forall 1 \leq i \leq n_{\text{train}}, \quad (15)$$

where  $K \in \mathbb{N}^*$  is the maximum sparsity of the approximation for *any*  $\mathbf{y}_i$  from the training set.

The K-SVD algorithm essentially relies on an alternating minimization strategy and consists of two steps. In a Sparse Coding (SC) step, the matrix of coefficients  $X$  is updated solving Eq. (15) given  $D$ . While different methods can be used for this step, the original implementation of K-SVD relies on the OMP algorithm. By construction, each column vector of  $X$  is, at most,  $K$ -sparse. Once the coefficients are estimated, the dictionary  $D$  is updated in a Codebook Update (CU) step, exploiting the sparse structure of  $X$  and one iterates between these two steps until a convergence criterion is satisfied, see Aharon *et al.* (2006) for details.

---

<sup>2</sup>Strictly speaking, this is then not a basis since the elements  $\{\mathbf{d}_i\}_i$  are not linearly independent and constitute an overcomplete set.

## 2.6. Discussion

Two approximation bases have been briefly discussed above. If the correlation kernel of the QoI is known, or the empirical covariance of a representative training set can be estimated, the (Hilbert-) Karhunen-Loève decomposition can be used and leads to an approximation which is optimal in the sense of Eq. (8). Minimization of the data misfit allows to estimate the associated vector of coefficients  $\mathbf{x}$  from the measurements. This estimation method is hereafter referred to as the Principal Component Analysis-based (PCA) method. A major limitation of this approach is that, without additional hypotheses on the solution, the estimation problem is only well-posed for a coefficient vector of size  $n_D \leq n_s$ . Since measurements are typically scarce, this can severely limit the versatility of the approximation basis and the accuracy of this approach.

The additional sparsity constraint introduced in the K-SVD approach, and its variants, allows to use a *redundant* dictionary in the sense that  $\{\mathbf{d}_l\}_l$  constitutes a frame for  $\mathcal{H}_h$ . This redundancy property is key as it allows to rely on dictionary with more modes than measurements while still dealing with a solvable problem for  $\mathbf{x}$  thanks to the suitable sparsity constraint that one should have  $K \leq n_s \leq n_D$ . While a K-SVD-based dictionary is not optimal at approximating elements of the training set  $\{\mathbf{y}_i\}_i$  for a given size  $n_D$ , it however potentially allows to represent a QoI  $\mathbf{y}_i$  with elements from a *larger* dictionary,  $n_D \geq n_s$  than with the PCA approach where, in practice, one should have  $n_D \leq n_s$ . A potentially better approximation can hence be achieved.

## 3. Observability-constrained dictionary learning

### 3.1. Observability issue

While enjoying optimality in a different sense, dictionaries derived with the PCA or the K-SVD approaches discussed in the previous section, and more generally most dictionary learning techniques, rely on a minimization of some norm of the approximation residual of the (centered) training set,  $\|Y - DX\|$ . Once the dictionary is determined, the unknown QoI  $\mathbf{y}$  is estimated as  $\hat{\mathbf{y}} = D\mathbf{x}$ , with  $\mathbf{x}$  inferred from the measurements  $\mathbf{s}$  by minimization of the data misfit.

This is the standard approach but it nonetheless suffers an important drawback. Indeed, the dictionary is determined *a priori* from the training *fields*  $\{\mathbf{y}_i\}_i$  solely, while the actual QoI is, *in fine*, estimated from the *measurements*  $\mathbf{s}$ :  $\mathbf{y} \approx \hat{\mathbf{y}} = D\mathbf{x}(\mathbf{s})$  where the dependance of  $\mathbf{x}$  upon  $\mathbf{s}$  has been explicitly written. Instead, standard dictionary learning approaches rely on a training step where a given field is implicitly approximated as  $\mathbf{y} \approx \hat{\mathbf{y}} = D\mathbf{x}(\mathbf{y})$ . For instance, in the CU step of K-SVD, modes  $\{\mathbf{d}_l\}_l$  are updated given *exact* coefficient vectors  $\{\mathbf{x}_i\}_i$  as evaluated from the SC step. This is significantly different from the actual situation where the coefficient vector can only be evaluated from the data misfit.

From this situation, it results that the modes of the dictionary may be poorly informed by the available data in practice. The modes of the dictionary are determined without accounting for the observation operator  $F$ . In particular, the standard approach essentially consists in looking for  $D$  such that  $Y \approx DX$  (CU-like step) with  $X = D^+Y$  (SC-like step) and it can be interpreted as looking for a  $n_D$ -dimensional autoencoder  $D$  such that

$$Y \approx DD^+Y, \quad (16)$$

where the low-rank approximation is understood in the sense of (6) or (8).

Instead, a more consistent approach is to look for a  $n_D$ -dimensional autoencoder  $D$  such that  $Y \approx DX$  (CU-like step) with  $X = (FD)^+S$  (SC-like step), *i.e.*,

$$Y \approx D(FD)^+FY. \quad (17)$$

As an illustration of the potential issues of standard approaches of Sec. 2.3 and 2.4, consider the situation where the observation operator  $F$  is simply a restriction of the field of interest to a small set of points (point measurements). The vector of observations  $\mathbf{s}$  is then a subset of  $\mathbf{y}$  at indices  $\mathcal{J} \subseteq \{1, 2, \dots, n_y\} \in \mathbb{N}^*$ ,  $|\mathcal{J}| = n_s$ :  $\mathbf{s} = \mathbf{y}^{\mathcal{J}}$ . Since the dictionary is learned from  $Y$  only, it may happen that the modes vanish over the support of the sensors, so that  $FD = D^{\mathcal{J}} \approx [0]_{n_s \times n_D}$ . In this situation, the coefficients  $\mathbf{x}$  cannot be estimated from minimization of the data misfit in (11): the dictionary is *not observable* from the sensors.



### 3.2. Goal-oriented basis learning

To address the limitations of standard estimation techniques discussed above, we now introduce a Goal-Oriented Basis Learning (GOBAL) solution method. In a general, possibly infinite-dimensional, setting, we look for a redundant dictionary  $\Phi(\tau) = (\varphi_1 \varphi_2 \dots \varphi_{n_D})$ , and an estimation technique providing the function  $\psi$  from the measurements, minimizing the associated Bayes risk:

$$\begin{aligned} \{\Phi, \psi\} &\in \arg \min_{\tilde{\Phi}, \tilde{\psi}} \mathcal{E}_\mu \left[ \left\| y(\tau, \theta) - y^{(n_D)}(\tau, \theta; \mathbf{s}(\theta), \tilde{\Phi}, \tilde{\psi}) \right\|_{\mathcal{H}}^2 \right], \\ &= \arg \min_{\tilde{\Phi}, \tilde{\psi}} \mathcal{E}_\mu \left[ \left\| y(\tau, \theta) - \sum_{l=1}^{n_D} \tilde{\varphi}_l(\tau) \tilde{\psi}_l(\theta; \mathbf{s}(\theta)) \right\|_{\mathcal{H}}^2 \right], \end{aligned} \quad (18)$$

with  $y$  the centered data,  $\mathcal{E}_\mu[y(\cdot, \theta)] = 0$ .

**Remark 1.** This definition of the Bayes risk implicitly assumes perfect measurements  $\mathbf{s} = F y$ . In practice, noise affects the measurements so that this equation does not hold anymore and Eq. (2) should be accounted for instead. The definition of the Bayes risk above can be readily extended to account for a random noise with a given measure. Assuming the noise is statistically independent from the realizations  $\theta$ , one can introduce the probability space it belongs to, with an associated set of random elementary events  $\theta_\eta$  and probability measure  $\mu_\eta$ . Accordingly extending the definition of the expectation operator, the Bayes risk is then defined as

$$\mathcal{E}_{\mu, \mu_\eta} \left[ \left\| y(\tau, \theta) - y^{(n_D)}(\tau, \theta; \mathbf{s}(\theta, \theta_\eta)) \right\|_{\mathcal{H}}^2 \right]. \quad (19)$$

Restricting to the finite-dimensional framework, we let  $D \equiv \Phi = (\varphi_1 \varphi_2 \dots \varphi_{n_D})$ . From a collection of a finite number  $n_{\text{train}}$  of independent identically distributed (*iid*) realizations  $\{\mathbf{y}_i\}_i$  and associated (noisy) measurements  $\{\mathbf{s}_i\}_i$ , problem (18) can reformulate as looking for the dictionary that minimizes the *empirical Bayes risk*, with the Bayes risk definition (19):

$$\{D, X\} \in \arg \min_{\tilde{D}, \tilde{X}} \left\| Y - \tilde{D} \tilde{X}(S) \right\|_F^2, \quad (20)$$

where  $Y = (\mathbf{y}_1 \mathbf{y}_2 \dots \mathbf{y}_{n_{\text{train}}})$  and  $S = (\mathbf{s}_1 \mathbf{s}_2 \dots \mathbf{s}_{n_{\text{train}}})$ . The dictionary learning problem then admits a formulation similar to those discussed in Sec. 2. The coefficient vector  $\mathbf{x}$  associated with the estimation  $\hat{\mathbf{y}} = D \mathbf{x}$  is determined by minimization of the data misfit. Since the dictionary is redundant,  $n_D \geq n_s$ , constraints have to be considered for the estimation problem to be well-posed. Similarly to the parsimonious approach of Sec. 2.4, a sparsity constraint on  $\mathbf{x}$  is enforced and Eq. (20) finally reformulates as:

$$\left\{ \begin{array}{l} X \in \arg \min_{\tilde{X} \in \mathbb{R}^{n_D \times n_{\text{train}}}} \left\| S - F D \tilde{X} \right\|_F^2 \quad \text{s.t.} \quad \|\mathbf{x}_i\|_0 \leq n_s, \quad \forall 1 \leq i \leq n_{\text{train}}, \\ D \in \arg \min_{\tilde{D} \in \mathbb{R}^{n_y \times n_D}} \left\| Y - \tilde{D} X \right\|_F^2. \end{array} \right. \quad (21)$$

The decomposition (22) does not impose the basis elements to be orthogonal, in contrast with other methods such as PCA. This brings more flexibility in deriving an efficient approximation basis and results in superior potential performance. The dictionary  $D$  minimizes the training field error given the coefficients  $X$  (CU step) while the coefficients are such that they minimize the data misfit with sparsity constraints for a given dictionary (SC step).

As such, this formulation would however lead to a poor recovery procedure. Indeed, in the general case, it is unlikely that  $\mathbf{s}_i$  admits a good  $n_s$ -sparse approximation in the frame defined by  $F D$ . The naive approach of Eqs. (21) and (22) would hence lead to a large residual norm in the SC step and a poor overall

performance. Instead, one could explicitly introduce a sparsifying predictor operator  $P = (\mathbf{p}_1 \mathbf{p}_2 \dots \mathbf{p}_{n_D}) : \mathbb{R}^{n_D} \rightarrow \mathcal{S}_h \equiv L^2(\mathbb{R}^{n_s})$  learned from the training set of measurements such that  $\mathbf{s}_i$  admits a good  $n_s$ -sparse approximation in the linear span of  $\{\mathbf{p}_l\}_l, \forall 1 \leq i \leq n_{\text{train}}$ . The recovery procedure then relies on a pair of dictionaries  $\{P, D\}$  learned from training data with

$$\left\{ \begin{array}{l} \{P, X\} \in \arg \min_{\tilde{P}, \tilde{X}} \left\| S - \tilde{P} \tilde{X} \right\|_F^2 \quad \text{s.t.} \quad \|\mathbf{x}_i\|_0 \leq n_s, \quad \forall 1 \leq i \leq n_{\text{train}}, \\ D \in \arg \min_{\tilde{D}} \left\| Y - \tilde{D} X \right\|_F^2. \end{array} \right. \quad (23)$$

$$\left\{ \begin{array}{l} \{P, X\} \in \arg \min_{\tilde{P}, \tilde{X}} \left\| S - \tilde{P} \tilde{X} \right\|_F^2 \quad \text{s.t.} \quad \|\mathbf{x}_i\|_0 \leq n_s, \quad \forall 1 \leq i \leq n_{\text{train}}, \\ D \in \arg \min_{\tilde{D}} \left\| Y - \tilde{D} X \right\|_F^2. \end{array} \right. \quad (24)$$

The general procedure discussed above is hereafter termed the GLOBAL method and is now discussed in more details.

### 3.3. Implementation

As seen above, a given  $n_y$ -dimensional field  $\mathbf{y}_i$  is approximated in a  $n_D$ -dimensional space,  $\mathbf{y}_i \approx D \mathbf{x}_i$ , with  $n_D \leq n_y$ . The  $n_D$ -dimensional coefficient vector  $\mathbf{x}_i$  is informed by the measurements  $\mathbf{s}_i \in \mathbb{R}^{n_s}$ , with  $n_D \geq n_s$ . For a well-posed problem,  $\mathbf{x}_i$  is then imposed to be  $n_s$ -sparse, at most. For a given predictor  $P$ , it is then the solution of a sparse coding problem

$$\mathbf{x}_i \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^{n_D}} \|\mathbf{s}_i - P \tilde{\mathbf{x}}\|_2^2 \quad \text{s.t.} \quad \|\mathbf{x}_i\|_0 \leq n_s, \quad \forall 1 \leq i \leq n_{\text{train}}. \quad [\text{Sparse Coding}] \quad (25)$$

Since  $X$  is entirely determined from  $S$  for a given  $P$ , the solution to (24) immediately follows:

$$D = Y X^+. \quad [\text{Estimation Codebook Update}] \quad (26)$$

While  $D$  as given above minimizes the Frobenius norm of the residual  $Y - D X$  among all matrices of rank  $n_D \leq \min(n_{\text{train}}, n_y)$ , the Bayes risk can still be large. One should then exploit the remaining DOFs offered by  $P$  to improve upon this situation.

From (23),  $P$  is seen to minimize  $\left\| S - P \tilde{X} \right\|_F^2$  for a given  $\tilde{X}$ . We here adopt an approach similar to the Codebook Update step in K-SVD. For a given atom index  $l \in [1, n_D] \in \mathbb{N}^*$ , let  $\mathcal{I}^{(l)}$  be the set of column indices of non-vanishing coefficients,  $\mathcal{I}^{(l)} = \left\{ i, 1 \leq i \leq n_{\text{train}}; \tilde{X}_i^l \neq 0 \right\}$ . Let  $E^{(Y,l)} := Y_{\mathcal{I}^{(l)}} - D_{\setminus l} \tilde{X}_{\mathcal{I}^{(l)}}^l \in \mathbb{R}^{n_y \times |\mathcal{I}^{(l)}|}$ . The vector  $\tilde{\mathbf{x}}_{\mathcal{I}^{(l)}}^l$  can then be obtained from the rank-1 approximation of  $E^{(Y,l)}$ :

$$E^{(Y,l)} \stackrel{\text{SVD}}{=} \Phi \Sigma \Psi^H, \quad \tilde{\mathbf{x}}_{\mathcal{I}^{(l)}}^l \propto \Psi^{H,1}. \quad (27)$$

Iterating at random over the atoms  $1 \leq l \leq n_D$ , one can estimate the coefficients  $\tilde{X}$  informed from the QoI recovery error  $E^{(Y,l)}$ .

As for the coefficients  $X$  above, the predictor  $P$ , solution to (23), is then determined in the spirit of the K-SVD algorithm. For any atom index  $l$ , let the measurement error matrix be  $E^{(S,l)} := S_{\mathcal{I}^{(l)}} - P_{\setminus l} X_{\mathcal{I}^{(l)}}^l \in \mathbb{R}^{n_s \times |\mathcal{I}^{(l)}|}$ . The predictor  $\mathbf{p}_l$  is then determined as the minimizer of  $\left\| E^{(S,l)} - \mathbf{p}_l \tilde{\mathbf{x}}_{\mathcal{I}^{(l)}}^l \right\|_F$ , with the coefficient vector  $\tilde{\mathbf{x}}_{\mathcal{I}^{(l)}}^l$  achieving a low estimation error  $\|E^{(Y,l)}\|_F$  as given by (27). It yields

$$\mathbf{p}_l \propto E^{(S,l)} \left( \tilde{\mathbf{x}}_{\mathcal{I}^{(l)}}^l \right)^\top, \quad \|\mathbf{p}_l\|_2 = 1. \quad [\text{Features Codebook Update}] \quad (28)$$

This procedure leads to a predictor  $P$  minimizing the data misfit while accounting for associated coefficients  $X$  achieving a low QoI estimation error.

### 3.4. Computational aspects

While the procedure described above solely relies on simple linear algebra operations, it involves potentially large scale vectors, typically the training elements  $\mathbf{y}_i \in \mathbb{R}^{n_y}$ . In order to improve its computational efficiency, a number of modifications to the plain procedure is now discussed.

The estimation of the dictionary  $D$  is given by the product of the training field matrix  $Y$  and the Moore-Penrose pseudo-inverse of the real-valued  $n_D \times n_{\text{train}}$  matrix of coefficients  $X$ , Eq. (26). Efficient algorithms are available to compute  $X^+$ . However,  $X$ , as given by the Sparse Coding step in (25), is sparse, *i.e.*, each of its columns only has  $n_s$  non-zero entries, at most. This structure can be exploited when computing its pseudo-inverse. In this work, we rely on the QRGINV procedure introduced in Katsikis *et al.* (2011) which explicitly accounts for the sparse structure of  $X$  to speed up computations.

In the same objective of handling large matrices, it is important to notice that  $D$ , as given by Eq. (26), lies in the linear span of the columns of  $Y$ . Each dictionary atom  $\mathbf{d}_l$  then lies in the linear span of  $\{\mathbf{y}_i\}_{i=1}^{n_{\text{train}}}$ . Let consider the QR decomposition of the training matrix,  $Y \stackrel{\text{QR}}{=} Q^{(Y)} R^{(Y)}$ . Since  $\mathbf{d}_l \in \text{span}\{\mathbf{y}_i\}_i$ , there exists a matrix  $B^{(D)}$  such that  $D = Y B^{(D)}$ . The dictionary is the solution to Eq. (24), which reformulates as

$$B^{(D)} \in \arg \min_{\tilde{B}^{(D)} \in \mathbb{R}^{n_{\text{train}} \times n_D}} \left\| Q^{(Y)} R^{(Y)} - Q^{(Y)} R^{(Y)} \tilde{B}^{(D)} X \right\|_F^2, \quad D = Y B^{(D)}, \quad (29)$$

$$(30)$$

or, with  $C^{(D)}$  such that  $D = Q^{(Y)} C^{(D)}$ ,

$$C^{(D)} \in \arg \min_{\tilde{C}^{(D)} \in \mathbb{R}^{n_{\text{train}} \times n_D}} \left\| R^{(Y)} - \tilde{C}^{(D)} X \right\|_F^2. \quad (31)$$

Since  $n_{\text{train}} \leq n_y$ , the residual in Eq. (31) involves smaller matrices than in Eq. (24) and is hence faster to evaluate. The dimension  $n_y$  of  $\mathbf{y}$  then entirely vanishes from the algorithm, allowing to handle QoIs of arbitrary size, provided their QR decomposition can be computed in a prior (offline) step.

**Remark 2.** *The idea above is to decompose the, potentially very large, training matrix  $Y$  in a unitary matrix  $Q^{(Y)}$  and a smaller  $n_{\text{train}} \times n_{\text{train}}$  matrix  $R^{(Y)}$ , so that the estimation procedure involves  $R^{(Y)}$  only. A potentially faster alternative to the QR decomposition discussed above relies on an eigenvalue decomposition of a small  $n_{\text{train}} \times n_{\text{train}}$  symmetric matrix:*

$$Y^T Y \stackrel{\text{eigen}}{=} V^{(Y)} \left( \Sigma^{(Y)} \right)^2 \left( V^{(Y)} \right)^T, \quad (32)$$

so that there exists  $R^{(Y)}$  such that  $Y = Q^{(Y)} R^{(Y)}$ , with  $Q^{(Y)}$  a unitary matrix:

$$Q^{(Y)} \leftarrow Y V^{(Y)} \left( \Sigma^{(Y)} \right)^{-1}, \quad R^{(Y)} \leftarrow \Sigma^{(Y)} \left( V^{(Y)} \right)^T, \quad (33)$$

with  $\Sigma^{(Y)}$  a diagonal matrix with non-negative entries.

Disregarding the specific decomposition used, the generic terms  $Q^{(Y)}$  and  $R^{(Y)}$  will be employed in the remainder of the paper.

The last implementation aspect concerns the SVD in Eq. (27) as part of the Features Codebook Update step. The whole GLOBAL procedure relies on a block-coordinate descent approach, with alternate minimizations between the three steps. As such, high accuracy is not necessarily required at early steps of the iterative procedure and the SVD can safely be approximated. Further, one should notice that only the dominant (right) singular vector  $\Psi^{\text{H},1}$  is required. An inexpensive approximation of the solution to Eq. (27) is then given by a 1-iteration alternate minimization:

$$\mathbf{w} := E^{(Y,l)} \left( \mathbf{x}_{\mathcal{I}^{(l)}}^l \right)^T, \quad \tilde{\mathbf{x}}_{\mathcal{I}^{(l)}}^l \propto \mathbf{w}^T E^{(Y,l)}. \quad (34)$$

Using the  $Y = Q^{(Y)} R^{(Y)}$  decomposition discussed above, the approximation of the solution to (27) can be derived even more efficiently:

$$\mathbf{w} := E^{(R,l)} (\mathbf{x}_{\mathcal{I}^{(l)}}^l)^\top, \quad \tilde{\mathbf{x}}_{\mathcal{I}^{(l)}}^l \propto \mathbf{w}^\top E^{(R,l)}. \quad (35)$$

From this coefficient vector  $\tilde{\mathbf{x}}_{\mathcal{I}^{(l)}}^l$  as estimated from the residual  $E^{(R,l)}$ , the predictor  $\mathbf{p}_l$  follows from (28).

### 3.5. Sparse Bayesian Learning

#### 3.5.1. A probabilistic estimate

The GLOBAL method as discussed so far allows to learn a dictionary leading to a good sparse approximation of a set of data. From given observations  $\mathbf{s}$ , one gets a deterministic estimation  $\hat{\mathbf{y}} = D \mathbf{x}(\mathbf{s})$ . Instead, it is often useful to know the uncertainty associated with an outcome, in a broad sense. It allows, for instance, to define confidence intervals for the estimated QoI and identify spatial regions where the estimation is likely to be poor, hence possibly asking for more local information, such as additional sensors. To extend the general framework of the GLOBAL approach introduced in the previous section, we now adopt a Bayesian viewpoint.

In the general case where the sparsity of the solution is not known beforehand, the sparse coding step given the predictor operator  $P$  can formulate under the form of the BPDN problem, see also Eq. (14), where the compressible coefficient vector is given by:

$$\mathbf{x} \in \arg \min_{\tilde{\mathbf{x}} \in \mathbb{R}^{n_D}} \|\mathbf{s} - P \tilde{\mathbf{x}}\|_2^2 + \rho \|\tilde{\mathbf{x}}\|_1, \quad \rho > 0. \quad (36)$$

Observations are assumed to be affected with additive errors  $\boldsymbol{\eta}$  independent from the prediction:

$$\mathbf{s}(\theta, \theta_\eta) = P \mathbf{x}(\theta) + \boldsymbol{\eta}(\theta_\eta). \quad (37)$$

The components of the error are modeled as zero-mean *iid* Gaussian random variables. The measurement noise is assumed to affect similarly and independently each component of  $\mathbf{s}$  and the error  $\boldsymbol{\eta}$  is modeled as a multivariate, zero-mean, Gaussian random variable with diagonal covariance matrix  $\beta^{-1} I_{n_s}$ :  $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \beta^{-1} I_{n_s})$ .

The coefficient vector is modeled as a random vector endowed with a Laplace prior density:

$$p_{\mathbf{x}}(\mathbf{x} | \lambda) = \frac{\lambda}{2} \exp\left(-\frac{\lambda}{2} \|\mathbf{x}\|_1\right), \quad (38)$$

with  $\lambda > 0$ . In a Bayesian perspective, the coefficient vector as estimated from the observations is described by its posterior density

$$p_{\mathbf{x}}(\mathbf{x} | \mathbf{s}, \beta, \lambda) \propto p_{\mathbf{s}}(\mathbf{s} | \mathbf{x}, \beta) p_{\mathbf{x}}(\mathbf{x} | \lambda). \quad (39)$$

This Bayesian formulation provides a complete description of the coefficient vector. In particular, since  $\arg \min_{\tilde{\mathbf{x}}} \|\mathbf{s} - P \tilde{\mathbf{x}}\|_2^2 + \rho \|\tilde{\mathbf{x}}\|_1 = \arg \max_{\tilde{\mathbf{x}}} p_{\mathbf{s}}(\mathbf{s} | \tilde{\mathbf{x}}, \beta) p_{\mathbf{x}}(\tilde{\mathbf{x}} | \lambda)$ , with  $\lambda = \rho \beta$ , the solution to the BPDN problem (36) is equivalently obtained as the maximum a posteriori (MAP).

#### 3.5.2. Hierarchical prior

A priori values for the parameters  $\beta$  and  $\lambda$  are difficult to set and they should be learned from the data instead. We then adopt a hyperprior formulation and rely on an approach inspired from the Sparse Bayesian Learning (SBL) described in Tipping & Faul (2006) and Babacan *et al.* (2010). In particular, while the Laplace prior (38) is not conjugate to the noise model (37), a suitable three-stage hierarchical prior formulation allows to derive a closed-form expression of the MAP *via* a constructive sequential procedure.

To this end, the prior (38) is substituted with a product of univariate Gaussian random variables:

$$p_{\mathbf{x}}(\mathbf{x} | \boldsymbol{\gamma}) = \prod_{l=1}^{n_D} \mathcal{N}(x^l | 0, \gamma_l), \quad (40)$$

with  $\boldsymbol{\gamma} := (\gamma_1 \gamma_2 \dots \gamma_{n_D})$  and the variances  $\gamma_l \geq 0$ ,  $1 \leq l \leq n_D$ , modeled as

$$p_{\gamma_l}(\gamma_l | \lambda) = \Gamma(\gamma_l | 1, \lambda/2) = \frac{\lambda}{2} \exp\left(-\frac{\lambda \gamma_l}{2}\right), \quad (41)$$

where  $\lambda \geq 0$  has a Gamma prior with parameter  $\nu > 0$

$$p_{\lambda}(\lambda | \nu) = \Gamma(\lambda | \nu/2, \nu/2). \quad (42)$$

The hyperparameter  $\lambda$  has the same role as in Eq. (38) since

$$p_{\mathbf{x}}(\mathbf{x} | \lambda) = \int_{(\mathbb{R}^+)^{n_D}} p_{\mathbf{x}}(\mathbf{x} | \boldsymbol{\gamma}) p_{\boldsymbol{\gamma}}(\boldsymbol{\gamma} | \lambda) d\boldsymbol{\gamma} = (\lambda/2)^{n_D} \exp\left(-\lambda^{1/2} \sum_{l=1}^{n_D} |x^l|\right), \quad (43)$$

is of similar form as in Eq. (38). The hierarchical prior formulation then results in a Laplace prior on  $(\mathbf{x} | \lambda)$ . It is important to note that  $\boldsymbol{\gamma}$  is directly responsible for the sparsity of the solution vector  $\mathbf{x}$  since, for  $\gamma_l \rightarrow 0$ , the prior variance of  $x^l$  tends to zero and the posterior  $p_{x^l}(x^l | \mathbf{s})$  will hence be zero, almost surely.

The posterior then follows a multivariate Gaussian distribution which mean and variance admit a closed-form expression, Babacan *et al.* (2010),

$$\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma), \quad \boldsymbol{\mu} = \beta \Sigma P^T \mathbf{s}, \quad \Sigma = (\beta P^T P + \Lambda)^{-1}, \quad (44)$$

with  $\Lambda := \text{diag}(1/\gamma_l)$ ,  $1 \leq l \leq n_D$ .

Hyperparameters  $\{\beta, \boldsymbol{\gamma}, \lambda, \nu\}$  are determined from the data  $\mathbf{s}$  via a type-II evidence procedure by maximizing the (log)-marginal likelihood, *i.e.*, by integrating out  $\mathbf{x}$ :

$$\log p(\mathbf{s}, \boldsymbol{\gamma}, \lambda, \beta, \nu) = \log \int_{\mathbb{R}^{n_D}} p_{\mathbf{s}}(\mathbf{s} | \mathbf{x}, \beta) p_{\mathbf{x}}(\mathbf{x} | \boldsymbol{\gamma}) p_{\boldsymbol{\gamma}}(\boldsymbol{\gamma} | \lambda) p_{\lambda}(\lambda | \nu) p_{\beta}(\beta) d\mathbf{x}. \quad (45)$$

A careful analysis of this optimization problem allows an analytical treatment and a sequential solution procedure where the sparse quantities  $\boldsymbol{\mu}$  and  $\Sigma$  are iteratively updated, in a memory and computationally efficient way. In particular, manipulations only act on the active set, *i.e.*, predictors  $\mathbf{p}_l$  associated with a non-vanishing entry  $\mu^l$  through principled sequential additions or deletions of basis element candidates. Our implementation heavily borrows from the constructive Sparse Bayesian Learning (SBL) algorithm, Tipping & Faul (2006) and Babacan *et al.* (2010).

Once the coefficient random vector  $\mathbf{x}$  is determined, the full probability distribution  $p(\hat{\mathbf{y}} | \mathbf{s})$  of the inferred QoI can then be estimated and leads to the probabilistic estimate

$$\hat{\mathbf{y}} \sim \mathcal{N}(D \boldsymbol{\mu}, D \Sigma D^T). \quad (46)$$

### 3.6. Algorithm

The dictionary learning procedure is summarized in Algorithm 1.

In the online stage, *i.e.*, once *in situ* with only some measurements  $\mathbf{s}$  are available to estimate the QoI  $\mathbf{y}$ , dictionaries  $P$  and  $D$  learned in the offline (training) stage are used. The Sparse Bayesian Learning (SBL) procedure described in Sec. 3.5 can be employed to estimate the posterior mean  $\boldsymbol{\mu}$  and associated covariance  $\Sigma$  of the coefficients  $\mathbf{x}$  associated with the atoms of the dictionary  $D$ . The estimate  $\hat{\mathbf{y}}(\mathbf{s}; P, D)$  is finally given as in Eq. (46) in the form of a random field.

**Remark 3.** *If one is not interested in the statistics associated with the estimation, the (deterministic) coefficients  $\mathbf{x}(\mathbf{s}; P)$  can be more directly estimated as the solution to the Sparse Coding problem, Eq. (25), typically obtained with an Orthogonal Matching Pursuit (OMP) algorithm. The then deterministic estimate is hence  $\hat{\mathbf{y}}(\mathbf{s}; P, D) = D \mathbf{x}$ .*

---

**Algorithm 1:** Sketch of the GLOBAL algorithm for learning observable dictionaries – [Offline]

---

**Input:** A training set of  $n_{\text{train}}$  pairs  $(\mathbf{y}_i, \mathbf{s}_i)$ :  $Y$  and  $S$ .  
**Output:** Features dictionary  $P$  and QoI dictionary  $D$ .

- 1 Let  $n_s = |\mathbf{s}_i|$  for some  $i$  and  $n_D$  be the size of the dictionary, chosen such that  $n_s \leq n_D \leq n_{\text{train}}$ .  
Choose  $n_{\text{up}} \in [1, n_D]$  and  $r_{\text{max}}$ .
- 2 **Initialization:** Let  $Y = Q^{(Y)} R^{(Y)}$  (QR decomposition or eigenvalue problem, see Sec. 3.4).
- 3 From the  $n_D$ -truncated SVD of  $R^{(Y)}$ , initialize  $C^{(D)}$  with the dominant left singular vectors and  $X$  such that  $R^{(Y)} \stackrel{\text{trunc. SVD}}{\approx} C^{(D)} X$ . Initialize  $P = F Q^{(Y)} C^{(D)}$
- 4 Set  $\varepsilon_{\text{best}} \leftarrow \|R^{(Y)} - C^{(D)} X\|_F$  and  $r \leftarrow 0$ .
- 5 **while**  $r < r_{\text{max}}$  *and*  $\varepsilon$  *does not converge* **do**
- 6      $r \leftarrow r + 1$ .
- 7     **[Features Codebook Update]**  $\rightarrow P$
- 8     **for**  $1 \leq j \leq n_{\text{up}}$  **do**
- 9         Choose  $l \in [1, n_D] \subset \mathbb{N}^*$  at random,
- 10          $\mathcal{I}^{(l)} \leftarrow \{i, 1 \leq i \leq n_{\text{train}}; X_i^l \neq 0\}$ ,
- 11         Error matrices:  $E^{(R,l)} \leftarrow R_{\mathcal{I}^{(l)}}^{(Y)} - C_{\setminus l}^{(D)} X_{\mathcal{I}^{(l)}}^{\setminus l}$ ,  $E^{(S,l)} \leftarrow S_{\mathcal{I}^{(l)}} - P_{\setminus l} X_{\mathcal{I}^{(l)}}^{\setminus l}$ ,
- 12         Let  $\mathbf{w} = E^{(R,l)} (\mathbf{x}_{\mathcal{I}^{(l)}}^l)^\top$ ,
- 13         Fast estimation of temporary coefficients, see Eq. (35):  $\tilde{\mathbf{x}}_{\mathcal{I}^{(l)}} \leftarrow \mathbf{w}^\top E^{(R,l)}$ ,
- 14         Update the features dictionary:  $\mathbf{p}_l \leftarrow E^{(S,l)} \tilde{\mathbf{x}}_{\mathcal{I}^{(l)}}^\top$  and normalize to  $\|\mathbf{p}_l\|_2 = 1$ .
- 15     **[Sparse Coding]**  $\rightarrow X$
- 16     **for**  $1 \leq i \leq n_{\text{train}}$  [embarrassingly parallelizable] **do**
- 17         Use the SBL algorithm to evaluate  $\mathbf{x}_i$  from  $\mathbf{s}_i$  given  $P$ , see Sec. 3.5.
- 18         Alternatively, use OMP with sparsity constraint  $K = n_s$ .
- 19     Evaluate  $\varepsilon = \|R^{(Y)} - C^{(D)} X\|_F$ .
- 20     **if**  $\varepsilon < \varepsilon_{\text{best}}$  **then**
- 21          $\varepsilon_{\text{best}} \leftarrow \varepsilon$ ,  $P_{\text{best}} \leftarrow P$ ,  $C_{\text{best}}^{(D)} \leftarrow C^{(D)}$ .
- 22     **[Estimation Codebook Update]**  $\rightarrow C^{(D)}$
- 23     Update the dictionary:  $B^{(D)} \leftarrow X^+$  [possibly with QRGINV] and  $C^{(D)} \leftarrow R^{(Y)} B^{(D)}$ .
- 24 Learned dictionaries are finally  $P \leftarrow P_{\text{best}}$ ,  $C^{(D)} \leftarrow C_{\text{best}}^{(D)}$ .
- 25 If desired, assemble the physical QoI dictionary:  $D \leftarrow Q^{(Y)} C^{(D)}$ .

---

### 3.7. Numerical complexity

The complexity of the GLOBAL method, in its deterministic version, is now discussed.

We assume that the action of a  $m \times n$  matrix on a vector is of complexity  $2mn$  and denote the number of iterations of the algorithm by  $r$ . We let  $n_{\text{up}} \leftarrow n_D$ .

First, the training set matrix  $Y$  is decomposed into  $Q^{(Y)}$  and  $R^{(Y)}$  via the eigenvalue decomposition of  $Y^\top Y$ , see Eq. (32). The cost associated with forming the symmetric matrix is  $n_y n_{\text{train}}^2$ , while the eigenvalue decomposition takes  $\mathcal{O}(n_D n_{\text{train}}^2)$  operations. The cost of the initial decomposition of the training set is then

$$\mathcal{J}_{\text{decomp}} = n_y n_{\text{train}}^2 + \mathcal{O}(n_D n_{\text{train}}^2). \quad (47)$$

In the Estimation Codebook Update step, the numerical complexity associated with computing the pseudo-inverse of the full row-rank  $n_D \times n_{\text{train}}$ -matrix  $X$  is:

$$\mathcal{J}_{\text{ECU}} = \frac{3}{2} n_D^2 n_{\text{train}} + \frac{1}{3} n_D^3, \quad (48)$$



at leading order, Möller (1999). It is cubic in the small number of atoms and only linear in the number of training elements, hence allowing a fast evaluation. Note that this is the complexity of the plain vanilla pseudo-inverse and that the used QRGINV algorithm has a lower complexity and enjoys superior computational performance.

Relying on a Cholesky-based variant of the OMP algorithm for the Sparse Coding step, Cotter *et al.* (1999); Blumensath & Davies (2008), the number of operations is found to be, Rubinstein *et al.* (2008),

$$\mathcal{J}_{\text{SC}} = 2 n_s n_D n_{\text{train}} + 2 n_s^2 n_{\text{train}} + 2 n_s (n_D + n_{\text{train}}) + n_s^3. \quad (49)$$

Matrix operations involved in the Features Codebook Update step are of similar nature to those of the approximate K-SVD discussed in Rubinstein *et al.* (2008), so that one can borrow from their analysis. The number of operations per signal is here:

$$\mathcal{J}_{\text{FCU}} = 4 n_s n_{\text{train}}^2 + 6 n_{\text{train}} n_s n_D + 2 n_{\text{train}} n_s^2 + 2 n_D^2 (n_s + n_{\text{train}}). \quad (50)$$

The framework of this study is limited observations compared to the size of training set. Further, the size of the dictionary is a small multiple  $m \in \mathbb{N}^*$  of the number of sensors. One then has  $n_s \leq n_D \leq n_{\text{train}} \leq n_y$  and the total cost  $\mathcal{J}$  approximates as

$$\begin{aligned} \mathcal{J} &= \mathcal{J}_{\text{decomp}} + r (\mathcal{J}_{\text{ECU}} + n_{\text{train}} \mathcal{J}_{\text{SC}} + \mathcal{J}_{\text{FCU}}), \\ &\approx n_{\text{train}}^2 (n_y + 2 r n_s^2 (m + 1)) + \mathcal{O}(n_s n_{\text{train}}^2). \end{aligned} \quad (51)$$

Initialization of the training set aside (line 2 of the algorithm), the complexity of our core dictionary learning procedure is then seen to mainly scale with the square of the number of sensors and the square of the size of the training set. The bottleneck of the numerical complexity is the Sparse Coding step, which scales with  $n_{\text{train}}^2$ . Note however that this step is embarrassingly parallelizable among the  $n_{\text{train}}$  different elements of the training set, so that the walltime cost of this step, carried-out over  $n_{\text{train}}$  cores, only scales with  $n_{\text{train}}$ .

For a very large training set, alternatives to the present approach can be derived, relying on an online, or mini-batch, treatment of the training set, *e.g.*, inspired from the work of Mairal *et al.* (2010).

## 4. Numerical experiments

### 4.1. Configuration

The methodology introduced above is now demonstrated on a numerical simulation of the two-dimensional incompressible flow over an open cavity using our in-house code OLORIN. The cavity aspect ratio is 2 (deep cavity). The numerical grid comprises  $296 \times 128$  nodes on a structured, non-uniform, staggered mesh. The flow is from left to right on the plots and a laminar boundary layer flow profile is prescribed at the entrance of the numerical domain. The flow Reynolds number  $Re$ , based on the cavity length and the uniform flow rate velocity, is  $Re = 7500$ . The Navier-Stokes equations are solved using a prediction-projection method and a finite-volume second-order centered scheme in space. Viscous terms are evaluated implicitly while convective fluxes are estimated with a linear Adams-Bashford scheme. Details about the numerical settings may be found in Rizi *et al.* (2014); Podvin *et al.* (2006).

In the aim of being as realistic as possible, the retained sensors mimic actual devices routinely used in the flow control community. In particular, the sensors should be physically mounted on a solid wall and hence cannot be located in the bulk flow. They are here modeled as noisy point devices measuring the local wall shear stress, *i.e.*:

$$s^j(\mathbf{z}_s^{(j)}) \propto \left. \frac{\partial \|\mathbf{u}(\mathbf{z})\|_2}{\partial n} \right|_{\mathbf{z}_s^{(j)}}, \quad 1 \leq j \leq n_s, \quad (52)$$

with  $\mathbf{z}_s^{(j)}$  the location of the  $j$ -th sensor  $\mathbf{u}(\mathbf{z})$  the fluid velocity vector and  $n$  the distance to the wall along the local normal direction.

Sensor placement is of critical importance for the efficiency of any recovery procedure and is a well known problem. Many methods have been introduced to address this issue, as briefly discussed in Sec. 1. Here, no sensor placement method is employed. In an attempt to convincingly demonstrate the performance of our GOBAL method against alternatives approaches such as PCA and K-SVD, the sensors all these methods rely on are selected such that *the recovery performance of K-SVD is maximized*. Sets of  $n_s = 10$  sensors locations  $\{\mathbf{z}_s^{(j)}\}_{j=1}^{n_s}$  are sampled uniformly at random from the geometry of the cavity wall. For each realization of the sensor set, the recovery performance of the K-SVD-based approach is evaluated. The sensor location finally retained is the one leading to the best recovery performance. The performance of the K-SVD-based approach is thus significantly favored by this choice.

The objective of the estimation procedure is to recover the field of velocity magnitude. The training set comprises  $n_{\text{train}} = 1000$  instances,  $\{\mathbf{s}_i, \mathbf{y}_i\}_{i=1}^{n_{\text{train}}}$ , and the recovery performance is quantified in terms of relative error norm:

$$\varepsilon := \left\| Y^{\text{CV}} - \widehat{Y}^{\text{CV}}(S^{\text{CV}}) \right\|_F \left\| Y^{\text{CV}} \right\|_F^{-1}, \quad (53)$$

with  $\{\mathbf{s}_i^{\text{CV}}, \mathbf{y}_i^{\text{CV}}\}_{i=1}^{n_{\text{CV}}}$  a set of  $n_{\text{CV}} = 246$  instances not contained in the training set.

Unless explicitly specified, the GOBAL-based estimate is the posterior mean.

#### 4.2. Dictionaries

In addition to our method GOBAL, the performance of a recovery using both a PCA- and a K-SVD-based dictionary is studied. The PCA dictionary is given by the set of  $n_D$  dominant eigenvectors of the empirical correlation matrix  $Y^T Y$ , see Eq. (9), while the K-SVD dictionary is determined as the solution of Eq. (15), with  $K = n_s$  since this is the maximum number of atoms which can be informed by the available measurements. With either of these dictionaries, determined solely from  $Y$ , and not from  $\{S, Y\}$ , the estimation is given by  $\widehat{\mathbf{y}}_i = D \widehat{\mathbf{x}}_i$ ,  $1 \leq i \leq n_{\text{CV}}$ , with  $\widehat{\mathbf{x}}_i$  minimizing the data misfit in the sense of Eq. (11) (PCA approach) or *via* a Sparse Coding step, Eq. (21) (K-SVD approach).

Since K-SVD and GOBAL are iterative procedures, the PCA-based dictionary is used as initial condition in both these methods. Note that the mean field  $\overline{\mathbf{y}}$  has been subtracted from the training set so that dictionaries are learned from centered data.

In our cavity flow simulation, the flow is horizontal, from left to right. Spatial oscillations develop in the shear layer above the cavity, while the left-most part of the domain essentially remains still and the flow varies very little about its mean  $\overline{\mathbf{y}}$ . A large-scale structure within the cavity as produced by the flow recirculation. One can refer to Rizi *et al.* (2014) for a detailed discussion of the physics of this flow configuration.

The first and 10-th modes of the dictionaries from the different learning strategies are plotted in Fig. 1 for illustration. The main structures of the modes tend to reproduce the salient features of the flow. While the PCA and K-SVD dictionary look alike, at least for these two modes, the GOBAL-based dictionary significantly departs from these two, while still exhibiting similar trends.

#### 4.3. Dependence on $n_D$

A distinct feature of both the K-SVD and the present GOBAL approaches is that they rely on *redundant* approximations and allow to consistently inform a dictionary with more atoms than the number of sensors. At most  $n_s$  atoms are inferred from the data while the other atoms are associated with a vanishing coefficient. In contrast, the PCA-based approach informs the modes via Eq. (11) and, among the infinite number of solutions when  $n_D > n_s$ , selects the one with the least  $L^2$ -norm.

The recovery error  $\varepsilon$  achieved by these different methods is plotted in Fig. 2 as a function of the size  $n_D$  of the dictionary for  $n_s = 10$  sensors. When  $n_D = 10 = n_s$ , all dictionary modes can be informed from the available data. Yet, both the PCA- and the K-SVD-based approaches achieve a poor performance, with  $\varepsilon \approx 0.8$ . The conditioning of the observation of the dictionary atoms *via*  $FD$  is poor and prevents to accurately estimate  $X^{\text{CV}}$  from  $S^{\text{CV}}$ . In this particular configuration, the sparsity constraint  $\|\mathbf{x}_i\|_0 \leq n_s$  is always inactive and the sparse coding step, Eq. (21), essentially reduces to the minimization of the data misfit

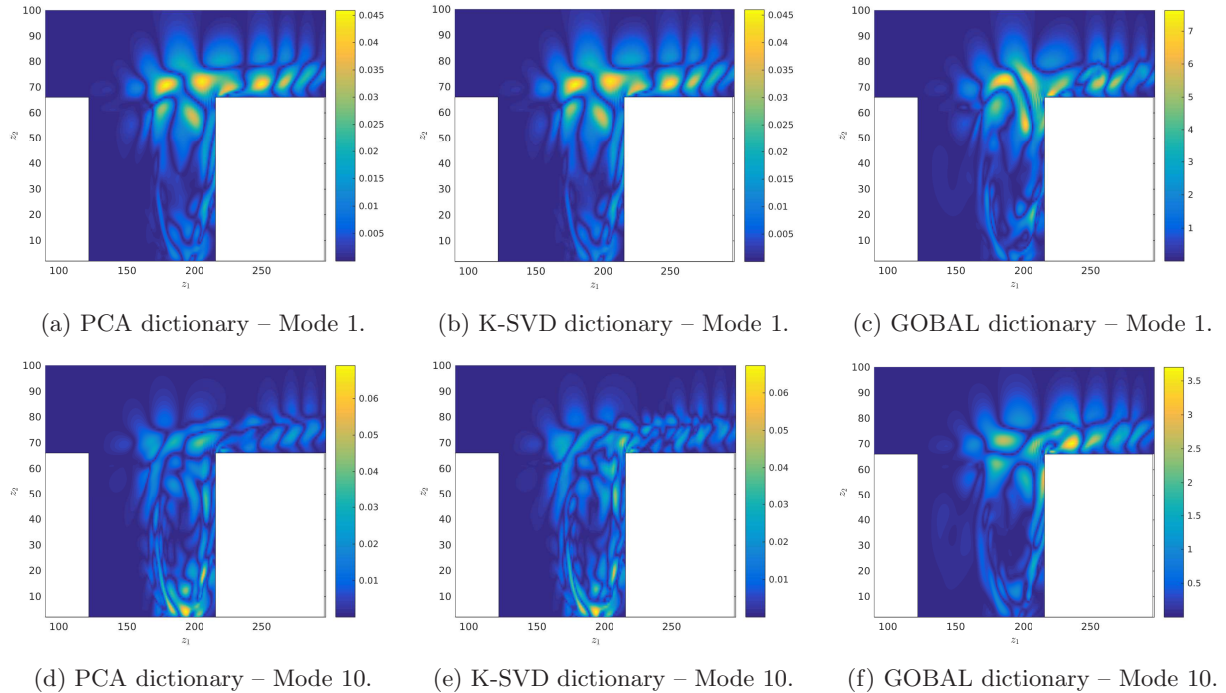


Figure 1: First (1st row) and 10-th (2nd row) modes of the different dictionaries. From left to right: PCA, K-SVD and GOBAL.




Figure 2: Influence of the size of the dictionary on the estimation error  $\varepsilon$ .  $n_s = 10$  sensors.

(11). However, this minimization is achieved *via* the greedy OMP algorithm and potentially deteriorates the recovery performance compared to (11). It results that, for  $n_D = n_s$ , the K-SVD approach performs worst. In contrast, the present GOBAL approach enjoys a well conditioned predictor operator  $P$  and accurately estimates  $X^{CV}$  and  $Y^{CV}$ , with  $\varepsilon \approx 0.2$ .

When the size of the dictionary grows beyond  $n_D = n_s$ , the performance of the PCA-based approach essentially gradually deteriorates as estimating  $\mathbf{x}_i$  from  $\mathbf{s}_i^{CV}$  *via* (11) becomes increasingly ill-posed. In




Figure 3: Example of field estimation using different dictionary learning strategies. The truth is in the top left and the sensors are denoted with large red dots. Plotted are the fluctuations of the flow field from its (empirical) mean as estimated from the training set.

contrast, the performance of the K-SVD-based approach remains roughly constant,  $\varepsilon \approx 0.7$ . However, since  $D$  stems from  $Y$  solely, observations  $\mathbf{s}_i^{\text{CV}}$  do not necessarily admit a sparse representation in the linear span of the columns of  $FD$  and the estimation of the coefficients *via* the sparse coding step (21) is likely to lead to a poor set of coefficients, hence a poor  $\varepsilon$ . The GOBAL approach, however, precisely addresses this issue and explicitly determines the predictor operator such that  $S^{\text{CV}}$  is likely to admit a compressible representation in  $\text{span}\{\mathbf{p}_l\}_l$ . The performance is seen to be good and to (slightly) improve when the size of the dictionary increases, allowing for better tailored atoms to be used for reconstructing a given field  $\mathbf{y}^{\text{CV}}$  from observations  $\mathbf{s}^{\text{CV}}$ .

The recovery performance is also illustrated in Fig. 3 where the different methods are employed to infer a field chosen at random from  $\{\mathbf{y}_i^{\text{CV}}\}_i$  from associated observations, as given by the shear stress measurement at the locations of the sensors. The sensors are denoted with red dots in the figure. The recovered field is plotted for the PCA- (top right), K-SVD- (bottom right) and GOBAL-based (bottom left) methods, along with the actual field (top left). The superior performance of the GOBAL approach is clearly visible.

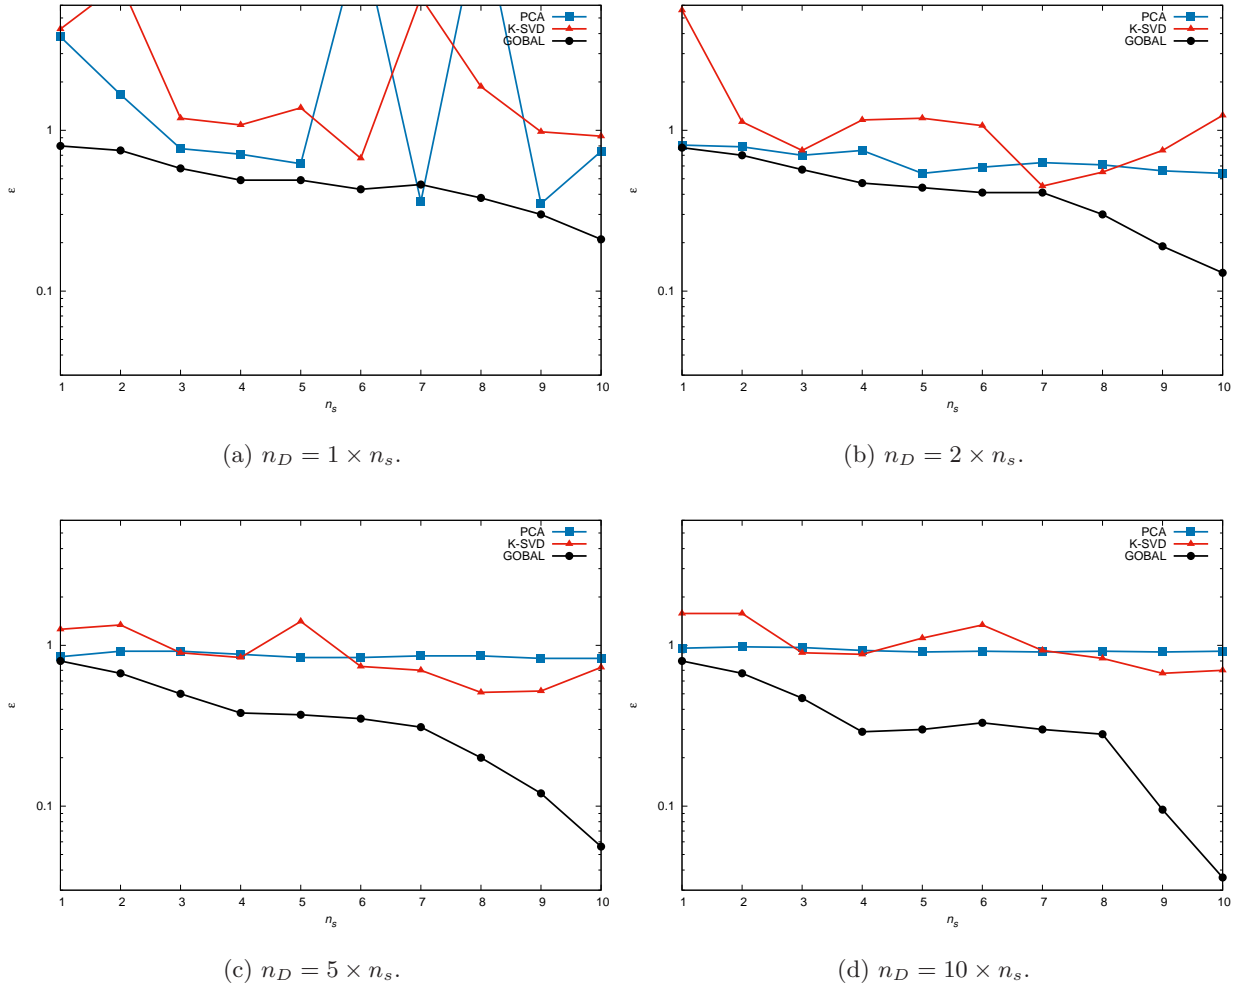


Figure 4: Influence of the number of sensors on the estimation error  $\varepsilon$  for several relative size of the dictionary.

#### 4.4. Dependence on $n_s$

The quality and quantity of information is pivotal to a reliable estimation. The impact of the quantity is here studied in terms of number of sensors. From the 10 sensors considered in the above section, the different recovery methods are now applied only relying on the first  $1 \leq n_s \leq 10$  sensors from the full set of 10 sensors. Note that the set of 10 sensors considered so far was determined, from a large set of sensor placement realizations, as that giving the best recovery performance for the K-SVD-based approach. Within this “best set”, they are not sorted or ranked in any way.

The influence of the number of sensors on the recovery accuracy is illustrated in Fig. 4 for different sizes of dictionary. It is important to note that the size of the dictionary here increases with the number of sensors, keeping the ratio  $m = n_D/n_s$  constant. For a given dictionary, the performance of the GOBAL approach is seen to significantly improve when more information becomes available (increasing  $n_s$ ). In contrast, the performance of the K-SVD-based approach hardly improves and remains rather poor, while the PCA-based performance does not improve and remains around a relative error  $\varepsilon \approx 1$ . Again, the GOBAL method leads to a situation where the observations are likely to admit a sparse representation in the linear span of the predictor operator, hence allowing an accurate recovery, while no such situation occurs with the other two approaches. It is remarkable that this superior performance of GOBAL is achieved despite the fact that the sensors are located to be “optimal” for *another* method, namely K-SVD.




Figure 5: Influence of the noise in the training and online measurements on the recovery accuracy.  $n_s = 10$ ,  $n_D = 50$ . Best seen in color (online version).

#### 4.5. Robustness

An important point in most, if not all, numerical methods is robustness. To be applicable in practice, solution methods should gracefully deteriorate the solution accuracy when input data are affected by a weak deviation from their nominal values. These deviations may come from round-off and truncation errors in the numerics. However, in the present context, the largest (by far) source of error, also termed noise, is expected to originate from the sensor measurements. Characterizing the influence of measurement noise onto the solution accuracy is necessary to demonstrate the benefit and performance of our method.

The noise is assumed to be a Gaussian vector with zero-mean and diagonal variance matrix, as discussed in Sec. 3.5.1,  $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \delta^2 I_{n_s})$ . Our GLOBAL method determines the posterior density of the estimate under an *additive* noise assumption, see Eq. (37). However, we here test our method in a context *different* from the one it is derived for: the accuracy of the solution is studied as the variance  $\delta^2 I_{n_s}$  of a *multiplicative* measurement noise varies, that is, the measurement data are actually given by

$$\mathbf{s} = (\mathbf{1}_{n_s} + \mathcal{N}(\mathbf{0}, \delta^2 I_{n_s})) \odot P \mathbf{x}, \quad (54)$$

with  $\odot$  the Hadamard product. Noise affects both the training data and the *in situ* measurements and a given QoI is then associated with *several* measurements. The GLOBAL algorithm is then run on an extended training set  $\left\{ Y, \left\{ S \left( \theta_{\eta}^{(j)} \right) \right\}_j \right\}$ . In this work, five realizations,  $j = 1, \dots, 5$ , of each measurement vector are considered.

The Figure 5 shows the evolution of the recovery accuracy  $\varepsilon$  as a function of the measurement noise standard deviation  $\delta$  in the training set  $S$ , and the noise standard deviation  $\delta^{\text{CV}}$  in the actual measurements  $S^{\text{CV}}$ . Several comments are in order. The accuracy of the recovery decreases when the noise variance  $\delta^{\text{CV}}$  increases. However, when the noise is strong ( $\delta^{\text{CV}} = 0.5$ ), the best performance is achieved with dictionaries trained with strong noise as well,  $\delta = 0.3$  and  $\delta = 0.5$ . If properly trained, the dictionaries learned from noisy data are hence more robust to noise in the measurements. However, this is at the price of a reduced performance if the measurement noise is low, in which case dictionaries trained without noise perform best. This shows how the learning can be made robust to the expected level of noise of *in situ* data and that the best performance is obtained when noise in the training step is as similar to that actually affecting the sensors as possible.

Another observation is that, for any noise level, the SBL-based version of GLOBAL always performs significantly better than its OMP-based counterpart. There are two reasons for this. First, the OMP-



variant of GLOBAL is here used with  $K = n_s$ , not allowing any redundancy in the data to inform the non-vanishing mode coefficients, and hence being prone to inaccuracy. Second, OMP is a greedy approach which converges to a local minimum. To address this issue, several initial conditions are usually used, and the best solution is retained. However, in practical situations as mimicked here, the “best” solution cannot be known, nor *a posteriori* verified. All solutions from the OMP step are hence equivalent, with no mean of discriminating the “good” from the “poor” solutions. We hence use a *unique* initial condition for OMP, identical for the training and the online step, with the risk for the OMP-based SC step to reach a poor local minimum. The Sparse Bayesian Learning variant of our SC step does not suffer these two limitations and potentially achieves a better overall recovery performance.

#### 4.6. Probabilistic estimate

The probabilistic estimate provided by our GLOBAL method is illustrated in Fig. 6. The posterior mean field,  $\hat{\mathbf{y}} = D \boldsymbol{\mu}$  (bottom left plot) is seen to compare well with the truth (top left plot). The diagonal part of the posterior standard deviation,  $[\text{diag}(D \boldsymbol{\Sigma} D^T)]^{1/2}$ , (bottom right plot) illustrates the spatial regions where noise in the measurements, and more generally the limitation introduced by the likelihood model, leads to uncertainty in the estimate. Clearly, regions with the largest estimated standard deviation are those where the flow exhibits the most fluctuations, *i.e.*, the downstream part of the shear layer and the main recirculation vortex within the cavity. In the top right plot, the (absolute value) reconstruction error of the posterior mean,  $|\mathbf{y} - \hat{\mathbf{y}}|$ , also exhibits the same general pattern as the posterior standard deviation. That is, regions of the estimated field identified as the most uncertain are indeed roughly those where the error is the largest. It is however important to note that the error is low, even its maximum  $\|\mathbf{y} - \hat{\mathbf{y}}\|_\infty \approx 1.6 \cdot 10^{-2}$ , compared to the field itself which lies between 0 and about 1.4. This is remarkable since the estimation method targets a reconstruction in the  $L^2$ -sense rather than in the  $L^\infty$ -sense.

## 5. Closing remarks

In this paper, an estimation technique for inferring high-dimensional quantities is introduced. It relies on learning an approximation basis, the dictionary, in which the quantity of interest is estimated from a set of observations by minimizing the data misfit. The original contribution of the present work is that the dictionary is learned with the specific constraint that it has to be *observable* from the available sensors. This makes estimation of the associated coefficients more accurate and reliable. An algorithm is presented and allows to derive an observable dictionary from a training set of observations and fields to be inferred. Estimation in the online phase, when *in situ*, is achieved via an algorithm combining Bayesian estimation and sparse recovery. It results in a probabilistic estimation of the quantity to be inferred. The formulation of the proposed recovery procedure does not involve the size of the high-dimensional quantity of interest but instead scales with the size of the training set, usually much smaller. It results in computationally efficient learning. The method is illustrated in the case of the two-dimensional flow over an open cavity whose velocity field is to be estimated from a small set of point sensors located on the wall. Comparisons with established PCA- and K-SVD-based estimation methods show the superior performance of the present approach.

This study clearly emphasizes the pivotal role of the observations quality onto the recovery performance. While sensor placement is a well-known aspect, the need for a *joint* derivation of the estimation procedure and determination of an approximation basis is clearly demonstrated in the present work. We strongly advocate learning *both* the representation basis and the recovery procedure. While not done here, sensor placement could, and actually should, be learned at the same time, as it strongly impacts the resulting observability properties of the dictionary. An integrated approach, accounting for all these aspects as early in the estimation procedure as possible, is key to an efficient and reliable method.

The present work is a first step towards a more general framework. In addition to coupling dictionary learning with sensor placement, future extensions include nonlinear observers where the measurements are not well approximated in the image space of a linear predictor operator. Derivation of the learning step




Figure 6: Probabilistic estimation with GOBAL. The truth is in the top left.

in this more general framework leads to a bi-level optimization problem and will be presented in a future paper.

The estimation procedure introduced here relies on the minimization of the reconstruction residual  $L^2$ -norm. While a common approach, this choice is questionable in some situations. For instance, while suitable for minimizing the estimation error in amplitude, this is a poor choice for minimizing the error in *phase*. As an example, consider a spatially-located feature well estimated but suffering a slight spatial shift. The associated  $L^2$ -norm estimation error would then be large while the recovery is “good” in the eyeball norm. This motivates alternative norms to be used. Among other choices, Wasserstein distances are commonly used in some other contexts, *e.g.*, image processing, and may be a more appropriate choice. To alleviate the significant computational overhead arising from working with Wasserstein distances, one could instead rely on a framework stable to deformations as, say, under a scattering transform preserving some invariants, to some extent, Mallat (2012).

More generally, an efficient estimation procedure first learns a suitable transformation such that, in the resulting framework, the inverse problem formulates as a problem easier to solve than its original counterpart.

Here, transformations promoting sparse solutions were exploited, allowing efficient sparse recovery tools to be used. Reformulating the estimation problem in a reproducing kernel Hilbert space (RKHS), with a suitable kernel and inner product learned from the data, is a more general approach and may allow to “untangle” a complex relationship between, say observations and approximation coefficients, paving the way for an efficient use of tools from the world of linear methods. Learning the manifold onto which the data lie, and a suitable kernel, allows application of techniques presented in this work to a more general context and provides new perspectives for inverse problems. Restricting to only one example, the current work may readily be extended into a *filter* for observing dynamical systems. Relying on *extended* observations, as collected over a given time-horizon, the present observable dictionary learning method can be reformulated in a suitable RKHS learned from the data. Since the estimation problem then relies on an extended set of observations, we expect superior performance as compared to the present static estimator.

These developments are the subject of on-going efforts.

## Acknowledgement

The first author (LM) gratefully acknowledges stimulating discussions with Youssef M Marzouk and Alessio Spantini.

## Conflict of interest

On behalf of all authors, the corresponding author states that there is no conflict of interest.

## References

- AHARON M., ELAD M. & BRUCKSTEIN A.M., K-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation, *IEEE Trans. Signal Proc.*, **54** (11), p. 4311–4322, 2006.
- ALONSO A.A., FROUZAKIS C.E. & KEVREKIDIS I.G., Optimal sensor placement for state reconstruction of distributed process systems, *AIChE Journal*, **50** (7), p. 1438–1452, 2004.
- BABACAN S.D., MOLINA R. & KATSAGGELOS A.K., Bayesian compressive sensing using Laplace priors, *IEEE Trans. Image Proc.*, **19** (1), p. 53–63, 2010.
- BAKIR P.G., Evaluation of Optimal Sensor Placement Techniques for Parameter Identification in Buildings, *Math. Comput. Appl.*, **16** (2), p. 456–466, 2011.
- BLUMENSATH T. & DAVIES M.E., Gradient pursuits, *IEEE Trans. Signal Proc.*, **56** (6), p. 2370–2382, 2008.
- BLUMENSATH T. & DAVIES M.E., Iterative hard thresholding for compressed sensing, *Appl. Comput. Harmonic Anal.*, **27** (3), p. 265–274, 2009.
- BRIGHT I., LIN G. & KUTZ J.N., Compressive sensing based machine learning strategy for characterizing the flow around a cylinder with limited pressure measurements, *Phys. Fluids*, **25** (12), p. 127102, 2013.
- BROOKS S., GELMAN A., JONES G.L. & MENG X.-L., ed., *Handbook of Markov Chain Monte Carlo*, Chapman & Hall CRC, 3rd edition, 2013.
- BRUNTON B.W., BRUNTON S.L., PROCTOR J.L. & KUTZ J.N., Sparse sensor placement optimization, *SIAM J. Appl. Math.*, **76** (5), p. 2099–2122, 2016.
- BUFFONI M., CAMARRI S., IOLLO A., LOMBARDI E. & SALVETTI M.-V., A non-linear observer for unsteady three-dimensional flows, *J. Comput. Phys.*, **227** (4), p. 2626–2643, 2008.
- CANDÈS E.J., ROMBERG J. & TAO T., Stable signal recovery from incomplete measurements, *Comm. Pure Appl. Math.*, **59**, p. 1207–1223, 2005.
- CANDÈS E.J., ROMBERG J. & TAO T., Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information, *IEEE Trans. Inform. Theor.*, **52** (2), p. 489–509, 2006.
- CANDÈS E.J. & TAO T., Decoding by linear programming, *IEEE Trans. Inform. Theory*, **51**, p. 4203–4215, 2004.
- COTTER S.F., ADLER J., RAO B.D. & KREUTZ-DELGADO K., Forward sequential algorithms for best basis selection, *IEE Proc. - Vis. Image Signal Process*, **146** (5), p. 235–244, 1999.
- DONOHO D.L., Compressed sensing, *IEEE Trans. Infor. Theo.*, **52** (4), p. 1289–1306, 2006.
- ECKART C. & YOUNG G., The approximation of one matrix by another of lower rank, *Psychometrika*, **1** (3), p. 211–218, 1936.
- EFENDIEV Y., HOU T. & LUO W., Preconditioning Markov chain Monte Carlo simulations using coarse-scale models, *SIAM J. Sci. Comput.*, **28** (2), p. 776–803, 2006.
- ELLAM L., ZABARAS N. & GIROLAMI M., A Bayesian approach to multiscale inverse problems with on-the-fly scale determination, *J. Comput. Phys.*, **326**, p. 115–140, 2016.
- EVANS S.N. & STARK P.B., Inverse problems as statistics, *Inverse Problems*, **18** (4), p. 55–97, 2002.
- HOLMES P., LUMLEY J.L., BERKOOZ G. & ROWLEY C.W., *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge: Cambridge University Press, 2nd edition, 402 p., 2012.

- KAIPIO J. & SOMERSALO E., *Statistical and Computational Inverse Problems*, Springer-Verlag New York, 2005.
- KAMMER D.C., Sensor placement for on-orbit modal identification and correlation of large space structures, *J. Guidance Control Dyn.*, **14** (2), p. 251–259, 1991.
- KASPER K., MATHELIN L. & ABOU-KANDIL H., A statistical learning approach for constrained sensor placement, In *American Control Conference*, Chicago, IL, USA, 6 p., July 1-3, 2015.
- KATSIKIS V.N., PAPPAS D. & PETRALIAS A., An improved method for the computation of the Moore-Penrose inverse matrix, *Appl. Math. Comput.*, **217** (23), p. 9828–9834, 2011.
- KHANINEZHAD M.M., JAFARPOUR B. & LI L., Sparse geologic dictionaries for subsurface flow model calibration: Part I. Inversion formulation, *Advances in Water Resources*, **39**, p. 106–121, 2012a.
- KHANINEZHAD M.M., JAFARPOUR B. & LI L., Sparse geologic dictionaries for subsurface flow model calibration: Part II. Robustness to uncertainty, *Advances in Water Resources*, **39**, p. 122–136, 2012b.
- LEROUX R., CHATELLIER L. & DAVID L., Maximum likelihood estimation of missing data applied to flow reconstruction around NACA profiles, *Fluid Dyn. Res.*, **47** (5), p. 051406, 2015.
- LEVY A. & RUBINSTEIN J., Hilbert-space Karhunen-Loève transform with application to image analysis, *J. Opt. Soc. Am. A*, **16** (1), p. 28–35, 1999.
- MAIRAL J., BACH F., PONCE J. & SAPIRO G., Online Learning for Matrix Factorization and Sparse Coding, *J. Mach. Learning Res.*, **11**, p. 19–60, 2010.
- MALLAT S., Group Invariant Scattering, *Comm. Pure Appl. Math.*, **65** (10), p. 1331–1398, 2012.
- MALLAT S. & ZHANG Z., Matching pursuits with time-frequency dictionaries, *IEEE Trans. Signal Process*, **41** (12), p. 3397–3415, 1993.
- MARZOUK Y.M. & NAJM H.N., Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems, *J. Comput. Phys.*, **228** (6), p. 1862–1902, 2009.
- MEO M. & ZUMPARO G., On the optimal sensor placement techniques for a bridge structure, *Engineering Structures*, **27** (10), p. 1488–1497, 2005.
- MÖLLER H. MICHAEL, Exact computation of the generalized inverse and least-squares solution, Universität Dortmund, Germany, *Tech. Rep.*, 1999.
- NEEDEL D. & TROPP J.A., CoSaMP: iterative signal recovery from incomplete and inaccurate samples, *Appl. Comput. Harmon. Anal.*, **26** (3), p. 301–321, 2008.
- PODVIN B., FRAIGNEAU Y., LUSSEYRAN F. & GOUGAT P., A reconstruction method for the flow past an open cavity, *J. Fluids Eng.*, **128** (3), p. 531–540, 2006.
- RANIERI J., CHEBIRA A. & VETTERLI M., Near-Optimal Sensor Placement for Linear Inverse Problems, *IEEE Trans. Signal Proc.*, **62** (5), p. 1135–1146, 2014.
- RIZI M.-Y., PASTUR L., ABBAS-TURKI M., FRAIGNEAU Y. & ABOU-KANDIL H., Closed-loop analysis and control of cavity shear layer oscillations, *Int. J. Flow Control*, **6** (4), p. 171–187, 2014.
- RUBINSTEIN R., ZIBULEVSKY M. & ELAD M., Efficient Implementation of the K-SVD Algorithm using Batch Orthogonal Matching Pursuit, Technion, *Tech. Rep.*, Haifa, Israël, 2008.
- SARGSYAN S., BRUNTON S.L. & KUTZ J.N., Nonlinear model reduction for dynamical systems using sparse sensor locations from learned libraries, *Phys. Rev. E*, **92** (3), p. 033304, 2015.
- SIROVICH L., Turbulence and the dynamics of coherent structures. Part I: coherent structures, *Quarterly Appl. Math.*, **45** (3), p. 571–591, 1987.
- STREBELLE S., Conditional Simulation of Complex Geological Structures Using Multiple-Point Statistics, *Math. Geology*, **34** (1), p. 1–21, 2002.
- STUART A.M., Inverse problems: a Bayesian perspective, *Acta Numerica*, **19**, p. 451–559, 2010.
- TIPPING M.E. & FAUL A.C., Fast marginal likelihood maximisation for sparse Bayesian models, In *Proc. 9th Int. Workshop on Artif. Int. Stat.*, Key West, FL, USA, 13 p., 3-6 Jan., 2006.
- YANG X., VENTURI D., CHEN C., CHRYSOSTOMIDIS C. & KARNIADAKIS G.E., EOF-based constrained sensor placement and field reconstruction from noisy ocean measurements: Application to Nantucket Sound, *J. Geophys. Res.*, **115** (C12), p. C12072, 2010.