

ON OPTIMAL EXPERIMENTAL DESIGNS FOR SPARSE POLYNOMIAL CHAOS EXPANSIONS

N. Fajraoui, S. Marelli, B. Sudret



Data Sheet

Journal: -

Report Ref.: RSUQ-2017-001

Arxiv Ref.: <http://arxiv.org/abs/1703.05312> - [stat.CO]

DOI: -

Date submitted: 14/11/2016

Date accepted: -

On optimal experimental designs for sparse polynomial chaos expansions

N. Fajraoui¹, S. Marelli¹, and B. Sudret¹

¹*Chair of Risk, Safety and Uncertainty Quantification,
ETH Zurich, Stefano-Frascini-Platz 5, 8093 Zurich, Switzerland*

Abstract

Uncertainty quantification (UQ) has received much attention in the literature in the past decade. In this context, Sparse Polynomial chaos expansions (PCE) have been shown to be among the most promising methods because of their ability to model highly complex models at relatively low computational costs. A least-square minimization technique may be used to determine the coefficients of the sparse PCE by relying on the so called experimental design (ED), *i.e.* the sample points where the original computational model is evaluated. An efficient sampling strategy is then needed to generate an accurate PCE at low computational cost. This paper is concerned with the problem of identifying an optimal experimental design that maximizes the accuracy of the surrogate model over the whole input space within a given computational budget. A novel sequential adaptive strategy where the ED is enriched sequentially by capitalizing on the sparsity of the underlying metamodel is introduced. A comparative study between several state-of-the-art methods is performed on four numerical models with varying input dimensionality and computational complexity. It is shown that the optimal sequential design based on the S-value criterion yields accurate, stable and computationally efficient PCE.

Keywords: Optimal Experimental Design – Sparse Polynomial Chaos Expansions – Sequential Experimental design

1 Introduction

Prediction of the behavior of complex physical and engineering systems is affected by diverse types of uncertainties, including imperfect knowledge of the system parameters and their variability. These uncertainties need to be appropriately considered and their impact on the quality

of model predictions needs to be quantified in a rigorous way. In this context, uncertainty quantification (UQ) is a flexible framework within which parametric uncertainties can be effectively tackled. However, UQ may become intractable in the presence of expensive-to-evaluate computational models. In this context, surrogate modelling techniques have been increasingly adopted to enable complex uncertainty quantification analyses that would otherwise be impossible.

Amongst the available surrogate modeling techniques, polynomial chaos expansions (PCE) are one of the most promising methods for uncertainty quantification (Ghanem and Spanos, 2003; Xiu and Hesthaven, 2005). In PCE, the model response is expanded in terms of a set of orthonormal multivariate polynomials that are orthogonal with respect to a suitable probability measure. They allow one to uncover the relationships between different input parameters and how they affect the model responses. Among the available strategies to construct PCE approximations, this paper focuses on non-intrusive least square minimization and sparse-regression-based approaches, meaning that the computational model is considered as a black-box model. They have been extensively proven to be particularly effective in the presence of complex models (Blatman and Sudret, 2011a). The least-square calculation of the PCE coefficients relies on the so-called *experimental design* (ED), a sample from the joint distribution of the input parameters, and the corresponding model responses. A common property of least-square techniques is that their accuracy and efficiency (in terms number of required full-model evaluations) depends on how the *experimental design* is sampled. In this context, different sampling strategies have been extensively investigated in the literature (Giunta et al., 2003; Goel et al., 2008; Simpson et al., 2001,?). This is the focus of the design of experiments (DOE) field of research.

Two different classes of DOE sampling strategies are of particular interest in this paper. The first class involves space-filling designs that aim at filling the space of the input parameters uniformly according to some appropriate measure. Several authors have shown that space filling designs are well suited to perform metamodelling (Fang et al., 2005; Simpson et al., 2001). Among such designs, Latin hypercube sampling (LHS) methods have become very popular. An oversampling of $2 \sim 3$ times the number of unknown coefficients is recommended to get reliable results. These designs do not satisfy any specific optimality measure.

The second class of interest in this paper comprises optimal experimental design strategies which rely on existing prior knowledge about the surrogate model properties. They aim at maximizing the quality of a given metamodel, with respect to some optimality criteria. A number of optimality criteria have been proposed and discussed in the literature in the context of regression applications (Myers and Montgomery, 2002). Among recent examples involving optimal experimental designs in the context of metamodelling Shin and Xiu (2016) proposed a simple technique to generate efficiently quasi-optimal experimental designs for ordinary least squares (OLS) regression. It is an optimization algorithm that aims at finding the design of

experiments of fixed size N such that the obtained results are as close as possible to those obtained with a larger sample, prior to performing model evaluations. The optimality criterion is related to the orthogonality of the columns of the information matrix. An alternative method was introduced in Zein et al. (2013), where a D-optimal criterion is optimized to make the PCE construction computationally feasible.

In this paper, we are interested in selecting the best design of experiments in order to ensure the accuracy of the surrogate model over the whole input parameter space with a minimum number of evaluations of the computational model. Furthermore, recent literature has shown that it is more appropriate to build up the ED sequentially by iteratively enriching an initial ED, especially when model evaluations are time consuming. Recently, in metamodeling context, sequential sampling methodologies have been identified as a more efficient strategy than one-stage sampling methods, where sample points are generated all at once (Jin et al., 2002; Lin et al., 2008). dos Santos and dos Santos (2007) proposed a sequential design to improve the overall accuracy of nonlinear regression metamodels for one-dimensional and two-dimensional inputs. The approach improves the accuracy of the resulting metamodel compared to space filling designs. Crombecq et al. (2011) performed a comparison and analysis of different space filling sequential strategies for simulation-based modelling. A review of sequential experimental designs for metamodeling in engineering can be found in Jin et al. (2002).

In this paper, we propose a novel adaptive design strategy for PCE that iteratively updates the experimental design by capitalizing on the sparsity of the underlying metamodel. The idea is to generate new sample points according to an optimality criterion based on the current metamodel estimate, so as to improve its accuracy. The new sampling points are then added to the ED and the metamodel is updated.

The outline of the paper is as follows. In Section 2, sparse polynomial chaos expansions are introduced. In Section 3, several commonly adopted sampling strategies are reviewed. In Section 4, several sequential design strategies are reviewed. In Section 5, a novel sequential approach optimal for sparse PCE is proposed. Finally, in Section 6, a comprehensive comparative study of the different sampling strategies is performed on four numerical models with varying input dimensionality and computational complexity.

2 Sparse Polynomial Chaos Expansions

2.1 Polynomial Chaos Expansions

Let us consider a physical system whose behaviour is represented by a computational model \mathcal{M} . We suppose that this model is a function of M uncertain input parameters that are represented by independent random variables $\mathbf{X} = \{X_1, \dots, X_M\}$ with marginal probability density functions

PDFs $\{f_{X_i}(x_i), i = 1, \dots, M\}$. As a consequence, the model scalar response denoted by Y is also a random variable:

$$Y = \mathcal{M}(\mathbf{X}). \quad (1)$$

Provided that the output random variable Y has a finite variance, it can be represented as a polynomial chaos expansion as follows:

$$Y = \mathcal{M}(\mathbf{X}) = \sum_{\boldsymbol{\alpha} \in \mathbb{N}^M} y_{\boldsymbol{\alpha}} \Psi_{\boldsymbol{\alpha}}(\mathbf{X}). \quad (2)$$

where $y_{\boldsymbol{\alpha}} \in \mathbb{R}$ are the expansion coefficients to be determined, $\Psi_{\boldsymbol{\alpha}}(\mathbf{X})$ are multivariate polynomials, $\boldsymbol{\alpha} \in \mathbb{N}^M$ is a multi-index $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_M\}$ that identifies the degree of the multivariate polynomials $\Psi_{\boldsymbol{\alpha}}$ in each of the input variables X_i . The multivariate polynomials $\Psi_{\boldsymbol{\alpha}}$ are assembled as the tensor product of univariate polynomials orthogonal w.r.t. the marginal PDF of the corresponding variable, *i.e.*:

$$\Psi_{\boldsymbol{\alpha}}(\mathbf{x}) = \prod_{i=1}^M \phi_{\alpha_i}^{(i)}(x_i). \quad (3)$$

with

$$\langle \phi_{\alpha}^{(i)}, \phi_{\beta}^{(i)} \rangle \stackrel{\text{def}}{=} \int_{\mathcal{D}_{X_i}} \phi_{\alpha}^{(i)}(x_i) \times \phi_{\beta}^{(i)}(x_i) f_{X_i}(x_i) dx_i = \delta_{\alpha\beta}. \quad (4)$$

where $\phi_{\alpha_i}^{(i)}$ is an orthogonal polynomial in the i -th variable of degree α_i and $\delta_{\alpha\beta} = 1$ if $\alpha = \beta$ and $\delta_{\alpha\beta} = 0$ otherwise. Due to this construction, it follows that the multivariate polynomials $\Psi_{\boldsymbol{\alpha}}$ are orthonormal w.r.t. the input random vector \mathbf{X} . For instance, if the input random variable \mathbf{X} are standard normal, a possible basis is the family of multivariate Hermite polynomials, which are orthogonal with respect to the Gaussian measure. Other common distributions can be used together with basis functions from the Askey scheme. A more general case can be treated through an isoprobabilistic transformation of the input random vector \mathbf{X} into a standard random vector. For computational purposes, the series in Eq. (2) is truncated after a finite number of terms P , yielding the following PC approximation:

$$Y^{PCE} = \sum_{\boldsymbol{\alpha} \in \mathcal{A}} y_{\boldsymbol{\alpha}} \Psi_{\boldsymbol{\alpha}}(\mathbf{X}) \quad (5)$$

where \mathcal{A} is a finite set of multi-indices of cardinality P . A standard truncation scheme, which corresponds to selecting all polynomials in the M input variables of total degree not exceeding p , is commonly used:

$$\mathcal{A} = \{\boldsymbol{\alpha} \in \mathbb{N}^M : |\boldsymbol{\alpha}| \leq p\}. \quad (6)$$

The cardinality of the set \mathcal{A} reads:

$$\text{card } \mathcal{A} \equiv P = \binom{M+p}{p}. \quad (7)$$

The next step is the computation of the polynomial chaos coefficients $\mathbf{y} = \{y_{\boldsymbol{\alpha}}, \boldsymbol{\alpha} \in \mathcal{A}\}$. Several intrusive approaches (*e.g.* Galerkin scheme) or non-intrusive approaches (*e.g.* projection,

least square regression) (Sudret, 2008; Xiu, 2010) have been proposed in the literature. We herein focus our analysis on least-square methods, originally introduced by Berveiller et al. (2006); Choi et al. (2004) under the generic name regression. In this approach, the exact expansion can be written as the sum of a truncated series and a residual:

$$Y = \sum_{\alpha \in \mathcal{A}} y_{\alpha} \Psi_{\alpha}(\mathbf{X}) + \epsilon. \quad (8)$$

where ϵ corresponds to the truncated terms.

The term non-intrusive indicates that the polynomial chaos coefficients are investigated over a set of input realizations $\mathcal{X} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}\}$, referred to as the experimental design (ED). One can employ Monte Carlo sampling methods, Latin Hypercube sampling (LHS, see McKay et al. (1979)), or quasi-random sequences such as the Sobol' or Halton sequence to sample points of the ED. Denoting by $\mathcal{Y} = \{\mathcal{M}(\mathbf{x}^{(1)}), \dots, \mathcal{M}(\mathbf{x}^{(N)})\}$ the set of outputs of the computational model \mathcal{M} for each point in the ED \mathcal{X} , the set of coefficients $\mathbf{y} = \{y_{\alpha}, \alpha \in \mathcal{A}\}$ is then computed by minimizing the least square residual of the polynomial approximation over the ED \mathcal{X} :

$$\hat{\mathbf{y}} = \underset{\mathbf{y}_{\alpha} \in \mathbb{R}^P}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^N \left(\mathcal{M}(\mathbf{x}^{(i)}) - \sum_{\alpha \in \mathcal{A}} y_{\alpha} \Psi_{\alpha}(\mathbf{x}^{(i)}) \right)^2. \quad (9)$$

The ordinary least-square solution of Eq. (9) can be formulated equivalently in matrix notation as follows:

$$\hat{\mathbf{y}} = (\mathbf{A}^{\top} \mathbf{A})^{-1} \mathbf{A}^{\top} \mathcal{Y}. \quad (10)$$

where

$$\mathbf{A} = \{A_{ij} = \psi_j(\mathbf{x}^{(i)}), i = 1, \dots, N; j = 0, \dots, P\}.$$

in which \mathbf{A} is the model matrix that contains the values of all the polynomial basis evaluated at the experimental design points.

Solving Eq. (10) requires a minimum number of sample points. One typically selects $N > \operatorname{card} \mathcal{A}$ to ensure the well-posedness of the regression problem. As a rule of thumb, the number of samples is often chosen such that $N \approx 2 - 3 \operatorname{card} \mathcal{A}$. Due to the sharp increase in the number of terms in Eq. (5) (see Eq. (7)), the required number of model evaluations becomes prohibitive when the polynomial degree p is large. The truncation scheme proposed in Eq. (6) may thus lead to unaffordable calculations especially in high dimensions. This problem is known as the curse of dimensionality. As a way to mitigate this limitation, a so-called *hyperbolic* truncation scheme was proposed in Blatman and Sudret (2010). It corresponds to selecting all multi-indices that satisfy:

$$\mathcal{A}^{M,p,q} = \{\alpha \in \mathcal{A} : \|\alpha\|_q \leq p\}, \quad (11)$$

where:

$$\|\boldsymbol{\alpha}\| = \left(\sum_{i=1}^M \alpha_i^q \right)^{1/q}, \quad (12)$$

and $0 < q < 1$. Lower values of q limit the number of high order interaction terms, leading to *sparse* solutions. The value $q = 1$ corresponds to the standard truncation shown in Eq. (6).

2.2 Sparse Polynomial Chaos Expansions

To further sparsify the solution, without sacrificing possibly important interaction terms, Blatman and Sudret (2011a) proposed an adaptive sparse PCE calculation strategy based on the least angle regression (LAR) algorithm (Efron et al., 2004). Based on a given experimental design, only regressors that have the greatest impact on the model response are retained. LAR consists in directly modifying the least-square minimization problem in Eq. (9) by adding a penalty term of the form $\lambda \|\mathbf{y}\|_1$, *i.e.* solving:

$$\mathbf{y} = \underset{\mathbf{y}_{\boldsymbol{\alpha}} \in \mathbb{R}^P}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^N \left(\mathcal{M}(\mathbf{x}^{(i)}) - \sum_{\boldsymbol{\alpha} \in \mathcal{A}} y_{\boldsymbol{\alpha}} \Psi_{\boldsymbol{\alpha}}(\mathbf{x}^{(i)}) \right)^2 + \lambda \|\mathbf{y}\|_1, \quad (13)$$

in which $\|\mathbf{y}\|_1 = \sum_{\boldsymbol{\alpha} \in \mathcal{A}} |y_{\boldsymbol{\alpha}}|$ is the regularization term that forces the minimization to favour low-rank solutions. In this study, we adopt the hybrid Least Angle Regression (LAR) method for building sparse PCE. For further details, the reader is referred to Blatman and Sudret (2011b).

2.3 Measures of accuracy

Once the metamodel is built, it is important to assess its quality. A natural measure is provided by the mean square error *MSE*, defined as:

$$MSE = \mathbb{E} [(\mathcal{M}(\mathbf{X}) - \mathcal{M}^{PC}(\mathbf{X}))^2] \approx \frac{1}{N_{val}} \sum_{i=1}^{N_{val}} \left(\mathcal{M}(\mathbf{x}^{(i)}) - \mathcal{M}^{PC}(\mathbf{x}^{(i)}) \right)^2 \quad (14)$$

where the validation set $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N_{val})}\}$ consists of samples of the input vector that do not belong to the experimental design and \mathcal{M}^{PC} is the surrogate model. In practice, one needs a sufficiently large validation set (typically in the order of $N_{val} \sim 10^{4-5}$) to achieve a stable estimate of the *MSE*. A low *MSE* corresponds to an accurate metamodel. A common variant of the MSE error in Eq. (14) is the relative mean square error *RMSE*, defined as:

$$RMSE = \frac{MSE}{\sigma_Y^2}, \quad (15)$$

where σ_Y^2 is the variance of the model response or its estimator:

$$\sigma_Y^2 = \frac{1}{N_{val}} \sum_{i=1}^{N_{val}} \left(\mathcal{M}(\mathbf{x}^{(i)}) - \mu_{\mathcal{M}} \right)^2, \quad \mu_{\mathcal{M}} = \frac{1}{N_{val}} \sum_{i=1}^{N_{val}} \mathcal{M}(\mathbf{x}^{(i)}). \quad (16)$$

In practice, this error is only affordable for analytical functions, since it requires evaluating a large number of model responses. As an alternative, an error estimate based on the ED may be used instead. In particular, the leave-one-out (LOO) error, denoted as ϵ_{LOO} , has been proposed (Geisser, 1975; Stone, 1974) in the context of PCE. The LOO error is a cross-validation technique developed in statistical learning theory. It consists in rebuilding sequentially N metamodels, denoted $\mathcal{M}^{PC \setminus i}$, using the original experimental design excluding one point $\mathcal{X} \setminus \mathbf{x}^{(i)} = \{\mathbf{x}^{(j)}, j = 1, \dots, N, j \neq i\}$, and then calculating the prediction error at the withheld point $\mathbf{x}^{(i)}$ (see, e.g., Blatman and Sudret (2011a)). The general formulation of the leave-one-out error is:

$$\epsilon_{LOO} = \frac{\sum_{i=1}^N (\mathcal{M}(\mathbf{x}^{(i)}) - \mathcal{M}^{PC \setminus i}(\mathbf{x}^{(i)}))^2}{\sum_{i=1}^N (\mathcal{M}(\mathbf{x}^{(i)}) - \mu_{\mathcal{M}})^2}. \quad (17)$$

When the model is a linear superimposition of orthogonal terms, as in the case for PCE, the LOO error does not require the actual calculation of N independent surrogates on $N - 1$ experimental design points each. Instead, it can be calculated analytically from a single surrogate based on all the N points according to the following equation:

$$\epsilon_{LOO} = \sum_{i=1}^N \left(\frac{\mathcal{M}(\mathbf{x}^{(i)}) - \mathcal{M}^{PC}(\mathbf{x}^{(i)})}{1 - h_i} \right)^2 \bigg/ \sum_{i=1}^N (\mathcal{M}(\mathbf{x}^{(i)}) - \mu_{\mathcal{M}})^2, \quad (18)$$

where h_i is the i^{th} diagonal term of matrix $\mathbf{A}(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T$.

3 Space-filling designs of experiments

The LAR algorithm allows one to detect automatically the significant terms in the PC expansion. However, its accuracy is greatly dependent on the sampling strategy adopted to generate the ED. Space-filling techniques are a common choice in this context. This class of designs aims at filling the input space $\mathcal{D}_{\mathbf{X}}$ as uniformly as possible for a given number of samples N . In this section, we briefly review the space filling sampling techniques adopted in our study. Without loss of generality, we will consider independent uniform sampling of the unit hypercube. The more general case (for example sampling of non-uniform distributions) can be achieved by considering isoprobabilistic transforms (e.g., Nataf or Rosenblatt transformation, (Rosenblatt, 1952)).

3.1 Monte Carlo sampling

Monte Carlo is a popular method for generating samples from a random vector. Samples are drawn randomly according to the probability distributions of the input variables, using a random number generator. However, the MC sampling method is not particularly efficient. A problem of clustering arises when small sample sizes are considered. In other words, it is likely that there are regions of the parameter space that are scarcely covered by the sampling, as well as regions in which random points are clustered together. Therefore, to reliably cover the entire parameter space, a large number of simulations is often needed, hence resulting in extensive computational costs.

3.2 Quasi Monte Carlo sampling

Quasi Monte Carlo sampling, also known as low discrepancy sequences (LDS), is a family of sampling strategies based on the deterministic placement of sample points as uniformly as possible, avoiding large gaps or clusters. Uniformity is quantified in terms of *discrepancy* from the uniform distribution. The advantage of using low discrepancy sequences is in general a faster convergence rate of stochastic estimators with respect to the standard MC method. Several sequences are available in the literature, *e.g.* Halton (Halton, 1960), Faure (Faure, 1982) and Sobol (Sobol, Sobol) sequences, all of which are based on the Van der Corput sequence. Several practical studies have proven that the Sobol' sequence outperforms the other LDS in regression applications (Kucherenko et al., 2015), therefore it is the one of choice in this paper.

3.3 Latin hypercube sampling

Latin hypercube sampling is widely applied in computational engineering. It was originally introduced by McKay et al. (1979) as an improvement of the Monte Carlo sampling technique. It is a form of stratified sampling technique that improves convergence w.r.t. MCS thanks to a better representation of the sample space. It was shown in Stein (1987) that the variance of the LHS sample mean is asymptotically (i.e. in the limit of a large number of samples) lower than traditional MCS. In its original formulation in McKay et al. (1979), LHS consists in dividing the domain of each input variable in N intervals of equal probability, followed by drawing a random sample inside each interval. A Latin hypercube design of N points in M dimensions can be built from a $N \times M$ matrix π_{ij} , where each column is a random permutation of the sequence $\{1, 2, \dots, N\}$. The design points are then defined as follows:

$$\mathcal{X}_{LHS} = \{x_j^{(i)} = \frac{\pi_{ij} - 1 + u_{ij}}{N}, i = 1, \dots, N; j = 1, \dots, M\}. \quad (19)$$

where u_{ij} is a realization of a uniform random variable $U \sim \mathcal{U}(0, 1)$. It is easy to verify that one dimensional projections of the LHS design are uniformly spaced. However, this property is not

guaranteed in principle in M -dimensional space. Because of the random nature of LHS, some configurations exist such that the input variables are correlated, resulting in poor space filling qualities. Several modifications of the sampling scheme have been proposed in the literature leading to the so-called optimal Latin hypercube sampling designs (OLHS). One common strategy is to combine standard LHS with a space filling criterion such as entropy, maximin distance (Pronzato and Müller, 2012), low discrepancy or minimum correlation between the samples. In general, an OLHS is achieved by creating several independent LHS designs of the desired size, then choosing the one that best satisfies the chosen optimality criterion.

The main difficulty with LHS designs is that the sampling size has to be selected *a priori*. To circumvent this problem, Blatman and Sudret (2010); Rennen (2009); Wang (2003) developed the so-called nested-LHS designs (NLHS). The main idea is to enrich an existing LHS design such that the resulting combined design meets the LHS requirements as closely as possible. Enriching an existing LHS of size N with N_{en} points can be achieved by substituting the partitioning of the unit hypercube π_{ij} in Eq. (19) with a finer one of size $N + N_{en}$. Then, N_{en} empty partitions are identified and a random sample is drawn in each of them. While this procedure does not guarantee that the final design is still an LHS (due to the possibility that multiple points of the original design lie in the same stratum after the new partitioning), the properties of the final design tend to be closer to those of an LHS. For more details, see Blatman and Sudret (2010); Wang (2003).

3.4 Measures of quality for space filling designs

A number of different criteria for assessing the space filling quality of a particular experimental design can be found in the literature. In the following, we briefly recall two widely used space-filling criteria, allowing for enhanced discrimination when choosing among candidate designs. The first one is the maximin distance criterion (Johnson et al., 1990), which aims at maximizing the smallest distance between neighboring points, *i.e.*:

$$\text{maximin}(\mathcal{X}) \stackrel{\text{def}}{=} \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y} \in \mathcal{X}} d(\mathbf{x}, \mathbf{y}) \quad (20)$$

where $d(\mathbf{x}, \mathbf{y})$ is the Euclidean distance between two sample points \mathbf{x} and \mathbf{y} . A second criterion was proposed to assess the uniformity quality of the design based on a discrepancy measures. To this purpose, a common and computationally efficient tool is the centred D^2 discrepancy

criterion, which reads:

$$D^2(\mathcal{X}) = \left(\frac{13}{12}\right)^M - \frac{2}{N} \sum_{i=1}^N \prod_{k=1}^M \left(1 + \frac{1}{2}|x_k^{(i)} - \frac{1}{2}| - \frac{1}{2}|x_k^{(i)} - \frac{1}{2}|^2\right) + \frac{1}{N^2} \sum_{i,j=1}^N \prod_{k=1}^M \left(1 + \frac{1}{2}|x_k^{(i)} - \frac{1}{2}| - \frac{1}{2}|x_k^{(j)} - \frac{1}{2}| - \frac{1}{2}|x_k^{(i)} - x_k^{(j)}|\right) \quad (21)$$

A design is called uniform if it shows a low centered D^2 discrepancy.

4 Optimal designs of experiments

Optimal designs leverage on the knowledge about the properties of the surrogate model that is being considered. This can be achieved by searching a design \mathcal{X} that optimizes a given optimality criterion. For instance, in the context of regression-based surrogate models (*e.g.* PCE), Zein et al. (2013) proposed a sampling method for regression-based polynomial chaos expansion that consists in maximizing the determinant of the information matrix $\mathbf{A}^\top \mathbf{A}$ in Eq. (10) for a given number of samples, by using the D-optimality criterion. Recently, Shin and Xiu (2016) proposed a point selection approach to construct an optimal experimental design with the aim of providing an accurate OLS estimation. The optimal design is determined by maximizing a quantity that combines both the determinant and the column orthogonality of the information matrix, denoted S -value criterion. In the following section, we will briefly review the D-optimality criterion and S -value criterion for regression-based PCE.

4.1 D-optimal designs

The D-optimality criterion is aimed at minimizing the variance of the regression parameters estimate. The latter is derived from Eq. (10) as follows:

$$\begin{aligned} \text{Var}(\hat{\mathbf{y}}) &= \text{Var}((\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathcal{Y}) \\ &= (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \text{Var}(\mathcal{Y}) ((\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A})^\top. \end{aligned} \quad (22)$$

Under the assumption of independent normal model errors with constant variance ($\text{Var}(\mathcal{Y}) = \sigma^2 \mathbf{I}$), we obtain:

$$\text{Var}(\hat{\mathbf{y}}) = (\mathbf{A}^\top \mathbf{A})^{-1} \sigma^2. \quad (23)$$

Reducing the variance of the PCE coefficients leads to the minimization of the determinant $\det(\mathbf{A}^\top \mathbf{A})^{-1}$. To avoid the inversion, one maximizes the determinant of the information matrix $\det(\mathbf{A}^\top \mathbf{A})$. The D -optimal design, denoted \mathcal{X}_D , is then the solution of the following optimization problem over the input design space $\mathcal{D}_{\mathbf{X}}$:

$$\mathcal{X}_D = \arg \max_{\mathcal{X} \in \mathcal{D}_{\mathbf{X}}} \det(\mathbf{A}^\top \mathbf{A}). \quad (24)$$

4.2 S-optimal designs

We herein briefly review the optimal approach developed in Shin and Xiu (2016). The goal is to find the optimal N -subset within a large-size sample \mathcal{X}_L of size N_L (typically in the order of $N_L = 10^{5-6}$) that provides the most accurate PCE. It is a subset selection problem where potential solutions are viewed as a $N \times M$ matrix \mathcal{X} consisting of exactly N different rows of \mathcal{X}_L . The minimization problem in Eq. (9) over the ED \mathcal{X} reads:

$$\hat{\mathbf{y}} = \operatorname{argmin}_{\mathbf{y} \in \mathbb{R}^{\operatorname{card} \mathcal{A}}} \frac{1}{N} \sum_{i=1}^N \left(\mathcal{M}(\mathbf{x}^{(i)}) - \sum_{\alpha \in \mathcal{A}} y_{\alpha} \Psi_{\alpha}(\mathbf{x}^{(i)}) \right)^2. \quad (25)$$

Equivalently, the minimization problem in Eq. (9) over the large ED, \mathcal{X}_L , reads:

$$\hat{\mathbf{y}}_L = \operatorname{argmin}_{\mathbf{y} \in \mathbb{R}^{\operatorname{card} \mathcal{A}}} \frac{1}{N_L} \sum_{i=1}^{N_L} \left(\mathcal{M}(\mathbf{x}_L^{(i)}) - \sum_{\alpha \in \mathcal{A}} y_{\alpha} \Psi_{\alpha}(\mathbf{x}_L^{(i)}) \right)^2. \quad (26)$$

The aim is to find a subset such that the estimated PCE coefficients based on the optimal design are as close as possible to those based on the large ED, \mathcal{X}_L . This is mathematically equivalent to the following optimization problem:

$$\mathcal{X}^{\dagger} = \operatorname{argmin}_{\mathcal{X} \subset \mathcal{X}_L} (\|\hat{\mathbf{y}}(\mathcal{X}) - \hat{\mathbf{y}}_L\|). \quad (27)$$

This can be intuitively interpreted as a row selection problem, such that:

$$R^{\dagger} = \operatorname{argmin}_R (\|\hat{\mathbf{y}}(R\mathcal{X}_L) - \hat{\mathbf{y}}_L\|), \quad (28)$$

where $R = [R_{(1)}; \dots; R_{(N)}] \in \mathbb{R}^{N \times N_L}$ denotes a row selection matrix, where each row $R_{(i)} = [0, \dots, 1, 0, \dots, 0]$ is a vector of length- N_L with all zeros but a one at k_i location, $1 \leq k_i \leq N_L$. As is the case in many subset selection algorithms, QR decomposition is used to find the most representative rows. Let \mathbf{A}_L denote the model matrix obtained when considering \mathcal{X}_L and the corresponding model evaluations \mathcal{Y}_L . let $\mathbf{A}_L = Q_L R_L$ be the QR factorization of \mathbf{A}_L . It is then possible to prove that:

$$\|\hat{\mathbf{y}} - \hat{\mathbf{y}}_L\| \leq \|R_L^{-1}\| \cdot \|g\|, \quad (29)$$

where \hat{g} is the least square solution to

$$RQ_L g = RP^{\perp} \mathcal{Y}_L, \quad (30)$$

where $P^{\perp} = (I - Q_L Q_L^T)$ is the projection operator to the orthogonal complement of the column space of Q_L .

It is then clear that achieving a smaller $\|g\|$ leads to the minimization of the difference $\|\hat{\mathbf{y}} - \hat{\mathbf{y}}_L\|$. To do this, one seeks to: (i) maximize the determinant of RQ_L so that the variance

of g is minimized and (ii) maximize the column orthogonality of RQ_L , so that the correlation between design points is minimized. For this purpose, a selection criterion combining these conditions was proposed in Shin and Xiu (2016). Such criterion is hereafter called S -value. For the sake of simplicity, the S -value will be defined for a matrix \mathbf{A} of size $N \times \text{card}(\mathcal{A})$ as follows:

$$\mathcal{S}(\mathbf{A}) = \left(\frac{\sqrt{\det \mathbf{A}^\top \mathbf{A}}}{\prod_{i=1}^{\text{card} \mathcal{A}} \|\mathbf{A}^{(i)}\|} \right)^{\frac{1}{\text{card} \mathcal{A}}}, \quad (31)$$

where $\mathbf{A}^{(i)}$ is the i -th column of \mathbf{A} . The optimal sample is thus the solution of the following optimization problem:

$$R^\dagger = \underset{R}{\text{argmax}}(\mathcal{S}(RQ_L)). \quad (32)$$

It is worth noting that in the context of PCE, the columns of the model matrix \mathbf{A} consist of orthonormal basis elements. The QR decomposition is then not necessary and the matrix \mathbf{A}_L can be used instead of Q_L . Thus, we seek:

$$R^\dagger = \underset{R}{\text{argmax}}(\mathcal{S}(R\mathbf{A}_L)). \quad (33)$$

The optimality condition Eq. (33) can then be equivalently rewritten as follows:

$$\mathcal{X}^\dagger = \underset{\mathcal{X} \in \mathbb{X}}{\text{arg max}}(\mathcal{S}(\mathbf{A}(\mathcal{X}))). \quad (34)$$

In their recent study, Shin and Xiu (2016) considered a standard truncation scheme to select the set of basis elements. Although their optimal design has proven powerful, it may still face challenges in high dimensional cases. As already mentioned, one can reduce the size of the candidate set by considering other truncation strategies, *e.g.* the hyperbolic truncation. However, it may not be known in advance which candidate set will perform best. This motivates the need of an adaptive strategy that selects iteratively the optimal candidate basis and thus the associated optimal design.

5 Sequential experimental designs

In the presence of expensive computational models, it can be preferable to adopt a greedy iterative strategy in the construction of an experimental design for surrogate modelling. The class of so-called sequential experimental designs is based on constructing a relatively small initial space-filling design, training a surrogate model and assessing its accuracy. If needed, the initial design can then be enriched sequentially until some accuracy measure of the surrogate model (*e.g.* the number of desired points or the $RMSE$ in Eq. (15)) is satisfied. In this paper we explore two different strategies to address this problem. The first is to enrich the initial

experimental design according to a space-filling criterion without assuming any knowledge on the functional form of the surrogate model. The second assumes some degree of knowledge on the behaviour of the surrogate model. Note that, in any case, none of the criteria used in this study depend on the model responses on the experimental design points.

In the following, the initial experimental design \mathcal{X}_{init} of size N_{init} is enriched at each iteration with N_{new} data points. The new points are to be selected, according to a specific criterion, from an available ED \mathcal{X}_L of size N_L such that $N_L \gg N_{init}$.

5.1 Sequential selection using space filling criteria

In this section we review a sequential approach for enriching the initial experimental design according to a maximin distance criterion. This strategy was proposed by Kersaudy (2015). Note that such idea was originally developed by Ravi et al. (1991) as a method for the dispersion problem and applied later by Rennen (2009) to subset selection for Kriging metamodels, thereby referred to as greedy maximin selection. At each iteration, the algorithm selects the point $\mathbf{x}^* \in \mathcal{X}_L$ that maximizes the minimum distance with all the points in the current experimental design \mathcal{X} . The obtained experimental design \mathcal{X} of size $N = N_{init} + N_{new}$ is therefore generated according to the following algorithm:

1. Let $\mathcal{X} = \mathcal{X}_{init}$ (note that $N_{init} < N$)
2. find $\mathbf{x}^* = \max_{\mathbf{x}_L \in \mathcal{X}_L} \min_{\mathbf{x} \in \mathcal{X}} (\|\mathbf{x} - \mathbf{x}_L\|)$
3. update $\mathcal{X} = \mathcal{X} \cup \mathbf{x}^*$ and $\mathcal{X}_L = \mathcal{X}_L \setminus \mathbf{x}^*$
4. repeat step ii until the set \mathcal{X} contains the desired number of points N .

The main advantage of this method is that it is easy to implement. However, the main disadvantage lies in that one needs to compute all the distances between every point of $\mathbf{x}_L^{(i)} \in \mathcal{X}_L$ and every point $\mathbf{x}^{(j)} \in \mathcal{X}$, which can be challenging in high dimension. In the following, the obtained design is referred to as sequential maximin design. A similar idea has also been proposed in Iooss et al. (2010) where the centred discrepancy measure D^2 in Eq. (21) is used instead of the maximin criterion. In this case, the algorithm enriches the initial experimental design \mathcal{X}_{init} with points that minimize its D^2 discrepancy. It operates in the same way as the greedy maximin algorithm. In this paper, we choose to limit our study to the maximin criterion based on Euclidian distance.

5.2 Adaptive selection of the optimal design

The optimal design strategies presented until now rely on the assumption that the regressors in an OLS problem are known. However, in the context of sparse PCE, the optimal polynomial basis is

in general not known a priori. To address this problem, we propose an iterative strategy in which the estimated sparse PCE basis is automatically updated and the experimental design is enriched. To construct a sparse PCE, we propose the use of a basis selection method based on compressive sensing techniques. In particular, we rely herein on the least angle regression algorithm (LAR) (Efron et al., 2004). Other alternatives (*e.g.* orthogonal matching pursuit (OMP) (Jakeman et al., 2015; Mallat and Zhang, 1993) or stepwise regression (Hesterberg et al., 2008)) are also directly compatible with our procedure.

5.2.1 Sparse PCE: Least Angle regression

The LAR algorithm is an iterative regression method applied in the context of PCE for basis selection (Blatman and Sudret, 2011a). The basic idea of LAR is to select a parsimonious number of regressors that have the greatest impact on the model response, among a possibly large set of candidates. LAR eventually provides a sparse polynomial chaos expansion, *i.e.* one that only contains a small number of regressors compared to a classical full representation. It consists on iteratively moving regressors from a *candidate set* to an *active set*. The next regressor is chosen based on its correlation with the current residual. At each iteration, analytical relations are used to identify the best set of regression coefficients for that particular active set, by imposing that every active regressor is equicorrelated with the current residual. At the end of the process, the optimal sparse PC basis is chosen according to a suitable measure of accuracy. In particular, the leave-one-out (LOO) error, ϵ_{LOO} , is a natural candidate in the context of PCE.

The LAR algorithm in the context of PCE (Blatman and Sudret, 2011a) can be summarized as follows:

1. Initialize the set of candidate regressors to the full basis and the active set of regressors to \emptyset .
2. Initialize all coefficients equal to 0. Set the residual equal to the output vector.
3. Find the regressor Ψ_{α_j} that is most correlated with the current residual
4. Move all the coefficients of the current active set towards their least-square value until their regressors are equicorrelated to the residual as some other regressor in the candidate set. This regressor will also be the most correlated to the residual in the next iteration.
5. Calculate and store the accuracy estimates, ϵ_{LOO}^j , for the current iteration
6. Update all the active coefficients and move Ψ_{α_j} from the candidate set to the active set
7. Repeat the previous steps (iii-v) until the size of the active set is equal to $m = \min(P, N-1)$
8. Choose the best iteration based on their prescribed error estimate, ϵ_{LOO} .

Note that after detecting the relevant basis by LAR, the coefficients are recomputed using the ordinary least-square method (hybrid-LAR). Readers are referred to Blatman and Sudret (2010) for further details.

5.2.2 Adaptive selection of the optimal design

The algorithm proposed in the last subsection allows one to detect the most relevant regressors in the PC expansion on the basis of the current experimental design. However, it is based on an arbitrarily fixed experimental design. To address this problem, we propose an adaptive strategy that updates iteratively both the experimental design and the estimated sparse PCE basis. The initial ED generated using a space filling techniques will be enriched sequentially according to an optimality criterion. At each iteration, the optimality criterion is reformulated based on the updated sparse basis. The proposed algorithms works as follows:

1. **Initial experimental design:** we start with a small sample set \mathcal{X}_{init} of a size N_{init} , normally produced by LHS or Quasi Monte Carlo sequences.
2. **Basis selection:** a degree-adaptive sparse PCE (Blatman and Sudret, 2011b) is then constructed using LARS to identify the sparse set of polynomials (out of a candidate set) that best approximates the computational model based on the current experimental design.
3. **Formulation of the model matrix:** This step is the core of the sequential process. The information matrix is updated according to the basis obtained in the previous step. The optimality criterion (D -optimal or S -optimal, see Sections 4.1 and 4.2, respectively) is then re-evaluated on $\mathcal{X}_L \setminus \mathcal{X}$ to identify the next set of candidate points X_{new} .
4. **Experimental design augmentation:** The identified N_{new} S - or D - optimal data points are added to the current experimental design, in which the computational model is evaluated. The sparse PCE is then updated.

The procedure (steps ii, iii and iv) is repeated until a stopping criterion is fulfilled. In particular, the accuracy of the metamodel (estimated *e.g.* with cross-validation techniques) or the maximum number of sample points can be used. For easier comparisons between the methods, we choose to fix the maximum number of sample points as the stopping criterion in the following.

The effectiveness of the adaptive methods is dependent on the optimization algorithm that searches the optimal sample points. Various search strategies have been proposed in the literature. In this work, the D -optimal experimental design is generated using the built-in MATLAB function *candexch* from the statistics toolbox. The *candexch* routine implements the Fedorov exchange algorithm (Miller and Nguyen, 1994). The S -optimal experimental design is generated

using the greedy algorithm proposed in Shin and Xiu (2016). For a comprehensive overview of the implementation of the greedy algorithm, the reader is referred to Shin and Xiu (2016).

6 Results and Discussion

In this section, we investigate the efficiency of the proposed adaptive sequential sampling methods, when either the D -optimality or the S -value criteria are used to generate the sequential samples. For the sake of clarity, the sequential sampling corresponding to these two sampling settings are denoted as *Seq D-optimal* and *Seq S-optimal*, respectively. A comparative study with other commonly used sampling methods is performed. Four different sampling techniques are considered: *maximin* LHS, *min – corr* LHS, Sobol’ sequences and Sequential *maximin* design, denoted *Seq maximin*. *Maxmin* LHS is obtained by selecting an LHS design from a set of five different designs according to the maximin distance criterion, while *min – corr* LHS is obtained similarly, but with a minimum-correlation based selection criterion. For this purpose, we use the MATLAB function *lhsdesign*. As already mentioned, D-optimal designs are generated with the MATLAB built-in *candexch* function. For comparison of the space-filling properties of the sample points, the measures described in Section 3.4 are used. In addition, the condition number (CN) of the information matrix is considered, as a measure of orthogonality.

In this section, the proposed adaptive sampling methods are validated on four models with varying input dimensionality and complexity. They comprise two analytical functions, namely the Ishigami function and the Sobol’ g-function, and two finite element models, namely a truss structure and a 1D diffusion problem. Because all of the presented benchmarks are inexpensive to evaluate, the relative generalization error (RMSE) on a validation set is then employed to assess the accuracy of the resulting polynomial chaos expansions. All of the PCE metamodels, as well as the LARS basis selection, are developed within the UQLAB software framework (Marelli and Sudret, 2014, 2015).

Since the performance of the ED is sensitive to the random nature of the sampling methods, each analysis is replicated 50 times. To generate a randomized Sobol’ sequences, several parallel sequences are generated using different starting points. The results are then summarized in box-plots. These replications aim at assessing the effect of stochastic variations in the generation of the designs.

6.1 Ishigami function

The Ishigami function (Ishigami and Homma, 1990) is a well-known benchmark for polynomial chaos expansions. It is a 3-dimensional analytical function characterized by non-monotonicity

and high non-linearity, given by the following equation:

$$f(\mathbf{x}) = \sin(x_1) + a \sin^2(x_2) + bx_3^4 \sin(x_1). \quad (35)$$

The input random vector consists of three *i.i.d.* uniform random variables $X_i \sim \mathcal{U}(-\pi, \pi)$. The numerical values $\{a = 7, b = 0.1\}$ are chosen for this example. Note that this model is sparse in nature.

A sparse, degree-adaptive PCE approach with maximum degree in the range 3-20 is chosen. The PCE coefficients are computed with the hybrid LAR method (Blatman and Sudret, 2010). The accuracy of the PCE is assessed with a MCS-based validation set of size $N_{val} = 10^6$. The sequential sampling methods (*Seq* S-optimal, *Seq* D-optimal and *Seq* maximin) are initialized with an experimental design generated by maximin LHS of size 40, iteratively augmented with $N_{add} = 20$ additional points. For the other sampling methods (*maximin* LHS, *min - corr* LHS), a predefined size is considered. As already mentioned, the analysis is replicated 50 times to assess the statistical uncertainty. The performance of the different sampling methods is compared in terms of their RMSE on the validation set for varying ED sizes.

As a first step, the performances of the various design strategies with respect to the space-filling quality measures introduced in Section 3.4 are compared. The average values of the different measures calculated from the 50 repetitions for various sizes of the ED are plotted in Figure 1. The top left panel of Figure 1 reports results for the maximin distance D_{min} criterion. As expected, the sequential *maximin* experimental design systematically shows the maximal minimum distance across all designs. All the other designs show a consistent decrease in the minimal distance with increasing ED size, leading to a deterioration of their space filling properties. The worst results are obtained with D-optimal design.

In the same format, the top-right panel of Figure 1 reports results for the *discrepancy* criterion. All the non-sequential strategies tested (LHS and QMC) show an approximately constant discrepancy regardless of the experimental design size. However, all the sequential strategies show instead a constant increase with the experimental design size. This means that all the sequential strategies tend to produce samples that deviate from the uniform distribution. In other words, design strategies that are optimal for regression problems do not necessarily exhibit a low-discrepancy behaviour.

The condition number of the model matrices are plotted with respect to the size of the ED in the lower panel of Figure 1. As expected, *Seq* S-optimal designs show a consistently smaller condition number for all experimental designs sizes. *Seq* D-optimal, *maximin* LHS and *min - corr* LHS exhibit comparable behaviours with a monotone increase in condition number. The remaining *Seq* maximin ED and Sobol' sequences show instead a somewhat more erratic




Figure 1: The average, over 50 trials of the *maximin* distance, discrepancy measure and the condition number for the 5 tested experimental designs for the Ishigami function. Top left (a): Variation of *maximin* distance with increasing size of the ED; Top right (b): Variation of Discrepancy measure with increasing size of the ED; Bottom (c): Variation of condition number criterion with increasing size of the ED;

behaviour. Overall, Figure 1 shows that *Seq* S-optimal designs provides an overall good balance between the three criteria.

To gain an intuitive understanding of the different experimental designs, Figure 2 compares 2D scatter plots and histograms for the final ED of size $N = 160$ generated by *maximin* LHS and the three sequential designs *Seq* maximin, *Seq* D-optimal and *Seq* S-optimal. Uniform space coverage is obtained only with *maximin* LHS. The sequential design obtained with *Seq* maximin and *Seq* S-optimal designs appear to be fairly uniformly distributed, with some degree of accumulation on the boundaries. On the contrary, *Seq* D-optimal sample points tend to cluster near the boundaries of the design space.

The metamodeling performance of the different sampling strategies are compared in Figure 3. The boxplots represent the median RMSE error, the 25% and 75% percentiles and the extreme error values (outliers) of the 50 independent runs for varying ED sizes, ranging from $N = 40$ to $N = 160$. We observe that *Seq* S-optimal design consistently outperforms other sampling strategies. In addition, it generally shows a more stable behaviour with smaller variability between repetitions, especially as the size of the experimental design increases. Accurate results are also achieved with *Seq* maximin design. The prediction improvement seems to be more important for larger experimental designs. We note that comparable results are obtained with *maximin* LHS, *min - corr* LHS and Sobol' sequences.

6.2 Sobol function

Let us consider now the so called Sobol' g-function:

$$f(\mathbf{x}) = \prod_{i=1}^q \frac{|4x_i - 2| + a_i}{1 + a_i} \quad (36)$$

The input parameters $\{X_i, i = 1, \dots, q\}$ are *i.i.d.* uniform distributions over $[0, 1]$. The $\{a_i, i = 1, \dots, q\}$ are parameters controlling the complexity of the function. For our numerical application, we choose to fix $q = 8$ and $a = [1, 2, 5, 10, 20, 50, 100, 500]$. The same analysis presented for the Ishigami function was conducted to illustrate the performance of the proposed approach for this benchmark. The Sobol' function is highly complex (strongly nonlinear, non-monotonic and non-polynomial), therefore requiring high polynomial degree even for low dimensions. The truncation options $p = 30$ and $q = 0.5$ are used for the PCE calculation. The initial experimental design is generated by a maximin LHS of size 100, augmented iteratively with $N_{add} = 50$ additional points.

The performances of the various design strategies with respect to the space-filling quality measures have been compared, but they are not shown herein because the general conclusions drawn for the Ishigami function still hold. The metamodeling performance of the different




Figure 2: An example of 2D plots of a 3-dimensional designs of size 160. Top left: Experimental design obtained with maximin LHS; Top right: Experimental design obtained with Seq maximin; Bottom left: Experimental design obtained with maximin sequential D-optimal; Bottom right: Experimental design obtained with sequential S -optimal;




Figure 3: Evolution of boxplots of the validation error for varying sizes of the experimental design for the Ishigami function. The boxplots are obtained from 50 replications.

sampling strategies are compared in Figure 4. As in the previous example, the *Seq* S-optimal design outperforms the others sampling methods (Figure 4). Good performance are also achieved with *Seq* maximin design as in the Ishigami function. Furthermore, much smaller variances (boxplots are smaller) are shown for *Seq* S-optimal design and lead to the conclusion that this class of designs is robust. Sobol' sequences give the least accurate results. This behaviour may be due to the increase in the dimensionality of the problem, coupled with the tendency of Sobol' sequences to cluster on regular sub-manifolds in high dimension (Morokoff and Caffisch, 1994).

6.3 Maximum deflection of a truss structure

Let us consider the truss structure shown in Figure 5. Ten independent variables are considered, namely: the vertical loads, P_1, \dots, P_6 ; the cross-sectional area and Young's modulus of the horizontal bars, A_1 and E_1 , respectively; the cross-sectional area and Young's modulus of the vertical bars, A_2 and E_2 , respectively (Figure 5). The model input parameters are described by independent random variables whose distributions are reported in Table 1. An isoprobabilistic transformation of \mathbf{X} into standard normal variables is employed. In the underlying deterministic problem, the truss deflection u is computed with a finite-element MATLAB code. The same study as in the two previous examples is conducted for this 10-dimensional function.

A sparse degree-adaptive PCE approach with maximum degree in the range 3 – 20 is chosen.




Figure 4: Evolution of boxplots of the validation error for varying sizes of the experimental design for the Sobol function. The boxplots are obtained from 50 replications.




Figure 5: Truss structure with 23 members.

The accuracy of the PCE is assessed with a MCS-based validation set of size $N_{val} = 10^6$. The sequential sampling methods (*Seq S-optimal*, *Seq D-optimal* and *Seq maximin*) are initialized with an experimental design generated by maximin LHS of size 100, augmented iteratively with $N_{add} = 50$ additional points. For the other sampling methods (*maximin LHS*, *min-corr LHS*), a predefined size is considered. The performance of the different sampling methods is compared in terms of their RMSE on the validation set for varying ED sizes.

As in the previous cases, the performances of the various design strategies with respect to the space-filling quality measures were compared but the results are not shown though for the sake of conciseness, as they are in line with the previous two cases. The metamodeling performance of the different sampling strategies are compared in Figure 6. As in the two previous examples, the *Seq S-optimal* design performs better than the other sampling methods, with a faster decrease of

Table 1: Truss-deflection problem: random input 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
$P_1 \cdots P_6$ [KN]	Gumbel	50	0.15

the validation errors. One sees that the computational budget gain increases exponentially with the size of the ED. Fairly accurate results are also obtained with *Seq* maximin design. Similar accuracies are obtained with *maximin*, *min - corr* LHS. Again, a poor performance is observed for Sobol' sequences and *D*-optimal designs, especially for large experimental designs.




Figure 6: Evolution of boxplots of the validation error for varying sizes of the experimental design for the truss function.

6.4 One-dimensional diffusion problem

In this section, we consider a standard benchmark in UQ computations (Shin and Xiu, 2016), a one-dimensional stochastic diffusion problem subject to uncertainty in the diffusion coefficient.

Let us consider the following boundary-value problem defined over a domain $\mathcal{D} = [0; L]$:

$$\begin{aligned} -\frac{d}{dx} \left[E(x, \omega) \frac{du}{dx}(x, \omega) \right] + f(x) &= 0 \\ u(0) &= 0 \\ \frac{du}{dx}(L) &= F. \end{aligned} \tag{37}$$

Furthermore, assume that the diffusion coefficient is a log-normal random field described by:

$$E(x, \omega) = \exp(\lambda_E + \zeta_E g(x, \omega)), \tag{38}$$

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

$$\rho(x, x') = e^{-|x' - x|/l}. \tag{39}$$

The Gaussian random field $g(x)$ in (38) is approximated using the Karhunen-Loève Expansion (Loève, 1977):

$$g(x, \omega) = \sum_{k=1}^M \sqrt{\iota_k} \varphi_k(x) \xi_k(\omega), \tag{40}$$

Where $\{(\iota_k; \varphi_k), k = 1, \dots, M\}$ are the solutions of the Fredholm equation associated to the exponential kernel, see Ghanem and Spanos (1991); Sudret and Der Kiureghian (2000) for the detailed analytical solution.

For the numerical test we set the following numerical values: $L = 1$, $F = 1$, $f = 0.5$. We set: $\lambda_E = 10$, $\zeta_E = 3$ and $l = 1/3$. The number of terms to be retained in (40) is selected such that:

$$\sum_{k=1}^M \iota_k / \sum_{k=1}^{\infty} \iota_k \geq 0.99. \tag{41}$$

$M = 62$ are therefore required, making this the highest dimensional example considered.. The quantity of interest (model response) is $u(L)$.

A sparse, degree-adaptive PCE approach with maximum degree in the range 1-4 is chosen with $q = 0.5$. It is worth noting that the model is extremely sparse with a low effective dimension, although its high dimensionality. For the sequential sampling methods, the following conditions are then set: an initial experimental design of size $N_{init} = 100$ generated with *maximin* LHS, $N_{add} = 20$ are considered.

We now investigate the metamodeling performance of the different sampling strategies. Figure 7 shows boxplots of the RMSE error. As in the previous examples, the *Seq* S-optimal performs better than the other sampling methods. Unlike the previous examples, *Seq* maximin

designs do not perform better than *maximin* LHS and *min-corr* LHS. It is worth to note that *Seq* maximin designs search the design in 62 dimensions and does not capture the sparsity of the model (i.e.; the model has a few variables that are significantly more influential than the others). Again, a poor performance is observed for the *Seq D*-optimal designs.




Figure 7: Evolution of boxplots of the validation error for varying sizes of the experimental design for the 1D diffusion function.

6.5 Effects of the choice of N_{init} and N_{Add}

Determining the optimal sample size of the initial design in the sequential construction is a difficult task. In the context of sequential sampling, previous literature does not have a well-defined rule of thumb for the number of initial runs, given a total budget N . According to Bernardo et al. (1992), the initial design should at least contain three observations per input variable. Jones et al. (1998) recommended larger initial designs with up to ten observations per input variable. A similar question holds on the choice of N_{add} . It is likely that the optimal N_{init} and N_{add} depend on the nature of the model that is to be approximated and on its degree of sparsity.

To provide some insight into the effects of those choices on the *Seq S*-optimal design class, a parametric study was performed on the examples presented in the previous sections. For each of the examples provided, several choices of N_{init} and N_{add} were compared in terms of the final surrogate RMSE prediction. The most significant results are reported in Figure 8 for *Seq S*-

optimal designs. The results for the Ishigami function suggest that indeed the choice of N_{init} and N_{add} plays a significant role in the accuracy of PCE for all of the considered ED sizes, with best performances achieved with a relatively large initial experimental design. However, no comparable trends were observed for either of the remaining benchmarks in the other panels of Figure 8.

These results show that no reliable rule of thumb can be determined to choose N_{init} and N_{add} . The idea is to start with a relatively small sample to obtain initial information about the global behaviour of the metamodel. This information will be updated through iterations. However, if the initial design is too small, the obtained information is not sufficient to achieve an accurate estimate of the optimal basis from LARS, leading to a poor optimization of the S -value criterion in the early iterations. A rather large initial design will not ameliorate the results despite an increased computational effort.

7 Conclusions

The paper handles the topic of selecting the best design of experiments in order to ensure the optimal accuracy of sparse polynomial chaos expansions over the whole input space within a given computational budget. The coefficients of the PCE are computed based on least-square analysis by relying on the so-called experimental design. An efficient approach for constructing an optimal experimental design for sparse polynomial chaos expansions was proposed. The complete approach includes two major elements: (a) the use of degree-adaptive sparse PCE based on LARS for selecting the optimal basis, and (b) the choice of an efficient criterion for selecting the optimal sample points. The ED is build up sequentially by enriching an initial design, which is generated using a space-filling technique. The added points are chosen according to an optimality criterion that depends on the current sparse basis. In this work, two optimality criteria are used, that is: (i) D-optimality and (ii) S-value. The D-optimal criterion is related to the determinant of the information matrix whereas the S-value criterion is related to the orthogonality of the columns of the information matrix.

A comprehensive comparative study of the different sampling strategies was performed on four numerical models with varying input dimensionality and complexity. To complete our investigation, a sequential strategy based on optimizing the well-known maximin distance criterion was also studied. Our results lead to the following major conclusions.

- The use of a sequential adaptive design according to the S -value criterion, *Seq S-optimal*, ensures optimal performances in terms of the relative generalization error. The proposed approach provides a significant accuracy and stability improvement compared to other sampling strategies.

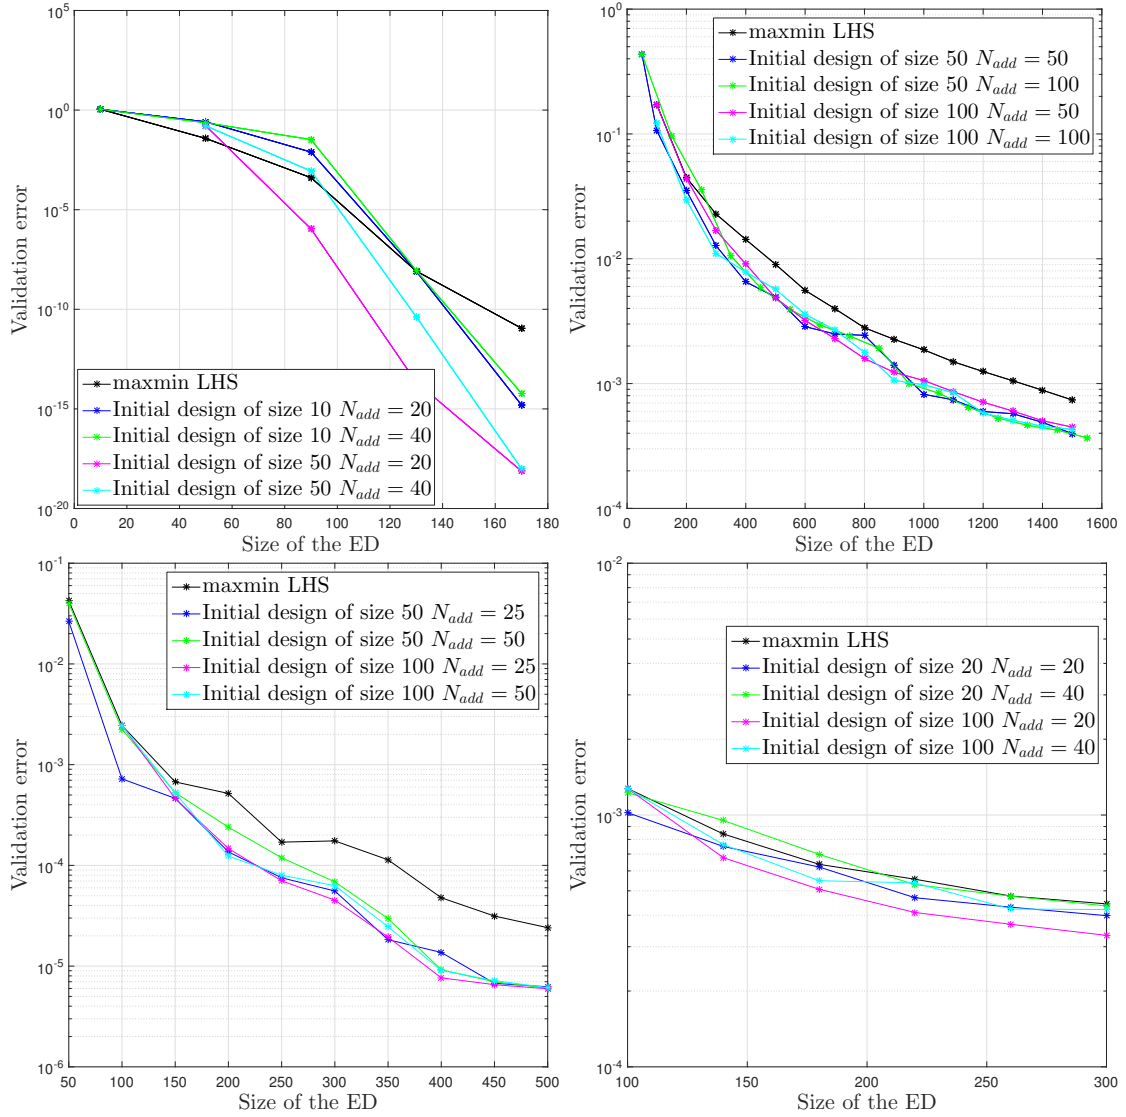


Figure 8: Effect of the size of the initial design and N_{add} points. Top left: Ishigami function; Top right: Sobol function; Bottom left: Truss function; Bottom right: 1D diffusion problem

- The *Seq* D-optimal approach has shown poor performances.
- The sequential maximin designs performed well when the number of random input parameters is low to moderate, say less than 10. However, an erratic behaviour is observed for high dimensional cases.

Finally, further work may include using other basis selection approaches such as Orthogonal Matching Pursuit (OMP) (Jakeman et al., 2015; Mallat and Zhang, 1993) as well as metamodelling techniques such as Low Rank Approximations (LRA) (Chevreuil et al., 2015; Konakli and Sudret, 2016a,b,c) to identify the optimal ED. Future research may also investigate the use of the sequential adaptive designs in the framework of reliability analysis.

References

- Bernardo, M. C., R. Buck, L. Liu, W. A. Nazaret, J. Sacks, and W. J. Welch (1992). Integrated circuit design optimization using a sequential strategy. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.* *11*(3), 361–372.
- Berveiller, M., B. Sudret, and M. Lemaire (2006). Stochastic finite elements: a non intrusive approach by regression. *Eur. J. Comput. Mech.* *15*(1-3), 81–92.
- Blatman, G. and B. Sudret (2010). An adaptive algorithm to build up sparse polynomial chaos expansions for stochastic finite element analysis. *Prob. Eng. Mech.* *25*, 183–197.
- Blatman, G. and B. Sudret (2011a). Adaptive sparse polynomial chaos expansion based on Least Angle Regression. *J. Comput. Phys.* *230*, 2345–2367.
- Blatman, G. and B. Sudret (2011b). Adaptive sparse polynomial chaos expansion based on Least Angle Regression. *J. Comput. Phys* *230*, 2345–2367.
- Chevreuil, M., R. Lebrun, A. Nouy, and P. Rai (2015). A least-squares method for sparse low rank approximation of multivariate functions. *SIAM/ASA J. Uncertain. Quantif* *3*(1), 897–921.
- 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.
- Crombecq, K., E. Laermans, and T. Dhaene (2011). Efficient space-filling and non-collapsing sequential design strategies for simulation-based modeling. *Eur. J. Oper. Res.* *214*(3), 683–696.

- dos Santos, M. I. R. and P. M. R. dos Santos (2007). Sequential designs for simulation experiments: nonlinear regression metamodeling. In *Proc. 26th IASTED Int. Conf. Modell. Identif. Control*, pp. 88–93. ACTA Press.
- Efron, B., T. Hastie, I. Johnstone, and R. Tibshirani (2004). Least angle regression. *Ann. Stat.* 32, 407–499.
- Fang, K.-T., R. Li, and A. Sudjianto (2005). *Design and modeling for computer experiments*. CRC Press.
- Faure, H. (1982). Discrépance de suites associées à un système de numération (en dimension s). *Acta Arith.* 41(4), 337–351.
- Geisser, S. (1975). The predictive sample reuse method with applications. *J. Amer. Stat. Assoc.* 70, 320–328.
- Ghanem, R. and P. Spanos (1991). *Stochastic finite elements – A spectral approach*. Springer Verlag, New York. (Reedited by Dover Publications, Mineola, 2003).
- Ghanem, R. G. and P. D. Spanos (2003). *Stochastic finite elements: a spectral approach*. Courier Corporation.
- Giunta, A. A., S. F. Wojtkiewicz, M. S. Eldred, et al. (2003). Overview of modern design of experiments methods for computational simulations. In *Proc. 41st AIAA Aerospace Sciences Meeting and Exhibit, AIAA-2003-0649*.
- Goel, T., R. T. Haftka, W. Shyy, and L. T. Watson (2008). Pitfalls of using a single criterion for selecting experimental designs. *Int. J. Numer. Meth. Eng.* 75(2), 127–155.
- Halton, J. H. (1960). On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numer. Math.* 2(1), 84–90.
- Hesterberg, T., N. H. Choi, L. Meier, C. Fraley, et al. (2008). Least angle and l1 penalized regression: A review. *Stat. Surv.* 2, 61–93.
- Iooss, B., L. Boussouf, V. Feuillard, and A. Marrel (2010). Numerical studies of the metamodel fitting and validation processes. *Int. J. Adv. Sys. Meas.* 3, 11–21.
- Ishigami, T. and T. Homma (1990). An importance quantification technique in uncertainty analysis for computer models. In *Proc. ISUMA'90, First Int. Symp. Unc. Mod. An.*, pp. 398–403. University of Maryland.
- Jakeman, J., M. Eldred, and K. Sargsyan (2015). Enhancing ℓ_1 -minimization estimates of polynomial chaos expansions using basis selection. *J. Comput. Phys.* 289, 18–34.

- Jin, R., W. Chen, and A. Sudjianto (2002). On sequential sampling for global metamodeling in engineering design. In *ASME 2002 Proc. Int. Des. Eng. Tech. Conf. Comput. Inform. Eng. Conf.*, pp. 539–548. American Society of Mechanical Engineers.
- Johnson, M. E., L. M. Moore, and D. Ylvisaker (1990). Minimax and maximin distance designs. *J. Stat. Plan. Infer.* 26(2), 131–148.
- Jones, D. R., M. Schonlau, and W. J. Welch (1998). Efficient global optimization of expensive black-box functions. *J. Global Optim.* 13(4), 455–492.
- Kersaudy, P. (2015). *Modélisation statistique de l'exposition humaine aux ondes radiofréquences*. Ph. D. thesis, Université Paris-Est.
- Konakli, K. and B. Sudret (2016a). Global sensitivity analysis using low-rank tensor approximations. *Reliab. Eng. Syst. Safe* 156, 64 – 83.
- Konakli, K. and B. Sudret (2016b). Polynomial meta-models with canonical low-rank approximations: Numerical insights and comparison to sparse polynomial chaos expansions. *J. Comput. Phys.*
- Konakli, K. and B. Sudret (2016c). Reliability analysis of high-dimensional models using low-rank tensor approximations. *Probab. Eng. Mech.* 46, 18–36.
- Kucherenko, S., D. Albrecht, and A. Saltelli (2015). Exploring multi-dimensional spaces: a comparison of Latin Hypercube and quasi Monte Carlo sampling techniques. *arXiv preprint arXiv:1505.02350*.
- Lin, Y., D. Luo, T. Bailey, R. Khire, J. Wang, and T. W. Simpson (2008). Model validation and error modeling to support sequential sampling. In *ASME 2008 Proc. Int. Des. Eng. Tech. Conf. Comput. Inform. Eng. Conf.*, pp. 1255–1264. American Society of Mechanical Engineers.
- Loève, M. (1977). *Probability theory*. 1977.
- Mallat, S. G. and Z. Zhang (1993). Matching pursuits with time-frequency dictionaries. *IEEE Trans. Signal Process* 41(12), 3397–3415.
- Marelli, S. and B. Sudret (2014). UQLab: A framework for uncertainty quantification in Matlab. In *Vulnerability, Uncertainty, and Risk (Proc. 2nd Int. Conf. on Vulnerability, Risk Analysis and Management (ICVRAM2014), Liverpool, United Kingdom)*, pp. 2554–2563.
- 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.

- McKay, M. D., R. J. Beckman, and W. J. Conover (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 2, 239–245.
- Miller, A. J. and N.-K. Nguyen (1994). Algorithm as 295: A fedorov exchange algorithm for d-optimal design. *J. R. Stat Soc. series c (applied statistics)* 43(4), 669–677.
- Morokoff, W. J. and R. E. Caflisch (1994). Quasi-random sequences and their discrepancies. *Siam. J. Sci. Comput.* 15(6), 1251–1279.
- Myers, R. and D. Montgomery (2002). *Response surface methodology: process and product optimization using designed experiments* (2nd ed.). J. Wiley & Sons.
- Prinzato, L. and W. G. Müller (2012). Design of computer experiments: space filling and beyond. *Stat. Comput* 22(3), 681–701.
- Ravi, S., D. J. Rosenkrantz, and G. K. Tayi (1991). Facility dispersion problems: Heuristics and special cases. In *Algorithms and Data Structures*, pp. 355–366. Springer.
- Rennen, G. (2009). Subset selection from large datasets for Kriging modeling. *Struct. Multidiscip. O.* 38(6), 545–569.
- Rosenblatt, M. (1952). Remarks on a multivariate transformation. *Ann. Math. Stat.* 23(3), 470–472.
- Shin, Y. and D. Xiu (2016). Nonadaptive quasi-optimal points selection for least-squares linear regression. *Siam. J. Sci. Comput.* 38(1), A385–A411.
- Simpson, T. W., D. K. Lin, and W. Chen (2001). Sampling strategies for computer experiments: design and analysis. *Int. J. Rel. App.* 2(3), 209–240.
- Simpson, T. W., J. Poplinski, P. N. Koch, and J. K. Allen (2001). Metamodels for computer-based engineering design: survey and recommendations. *Eng. Comput.* 17(2), 129–150.
- Sobol, I. The distribution of points in a cube and the approximate evaluation of integrals, zh. vychisl. mat. i mat. fiz. 7 (1967), 784–802.
- Stein, M. (1987). Large sample properties of simulations using Latin Hypercube sampling. *Technometrics* 29(2), 143–151.
- Stone, M. (1974). Cross-validated choice and assessment of statistical predictions. *J. Royal Stat. Soc., Series B* 36, 111–147.

- Sudret, B. (2008). Global sensitivity analysis using polynomial chaos expansions. *Reliab. Eng. Sys. Safety* 93, 964–979.
- Sudret, B. and A. Der Kiureghian (2000). Stochastic finite elements and reliability: a state-of-the-art report. Technical Report UCB/SEMM-2000/08, University of California, Berkeley (173 pages).
- Wang, G. (2003). Adaptive response surface method using inherited Latin Hypercube design points. *J. Mech. Design*. 125, 210–220.
- Xiu, D. (2010). *Numerical methods for stochastic computations – A spectral method approach*. Princeton University press.
- Xiu, D. and J. S. Hesthaven (2005). High-order collocation methods for differential equations with random inputs. *Siam. J. Sci. Comput.* 27(3), 1118–1139.
- Zein, S., B. Colson, and F. Glineur (2013). An efficient sampling method for regression-based polynomial chaos expansion. *Commun. Comput. Phys.* 13(04), 1173–1188.