Sparse polynomial chaos expansions of vector-valued response quantities

Conference Paper

Author(s):
Blatman, Géraud; Sudret, Bruno

Sparse polynomial chaos expansions of vector-valued response quantities

G. Blatman
Materials and Mechanics of Components Department
EDF R&D, Moret-sur-Loing, France.

B. Sudret
Institute of Structural Engineering
Chair of Risk, Safety and Uncertainty Quantification
ETH Zürich, Switzerland.

ABSTRACT: Sparse polynomial chaos expansions have recently emerged in uncertainty quantification analysis as a tool to solve high dimensional problems, *e.g.* stochastic problems involving a few dozens to a few hundred random variables. Based on penalized regression analysis and the so-called *least angle regression* algorithm the method has proven efficiency in a number of applications. The approach was so far rather limited to *scalar* output quantities, *e.g. quantities of interest* that are post-processed from the solution of a stochastic partial differential equation (SPDE). In this paper we extend this approach to *vector* output quantities in order to obtain the complete solution field. This is carried out by using principal component analysis before computing the PC expansions of the various components. As a whole a complete non intrusive framework is obtained that is only based on a set of deterministic solutions of the underlying deterministic problem. The approach is illustrated by the computation of the displacement field of a tension rod with lognormal, spatially variable Young's modulus. The problem exhibits 62 stochastic dimensions.

1 INTRODUCTION

Polynomial chaos (PC) expansions allow one to represent explicitly the random response of a mechanical system whose input parameters are modelled by random variables. The PC coefficients may be efficiently computed using non intrusive techniques such as least squares (Berveiller, Sudret, and Lemaire 2006). However, the required number of model evaluations (*i.e.* the computational cost) increases with the PC size, which itself dramatically increases with the number of input variables when the common truncation scheme of the PC expansion is applied (*i.e.* retain all the multivariate polynomials of total degree not greater than a prescribed p).

To overcome this problem, an iterative procedure based on Least Angle Regression for building up a *sparse* PC approximation (i.e. a PC representation containing a small number of significant coefficients) was devised in (Blatman and Sudret 2010, Blatman and Sudret 2011). This method allows an automatic enrichment of both the PC basis and the experimental design (i.e. the set of model evaluations to be performed). It has to be noticed though that LAR-PC was mainly devoted to *scalar* model responses, e.g. the

maximal Von Mises stress in a finite element calculation. In case of *vector-valued* model responses, it is necessary to build up separate LAR-PC metamodels for each output variable, which may reveal cumbersome when the model response has a large dimension, e.g. a discretized field such as the Von Mises stresses at all the integration points of all elements.

A two-step strategy is described in this paper in order to decrease the computational effort. First a *principal component analysis* (PCA) of the vector random response is carried out, allowing one to capture the main stochastic features of the response by means of a small number of (non physical) variables compared to the original number of output components. Then the LAR-PC scheme is applied to each non physical variable.

2 SPARSE POLYNOMIAL CHAOS APPROXIMATION

2.1 *Polynomial chaos approximation*

Consider a mechanical system described by a numerical model $\mathcal{M}$ which can be analytical or more generally algorithmic (*e.g.* a finite element model). Sup-

pose that this model has M uncertain input parameters which are represented by *independent* random variables $\{X_1, \ldots, X_M\}$ gathered in a random vector $\mathbf{X}$ with prescribed joint probability density function $f_{\mathbf{X}}(\mathbf{x})$. Hence the model response denoted by $\mathbf{Y} = \mathcal{M}(\mathbf{X})$ is also random. Throughout this paper, a *vector valued* response $\mathbf{Y} = \{Y_1, \ldots, Y_K\}$, $K \geq 2$ is considered.

Provided that the quantity $\mathbb{E}\left[\|\mathbf{Y}\|^2\right]$ is finite (where $\|\cdot\|$ denotes the usual Euclidean norm), random vector $\mathbf{Y}$ may be recast as follows (Soize and Ghanem 2004):

$$\mathbf{Y} = \mathcal{M}(\mathbf{X}) = \sum_{\boldsymbol{\alpha} \in \mathbb{N}^M} \boldsymbol{a}_{\boldsymbol{\alpha}} \psi_{\boldsymbol{\alpha}}(\mathbf{X}) \quad ,$$
(1)

$$\boldsymbol{a}_{\boldsymbol{\alpha}} = \{a_{\boldsymbol{\alpha}}^{(1)}, \ldots, a_{\boldsymbol{\alpha}}^{(K)}\}^{\mathsf{T}}$$

This expansion is referred to as the *polynomial chaos* (PC) *representation* of Y. The $\boldsymbol{a}_{\boldsymbol{\alpha}}$'s are unknown deterministic vectors and the $\psi_{\boldsymbol{\alpha}}$'s are multivariate polynomials which are orthonormal with respect to the joint PDF $f_{\mathbf{X}}$ of $\mathbf{X}$, *i.e.* $\mathbb{E}\left[\psi_{\boldsymbol{\alpha}}(\mathbf{X})\psi_{\boldsymbol{\beta}}(\mathbf{X})\right] = 1$ if $\boldsymbol{\alpha} = \boldsymbol{\beta}$ and 0 otherwise. For instance, if $\mathbf{X}$ is a standard normal random vector, the $\psi_{\boldsymbol{\alpha}}$ are normalized multivariate Hermite polynomials.

2.2 *Truncation of the polynomial chaos series*

For computational purpose, the PC series is truncated after a finite number of terms P. Typically one retains only the polynomials $\psi_{\boldsymbol{\alpha}}$ such that a given q-norm of $\boldsymbol{\alpha}$ is not greater than p, *i.e.* for a given $q \in (0, 1]$:

$$\|\boldsymbol{\alpha}\|_q = \left(\sum_{i=1}^{M} \alpha_i^q\right)^{1/q} \leq p$$
(2)

which yields the following PC approximation:

$$\mathbf{Y} \approx \sum_{\|\boldsymbol{\alpha}\|_q \leq p} \boldsymbol{a}_{\boldsymbol{\alpha}} \psi_{\boldsymbol{\alpha}}(\mathbf{X})$$
(3)

Classically q is set equal to 1, which leads to retain those polynomials with *total degree* not greater than p. In this case the number of terms P is given by:

$$P = \frac{(M + p)!}{M! \, p!}$$
(4)

This formula shows that the PC size P increases rapidly with both M and p. As shown later, the computational cost required to estimate the coefficients increases itself with P. Then the cost may become prohibitive in the presence of high-dimensional problems. In order to moderate the inflation of PC terms, a so-called *hyperbolic* truncation scheme based on parameters q strictly less than 1 was proposed (Blatman and Sudret 2010, Blatman and Sudret 2011).

2.3 *Estimation of the PC coefficients by ordinary least squares*

The PC coefficients in Eq.(3) can be estimated using a least-squares approach (Berveiller, Sudret, and Lemaire 2006). Let us denote their indices by $(\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_{P-1})$ and let us introduce the P-by-K PC coefficient matrix $\boldsymbol{A} = \{\boldsymbol{a}_{\boldsymbol{\alpha}_0}, \ldots, \boldsymbol{a}_{\boldsymbol{\alpha}_{P-1}}\}^{\mathsf{T}}$. Considering a set of realizations of $\mathbf{X}$ denoted by $\mathcal{X} = \{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(N)}\}$ as well as the corresponding model evaluations $\mathcal{Y} = \{\boldsymbol{y}^{(1)} = \mathcal{M}(\boldsymbol{x}^{(1)}), \ldots, \boldsymbol{y}^{(N)} = \mathcal{M}(\boldsymbol{x}^{(N)})\}$, the coefficient estimates are then the solution of the minimization problem:

$$\min_{\boldsymbol{A} \in \mathbb{R}^{P \times K}} \left\{ \sum_{i=1}^{N} \left(\mathcal{M}(\boldsymbol{x}^{(i)}) - \sum_{\|\boldsymbol{\alpha}\|_q \leq p} \boldsymbol{a}_{\boldsymbol{\alpha}} \psi_{\boldsymbol{\alpha}}(\boldsymbol{x}^{(i)}) \right)^2 \right\}$$
(5)

The least squares problem admits a closed form solution:

$$\widehat{\boldsymbol{A}} = \left(\boldsymbol{\Psi}^{\mathsf{T}} \boldsymbol{\Psi}\right)^{-1} \boldsymbol{\Psi}^{\mathsf{T}} \mathcal{Y}$$
(6)

where $\boldsymbol{\Psi}$ is the matrix containing the evaluations of the basis polynomials, that is $\boldsymbol{\Psi} = \{\psi_{\boldsymbol{\alpha}_j}(\boldsymbol{x}^{(i)}), 1 \leq i \leq N, 0 \leq j \leq P - 1\}$.

The sample size N (*i.e.* the number of model evaluations) must be greater than P to make this problem well-posed (in practice the thumb rule $N = 3P$ is often used). In the presence of a time-demanding model, the computational budget N is often limited and the number P of unknown PC coefficients may be close to or even greater than the number of model evaluations N, leading to a numerically ill-conditioned or even ill-posed problem.

2.4 *Estimation of the PC coefficients by Least Angle Regression*

To overcome the abovementioned limitation, it is assumed that the model response can be correctly approximated by a *sparse* PC approximation, *i.e.* a PC that only contains a small number of non zero coefficients. Such an assumption is reasonable as most of physical systems are governed by main effects and low order interactions (*sparsity of effects principle* (Montgomery 2004)).

Accordingly, the PC coefficients are sought as the solution of a particular *regularized least squares* problem, referred to as the *LASSO* (Tibshirani 1996). As the LASSO method is dedicated to *scalar* response quantities, it could be applied successively to every response component. The PC coefficients $A = (a_{\boldsymbol{\alpha}_0}^{(k)}, \ldots, a_{\boldsymbol{\alpha}_{P-1}}^{(k)})$ of some response component

$Y_k, k = 1, \ldots, K$ are obtained by solving:

$$\min_{A \in \mathbb{R}^P} \left\{ \sum_{i=1}^N \left(Y_k(\boldsymbol{x}^{(i)}) - \sum_{\|\boldsymbol{\alpha}\|_q \leq p} a_{\boldsymbol{\alpha}}^{(k)} \psi_{\boldsymbol{\alpha}}(\boldsymbol{x}^{(i)}) \right)^2 \right\} + \lambda \|A\|_1$$

$$(7)$$

where λ is a positive real number that adjusts the trade-off between the goodness-of-fit and the complexity of the solution. The choice of a ℓ_1-penalty on the PC coefficients guarantees a *sparse* solution, the latter being all the sparser as the parameter λ is large. Note that the optimal value of λ is unknown. A brute-force approach would consist in applying the LASSO for many values of λ, and to retain the most accurate PC approximation.

A much more efficient approach consists in determining a whole set of LASSO solutions using a variable selection algorithm known as *least angle regression* (LAR) (Efron, Hastie, Johnstone, and Tibshirani 2004). The LAR algorithm is as follows:

1. Initialize the coefficients to $a_{\boldsymbol{\alpha}_0}, \ldots, a_{\boldsymbol{\alpha}_{P-1}} = 0$. Set the initial residual equal to the vector of observations $\boldsymbol{\mathcal{Y}}_k$ of Y_k.

2. Find the vector $\boldsymbol{\psi}_{\boldsymbol{\alpha}_j}$ which is most correlated with the current residual.

3. Move $a_{\boldsymbol{\alpha}_j}$ from 0 toward the least-square coefficient of the current residual on $\boldsymbol{\psi}_{\boldsymbol{\alpha}_j}$, until some other predictor $\boldsymbol{\psi}_{\boldsymbol{\alpha}_k}$ has as much correlation with the current residual as does $\boldsymbol{\psi}_{\boldsymbol{\alpha}_j}$.

4. Move jointly $\{a_{\boldsymbol{\alpha}_j}, a_{\boldsymbol{\alpha}_k}\}^{\mathsf{T}}$ in the direction defined as their joint least-square coefficient of the current residual on $\{\boldsymbol{\psi}_{\boldsymbol{\alpha}_j}, \boldsymbol{\psi}_{\boldsymbol{\alpha}_k}\}$, until some other predictor $\boldsymbol{\psi}_{\boldsymbol{\alpha}_l}$ has as much correlation with the current residual.

5. Continue this way until $m \equiv \min(P, N-1)$ predictors have been entered.

Steps 3 and 4 mention a "move" of the *active* coefficients toward their least-square value. It corresponds to an updating of the form $\hat{\boldsymbol{a}}^{(k+1)} = \hat{\boldsymbol{a}}^{(k)} + \gamma^{(k)} \tilde{\boldsymbol{w}}^{(k)}$. Vector $\tilde{\boldsymbol{w}}^{(k)}$ and coefficient $\gamma^{(k)}$ are referred to as the LAR *descent direction* and *step*, respectively. Both quantities may be derived algebraically as shown in (Efron, Hastie, Johnstone, and Tibshirani 2004) and (Blatman 2009, Appendix E). Eventually LAR returns a collection of m PC metamodels that respectively contain $1, 2, \ldots, m$ non zero coefficients. The optimal metamodel is determined using a cross-validation technique (Blatman and Sudret 2010).

Note that a limitation of LAR lies in the requirement of given truncated PC basis such as in Eq.(3) and design of experiments (matrices $\mathcal{X}$ and $\mathcal{Y}$). An iterative strategy was deviced in (Blatman and Sudret 2010) in order to automatically select an optimal truncation level (integer p introduced in Section 2.2) as well as the minimal sample size ensuring a target approximation accuracy.

3 REDUCED CHAOS EXPANSIONS BASED ON PRINCIPAL COMPONENT ANALYSIS

In the presence of a vector-valued response **Y**, the direct LAR procedure as described above should be applied *componentwise*. This may reveal cumbersome in case of a large number K of response components (*e.g.* a displacement vector in a finite element model involving millions of DOF). In order to decrease the computational effort, one proposes to perform a *principal component analysis* (Jolliffe 1986) of the vector random response, allowing one to capture the main stochastic features of the response by means of a small number K' of (non physical) variables compared to the original number K of output components.

3.1 *Principal component analysis (PCA)*

Let us define the *sample covariance matrix* of the model response as:

$$\mathcal{S}_{i,j} := \frac{1}{N-1} \sum_{k=1}^N \left(y_i^{(k)} - \bar{y}_i \right) \left(y_j^{(k)} - \bar{y}_j \right) \quad (8)$$

where $i, j = 1, \ldots, K$ and $\bar{\boldsymbol{y}} = \{\bar{y}_1, \ldots, \bar{y}_K\}^{\mathsf{T}}$ denotes the sample mean:

$$\bar{\boldsymbol{y}} := \frac{1}{N} \sum_{i=1}^N \boldsymbol{y}^{(i)} \quad (9)$$

An eigenvalue decomposition of $\mathcal{S}$ is performed:

$$\mathcal{S} = \underbrace{(\boldsymbol{w}_1, \ldots, \boldsymbol{w}_K)}_{= \boldsymbol{W}} \underbrace{\begin{pmatrix} l_1 & & 0 \\ & \ddots & \\ 0 & & l_K \end{pmatrix}}_{= \boldsymbol{L}} \underbrace{\begin{pmatrix} \boldsymbol{w}_1^{\mathsf{T}} \\ \vdots \\ \boldsymbol{w}_K^{\mathsf{T}} \end{pmatrix}}_{= \boldsymbol{W}^{\mathsf{T}}}$$

$$(10)$$

where the l_i's and the $\boldsymbol{w}_i$'s are the eigenvalues and eigenvectors, respectively. We denote by $\bar{\mathcal{Y}}$ the $(N \times K)$-matrix obtained by replicating N times the sample mean vector $\bar{\boldsymbol{y}}$:

$$\bar{\mathcal{Y}} := \underbrace{\{\bar{\boldsymbol{y}}, \ldots, \bar{\boldsymbol{y}}\}}_{N \text{ times}}^{\mathsf{T}} \quad (11)$$

As the $\boldsymbol{w}_i$'s form an orthonormal basis of $\mathbb{R}^K$, the sample matrix $\mathcal{Y}$ can be cast as:

$$\mathcal{Y} = \bar{\mathcal{Y}} + \sum_{i=1}^{K} \boldsymbol{b}_i \, \boldsymbol{w}_i^\mathsf{T} \qquad (12)$$

where the $\boldsymbol{b}_i$'s are vectors of size N given by:

$$\boldsymbol{b}_i = (\mathcal{Y} - \bar{\mathcal{Y}}) \, \boldsymbol{w}_i \qquad (13)$$

The $\boldsymbol{b}_i$'s are empirically centered, have empirical variance equal to l_i and are mutually empirically uncorrelated, that is:

$$\frac{1}{N} \sum_{k=1}^{N} b_{i,k} = 0 \quad , \quad \frac{1}{N-1} \sum_{k=1}^{N} b_{i,k} \, b_{j,k} = l_i \, \delta_{i,j} \qquad (14)$$

for $i, j = 1, \ldots, K$. It is possible to approximate the sample matrix $\mathcal{Y}$ by the K'-term truncation:

$$\widehat{\mathcal{Y}}^{(K')} = \bar{\mathcal{Y}} + \sum_{i=1}^{K'} \boldsymbol{b}_i \, \boldsymbol{w}_i^\mathsf{T} \qquad (15)$$

where $K' < K$. The previous expression is referred to as the *principal component approximation* of the sample matrix $\mathcal{Y}$. The number K' of retained principal components can be selected with respect to a prescribed approximation error ε_{PCA}:

$$K' = \min \left\{ Q \in \{1, \ldots, K\} : \right.$$
$$\left. \frac{\sum_{i=1}^{Q} l_i}{\operatorname{trace}(\mathcal{S})} \geq 1 - \varepsilon_{PCA} \right\} \qquad (16)$$

Eq.(15) suggests the following sample-based principal component decomposition of $\mathbf{Y}$:

$$\widehat{\mathbf{Y}}^{(K')} = \bar{\mathcal{Y}} + \sum_{i=1}^{K'} \mathsf{B}_i \, \boldsymbol{w}_i \qquad (17)$$

In the previous formula, each B_i is the projection of the actual response $\mathbf{Y}$ onto the vector $\boldsymbol{w}_i$, that is:

$$\mathsf{B}_i = \boldsymbol{w}_i^\mathsf{T} (\mathbf{Y} - \mathbb{E}[\mathbf{Y}]) \qquad (18)$$

However, as $\mathbf{Y}$ is unknown, so are the random variables B_i. Therefore the latter have to be replaced with some estimates $\widehat{\mathsf{B}}_i$ based on the realizations $\boldsymbol{b}_i$. Hence the following approximation:

$$\widehat{\widehat{\mathbf{Y}}}^{(K')} = \bar{\mathcal{Y}} + \sum_{i=1}^{K'} \widehat{\mathsf{B}}_i \, \boldsymbol{w}_i \qquad (19)$$

Note that the double hat notation reflects the double level of approximation, namely the sample-based PCA approximation and the estimation of the B_i's. In the sequel the estimators $\widehat{\mathsf{B}}_i$ are sought under the form of truncated sparse polynomial chaos expansions of the unknown B_i's.

3.2 *PCA-based vector-valued polynomial chaos expansions*

The actual random variables B_i arising from Eq.(18) are regarded as functions of the input random vector $\mathbf{X}$, and more specifically *linear combinations* of the latter:

$$\mathsf{B}_i = \boldsymbol{w}_i^\mathsf{T} (\mathcal{M}(\mathbf{X}) - \mathbb{E}[\mathcal{M}(\mathbf{X})]) \equiv \varphi_i(\mathbf{X}) \qquad (20)$$

with $i = 1, \ldots, K'$ and these functional relationships are approximated using the LAR approach summarized in Section 2.4.

The following calculations are performed in the algorithm iterations:

1. Estimate the sample covariance matrix $\mathcal{S}$ of the response vector (Eq.(8)).

2. Compute its spectral decomposition (Eq.(10)).

3. Retain the appropriate number K' of eigenvalues according to the criterion (16) for a prescribed error ε_{PCA}.

4. Compute the vectors $\{\boldsymbol{b}_1, \ldots, \boldsymbol{b}_{K'}\}$ (Eq.(13)).

5. For $i = 1, \ldots, K'$, apply LAR to the input-output sample set $\{\mathcal{X}, \boldsymbol{b}_i\}$ in order to construct a sparse PC approximation of the random variable B_i. Eventually a PC approximation $\widehat{\widehat{\mathbf{Y}}}^{(K')}$ is obtained.

The algorithm is stopped if some estimate $\widehat{\varepsilon}$ of the approximation error is less than a target error ε_{tgt}. Otherwise, the experimental design is enriched by means of additional model evaluations and the previous steps are conducted again.

4 ERROR ESTIMATION

Let us recall the expression of the L_2-approximation error:

$$\Sigma = \left\| \mathbf{Y} - \widehat{\widehat{\mathbf{Y}}}^{(K')} \right\|^2 = \mathbb{E}\left[\left\| \mathbf{Y} - \widehat{\widehat{\mathbf{Y}}}^{(K')} \right\|_2^2 \right] \qquad (21)$$

This error can be bounded by a sum of three different contributions due to the triangular inequality:

$$\Sigma \leq \left(\sqrt{\Sigma_I} + \sqrt{\Sigma_{II}} + \sqrt{\Sigma_{III}} \right)^2 \qquad (22)$$

where:

$$\begin{cases} \Sigma_I & := & \left\| \mathbf{Y} - \mathbf{Y}^{(K')} \right\|^2 \\[2mm] \Sigma_{II} & := & \left\| \mathbf{Y}^{(K')} - \widehat{\mathbf{Y}}^{(K')} \right\|^2 \\[2mm] \Sigma_{III} & := & \left\| \widehat{\mathbf{Y}}^{(K')} - \widehat{\widehat{\mathbf{Y}}}^{(K')} \right\|^2 \end{cases} \qquad (23)$$

The term Σ_I is the error due to the truncation of the theoretical principal component decomposition. It is recalled that this error is equal to the sum of the non retained eigenvalues of the actual covariance matrix of $\mathbf{Y}$:

$$\Sigma_I = \sum_{k=K'+1}^{K} \lambda_k \qquad (24)$$

The term Σ_{II} is the error induced by the substitution of the actual covariance matrix by its sample counterpart. Thus it accounts for the approximation of the actual eigenvalues and eigenvectors by the empirical ones. Lastly, the term Σ_{III} corresponds to the error of approximation of the random variables B_i by sparse PC expansions $\widehat{\mathsf{B}}_i$.

In this work, the PCA error Σ_I is estimated by the following cumulated sum of sample-based eigenvalues:

$$\widehat{\Sigma}_I = \sum_{k=K'+1}^{K} l_k \qquad (25)$$

Actually such a choice leads to neglect the term Σ_{II}. This is all the more relevant since the size N of the experimental design increases as the observed eigenvalues converge to the theoretical ones. Note that Σ_{II} might be estimated though by resorting to a *bootstrap* procedure (Beran and Srivastava 1985). On the other hand, the term Σ_{III} can be recast in terms of random variables B_i and $\widehat{\mathsf{B}}_i$ by exploiting Eqs.(17),(19) and the orthonormality of the eigenvectors $\boldsymbol{w}_i$:

$$\Sigma_{III} = \sum_{i=1}^{K'} \mathbb{E}\left[\left(\mathsf{B}_i - \widehat{\mathsf{B}}_i \right)^2 \right] \qquad (26)$$

Thus Σ_{III} can be estimated from the *cross-validation* error estimates $\widehat{\varepsilon}_{LOO}^{(i)}$ related to the various *scalar-valued* PC metamodels $\widehat{\mathsf{B}}_i$:

$$\widehat{\Sigma}_{III} = \sum_{i=1}^{K'} \widehat{\varepsilon}_{LOO}^{(k)} \qquad (27)$$

Eventually the upper bound of the error is estimated by:

$$\widehat{\Sigma}^+ \equiv \left(\sqrt{\sum_{k=K'+1}^{K} l_k} + \sqrt{\sum_{k=1}^{K'} \widehat{\varepsilon}_{LOO}^{(k)}} \right)^2 \qquad (28)$$

In practice it is often convenient to evaluate a *relative* error measure. In this respect the sources of error can be divided by an estimate of $\|\mathbf{Y}\|^2$, such as the sum of all the eigenvalues l_i (*i.e.* the trace) of the sample covariance matrix $\mathcal{S}$. This yields the following estimate of the scaled error bound:

$$\varepsilon \equiv \frac{\widehat{\Sigma}^+}{\text{trace}(\mathcal{S})} = \left(\sqrt{\varepsilon_1} + \sqrt{\varepsilon_2} \right)^2 \qquad (29)$$

where:

$$\varepsilon_1 := \frac{\sum_{k=K'+1}^{K} l_k}{\text{trace}(\mathcal{S})} \quad , \quad \varepsilon_2 := \frac{\sum_{k=1}^{K'} \widehat{\varepsilon}_{LOO}^{(k)}}{\text{trace}(\mathcal{S})} \qquad (30)$$

It has been shown that the PCA-LAR algorithm is stopped as soon as the error estimate ε is less than a prescribed target error ε_{tgt}. According to Eq.(29), it is then relevant to respectively define target values $\varepsilon_{1,tgt}$ and $\varepsilon_{2,tgt}$ for parameters ε_1 and ε_2 in such a way that:

$$\varepsilon_{tgt} \equiv \left(\sqrt{\varepsilon_{1,tgt}} + \sqrt{\varepsilon_{2,tgt}} \right)^2 \qquad (31)$$

Note that setting $\varepsilon_{1,tgt}$ equal to zero leads to a *pure LAR strategy*, *i.e.* to apply the LAR algorithm to every component of the vector-valued model response $\mathbf{Y}$ without having reduced the number of outputs by means of PCA. In contrast, a too high value of $\varepsilon_{1,tgt}$ would have to be counterbalanced by a significant computational cost to construct a very accurate PC approximation, *i.e.* to reach a very low value of ε_2.

In practice the assigned value of $\varepsilon_{1,tgt}$ determines the number K' of retained eigenvalues in PCA according to Eq.(16) (parameter ε_{PCA} is then equal to $\varepsilon_{1,tgt}$). Then $\varepsilon_{2,tgt}$ determines the accuracy of the PC approximations of the random variables B_i.

5 APPLICATION TO A STOCHASTIC DIFFUSION PROBLEM

Let us consider a one-dimensional stochastic diffusion problem defined by the following boundary-value problem defined over a domain $\mathcal{D} = [0, L]$ (see Ernst, Mugler, Starkloff, and Ullmann (2012)):

$$\begin{array}{rcl} [E(x,\omega)\, u'(x,\omega)]' + f(x) & = & 0 \\ u(0) & = & 0 \\ (E\, u')(L) & = & F \end{array} \qquad (32)$$

This PDE describes the displacement field $u(x,\omega)$ of a tension rod having unit cross-section $S = 1$ and spatially varying Young's modulus $E(x,\omega)$ which is clamped at $x = 0$ and is submitted to a uniform axial load $f(x)$ and pinpoint load F at its right-end $x = L$.

In Eq.(32) the diffusion coefficient $E(x,\omega)$ is a lognormal stationary random field obtained by the exponentiation of a Gaussian random field as follows:

$$E(x,\omega) = \exp\left[\lambda_E + \zeta_E\, g(x,\omega) \right] \qquad (33)$$

where $g(x, \omega)$ is a standard normal stationary Gaussian random field with an exponential autocorrelation function:

$$\text{Cov}\,[g(x)\,g(x')] = e^{-|x'-x|/\ell} \tag{34}$$

We select the following numerical values for the numerical experimentation: $L = 1$, $F = 1$, $f = 0.5$ for the boundary-value problem. The parameters λ_E, ζ_E are chosen so that the mean value and standard deviation of $E(x, \omega)$ are 10 and 3 respectively and $\ell = 1/3$.

Random field g is represented by its truncated Karhunen-Loève expansion:

$$g(x, \omega) = \sum_{k=1}^{M} \sqrt{\mathfrak{l}_k}\, \varphi_k(x)\xi_k(\omega) \tag{35}$$

where $\{(\mathfrak{l}_k\,\varphi_k), k = 1, \ldots, M\}$ are the solution of the Fredholm equation associated to the exponential kernel, see Ghanem and Spanos (2003, Sudret and Der Kiureghian (2000) for the detailed analytical solution. The 5 first eingenmodes $\varphi_k(x)$ are plotted in Figure 1. We select $M = 62$ so as to have a mean-
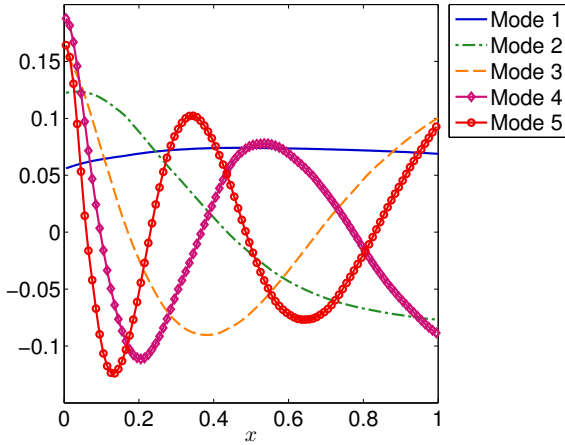

Figure 1: Five first Karhunen-Loève eingenmodes $\varphi_k(x)$ of the diffusion coefficient $u(x, \omega)$

square error less than 1%:

$$\sum_{k=1}^{M=62} \mathfrak{l}_k / \sum_{k=1}^{\infty} \mathfrak{l}_k \geq 0.99 \tag{36}$$

For a given realization of the random field denoted by $e_0(x) \stackrel{\text{def}}{=} E(x, \omega_0)$ and assuming a constant uniform load $f(x) \equiv f$, Eq.(32) can be solved as follows:

$$u(x, \omega_0) = \int_0^x \frac{F + f(L - y)}{e_0(y)} dy \tag{37}$$

Substituting for Eq.(33), Eq.(35) in Eq.(37) one finally gets the *analytical* solution related to a truncated KL expansion of order M:

$$u_M(x, \omega_0) = \int_0^x \frac{F + f(L - y)}{\exp\left[\lambda_E + \zeta_E\,\Phi(x)^{\mathsf{T}} \cdot \Xi_0\right]} dy \tag{38}$$

where $\Xi_0 \stackrel{\text{def}}{=} \{\xi_1(\omega_0), \ldots, \xi_M(\omega_0)\}^{\mathsf{T}}$ is a realization of a standard random vector of size M and $\Phi(x) \stackrel{\text{def}}{=} \left\{\sqrt{\mathfrak{l}_1}\,\varphi_1(x), \ldots, \sqrt{\mathfrak{l}_M}\,\varphi_M(x)\right\}^{\mathsf{T}}$. For each realization Ξ_0 and due to the analytical expression of the eigenfunctions gathered in $\Phi(x)$ the integral in Eq.(38) can be evaluated with arbitrary precision. In the sequel, we use it as a reference using Matlab's built-in integration routine `quad` and a tolerance of 10^{-10}. As an illustration, Figure 2 represents 5 realizations of $E(x, \omega)$ and the associated solution obtained by Eq.(38).

(a) Coefficient of diffusion
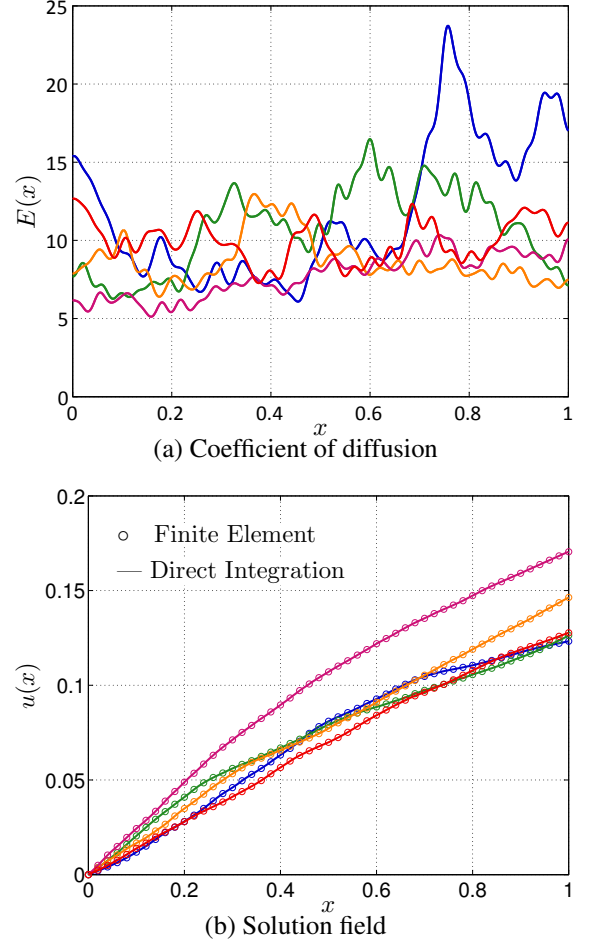

(b) Solution field

Figure 2: Five realizations of the coefficient of diffusion $E(x, \omega)$ and associated solution field $u(x, \omega)$

In the sequel the various results are compared to a large Monte Carlo simulation ($N_{MCS} = 50,000$) of the analytical solution in Eq.(38). The PC-based numerical solution is carried out using a finite element model with 200 linear elements over $[0, 1]$, which appeared sufficient to get accurate deterministic solutions. The experimental design is of fixed size ($N = 500$) and it is drawn by Latin Hypercube Sampling. The mean value and standard deviation of the solution field $u(x, \omega)$ are plotted in Figure 3 using $K' = 1, 2, 3$ principal components (out of $K = 201$) respectively.

It can be observed that the mean solution is actually the mere empirical mean of the trajectories obtained from the experimental design. As far as stan-
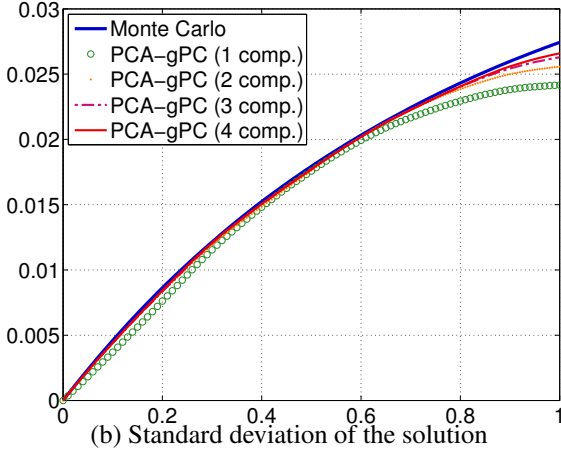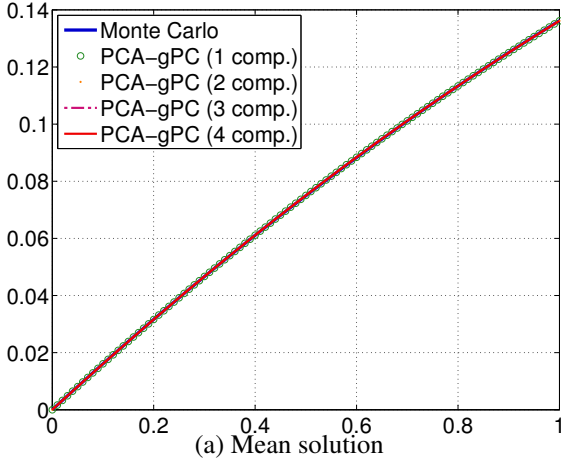
(a) Mean solution

(b) Standard deviation of the solution

Figure 3: Mean and standard deviation of the solution field $u(x, \omega)$ (finite element analysis using 200 linear elements)

dard deviation is concerned a rather accurate result is obtained for $x \in [0 ; 0.8L]$ as soon as two components are used. The convergent over the full domain is observed when increasing the number of principal components K'.

In order to better understand the convergence of the PCA-gPC procedure for this example, the error estimates proposed in Eq.(30) are computed as a function of the number of principal components $K' = 1, \ldots, 10$. A cross-validation technique is also used based on the data base built from the Monte Carlo simulation: for each realizationg ω_0 of the 62 input variables used in the KL expansion, the mean-square error between the solution field that is computed by Eq.(38) and the finite element solution ($K = 201$) is computed. This error is then averaged over the 50,000 available realizations. This so-called *cross-validation* error is plotted together with ε, ε_1, ε_2 in Figure 4.

First it is observed that the PCA error ε_1 decreases rather fast with the number of components. This is due to the fact that the components of a solution field that is discretized by finite elements are rather correlated, so that a few number of principal components already captures much of the spectral content.

Using a fixed experimental design ($N = 500$) it is observed that the accuracy of the PC expansion of each B_k decreases when k increases. When using more and more components the metamodel er-
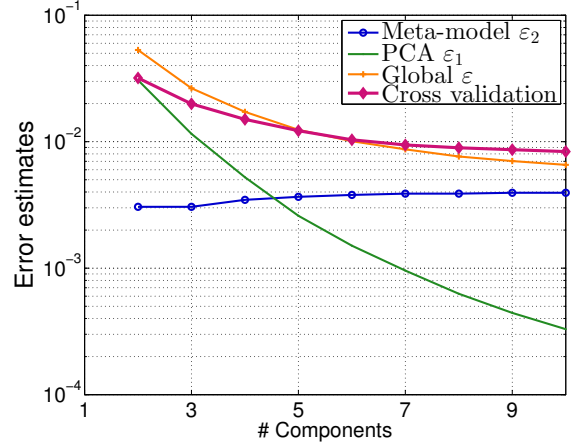

Figure 4: Evolution of the various error estimates with the number of components $K' = 1, \ldots, 10$

ror ε_2 thus increases. The global error Eq.(29) which combines ε_1 and ε_2 first decreases and then attains a plateau. Note that the proposed error estimate is indeed an upper bound of the "true" cross-validation error as long as the accuracy of each single PC expansion is not too bad, *i.e.* for $K' \leq 5$ in the present example. In order to have a better global control on the error it is necessary to augment the size of the experimental design dynamically so as to capture the details of the higher order principal components.

6 CONCLUSION

Polynomial chaos expansions are more and more used in the field of uncertainty quantification in engineering. In order to tackle large-dimensional problems such as those appearing when spatial variability is accounted for by using input random fields, it is necessary to resort to so-called sparse PC expansions.

The present paper extends the approach developed by Blatman and Sudret based on least angle regression to the case of multidimensional-output computational models. The procedure is based on a principal component analysis of the experimental design matrix. By keeping a few components (typically less than 10 for a finite element problem) the problem is reduced to computing a series of sparse PC expansions of scalar variables which appear as "dual" to the eigenvectors of the experimental matrix.

Due to the well-posedness of the approximation problem in the $\mathcal{L}^2$ sense, errors made either by truncating the PCA analysis or in the various truncated PC expansions can be bounded. The illustration on a 1D stochastic diffusion problem involving 62 input random variables has shown that the PCA convergence is extremely fast and the derived error estimate is an upper bound to the real mean-square error provided each component is properly approximated. More generally, the derivation of a strict upper bound remains a challenging problem and will deserve further numerical and theoretical investigations.

REFERENCES

Beran, R. and M. Srivastava (1985). Bootstrap tests and confidence regions for functions of a covariance matrix. *Ann. Stat. 13*(1), 96–115.

Berveiller, M., B. Sudret, and M. Lemaire (2006). Stochastic finite elements: a non intrusive approach by regression. *Eur. J. Comput. Mech. 15*(1-3), 81–92.

Blatman, G. (2009). *Adaptive sparse polynomial chaos expansions for uncertainty propagation and sensitivity analysis*. Ph. D. thesis, Université Blaise Pascal, Clermont-Ferrand.

Blatman, G. and B. Sudret (2010). Efficient computation of Sobol' sensitivity indices using sparse polynomial chaos expansions. *Reliab. Eng. Sys. Safety 95*(11), 1216–1229.

Blatman, G. and B. Sudret (2011). Adaptive sparse polynomial chaos expansion based on Least Angle Regression. *J. Comp. Phys. 230*(6), 2345–2367.

Efron, B., T. Hastie, I. Johnstone, and R. Tibshirani (2004). Least angle regression. *Annals of Statistics 32*, 407–499.

Ernst, O., A. Mugler, H.-J. Starkloff, and E. Ullmann (2012). On the convergence of generalized polynomial chaos expansions. *ESAIM: Mathematical Modelling and Numerical Analysis 46*(02), 317–339.

Ghanem, R. and P. Spanos (2003). *Stochastic Finite Elements : A Spectral Approach*. Courier Dover Publications.

Jolliffe, I. (1986). *Principal component analysis*. Springer Verlag. (Second Edition, 2002).

Montgomery, D. (2004). *Design and analysis of experiments*. John Wiley and Sons, New York.

Soize, C. and R. Ghanem (2004). Physical systems with random uncertainties: chaos representations with arbitrary probability measure. *SIAM J. Sci. Comput. 26*(2), 395–410.

Sudret, B. and A. Der Kiureghian (2000). Stochastic finite elements and reliability: a state-of-the-art report. Technical Report UCB/SEMM-2000/08, University of California, Berkeley. 173 pages.

Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *J. Royal Stat. Soc., Series B 58*, 267–288.