



Polynomial meta-models with canonical low-rank approximations: numerical insights and comparison to sparse polynomial chaos expansions

Katerina Konakli, Bruno Sudret

► To cite this version:

Katerina Konakli, Bruno Sudret. Polynomial meta-models with canonical low-rank approximations: numerical insights and comparison to sparse polynomial chaos expansions . Journal of Computational Physics, 2016. <hal-01432141>

HAL Id: hal-01432141

<https://hal.science/hal-01432141v1>

Submitted on 11 Jan 2017

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

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



HAL Authorization

POLYNOMIAL META-MODELS WITH CANONICAL LOW-RANK APPROXIMATIONS: NUMERICAL INSIGHTS AND COMPARISON TO SPARSE POLYNOMIAL CHAOS EXPANSIONS

K. Konakli, B. Sudret



Data Sheet

Journal:	-
Report Ref.:	RSUQ-2015-007
Arxiv Ref.:	http://arxiv.org/abs/1511.07492 [math.NA]
DOI:	-
Date submitted:	15/04/2016
Date accepted:	-

Polynomial meta-models with canonical low-rank approximations: numerical insights and comparison to sparse polynomial chaos expansions

K. Konakli¹ and B. Sudret¹

¹*Chair of Risk, Safety and Uncertainty Quantification, Department of Civil Engineering, ETH Zurich, Stefano-Franscini-Platz 5, 8093 Zurich, Switzerland*

Abstract

The growing need for uncertainty analysis of complex computational models has led to an expanding use of meta-models across engineering and sciences. The efficiency of meta-modeling techniques relies on their ability to provide statistically-equivalent analytical representations based on relatively few evaluations of the original model. Polynomial chaos expansions (PCE) have proven a powerful tool for developing meta-models in a wide range of applications; the key idea thereof is to expand the model response onto a basis made of multivariate polynomials obtained as tensor products of appropriate univariate polynomials. The classical PCE approach nevertheless faces the “curse of dimensionality”, namely the exponential increase of the basis size with increasing input dimension. To address this limitation, the sparse PCE technique has been proposed, in which the expansion is carried out on only a few relevant basis terms that are automatically selected by a suitable algorithm. An alternative for developing meta-models with polynomial functions in high-dimensional problems is offered by the newly emerged low-rank approximations (LRA) approach. By exploiting the tensor-product structure of the multivariate basis, LRA can provide polynomial representations in highly compressed formats. Through extensive numerical investigations, we herein first shed light on issues relating to the construction of canonical LRA with a particular greedy algorithm involving a sequential updating of the polynomial coefficients along separate dimensions. Specifically, we examine the selection of optimal rank, stopping criteria in the updating of the polynomial coefficients and error estimation. In the sequel, we confront canonical LRA to sparse PCE in structural-mechanics and heat-conduction applications based on finite-element solutions. Canonical LRA exhibit smaller errors than sparse PCE in cases when the number of available model evaluations is small with respect to the input dimension, a situation that is often encountered in real-life problems. By introducing the conditional generalization error, we further demonstrate that canonical LRA tend to outperform sparse PCE in the prediction of extreme model responses, which is critical in reliability analysis.

Keywords: uncertainty quantification – meta-modeling – sparse polynomial chaos expansions – canonical low-rank approximations – rank selection – meta-model error

1 Introduction

It is nowadays common practice to study the behavior of physical and engineering systems through computer simulation. Proper analysis of the system response must account for the prevailing uncertainties in the system model and the underlying phenomena, which requires repeated simulations under varying scenarios for the input parameters. Modern advances in computer science combined with the improved understanding of physical laws are leading to computational models of increasing complexity. Uncertainty propagation through such models may become intractable in cases when a single simulation is computationally demanding. A remedy is to substitute a complex model with a *meta-model* that possesses similar statistical properties, but has a simple functional form.

The focus of the present work is on meta-models that are built with polynomial functions due to the simplicity and versatility they offer. A popular class of meta-models thereof are the so-called polynomial chaos expansions (PCE) (Xiu and Karniadakis, 2002; Ghanem and Spanos, 2003). The key idea of PCE is to expand the model response onto an appropriate basis made of orthonormal multivariate polynomials, the latter obtained as tensor products of univariate polynomials in each of the input parameters. In non-intrusive approaches that are of interest herein, the coefficients of the expansion are evaluated in terms of the response of the original model at a set of points in the input space, called the experimental design (Choi et al., 2004; Berveiller et al., 2006; Xiu, 2009). Although PCE have proven powerful in a wide range of applications, they face limitations in cases with high-dimensional input. This is because the number of the basis terms, and thus of the unknown expansion coefficients, grows exponentially with the number of input parameters, which is commonly referred to as the “curse of dimensionality”. As shown in Blatman and Sudret (2011); Doostan and Owhadi (2011), the efficiency of the PCE approach can be significantly improved by using a sparse basis.

A promising alternative for developing meta-models with polynomial functions in high-dimensional spaces is provided by canonical decompositions. In canonical decompositions, also known as separated representations, a tensor is expressed as a sum of rank-one components. This type of representation constitutes a special case of tensor decompositions, which are typically used to compress information or extract a few relevant modes of a tensor; a survey on different types of tensor decompositions can be found in Kolda and Bader (2009). The original idea of canonical decomposition dates back to 1927 (Hitchcock, 1927), but became popular in the second half of the 20th century after its introduction to psychometrics (Carroll and Chang, 1970; Harshman, 1970). Since then, it has been used in a broad range of fields, including chemometrics (Appellof and Davidson, 1981; Bro, 1997), neuroscience (Mocks, 1988; Andersen and Rayens, 2004), fluid mechanics (Felippa and Ohayon, 1990; Ammar et al., 2006), signal processing (Sidiropoulos et al., 2000; De Lathauwer and Castaing, 2007), image analysis (Shashua and Levin, 2001; Furukawa et al., 2002) and data mining Acar et al. (2006); Beylkin et al. (2009). More recently, canonical decompositions are attracting an increasing interest in the field of uncertainty quantification (Nouy, 2010; Chevreuil et al., 2013; Doostan et al., 2013; Hadigol et al., 2014; Mathelin, 2014; Rai, 2014; Validi, 2014; Chevreuil et al., 2015).

By exploiting the tensor-product structure of the multivariate polynomial basis, canonical decompositions can provide equivalent to PCE representations in highly-compressed formats. It is emphasized that the number of parameters in canonical decompositions grows only linearly with the input dimension, which, in cases of high-dimensional problems, results

in a drastic reduction of the number of unknowns compared to PCE. Naturally, canonical decompositions with a few rank-one components are of interest, thus leading to the name low-rank approximations (LRA). We note that although the present study is constrained to the use of polynomial bases, different basis functions may be considered for the construction of LRA in a general case (see, *e.g.* Chevreuil et al. (2015)).

Recently proposed methods for building canonical LRA meta-models in a non-intrusive manner rely on the sequential updating of the polynomial coefficients along separate dimensions. The underlying algorithms require solving a series of minimization problems of small size, independent of the input dimension, which can be easily handled using standard techniques. However, the LRA construction involves open questions that call for further investigations. In particular, stopping criteria in the sequential updating of the polynomial coefficients as well as criteria for selection of the optimal rank and polynomial degree are not yet well established. Considering a particular greedy algorithm for building canonical LRA meta-models, we herein shed light on the aforementioned issues through extensive numerical investigations. In the sequel, we assess the comparative accuracy of canonical LRA and sparse PCE in applications involving finite-element models pertinent to structural mechanics and heat conduction. In these applications, sparse PCE are built with a state-of-art method, where a candidate basis is defined by means of a hyperbolic truncation scheme and the final sparse basis is determined using least angle regression. Comparisons between the meta-model errors are carried out for experimental designs of varying sizes drawn with Sobol sequences and Latin hypercube sampling.

The organization of the paper is as follows: In Section 2, we present the mathematical setup of non-intrusive meta-modeling and describe corresponding error measures. Sections 3 and 4 respectively describe the sparse PCE and canonical LRA approaches. After investigating open questions in the construction of LRA in Section 5, we confront the two types of polynomial meta-models in Section 6. The paper concludes with a summary of the main findings and respective outlooks in Section 7.

2 Non-intrusive meta-modeling

2.1 Mathematical setup

We consider a physical or engineering system whose behavior is represented by a computational model \mathcal{M} . Let $\mathbf{X} = \{X_1, \dots, X_M\}$ and $\mathbf{Y} = \{Y_1, \dots, Y_N\}$ respectively denote the M -dimensional input vector and the N -dimensional response vector of the model. In order to account for the uncertainty in the input and the resulting uncertainty in the response, the elements of \mathbf{X} and \mathbf{Y} are described by random variables. For the sake of simplicity, we hereafter restrain our analysis to the case of a scalar model response, *i.e.* $N = 1$. Note that the case of a vector model response can be addressed by separately treating each element of \mathbf{Y} as in the case of a scalar response. The above are summarized in the mathematical formalism:

$$\mathbf{X} \in \mathcal{D}_{\mathbf{X}} \subset \mathbb{R}^M \mapsto Y = \mathcal{M}(\mathbf{X}) \in \mathbb{R}, \quad (1)$$

where $\mathcal{D}_{\mathbf{X}}$ denotes the support of \mathbf{X} .

In a typical real-life application, the model \mathcal{M} is not known in a closed analytical form and may represent a complex computational process. A meta-model $\widehat{\mathcal{M}}$ is an analytical function such that $\widehat{Y} = \widehat{\mathcal{M}}(\mathbf{X})$ possesses similar statistical properties with Y . Non-intrusive methods for building meta-models rely on a series of calls to the original model \mathcal{M} , which may be

used as a “black-box” without any modification. Building a meta-model in a non-intrusive manner requires an experimental design (ED), *i.e.* a set of realizations of the input vector $\mathcal{E} = \{\boldsymbol{\chi}^{(1)}, \dots, \boldsymbol{\chi}^{(N)}\}$, and the corresponding model evaluations $\mathcal{Y} = \{\mathcal{M}(\boldsymbol{\chi}^{(1)}), \dots, \mathcal{M}(\boldsymbol{\chi}^{(N)})\}$.

2.2 Error measures

For a set of realizations of the input vector $\mathcal{X} = \{\boldsymbol{x}_1, \dots, \boldsymbol{x}_n\} \subset \mathcal{D}_{\mathbf{X}}$ and two real-valued functions a and b with common domain $\mathcal{D}_{\mathbf{X}}$, we define the semi-inner product:

$$\langle a, b \rangle_{\mathcal{X}} = \frac{1}{n} \sum_{i=1}^n a(\boldsymbol{x}_i) b(\boldsymbol{x}_i). \quad (2)$$

Eq. (2) leads to the semi-norm $\|a\|_{\mathcal{X}} = \sqrt{\langle a, a \rangle_{\mathcal{X}}}$, which is employed in the sequel to describe meta-model error measures.

A good measure of the accuracy of a meta-model is the generalization error Err_G , which represents the mean-square of the residual $\varepsilon = Y - \hat{Y}$:

$$Err_G = \mathbb{E}[\varepsilon^2] = \mathbb{E}\left[\left(Y - \hat{Y}\right)^2\right]. \quad (3)$$

In most practical situations, it is not possible to evaluate the generalization error analytically. An estimator \widehat{Err}_G of this error may be computed via Monte Carlo Simulation (MCS) by using the exact-model and meta-model responses at a sufficiently large set of points in the input space $\mathcal{X}_{\text{val}} = \{\boldsymbol{x}_1, \dots, \boldsymbol{x}_{n_{\text{val}}}\}$, called *validation set*:

$$\widehat{Err}_G = \left\| \mathcal{M} - \widehat{\mathcal{M}} \right\|_{\mathcal{X}_{\text{val}}}^2. \quad (4)$$

The corresponding estimator of the relative generalization error, denoted by \widehat{err}_G , is obtained by normalizing \widehat{Err}_G with the empirical variance of $\mathcal{Y}_{\text{val}} = \{Y(\boldsymbol{x}_1), \dots, Y(\boldsymbol{x}_{n_{\text{val}}})\}$, the latter representing the set of model responses at the validation set.

Unfortunately, a validation set is not available in typical meta-modeling applications, where a large number of evaluations of the original model is non-affordable. An alternative estimator is the empirical error \widehat{Err}_E :

$$\widehat{Err}_E = \left\| \mathcal{M} - \widehat{\mathcal{M}} \right\|_{\mathcal{E}}^2, \quad (5)$$

where the subscript \mathcal{E} indicates that the semi-norm is evaluated at the points of the ED. The corresponding relative error, denoted by \widehat{err}_E , is obtained by normalizing \widehat{Err}_E with the empirical variance of $\mathcal{Y} = \{\mathcal{M}(\boldsymbol{\chi}^{(1)}), \dots, \mathcal{M}(\boldsymbol{\chi}^{(N)})\}$, the latter representing the set of model responses at the ED. The empirical error does not require any additional evaluations of the exact model than those already used to build the meta-model. It therefore serves the goal of limiting the number of runs of an expensive computational model to the smallest possible. However, it tends to underestimate the actual generalization error, which might be severe in cases of overfitting.

Relying on the ED only, a fair approximation of the generalization error can be obtained by means of cross-validation (CV) techniques. In k -fold CV, the ED is randomly partitioned into k sets of approximately equal size. A meta-model is built considering all but one of the partitions, which comprises the training set, while the excluded or testing set, is used to evaluate the generalization error. By alternating through the k sets, k meta-models are obtained in this way; their average generalization error provides an estimate of the error of the meta-model built with the full ED.

3 Polynomial Chaos Expansions

3.1 Spectral representation

Assuming that Y in Eq. (1) has a finite variance, the following representation is possible (Soize and Ghanem, 2004):

$$Y = \sum_{j=0}^{\infty} y_j \phi_j, \quad (6)$$

in which $\{\phi_j, j \in \mathbb{N}\}$ is a set of random variables that forms a Hilbertian basis in the space of second-order variables and $\{y_j, j \in \mathbb{N}\}$ are the coordinates of Y in this basis. Eq. (6) constitutes a *spectral representation* of the random response Y .

Let us consider the Hilbert space \mathcal{H} of square-integrable real-valued functions of \mathbf{X} equipped with the inner product:

$$\langle u, v \rangle_{\mathcal{H}} = \int_{\mathcal{D}_{\mathbf{X}}} u(\mathbf{x})v(\mathbf{x})f_{\mathbf{X}}(\mathbf{x})d\mathbf{x}, \quad (7)$$

where $f_{\mathbf{X}}$ denotes the joint probability density function (PDF) of \mathbf{X} . Let us also consider the Hilbert space \mathcal{H}_i of square-integrable real-valued functions of X_i equipped with the inner product:

$$\langle u, v \rangle_{\mathcal{H}_i} = \int_{\mathcal{D}_{X_i}} u(x)v(x)f_{X_i}(x)dx_i, \quad (8)$$

where f_{X_i} denotes the marginal PDF of X_i defined over the support \mathcal{D}_{X_i} .

Under the assumption that the components of \mathbf{X} are independent, it is straightforward to show that the Hilbert space defined by the tensor product $\bar{\mathcal{H}} = \otimes_{i=1}^M \mathcal{H}_i$ and equipped with the inner product:

$$\langle u, v \rangle_{\bar{\mathcal{H}}} = \int_{\mathcal{D}_{\mathbf{X}}} u(x_1, \dots, x_M)v(x_1, \dots, x_M)f_{X_1}(x_1) \dots f_{X_M}(x_M)dx_1 \dots dx_M \quad (9)$$

is isomorphic with \mathcal{H} . Accordingly, if $\{\psi_{\alpha_i}^{(i)}, \alpha_i \in \mathbb{N}\}$ is a basis in the Hilbert space \mathcal{H}_i , then a basis in the Hilbert space \mathcal{H} is defined by the set $\{\Psi_{\alpha}, \alpha = (\alpha_1, \dots, \alpha_M) \in \mathbb{N}^M\}$ with:

$$\Psi_{\alpha}(\mathbf{X}) = \prod_{i=1}^M \psi_{\alpha_i}^{(i)}(X_i). \quad (10)$$

It follows that the problem of specifying a Hilbertian basis in \mathcal{H} reduces to specifying Hilbertian bases in $\mathcal{H}_i, i = 1, \dots, M$. This proposition will be employed in the sequel for building the bases of polynomial chaos expansions. Although the above analysis has been constrained to the case when the components of \mathbf{X} are independent, it will be seen that cases with mutually dependent input variables may be treated similarly after an appropriate isoprobabilistic transformation.

3.2 Construction of polynomial basis

Polynomial chaos expansions (PCE) are spectral representations in which the basis consists of multivariate polynomials in \mathbf{X} that are orthonormal with respect to $f_{\mathbf{X}}$. According to Section 3.1, these can be obtained as tensor products of univariate polynomials that constitute Hilbertian bases in the spaces of $X_i, i = 1, \dots, M$.

A Hilbertian basis $\{\psi_k^{(i)}, k \in \mathbb{N}\}$ of \mathcal{H}_i satisfies the orthonormality condition:

$$\langle \psi_j^{(i)}, \psi_k^{(i)} \rangle_{\mathcal{H}_i} = \delta_{jk}, \quad (11)$$

where δ_{jk} is the Kronecker delta symbol. Classical algebra allows one to build a family of orthogonal polynomials $\{Q_k^{(i)}, k \in \mathbb{N}\}$, where k denotes the polynomial degree, so that they satisfy (Abramowitz and Stegun, 1970):

$$\langle Q_j^{(i)}, Q_k^{(i)} \rangle_{\mathcal{H}_i} = c_k^{(i)} \delta_{jk}. \quad (12)$$

In the above equation, the constant $c_k^{(i)}$ represents the squared L_2 norm of the k -th degree polynomial:

$$c_k^{(i)} = \|Q_k^{(i)}\|_2^2 = \langle Q_k^{(i)}, Q_k^{(i)} \rangle_{\mathcal{H}_i}. \quad (13)$$

The corresponding orthonormal polynomial families are thus obtained through the normalization:

$$P_k^{(i)} = Q_k^{(i)} / \sqrt{c_k^{(i)}}. \quad (14)$$

Once the families of univariate orthonormal polynomials associated with the elements of \mathbf{X} have been determined, the multivariate polynomial basis can be obtained through the tensorization shown in Eq. (10), after substituting $\psi_{\alpha_i}^{(i)}$ with $P_{\alpha_i}^{(i)}$.

For standard distributions, *i.e.* uniform, Gaussian, Gamma, Beta, the associated families of orthogonal polynomials are well-known (Xiu and Karniadakis, 2002). For instance, a uniform variable with support $[-1, 1]$ is associated with the family of Legendre polynomials, whereas a standard normal variable is associated with the family of Hermite polynomials. However, it is common in practical situations that the input variables do not follow standard distributions. In such cases, the random vector \mathbf{X} is first transformed into a basic random vector \mathbf{X}' (*e.g.* a standard normal or standard uniform random vector) through an isoprobabilistic transformation $\mathbf{X} = T^{-1}(\mathbf{X}')$ and then, the model response $\mathcal{M}(T^{-1}(\mathbf{X}'))$ is expanded onto the polynomial basis associated with \mathbf{X}' . This approach also allows dealing with mutually dependent input variables through an isoprobabilistic transformation into a vector of independent variables (*e.g.* Nataf transformation in the case of joint PDF with Gaussian copula).

The exact representation of the random response requires an infinite number of basis terms. However, in practical implementation of PCE, an approximation containing a finite number is considered:

$$\hat{Y}^{\text{PCE}} = \hat{\mathcal{M}}^{\text{PCE}}(\mathbf{X}) = \sum_{\alpha \in \mathcal{A}} y_{\alpha} \Psi_{\alpha}(\mathbf{X}), \quad (15)$$

where the set of retained multi-indices \mathcal{A} is determined according to an appropriate truncation scheme. A typical truncation scheme consists in selecting multivariate polynomials up to a total degree p^t , *i.e.* $\mathcal{A} = \{\alpha \in \mathbb{N}^M : \|\alpha\|_1 \leq p^t\}$, with $\|\alpha\|_1 = \sum_{i=1}^M \alpha_i$. The corresponding number of terms in the truncated series is:

$$\text{card}(\mathcal{A}) = \binom{M + p^t}{p^t} = \frac{(M + p^t)!}{M! p^t!}. \quad (16)$$

This number increases exponentially with the input dimension M giving rise to the “curse of dimensionality”. To limit the number of basis terms that include interactions between input variables, which are usually less significant, (Blatman and Sudret, 2010) proposed the use of a hyperbolic truncation scheme. In the latter, the set of retained multi-indices is defined as $\mathcal{A} = \{\alpha \in \mathbb{N}^M : \|\alpha\|_q \leq p^t\}$, with:

$$\|\alpha\|_q = \left(\sum_{i=1}^M \alpha_i^q \right)^{1/q}, \quad 0 < q \leq 1. \quad (17)$$

According to Eq. (17), lower values of q correspond to a smaller number of interaction terms in the PCE basis. At the limit $q \rightarrow 0$, the expansion becomes additive, *i.e.* a sum of univariate functions in the X_i 's.

3.3 Computation of polynomial coefficients

Next, we briefly review a non-intrusive approach for computing the coefficients y_α based on least-square analysis, originally introduced by Choi et al. (2004); Berveiller et al. (2006) under the name *regression method*. In this approach, the exact expansion is viewed as the sum of a truncated series and a residual:

$$Y = \mathcal{M}(\mathbf{X}) = \sum_{\alpha \in \mathcal{A}} y_\alpha \Psi_\alpha(\mathbf{X}) + \varepsilon, \quad (18)$$

where ε corresponds to the truncated terms. Then, the set of coefficients $\mathbf{y} = \{y_\alpha, \alpha \in \mathcal{A}\}$ can be obtained by minimizing the mean-square error of the residual over the ED:

$$\mathbf{y} = \arg \min_{\mathbf{v} \in \mathbb{R}^{\text{card} \mathcal{A}}} \left\| \mathcal{M} - \sum_{\alpha \in \mathcal{A}} v_\alpha \Psi_\alpha \right\|_{\mathcal{E}}^2, \quad (19)$$

leading to:

$$\mathbf{y} = (\Psi^T \Psi)^{-1} \Psi^T \mathcal{Y}, \quad (20)$$

where $\Psi = \{\Psi_{ij} = \Psi_j(\chi^{(i)}), i = 1, \dots, N, j = 1, \dots, \text{card} \mathcal{A}\}$ and \mathcal{Y} is the set of model responses evaluated at the ED, as defined earlier.

For high-dimensional problems, the number of coefficients to be evaluated can be very large. For instance, in typical engineering problems where M varies in the range $10 - 50$, by considering only low-degree polynomials with, say, $p^t = 3$, Eq. (16) results in $286 - 23,426$ unknown coefficients. Obviously, the number of model evaluations required to solve Eq. (19), which is typically 2-3 times the number of unknowns, becomes prohibitively large in cases with high-dimensional input. This limitation constitutes a bottleneck in the classical PCE approach.

More efficient schemes for evaluating the PCE coefficients can be devised by considering the respective regularized problem:

$$\mathbf{y} = \arg \min_{\mathbf{v} \in \mathbb{R}^{\text{card} \mathcal{A}}} \left\| \mathcal{M} - \sum_{\alpha \in \mathcal{A}} v_\alpha \Psi_\alpha \right\|_{\mathcal{E}}^2 + \lambda \mathcal{P}(\mathbf{v}), \quad (21)$$

in which $\mathcal{P}(\mathbf{v})$ is an appropriate regularization functional of $\mathbf{v} = \{v_1, \dots, v_{\text{card} \mathcal{A}}\}$. If $\mathcal{P}(\mathbf{v})$ is selected as the L_1 norm of \mathbf{v} , *i.e.* $\mathcal{P}(\mathbf{v}) = \sum_{i=1}^{\text{card} \mathcal{A}} |v_i|$, insignificant terms may be disregarded from the set of predictors, leading to *sparse* solutions. (Blatman and Sudret, 2011) proposed to use the hybrid least angle regression (LAR) method for building sparse PCE. This method employs the LAR algorithm (Efron et al., 2004) to select the best set of predictors and subsequently, estimates the coefficients using ordinary least squares (OLS), as described above in Eq. (20).

It will be seen in Section 4 that canonical low-rank approximations offer an alternative approach for dealing with high-dimensional problems. By exploiting the tensor-product form of the polynomial basis, such representations may reduce the number of unknown coefficients by orders of magnitude.

3.4 Accuracy estimation

A good measure of the PCE accuracy is the leave-one-out (LOO) error (Allen, 1971), corresponding to the CV error for the case $k = N$ (see Section 2.2 for details on the CV technique). Using algebraic manipulations, this can be computed based on a *single* PCE built with the full ED. Let $h(\chi^{(i)})$ denote the i -th diagonal term of matrix $\Psi(\Psi^T \Psi)^{-1} \Psi^T$. The LOO error can then be computed as (Blatman, 2009):

$$\widehat{Err}_{\text{LOO}} = \left\| \frac{\mathcal{M} - \widehat{\mathcal{M}}^{\text{PCE}}}{1 - h} \right\|_{\mathcal{E}}^2. \quad (22)$$

The relative LOO error is obtained by normalizing $\widehat{Err}_{\text{LOO}}$ with the empirical variance of the model responses at the ED, denoted by $\widehat{\text{Var}}[\mathcal{Y}]$. Because this error tends to be too optimistic, the following corrected estimate is used instead (Chapelle et al., 2002):

$$\widehat{err}_{\text{LOO}} = \frac{\widehat{Err}_{\text{LOO}}}{\widehat{\text{Var}}[\mathcal{Y}]} \left(1 - \frac{\text{card} \mathcal{A}}{N} \right)^{-1} \left(1 + \text{tr}((\Psi^T \Psi)^{-1}) \right). \quad (23)$$

4 Canonical low-rank approximations

4.1 Formulation with polynomial bases

We herein consider again the mapping in Eq. (1). A rank-one function of the input vector \mathbf{X} has the form:

$$w(\mathbf{X}) = \prod_{i=1}^M v^{(i)}(X_i), \quad (24)$$

where $v^{(i)}$ denotes a univariate function of X_i . A representation of the model response $Y = \mathcal{M}(\mathbf{X})$ as a finite sum of rank-one functions constitutes a canonical decomposition with rank equal to the number of rank-one components. A rank- R decomposition of $Y = \mathcal{M}(\mathbf{X})$ therefore reads:

$$\widehat{Y}_R = \widehat{\mathcal{M}}_R(\mathbf{X}) = \sum_{l=1}^R b_l \left(\prod_{i=1}^M v_l^{(i)}(X_i) \right), \quad (25)$$

where $v_l^{(i)}$ denotes a univariate function of X_i in the l -th rank-one component and $\{b_l, l = 1, \dots, R\}$ are scalars that can be viewed as normalizing constants.

An exact canonical decomposition represents a *rank decomposition*. In general, the rank decomposition of a given tensor is not unique; conditions of uniqueness are discussed in Kolda and Bader (2009). The lowest rank of a rank decomposition of a given tensor, called the *tensor rank* (Hitchcock, 1927; Kruskal, 1977), can be determined numerically by fitting various canonical decomposition models (Kolda and Bader, 2009). Naturally, of interest are decompositions where the exact response is approximated with sufficient accuracy by using a relatively small number of terms R . Such decompositions are hereafter called low-rank approximations (LRA).

The focus of the present work is on canonical LRA made of polynomial functions due to the combination of simplicity and versatility these offer; in a general case however, the use of polynomial functions is not a constraint (see *e.g.* Chevreuil et al. (2015)). By expanding $v_l^{(i)}$ onto a polynomial basis that is orthonormal with respect to the marginal distribution f_{X_i} , Eq. (25) takes the form:

$$\widehat{Y}_R = \widehat{\mathcal{M}}_R(\mathbf{X}) = \sum_{l=1}^R b_l \left(\prod_{i=1}^M \left(\sum_{k=0}^{p_i} z_{k,l}^{(i)} P_k^{(i)}(X_i) \right) \right), \quad (26)$$

where $P_k^{(i)}$ denotes the k -th degree univariate polynomial in the i -th input variable, p_i is the maximum degree of $P_k^{(i)}$ and $z_{k,l}^{(i)}$ is the coefficient of $P_k^{(i)}$ in the l -th rank-one component. Appropriate families of univariate polynomials according to the distributions of the respective input variables are determined as discussed in Section 3.2. Similarly to PCE, the case of dependent input can be treated through an isoprobabilistic transformation of the input variables $\{X_i, i = 1, \dots, N\}$ into independent reduced variables.

Disregarding the redundant parameterization arising from the normalizing constants, the number of unknowns in Eq. (26) is $R \cdot \sum_{i=1}^M (p_i + 1)$, which grows *only linearly* with the input dimension M . Thus, a representation of the model response in the form of canonical LRA results in a drastic reduction of the number of unknowns as compared to PCE. To emphasize this, we consider PCE with the candidate basis determined by the truncation scheme $\mathcal{A} = \{\boldsymbol{\alpha} \in \mathbb{N}^M : \alpha_i \leq p_i, i = 1, \dots, M\}$, so that the expansion relies on the same polynomial functions as those used in Eq. (26). For the case when $p_i = p, i = 1, \dots, M$, the resulting number of unknowns is $(p+1)^M$ in PCE versus $(p+1) \cdot M \cdot R$ in LRA. Assuming a typical engineering problem with $M = 10$ and low-degree polynomials with $p = 3$, these formulas yield 1,048,576 unknowns in PCE versus $40R$ unknowns in LRA; for a low rank, say $R \leq 10$, the latter number does not exceed a mere 400. The reduction in the number of unknowns achieved with the compressed LRA representation becomes even more pronounced in cases with high-degree polynomials.

4.2 Greedy construction

Different algorithms have been proposed in the literature for building a decomposition in the form of Eq. (26) in a non-intrusive manner (Chevreuil et al., 2013; Doostan et al., 2013; Mathelin, 2014; Rai, 2014; Validi, 2014; Chevreuil et al., 2015). A common point is that the polynomial coefficients are determined by means of an alternated least-squares (ALS) minimization. The ALS approach consists in sequentially solving a least-squares minimization problem along a single dimension $i \in \{1, \dots, M\}$, while “freezing” the coefficients in all remaining dimensions. In the present study, we adopt the skeleton of the greedy algorithm proposed in Chevreuil et al. (2013, 2015). This algorithm involves a progressive increase of the rank by successively adding rank-one components up to a prescribed maximal. It thus results in a set of candidate decompositions with varying ranks and requires appropriate criteria for selecting the optimal one. In the following, we describe the construction of the approximation for a prescribed rank and discuss rank selection criteria.

4.2.1 Approximation for a prescribed rank

Let us denote by $\widehat{\mathcal{M}}_r$ the rank- r approximation of \mathcal{M} . The corresponding model response is then approximated by:

$$\widehat{Y}_r = \widehat{\mathcal{M}}_r(\mathbf{X}) = \sum_{l=1}^r b_l w_l(\mathbf{X}), \quad (27)$$

where w_l represents the l -th rank-one component:

$$w_l(\mathbf{X}) = \prod_{i=1}^M \left(\sum_{k=0}^{p_i} z_{k,l}^{(i)} P_k^{(i)}(X_i) \right). \quad (28)$$

The employed algorithm involves a sequence of pairs of a *correction step* and an *updating step*, so that in the r -th correction step, the rank-one tensor w_r is built, while in the r -th

updating step, the set of normalizing coefficients $\{b_1, \dots, b_r\}$ is determined. These steps are detailed next.

Correction step: In the r -th correction step, the rank-one tensor w_r is obtained as the solution to the minimization problem:

$$w_r(\mathbf{X}) = \arg \min_{\omega \in \mathcal{W}} \left\| \mathcal{M} - \widehat{\mathcal{M}}_{r-1} - \omega \right\|_{\mathcal{E}}^2, \quad (29)$$

where \mathcal{W} represents the space of rank-one tensors. The sequence is initiated by setting $Y_0 = \widehat{\mathcal{M}}_0(\mathbf{X}) = 0$. Eq. (29) is solved by means of an ALS scheme that involves successive minimizations along each dimension $i = 1, \dots, M$. In the minimization along dimension j , the polynomial coefficients in all other dimensions are “frozen” at their current values and the coefficients $\mathbf{z}_r^{(j)} = \{z_{1,r}^{(j)} \dots z_{p_j,r}^{(j)}\}$ are determined as:

$$\mathbf{z}_r^{(j)} = \arg \min_{\zeta \in \mathbb{R}^{p_j+1}} \left\| \mathcal{M} - \widehat{\mathcal{M}}_{r-1} - \left(\prod_{i \neq j} v_r^{(i)} \right) \left(\sum_{k=0}^{p_j} \zeta_k P_k^{(j)} \right) \right\|_{\mathcal{E}}^2, \quad (30)$$

where:

$$v_r^{(i)}(X_i) = \sum_{k=0}^{p_i} z_{k,r}^{(i)} P_k^{(i)}(X_i). \quad (31)$$

To initiate the r -th correction step, one needs to assign arbitrary values to $v_r^{(i)}$, $i = 1, \dots, M$; in the subsequent example applications, we use $v_r^{(1)}(X_1) = \dots = v_r^{(M)}(X_M) = 1$. We underline that a correction step may involve several iterations over the set of dimensions $\{1, \dots, M\}$. This aspect of the algorithm, which is not detailed in some of the aforementioned studies, can be critical for the LRA accuracy, as shown in the numerical investigations in Section 5. Chevreuil et al. (2015) proposed a stopping criterion that combines thresholds on the number of iterations I_r and on the empirical error:

$$\widehat{Err}_r = \left\| \mathcal{M} - \widehat{\mathcal{M}}_{r-1} - w_r \right\|_{\mathcal{E}}^2. \quad (32)$$

A threshold on the empirical error was also imposed by Doostan et al. (2013) and Validi (2014). However, appropriate values for these thresholds as well as their effects on the LRA accuracy were not examined. The stopping criterion employed in the present study involves the number of iterations I_r and the decrease in the relative empirical error in two successive iterations, denoted by $\Delta \widehat{err}_r$. The relative empirical error is computed by normalizing \widehat{Err}_r in Eq. (32) with $\widehat{\text{Var}}[\mathcal{Y}]$, *i.e.* with the empirical variance of the model responses at the ED. Accordingly, the algorithm exits the r -th correction step if either I_r reaches a maximum allowable value I_{\max} or $\Delta \widehat{err}_r$ becomes smaller than a prescribed threshold $\Delta \widehat{err}_{\min}$. Appropriate values for I_{\max} and $\Delta \widehat{err}_{\min}$ will be later discussed based on the numerical investigations in Section 5.

Updating step: After the completion of a correction step, the algorithm moves to an updating step, in which the set of coefficients $\mathbf{b} = \{b_1 \dots b_r\}$ is obtained as the solution of the minimization problem:

$$\mathbf{b} = \arg \min_{\beta \in \mathbb{R}^r} \left\| \mathcal{M} - \sum_{l=1}^r \beta_l w_l \right\|_{\mathcal{E}}^2. \quad (33)$$

Note that in each updating step, the size of vector \mathbf{b} is increased by one. In the r -th updating step, the value of the new element b_r is determined for the first time, whereas the values of the existing elements $\{\beta_1, \dots, \beta_{r-1}\}$ are updated.

Construction of a rank- R decomposition in the form of Eq. (26) requires repeating pairs of a correction and an updating step for $r = 1, \dots, R$. The algorithm is summarized below.

Algorithm 1 Non-intrusive construction of a polynomial rank- R approximation of $Y = \mathcal{M}(\mathbf{X})$ with an experimental design $\mathcal{E} = \{\chi^{(1)}, \dots, \chi^{(N)}\}$:

1. Set $\widehat{\mathcal{M}}_0(\chi^{(q)}) = 0$, $q = 1, \dots, N$.
2. For $r = 1, \dots, R$, repeat steps (a)-(d):
 - (a) Initialize: $v_r^{(i)}(\chi^{(q)}) = 1$, $i = 1, \dots, M$, $q = 1, \dots, N$; $I_r = 0$; $\widehat{\Delta err}_r = \epsilon > \widehat{\Delta err}_{\min}$.
 - (b) While $\widehat{\Delta err}_r > \widehat{\Delta err}_{\min}$ and $I_r < I_{\max}$, repeat steps i-iv:
 - i. Set $I_r \leftarrow I_r + 1$.
 - ii. For $i = 1, \dots, M$, repeat steps A-B:
 - A. Determine $\mathbf{z}_r^{(i)} = \{z_{1,r}^{(i)} \dots z_{p_i,r}^{(i)}\}$ using Eq. (30).
 - B. Update $v_r^{(i)}$, using Eq. (31).
 - iii. Update w_r using Eq. (28).
 - iv. Compute \widehat{Err}_r using Eq. (32) and update $\widehat{\Delta err}_r$.
 - (c) Determine $\mathbf{b} = \{b_1 \dots b_r\}$ using Eq. (33).
 - (d) Evaluate $\widehat{\mathcal{M}}_r(\chi^{(q)})$, $q = 1, \dots, N$, using Eq. (27).

We emphasize that the above algorithm relies on the solution of several small-size minimization problems; in particular, one needs to solve M minimization problems of size $\{p_i + 1, i = 1, \dots, M\}$ in each iteration of a correction step (note that $p_i < 20$ in typical applications) and one minimization problem of size r in the r -th updating step (recall that small ranks are of interest in LRA). Thus, the LRA construction substitutes the single large-size minimization problem involved in the PCE construction with a series of small-size ones; in high-dimensional applications, this can also offer a significant advantage in terms of required computer memory.

Because of the small number of unknowns involved in Eq. (30) and Eq. (33), these minimization problems can be efficiently solved with OLS, as shown in the subsequent example applications. An alternative approach employed in (Chevreuil et al., 2013; Rai, 2014; Chevreuil et al., 2015) is to substitute these equations with the respective regularized problems:

$$\mathbf{z}_r^{(j)} = \arg \min_{\boldsymbol{\zeta} \in \mathbb{R}^{p_j}} \left\| \mathcal{M} - \widehat{\mathcal{M}}_{r-1} - \left(\prod_{i \neq j} v_r^{(i)} \right) \left(\sum_{k=0}^{p_j} \zeta_k P_k^{(j)} \right) \right\|_{\mathcal{E}}^2 + \lambda \mathcal{P}(\boldsymbol{\zeta}) \quad (34)$$

and

$$\mathbf{b} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^r} \left\| \mathcal{M} - \sum_{l=1}^r \beta_l w_l \right\|_{\mathcal{E}}^2 + \lambda \mathcal{P}(\boldsymbol{\beta}), \quad (35)$$

where $\mathcal{P}(\mathbf{u})$ is selected as either the L_1 or the L_2 norm of \mathbf{u} .

We conclude this section by briefly referring to alternative algorithms for developing LRA non-intrusively. The algorithms in Doostan et al. (2013); Mathelin (2014); Validi (2014) involve a progressive increase of the rank as well. In Doostan et al. (2013); Validi (2014), when the r -th rank-one component is added, the polynomial coefficients in the $(r - 1)$ previously built rank-one terms are also updated. Thus, the r -th step requires the solution of minimization problems of size $\{(p_i + 1)r, i = 1, \dots, M\}$; in applications with high-dimensional input, the size of these minimization problems remains orders of magnitude smaller compared

to those involved in the computation of the PCE coefficients. Conversely, the algorithm in Mathelin (2014) does not require the updating of coefficients in previously added components when a new rank-one term is added. It is noteworthy that in the latter, the most-relevant dimensions to be included in the ALS scheme are determined by means of a LAR-based selection technique and CV error criteria.

4.2.2 Rank selection

In a typical application, the optimal rank R is not known *a priori*. As noted earlier, the progressive construction of LRA described in Section 4.2.1 results in a set of decompositions of increasing rank. Thus, one may set $r = 1, \dots, r_{\max}$ in Step 2 of Algorithm 1, where r_{\max} is a maximum allowable candidate rank, and at the end, select the optimal-rank decomposition using error-based criteria. Two approaches for rank selection are described below and investigated extensively in Section 5.

Chevreur et al. (2015) proposed to select the optimal rank by means of 3-fold CV (see Section 2.2); they did not however examine the accuracy of this approach. In the general case of k -fold CV, the procedure requires building LRA of increasing rank $r = 1, \dots, r_{\max}$ for each of the k training sets. For $r = 1, \dots, r_{\max}$, the generalization error of each of the k meta-models is estimated using the respective testing set. The rank $R \in \{1, \dots, r_{\max}\}$ yielding the smallest average generalization error over the k meta-models is identified as optimal; then, a new decomposition of rank R is built using the full ED. The average generalization error corresponding to the selected rank provides an estimate of the actual generalization error of the final meta-model.

Although the above procedure for rank selection relies on the set of model evaluations at the ED only, it requires repeating Algorithm 1 $(k+1)$ times, namely k times for rank r_{\max} and another time for rank R . Because it is of interest to limit the computational effort required for rank selection, we consider an alternative approach based on the LOO error in the updating step. As described in Section 3.4, the LOO error represents the average generalization error in k -fold CV with $k = N$, but can be computed *from a single meta-model* built with the full ED.

Let \mathbf{W} denote the information matrix in the minimization problem represented by Eq. (33), *i.e.* $\mathbf{W} = \{W_{ij} = w_j(\boldsymbol{\chi}^{(i)}), i = 1, \dots, N, j = 1, \dots, r\}$, and let $h(\boldsymbol{\chi}^{(i)})$ denote the i -th diagonal term of matrix $\mathbf{W}(\mathbf{W}^T \mathbf{W})^{-1} \mathbf{W}^T$. Then, the LOO error can be computed as:

$$\widehat{Err}_{\text{LOO}} = \left\| \frac{\mathcal{M} - \widehat{\mathcal{M}}_r}{1 - h} \right\|_{\mathcal{E}}^2. \quad (36)$$

The relative LOO error is obtained by normalizing $\widehat{Err}_{\text{LOO}}$ with $\widehat{\text{Var}}[\mathcal{Y}]$, *i.e.* with the empirical variance of the model responses at the ED. As noted in Section 3.4, this error may be too optimistic. For this reason, a corrected error may be used in the case of PCE (see Eq. (23)). A similar correction may not be formally applied in the present case, because this is based on the assumption of orthogonality between the regressors, which does not hold for Eq. (33). However, it is of interest to examine the performance of the corrected error estimate, given by:

$$\widehat{err}_{\text{LOO}} = \frac{\widehat{Err}_{\text{LOO}}}{\widehat{\text{Var}}[\mathcal{Y}]} \left(1 - \frac{r}{N}\right)^{-1} \left(1 + \text{tr}((\mathbf{W}^T \mathbf{W})^{-1})\right), \quad (37)$$

used in an approximate sense.

5 Example applications of canonical low-rank approximations

In the following, we investigate the construction of LRA according to Algorithm 1 in Section 4.2.1 considering four models with different characteristics and input dimensionality. The first model is an analytical function of rank-1 structure with a 5-dimensional input, representing the deflection of a simply-supported beam under a static load. The subsequent three examples involve finite-element models with input dimensionality 10, 21 and 53. In particular, we develop LRA for the deflection of a truss structure subjected to vertical static loads, the top horizontal displacement of a three-span five-story frame subjected to horizontal static loads and the temperature response in stationary heat conduction with thermal conductivity described by a random field. In each application, (i) we assess the performance of the two criteria for rank selection discussed in Section 4.2.2 and evaluate the accuracy of the respective error measures and (ii) we investigate optimal stopping criteria in the correction step. The minimization problems in both the correction and the updating steps are solved with the OLS method. In all applications, effects of varying ED size are examined. The EDs are herein obtained using Sobol quasi-random sequences (Niederreiter, 1992) (LHS designs will be considered in the following section). In order to assess the accuracy of the different rank-selection methods and stopping criteria, we use sufficiently large validation sets drawn with MCS. We note however that such large validation sets are not available in real-life problems, in which the analyst typically needs to rely solely on the ED.

5.1 Beam deflection

In the first example, we consider a simply-supported beam with a constant rectangular cross-section subjected to a concentrated load at the midpoint of the span. The response quantity of interest is the mid-span deflection given by:

$$u = \frac{PL^3}{4Ebh^3}, \quad (38)$$

where b and h respectively denote the width and height of the cross-section, L is the length of the beam, E is the Young's modulus and P is the magnitude of the load. The aforementioned parameters are described by independent random variables, thus leading to an uncertainty propagation problem of dimension $M = 5$. The distributions of the input parameters are listed in Table 1. Employing an isoprobabilistic transformation of $\mathbf{X} = \{b, h, L, E, P\}$ into a vector of standard normal variables, we develop LRA representations of the random response $U = \mathcal{M}(\mathbf{X})$ with basis functions made of Hermite polynomials.

Table 1: Beam-deflection problem: Distributions of input random variables.

Variable	Distribution	Mean	CoV
b [m]	Lognormal	0.15	0.05
h [m]	Lognormal	0.3	0.05
L [m]	Lognormal	5	0.01
E [MPa]	Lognormal	30,000	0.15
P [MN]	Lognormal	0.01	0.20

5.1.1 Rank selection and error measures

We begin our analysis by investigating the selection of the optimal rank among a set of candidate values $\{1, \dots, 20\}$. After preliminary investigations, we set a common polynomial degree $p_i = p = 5$ for $i = 1, \dots, 5$. The stopping criterion in the correction step is defined by setting $I_{\max} = 50$ and $\Delta\widehat{err}_{\min} = 10^{-8}$ (the selection of $\Delta\widehat{err}_{\min}$ will be explained in the following subsection). We use EDs of size N varying from 50 to 5,000 and a validation set of size $n_{\text{val}} = 10^6$. The actual optimal rank, denoted by R^{opt} , is identified as the one yielding the minimum relative generalization error \widehat{err}_G , the latter estimated with the validation set as described in Section 2.2. It is of interest to compare R^{opt} with the rank selected with the 3-fold CV approach, denoted by R^{CV3} , and the rank selected with the LOO-error criterion, denoted by R^{LOO} . Note that for a given ED, the value of R^{CV3} is not fixed because of the random partition of the ED into training and testing sets. This aspect of randomness is considered in the following analysis.

In Figure 1, we investigate rank selection with the 3-fold CV approach. The left graph shows boxplots of R^{CV3} for 20 random partitions of each considered ED into training and testing sets, while the right graph shows boxplots of the relative generalization errors of the resulting meta-models. In order to assess the accuracy of the 3-fold CV approach, we also plot the actual optimal rank R^{opt} and corresponding relative generalization error for each ED. The rank-1 structure of the original model (see Eq. (38)) matches the structure of the optimal LRA for all N except for $N = 100$ ($R^{\text{opt}} = 2$ in this case). The median R^{CV3} coincides with R^{opt} for most EDs, while selection of a non-optimal rank appears to have a smaller effect for the larger EDs. We underline that highly accurate meta-models are obtained in all cases, with \widehat{err}_G not exceeding 10^{-4} . For $N \geq 200$, \widehat{err}_G is of the order of 10^{-8} or smaller even for cases with non-optimal ranks. The LOO-error criterion yields $R^{\text{LOO}} = 1$ for all considered EDs, which coincides with R^{opt} except for $N = 100$.

To gain further insight into the criteria for rank selection, we examine the estimation of the generalization error by the ED-based error measures. In the left graph of Figure 2, we consider meta-models with rank R^{CV3} and compare the 3-fold CV error $\widehat{err}_{\text{CV3}}$ with the corresponding generalization error \widehat{err}_G for one example partition of each ED. In the right graph of the same figure, we consider meta-models with rank R^{LOO} and compare the LOO error $\widehat{err}_{\text{LOO}}$ with the corresponding generalization error. In these graphs, the 3-fold CV error appears to be a better estimator of the generalization error compared to the LOO error. It must be noted that contrary to the 3-fold CV error, the estimation of the LOO error largely deteriorates for higher ranks. The lack of orthogonality of the regressors leads to overly high correction factors, whereas the non-corrected errors are too optimistic. Overall, $\widehat{err}_{\text{LOO}}$ is deemed inappropriate for estimating \widehat{err}_G in the present example.

5.1.2 Stopping criterion in the correction step

We herein investigate the effects of different stopping criteria in the correction step on the LRA accuracy. For $N \in \{50; 200; 1,000; 5,000\}$, the left graph of Figure 3 shows the relative generalization errors of the LRA meta-models with optimal rank, while the parameter $\Delta\widehat{err}_{\min}$ that drives the loops over the set of dimensions in ALS varies between 10^{-9} and 10^{-4} ; other parameters are fixed to their values in Section 5.1.1. The right graph of the same figure indicates the maximum number of iterations I_r performed in a correction step. Except for the largest ED, the accuracy of LRA strongly depends on $\Delta\widehat{err}_{\min}$, with decreasing values leading to orders-of-magnitude smaller \widehat{err}_G . The right graph indicates that a larger

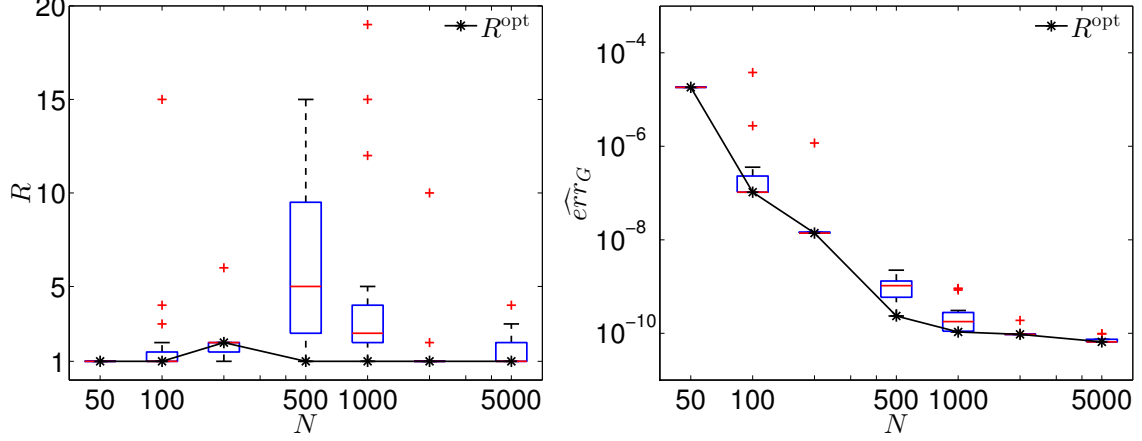


Figure 1: Beam-deflection problem: Comparison of ranks selected with 3-fold CV (20 replications) to optimal ranks (left) and corresponding relative generalization errors (right).




Figure 2: Beam-deflection problem: Comparison of ED-based errors to corresponding relative generalization errors for ranks selected with 3-fold CV (left) and with the LOO-error criterion (right).

number of iterations is required for the smallest ED.

5.2 Truss deflection

The second example also derives from structural mechanics, but involves a finite-element model. In particular, the truss shown in Figure 4 is considered (also studied in Blatman and Sudret (2008)), with the mid-span deflection u representing the response quantity of interest. The random input herein comprises $M = 10$ independent variables: the vertical loads P_1, \dots, P_6 , the cross-sectional area and Young's modulus of the horizontal bars, respectively denoted by A_1 and E_1 , and the cross-sectional area and Young's modulus of the diagonal bars, respectively denoted by A_2 and E_2 . The distributions of the input variables are listed in Table 2. After employing an isoprobabilistic transformation of $\mathbf{X} = \{P_1, \dots, P_6, A_1, A_2, E_1, E_2\}$ into a vector of standard normal variables, we develop LRA representation of the random response $U = \mathcal{M}(\mathbf{X})$ with basis functions made of Hermite polynomials. In the underlying deterministic problem, the mid-span deflection is com-




Figure 3: Beam-deflection problem: Relative generalization error (left) and corresponding maximum number of iterations in a correction step (right) versus threshold of differential error in the stopping criterion.

puted with an in-house finite-element analysis code developed in the Matlab environment.




Figure 4: Truss structure.

Table 2: Truss-deflection problem: Distributions of input random variables.

Variable	Distribution	Mean	CoV
A_1 [m]	Lognormal	0.002	0.10
A_2 [m]	Lognormal	0.001	0.10
E_1, E_2 [MPa]	Lognormal	$2.1 \cdot 10^5$	0.10
P_1, \dots, P_6 [kN]	Gumbel	50	0.15

5.2.1 Rank selection and error measures

Similarly to the previous example, we investigate rank selection based on the 3-fold CV and LOO errors, considering the candidate values $\{1, \dots, 20\}$. After preliminary investigations, the polynomial degree is set to $p_i = p = 3$ for $i = 1, \dots, 10$, while the parameters of the stopping criterion in the correction step are set to $I_{\max} = 50$ and $\Delta \widehat{err}_{\min} = 10^{-6}$. Again, we use EDs of size N varying between 50 and 5,000 to build the LRA meta-models and a validation set of size $n_{\text{val}} = 10^6$ to estimate the relative generalization errors and identify the actual optimal ranks.

The left graph of Figure 5 shows boxplots of $R^{\text{CV}3}$, *i.e.* the rank selected with 3-fold CV, for 20 random partitions of each considered ED, together with the corresponding actual optimal rank R^{opt} . The right graph shows boxplots of the relative generalization errors of

the meta-models with rank R^{CV3} as well as the relative generalization errors of the meta-models with rank R^{opt} . We observe that the actual optimal rank is equal to 1 for $N \leq 500$, but attains higher values for larger EDs. This is because the larger amount of information contained in the latter allows the estimation of a larger number of coefficients with higher accuracy. This trend is captured by the 3-fold CV approach, which yields $R^{\text{CV3}} = 1$ for all random partitions of the ED when $N \leq 500$ and varying higher ranks for the larger EDs. The median of R^{CV3} coincides with R^{opt} for $N = 5,000$, but is smaller than R^{CV3} for $N = 1,000$ and $N = 2,000$. As in the beam-deflection problem, effects of non-optimal rank selection on the LRA accuracy tend to be less significant as the ED size increases. The LOO-error criterion yields $R^{\text{LOO}} = 1$ for all considered EDs, which coincides with R^{opt} only for $N \leq 500$.

We next assess the estimation of the generalization error by the ED-based error measures. In the left graph of Figure 6, we compare $\widehat{err}_{\text{CV3}}$ to \widehat{err}_G for LRA meta-models with rank R^{CV3} . In the right graph, we compare $\widehat{err}_{\text{LOO}}$ to \widehat{err}_G for LRA meta-models with rank R^{LOO} . Both ED-based error measures appear to approximate the corresponding generalization errors fairly well, particularly for the larger EDs. However, similarly to the beam-deflection problem, $\widehat{err}_{\text{LOO}}$ is found to be an inappropriate estimator of \widehat{err}_G for higher-rank meta-models. On the other hand, for sufficiently large EDs, $\widehat{err}_{\text{CV3}}$ provides fair estimates of \widehat{err}_G for all ranks.

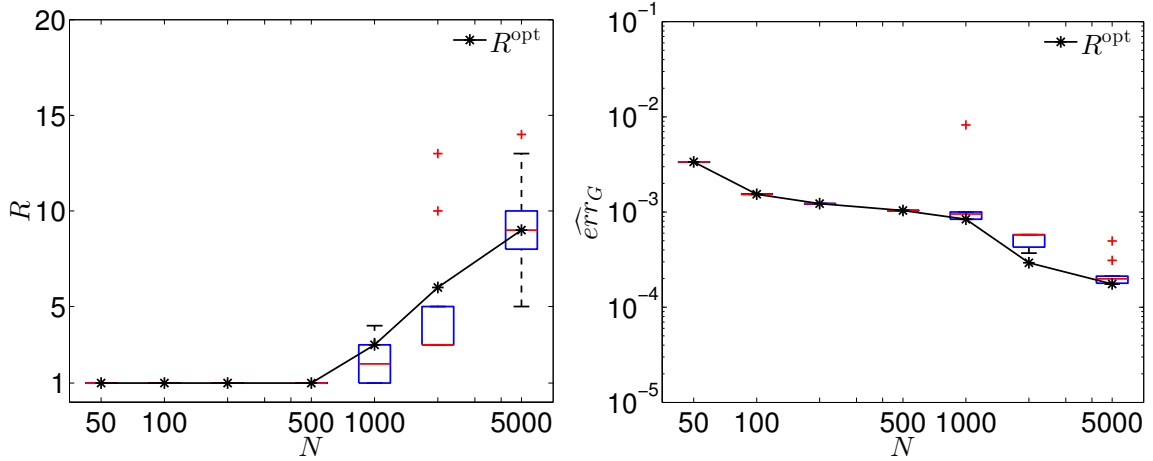


Figure 5: Truss-deflection problem: Comparison of ranks selected with 3-fold CV (20 replications) to optimal ranks (left) and corresponding relative generalization errors (right).

5.2.2 Stopping criterion in the correction step

We now examine effects of the differential error threshold in the correction step, while other parameters are fixed to their values above. For $N \in \{50; 200; 1,000; 5,000\}$, the left graph of Figure 7 shows the relative generalization errors for the LRA meta-models with optimal ranks, while $\Delta\widehat{err}_{\text{min}}$ varies from 10^{-7} to 10^{-2} . The right graph of the same figure shows the corresponding maximum number of iterations I_r in a correction step. While selecting a sufficiently small threshold $\Delta\widehat{err}_{\text{min}}$ appears critical for the smaller ED, it does not essentially affect the meta-model accuracy when $N = 200$ or $N = 1,000$. For $N = 5,000$, variations of \widehat{err}_G within the same order of magnitude are observed with decreasing $\Delta\widehat{err}_{\text{min}}$, following a decreasing but non-monotonic trend. Thus, small values of $\Delta\widehat{err}_{\text{min}}$ may lead to unnecessary




Figure 6: Truss-deflection problem: Comparison of ED-based errors to corresponding relative generalization errors for ranks selected with 3-fold CV (left) and with the LOO-error criterion (right).

iterations in the correction step for the larger EDs.

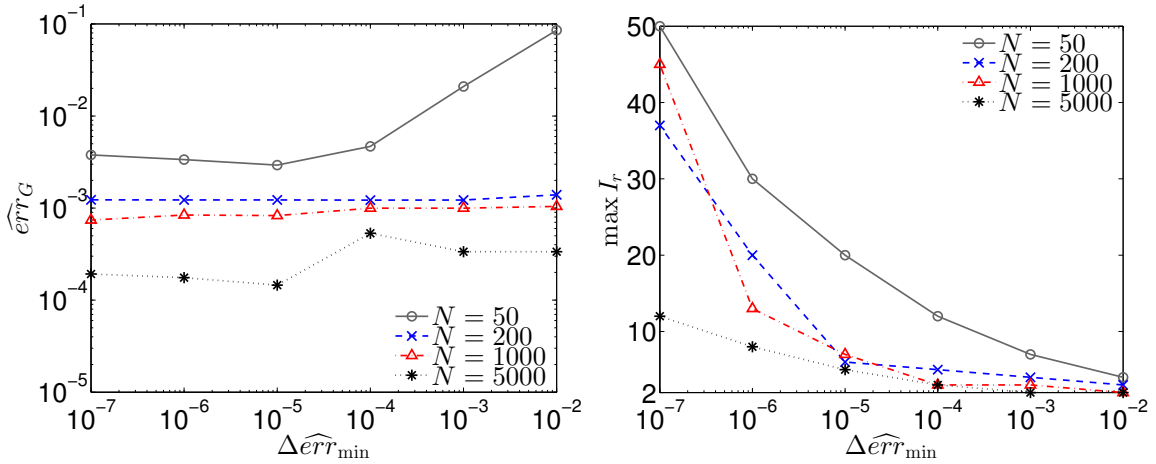


Figure 7: Truss-deflection problem: Relative generalization error (left) and corresponding maximum number of iterations in a correction step (right) versus threshold of differential error in the stopping criterion.

5.3 Frame displacement

The present example involves a finite-element model representing the three-span five-story frame shown in Figure 8 (also studied in Liu and Der Kiureghian (1991); Blatman and Sudret (2010)). The response quantity of interest is the horizontal displacement u at the top right corner of the top floor, under the depicted horizontal loads acting at the floor levels. Table 3 lists the properties (Young's modulus, moment of inertia, cross-sectional area) of the different elements according to their labels in Figure 8. These element properties together with the values of the horizontal loads comprise the random input of the problem of dimension $M = 21$. The distributions of the input random variables are listed in Table 4. The truncation of the Gaussian distributions used to model the element properties is a

modification of the original example in Liu and Der Kiureghian (1991) that was introduced by Blatman and Sudret (2010). The input variables are correlated, with their dependence structure defined using a Gaussian copula (see Nelsen (2006) for modeling of probabilistic dependence with copulas). The elements of the associated linear correlation matrix are defined as follows:

- the correlation coefficient between the cross-sectional area A_i and the moment of inertia I_i of a certain element i is $\rho_{A_i, I_i} = 0.95$;
- the correlation coefficient between the geometric properties of two distinct elements i and j are $\rho_{A_i, I_j} = \rho_{I_i, I_j} = \rho_{A_i, A_j} = 0.13$;
- the correlation coefficient between the two Young's moduli is $\rho_{E_1, E_2} = 0.90$;
- the remaining correlation coefficients are zero.

Note that in the original example, the above values are considered for the corresponding linear correlation coefficients in the standard normal space; however, as noted in Blatman and Sudret (2010), the differences between the two are insignificant. After an isoprobabilistic transformation of $\mathbf{X} = \{P_1, P_2, P_3, E_1, E_2, I_1, \dots, I_8, A_1, \dots, A_8\}$ into a vector of independent standard normal variables, we develop LRA representations of the random response $U = \mathcal{M}(\mathbf{X})$ with the basis functions made of Hermite polynomials. In the underlying deterministic problem, the displacement of interest is computed with an in-house finite-element analysis code developed in the Matlab environment.




Figure 8: Frame structure.

5.3.1 Rank selection and error measures

We herein investigate rank selection, while setting the polynomial degree to $p_i = p = 3$ for $i = 1, \dots, 21$. We define the stopping criterion in the correction step by requiring $I_{\max} = 50$ and $\Delta \widehat{err}_{\min} = 10^{-6}$. As in the two preceding examples, we use EDs of sizes N varying from 50 to 5,000 to build the meta-models and a validation set of size $n_{\text{val}} = 10^6$ to identify the actual optimal rank among the candidate values $\{1, \dots, 20\}$.

Table 3: Frame-displacement problem: Element properties.

Element	Young's modulus	Moment of inertia	Cross-sectional area
C_1	E_1	I_1	A_1
C_2	E_1	I_2	A_2
C_3	E_1	I_3	A_3
C_4	E_1	I_4	A_4
B_1	E_2	I_5	A_5
B_2	E_2	I_6	A_6
B_3	E_2	I_7	A_7
B_4	E_2	I_8	A_8

Table 4: Frame-displacement problem: Distributions of input random variables.

Variable	Distribution	Mean	Standard deviation
P_1 [KN]	Lognormal	133.45	40.04
P_2 [KN]	Lognormal	88.97	35.59
P_3 [KN]	Lognormal	71.17	28.47
E_1 [KN/m ²]	Truncated Gaussian over $[0, \infty)$	$2.3796 \cdot 10^7$	$1.9152 \cdot 10^6$
E_2 [KN/m ²]	Truncated Gaussian over $[0, \infty)$	$2.1738 \cdot 10^7$	$1.9152 \cdot 10^6$
I_1 [m ⁴]	Truncated Gaussian over $[0, \infty)$	$8.1344 \cdot 10^{-3}$	$1.0834 \cdot 10^{-3}$
I_2 [m ⁴]	Truncated Gaussian over $[0, \infty)$	$1.1509 \cdot 10^{-2}$	$1.2980 \cdot 10^{-3}$
I_3 [m ⁴]	Truncated Gaussian over $[0, \infty)$	$2.1375 \cdot 10^{-2}$	$2.5961e - 03$
I_4 [m ⁴]	Truncated Gaussian over $[0, \infty)$	$2.5961 \cdot 10^{-2}$	$3.0288 \cdot 10^{-3}$
I_5 [m ⁴]	Truncated Gaussian over $[0, \infty)$	$1.0811 \cdot 10^{-2}$	$2.5961 \cdot 10^{-3}$
I_6 [m ⁴]	Truncated Gaussian over $[0, \infty)$	$1.4105 \cdot 10^{-2}$	$3.4615 \cdot 10^{-3}$
I_7 [m ⁴]	Truncated Gaussian over $[0, \infty)$	$2.3279 \cdot 10^{-2}$	$5.6249 \cdot 10^{-3}$
I_8 [m ⁴]	Truncated Gaussian over $[0, \infty)$	$2.5961 \cdot 10^{-2}$	$6.4902 \cdot 10^{-3}$
A_1 [m ²]	Truncated Gaussian over $[0, \infty)$	$3.1256 \cdot 10^{-1}$	$5.5815 \cdot 10^{-2}$
A_2 [m ²]	Truncated Gaussian over $[0, \infty)$	$3.7210 \cdot 10^{-1}$	$7.4420 \cdot 10^{-2}$
A_3 [m ²]	Truncated Gaussian over $[0, \infty)$	$5.0606 \cdot 10^{-1}$	$9.3025 \cdot 10^{-2}$
A_4 [m ²]	Truncated Gaussian over $[0, \infty)$	$5.5815 \cdot 10^{-1}$	$1.1163 \cdot 10^{-1}$
A_5 [m ²]	Truncated Gaussian over $[0, \infty)$	$2.5302 \cdot 10^{-1}$	$9.3025 \cdot 10^{-2}$
A_6 [m ²]	Truncated Gaussian over $[0, \infty)$	$2.9117 \cdot 10^{-1}$	$1.0232 \cdot 10^{-1}$
A_7 [m ²]	Truncated Gaussian over $[0, \infty)$	$3.7303 \cdot 10^{-1}$	$1.2093 \cdot 10^{-1}$
A_8 [m ²]	Truncated Gaussian over $[0, \infty)$	$4.1860 \cdot 10^{-1}$	$1.9537 \cdot 10^{-1}$

Figure 9 presents results of rank selection with the 3-fold CV approach, considering 20 random partitions of each ED. The left and right graphs of this figure respectively show boxplots of R^{CV3} and the relative generalization errors of the corresponding meta-models. The accuracy of the approach is assessed through comparisons to LRA with the actual optimal rank R^{opt} , identified by means of the validation set. As in the truss-deflection problem, the optimal rank is $R^{\text{opt}} = 1$ for the smaller EDs and increases for larger N ,

without however exceeding $R^{\text{opt}} = 2$ in the present example. Except for $N = 50$, the ED-based rank R^{CV3} exhibits small or no dispersion for different random selections of the training and testing sets, while its median value coincides with R^{opt} in most cases. Selection of a rank higher than unity may lead to gross generalization errors for $N = 50$, because the few points comprising the ED are insufficient to accurately estimate a large number of coefficients. Again, the LOO-error criterion yields $R^{\text{LOO}} = 1$ for all considered EDs.

Figure 10 provides further insight into rank selection by comparing the ED-based errors $\widehat{err}_{\text{CV3}}$ and $\widehat{err}_{\text{LOO}}$ with the corresponding relative generalization errors \widehat{err}_G . The left graph plots $\widehat{err}_{\text{CV3}}$ together with \widehat{err}_G for LRA with rank R^{CV3} , while the right graph plots $\widehat{err}_{\text{LOO}}$ together with \widehat{err}_G for LRA with rank R^{LOO} . The ED-based error measures appear to approach the corresponding generalization errors with increasing ED size, with $\widehat{err}_{\text{CV3}}$ providing accurate estimates with only $N = 500$. As also observed in the previous examples, $\widehat{err}_{\text{LOO}}$ may largely overestimate the errors of meta-models with higher ranks.




Figure 9: Frame-displacement problem: Comparison of ranks selected with 3-fold CV (20 replications) to optimal ranks (left) and corresponding relative generalization errors (right).




Figure 10: Frame-displacement problem: Comparison of ED-based errors to corresponding relative generalization errors for ranks selected with 3-fold CV (left) and with the LOO-error criterion (right).

5.3.2 Stopping criterion in the correction step

We next investigate effects of the stopping criteria on the LRA accuracy considering the cases with $N \in \{50; 200; 1,000; 5,000\}$. The left graph of Figure 10 shows the relative generalization errors of LRA with optimal ranks while $\Delta \widehat{err}_{\min}$ varies between 10^{-6} and 10^{-1} and other parameters are fixed to their values above. The right graph shows the corresponding maximum number of iterations in a correction step. For $N = 200$, the considered variation of $\Delta \widehat{err}_{\min}$ corresponds to a variation of \widehat{err}_G of about one order of magnitude. For the other EDs, the value of $\Delta \widehat{err}_{\min}$ has a negligible effect on the LRA accuracy and thus, repeated iterations in the correction step add an unnecessary computational effort.

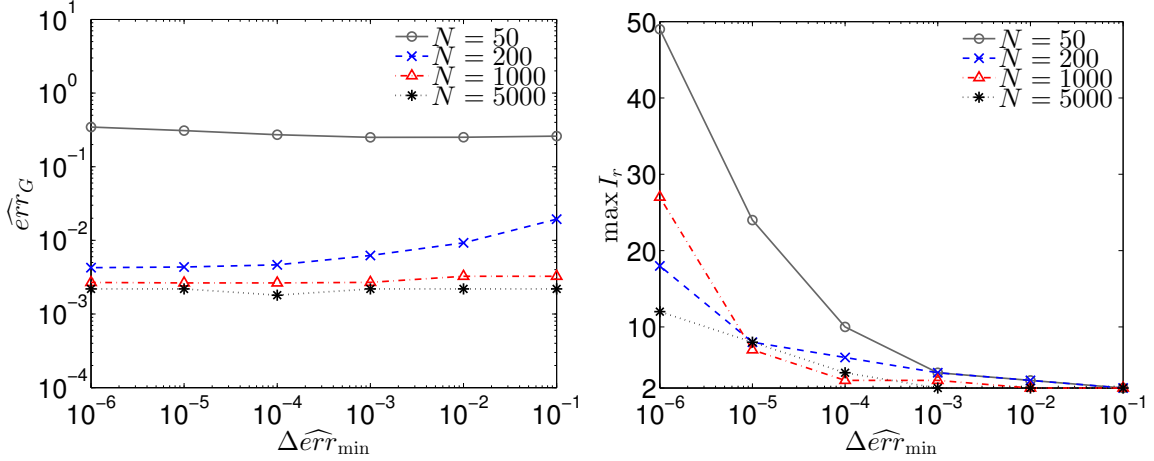


Figure 11: Frame-displacement problem: Relative generalization error (left) and corresponding maximum number of iterations in a correction step (right) versus threshold of differential error in the stopping criterion.

5.4 Heat conduction

The last example application (inspired by a problem studied in Nouy (2010)) concerns stationary heat conduction in the two-dimensional square domain $D = (-0.5, 0.5)\text{m} \times (-0.5, 0.5)\text{m}$ shown in the left graph of Figure 12. The temperature field $T(\mathbf{z})$, $\mathbf{z} \in D$, is described by the partial differential equation:

$$-\nabla(\kappa(\mathbf{z})\nabla T(\mathbf{z})) = I_A(\mathbf{z})Q, \quad (39)$$

with boundary conditions $T = 0$ on the top boundary and $\nabla T \cdot \mathbf{n} = 0$ on the left, right and bottom boundaries, where \mathbf{n} denotes the vector normal to the boundary. On the right side of Eq. (39), $Q = 2 \cdot 10^3 \text{W/m}^3$ and I_A is an indicator function equal to unity if $\mathbf{z} \in A$, where $A = (0.2, 0.3)\text{m} \times (0.2, 0.3)\text{m}$ is a square domain within D (see left graph of Figure 12). The thermal conductivity $\kappa(\mathbf{z})$ is a lognormal random field described by:

$$\kappa(\mathbf{z}) = \exp[a_\kappa + b_\kappa g(\mathbf{z})], \quad (40)$$

in which $g(\mathbf{z})$ represents a standard Gaussian random field with a square-exponential auto-correlation function:

$$\rho(\mathbf{z}, \mathbf{z}') = \exp(-\|\mathbf{z} - \mathbf{z}'\|^2/\ell^2). \quad (41)$$

In Eq. (40), the values of a_κ and b_κ are such that the mean and standard deviation of κ are $\mu_\kappa = 1 \text{ W/C} \cdot \text{m}$ and $\sigma_\kappa = 0.3 \text{ W/C} \cdot \text{m}$ respectively, while in Eq. (41), $\ell = 0.2 \text{ m}$. The response quantity of interest is the average temperature in the square domain $B = (-0.3, -0.2) \text{ m} \times (-0.3, -0.2) \text{ m}$ within D (see left graph of Figure 12):

$$\tilde{T} = \frac{1}{|B|} \int_{\mathbf{z} \in B} T(\mathbf{z}) d\mathbf{z}. \quad (42)$$

To solve Eq. (39), the Gaussian random field $g(\mathbf{z})$ is first discretized using the expansion optimal linear estimation (EOLE) method (Li and Der Kiureghian, 1993), as described next. Let $\{\zeta_1, \dots, \zeta_n\}$ denote the points of an appropriately defined grid in D . By retaining the first M terms in the EOLE series, $g(\mathbf{z})$ is approximated by:

$$\hat{g}(\mathbf{z}) = \sum_{i=1}^M \frac{\xi_i}{\sqrt{l_i}} \phi_i^T \mathbf{C}_{\mathbf{z}\zeta}(\mathbf{z}), \quad (43)$$

where $\{\xi_1, \dots, \xi_M\}$ are independent standard normal variables, $\mathbf{C}_{\mathbf{z}\zeta}$ is a vector with elements $\mathbf{C}_{\mathbf{z}\zeta}^{(k)} = \rho(\mathbf{z}, \zeta_k)$, for $k = 1, \dots, n$, and (l_i, ϕ_i) are the eigenvalues and eigenvectors of the correlation matrix $\mathbf{C}_{\zeta\zeta}$ with elements $\mathbf{C}_{\zeta\zeta}^{(k,l)} = \rho(\zeta_k, \zeta_l)$, for $k, l = 1, \dots, n$. Sudret and Der Kiureghian (2000) recommend that for a square-exponential autocorrelation function, the size of the element in the EOLE grid must be $1/2 - 1/3$ of ℓ . In the present numerical application, we use a square grid with element size 0.1 m ($1/2$ of ℓ), thus comprising $n = 121$ points. The number of terms in the EOLE series is determined according to the rule:

$$\sum_{i=1}^M l_i / \sum_{i=1}^n l_i \geq 0.99, \quad (44)$$

herein leading to $M = 53$. The shapes of the first 20 basis functions $\{\phi_i^T \mathbf{C}_{\mathbf{z}\zeta}(\mathbf{z}), i = 1, \dots, 20\}$ are shown in Figure 13.

The underlying deterministic problem is solved with an in-house finite-element analysis code developed in the Matlab environment. The domain D is discretized into 16,000 triangular T3 elements, as shown in the right graph of Figure 12, using the software *Gmsh* (Geuzaine and Remacle, 2009). The temperature field $T(\mathbf{z})$ for two realizations of the conductivity random field is depicted in Figure 14. Because the input vector $\mathbf{X} = \{\xi_1, \dots, \xi_M\}$ consists of independent standard normal variables, we develop LRA representations of the random response $\tilde{T} = \mathcal{M}(\mathbf{X})$ with basis functions made of Hermite polynomials.

5.4.1 Rank selection and error measures

In examining the accuracy of the two rank-selection criteria and associated error measures, we consider EDs of size varying between $N = 50$ and $N = 5,000$ and a validation set of size $n_{\text{val}} = 10^4$. After preliminary investigations, the maximum polynomial degree is set to $p_i = p = 2$ for $\{i = 1, \dots, 53\}$ and the stopping criterion in the correction step is defined by $I_{\text{max}} = 50$ and $\Delta \widehat{err}_{\text{min}} = 10^{-5}$. As in the previous examples, candidate ranks are selected among $\{1, \dots, 20\}$.

The left graph of Figure 15 shows boxplots of the ranks selected with 3-fold CV for 20 random partitions of each ED. The optimal rank, also indicated in the graph, is equal to unity in all cases. For $N \geq 200$, the 3-fold CV approach consistently selects the rank-1 meta-models in all trials. For the two smaller EDs however, it may erroneously select higher ranks, with the effect on the LRA accuracy depicted in the right graph of the same figure.





Figure 12
mesh (rig]

e-element




Figure 13: Heat-conduction problem: Shapes of the first 20 basis functions in the EOLE discretization (from left-top to bottom-right row-wise).

The LOO error criterion identifies the optimal rank $R^{\text{LOO}} = R^{\text{opt}} = 1$ for all considered EDs.

In Figure 16, we compare \widehat{err}_G with $\widehat{err}_{\text{CV3}}$ and $\widehat{err}_{\text{LOO}}$ for all candidate ranks and two example EDs of size $N = 100$ (left) and $N = 5,000$ (right). The $\widehat{err}_{\text{CV3}}$ errors for one example partition of the ED are shown. Note in these graphs that both error criteria identify the unity rank as optimal. In this example, $\widehat{err}_{\text{LOO}}$ can provide reasonably good estimates of \widehat{err}_G even for higher ranks; however $\widehat{err}_{\text{CV3}}$ overall outperforms $\widehat{err}_{\text{LOO}}$ in the estimation of \widehat{err}_G at $R^{\text{opt}} = 1$, which is herein evident in the case of $N = 100$.

5.4.2 Stopping criterion in the correction step

Finally, we investigate effects of the parameters of the stopping criterion on the LRA accuracy. The left graph of Figure 17 shows \widehat{err}_G for LRA with rank $R^{\text{opt}} = 1$ versus N , while the maximum allowable number of iterations takes the values $I_{\text{max}} = 1, 2, 10, 20$ and the differential error threshold is fixed to $\Delta\widehat{err}_{\text{min}} = 10^{-5}$. The right graph of the same figure




Figure 14: Heat-conduction problem: Example realizations of the temperature field.

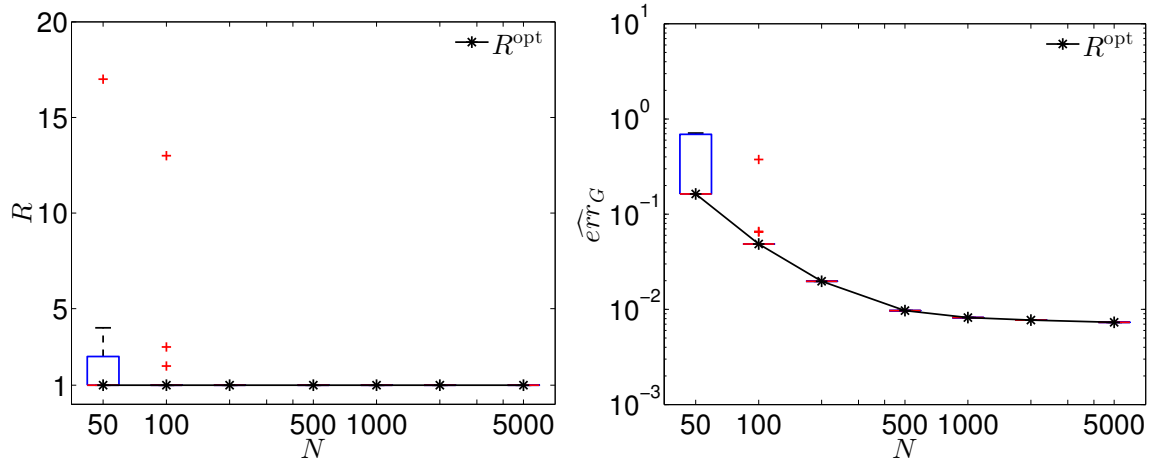


Figure 15: Heat-conduction problem: Comparison of ranks selected with 3-fold CV (20 replications) to optimal ranks (left) and corresponding relative generalization errors (right).

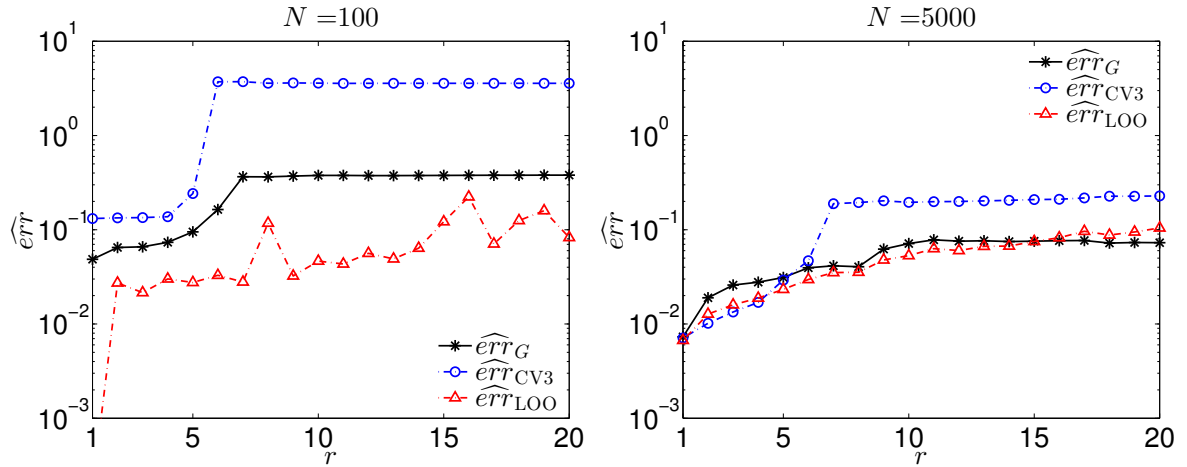


Figure 16: Heat-conduction problem: Comparison of error measures in rank-selection criteria with the relative generalization error based on the validation set.

shows \widehat{err}_G for LRA with rank $R^{\text{opt}} = 1$ versus N , while the differential error threshold takes the values $\Delta\widehat{err}_{\min} = 10^{-1}, 10^{-3}, 10^{-5}, 10^{-7}$ and the maximum allowable number of

iterations is fixed to $I_{\max} = 50$. Overall, both I_{\max} and $\Delta\widehat{err}_{\min}$ have a relatively small effect on the LRA accuracy, which becomes negligible for $N > 500$.




Figure 17: Heat-conduction problem: Relative generalization error for varying stopping criteria in the correction step versus size of experimental design.

5.5 Discussion on the construction of canonical low-rank approximations

In the above numerical investigations, the 3-fold CV error proves a more reliable criterion for rank selection as compared to the LOO error. The lack of orthogonality of the regressors may lead to excessive correction factors in the latter, particularly for higher ranks, whereas its uncorrected counterpart tends to be overly optimistic. On the other hand, the accuracy of the 3-fold CV error consistently improves with increasing size of the ED. For a certain ED, the median generalization error of LRA with rank selected via 3-fold CV is fairly close to the generalization error of the LRA with optimal rank. Divergence of the selected rank from the optimal one tends to have a smaller effect on the LRA accuracy as the ED size increases. In the examples involving finite-element models, EDs of size $N \geq 10 \cdot M$ lead to LRA with generalization errors of the order of 10^{-3} or smaller, even in cases with non-optimal rank selection. Note that this level of accuracy is typically sufficient in several meta-modeling applications, including sensitivity analysis (Konakli and Sudret, Konakli and Sudret). For the analytical beam-deflection function with underlying rank-1 structure, the same condition on the ED size leads to even higher meta-model accuracy.

Optimal values for the parameters in the stopping criterion of the correction step appear to strongly depend on the specific application. Overall, setting a sufficiently low differential error threshold $\Delta\widehat{err}_{\min}$ tends to be more critical when small EDs are considered ($N \leq 10 \cdot M$ in the examined applications). Based on the above results, we recommend the use of $\Delta\widehat{err}_{\min}$ values in the range $10^{-5} - 10^{-6}$ and caution that lower values may lead to numerical instabilities in certain cases. By imposing a simultaneous constraint on the maximum number of iterations I_{\max} , excessive unnecessary iterations are avoided. In the above case studies, setting $I_{\max} = 50$ allows the required number of iterations for achieving nearly the maximum or a sufficient meta-model accuracy. Note that depending on the application, the required number of iterations might be considerably smaller than that imposed by the recommended thresholds. Nevertheless, because these iterations involve minimization problems of only

small size, they are relatively inexpensive from a computational viewpoint. We underline that in a typical realistic meta-modeling application, the main computational effort lies in the evaluation of the model responses at the ED.

6 Canonical low-rank approximations versus sparse polynomial chaos expansions

We next confront canonical LRA to sparse PCE in the same meta-modeling applications considered in Section 5. The focus of the comparison is set on the applications involving finite-element models. The beam-deflection problem is only briefly examined in order to confirm the superior performance of LRA when the original model has a low-rank structure. The two types of meta-models are built using polynomials from the same family. Consistently with Section 5, Hermite polynomials are herein used in all four applications. Following the discussion above, we build LRA with rank $R^{\text{CV}3}$ (one random partition per ED is considered). A common polynomial degree is set in all dimensions, with its value $p \in \{1, \dots, 20\}$ also selected via 3-fold CV (one random partition per ED is considered). The associated ED-based error estimate of the meta-model is denoted by $\widehat{err}_{\text{CV}3}^{\text{LRA}}$. In building sparse PCE, we determine the candidate basis with a hyperbolic truncation scheme (see Section 3.2) and compute the coefficients with hybrid LAR (see Section 3.3). An optimal combination of the parameters $p^t \in \{1, \dots, 20\}$ and $q \in \{0.25, 0.50, 0.75, 1\}$ is selected in terms of the minimum corrected LOO error, hereafter denoted by $\widehat{err}_{\text{LOO}}^{\text{PCE}}$. The PCE computations are performed with the software UQLab (Marelli and Sudret, 2014, 2015). Comparisons between LRA and PCE are based on EDs obtained with Sobol sequences as well as Latin Hypercube Sampling (LHS). Of interest are cases with small ED sizes, which are typically encountered in real-life meta-modeling applications. Similarly to Section 5, large MCS validation sets are used to determine the actual errors of the meta-models. Because however such validation sets are not available in real-life applications, we further examine the accuracy of the ED-based error measures $\widehat{err}_{\text{CV}3}^{\text{LRA}}$ and $\widehat{err}_{\text{LOO}}^{\text{PCE}}$ in estimating the corresponding generalization errors $\widehat{err}_G^{\text{LRA}}$ and $\widehat{err}_G^{\text{PCE}}$.

6.1 Beam deflection

We herein develop LRA and PCE meta-models of the beam-deflection function described in Section 5.1, using EDs of size $30 \leq N \leq 500$ based on Sobol sequences. We assess the accuracy of the meta-models using a validation set of size $n_{\text{val}} = 10^6$. The left graph of Figure 18 compares the maximum polynomial degree p of the univariate polynomials in LRA with the maximum total polynomial degree p^t in PCE. In addition to the degree p selected by means of 3-fold CV, the graph also shows the actual optimal degree based on the validation set; the two values either agree or differ by one. The ED-based and generalization errors of the two types of meta-models are shown in the right graph of the figure. Note that the depicted errors of LRA correspond to the degree selected by means of 3-fold CV, *i.e.* we assess the accuracy of meta-models developed by using the information contained in the ED only. In the present case where the underlying model has a rank-1 structure, the LRA meta-models are 2-3 orders of magnitude more accurate than the PCE ones. It is noteworthy that EDs of size as small as $N = 30 - 50 \leq 10 \cdot M$ yield highly accurate LRA meta-models with errors of order $10^{-4} - 10^{-6}$. The generalization errors of both LRA and PCE are approximated fairly well by the corresponding ED-based measures.

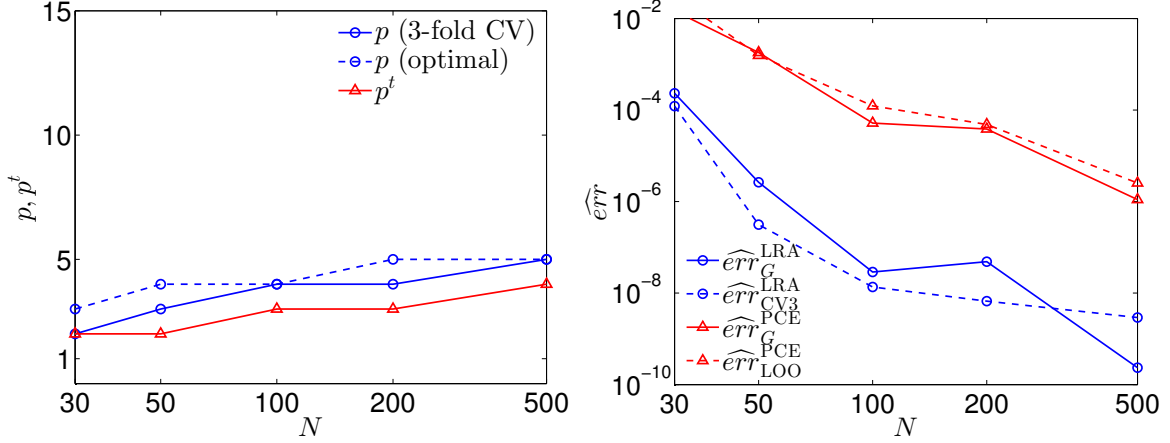


Figure 18: Beam deflection: Polynomial degrees (left) and corresponding error measures (right) of LRA and PCE meta-models based on Sobol sequences.

6.2 Truss deflection

We next assess the comparative accuracy of LRA and PCE in representing the truss-deflection model described in Section 5.2. To this end, we use EDs of size $30 \leq N \leq 500$, obtained with Sobol sequences as well as LHS, and a validation set of size $n_{\text{val}} = 10^6$. We first examine the results based on Sobol sequences, presented in Figure 19. Similarly to the previous example, the left graph shows the polynomial degrees, while the right graph shows the corresponding ED-based and generalization errors. Again, the selected and optimal degrees p of the LRA meta-models either coincide or differ by one. LRA are more accurate than PCE for the smaller EDs, but $\widehat{err}_G^{\text{PCE}}$ decreases faster than $\widehat{err}_G^{\text{LRA}}$ with increasing N , rendering PCE superior for the larger EDs. This trend is captured by the ED-based error measures, even though these are overall more pessimistic for PCE. The generalization errors of LRA and PCE meta-models built with LHS designs are shown in Figure 20. For each N , the corresponding generalization errors of the meta-models obtained with Sobol sequences are also shown for comparison reasons. The depicted boxplots correspond to 20 EDs with each representing the best among 5 random LHS designs, where the selection criterion is the maximum of the minimum distance between the points, so-called *maximin* LHS designs. Clearly, the PCE errors exhibit a larger dispersion than the LRA ones. The median errors of the LHS-based LRA are very close to the corresponding errors of the LRA based on Sobol sequences. In the case of PCE, the median errors for LHS designs are either similar or larger than the corresponding errors for Sobol sequences.

To gain further insight into the behaviors of the two types of meta-models, we plot in Figure 21 the meta-model versus the actual-model responses at the points of the validation set for the case of Sobol sequences with $N = 100$. Note in Figure 19 that the corresponding LRA and PCE are characterized by similar generalization errors ($\widehat{err}_G^{\text{LRA}} = 2.10 \cdot 10^{-3}$, $\widehat{err}_G^{\text{PCE}} = 2.56 \cdot 10^{-3}$). Obviously, LRA (left graph) provide better predictions of extreme responses, which are systematically underestimated by PCE (right graph). Because the extreme responses represent only a small fraction of the validation set, the observed differences have only a minor influence on the generalization error. In order to capture the meta-model performance in particular regions of interest, we introduce the *conditional generalization error*:

$$\widehat{Err}_G^{\text{C}} = \left\| \mathcal{M} - \widehat{\mathcal{M}} \right\|_{\mathcal{X}_{\text{val}}^{\text{C}}}^2. \quad (45)$$




Figure 19: Truss deflection: Polynomial degrees (left) and corresponding error measures (right) of LRA and PCE meta-models based on Sobol sequences.

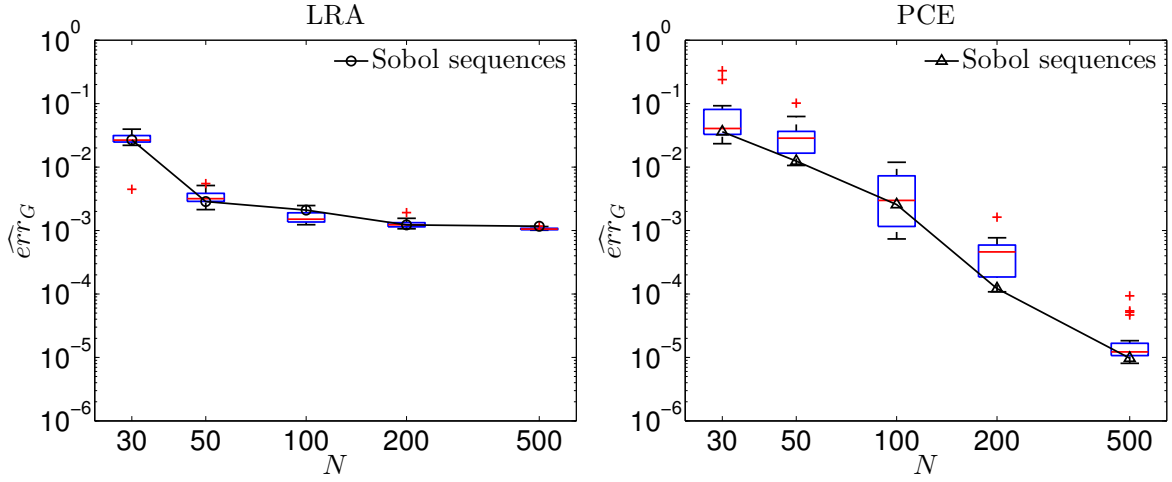


Figure 20: Truss deflection: Comparison of relative generalization errors of meta-models based on LHS (20 replications) to meta-models based on Sobol sequences.

The above equation indicates that \widehat{Err}_G^C is computed similarly to \widehat{Err}_G (see Eq. (4)), but by considering only a subset $\mathcal{X}_{\text{val}}^C$ of the validation set \mathcal{X}_{val} , defined by an appropriate condition. The subsets $\mathcal{X}_{\text{val}}^C$ of interest in the prediction of extreme responses are defined by:

$$\mathcal{X}_{\text{val}}^C = \{\mathbf{x} \in \mathcal{X}_{\text{val}} : Y = \mathcal{M}(\mathbf{x}) \geq y_{\text{lim}}\}. \quad (46)$$

The corresponding relative error \widehat{err}_G^C is obtained after normalization with the empirical variance of $\mathcal{Y}_{\text{val}}^C$, the latter denoting the set of model responses at $\mathcal{X}_{\text{val}}^C$. The left graph of Figure 22 shows the evolution of \widehat{err}_G^C versus the response threshold u_{lim} for the same meta-models considered in Figure 22. The conditional error increases faster for PCE than for LRA with the former being about an order of magnitude larger at the highest response threshold. The right graph of Figure 22 compares the LRA and PCE boxplots of \widehat{err}_G^C for the 20 LHS designs of size $N = 100$. Note again the faster increase of the median PCE error, which is about an order of magnitude larger than the median LRA error at the highest response threshold.




Figure 21: Truss deflection: Comparison of the exact model responses at the validation set with the respective responses of the LRA meta-model (left) and the PCE meta-model (right) for $N = 100$.




Figure 22: Truss deflection: Comparison between the LRA and PCE relative conditional generalization errors for EDs based on Sobol sequences (left) and LHS (right).

6.3 Frame displacement

For the frame-displacement problem, described in Section 5.3, we again examine the comparative accuracy of LRA and PCE by drawing EDs with Sobol sequences and LHS designs. The EDs are of size $100 \leq N \leq 2,000$, while the validation set comprises $n_{\text{val}} = 10^6$ points. Figures 23 and 24 respectively show results for EDs drawn with Sobol sequences and 20 maximin LHS designs (each being the best among 5 random ones), in a manner similar to Figures 19 and 20 in the previous example. The polynomial degrees for the case of Sobol sequences are shown in the left graph of Figure 23. The selected degrees for LRA match the optimal degrees except for $N = 200$. The increasing trend in the PCE degree with increasing ED size is interrupted at $N = 1,000$ due to the change of the optimal truncation parameter q from 0.50 to 1.0. Note in Figure 23 that similarly to the truss-deflection problem, LRA exhibit smaller errors than PCE for the smaller EDs, but the PCE errors decrease faster with increasing ED size. The LRA generalization error is rather accurately estimated with the 3-fold CV approach, whereas the PCE generalization error is slightly underestimated by

the LOO error. Figure 24 shows that the LRA errors for the LHS designs exhibit a slightly smaller dispersion than the PCE ones. In all cases, the median errors for the LHS designs are fairly close to the corresponding errors for the Sobol sequences.




Figure 23: Frame displacement: Polynomial degrees (left) and corresponding error measures (right) of LRA and PCE meta-models based on Sobol sequences.

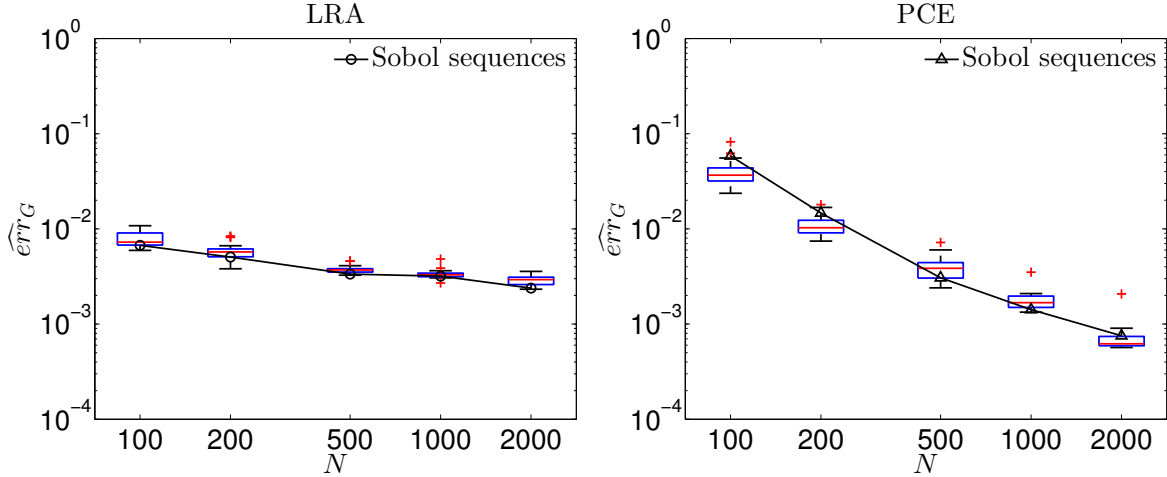


Figure 24: Frame displacement: Comparison of relative generalization errors of meta-models based on LHS (20 replications) to meta-models based on Sobol sequences.

Note in Figures 23 and 24 that the two types of meta-models exhibit similar generalization errors at $N = 500$. However, as in the truss-deflection problem, LRA outperform PCE in the prediction of extreme responses. This can be observed in Figure 25, which shows the meta-model versus the actual-model responses at the validation set for an example LHS design of size $N = 500$. The superior performance of LRA at the upper tail of the response distribution is not reflected on the generalization errors ($\widehat{err}_G^{LRA} = 3.68 \cdot 10^{-3}$, $\widehat{err}_G^{PCE} = 3.47 \cdot 10^{-3}$), but can be captured by the conditional generalization error (Eq. (45)-(46)). For the meta-models developed with EDs of size $N = 500$, Figure 26 depicts the evolution of \widehat{err}_G^C with increasing response threshold; the left graph of the figure shows this error for PCE and LRA based on Sobol sequences, while the right graph shows respective boxplots for the 20 LHS designs. In both graphs, we observe that the conditional errors of LRA and PCE are similar for the

lower responses thresholds, but the PCE errors are larger for the higher ones.




Figure 25: Frame displacement: Comparison of the exact model responses at the validation set with the respective responses of the LRA meta-model (left) and the PCE meta-model (right) for $N = 500$.




Figure 26: Frame displacement: Comparison between the LRA and PCE relative conditional generalization errors for EDs based on Sobol sequences (left) and LHS (right).

6.4 Heat conduction

In this last example, we consider again the heat-conduction problem investigated in Section 5.4. We develop LRA and PCE meta-models for the average temperature \tilde{T} in domain B (see Figure 12) using EDs of size $100 \leq N \leq 2,000$ based on Sobol sequences. To assess the comparative accuracy of the meta-models, we use a validation set of size $n_{\text{val}} = 10^4$. The left and right graphs of Figure 27 respectively show the polynomial degrees and error measures for the two types the meta-models. Note that the optimal LRA degree p is accurately identified by 3-fold CV for all considered EDs. The increasing trend in the PCE degree p^t with increasing N is interrupted at $N = 1,000$ due to the change of the optimal truncation parameter q from 0.50 to 0.75. The errors demonstrate similar trends with those observed for the previously examined finite-element models, *i.e.* $\widehat{err}_G^{\text{LRA}}$ is smaller than $\widehat{err}_G^{\text{PCE}}$ for

the smaller EDs, but the latter decreases faster with increasing N . For both types of meta-models, the generalization errors are estimated fairly well by the corresponding ED-based measures.




Figure 27: Heat conduction: Polynomial degrees (left) and corresponding error measures (right) of LRA and PCE meta-models based on Sobol sequences.

For the ED of size $N = 500$, the generalization errors of LRA and PCE are nearly equal ($\widehat{err}_G^{\text{LRA}} = 9.68 \cdot 10^{-3}$, $\widehat{err}_G^{\text{PCE}} = 9.15 \cdot 10^{-3}$); however, as in the previous examples, LRA provide better predictions of the extreme responses. The left and right graphs of Figure 28 respectively show the LRA and PCE responses versus the actual model responses at the validation set. The PCE predictions of the extreme responses exhibit a negative bias, which does not appear in the LRA predictions. The lower accuracy of the PCE meta-model at the tail of the response distribution is captured by the conditional generalization error, which is plotted in Figure 29 for varying response thresholds.




Figure 28: Heat conduction: Comparison of the exact model responses at the validation set with the respective responses of the LRA meta-model (left) and the PCE meta-model (right) for $N = 500$.




Figure 29: Heat conduction: Comparison between the LRA and PCE relative conditional generalization errors.

7 Conclusions

In this study, a newly emerged class of meta-models based on canonical low-rank approximations (LRA) was confronted to the popular sparse polynomial chaos expansions (PCE). Both meta-modeling approaches hold strong promise against the curse of dimensionality that often poses a major challenge in real-life uncertainty propagation problems. We examined LRA and PCE meta-models developed in a non-intrusive manner by relying on the same polynomial families to build the basis functions; the polynomial coefficients in both meta-model types were obtained by solving error-minimization problems. In the considered sparse PCE approach, a single large-size minimization problem is efficiently solved by retaining only the significant basis terms, as identified with the least angle regression algorithm. In the canonical LRA approach, the tensor-product form of the basis is retained, leading to a series of small-size minimization problems that were herein solved using ordinary least-squares (OLS).

We first shed light on issues pertinent to the construction of LRA through extensive numerical investigations. To this end, we considered a particular greedy algorithm comprising a series of pairs of a correction and an updating step. In a correction step, a rank-one component is built based on the sequential updating of the polynomial coefficients in different dimensions, whereas in an updating step, the coefficients of the new set of rank-one components are determined. Two approaches for rank selection were examined; the first relies on error estimation with 3-fold cross-validation (CV), whereas the second, based on a simpler computation, uses a corrected version of the leave-one-out (LOO) error in the updating step. The lack of orthogonality of the regressors led to excessive correction factors of the LOO error in certain cases, particularly for higher ranks. On the other hand, the 3-fold CV error was found an overall reliable estimator of the generalization error, also appropriate for the selection of the polynomial degree. Finally, a stopping criterion for the correction step was examined, combining a threshold on the differential empirical error in two successive iterations with a maximum allowable iteration number. Effects of the two thresholds on the LRA accuracy were investigated and appropriate values were proposed.

The comparative accuracy of the particular LRA and PCE approaches was investigated in three problems involving finite-element models of varying dimensionality M . Experimental designs (EDs) of varying sizes were drawn with Sobol sequences and Latin hypercube

sampling (LHS). The errors obtained by using Sobol sequences as well as the median errors obtained by using LHS designs were lower for LRA when small EDs (approximately up to size $10M - 20M$) were considered. However, the PCE errors decreased faster with increasing ED size. In cases when the two types of meta-models exhibited similar generalization errors, LRA provided better estimates of the extreme responses. The superiority of LRA in predicting extreme responses was quantified by means of the introduced conditional generalization error. This finding renders LRA particularly promising for problems where the accuracy in estimating the tail of the response distribution is important, such as reliability applications (Konakli and Sudret, 2015b,a). It is further emphasized that by relying on a series of OLS minimizations, the construction of LRA involved simpler computations compared to PCE. Because of the small size of these minimization problems, the LRA construction also required far less computer memory.

References

- Abramowitz, M. and I. Stegun (1970). *Handbook of mathematical functions*. Dover Publications, Inc.
- Acar, E., S. A. Camtepe, and B. Yener (2006). Collective sampling and analysis of high order tensors for chatroom communications. In *Intelligence and security informatics*, pp. 213–224. Springer.
- Allen, D. (1971). *The prediction sum of squares as a criterion for selecting predictor variables*. Number 23. Technical report, Dept. of Statistics, University of Kentucky.
- Ammar, A., B. Mokdad, F. Chinesta, and R. Keunings (2006). A new family of solvers for some classes of multidimensional partial differential equations encountered in kinetic theory modeling of complex fluids. *J. Non-Newton. Fluid* 139(3), 153–176.
- Andersen, A. H. and W. S. Rayens (2004). Structure-seeking multilinear methods for the analysis of fMRI data. *NeuroImage* 22(2), 728–739.
- Appelhof, C. J. and E. Davidson (1981). Strategies for analyzing data from video fluorometric monitoring of liquid chromatographic effluents. *Anal. Chem.* 53(13), 2053–2056.
- 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.
- Beylkin, G., J. Garcke, and M. J. Mohlenkamp (2009). Multivariate regression and machine learning with sums of separable functions. *SIAM J. Sci. Comput.* 31(3), 1840–1857.
- 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 (2008). Sparse polynomial chaos expansions and adaptive stochastic finite elements using a regression approach. *Comptes Rendus Mécanique* 336(6), 518–523.
- Blatman, G. and B. Sudret (2010). An adaptive algorithm to build up sparse polynomial chaos expansions for stochastic finite element analysis. *Prob. Eng. Mech.* 25(2), 183–197.

- Blatman, G. and B. Sudret (2011). Adaptive sparse polynomial chaos expansion based on least angle regression. *J. Comput. Phys.* 230, 2345–2367.
- Bro, R. (1997). PARAFAC. tutorial and applications. *Chemometr. Intell. Lab.* 38(2), 149–171.
- Carroll, J. D. and J.-J. Chang (1970). Analysis of individual differences in multidimensional scaling via an n-way generalization of “Eckart-Young” decomposition. *Psychometrika* 35(3), 283–319.
- Chapelle, O., V. Vapnik, and Y. Bengio (2002). Model selection for small sample regression. *Mach. Learn.* 48(1), 9–23.
- Chevreuril, M., R. Lebrun, A. Nouy, and P. Rai (2015). A least-squares method for sparse low rank approximation of multivariate functions. *SIAM/ASA J. Uncertainty Quantification* 3(1), 897–921.
- Chevreuril, M., P. Rai, and A. Nouy (2013). Sampling based tensor approximation method for uncertainty propagation. In *Proc. 11th Int. Conf. on Structural Safety and Reliability (ICOSSAR2013)*, New York.
- Choi, S., R. Grandhi, R. Canfield, and C. Pettit (2004). Polynomial chaos expansion with Latin Hypercube sampling for estimating response variability. *AIAA Journal* 45, 1191–1198.
- De Lathauwer, L. and J. Castaing (2007). Tensor-based techniques for the blind separation of DS-CDMA signals. *Signal Process.* 87(2), 322–336.
- Doostan, A. and H. Owhadi (2011). A non-adapted sparse approximation of PDEs with stochastic inputs. *J. Comput. Phys.* 230(8), 3015–3034.
- Doostan, A., A. Validi, and G. Iaccarino (2013). Non-intrusive low-rank separated approximation of high-dimensional stochastic models. *Comput. Methods Appl. Mech. Eng.* 263, 42–55.
- Efron, B., T. Hastie, I. Johnstone, and R. Tibshirani (2004). Least angle regression. *Ann. Stat.* 32, 407–499.
- Felippa, C. and R. Ohayon (1990). Mixed variational formulation of finite element analysis of acoustoelastic/slosh fluid-structure interaction. *J. Fluid. Struct.* 4(1), 35–57.
- Furukawa, R., H. Kawasaki, K. Ikeuchi, and M. Sakauchi (2002). Appearance based object modeling using texture database: Acquisition compression and rendering. In *Rendering Techniques*, pp. 257–266.
- Geuzaine, C. and J.-F. Remacle (2009). Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. *Int. J. Numer. Meth. Eng.* 79(11), 1309–1331.
- Ghanem, R. G. and P. D. Spanos (2003). *Stochastic finite elements: a spectral approach*. Courier Corporation.
- Hadigol, M., A. Doostan, H. G. Matthies, and R. Niekamp (2014). Partitioned treatment of uncertainty in coupled domain problems: A separated representation approach. *Comput. Methods Appl. Mech. Eng.* 274, 103–124.

- Harshman, R. A. (1970). *Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis*. University of California at Los Angeles Los Angeles.
- Hitchcock, F. (1927). The expression of a tensor or a polyadic as a sum of products. *J. Math. Phys.* 6, 164–189.
- Kolda, T. G. and B. W. Bader (2009). Tensor decompositions and applications. *SIAM Rev.* 51(3), 455–500.
- Konakli, K. and B. Sudret. Global sensitivity analysis using low-rank tensor approximations. *Reliab. Eng. Sys. Safety* (submitted).
- Konakli, K. and B. Sudret (2015a). Addressing high dimensionality in reliability analysis using low-rank tensor approximations. In *Proc. annual European Safety and Reliability Conf. (ESREL2015), Zürich, Switzerland*.
- Konakli, K. and B. Sudret (2015b). Low-rank tensor approximations for reliability analysis. In *Proc. 12th Int. Conf. on Applications of Statistics and Probability in Civil Engineering (ICASP12), Vancouver, Canada*.
- Kruskal, J. B. (1977). Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear Algebra Appl.* 18(2), 95–138.
- Li, C.-C. and A. Der Kiureghian (1993). Optimal discretization of random fields. *J. Eng. Mech.* 119(6), 1136–1154.
- Liu, P.-L. and A. Der Kiureghian (1991). Optimization algorithms for structural reliability. *Structural Safety* 9, 161–177.
- Marelli, S. and B. Sudret (2014). UQLab: a framework for uncertainty quantification in MATLAB. In *Proc. 2nd Int. Conf. on Vulnerability, Risk Analysis and Management (ICVRAM2014), Liverpool, United Kingdom*.
- Marelli, S. and B. Sudret (2015). Uqlab user manual – polynomial chaos expansions. Technical report, Chair of Risk, Safety & Uncertainty Quantification, ETH Zurich. Report # UQLab-V0.9-104.
- Mathelin, L. (2014). Quantification of uncertainty from high-dimensional scattered data via polynomial approximation. *Int. J. Uncertainty Quantification* 4(3).
- Mocks, J. (1988). Topographic components model for event-related potentials and some biophysical considerations. *IEEE Trans. Bio-Med. Eng.* 35(6), 482–484.
- Nelsen, R. (2006). *An introduction to copulas* (2nd ed.), Volume 139 of *Lecture Notes in Statistics*. Springer-Verlag, New York.
- Niederreiter, H. (1992). *Random number generation and quasi-Monte Carlo methods*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.
- Nouy, A. (2010). Proper generalized decompositions and separated representations for the numerical solution of high dimensional stochastic problems. *Arch. Comput. Meth. Eng.* 17, 403–434.

- Rai, P. (2014). *Sparse Low Rank Approximation of Multivariate Functions – Applications in uncertainty quantification*. Ph. D. thesis, Engineering Sciences [physics]. Ecole Centrale Nantes.
- Shashua, A. and A. Levin (2001). Linear image coding for regression and classification using the tensor-rank principle. In *Proc. 2001 IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (CVPR)*, Volume 1, pp. I–42. IEEE.
- Sidiropoulos, N. D., R. Bro, and G. B. Giannakis (2000). Parallel factor analysis in sensor array processing. *IEEE Trans. Signal Proces.* 48(8), 2377–2388.
- 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 element methods and reliability: a state-of-the-art report*. Department of Civil and Environmental Engineering, University of California.
- Validi, A. (2014). Low-rank separated representation surrogates of high-dimensional stochastic functions: Application in Bayesian inference. *J. Comput. Phys.* 260, 37–53.
- Xiu, D. (2009). Fast numerical methods for stochastic computations: a review. *Comm. Comput. Phys.* 5(2-4), 242–272.
- Xiu, D. and G. E. Karniadakis (2002). The Wiener–Askey polynomial chaos for stochastic differential equations. *SIAM J. Sci. Comput.* 24(2), 619–644.