# QUANTIFICATION OF UNCERTAINTY FROM HIGH-DIMENSIONAL SCATTERED DATA VIA POLYNOMIAL APPROXIMATION 

Lionel Mathelin*

LIMSI-CNRS, BP 133, 91403 Orsay, France; Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA

Original Manuscript Submitted: 06/24/2013; Final Draft Received: 01/27/2014


#### Abstract

This paper discusses a methodology for determining a functional representation of a random process from a collection of scattered pointwise samples. The present work specifically focuses onto random quantities lying in a high-dimensional stochastic space in the context of limited amount of information. The proposed approach involves a procedure for the selection of an approximation basis and the evaluation of the associated coefficients. The selection of the approximation basis relies on the a priori choice of the high-dimensional model representation format combined with a modified least angle regression technique. The resulting basis then provides the structure for the actual approximation basis, possibly using different functions, more parsimonious and nonlinear in its coefficients. To evaluate the coefficients, both an alternate least squares and an alternate weighted total least squares methods are employed. Examples are provided for the approximation of a random variable in a high-dimensional space as well as the estimation of a random field. Stochastic dimensions up to 100 are considered, with an amount of information as low as about 3 samples per dimension, and robustness of the approximation is demonstrated with respect to noise in the dataset. The computational cost of the solution method is shown to scale only linearly with the cardinality of the a priori basis and exhibits a $\left(N_{q}\right)^{s}, 2 \leq s \leq 3$, dependence with the number $N_{q}$ of samples in the dataset. The provided numerical experiments illustrate the ability of the present approach to derive an accurate approximation from scarce scattered data even in the presence of noise.


KEY WORDS: uncertainty quantification, least angle regression, high-dimensional model reduction, total least squares, alternate least squares, polynomial chaos

## 1. INTRODUCTION

With the growing available computational power, and as more efficient numerical methods become available, domains as diverse as engineering, chemistry, psychometrics, medicine, finance, or social sciences now heavily rely on simulation for the prediction of more and more complex phenomena, often combining multimodels and high accuracy requirement. The prediction capability of modern simulations is often such that a new bottleneck for accuracy has emerged from the lack of relevant boundary and/or initial conditions (BICs) as well as parameters intrinsic to the model of the system at hand, e.g., diffusivity, viscosity, etc. These sources of uncertainty are hereafter simply referred to as BICs. They are often poorly known and have to be estimated or modeled. This introduces modeling errors which often constitute the main source of lack of accuracy in the simulation chain. This situation has triggered a renewed interest for stochastic modeling where it is explicitly accounted for uncertainty in the model. The BICs may sometimes be modeled from first principles but are often approximated in a functional form involving a set of influencing parameters and identified from experimental measurements. However, more often than not, only relatively few

[^0]measurements are available, in particular when a significant number of parameters is of influence so that representing the BICs takes the form of a high-dimensional approximation problem.

If the random process, which output is to be represented in closed form, is driven by known equations, efficient techniques may be used to determine its representation. In the specific case of high-dimensional quantities, tensorbased representations have proved to be effective when applicable. In particular, low-rank approximations based on an a priori chosen separated representation can be efficiently derived, see [1-4] in the context of uncertainty quantification (UQ). If a closed-form model description of the process at hand is not available, one is typically left with approximating it from a finite collection of instances, hereafter termed samples. When the process is known only from a closed numerical code used as a black-box or if measurements can be made arbitrarily (design of experiments), some properties of approximation theory can be exploited. For instance, measurements may be taken at some particular locations in the parameter space, possibly associating a weight to them, so that the random Quantity of Interest (QoI) can be represented in the retained approximation basis with good accuracy using (sparse) quadrature techniques, [5]; see also [6] for an application to UQ. Anisotropy in the QoI may be exploited by biasing the quadrature weights [79]. In [10], an alternate least squares (ALS) technique to estimate the coefficients has been considered with samples lying on a tensor-product grid. Another situation of design of experiment arises in importance sampling where the Markov-Chain Monte Carlo algorithm requires a new sample at a specific proposed location. This control over the samples usually brings efficiency and allows one to approximate a reasonably behaved QoI with accuracy.

A different situation occurs when the data are scattered, with no ability to choose the set of samples nor to add a measurement. This is a common situation, typically arising when samples come from a past experiment or are costly to acquire so that new samples cannot be taken. In this context, one has to resort to a regression-based approach and the coefficients of the approximation are then solution of an optimization problem. This type of approach was considered in [11-13].

In the present work, the focus is specifically put on deriving a closed-form approximation of a high-dimensional quantity of interest from a small, uncontrolled, collection of its samples. This requires one to determine an approximation basis finely tuned to the data at hand and an efficient way of evaluating the associated coefficients. To this aim, we rely on the fact that, as a counterpart of the curse of dimensionality associated with high-dimensional problems, real applications often reward with a blessing of dimensionality. Indeed, in many cases, the QoI can be well approximated in a low-dimensional subspace of the solution space, sometimes involving orders of magnitude fewer degrees-of-freedom. This typically occurs when the solution exhibits some degree of sparsity in the retained functional space. Efficient techniques have been proposed in the recent past to take advantage of this situation and essentially consist of matching the approximation with the observational data while promoting a sparse coefficient set. This class of methods works well in many different contexts and has been recently applied to the UQ framework [14, 15]. These techniques rely on the compressed sensing theory, e.g., [16, 17], and may seem well suited for the present problem as they promote a low cardinality approximation of the QoI. However, they require one to handle a potentially huge representation basis, or dictionary, and associated optimization problem, leading to severe memory and computation limitations in the present high-dimensional context.

In this paper, we present a solution method combining the strength of different techniques, taking advantage of the sparsity of the representation in a suitable basis and allowing an efficient approximation of a well-behaved multivariate function with a low number of degrees-of-freedom hence compatible with a small experimental dataset. The driving principle is first to consider a tight approximation basis based on a priori knowledge on the QoI at hand and to rely on the available data to further refine it. In a nutshell, an initial approximation basis is first considered in the highdimensional model representation format (HDMR, [18, 19]), assuming it is suitable for representing the QoI. This initial basis is hereafter referred to as a priori basis. Next, available data are used to refine it by retaining only its most relevant basis functions through a constructive subset selection procedure based on a modification of the Least Angle Regression approach proposed in [20]. This a posteriori basis defines a skeleton from which a final basis is built and the associated coefficients are evaluated with an alternate least squares technique. The solution method allows us to approximate random variables as well as random fields and is shown here to outperform both sparse grids and tensored-based techniques.

The paper is organized as follows. The representation of a random quantity is central to the methodology discussed in this paper. Standard techniques for deriving a closed-form approximation of a random variable from a finite set of
samples are briefly recalled in Section 2. Similarly, different representation formats of functions in high-dimensional spaces are subsequently heavily used in the paper and a short discussion is given in Section 3. The proposed solution method is introduced and discussed in Section 4 and an algorithm is given. Scalability of the proposed approach together with its robustness with respect to noise in the data is also discussed. In Section 5, the present methodology is illustrated on a stochastic diffusion equation involving up to 100 dimensions and on the space-dependent solution of the shallow water equations with random parameters. Accuracy, robustness, and scalability of the proposed approach are shown. Concluding remarks close the paper in Section 6.

## 2. QUANTIFICATION OF UNCERTAINTY

Thanks to its pivotal role in the rest of the paper, the representation of a random quantity and standard ways of evaluating it in closed form from a discrete set of samples is now briefly discussed.

### 2.1 General Framework

Random quantities are defined on a probability space $\left(\Theta, \mathcal{B}_{\Theta}, \mu_{\Theta}\right)$ where $\Theta$ is the space of elementary events $\theta \in \Theta$, $\mathcal{B}_{\Theta}$ a $\sigma$-algebra defined on $\Theta$, and $\mu_{\Theta}$ a probability measure on $\mathcal{B}_{\Theta}$. To make the description of the problem amenable to a tractable representation, it is convenient to introduce a finite set of statistically independent random variables $\left\{\xi_{i}\right\}_{i=1}^{d}: \Theta \rightarrow \Xi_{i} \subseteq \mathbb{R}, \theta \mapsto \xi_{i}(\theta)$. The set of these $d$ random variables is defined on a probability space $\left(\Xi, \mathcal{B}_{\Xi}, \mu_{\Xi}\right)$ with $\Xi=\times_{i=1}^{d} \Xi_{i}=\xi(\Theta) \subseteq \mathbb{R}^{d}, \xi:=\left(\xi_{1} \ldots \xi_{d}\right), \mathcal{B}_{\Xi} \subset 2^{\Xi}$ a $\sigma$-algebra on $\Xi$ and $\mu_{\Xi}=\mu_{\Theta} \circ \xi^{-1}$ the probability measure on $\mathcal{B}_{\Xi}$. Since the physical process at hand relies on random quantities belonging to $\left(\Theta, \mathcal{B}_{\Theta}, \mu_{\Theta}\right)$, a suitable description of its output, or its solution in case the physical process is described by a known mathematical model, may be determined in $\left(\Xi, \mathcal{B}_{\Xi}, \mu_{\Xi}\right)$ as justified by the Doob-Dynkin lemma.

In this work, we restrict ourselves to random variables of physical significance, i.e., real-valued second-order variables satisfying

$$
\begin{equation*}
\stackrel{\mathscr{E}}{\theta}\left[u(\theta)^{2}\right]:=\int_{\Theta} u(\theta)^{2} \mathrm{~d} \mu_{\Theta}(\theta)=\int_{\Xi} u(\boldsymbol{\zeta})^{2} \mathrm{~d} \mu_{\Xi}(\boldsymbol{\zeta})=: \underset{\xi}{\mathscr{E}}\left[u(\xi)^{2}\right]<+\infty \tag{1}
\end{equation*}
$$

where $\mathscr{E}$ denotes the expectation operator and $u$ is the quantity of interest $(\mathrm{QoI})$. It is then natural to consider the space of square integrable functions $\mathcal{S}$ for describing real-valued functions of the random quantities:

$$
\begin{equation*}
\mathcal{S}:=L^{2}\left(\Xi, \mu_{\Xi}\right)=\left\{v: \Xi \rightarrow \mathbb{R}, \boldsymbol{\xi} \mapsto v(\xi) ; \mathscr{E}\left[v(\boldsymbol{\xi})^{2}\right]<+\infty\right\} \tag{2}
\end{equation*}
$$

Upon introduction of a natural inner product of $\mathcal{S}:\langle v, w\rangle_{L^{2}\left(\Xi, \mu_{\Xi}\right)}:=\int_{\Xi} v(\zeta) w(\zeta) \mathrm{d} \mu_{\Xi}(\zeta), \forall v, w \in \mathcal{S}$, and the associated norm $\|v\|_{L^{2}\left(\Xi, \mu_{\Xi}\right)}^{2}:=\langle v, v\rangle_{L^{2}\left(\Xi, \mu_{\Xi)}\right.}, \mathcal{S}$ is a Hilbert space. Further, we define $\langle v\rangle_{L^{2}\left(\Xi, \mu_{\Xi)}\right)}:=\mathscr{E}_{\xi}[v(\xi)]$. One can now rely on functional analysis results and take advantage of approximation theory techniques to characterize the output $u$. Introducing a Hilbertian basis $\left\{\psi_{k}\right\}_{k \in \mathbb{N}}$ of $\mathcal{S}$, the output can then be uniquely represented as $u(\xi)=$ $\sum_{\alpha} c_{\alpha} \psi_{\alpha}(\xi)$.

The basis $\left\{\psi_{\alpha}\right\}_{\alpha \in \mathbb{N}}$ is typically chosen orthonormal with respect to the inner product $\langle v, w\rangle_{L^{2}\left(\Xi, \mu_{\Xi}\right)}$. Orthonormality of the basis leads to $\left\langle\psi_{\alpha}, \psi_{\alpha^{\prime}}\right\rangle_{L^{2}\left(\Xi, \mu_{\Xi)}\right.}=\delta_{\alpha \alpha^{\prime}}, \forall \alpha, \alpha^{\prime} \in \mathbb{N}$, with $\delta$ the Kronecker delta, and the decomposition coefficients $\left\{c_{\alpha}\right\}$ then express as

$$
\begin{equation*}
c_{\alpha}=\left\langle u, \psi_{\alpha}\right\rangle_{L^{2}\left(\Xi, \mu_{\Xi}\right)}=\int_{\Xi} u(\zeta) \psi_{\alpha}(\zeta) \mathrm{d} \mu_{\Xi}(\zeta), \quad \forall \alpha \in \mathbb{N} \tag{3}
\end{equation*}
$$

For a given representation basis $\left\{\psi_{\alpha}\right\}$ of $\mathcal{S}$, the output $u(\boldsymbol{\xi})$ is entirely characterized by the set of coefficients $\left\{c_{\alpha}\right\}$. For computational purpose, the infinite dimensional representation is substituted with a finite dimensional approximation relying on a subset $\mathcal{J} \subset \mathbb{N}$ of the representation basis:

$$
\begin{equation*}
u(\xi) \approx \sum_{\alpha \in \mathcal{J}} c_{\alpha} \psi_{\alpha}(\boldsymbol{\xi}) \tag{4}
\end{equation*}
$$

Volume 4, Number 3, 2014

### 2.2 Computing a Data-Driven Approximation

As seen above, in many situations, a closed-form model of the QoI is not available or not reliable enough to be used and one can only rely on the sole available input-output information to approximate the output $u$. The solution method then consists of using a set of outputs given some inputs, i.e., samples of the process. One then looks for a functional form of the map between the set of random variables $\xi^{(q)}$ and the output value $u\left(\xi^{(q)}\right)=: u^{(q)}, \forall 1 \leq q \leq N_{q}$, where $N_{q}$ is the size of the available experimental set. Approximating the output under the functional form of Eq. (4) results in evaluating the coefficients $\left\{c_{\alpha}\right\}$ from $\left\{\left(\xi^{(q)}, u^{(q)}\right)\right\}_{q=1}^{N_{q}}, \xi^{(q)}=\left(\xi_{1}^{(q)} \ldots \xi_{d}^{(q)}\right)$.

### 2.2.1 Direct Evaluation

If the sampling can be controlled, in the sense that samples can be drawn arbitrarily, the popular Monte Carlo approach can be followed and the approximation coefficients are then estimated from

$$
\begin{equation*}
c_{\alpha}=\int_{\Xi} u(\boldsymbol{\zeta}) \psi_{\alpha}(\boldsymbol{\zeta}) \mathrm{d} \mu_{\Xi}(\boldsymbol{\zeta}) \approx \sum_{q} u\left(\xi^{(q)}\right) \psi_{\alpha}\left(\xi^{(q)}\right) \tag{5}
\end{equation*}
$$

Monte Carlo-based estimation is very robust and easy to implement but suffers from a slow $\mathcal{O}\left(N_{q}{ }^{-1 / 2}\right)$ asymptotic convergence rate. However, since the convergence rate does not depend on the dimensionality of the integral, this is a wise choice for very high-dimensional problems where other methods fail. Alternatively, quasi-Monte Carlo methods generate a low-discrepancy sequence of samples improving the convergence rate of the evaluation for moderateto high-dimensional problems.

For low to moderate dimensionality problems, the $d$-dimensional integral arising in Eq. (3) may be advantageously evaluated with a quadrature rule:

$$
\begin{equation*}
c_{\alpha}=\int_{\Xi} u(\zeta) \psi_{\alpha}(\zeta) \mathrm{d} \mu_{\Xi}(\zeta) \approx \sum_{q} w^{(q)} u\left(\xi^{(q)}\right) \psi_{\alpha}\left(\xi^{(q)}\right) \tag{6}
\end{equation*}
$$

where $\left\{w^{(q)}\right\}$ are the weights associated with the quadrature points $\left\{\boldsymbol{\xi}^{(q)}\right\}$ [21].

### 2.2.2 Regression

The above methods require some kind of control over the samples. If no experimental design can be exploited, a solution method is then to reformulate the evaluation of the coefficients as a minimization problem:

$$
\begin{equation*}
\boldsymbol{c}=\underset{\widetilde{\boldsymbol{c}} \in \mathbb{R}^{|\mathcal{J}|}}{\arg \min }\|\boldsymbol{u}-\Psi \widetilde{\boldsymbol{c}}\|_{2} \tag{7}
\end{equation*}
$$

with $\boldsymbol{c}=\left(c_{1} \ldots c_{|\mathcal{J}|}\right)^{T}, \boldsymbol{u}=\left(u^{(1)} \ldots u^{\left(N_{q}\right)}\right)^{T}, \Psi \in \mathbb{R}^{N_{q} \times|\mathcal{J}|}, \Psi_{q \alpha}=\psi_{\alpha}\left(\boldsymbol{\xi}^{(q)}\right)$, and $|\mathcal{J}|$ the cardinality of the approximation basis $\left\{\psi_{\alpha}\right\}_{\alpha \in \mathcal{J}}$. For a full column rank $\Psi$, the solution is given by $\boldsymbol{c}=\Psi^{+} \boldsymbol{u}$ which is typically evaluated using the Cholesky decomposition of the symmetric positive definite matrix $\Psi^{T} \Psi$ or the QR decomposition of $\Psi$. When the size of the dataset grows, this standard least squares (LS) problem may become computationally involved. The quasi-regression solution alleviates the computational burden and is given by

$$
\begin{equation*}
c_{\alpha}=\boldsymbol{\psi}_{\alpha}^{T} \boldsymbol{u} /\left\|\psi_{\alpha}\right\|_{2}^{2}, \quad \boldsymbol{\psi}_{\alpha}=\left(\psi_{\alpha}\left(\boldsymbol{\xi}^{(1)}\right) \ldots \psi_{\alpha}\left(\xi^{\left(N_{q}\right)}\right)\right)^{T}, \quad 1 \leq \alpha \leq|\mathcal{J}| \tag{8}
\end{equation*}
$$

Standard LS formulation as considered in Eq. (7) treats all predictors $\left\{\psi_{\alpha}\right\}_{\alpha=1}^{|\mathcal{J}|}$ the same way and uses the available data to estimate all the coefficients to produce an estimate with a low bias but often a large variance. As will
be discussed in Section 4.3.1, additional properties of the QoI may be exploited or imposed to the approximation coefficients. This class of approaches trades some increase in bias with a decrease in variance and often results in an improved accuracy. A suitable solution method then typically formulates as a penalized LS problem

$$
\begin{equation*}
\boldsymbol{c}=\underset{\widetilde{\boldsymbol{c}} \in \mathbb{R}^{|\mathcal{J}|}}{\arg \min }\|\boldsymbol{u}-\Psi \widetilde{\boldsymbol{c}}\|_{2}+\mathscr{J}(\widetilde{\boldsymbol{c}}) \tag{9}
\end{equation*}
$$

The properties of the penalized LS solution are driven by the choice of the function $\mathscr{J}$, the flexibility of which leads to a variety of solution techniques; see [22,23]. Since we have no control over the sampling strategy, we will rely on regression to estimate the approximation coefficients. The discussion of an efficient LS formulation in the present context is postponed to Section 4.

## 3. FUNCTIONAL REPRESENTATION OF RANDOM VARIABLES

### 3.1 Tensored Bases

As seen above, a random quantity is conveniently approximated in a Hilbertian basis $\left\{\psi_{k}\right\}$. If the random quantity is known, or expected, to exhibit a certain degree of smoothness along the stochastic space, a suitable and popular choice is to take advantage of this smoothness using a spectral-based approximation relying on polynomials. Early efforts toward this direction are the pioneering works of [24], who used univariate Hermite polynomials $\psi_{\alpha}\left(\xi_{i}\right)$ of zero-centered, unit variance, normal random variables $\xi_{i} \sim \mathcal{N}(0,1)$. These polynomials define an orthogonal basis of $L^{2}\left(\Xi_{i}, \mu_{\Xi_{i}}\right), \mu_{\Xi_{i}} \propto e^{-(1 / 2) \xi_{i}^{2}}$. Tensorization of univariate Hermite polynomials $\psi$ then leads to an orthogonal basis of $L^{2}\left(\Xi, \mu_{\Xi}\right)$ :

$$
\begin{equation*}
\left\langle\psi_{\alpha}, \psi_{\alpha^{\prime}}\right\rangle_{L^{2}\left(\Xi, \mu_{\Xi}\right)} \propto \int_{\Xi} \psi_{\alpha}(\zeta) \psi_{\alpha^{\prime}}(\zeta) e^{-(1 / 2)\left(\zeta^{T} \zeta\right)} \mathrm{d} \zeta \propto \delta_{\alpha \alpha^{\prime}} \tag{10}
\end{equation*}
$$

This can be extended to polynomials orthogonal with respect to different measures [25-27], and constitutes the so-called (generalized) polynomial chaos (PC) basis. A common practice is to consider an approximation space $\mathcal{S}_{p}$ spanned by polynomials of given maximum total degree $p$ :

$$
\begin{equation*}
\mathcal{S}_{p}=\operatorname{span}\left(\left\{\psi_{\boldsymbol{\alpha}}(\boldsymbol{\xi})=\psi_{\alpha_{1}}\left(\xi_{1}\right) \ldots \psi_{\alpha_{d}}\left(\xi_{d}\right)\right\} ; \boldsymbol{\alpha}=\left(\alpha_{1} \ldots \alpha_{d}\right), \sum_{i=1}^{d} \alpha_{i} \leq p\right) \tag{11}
\end{equation*}
$$

and the number of terms to be determined in the approximation (4) is then $|\mathcal{J}|=\binom{d+p}{d}$. We adopt the convention $\psi_{1} \equiv 1$. When the random quantity is not smooth enough for a low degree polynomial fit to be accurate, approximation schemes such as $h / p$-type refinement or multi-resolution analysis may be applied, see [28].

Some alternative representation formats specifically exploit the tensor-product structure of the Hilbert stochastic space $\mathcal{S}$ and approximates a $d$-variate function with a series of products of lower dimensional functions. Efficient algorithms allow us to determine the approximation coefficients of the representation by solving a series of lowdimensional problems while never considering the full-dimensional problem at once. A general presentation of tensorstructured numerical methods can be found in [29] while application to the approximation of a high-dimensional random quantity is considered in [3, 4, 10, 30]. For instance, a $d$-variate quantity may be approximated under a CANDECOMP-PARAFAC (CP) format [31, 32], with a sum of rank-1 terms, the simplest form of tensored-structure format:

$$
\begin{equation*}
u(\boldsymbol{\xi}) \approx \sum_{r=1}^{n_{r}} f_{1, r}\left(\xi_{1}\right) \ldots f_{d, r}\left(\xi_{d}\right) \tag{12}
\end{equation*}
$$

with $n_{r}$ the retained rank of the decomposition and $\left\{f_{i, r}\right\}_{i=1}^{d}$ univariate functions. Assuming $p$ th-order polynomials for $\left\{f_{i, r}\right\}$, the resulting cardinality of the approximation is $d n_{r} p$. It thus exhibits a linear dependence with the number of dimensions, in contrast with the exponential dependence of the PC. Alternative decomposition techniques,
easier to evaluate and numerically more stable than decomposition (12), such as the Tucker or Tensor-Trains, can be considered, see [29]. A tensored-structure format then constitutes a method of choice for deriving memory- and CPU-efficient approximation of high-dimensional quantities. They also lead to a low-cardinality basis $|\mathcal{J}|$ so that the conditioning of the approximation method remains good, in the sense that $|\mathcal{J}| \leq N_{q}$, a crucial feature for deriving a good approximation from the scarce available data.

### 3.2 High-Dimensional Model Representation

An efficient alternative to these tensored-structure formats for representing high-dimensional quantities is discussed in $[18,19]$. It consists of representing a quantity $u(\xi)$ with a sum of lower-dimensional terms accounting for increasing levels of interaction between the constitutive variables:

$$
\begin{equation*}
u(\xi)=f_{\emptyset}+\sum_{i=1}^{d} f_{i}\left(\xi_{i}\right)+\sum_{\substack{i, j=1, j>i}}^{d} f_{i j}\left(\xi_{i}, \xi_{j}\right)+\ldots+f_{12 \ldots d}\left(\xi_{1}, \ldots, \xi_{d}\right)=\sum_{\boldsymbol{\gamma} \subseteq\{1, \ldots, d\}} f_{\boldsymbol{\gamma}} \tag{13}
\end{equation*}
$$

where $f_{\gamma}$ are functions of $\mathcal{S}$ and depend only on a subset of variables $\xi_{\gamma}=\left\{\xi_{i}\right\}_{i \in \gamma}$ and $\gamma$ is a multi-index. This decomposition is exact, unique, and does not introduce any approximation. An important property is that the modes $\left\{f_{\boldsymbol{\gamma}}\right\}$ are mutually orthogonal: $\left\langle f_{\boldsymbol{\gamma}}, f_{\boldsymbol{\gamma}^{\prime}}\right\rangle_{L^{2}(\Xi, \mu \Xi)}=0, \forall \boldsymbol{\gamma} \neq \boldsymbol{\gamma}^{\prime} \subseteq\{1, \ldots, d\}$. The zeroth-order term $f_{\emptyset}$ accounts for the mean and is invariant across the entire domain $\Xi$, while the other modes are zero-mean:

$$
\begin{equation*}
f_{\emptyset}=\langle u\rangle_{L^{2}\left(\Xi, \mu_{\Xi}\right)}, \quad\left\langle f_{\gamma}\right\rangle_{L^{2}\left(\Xi, \mu_{\Xi}\right)}=0, \quad \forall \gamma \subseteq\{1, \ldots, d\} \backslash \emptyset \tag{14}
\end{equation*}
$$

The rationale behind the expected success of this so-called high-dimensional model representation (HDMR) is that many quantities of interest exhibit a significant dependence on low-dimensional groups of variables only, hence having negligible high order interaction decomposition terms. This leads to an efficient approximation of $u$ with only a low $N_{l}$-order HDMR: $u(\xi) \approx \sum_{\boldsymbol{\gamma} \subseteq\{1, \ldots, d\}} f_{\boldsymbol{\gamma}}\left(\xi_{\gamma}\right),|\boldsymbol{\gamma}| \leq N_{l}$. We denote $\mathcal{J}_{f}$ the set of retained modes, $\mathcal{J}_{f}:=\left\{\boldsymbol{\gamma} \subseteq\{1, \ldots, d\} ;|\boldsymbol{\gamma}| \leq N_{l}\right\}$.

Functions $\left\{f_{\gamma}\right\}$ are evaluated with the application of a set of commuting projections $\left\{\mathcal{P}_{i}\right\}$ onto the output $u$. The projection $\mathcal{P}_{i}$ eliminates the effect of variable $\xi_{i}$ while leaving the effect of the others unchanged. Letting $\mathcal{P} \emptyset$ be the identity operator on $\mathcal{S}$, we define $\mathcal{P}_{\boldsymbol{\eta}}=\prod_{i \in \boldsymbol{\eta}} \mathcal{P}_{i}, \forall \boldsymbol{\eta} \subseteq\{1, \ldots, d\}$. Functions $\left\{f_{\boldsymbol{\gamma}}\right\}$ can then be written [33],

$$
\begin{equation*}
f_{\boldsymbol{\gamma} \subseteq\{1, \ldots, d\} \backslash \emptyset}=\mathcal{P}_{\{1, \ldots, d\} \backslash \gamma} u-\sum_{\boldsymbol{\gamma}^{\prime} \subseteq \boldsymbol{\gamma}} f_{\boldsymbol{\gamma}^{\prime}}=\sum_{\boldsymbol{\gamma}^{\prime} \subseteq \boldsymbol{\gamma}}(-1)^{|\boldsymbol{\gamma}|-\left|\boldsymbol{\gamma}^{\prime}\right|} \mathcal{P}_{\{1, \ldots, d\} \backslash \boldsymbol{\gamma}^{\prime}} u, \quad f_{\emptyset}=\mathcal{P}_{\{1, \ldots, d\}} u \tag{15}
\end{equation*}
$$

Defining projections as $\mathcal{P}_{i} u(\xi)=\int_{\Xi_{i}} u\left(\xi_{1}, \ldots, \xi_{i-1}, \zeta^{\prime}, \xi_{i+1}, \ldots, \xi_{d}\right) \mathrm{d} \mu\left(\zeta^{\prime}\right)$, the measure $\mu$ determines the form of the projection. A popular choice consists of using $\mu=\mu_{\Xi_{i}}$ so that the Analysis of Variance (ANOVA) decomposition is obtained. An example of application of the HDMR representation to the approximation of a random quantity is presented in [9].

Remark 1. These different functional representations are not totally distinct. For instance, the PC basis defined in Eq. (11) can also be interpreted as a particular case of both HDMR and tensor-based expansion. For illustration, consider the following PC basis approximation space $\mathcal{S}_{p}=\operatorname{span}\left(\left\{\psi_{1}(\equiv 1), \psi_{2}\left(\xi_{1}\right), \psi_{2}\left(\xi_{2}\right), \psi_{3}\left(\xi_{1}\right), \psi_{2}\left(\xi_{1}\right)\right.\right.$ $\left.\left.\psi_{2}\left(\xi_{2}\right), \psi_{3}\left(\xi_{2}\right)\right\}\right)$. This corresponds to a HDMR representation with $N_{l}=2$ and $f_{\emptyset} \in \operatorname{span}\left(\psi_{1}\right), f_{1} \in \operatorname{span}\left(\psi_{2}\left(\xi_{1}\right)\right.$, $\left.\psi_{3}\left(\xi_{1}\right)\right), f_{2} \in \operatorname{span}\left(\psi_{2}\left(\xi_{2}\right), \psi_{3}\left(\xi_{2}\right)\right)$, $f_{12} \in \operatorname{span}\left(\psi_{2}\left(\xi_{1}\right) \psi_{2}\left(\xi_{2}\right)\right)$. Further, this can also be reformatted in a $n_{r}=3$-rank CP format, say with $f_{1,1} \in \operatorname{span}\left(\psi_{1}\right), f_{2,1} \in \operatorname{span}\left(\psi_{1}, \psi_{2}\left(\xi_{2}\right), \psi_{3}\left(\xi_{2}\right)\right), f_{1,2} \in \operatorname{span}\left(\psi_{2}\left(\xi_{1}\right)\right)$, $f_{2,2} \in \operatorname{span}\left(\psi_{1}, \psi_{2}\left(\xi_{2}\right)\right), f_{1,3} \in \operatorname{span}\left(\psi_{3}\left(\xi_{1}\right)\right)$ and $f_{2,3} \in \operatorname{span}\left(\psi_{1}\right)$.

## 4. QUANTIFYING UNCERTAINTY OF SCATTERED DATA

### 4.1 Setting Up the Stage

In the following, we will consider that the quantity of interest $u$ is a scalar-valued random field, indexed by space and/or time $\boldsymbol{x} \in \mathbb{R}^{d_{\boldsymbol{x}}}$ and depending on a set of random variables $\xi \in \mathbb{R}^{d}$. To approximate it, the only available piece of information is a collection of scattered samples $\left\{\boldsymbol{x}^{(q)}, \boldsymbol{\xi}^{(q)}, u^{(q)}\right\}_{q=1}^{N_{q}}$. In case these data come from an experimental context, the coordinates $\xi^{(q)}$ are not directly measurable. They are then inferred from auxiliary observations and depend on the modelization. ${ }^{1}$ Since the underlying random quantity $u$ is only known through these samples, no governing equation for the QoI can be exploited and, say, Galerkin projection-based weak-formulation methods cannot be employed. Further, these samples are scattered and do not follow a deterministic rule so that no deterministic sampling strategy can be assumed. Quadrature-based techniques can then not be applied either and one has to resort to regression to estimate the coefficients of the approximation in the retained basis $\left\{\psi_{\alpha}\right\}$. Standard $L^{2}$-regression solves Eq. (7) which is only well-posed for a matrix $\Psi$ such that $\Psi^{T} \Psi$ is invertible so that it requires the number of observations to be larger than the cardinality of the approximation basis, $N_{q} \geq|\mathcal{J}|$.

The choice of a good approximation basis in a general setting largely remains an open question. If one is given a dictionary of approximation functions, a priori selecting the best terms so that they can be evaluated from the data is a combinatorial optimization problem which algorithmic complexity quickly becomes intractable when the size of the dictionary grows. Dictionary-learning techniques require a training while availability of an independent training set cannot be assumed here.

The proposed approach is as follows. We separate the determination of an efficient representation format from the evaluation of the coefficients. We first choose an a priori general format for the approximation of $u$, Section 4.2. The selection of particular terms to be included in the approximation basis is left to a dedicated subset selection procedure which will further refine the approximation basis and make it as tight as possible, Section 4.3. A good a priori basis is motivated by results from compressed sensing which show that the number of samples necessary for accurately selecting the dominant basis functions of a $K$-sparse QoI (i.e., having $K$ non-zero coefficients in the retained approximation basis) varies as $K \log (|\mathcal{J}|)$ [34], illustrating the fact that it becomes increasingly difficult to select the best terms when the size $|\mathcal{J}|$ of the a priori dictionary increases. The subset selection hence produces an $a$ posteriori basis suitable for the data at hand. However, this basis is linear in its predictors as required by the selection method. To circumvent this limitation, the a posteriori basis is used as a skeleton only, of the best structure, and the final approximation of the QoI is evaluated with a different basis, of the same skeleton, but possibly nonlinear in its predictors, Section 4.4. A sketch of the solution method is shown in Fig. 1.


FIG. 1: Sketch of the solution method.

### 4.2 A Priori Choice of Representation of a Random Variable

We first focus on approximating a random variable and will discuss approximation of a more general random process in Section 4.8. The QoI is hence here a random variable $u(\boldsymbol{\xi})$.

In this work, we want to take advantage of the low order interactions of constitutive variables for many quantities of practical interest as mentioned in Section 3.2. Previous works have shown evidence of this low interaction configuration in various situations [9, 18, 19], and the QoI is hence chosen to be approximated under the HDMR form,

[^1]Eq. (13). An example is considered in Appendix A and demonstrates that a general HDMR format approximation with a tensor-based description of the interaction modes $\left\{f_{\gamma}\right\}$ involved in the HDMR may compare favorably with a full tensor-based approximation in terms of required number $\left|\mathcal{J}_{\text {prior }}\right|$ of basis functions for a given reconstruction accuracy, even for reasonably large dimensional problems. This motivates our choice of an HDMR format for the $a$ priori, data-independent, basis.

### 4.3 Subset Selection

We now build upon from the a priori basis and further improve it with an a posteriori, data-driven, procedure.

### 4.3.1 A Direct Approach

As discussed in Section 2.2, different techniques may be used to compute the coefficients of an approximation. In the case considered in this paper, the available data are scarce while the cardinality $\left|\mathcal{J}_{\text {prior }}\right|$ of the a priori approximation basis may be large, in particular when the dimensionality $d$ of the problem is large. It can then result in an illposed problem where one has to estimate $\left|\mathcal{J}_{\text {prior }}\right|$ coefficients for each stochastic mode $\lambda_{n}$ from $N_{q} \ll\left|\mathcal{J}_{\text {prior }}\right|$ pieces of information. However, this situation often only reflects our lack of knowledge on the quantity at hand and how conservative this naive approximation method is. Indeed, high-dimensional problems are often intrinsically sparse and lower dimensional. In the present setting, it is likely that many dimensions actually hardly contribute to the approximation and that representing the dependence of the QoI along only a subset of the dimensions yields an acceptable accuracy. In our a priori HDMR representation, it means that many interaction modes $\left\{f_{\gamma}\right\}$ can be discarded without significantly affecting the accuracy. The challenge for an efficient solution method is then to reveal and exploit the low-dimensional manifold onto which a good approximation of the solution lies. As an illustration, if $u(\boldsymbol{\xi})=g\left(\xi_{i}\right)$ was depending only on one dimension $i, i \in\{1, \ldots, d\}$, information theory allows us to show that one only requires $m+1+\left\lceil\log _{2} d\right\rceil$ function evaluations to approximate a sufficiently smooth function $g \in C^{s}$, having $s$ continuous derivatives, so that $\|u-\widehat{u}\|_{C\left(\Xi_{i}\right)} \leq a h^{s}, h:=1 / m$, where $a \geq 0$ is related to a norm of $g$ [35]. This number of samples actually is directly related to the number of information bits required to represent the integer $i \in[1, d]$.

While determining which interaction modes are dominant is an NP-hard problem in general, recent results have shown that a good estimation of the best subset can be obtained as the solution of a convex optimization problem. In particular, the LASSO formulation [36] has been proven effective. One of its formulations, referred to as basis pursuit denoising, writes

$$
\begin{equation*}
\boldsymbol{c}=\underset{\widetilde{\boldsymbol{c}} \in \mathbb{R}^{\left|\mathcal{J}_{\text {prior } \mid}\right|}}{\arg \min }\|\widetilde{\boldsymbol{c}}\|_{1} \quad \text { s.t. } \quad\|\boldsymbol{u}-\Psi \widetilde{\boldsymbol{c}}\|_{2} \leq \epsilon \tag{16}
\end{equation*}
$$

with $\Psi$ the matrix of evaluations of the approximation basis and $\epsilon$ the approximation residual. Efforts from the signal processing community, where the theory supporting these results is termed compressed sensing, have demonstrated its good recovery properties in the case where $N_{q}<\left|\mathcal{J}_{\text {prior }}\right|$, e.g., [17, 37, 38]. In particular, this formulation achieves provable and robust recovery bounds. ${ }^{2}$

The compressed sensing technique was proved very effective and is now being applied in many areas, including uncertainty quantification, $[14,15]$. However, standard implementations of the algorithm require the sensing matrix $\Psi$ to be available. This bears an intrinsic limitation when it comes to high-dimensional problems as it requires the use of the whole dictionary at once from which to select the basis functions associated with the dominant coefficients. While effective, this approach is not deemed tractable for high-dimensional problems, neither in terms of storage requirement nor CPU burden.

[^2]
### 4.3.2 A Progressive Selection

To circumvent the issues identified above, we here use a bottom-to-top approach which achieves a forward stagewise regression by progressively revealing important basis functions. Introduced by [20, 23], the least angle regression selection (LARS) technique relies on analytical solutions to speed up computations and essentially follows the piecewise linear regularization path of the LASSO. ${ }^{3}$ One advantage of LARS over other techniques is that the potential dictionary is never stored nor used as a whole. A LARS approach in the UQ framework was also considered in [40].

We consider the following polynomial approximation $\widetilde{f}_{\gamma}$ of $f_{\gamma}$ :

$$
\begin{equation*}
f_{\boldsymbol{\gamma}}\left(\left\{\xi_{i}\right\}_{i \in \boldsymbol{\gamma}}\right) \approx \widetilde{f}_{\boldsymbol{\gamma}}\left(\left\{\xi_{i}\right\}_{i \in \boldsymbol{\gamma}}\right):=\sum_{\boldsymbol{\alpha},|\boldsymbol{\alpha}| \leq \widetilde{p}} c_{\boldsymbol{\gamma}, \boldsymbol{\alpha}} \psi_{\boldsymbol{\alpha}}\left(\left\{\xi_{i}\right\}_{i \in \boldsymbol{\gamma}}\right), \quad \psi_{\boldsymbol{\alpha}}=\prod_{i \in \boldsymbol{\gamma}} \psi_{\alpha_{i}}\left(\xi_{i}\right), \tag{17}
\end{equation*}
$$

with $\boldsymbol{\alpha}=\left(\alpha_{i}, i \in \gamma\right), \alpha_{i} \in\{1, \ldots, \widetilde{p}\}$. Interaction modes $\left\{f_{\gamma}\right\}$ are then approximated in $\mathbb{P}_{\widetilde{p}}$, the space of polynomials with maximum total degree $\widetilde{p}$, by modes $\left\{\widetilde{f}_{\gamma}\right\}$ linear in their coefficients.

In the present framework, the HDMR approximation format naturally leads to groups of predictors whose importance in describing the QoI $u$ follows a similar trend. These groups are defined by the subsets $\left\{\mathcal{J}_{\gamma}\right\}$ of predictors which belong to a given interaction mode $f_{\gamma}, \mathcal{J}_{\gamma}=\left\{\psi_{\alpha}\left(\left\{\xi_{i}\right\}_{i \in \gamma}\right)\right\}$, and are likely to be strongly correlated. For instance, if the QoI exhibits a strong dependence on a given dimension $\xi_{j}$, one then wants to incorporate the whole set of predictors $\left\{\psi_{\alpha}\left(\left\{\xi_{i}\right\}_{i \in \gamma}\right)\right\}, \gamma: j \in \gamma$ without evaluating their relevance individually. One then looks for an approximation which is sparse at the level of groups of functions. Note that grouping predictors significantly alleviates the computational cost associated with the subset selection as further discussed in Section 4.7.

It is important to recall that this approximation format is made only for the subset selection step and is independent of the format the QoI will finally be approximated in. The selection of groups reduces to selection of interaction modes $f_{\gamma}$ and leaves the possibility for using different formats between the subset selection step and the coefficients evaluation step: an interaction mode found to be dominant is incorporated to the active dictionary $\mathcal{J}_{f, \text { post }}$ independently of the way its contribution to the approximation of $u$ is actually determined in the end. Indeed, since the LARS technique only applies to predictors linear in their coefficients, an approximation $\widetilde{f}_{\gamma}$ of the form (17) is suitable for the selection of the dominant groups. However, the final approximation $\widehat{f}_{\gamma}$ of the retained $f_{\gamma}$ may rely on predictors nonlinear in their coefficients: the subset selection step only serves to determine which interaction modes will be considered in the a posteriori approximation basis, the "skeleton" $\left\{f_{\gamma}: \gamma \in \mathcal{J}_{f, \mathrm{post}}\right\}$.

The selection is made using a modified LARS approach and the following optimization problem is solved:

$$
\begin{equation*}
\boldsymbol{c}=\underset{\widetilde{\boldsymbol{c}} \in \mathbb{R}^{|\mathcal{J}|}}{\arg \min }\|\boldsymbol{u}-\Psi \widetilde{\boldsymbol{c}}\|_{2}^{2}+\tau \sum_{\boldsymbol{\gamma} \in \mathcal{J}_{f}}\left\|\widetilde{\boldsymbol{c}}_{\boldsymbol{\gamma}}\right\|_{K_{\boldsymbol{\gamma}}} \tag{18}
\end{equation*}
$$

with $\tau>0$ the regularization parameter and $\|\cdot\|_{K_{\gamma}}$ a norm induced by a positive definite matrix $K_{\gamma}$. All predictors within a group $\gamma$ are here weighted similarly so that we use a scaled identity matrix $K_{\gamma}=\mathrm{I}_{\left|\mathcal{J}_{\gamma}\right|} / \mathcal{J}_{\boldsymbol{\gamma}} \mid, \forall \gamma \in \mathcal{J}_{f}$. The regularization term is a combination of $L^{2}$ - and $L^{1}$-norms and penalizes the $L^{1}$-norm of the "group" vector to promote a collective behavior: either a group is basically active (nonzero $K_{\gamma}$-norm) or inactive, essentially disregarding the detailed behavior within the group. This group LARS (gLARS) strategy was first proposed in [41] and the algorithm presented in [42] was modified to solve the optimization problem (18).

[^3]The set of dominant modes $\left\{f_{\gamma}\right\}$ is first determined by the gLARS approach with a low approximation order $\widetilde{p}$ and the basis is subsequently further refined by a LARS step, using $L^{1}$-regularization, onto these selected modes only now approximated with a higher $\widetilde{p}$ for improved accuracy.

### 4.4 Functional Spaces for the Final Approximation Basis

We now discuss the general methodology for approximating a random variable $u(\boldsymbol{\xi})$, from a finite set of its realizations. An a priori choice of representation format was first made, Section 4.2, and was adjusted based on the data through the subset selection procedure, the a posteriori step, Section 4.3. This has selected a set of groups, or interaction modes, $\left\{f_{\gamma}\right\}_{\gamma \in \mathcal{J}_{f, \text { post }}}$ deemed to most contribute to the HDMR representation of the QoI $u$. The actual approximation of $u$ will rely on these selected groups but does not bear restriction on the linearity with respect to the coefficients so that different suitable formats, possibly nonlinear, can then be considered.

Many possibilities exist to determine an approximation of $\left\{f_{\boldsymbol{\gamma}}, \gamma \in \mathcal{J}_{f, \text { post }}\right\}$ in a polynomial space, e.g., maximum partial degree, maximum total degree, hyperbolic cross, etc. For sake of simplicity, the space $\mathbb{P}_{p}$ of polynomials with maximum total degree $p$ is retained as a reasonable compromise between cardinality $\left|\mathcal{J}_{\gamma}\right|$ and expected accuracy of the approximation $\widehat{f}_{\gamma}$ :

$$
\begin{align*}
& f_{\boldsymbol{\gamma}}\left(\left\{\xi_{i}\right\}_{i \in \boldsymbol{\gamma}}\right) \approx \widehat{f}_{\boldsymbol{\gamma}}\left(\left\{\xi_{i}\right\}_{i \in \boldsymbol{\gamma}}\right)=\sum_{\boldsymbol{\alpha},|\boldsymbol{\alpha}| \leq p} c_{\boldsymbol{\gamma}, \boldsymbol{\alpha}} \psi_{\boldsymbol{\alpha}}\left(\left\{\xi_{i}\right\}_{i \in \boldsymbol{\gamma}}\right), \quad \psi_{\boldsymbol{\alpha}}\left(\left\{\xi_{i}\right\}_{i \in \boldsymbol{\gamma}}\right)=\prod_{i \in \boldsymbol{\gamma}} \psi_{\alpha_{i}}\left(\xi_{i}\right), \\
& \boldsymbol{\alpha}=\left(\alpha_{i}, i \in \boldsymbol{\gamma}\right), \quad \alpha_{i} \in\{1, \ldots, p\}, \quad 1 \leq|\boldsymbol{\gamma}| \leq N_{l}^{(\mathrm{PC})} \leq \min \left(N_{l}, p\right) \tag{19}
\end{align*}
$$

The cardinality associated with this approximation of $f_{\gamma}$ at a given iteration level $l=|\gamma|$ is $\left|\mathcal{J}_{\gamma}\right|=p!/(l!(p-l)!)$ and usually provides an accurate approximation with a low number of coefficients for low dimensions $|\gamma|$.

When the dimension $|\boldsymbol{\gamma}|$ increases, the number of terms in $\widehat{f}_{\boldsymbol{\gamma}}$ decreases and eventually degenerates for $|\boldsymbol{\gamma}|>p$. For modes of interaction order highe-than a prescribed threshold $N_{l}^{(\mathrm{PC})}$, a low-rank canonical decomposition is instead considered:

$$
\begin{equation*}
f_{\gamma}\left(\left\{\xi_{i}\right\}_{i \in \gamma}\right) \approx \widehat{f}_{\boldsymbol{\gamma}}\left(\left\{\xi_{i}\right\}_{i \in \gamma}\right)=\sum_{r=1}^{n_{r}} \prod_{i \in \boldsymbol{\gamma}} \sum_{\alpha=1}^{p} c_{\boldsymbol{\gamma}, \alpha}^{r, i} \psi_{\alpha}\left(\xi_{i}\right), \quad N_{l}^{(\mathrm{PC})}<|\gamma| \leq N_{l} \leq d \tag{20}
\end{equation*}
$$

The maximum number of modes at a given interaction level $l$ is $d!/((d-l)!l!)$. Relying on an approximation in $\mathbb{P}_{p}$ for interaction modes of order $|\boldsymbol{\gamma}| \leq N_{l}^{(\mathrm{PC})}$ and on low-rank approximation for higher interaction order modes, with maximum rank $n_{r}$, the total cardinality of this approximation format is bounded from above by

$$
\begin{equation*}
\left|\mathcal{J}_{\mathrm{eff}}\right| \leq \sum_{l=0}^{N_{l}^{(\mathrm{PC})}} \frac{d!p!}{(d-l)!(l!)^{2}(p-l)!}+\sum_{l=N_{l}^{(\mathrm{PC})}+1}^{N_{l}} \frac{d!}{(d-l)!l!} n_{r} l p \tag{21}
\end{equation*}
$$

### 4.5 Algorithm for Approximating a Random Variable

We will denote $\mathcal{J}_{f \text {, eff }}$ the set of modes $\left\{\widehat{f}_{\gamma}\right\}_{\gamma \subseteq \mathcal{J}_{f, \text { post }}} \quad$ finally considered for the approximation of $u$ and $\mathcal{J}_{\text {eff }}$ the set of associated predictors $\left\{\psi_{\alpha}\right\}$. The interaction modes are estimated sequentially. Once a new mode is evaluated, the whole approximation may be updated by reevaluating the coefficients of the predictors $\left\{\psi_{\alpha}\right\}$ already evaluated of the current evaluation set $\mathcal{J}_{f \text {, eff }} \in \mathcal{J}_{f \text {, post }}$. Let $\boldsymbol{z}=\left(z^{(1)} \ldots z^{\left(N_{q}\right)}\right)^{T}$ be the residual vector after basis functions $\widehat{f}_{\boldsymbol{\gamma}}, \boldsymbol{\gamma} \in \mathcal{J}_{f \text {, eff }}$ have been evaluated. The coefficients involved in the next mode $\widehat{f}_{\boldsymbol{\gamma}}, \boldsymbol{\gamma} \in \mathcal{J}_{f \text {, post }} \backslash \mathcal{J}_{f, \text { eff }}$ to be evaluated are then determined. If $\boldsymbol{\gamma}$ is such that $|\boldsymbol{\gamma}| \leq N_{l}^{(\mathrm{PC})}$, they are computed from the following system of equations ${ }^{4}$ :

[^4]\[

$$
\begin{equation*}
\left\{\boldsymbol{c}_{\boldsymbol{\gamma}, \cdot}=\underset{\widetilde{\boldsymbol{c}} \in \mathbb{R}^{\left|\mathcal{J}_{\boldsymbol{\gamma}}\right|}}{\arg \min }\|\boldsymbol{z}-\Psi \widetilde{\boldsymbol{c}}\|_{2}, \quad \forall i \in \boldsymbol{\gamma}, \quad \boldsymbol{\gamma} \subseteq \mathcal{J}_{f, \text { post }}, \quad|\boldsymbol{\gamma}| \leq N_{l}^{(\mathrm{PC})}\right. \tag{22}
\end{equation*}
$$

\]

with $\boldsymbol{c}_{\boldsymbol{\gamma}, \cdot}=\left(\boldsymbol{c}_{\boldsymbol{\gamma}, \alpha_{i}}, i \in \boldsymbol{\gamma}\right)^{T}$ and

$$
\begin{align*}
z^{(q)} & =u^{(q)}-\sum_{\boldsymbol{\gamma}^{\prime} \subseteq \mathcal{J}_{f, \text { eff }} \backslash \boldsymbol{\gamma}} \widehat{f}_{\boldsymbol{\gamma}^{\prime}}\left(\left\{\xi_{i}^{(q)}\right\}_{i \in \boldsymbol{\gamma}^{\prime}}\right), \quad \boldsymbol{z}=\left(z^{(1)} \ldots z^{\left(N_{q}\right)}\right)^{T} \\
\Psi_{q \boldsymbol{\alpha}} & =\psi_{\boldsymbol{\alpha}}\left(\left\{\xi_{i}^{(q)}\right\}_{i \in \boldsymbol{\gamma}}\right), \quad \Psi=\left[\Psi_{q \boldsymbol{\alpha}}\right] . \tag{23}
\end{align*}
$$

To solve for the coefficients associated with predictors nonlinear in their coefficients, an alternate least squares (ALS) approach is used, reformulating the nonlinear problem into a set of coupled linear equations:

$$
\begin{equation*}
\left\{\boldsymbol{c}_{\boldsymbol{\gamma}, \cdot}^{r, i}=\underset{\widetilde{\boldsymbol{c}} \in \mathbb{R}^{p}}{\arg \min }\left\|\boldsymbol{z}_{i}-\Psi \widetilde{\boldsymbol{c}}\right\|_{2}, \quad \forall i \in \boldsymbol{\gamma} \subseteq \mathcal{J}_{f, \mathrm{post}}, \quad N_{l}^{(\mathrm{PC})}<|\boldsymbol{\gamma}| \leq N_{l}\right. \tag{24}
\end{equation*}
$$

with $\boldsymbol{c}_{\boldsymbol{\gamma}, \cdot}^{r, i}=\left(\boldsymbol{c}_{\boldsymbol{\gamma}, 1}^{r, i} \ldots \boldsymbol{c}_{\boldsymbol{\gamma}, p}^{r, i}\right)^{T}$ and

$$
\begin{align*}
& z_{i}^{(q)}=u^{(q)}-\sum_{\boldsymbol{\gamma}^{\prime} \subseteq \mathcal{J}_{f, \text { eff }}} \widehat{f}_{\boldsymbol{\gamma}^{\prime}}\left(\left\{\xi_{i}^{(q)}\right\}_{i \in \boldsymbol{\gamma}^{\prime}}\right)-\sum_{r^{\prime}=1}^{r-1} \prod_{i^{\prime} \in \boldsymbol{\gamma}} \sum_{\alpha^{\prime}=1}^{p} c_{\boldsymbol{\gamma}, \alpha^{\prime}}^{r^{\prime}, i^{\prime}} \psi_{\alpha^{\prime}}\left(\xi_{i^{\prime}}^{(q)}\right), \\
& \Psi_{q \alpha}=\psi_{\alpha}\left(\xi_{i}^{(q)}\right) \prod_{i^{\prime} \in \boldsymbol{\gamma}, i^{\prime} \neq i} \sum_{\alpha^{\prime}=1}^{p} c_{\boldsymbol{\gamma}, \alpha^{\prime}}^{r, i^{\prime}} \psi_{\alpha^{\prime}}\left(\xi_{i^{\prime}}^{(q)}\right), \quad \Psi=\left[\Psi_{q \boldsymbol{\alpha}}\right] . \tag{25}
\end{align*}
$$

This whole step is embedded in a loop over the modes $f_{\boldsymbol{\gamma}}, \boldsymbol{\gamma} \in \mathcal{J}_{f, \text { post }}$ retained by the subset selection procedure. The cross-validation error $(\mathrm{CV} \varepsilon)$ is estimated from $\widehat{N_{q}}$ validation samples $\left\{\xi^{(\widehat{q})}, u^{(\widehat{q})}\right\}_{\hat{q}=1}^{\widehat{N_{q}}}$ independent from the $N_{q}$ samples of the training set. ${ }^{5}$ If the cross-validation error has increased over the last two loops, the approximation basis is likely to have become too large with respect to the available data and iterations are stopped. The retained basis is then the one that has led to the lowest CV . If $\mathrm{CV} \varepsilon$ keeps decreasing, the next interaction mode as selected by the subset selection step is considered and added to the current active set $\mathcal{J}_{f, \text { eff }}$ and the whole iteration is carried out. Once the approximation is determined, the coefficients are updated with the same sequential technique using both the training and the validation points, $N_{q}+\widehat{N_{q}}$. The approximation accuracy is estimated by the relative $L^{2}$-norm $\varepsilon$ of the approximation error estimation evaluated from a $\widetilde{N_{q}}$-point test set $\left\{\left(x^{(\widetilde{q})}, \xi^{(\widetilde{q})}, u^{(\widetilde{q})}\right)\right\}_{\tilde{q}=1}^{\widetilde{N_{q}}}$, independent from the training set:

$$
\begin{equation*}
\varepsilon^{2}:=\|\boldsymbol{u}-\widehat{\boldsymbol{u}}\|_{2}^{2} /\|\boldsymbol{u}\|_{2}^{2}, \quad \boldsymbol{u}=\left(u^{(1)} \ldots u^{\left(\widetilde{N_{q}}\right)}\right), \quad \widehat{\boldsymbol{u}}=\left(\widehat{u}^{(1)} \ldots \widehat{u}^{\left(\widetilde{N_{q}}\right)}\right) \tag{26}
\end{equation*}
$$

The global methodology is summarized in Algorithm 1. Statistical moments can be readily evaluated from the present HDMR of the QoI, see Appendix B.

[^5]```
Algorithm 1: Sketch of the solution method for approximating a random variable \(u(\xi)\)
    Select an a priori basis in HDMR format. Choose \(p, N_{l}, N_{l}^{(\mathrm{PC})}, n_{r}\) and \(\widetilde{p}\). Initialize \(\boldsymbol{z}=\left(u^{(1)} \ldots u^{\left(N_{q}\right)}\right)^{T}\).
    Subset selection step. Solve the LASSO optimization problem with the gLARS algorithm \(\longrightarrow\) sequence of \(a\)
    posteriori approximation bases indexed by \(s\) with ordered groups \(\mathcal{J}_{f, \text { post }}=\left\{\boldsymbol{\gamma}^{(s)}\right\}\). Initialize \(s\) and \(\mathcal{J}_{f, \text { eff }}\) :
    \(s \leftarrow 0, \mathcal{J}_{f}\), eff \(\leftarrow \emptyset\).
    Solve the approximation problem:
    repeat
        \(s \leftarrow s+1\).
        Consider the next mode \(\widehat{f}_{\boldsymbol{\gamma}^{(s)}}\) from the set \(\mathcal{J}_{f, \text { post }}\) selected in (2): \(\mathcal{J}_{f, \text { eff }} \leftarrow \mathcal{J}_{f, \text { eff }} \cup \boldsymbol{\gamma}^{(s)}\).
        Solve for the approximation coefficients \(\left\{c_{\gamma}\right\}_{\boldsymbol{\gamma} \in \mathcal{J}_{\text {f eff }}}\) by alternately solving for the coefficients of modes
        \(\left\{\widehat{f}_{\gamma}\right\}_{\gamma \in \mathcal{J}_{f, \text { eff }}}\), Eqs. \((22,24)\). [Update step]
        Estimate the cross-validation error \(\mathrm{CV} \varepsilon\) and evaluate the current approximation \(\widehat{\boldsymbol{u}}=\left(\widehat{u}^{(1)} \ldots \widehat{u}^{\left(N_{q}\right)}\right)^{T}\).
        Update the residual \(\boldsymbol{z} \leftarrow \boldsymbol{u}-\widehat{\boldsymbol{u}}\).
    until CVE has increased over the last two passes s and \(s-1\).
\(11 \mathcal{J}_{f, \text { eff }} \leftarrow \mathcal{J}_{f, \text { eff }} \backslash\left\{\boldsymbol{\gamma}^{(s)}, \boldsymbol{\gamma}^{(s-1)}\right\}\).
12 Update the coefficients \(\left\{\boldsymbol{c}_{\boldsymbol{\gamma}}\right\}_{\boldsymbol{\gamma} \in \mathcal{J}_{\mathcal{J}, \text { eff }}}\) of the retained modes with the extended set of data \(\left\{\xi^{(q)}, u^{(q)}\right\}_{q=1}^{N_{q}+\widehat{N_{q}}}\). It
    finally yields \(\widehat{u}(\boldsymbol{\xi})\) expressed in the basis \(\left\{\widehat{f}_{\gamma}\right\}_{\gamma \in \mathcal{J}_{f, \text { eff }}}\).
```


### 4.6 Robust Estimation

An important concern when deriving a methodology is the robustness with respect to noise and a more robust alternative to the methodology discussed so far is now presented.

To evaluate the approximation coefficients once an approximation basis is determined from the subset selection step, a standard approach is to minimize a norm between target observations and reconstructed approximation as done in the previous section, Eqs. (22) and (24): the approximation coefficients of a given mode $\widehat{f}_{\gamma}$ are basically given by $\boldsymbol{c}_{\boldsymbol{\gamma}}=\arg \min _{\tilde{\boldsymbol{c}} \in \mathbb{R}^{\left|\mathcal{J}_{\mathcal{Y}}\right|}}\|\boldsymbol{z}-\Psi \widetilde{\boldsymbol{c}}\|_{2}$, with $\Psi \in \mathbb{R}^{N_{q} \times\left|\mathcal{J}_{\mathcal{\gamma}}\right|}$ the matrix of the $\boldsymbol{\gamma}$-group predictors evaluated in $\left\{\tilde{\xi}^{(q)}\right\}$ and $z$ the target residual vector. The solution to this least squares problem is equivalently obtained from

$$
\begin{equation*}
\{\boldsymbol{c}, \Delta \boldsymbol{z}\}=\underset{\widetilde{\boldsymbol{c}} \in \mathbb{R}^{\left|\mathcal{J}_{\boldsymbol{\gamma}}\right|}}{\arg \min }\|\widetilde{\Delta \boldsymbol{z}}\|_{F} \quad \text { s.t. } \quad \boldsymbol{z}+\widetilde{\Delta \boldsymbol{z}}=\Psi \widetilde{\boldsymbol{c}}, \tag{27}
\end{equation*}
$$

which minimizes the Frobenius norm of the residual vector. This implicitly assumes no error in the coordinates $\left\{\xi^{(q)}\right\}$ at which the target is evaluated. For instance, these coordinates may be known as the solution of auxiliary inference problems. This brings errors so that the actual coordinates vector is only estimated with an error $\Delta \dot{\xi}^{(q)}$. Since $\Psi$ depends on $\xi$, an error predictor matrix $\Delta \Psi(\xi, \Delta \xi):=\Psi(\xi+\Delta \xi)-\Psi(\xi)$ arises and the estimation problem (27) then rewrites as a total least squares problem [43]:

$$
\begin{equation*}
\{\boldsymbol{c}, \Delta \Psi, \Delta \boldsymbol{z}\}=\underset{\widetilde{c} \in \mathbb{R}^{\left|\mathcal{J}_{\boldsymbol{\gamma}}\right|}}{\arg \min }\|\widetilde{\Delta \Psi} \widetilde{\Delta \boldsymbol{z}}\|_{F} \quad \text { s.t. } \quad \boldsymbol{z}+\widetilde{\Delta \boldsymbol{z}}=(\Psi+\widetilde{\Delta \Psi}) \widetilde{\boldsymbol{c}} . \tag{28}
\end{equation*}
$$

The realizations of the error in the data $\left\{\Delta \xi^{(q)}(\theta), \Delta z^{(q)}(\theta)\right\}$ are modeled to follow the distribution of zeromean iid variables. Further, predictors may be correlated:

$$
\begin{equation*}
\mathscr{E}_{\theta}\left[\left(\Delta \psi_{\alpha}-\mathscr{E}\left[\Delta \psi_{\alpha}\right]\right)\left(\Delta \psi_{\alpha^{\prime}}-\mathscr{E}_{\theta}\left[\Delta \psi_{\alpha^{\prime}}\right]\right)\right] \neq 0 \tag{29}
\end{equation*}
$$

with $\Delta \psi_{\boldsymbol{\alpha}}=\Delta \psi_{\boldsymbol{\alpha}}\left(\xi^{(q)}, \Delta \xi^{(q)}\right)$. A general approach to solve the weighted total least squares (wTLS) problem of Eq. (28) consists of the minimization of the usual weighted residual sum of squares $\rho^{2}$ [44]:

$$
\begin{equation*}
\rho^{2}:=\operatorname{vec}(\Delta X)^{T} \Lambda^{-1} \operatorname{vec}(\Delta X), \quad \Delta X:=(\Delta \Psi \Delta \boldsymbol{z})^{T}, \quad X:=(\Psi \boldsymbol{z})^{T} \tag{30}
\end{equation*}
$$

where the 'vec' operator unfolds a generic $m \times n$ matrix into a $m n$ vector and $X$ is the data matrix. The covariance matrix for $X, \Lambda:=\left\langle\operatorname{vec}\left(X-\langle X\rangle_{N_{q}}\right) \operatorname{vec}\left(X-\langle X\rangle_{N_{q}}\right)^{T}\right\rangle_{N_{q}}$, is evaluated and the minimization problem (28) is solved using the ALS-based algorithm proposed in [45].

As will be shown in the numerical experiments examples, Section 5.1.4, the present total LS formulation allows one to improve the approximation quality from noisy data.

Remark 2. When a large amount $N_{q}$ of experimental information is available, the data matrix $X \in \mathbb{R}^{\left(\left|\mathcal{J}_{\gamma}\right|+1\right) \times N_{q}}$ can be large. The resulting correlation matrix $\Lambda$ then has potentially very large dimensions. However, since the noise is assumed independent from one sample to another, $\Lambda$ has a block diagonal structure. Further, it is a symmetric definite positive matrix, allowing for additional reduction of the storage requirement. The structure of $\Lambda$ is then exploited in solving the weighted total least squares problem above through sparse storage and operations.

### 4.7 Asymptotic Numerical Complexity

While the primary motivation for this work is to determine an accurate representation of a random quantity from a small set of its realizations, it is desirable that the solution method remains computationally tractable. As seen above, the algorithm for approximating a random variable is essentially two fold.

The selection process essentially consists in sequentially building a subset, Section 4.3. Each step of the sequence involves solving a LS problem of growing size and finding the basis function, or group of functions, within the a priori set $\mathcal{J}_{\text {prior }}$ most correlated with the current residual. The matrix of the LS problem is $\Psi \in \mathbb{R}^{N_{q} \times\left|\mathcal{J}_{\text {post }}\right|}$, with $\left|\mathcal{J}_{\text {post }}\right|$ the cardinality of the current set of selected basis functions. The LS problem is solved via a QR decomposition of $\Psi$ in $\mathcal{O}\left(N_{q}\left|\mathcal{J}_{\text {post }}\right|^{2}\right)$ operations. The iterative selection process is carried out with a growing active set $\mathcal{J}_{\text {post }}$ until the problem becomes ill-posed, i.e., until $\left|\mathcal{J}_{\text {post }}\right|$ is about $N_{q}$. We use grouped LARS and denote $\overline{\left|\mathcal{J}_{\gamma, \text { post }}\right|}$ the average cardinality of the retained group predictors, i.e., the average number of basis functions in the group added to the active set. The subset selection process retains $n_{f}$ groups of variables so that the total cost associated with the LS step of the subset selection is

$$
\begin{equation*}
\mathscr{J}_{\mathrm{LS}}=\mathcal{O}\left(\sum_{s=1}^{n_{f}} N_{q}\left(\overline{\left|\mathcal{J}_{\gamma, \mathrm{post}}\right|} s\right)^{2}\right) \tag{31}
\end{equation*}
$$

As groups of predictors are moved to the active set, the size of the remaining a priori set decreases, $\left|\mathcal{J}_{\text {prior }}\right|^{\text {(current) }}$ $\simeq\left|\mathcal{J}_{\text {prior }}\right|-s \overline{\mathcal{J}_{\gamma, \text { post }} \mid}$. The cost associated with the evaluation of the correlation for each predictor in the inactive set is then

$$
\begin{equation*}
\mathscr{J}_{\text {correl }}=\mathcal{O}\left(\sum_{s=1}^{n_{f}} N_{q}\left(\left|\mathcal{J}_{\text {prior }}\right|-s \overline{\left|\mathcal{J}_{\gamma, \text { post }}\right|}\right)\right) \propto N_{q} . \tag{32}
\end{equation*}
$$

In practice, the cost associated with the evaluation of the correlation of the predictors in the inactive set with the current residual dominates so that the whole cost of the subset selection finally approximates as

$$
\begin{equation*}
\mathscr{J}_{\text {subsel }}=\mathscr{J}_{\mathrm{LS}}+\mathscr{J}_{\text {correl }} \simeq \mathcal{O}\left(N_{q}\left|\mathcal{J}_{\text {prior }}\right| n_{f}-N_{q}{\left.\overline{\mathcal{J}_{\gamma, \text { post }} \mid} \frac{n_{f}\left(n_{f}+1\right)}{2}\right) . . . .}_{2}\right. \tag{33}
\end{equation*}
$$

The second step of the solution method deals with the evaluation of the approximation coefficients, Sections 4.44.5. The cost associated with evaluating the coefficients of an $l$ th interaction order mode, $1 \leq l \leq N_{l}^{(\mathrm{PC})}$, encompasses the matrix $\Psi$ assembly cost $\mathcal{O}\left(N_{q} p!/(l!(p-l)!)\right)$ and the LS solution $\mathcal{O}\left(N_{q}(p!/(l!(p-l)!))^{2}\right)$. Since modes
$\left\{\widehat{f}_{\gamma}\right\}_{\gamma \in \mathcal{J}_{f, \text { eff }}}$ already evaluated may be updated once an additional one from the selected set is considered, the total cost is the sum of an arithmetic sequence. Its exact formulation depends on the selected set and is difficult to derive in closed-form. As a simple example, updating all coefficients for each new mode $\widehat{f}_{\gamma}$ considered, neglecting the cost associated with first-order interaction modes and assuming only second-order interaction modes are retained in the $a$ posteriori set, an upper bound for the cost writes

$$
\begin{equation*}
\mathscr{J}_{\text {coef }} \leq \sum_{s=1}^{\left|\mathcal{J}_{f, \text { eff }}\right|}\left[\mathcal{O}\left(s N_{q} p^{l}\right)+\mathcal{O}\left(s N_{q} p^{2 l}\right)\right], \quad \text { with } l=2 \tag{34}
\end{equation*}
$$

where $\left|\mathcal{J}_{f, \text { eff }}\right|$ is the number of groups finally retained for the approximation by the CV test, see Algorithm 1 . A quantitative discussion of the numerical cost is given in Section 5.1 .5 with an illustrative example.

### 4.8 Approximation of a Random Process by a Separated Representation

The approximation of a random process, say, a space-dependent uncertain quantity $u(\boldsymbol{x}, \boldsymbol{\xi})$ is now considered in the form of separation of variables:

$$
\begin{equation*}
u(\boldsymbol{x}, \boldsymbol{\xi}) \approx w_{0}(\boldsymbol{x})+\sum_{n=1}^{N} w_{n}(\boldsymbol{x}) \lambda_{n}(\boldsymbol{\xi}) \equiv \sum_{n=0}^{N} w_{n}(\boldsymbol{x}) \lambda_{n}(\boldsymbol{\xi}), \quad \lambda_{0} \equiv 1 \tag{35}
\end{equation*}
$$

The 'spatial' modes are associated with all physical dimensions the random process may be indexed upon (space, time,$\ldots)$ so that $\boldsymbol{x}=\left(x_{1} x_{2} \ldots t \ldots\right) \subseteq \mathbb{R}^{d_{\boldsymbol{x}}}$. They are defined as: $w_{n}(\boldsymbol{x})=\sum_{l=1}^{\left|\mathcal{J}_{x}\right|} c_{l, n}^{(w)} \phi_{l}(\boldsymbol{x})$ with $\left\{\phi_{l}\right\}$ a chosen truncated basis of cardinality $\left|\mathcal{J}_{x}\right|$. The functional form of "stochastic" modes $\left\{\lambda_{n}\right\}$ and their evaluation was discussed in Sections 4.2-4.5.

The spatial and stochastic modes of the approximation (35) are sequentially determined in turn. Let $\|v\|_{N_{q}}:=$ $\langle v, v\rangle_{N_{q}}$ be the norm induced by the data-driven inner product: $\langle\cdot, \cdot\rangle_{N_{q}}: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R},(v, w) \mapsto\langle v, w\rangle_{N_{q}}:=$ $\sum_{q=1}^{N_{q}} v^{(q)} w^{(q)}$. Assuming $\left\{\lambda_{n}\right\}$ known and projecting Eq. (35) onto the space spanned by $\left\{\phi_{l}\right\}$, the coefficients $\left\{c_{l, n}^{(w)}\right\}_{l}$ of the deterministic mode $w_{n}$ are the solution of the following problem:

$$
\begin{align*}
& \quad\left\langle u, \phi_{k} \lambda_{n}\right\rangle_{N_{q}}=\left\langle\sum_{n^{\prime}=0}^{n-1} w_{n^{\prime}} \lambda_{n^{\prime}}+\sum_{l=1}^{\left|\mathcal{J}_{x}\right|} c_{l, n}^{(w)} \phi_{l} \lambda_{n}, \phi_{k} \lambda_{n}\right\rangle_{N_{q}}, \quad \forall 1 \leq k \leq\left|\mathcal{J}_{x}\right|, \\
& \Leftrightarrow \quad \boldsymbol{c}_{,, n}^{(w)}=\underset{\widetilde{\boldsymbol{c}} \in \mathbb{R}^{\prime} \mathcal{J}_{x} \mid}{\arg \min }\left\|\boldsymbol{u}-\sum_{n^{\prime}=0}^{n-1} \boldsymbol{w}_{n^{\prime}} \odot \boldsymbol{\lambda}_{n^{\prime}}-(\Phi \widetilde{\boldsymbol{c}}) \odot \boldsymbol{\lambda}_{n}\right\|_{2} \tag{36}
\end{align*}
$$

where $\Phi \in \mathbb{R}^{N_{q} \times\left|\mathcal{J}_{x}\right|}, \Phi_{q l}=\phi_{l}\left(\boldsymbol{x}^{(q)}\right)$, $\boldsymbol{w}_{n}=\left(w_{n}\left(x^{(1)}\right) \ldots w_{n}\left(x^{\left(N_{q}\right)}\right)\right)^{T}, \boldsymbol{\lambda}_{n}=\left(\lambda_{n}\left(\boldsymbol{\xi}^{(1)}\right) \ldots \lambda_{n}\left(\boldsymbol{\xi}^{\left(N_{q}\right)}\right)\right)^{T}$ and $\odot$ is the Hadamard product. Similarly, the stochastic mode $\lambda_{n}$ is evaluated by determining the set of coefficients $\left\{c_{n}^{(\lambda)}\right\}$ minimizing $\left\|\boldsymbol{u}-\sum_{n^{\prime}=0}^{n-1} \boldsymbol{w}_{n^{\prime}} \odot \boldsymbol{\lambda}_{n^{\prime}}-\boldsymbol{w}_{n} \odot \boldsymbol{\lambda}_{n}\left(\left\{c_{n}^{(\lambda)}\right\}\right)\right\|_{2}$ using Algorithm 1 presented in Section 4.5. The spatial mode $w_{n}$ is then evaluated from Eq. (36) given all the other information and the whole iteration is repeated until convergence of the pair $\left\{w_{n}, \lambda_{n}\right\}$. The next pair can then be determined with the same methodology with $n \leftarrow$ $n+1$. The algorithm is summarized in Algorithm 2.

Remark 3. If the separated approximation grows beyond a few modes, it is beneficial to update the coefficients of, say, the spatial modes for improved accuracy: solve for $\left\{w_{0}, \ldots, w_{n}\right\}$ given $\left\{\boldsymbol{u}, \lambda_{0}, \ldots, \lambda_{n}\right\}$.

```
Algorithm 2: Sketch of the solution method for approximating a random process
    Choose \(\left|\mathcal{J}_{x}\right|\). Initialize \(\boldsymbol{z}=\boldsymbol{u}\) and \(\boldsymbol{\lambda}_{0}=\mathbf{1}\) and set \(n \leftarrow 0\).
    repeat
        while \(\left\|\lambda_{n}\right\|_{N_{q}}\) not converged do
```

            Solve for coefficients \(\left\{c_{l, n}^{(w)}\right\}_{l=1}^{\left|\mathcal{J}_{x}\right|}\) of the deterministic mode given \(\boldsymbol{\lambda}_{n}\) and \(\boldsymbol{z}\) using Eq. (36) and
            normalize \(w_{n}(\boldsymbol{x})=\sum_{l} c_{l, n}^{(w)} \phi_{l}(\boldsymbol{x}),\left\|w_{n}\right\|_{N_{q}}=1\).
            Solve for the stochastic mode \(\lambda_{n}(\boldsymbol{\xi})\) using Algorithm 1 given \(\boldsymbol{w}_{n}\) and \(\boldsymbol{z}\). If \(n=0, \boldsymbol{\lambda}_{0} \leftarrow \mathbf{1}\).
        Set \(\boldsymbol{z} \leftarrow \boldsymbol{z}-\boldsymbol{w}_{n} \odot \boldsymbol{\lambda}_{n}\), and \(n \leftarrow n+1\).
    until a termination criterion is met (for instance, \(\left\|\lambda_{n}\right\|_{N_{q}}\) below a threshold or maximum rank \(n>N\) reached).
    
## 5. NUMERICAL EXPERIMENTS

The methodology developed in the previous sections is now demonstrated on a set of examples. Different aspects of the global solution method are illustrated on a 1D stochastic diffusion equation. A more computationally involved example is next considered with a shallow water problem with multiple sources of uncertainty.

### 5.1 Stochastic Diffusion Equation

We consider a steady-state stochastic diffusion equation on $\Omega \times \Xi, \Omega=\left[x_{-}, x_{+}\right] \subset \mathbb{R}$, with deterministic Dirichlet boundary conditions:

$$
\begin{equation*}
\nabla_{x}\left(\nu\left(x, \xi^{\prime}\right) \nabla_{x} u(x, \xi)\right)=F\left(x, \xi^{\prime \prime}\right), \quad u\left(x_{-}, \xi\right)=u_{-}, \quad u\left(x_{+}, \xi\right)=u_{+} \tag{37}
\end{equation*}
$$

The right-hand side $F$ is a random source field and $v$ is a space-dependent random diffusion coefficient modeled as

$$
\begin{align*}
v\left(x, \xi^{\prime}\right) & =v_{0}(x)+\sum_{k=1}^{d_{v}} \sqrt{\sigma_{v, k}} \omega_{v, k}(x) \xi_{k}^{\prime}, \\
F\left(x, \xi^{\prime \prime}\right)=F_{0}(x)+\sum_{k=1}^{d_{F}} \sqrt{\sigma_{F, k}} \omega_{F, k}(x) \xi_{k}^{\prime \prime}, & \xi^{\prime \prime}=\left(\xi_{1}^{\prime} \ldots \xi_{d_{v}}^{\prime}\right), \tag{38}
\end{align*}
$$

with $\nu_{0}(x)=1$ and $F_{0}(x)=-1$ the respective mean values. The random variables $\left\{\xi_{1}^{\prime}, \ldots, \xi_{d_{v}}^{\prime}, \xi_{1}^{\prime \prime}, \ldots, \xi_{d_{F}}^{\prime \prime}\right\}$ are chosen mutually independent and uniformly distributed on $[0,1]$. The spatial modes $\omega_{v, k}(x)$ and $\omega_{F, k}(x)$, and their associated amplitude $\sqrt{\sigma_{V, k}}$ and $\sqrt{\sigma_{F, k}}$, are the first dominant eigenfunctions of the following eigenproblems:

$$
\begin{array}{ll}
\int_{\Omega} K_{v}\left(x, x^{\prime}\right) \omega_{v, k}\left(x^{\prime}\right) \mathrm{d} x^{\prime}=\sigma_{v, k} \omega_{v, k}(x), & K_{v}\left(x, x^{\prime}\right)=\sigma_{v}^{2} e^{-\left[\left(x-x^{\prime}\right)^{2} / 2 L_{c, v}^{2}\right]} \\
\int_{\Omega} K_{F}\left(x, x^{\prime}\right) \omega_{F, k}\left(x^{\prime}\right) \mathrm{d} x^{\prime}=\sigma_{F, k} \omega_{F, k}(x), & K_{F}\left(x, x^{\prime}\right)=\sigma_{F}^{2} e^{-\left[\left(x-x^{\prime}\right)^{2} / 2 L_{c, F}^{2}\right]} \tag{39}
\end{array}
$$

with $K_{v}$ and $K_{F}$ the correlation kernels. The random fields properties are chosen as $\sigma_{v}=0.7, \sigma_{F}=0.7, L_{c, v}=0.3$, $L_{F, v}=0.3$. Note that $F\left(\cdot, \xi^{\prime \prime}\right)<0$ a.e. and $v\left(\cdot, \xi^{\prime}\right)>0$ a.e. so that the problem remains coercive. The spectra of the operators associated with these eigenproblems are here the same and decay quickly thanks to the high correlation length as can be appreciated from Table 1 where the dominant eigenvalues are given. The resulting problem is then anisotropic in $\Xi$ in the sense that the degree of dependence of the input random parameters along the dimensions $\left\{\xi_{1}, \ldots, \xi_{8}\right\}$ strongly varies.

TABLE 1: Upper part of the spectrum of both eigenproblems (39)

| $\sigma_{1}$ | $\sigma_{2}$ | $\sigma_{3}$ | $\sigma_{4}$ | $\sigma_{5}$ | $\sigma_{6}$ | $\sigma_{7}$ | $\sigma_{8}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0.1815 | 0.1396 | 0.0906 | 0.0450 | 0.0236 | 0.0097 | 0.0035 | 0.0011 |

Denoting $\xi=\left(\xi^{\prime} \xi^{\prime \prime}\right) \in \mathbb{R}^{d}, d=d_{v}+d_{F}$, the solution $u$ is approximated in a rank- $N$ separated form: $u(x, \xi) \approx$ $\hat{u}(x, \boldsymbol{\xi})=w_{0}(x)+\sum_{n=1}^{N} w_{n}(x) \lambda_{n}(\boldsymbol{\xi})$. The stochastic approximation basis relies on an HDMR format with a maximum interaction order $N_{l}=3$ and 1D Legendre polynomials $\left\{\psi_{\alpha}\right\}_{\alpha=1}^{p}$ of maximum degree $p=8$.

In this section, the focus is on approximating a purely random quantity, i.e., disregarding its spatial dependence. We then rely on samples of the solution $u(x, \xi)$ taken at a given spatial location $x^{\star}:\left\{u^{(q)}:=u\left(x^{\star}, \xi^{(q)}\right)\right\}_{q=1}^{N_{q}}$.

### 5.1.1 Influence of the Number of Samples

We first focus on the achieved accuracy in the approximation with a given budget $N_{q}+\widehat{N_{q}}$ samples. The number of test points $\widetilde{N_{q}}$ to estimate the approximation error $\varepsilon$, Eq. (26), is chosen sufficiently large so that $\varepsilon$ is well estimated, $\widetilde{N_{q}}=10,000$. In Fig. 2, the performance of the present gLARS-ALS methodology is compared with both a plain HDMR approximation, i.e., with no subset selection hence considering the whole a priori approximation basis, and a PC approximation with a sparse grid technique. The Smolyak scheme associated with a Gauss-Patterson quadrature rule is used as the sparse grid, with varying number of points in the 1 D quadrature rule and varying levels. The dimensionality of the stochastic space is $d=8$.

The sparse grid is seen to require a large number of samples to reach a given approximation accuracy. ${ }^{6}$ The HDMR-format approximation, with various interaction orders $N_{l}$, provides a better performance than PC/Smolyak but


FIG. 2: Convergence of the approximation with the number of samples $N_{q}+\widehat{N_{q}}$. Different approximation methods are compared: plain HDMR, PC/Smolyak scheme sparse grid spectral decomposition, and the present gLARS-ALS. The convergence is plotted in terms of $\varepsilon . d=8, p=8, N_{l}=3, N_{l}^{(\mathrm{PC})}=3$.

[^6]still requires more points to reach a given accuracy than the present gLARS-ALS method which performs significantly better in approximating the QoI from a given dataset. The gLARS-ALS approximation error is also seen to be smooth and monotonic when the amount of information varies. When $N_{q}$ is large enough, the subset selection step becomes useless as all $\left|\mathcal{J}_{\text {prior }}\right|$ terms of the a priori basis can be evaluated from the large amount of information and the gLARS-ALS performance is then similar to that of the HDMR. Note that the benefit of a subset selection step in terms of accuracy improvement increases with the dimension $d$ as the size $\left|\mathcal{J}_{\text {prior }}\right|$ of the potential dictionary then grows.

### 5.1.2 Influence of the Stochastic Dimension

The approximation accuracy of the present method is now studied when the dimension of the stochastic space varies. The same problem as above is considered but with various truncation orders of the source $F$ and the diffusion coefficient $v$ definitions, see Eq. (38). The solution of the diffusion problem (37) is of dimension $d=d_{F}+d_{v}$ and the dimensions $d_{F}$ and $d_{v}$ are varied together, $d_{F}=d_{v}$. The resulting approximation error is plotted in Fig. 3 for different $d$ when the number of available samples varies. From $d=8$ to $d=40$, the required number of points for a given accuracy is seen to increase significantly, between a 2 - and a 10 -fold factor. However this is much milder than the increase in the potential approximation basis cardinality, i.e., if not subset selection was done, as $\left|\mathcal{J}_{\text {prior }}\right|$ shifts from $10,565(d=8)$ to $1.7 \times 10^{6}(d=40)$, demonstrating the efficiency of the subset selection step which activates only a small fraction of the dictionary. When $d$ further increases from 40 to 100 for a given $N_{q}$, the performance remains essentially the same with hardly any loss of accuracy: T he solution method is able to capture the low-dimensional manifold onto which the solution essentially lies and an increase in the size of the solution space hardly affects the number of samples it requires. This capability is a crucial feature when available data are scarce and the solution space is very large. As an illustration, when $d=100$, and with the parameters retained, the potential cardinality of the approximation basis is about $27 \times 10^{6}$ while the number of available samples is $\mathcal{O}(100-10,000)$. It clearly illustrates the pivotal importance of the subset selection step. Note that if one substitutes a PC approximation to the present HDMR format, about $352 \times 10^{9}$ terms need be evaluated with the present settings, a clearly daunting task.


FIG. 3: Convergence of the approximation with the number of samples $N_{q}+\widehat{N_{q}}$ and for different dimensionality of the QoI. The present gLARS-ALS approach is compared with a CANDECOMP-PARAFAC-type technique (labeled "CP").

For sake of completeness, the approximation given by a CP format, Eq. (12), is also considered. The univariate functions $\left\{f_{i, r}\right\}$ are approximated with the same polynomial approximation as in the present gLARS-ALS approach and a Tikhonov-based regularized ALS technique is used to determine each $f_{i, r}$ in turn given the others. Upon convergence, the next set of modes $\left\{f_{1, r+1}, \ldots, f_{d, r+1}\right\}$ is evaluated until a maximum rank $n_{r}$ set by cross validation. At each rank $r$, the best approximation, as estimated by cross validation, is retained from a set of initial conditions and regularization parameter values. As can be appreciated from Fig. 3, the number of samples required for a given approximation error is significantly larger than with the present gLARS-HDMR method.

### 5.1.3 Subset Selection

To further illustrate the subset selection step, the set of second- and third-order interaction retained modes $\left\{f_{\boldsymbol{\gamma}}\right\}_{\boldsymbol{\gamma} \in \mathcal{J}_{f, \text { post }}}$ are plotted in Fig. 4 in the $d=40$ case. Each bullet represents one of the $d$ stochastic dimensions and each line connects two ( 2 nd order, left plot) or three (3rd order, right plot) dimensions, denoting a retained mode. The first $d_{F}=20$ of the 40 dimensions are associated with the source term $F$ in the stochastic equation and are represented as the solid bullets of the first two quadrants, $d \in[1,20]$. The other $d_{v}=20$ dimensions are associated with the uncertain diffusion coefficient $v$ and are plotted as open bullets in the third and fourth quadrants, $d \in[21,40]$. The dimensions introduced by these two quantities are sorted with the associated magnitude of the eigenvalues $\sigma_{F}$ and $\sigma_{v}$ of their kernel, see Eqs. (39), which decreases along the counterclockwise direction. Hence, the norm of the eigenvalues of the kernel associated with $F$ decreases when one goes counterclockwise from the first to the second quadrant. Likewise, the norm of the eigenvalues associated with dimensions introduced by $v$ decreases from the third to the fourth quadrant. Dominant dimensions of the stochastic space for the output $u$ approximation are thus expected to lie at the beginning of the first and/or third quadrant.

From the plot of second-order modes (left), the subset selection process is seen to retain interaction modes mainly associated with dominant eigenvalues of both $F$ and $v$ : they mainly link bullets from the first (dominant) dimensions associated with $F$ to the first (dominant) dimensions associated with $v$, as one might expect. Further, modes associated with two dimensions both introduced by $v$ are seen to be selected while two dimensions both associated with $F$ are


FIG. 4: Graphical representation of the interaction modes retained by the subset selection procedure. Left: secondorder modes are plotted as a line linking two dimensions (bullets). Right: third-order modes are represented as 3branch stars and connect three dimensions.
rarely connected: the subset selection procedure is able to capture the nonlinearity associated with $v$ in the QoI and retains corresponding interaction modes. Indeed, note from Eq. (37) that the source term $F$ interacts linearly with the solution $u$ while the diffusion coefficient is nonlinearly coupled with $u$ and hence, interaction modes between two dimensions introduced by $F$ do not contribute to the approximation. The third-order modes (right plot) also illustrate the nonlinearity associated with $v$ : The retained modes either connect dimensions associated with $v$ only or with one $F$-related and two $v$-related dimensions. Again, no two dimensions of $F$ are connected, consistently with the linear dependence of $u$ with $F$. These results illustrate the effectiveness of the procedure to unveil the dominant dependence structure and to discard unnecessary approximation basis functions.

### 5.1.4 Robustness

The robustness of the approximation against measurement noise is now investigated. The dataset is corrupted with noise. Denoting the nominal value with a star as superscript, noise in the coordinates is modeled as

$$
\begin{equation*}
\xi^{(q)}=\xi^{(q)^{\star}}+s \zeta^{(q)}, \quad \forall 1 \leq q \leq N_{q}, \quad s>0 \tag{40}
\end{equation*}
$$

The noise is modeled as an additive $d$-dimensional, zero-centered, unit variance, Gaussian random vector $\zeta$ biased so that $\xi^{(q)} \in[-1,1]^{d}, \forall q$. It is independent from one sample $q$ to another. Without loss of generality, measurements are here modeled as being corrupted with a multiplicative noise: $u^{(q)}=u^{(q)^{\star}}\left(1+s_{u} \zeta_{u}^{(q)}\right)$, with $s_{u}=0.2$ and $\zeta_{u} \sim \mathcal{N}(0,1)$.

The evolution of the approximation accuracy when the noise intensity $s$ in the coordinates varies is plotted in Fig. 5 in terms of error estimation $\varepsilon$. We compare gLARS-ALS using standard least squares (LS) with its "robust" counterpart relying on weighted total least squares (wTLS) as discussed in Section 4.6.

When the noise intensity increases, the error exponentially increases, quickly deteriorating the quality of the approximation with a noise standard deviation here as low as $s=3 \times 10^{-5}$. When the noise is strong (low signal-to-noiseratio), both the LS and the wTLS methods achieve poor accuracy. However, if the dataset is only mildly


FIG. 5: Robustness of the approximation with respect to noise in the data: approximation error $\varepsilon$ from the standard least squares (LS) and weighted total least squares (wTLS). $d=5, N_{q}=500$.
corrupted with noise, the wTLS approach is seen to achieve a significantly better accuracy than the standard LS, while the solution process is significantly slower than that using the standard LS. The present paper is based on the assumption that the critical part of the whole solution chain of determining a good approximation of the QoI is the data acquisition and that the cost of the post-processing part is not the main issue. However, while it is useful only on a range of SNR and somehow computationally costly, this feature is deemed important for a successful solution method in an experimental context where noise is naturally present.

If the noisy dataset is unbiased, possible further improvements upon the wTLS approach include lowering the complexity of the approximation model. Indeed, the well-known bias-variance tradeoff indicates that a more robust, while less accurate, approximation can be obtained when the complexity of the retained model decreases. To improve the robustness of our present approach, a natural way is hence to trade some accuracy for some additional robustness. For instance, a predictor selection within each retained groups $\left\{f_{\gamma \in \mathcal{J}_{f, \text { post }}}\right\}$ can be considered, further lowering the final number of coefficients involved in the approximation and likely improving its robustness with respect to noise in the data. This could be achieved by estimating the approximation coefficients via a penalized (total) LS problem as mentioned in Section 2.2.2.

### 5.1.5 Scaling of the Solution Algorithm

In this section, the numerical complexity associated with the different steps of the solution method is illustrated in terms of computational time. Numerical experiments are carried out with varying number of samples $N_{q}$ and solution space dimensions $d$. When one is varying, the other remains constant. The nominal parameters are $d=40$ (dimension of the stochastic space), $p=6$ (maximum total order of the Legendre polynomials), $N_{q}=1000$ (number of samples), $N_{l}=3$ (maximum interaction order of the truncated HDMR approximation), $\widetilde{p}=5$ (maximum total polynomial order in the subset selection step).

Numerical results are gathered in Fig. 6. The asymptotic behavior of the number $n_{f}$ of required subset selection iterations as introduced in Section 4.7 might be different according to which limit is considered. For the present stochastic diffusion problem, first and second interaction order modes tend to be selected first. Assuming the active set $\mathcal{J}_{f, \text { post }}$ is dominated by first and second interaction order modes, it can easily be shown that the number of retained groups then satisfies


FIG. 6: Numerical cost of the subset selection and coefficients evaluation steps as a function of the stochastic dimension $d$ and size of the dataset $N_{q}$. Approximation coefficients are fully updated for each new mode. Nominal parameters are $d=40, p=6, N_{q}=1000, N_{l}=3, \widetilde{p}=5, N_{l}^{(\mathrm{PC})}=3$.

$$
\begin{equation*}
n_{f} \leq 1+n_{f_{1}}+\min \left[\frac{d(d-1)}{2}, 2 \frac{N_{q}-n_{f_{1}} \widetilde{p}}{\widetilde{p}(\widetilde{p}+1)}\right], \quad n_{f_{1}} \leq \min \left[d, \frac{N_{q}}{\widetilde{p}}\right] \tag{41}
\end{equation*}
$$

In the present example, second-order interaction groups dominate the retained set so that the number of retained groups tends to scale as $n_{f} \propto N_{q} / \widetilde{p}^{2}$. From Eq. (33) and for the present nominal parameters, it results in the following limit behavior for the subset selection step:

$$
\begin{array}{rll}
\lim _{d \rightarrow+\infty} \mathscr{J}_{\text {subsel }} & \propto N_{q}{ }^{2}\left|\mathcal{J}_{\text {prior }}\right| / \widetilde{p}^{2} & \longrightarrow \text { here }: \propto N_{q}{ }^{2} d^{N_{l}} \widetilde{p}^{N_{l}-2} \\
\lim _{N_{q} \rightarrow+\infty} \mathscr{J}_{\text {subsel }} & \propto N_{q}{ }^{2}\left|\mathcal{J}_{\text {prior }}\right| / \widetilde{p}^{2} & \longrightarrow \text { here }: \tag{42}
\end{array} N_{q}{ }^{2} d^{N_{l}} \widetilde{p}^{N_{l}-2} .
$$

Similarly, the cost associated with the coefficients evaluation is considered. The number of interaction modes $\mathcal{J}_{f \text {, eff }}$ effectively varies between 1 and $\mathcal{O}\left(N_{q} / p^{2}\right)$ along the solution procedure, and, since the cost associated with solving the LS problem dominates that of the matrix assembly, the cost of their evaluation finally simplifies in $\mathscr{J}_{\text {coef }} \propto$ $\mathcal{O}\left(N_{q}{ }^{2} p\right)$ or $\mathscr{J}_{\text {coef }} \propto \mathcal{O}\left(N_{q}{ }^{3} / p\right)$ depending on whether the coefficients are updated whenever an additional group is considered or not, see Section 4.5 and step (7) in Algorithm 1. In the present regime, the cost is found not to depend on $d$.

These asymptotic behaviors are consistent with the numerical experiments as can be appreciated from Fig. 6. The coefficients are here updated whenever a new mode from the selected set is considered, hence $\mathscr{J}_{\text {coef }} \propto \mathcal{O}\left(N_{q}{ }^{3} / p\right)$. It is seen that the subset selection step scales less favorably than the coefficients evaluation step with the dimensionality of the random variable. This stresses the benefit of a carefully chosen a priori approximation basis to reduce as much as possible the cardinality $\left|\mathcal{J}_{\text {prior }}\right|$.

### 5.2 Approximation of the Solution Random Field

We now consider the approximation of the space-dependent random solution $u(x, \boldsymbol{\xi})$ under the form (35) using Algorithm 2. The approximation obtained from different number of samples $\left\{x^{(q)}, \mathfrak{\xi}^{(q)}, u^{(q)}\right\}$ is compared with the Karhunen-Loève modes, computed from a full knowledge of the QoI, hereafter referred to as the reference solution. ${ }^{7}$

The simulation relies on the following parameters: $\left|\mathcal{J}_{x}\right|=32, p=10, d=6, N_{l}=3, N_{l}^{(\mathrm{PC})}=3$. The potential approximation basis cardinality is about $\left|\mathcal{J}_{x}\right|\left|\mathcal{J}_{\text {prior }}\right| \simeq 10^{5}$. Fig. 7 shows the first and second spatial modes, $w_{1}(x)$ and $w_{2}(x)$ for different sizes of the dataset, $N_{q}=1000,3000,9000$, and 26,000 . The mean mode $w_{0}(x)$ is virtually indistinguishable from the reference solution mean mode for any of the dataset sizes and is not plotted. On the left plot $\left[w_{1}(x)\right]$, it is seen that the approximation is decent, even with as low as $N_{q}=1000$ samples. For $N_{q}=3000$, the approximation is good. This $(1+d)=7$-dimensional case corresponds to $N_{q}{ }^{1 /(1+6)} \simeq 3.1$ samples per solution space dimension only and about $N_{q} /\left(\left|\mathcal{J}_{x}\right|\left|\mathcal{J}_{\text {prior }}\right|\right) \simeq 3 \%$ of the potentially required information.

For approximating the second spatial mode (Fig. 7, right plot), more points are needed to reach a good accuracy but $N_{q}=26,000$ is seen to already deliver a good performance. Quantitative approximation error results are gathered in Table 2 for various separation ranks $N$ and number of samples $N_{q}$.

The satisfactory performance of the present method can be understood from the upper part of the Karhunen-Loève approximation (normalized) spectrum plotted in Table 3. The norm of the eigenvalues decays quickly so that the first two modes contribute more than $90 \%$ of the QoI $L^{2}$-norm, showing that this problem efficiently lends itself to the present separation of variables-based methodology.

[^7]

FIG. 7: First [ $w_{1}(x)$, left] and second [ $w_{2}(x)$, right] spatial approximation modes of the stochastic diffusion solution. The reference (Karhunen-Loève) solution is plotted for comparison (thick line).

TABLE 2: Evolution of the approximation error $\varepsilon$, as defined in Eq. (26), with the decomposition rank $N$ and the number of samples $N_{q}$

| $N_{q} \backslash N$ | 0 | 1 | 2 |
| ---: | :---: | :---: | :---: |
| 1000 | $5.5 \times 10^{-3}$ | $7.4 \times 10^{-4}$ | $7.4 \times 10^{-4}$ |
| 3000 | $5.5 \times 10^{-3}$ | $4.2 \times 10^{-4}$ | $2.7 \times 10^{-4}$ |
| 9000 | $5.5 \times 10^{-3}$ | $3.1 \times 10^{-4}$ | $1.0 \times 10^{-4}$ |
| 26,000 | $5.4 \times 10^{-3}$ | $2.8 \times 10^{-4}$ | $6.2 \times 10^{-5}$ |

TABLE 3: Normalized upper spectrum of the Karhunen-Loève approximation

| $i$ | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\sigma_{i}$ | 144 | 30.2 | 13.6 | 2.64 | 1.37 | 0.250 | 0.0817 | 0.0167 | 0.00891 | 0.00232 |

### 5.3 A Shallow Water Flow Example

The methodology is now applied to the approximation of the stochastic solution of a shallow water flow simulation with multiple sources of uncertainty. It is a simple model for the simulation of wave propagation on the ocean surface. Waves are here produced by the sudden displacement of the sea bottom at a given magnitude in time, extension, and location, all uncertain.

### 5.3.1 Model

The problem is governed by the following set of equations:

$$
\begin{align*}
\frac{D v_{1}}{D t} & =f_{C} v_{2}-g \frac{\partial h}{\partial x_{1}}-b v_{1}+S_{v_{1}}  \tag{43}\\
\frac{D v_{2}}{D t} & =-f_{C} v_{1}-g \frac{\partial h}{\partial x_{2}}-b v_{2}+S_{v_{2}} \tag{44}
\end{align*}
$$

$$
\begin{equation*}
\frac{\partial h}{\partial t}=-\frac{\partial\left(v_{1}(H+h)\right)}{\partial x_{1}}-\frac{\partial\left(v_{2}(H+h)\right)}{\partial x_{2}}+S_{h} \tag{45}
\end{equation*}
$$

where $\left(v_{1}(\boldsymbol{x}, \boldsymbol{\xi}, t) v_{2}(\boldsymbol{x}, \boldsymbol{\xi}, t)\right)$ is the velocity vector at the surface, $\boldsymbol{x}=\left(x_{1} x_{2}\right) \in \Omega \subset \mathbb{R}^{2}, h(\boldsymbol{x}, \boldsymbol{\xi}, t)$ the elevation of the surface from its position at rest, $H(\boldsymbol{x})$ the sea depth, $f_{C}$ models the Coriolis force, $b$ is the viscous drag coefficient, $g$ the gravity constant and $S_{v_{1}}(\boldsymbol{x}, \boldsymbol{\xi}, t), S_{v_{2}}(\boldsymbol{x}, \boldsymbol{\xi}, t), S_{h}(\boldsymbol{x}, \boldsymbol{\xi}, t)$ are the source fields. Without loss of generality, the drag $b$ and the Coriolis force $f_{C}$ are neglected. No slip boundary conditions apply for the velocity. The sources are modeled as acting on $h$ only, $S_{v_{1}} \equiv 0$ and $S_{v_{2}} \equiv 0 . S_{h}$ models the source term acting on $h$ due to, say, an underwater seismic event. The fluid density and the free surface pressure are implicitly assumed constant. Full details on the numerical implementation of a similar problem are given in [47].

### 5.3.2 Sources of Uncertainty

Let $\xi=\left(\xi^{\prime} \xi^{\prime \prime}\right)$. The source $S_{h}$ is uncertain and is modeled as a time-dependent, spatially distributed quantity:

$$
\begin{equation*}
S_{h}(\boldsymbol{x}, \boldsymbol{\xi}, t)=a_{t}\left(\boldsymbol{\xi}^{\prime}, t\right) a_{\xi}\left(\xi_{1}^{\prime \prime}\right) \exp \left(-\frac{\left(\boldsymbol{x}-\boldsymbol{x}_{S_{h}}\left(\xi_{3}^{\prime \prime}\right)\right)^{T}\left(\boldsymbol{x}-\boldsymbol{x}_{S_{h}}\left(\xi_{3}^{\prime \prime}\right)\right)}{\sigma_{S_{h}}\left(\xi_{2}^{\prime \prime}\right)^{2}}\right) \tag{46}
\end{equation*}
$$

where $a_{t}\left(\xi^{\prime}, t\right)$ is a given time envelope, $a_{\xi}\left(\xi_{1}^{\prime \prime}\right)$ the uncertain source magnitude, $\sigma_{S_{h}}\left(\xi_{2}^{\prime \prime}\right)$ drives the uncertain source spatial extension, and $\boldsymbol{x}_{S_{h}}\left(\xi_{3}^{\prime \prime}\right)$ is the uncertain spatial location. The time envelope $a_{t}\left(\xi^{\prime}, t\right)$ is described with an $N_{a}$-term expansion:

$$
\begin{equation*}
a_{t}\left(\xi^{\prime}, t\right)=\overline{a_{t}}(t)+\sum_{i=1}^{N_{a}} \sqrt{\lambda_{i}} \xi_{i}^{\prime}(\theta) \varphi_{i}^{a_{t}}(t), \tag{47}
\end{equation*}
$$

with $\xi^{\prime}=\left(\xi_{1}^{\prime} \ldots \xi_{N_{a}}^{\prime}\right)$ the stochastic germ associated with the uncertainty in $a_{t}$. Random variables $\left\{\xi_{i}^{\prime}\right\}_{i=1}^{N_{a}}$ are iid, uniformly distributed. The solution of the shallow water problem then lies in a ( $d=N_{a}+3$ )-dimensional stochastic space.

### 5.3.3 Approximation from an Available Database

As an illustration of the methodology, we aim at approximating the sea surface field at a fixed amount of time $t^{\star}$ after a seismic event. The QoI is then a random field $u(\boldsymbol{x}, \boldsymbol{\xi})=h\left(\boldsymbol{x}, \boldsymbol{\xi}, t^{\star}\right)$. An accurate description of this field is of importance for emergency plans in case of a seaquake. Sea level measurements of the surface at various spatial locations from past events constitute the dataset $\left\{\boldsymbol{x}^{(q)}, \boldsymbol{\xi}^{(q)}, h\left(\boldsymbol{x}^{(q)}, \boldsymbol{\xi}^{(q)}, t^{\star}\right)\right\}_{q=1}^{N_{q}}$ used to derive an approximation of $u$ under a separated form: $u(\boldsymbol{x}, \boldsymbol{\xi}) \approx\langle u\rangle_{N_{q}}(\boldsymbol{x})+\sum_{n=1}^{N} w_{n}(\boldsymbol{x}) \lambda_{n}(\boldsymbol{\xi})$.

The solution method here relies on a $N_{q}=37,000$-sample dataset complemented with $\widehat{N_{q}}=5000$ crossvalidation samples and a $\widetilde{N_{q}}=5000$ set for error estimation. We consider a $N_{a}=5$ expansion for the time envelope, leading to a stochastic dimension of $d=5+3=8$. The effective number of samples per dimension is then about $N_{q}{ }^{1 /\left(d_{x}+d\right)} \simeq 2.9$. The approximation is determined based on a $\left|\mathcal{J}_{x}\right|=484$ spatial discretization DOFs (spectral elements) at the deterministic level and $p=6$ th order Legendre polynomials $\left\{\psi_{\alpha}\right\}, N_{l}=3, N_{l}^{(\mathrm{PC})}=3$, for the stochastic modes. The cardinality of this a priori basis is then $\left|\mathcal{J}_{x}\right|\left|\mathcal{J}_{\text {prior }}\right| \simeq 770 \times 10^{3} \gg N_{q}$, again relying on an efficient subset selection step to make the approximation problem well-posed.

The approximation error when the rank $N$ varies is shown in Table 4. It is seen that estimating the mean spatial mode $w_{0}$ leads to a relative error of about 0.12 while adding the first $\left(w_{1}, \lambda_{1}\right)$ and second $\left(w_{2}, \lambda_{2}\right)$ pair drops it to about 0.05 . Further adding pairs does not lower the approximation error with this dataset and more samples are needed to accurately estimate them. Spatial modes $w_{0}$ and $w_{1}$ of the separated approximation are plotted in Fig. 8 for illustration.

TABLE 4: Relative approximation error $\varepsilon$ evolution with the decomposition rank $N . N_{q}=37,000$

| $N$ | 0 | 1 | 2 | 3 |
| :---: | :---: | :---: | :---: | :---: |
| $\varepsilon$ | 0.117 | 0.056 | 0.046 | 0.044 |



FIG. 8: Mean $\left[w_{0}(\boldsymbol{x}) \equiv\langle u\rangle_{N_{q}}(\boldsymbol{x})\right.$, left $]$ and first $\left[w_{1}(\boldsymbol{x})\right.$, right $]$ spatial modes.

## 6. CONCLUSION

In this paper, a methodology was proposed for deriving a functional representation of a random process only known through a collection of its pointwise evaluations. The proposed method essentially relies on an efficient determination of an approximation basis consistent with the available information. This involves the choice of an a priori canonical HDMR format combined with tuning the basis via a data-driven subset selection step. This subset selection is carried out in a bottom-to-top manner, as opposed to a top-to-bottom manner as done in the compressed sensing standard framework. It essentially sorts the HDMR modes (groups of predictors) by their contribution in approximating the quantity of interest. The final approximation can rely on a different functional description of the modes, typically of higher order and/or nonlinear in the coefficients.

The method is progressive, data-driven, and was shown to here outperform current approximation techniques in terms of accuracy for a given number of samples. Its efficiency was demonstrated in two examples which have shown its ability to achieve a good approximation accuracy from a small dataset, as long as the quantity at hand is essentially lying on a low-dimensional manifold. In particular, the dominant dimensions are naturally revealed so that all the available information can be dedicated to approximate relevant dependences only. Through a total least squares approach, it was also shown that some robustness can be achieved, an important feature if the dataset comes from experiments. Using a robust approximation was shown to bring up to a twofold improvement upon the approximation error using standard least squares, but at the price of a computational overhead. The global solution method scales reasonably well, exhibiting a linear dependence with the cardinality of the a priori basis dictionary and a quadratic or cubic dependence with the number of samples, depending on the coefficients update strategy.

The present work was focused on a general methodology, disregarding fine-tuning aspects. Among other things, a natural improvement would be to carry out a predictor selection within each retained groups $\left\{f_{\gamma \in \mathcal{J}_{f, \text { post }}}\right\}$, further lowering the number of coefficients involved in the approximation. Moreover, the tensor structure of the Hilbert stochastic space can be exploited and developments towards a data-driven multilinear algebra effective tool for highdimensional uncertainty quantification are currently carried out.

## ACKNOWLEDGMENTS

The author gratefully acknowledges Tarek El Moselhy and Faidra Stavropoulou for stimulating discussions and useful comments. This work is part of the TYCHE project (ANR-2010-BLAN-0904) supported by the French Research National Agency (ANR).

## REFERENCES

1. Nouy, A., A generalized spectral decomposition technique to solve a class of linear stochastic partial differential equations, Comput. Methods Appl. Mech. Eng., 196(45-48):4521-4537, 2007.
2. Nouy, A., A priori model reduction through proper generalized decomposition for solving time-dependent partial differential equations, Comput. Methods Appl. Mech. Eng., 199:1603-1626, 2010.
3. Nouy, A., Proper generalized decompositions and separated representations for the numerical solution of high dimensional stochastic problems, Arch. Comput. Methods Eng., 17:403-434, 2010.
4. Matthies, H. and Zander, E., Solving stochastic systems with low-rank tensor compression, Linear Algebra Appl., 436(10):3819-3838, 2012.
5. Novak, E. and Ritter, K., Simple cubature formulas with high polynomial exactness, Construct. Approx., 15:499-522, 1999.
6. Xiu, D. and Hesthaven, J., High-order collocation methods for differential equations with random inputs, SIAM J. Sci. Comput., 27:1118-1139, 2005.
7. Nobile, F., Tempone, R., and Webster, C., An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data, SIAM J. Numer. Anal., 46(5):2411-2442, 2007.
8. Ganapathysubramanian, B. and Zabaras, N., Sparse grid collocation methods for stochastic natural convection problems, J. Comput. Phys., 225:652-685, 2007.
9. Ma, X. and Zabaras, N., An adaptive high-dimensional stochastic model representation technique for the solution of stochastic partial differential equations, J. Comput. Phys., 10:3884-3915, 2010.
10. Doostan, A. and Iaccarino, G., A least-squares approximation of partial differential equations with high-dimensional random inputs, J. Comput. Phys., 228:4332-4345, 2009.
11. Choi, S.-K., Grandhi, R., Canfield, R., and Pettit, C., Polynomial chaos expansion with latin hypercube sampling for estimating response variability, AIAA J., 42(6):1191-1198, 2004.
12. Berveiller, M., Sudret, B., and Lemaire, M., Stochastic finite element: A non intrusive approach by regression, J. Eur. Meca. Num., 15(1-2-3):81-92, 2006.
13. Beylkin, G., Garcke, J., and Mohlenkamp, M., Multivariate regression and machine learning with sums of separable functions, SIAM J. Sci. Comput., 31(3):1840-1857, 2009.
14. Doostan, A. and Owhadi, H., A non-adapted sparse approximation of PDEs with stochastic inputs, J. Comput. Phys., 230(8):3015-3034, 2011.
15. Mathelin, L. and Gallivan, K., A compressed sensing approach for partial differential equations with random input data, Commun. Comput. Phys., 12:919-954, 2012.
16. Candès, E. and Tao, T., Decoding by linear programming, IEEE Trans. Inform. Theory, 51:4203-4215, 2004.
17. Donoho, D., Compressed sensing, IEEE Trans. Infor. Theo., 52(4):1289-1306, 2006.
18. Rabitz, H. and Alış, O., General foundations of high-dimensional model representations, J. Math. Chem., 25:197-233, 1999.
19. Alış, O. and Rabitz, H., Efficient implementation of high-dimensional model representations, J. Math. Chem., 29(2):127-142, 2001.
20. Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R., Least angle regression, Ann. Stat., 32:407-499, 2004.
21. Abramowitz, M. and Stegun, I., Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, 10th edition, Dover Publications, New York, 1972.
22. Hesterberg, T., Choi, N. H., Meier, L., and Fraley, C., Least angle and $\ell_{1}$ penalized regression: A review, Stat. Surveys, 2:6193, 2008.
23. Hastie, T., Tibshirani, R., and Friedman, J., The Elements of Statistical Learning, 2nd ed., Springer Verlag, Berlin, 2009.
24. Wiener, N., The homogeneous chaos, Am. J. Math., 60(4):897-936, 1938.
25. Ghanem, R. and Spanos, P., Stochastic Finite Elements. A Spectral Approach, rev. edition, Springer Verlag, Berlin, 2003.
26. Xiu, D. and Karniadakis, G., The Wiener-Askey polynomial chaos for stochastic differential equations, SIAM J. Sci. Comput., 24(2):619-644, 2002.
27. Soize, C. and Ghanem, R., Physical systems with random uncertainties: Chaos representations with arbitrary probability measure, SIAM J. Sci. Comput., 26(2):395-410, 2004.
28. Le Maître, O. and Knio, O., Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics, Springer, Berlin, 2010.
29. Khoromskij, B., Tensors-structured numerical methods in scientific computing: Survey on recents advances, Chemometrics Intell. Lab. Syst., 110(1):1-19, 2012.
30. Khoromskij, B. and Schwab, C., Tensor-structured Galerkin approximation of parametric and stochastic elliptic pdes, SIAM J. Sci. Comput., 33(1):1-25, 2011.
31. Harshman, R., Foundations of the parafac procedure: Models and conditions for an "explanatory" multimodal factor analysis, UCLA Working Papers in Phonetics, 16:1-84, 1970.
32. Carroll, J. and Chang, J.-J., Analysis of individual differences in multidimensional scaling via an n-way generalization of "Eckart-Young" decomposition, Psychometrika, 35:283-319, 1970.
33. Kuo, F., Sloan, I., Wasilkowski, G., and Woźniakoski, H., On decompositions of multivariate functions, Math. Comput., 79(270):953-966, 2009.
34. Candès, E. and Romberg, J., Sparsity and incoherence in compressive sampling, Inv. Problems, 23:969-985, 2006.
35. DeVore, R., Petrova, G., and Wojtaszczyk, P., Approximation of functions of few variables in high dimensions, Constr. Approx., 33:125-143, 2011.
36. Tibshirani, R., Regression shrinkage and selection via the lasso, J. R. Statist. Soc. Ser. B, 58:267-288, 1996.
37. Chen, S., Donoho, D., and Saunders, M., Atomic decomposition by basis pursuit, SIAM J. Sci. Comput., 20:33-61, 1999.
38. Candès, E. and Tao, T., Near-optimal signal recovery from random projections: Universal encoding strategies, IEEE Trans. Inform. Theory, 52:5406-5425, 2004.
39. Cai, T., Wang, L., and Xu, G., New bounds for restricted isometry constants, IEEE Trans. Inform. Theory, 56(9):4388-4394, 2010.
40. Blatman, G. and Sudret, B., Adaptive sparse polynomial chaos expansion based on least angle regression, J. Comput. Phys., 230(6):2345-2367, 2011.
41. Yuan, M. and Lin, Y., Model selection and estimation in regression with grouped variables, J. R. Statist. Soc. Ser. B, 68:49-67, 2006.
42. Xie, J. and Zeng, L., Group variable selection methods and their applications in analysis of genomic data, in Feng, J., Fu, W., and Sun, F. (Eds.), Frontiers in Computational and Systems Biology, Computational Biology, Vol. 15, pp. 231-248, Springer-Verlag, Berlin, 2010.
43. Golub, G. and van Loan, C., Matrix Computations, 4th edition, JHU Press, Baltimore, MD, USA, 2012.
44. Markovsky, I. and van Huffel, S., Overview of total least squares methods, Signal Processing, 87(10):2283-2302, 2007.
45. Wentzell, P., Andrews, D., Hamilton, D., Faber, K., and Kowalski, B., Maximum likelihood principal component analysis, J. Chemometrics, 11:339-366, 1997.
46. Grasedyck, L., Hierarchical singular value decomposition of tensors, SIAM J. Matrix Anal. Appl., 31(4):2029-2054, 2010.
47. Mathelin, L., Desceliers, C., and Hussaini, M., Stochastic data assimilation with a polynomial chaos parametric estimation, Comput. Mech., 47(6):603-616, 2011.
48. Sobol, I., Sensitivity estimates for nonlinear mathematical models, Math. Modell. Comput. Exper., 1(4):407-414, 1993.
49. Homma, T. and Saltelli, A., Importance measures in global sensitivity analysis of nonlinear models, Rel. Eng. Syst. Safety, 52(1):1-17, 1996.

## APPENDIX A. A MOTIVATING EXAMPLE

To assess the choice of our a priori functional form for approximating a random variable, and while choosing a good basis is problem-dependent, let us consider a simple motivating example in the form of the 1D stochastic diffusion equation presented in Section 5.1, briefly recalled here for sake of convenience:

$$
\begin{equation*}
\nabla_{x}\left(\nu(x, \boldsymbol{\xi}) \nabla_{x} u(x, \boldsymbol{\xi})\right)=F(x, \boldsymbol{\xi}), \quad u\left(x_{-}, \boldsymbol{\xi}\right)=u_{-}, u\left(x_{+}, \boldsymbol{\xi}\right)=u_{+} . \tag{A.1}
\end{equation*}
$$

The solution $u$ is approximated under a separated format $u(x, \boldsymbol{\xi}) \approx \sum_{n=0}^{N} w_{n}(x) \lambda_{n}(\boldsymbol{\xi})$. The approximation space for the spatial modes $\left\{w_{n}(x)\right\}$ is given and we here focus on the accuracy of the approximation with different representations for the stochastic modes $\left\{\lambda_{n}(\boldsymbol{\xi})\right\}$. Each stochastic mode is determined either in a CP-like format, Eq. (12), or as an HDMR decomposition Eq. (13). In the latter case, interaction modes $\left\{f_{\gamma}\right\}$ are approximated with a low-rank canonical decomposition on tensorized, unit-normed, univariate polynomials of maximum degree $p$ : $f_{\gamma}\left(\left\{\xi_{i}\right\}_{i \in \gamma}\right) \approx \sum_{r=1}^{n_{r}} \prod_{i \in \gamma} \sum_{\alpha=2}^{p} c_{\gamma, \alpha}^{r, i} \psi_{\alpha}\left(\xi_{i}\right)$. This approximation is hereafter referred to as a CP-HDMR decomposition. Similarly, univariate functions $\left\{f_{i, r}\right\}_{i=1}^{d}$ involved in the CP decomposition Eq. (12) are approximated with the same polynomials: $f_{i, r}\left(\xi_{i}\right) \approx \sum_{\alpha=1}^{p} c_{\alpha, i, r} \psi_{\alpha}\left(\xi_{i}\right)$.

The representation of the stochastic modes here relies on $p=8$ th-order univariate Legendre polynomials $\left\{\psi_{\alpha}\right\}$. The dimension of the problem is chosen to be $d_{v}=d_{F}=5$ so that $d=10$.

The CP-HDMR expansion is here built sequentially, starting with 0th and first-order interaction modes only. From this first approximation of the output, the set of dominant dimensions is estimated from the $L^{2}$-norm of univariate interaction modes $\left\{f_{\boldsymbol{\gamma}}\right\}_{|\boldsymbol{\gamma}|=1}$. Only second-order interaction modes $\left\{f_{\boldsymbol{\gamma}}\right\}_{|\boldsymbol{\gamma}|=2}$ in these dominant dimensions are next estimated and the set of dominant dimensions is then further refined based on both first and second-order interaction modes via the sensitivity Sobol indices, see Appendix B. Third-order modes are then computed for this new set of dominant dimensions only and the procedure is repeated until some stopping criterion is met, for instance a maximum interaction order $N_{l}$ or a maximum basis cardinality $|\mathcal{J}|$. The number of samples $N_{q}$ is here chosen sufficiently large so that full knowledge of $u$ can be assumed. The approximation error then only comes from the choice of the approximation basis format, allowing a comparison. This section is loose on details, focusing on the main conclusions and leaving more in-depth discussion for main text sections.

First, the accuracy of the CP-HDMR approximation as a function of the decomposition rank $N$ is studied in terms of $\varepsilon$, Eq. (26). Plotted in Fig. A.1, the approximation error estimation $\varepsilon$ decreases when the maximum interaction order $N_{l}$ increases from 1 to 3 and as the decomposition rank $N$ increases.

The approximation is seen to improve exponentially fast as the number of modes $N$ in the separated representation increases until it reaches a plateau. Increasing the interaction order leads to an improved approximation: increasing from first to second order brings more than a one-order-of-magnitude improvement in the approximation error estimation and an additional order of magnitude from $N_{l}=2$ to $N_{l}=3$. In this $d=10$ example, the approximation hence exhibits a high convergence rate with $N_{l}$, supporting our assumption that low-order interactions dominate the HDMR decomposition.

This CP-HDMR approximation of the stochastic modes is now compared with a CP-like approximation in the form of Eq. (12). To evaluate the CP decomposition, we use an algorithm similar to that in [3]. Both decompositions rely on the same approximation basis for the deterministic modes $\left\{w_{n}\right\}$. We focus on the accuracy of the reconstruction as a function of the cardinality of the whole approximation basis both for $d=10$ and $d=40$ when the maximum decomposition rank $N$ varies, see Fig. A.2. The total cardinality increases as more terms are considered in the decomposition series.

The accuracy of the representation is seen to improve as more terms are considered in the series expansion and both the CP and the CP-HDMR formats exhibit an exponential convergence with the total size $|\mathcal{J}|$ of the decomposition. The first-order CP-HDMR quickly reduces the error but plateaus as the functional space is small. The third order CPHDMR decomposition is more costly in number of coefficients to evaluate to reach a given error and, in the present settings, should only be considered if high accuracy is needed.


FIG. A.1: Convergence of the error estimation of the approximation with the order $N$ of the separated representation and the maximum interaction order $N_{l}$ of the HDMR expansion.


FIG. A.2: Convergence of the approximation error estimation $\varepsilon$ with the total cardinality $|\mathcal{J}|$ of the representation basis. Approximations of the stochastic modes with the CP-HDMR and CP-like format are compared. $d=10$ (left) and $d=40$ (right).

Unless the targeted accuracy is really high, this motivating example tends to indicate that, for a reasonable required accuracy, a CP-HDMR format involves fewer unknowns than a CP-like decomposition, both for a low $d=10$ and a moderate $d=40$-dimensional problem. This is an important point since the number of coefficients which can be evaluated with a reasonable accuracy from experimental data is directly related to the size of the available dataset. Finally, a CP-HDMR format allows a great flexibility in representing interaction modes $\left\{f_{\gamma}\right\}$. In particular,
a more parsimonious representation is used in the main text and achieves a similar accuracy with a lower number of terms.

## APPENDIX B. STATISTICS AND SENSITIVITY ANALYSIS

Once an approximation of a random variable $u(\xi)$ is obtained, it is easy to estimate its first statistical moments. From the HDMR format properties, the estimated mean is simply given by the first term of the decomposition: $\langle u\rangle_{L^{2}\left(\Xi, \mu_{\Xi)}\right.} \simeq f_{\emptyset}$.

Thanks to the orthogonality property of the modes $\left\{f_{\gamma}\right\}$, the variance $\operatorname{Var}(u):=\left\langle\left(u-\langle u\rangle_{L^{2}\left(\Xi, \mu_{\Xi}\right)}\right)^{2}\right\rangle_{L^{2}\left(\Xi, \mu_{\Xi)}\right.}$ approximates as the sum of the variance of the individual interaction modes:

$$
\begin{align*}
\operatorname{Var}(u) & \simeq \sum_{\boldsymbol{\gamma} \in \mathcal{J}_{f, \text { eff }} \backslash \emptyset} \operatorname{Var}\left(\widehat{f}_{\boldsymbol{\gamma}}\right), \\
& =\sum_{\substack{\boldsymbol{\gamma} \in \mathcal{J}_{f, \text { eff } \backslash \emptyset} \backslash \boldsymbol{\alpha},|\boldsymbol{\alpha}| \leq p \\
|\boldsymbol{\gamma}| \leq N_{l}^{(\mathrm{PC})}}} c_{\boldsymbol{\gamma}, \boldsymbol{\alpha}}^{2}\left\|\psi_{\boldsymbol{\alpha}}\right\|_{L^{2}\left(\Xi, \mu_{\Xi}\right)}^{2}+\sum_{\substack{\boldsymbol{\gamma} \in \mathcal{J}_{f, \text { eff }} \backslash \emptyset \\
N_{l}^{(\mathrm{PCC}}<|\boldsymbol{\gamma}| \leq N_{l}}} \sum_{r, r^{\prime}=1}^{n_{r}(\boldsymbol{\gamma})} \prod_{i \in \boldsymbol{\gamma}} \sum_{\alpha=1}^{p} c_{\boldsymbol{\gamma}, \boldsymbol{\alpha}}^{r, i} c_{\boldsymbol{\gamma}, \boldsymbol{\alpha}}^{r^{\prime}, i}\left\|\psi_{\alpha}\right\|_{L^{2}\left(\Xi, \mu_{\Xi}\right),}^{2}, \tag{B.1}
\end{align*}
$$

where use was made of the orthogonality of the Hilbertian basis $\left\{\psi_{\alpha}\right\}$.
Other standard statistical quantities are the sensitivity indices $\left\{S_{\gamma}\right\}:=\operatorname{Var}\left(\widehat{f}_{\gamma}\right) / \operatorname{Var}(\widehat{u})$ which essentially represent the relative part of the variance of the QoI due to the interaction of a given set of input random variables only, [48, 49]. From Eq. (B.1), it immediately follows that $\sum_{\gamma \subseteq\{1, \ldots, d\} \backslash \emptyset} S_{\gamma}=1$ and the explicit expression of the sensitivity indices is straightforward to derive from the HDMR format. In practice, it is often more useful to assess the influence of a given input onto the variance of the QoI with the total sensitivity indices $\left\{S_{T, i}\right\}_{i=1}^{d}$ :

$$
\begin{equation*}
S_{T, i}:=\frac{\sum_{\boldsymbol{\gamma} \subseteq\{1, \ldots, d\} \backslash \emptyset: i \in \boldsymbol{\gamma}} \operatorname{Var}\left(\widehat{f}_{\boldsymbol{\gamma}}\right)}{\operatorname{Var}(\widehat{u})}, \quad 1 \leq i \leq d \tag{B.2}
\end{equation*}
$$

Again, using Eq. (B.1), this quantity is straightforward to estimate once the approximation of $u(\boldsymbol{\xi})$ is available.


[^0]:    *Correspond to Lionel Mathelin, E-mail: mathelin@limsi.fr

[^1]:    ${ }^{1}$ For instance, in a fluid flow, the Reynolds number may be uncertain and modeled as a random variable parametrized by $\xi_{i}$. The value of $\xi_{i}$ in each sample $\xi^{(q)}$ is then auxiliary deduced from the measurement of the flow velocity $V$ and the model $V\left(\xi_{i}\right)$.

[^2]:    ${ }^{2}$ For a sufficiently incoherent set of approximation and test functions, a $K$-sparse solution $\boldsymbol{c}$ to Eq. (16) satisfies [39] $\left\|\boldsymbol{c}^{\star}-\boldsymbol{c}\right\|_{2} \lesssim$ $h\left(\epsilon+\left\|\boldsymbol{c}^{\star}-\boldsymbol{c}_{K}^{\star}\right\|_{1} / \sqrt{K}\right)$, where $h>0$ is a constant depending on the set of approximation and test functions and $\boldsymbol{c}_{K}^{\star}$ is the $K$-term approximation of $\boldsymbol{c}^{\star}$ given by an oracle, i.e., it is the best $K$-term approximation of $\boldsymbol{c}^{\star}$ if one was given full knowledge of it.

[^3]:    ${ }^{3}$ In a nutshell, it consists of selecting, from the a priori set $\mathcal{J}_{\text {prior }}$, the predictor (approximation function) which is most correlated with the current residual, move this predictor to the active set $\mathcal{J}_{\text {post }}$, compute the increment solution vector by minimizing the residual $L^{2}$-norm, and follow the descent direction along the increment vector until a predictor from the inactive set becomes as correlated with the residual as those from the active set. The whole process is then repeated and allows to sequentially build the optimal subset of approximation functions by exploring the Pareto front defined by the competition between the two terms of the unconstrained formulation of the optimization problem of Eq. (16).

[^4]:    ${ }^{4}$ While not found necessary here, the solution of the LS problem may be regularized by adding a generic term of the form $\beta\|L \widetilde{\boldsymbol{c}}\|_{2}$. A typical choice is $L=I_{\left|\mathcal{J}_{\gamma}\right|}$ but one may also want to consider nondiagonal matrices $L$.

[^5]:    ${ }^{5}$ A ratio $\widehat{N_{q}} / N_{q} \simeq 1 / 2$ is typically accepted as a reasonable splitting of the set of samples. We here use the simplest crossvalidation method but more sophisticated techniques ( $k$-fold, Leave-One-Out, etc.) are available, see for instance [23]. While more accurate, they are significantly more computationally expensive.

[^6]:    ${ }^{6}$ Note that the plain Smolyak scheme is used here, which does not exploit anisotropy in the response surface. More sophisticated Smolyak-based approximations have been developed, see [7], and are expected to provide better results.

[^7]:    ${ }^{7}$ The spatial $\left\{w_{n}\right\}$ and stochastic modes $\left\{\lambda_{n}\right\}$ are sequentially determined from (36) via an ALS approach. Since the decomposition is two-dimensional, $u(\boldsymbol{x}, \boldsymbol{\xi}) \approx \sum_{n=0}^{N} w_{n}(\boldsymbol{x}) \lambda_{n}(\boldsymbol{\xi})$, the approximation problem is convex, see for instance [46], and the ALS approach converges to the best rank-1 approximation of the matricized $\boldsymbol{u}$ in the Frobenius sense. If the data-driven inner product $\langle\cdot, \cdot\rangle_{N_{q}}$ was inducing a cross norm (it only induces a semi-norm), then $\langle w \boldsymbol{\lambda}, w \boldsymbol{\lambda}\rangle_{N_{q}}=\|\boldsymbol{w}\|_{2}^{2}\|\boldsymbol{\lambda}\|_{2}^{2}$ and the pair ( $\boldsymbol{w}, \boldsymbol{\lambda}$ ) would be the dominant rank-1 approximation of the matricized $\boldsymbol{u}$. The Karhunen-Loève decomposition of $u$ is thus the reference solution one should obtain in the particular case where the empirical inner product induces a cross norm and $N_{q} \rightarrow \infty$.

