Noname manuscript No.
(will be inserted by the editor)

Sparse Pseudo Spectral Projection Methods with Directional Adaptation for Uncertainty Quantification

J. Winokur · D. Kim · F. Bisetti · O.P. Le Maître · O.M. Knio

Received: date / Accepted: date

Abstract We investigate two methods to build a polynomial approximation of a model output depending on some parameters. The two approaches are based on pseudo-spectral projection (PSP) methods on adaptively constructed sparse grids, and aim at providing a finer control of the resolution along two distinct subsets of model parameters. The control of the error along different subsets of parameters may be needed for instance in the case of a model depending on uncertain parameters and deterministic design variables. We first consider a nested approach where an independent adaptive sparse grid PSP is performed along the first set of directions only, and at each point a sparse grid is constructed adaptively in the second set of directions. We then consider the application of aPSP in the space of all parameters, and introduce directional refinement criteria to provide a tighter control of the projection error along individual dimensions. Specifically, we use a Sobol decomposition of the projection surpluses to tune the sparse grid adaptation. The behavior and performance of the two approaches are compared for a simple two-dimensional test problem and for a shock-tube ignition model involving 22 uncertain parameters and 3 design parameters. The numerical experiments indicate that whereas both methods provide effective means for tuning the quality of the representation along distinct subsets of parameters, PSP in the global parameter space generally requires fewer model evaluations than the nested approach to achieve similar projection error. In addition, the global approach is better suited for generalization to more than two subsets of directions.

Keywords Uncertainty Quantification · Polynomial Chaos · Adaptive Sparse Grids · Pseudo-Spectral Approximation · Chemical Kinetics

1 Introduction

Recent computational experiences [29,28] have clearly put into evidence that the availability of accurate model surrogates can provide substantial advantages in inference or parameter problems. This stems from the fact that it is substantially more effective to evaluate the surrogate than the model itself, especially when the complexity, scale, or dimensionality of the model is large. These advantages have in fact been demonstrated in different settings, including stochastic approaches based on random sampling of the surrogate [41], or in variational approaches where the surrogate is evaluated iteratively [40].

J. Winokur
Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708, USA
E-mail: Justin.Winokur@Duke.edu

D. Kim
King Abdullah University of Science and Techology, Thuwal 23955-6900, Saudi Arabia
E-mail: daesang.kim@kaust.edu.sa

F. Bisetti
King Abdullah University of Science and Techology, Thuwal 23955-6900, Saudi Arabia
E-mail: fabrizio.bisetti@kaust.edu.sa

O.P. Le Maître
LIMSI-CNRS (UPR 3251), Orsay, France
Visiting Professor: Department of Mechanical Engineering and Materials Science, Duke University
E-mail: olm@limsi.fr

O.M. Knio
Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708, USA
Present address: King Abdullah University of Science and Techology, Thuwal 23955-6900, Saudi Arabia
E-mail: omar.knio@duke.edu

Another situation where the availability of suitable surrogates is advantageous concerns problems of optimal experimental design. Frequently, these aim at maximizing the information gain from a set of experiments [27], with estimates involving integrals over both the space of experimental design variables and the space of random parameters [2,18,36]. In these and other settings, it is highly desirable to have available *global* surrogates that provide suitable representations of quantities of interest (QoIs) as function of both deterministic design variables and of random model inputs. This topic is central to present work.

We specifically focus on a polynomial chaos methodology in our approach to such constructions. Polynomial chaos expansions (PCEs) are Fourier-like expansions that express the dependence of QoIs on random inputs in terms of a series of orthogonal basis functions. They are orthogonal with respect to the canonical random variables used to parametrize the uncertain model inputs [14,24]. PCEs have been effectively used in a variety of applications (e.g. [22,30,26]) and have proven to be effective in forward and inverse problems [19,28,41] and sensitivity analyses [42,1,6].

We focus on so-called non-intrusive techniques, which essentially amount to performing an ensemble of deterministic model simulations to estimate the expansion coefficients. Various methods have been used for this purpose, including regression techniques [17], compressed sensing [8], and non-intrusive spectral projection (NISP) [33,25,24]. Of these, we will exclusively consider NISP techniques, which rely on quadratures to project the QoIs onto the polynomial basis.

In this context, we explore two approaches for constructing a global polynomial representation of QoIs. The first is a nested approach, similar to that developed by *Eldred et al.* [10,9] who relied on nested iterations between two spaces ("aleatory-epistemic" uncertainty) with the goal of interval estimation. In our work, we seek for a global representation over the product space. Specifically, we perform an adaptation in a first set of directions ("outer directions"), and at each of the corresponding grid points perform NISP in the remaining directions ("inner directions"). To accommodate potentially high-dimensional parametrization, we rely on an adaptive pseudo-spectral construction (aPSP) with well-established sparse tensorization concepts [38,12]. The aPSP method [5,4] eliminates internal aliasing by applying a Smolyak-type sparse tensor-product approximation to the projection operator instead of the classical quadrature approximation. The aPSP also accommodates non-isotropic sparse grid adaptation wherein the tensorization set is iteratively refined using a greedy algorithm [13,4]. The aPSP has been shown to reduce the number of model evaluation needed for the construction while *increasing* accuracy [44] and has been applied in practical settings [18,41,40].

An alternative approach is also explored, namely by considering the product space of two subsets of parameters (*e.g.* design variables and stochastic parameters). In this second approach, we perform a single adaptive projection in a product space allowing for increased sparsity in the quadrature grids. However, unlike the nested approach, there is no longer direct control of the fidelity of the representation in the two subsets of parameters individually. To overcome this hurdle, we develop a Sobol decomposition [39,16] of the projection surpluses to derive directional indicators to tune adaptation and termination criteria.

This paper is as organized follows. In Section 2, we introduce polynomial chaos expansions (PCEs) and outline the construction of the aPSP algorithm. In Section 3, we present the two approaches for obtaining a global representation of QoIs in terms of two independent sets of parameters. In Section 4, we employ a simple test problem to assess the performance of the two techniques, and to analyze the quality of resulting representations. In Section 5 we perform a large-scale demonstration for a stiff chemical system involving three design variables and 22 stochastic parameters. Major conclusions are summarized in Section 6.

2 Background

2.1 Polynomial Chaos Expansions

One of our main objectives in the present work is to determine an approximation of a QoI, F, depending on a high-dimensional random vector, $\boldsymbol{\xi}$, in a PC basis given by multivariate polynomials in $\boldsymbol{\xi}$. This section discusses the framework used for this purpose.

2.1.1 Polynomial Chaos Basis

In this work, we restrict our attention to PCEs based on a suitable parameterization of model inputs in terms of a $d-$dimensional real-valued random vector $\boldsymbol{\xi} \doteq (\xi_1 \cdots \xi_d)$. The random vector $\boldsymbol{\xi}$ is assumed to have independent components; we denote $\Xi \doteq \Xi_1 \times \cdots \times \Xi_d \subseteq \mathbb{R}^d$ the range of $\boldsymbol{\xi}$ and $\rho : \Xi \mapsto \mathbb{R}_+$ its probability density function,

$$\rho(\boldsymbol{\xi}) \doteq \rho_1(\xi_1)\cdots\rho_d(\xi_d), \quad \int_{\Xi_i} \rho_i(\xi_i) d\xi_i = 1 \text{ for } i = 1,\cdots,d. \tag{2.1}$$

For given ρ and Ξ, we introduce $L_2(\Xi,\rho)$ the space of real-valued, second-order random variables in $\boldsymbol{\xi}$, such that

$$U(\boldsymbol{\xi}) \in L_2(\Xi,\rho) \Leftrightarrow \|U\|_{L_2(\Xi,\rho)} < \infty, \quad \|U\|^2_{L_2(\Xi)} \doteq \langle U,U \rangle, \tag{2.2}$$

where $\langle \cdot, \cdot \rangle$ denotes the natural inner product in $L_2(\Xi, \rho)$:

$$\forall U, V \in L_2(\Xi, \rho)^2, \ \langle U, V \rangle \doteq \int_{\Xi} U(\boldsymbol{\xi}) V(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}. \tag{2.3}$$

For $i = 1, \cdots, d$, let $\{\psi_0^i, \psi_1^i, \dots\}$ be the set of orthonormal polynomials based on the measure ρ_i, where the lower index, k, in ψ_k^i refers to the polynomial degree, such that $\psi_0^i = 1$. For $\boldsymbol{k} = (k_1 \cdots k_d) \in \mathbb{N}_0^d$, we define

$$\Psi_{\boldsymbol{k}}(\boldsymbol{\xi}) \doteq \prod_{i=1}^{d} \psi_{k_i}^i(\xi_i), \tag{2.4}$$

the d-variate polynomial in $\boldsymbol{\xi}$ whose total degree is $|\boldsymbol{k}| \doteq \sum_{i=1}^{d} k_i$. Therefore, the measure ρ defined in (2.1) induces the type of multivariate polynomials $\Psi_{\boldsymbol{k}}$ in (2.4). It is immediate to verify that $\{\Psi_{\boldsymbol{k}}, \boldsymbol{k} \in \mathbb{N}_0^d\}$ is an orthonormal polynomial basis of $L_2(\Xi, \rho)$.

The Polynomial Chaos (PC) approximation of $F \in L_2(\Xi, \rho)$ is

$$F(\boldsymbol{\xi}) \approx \sum_{\boldsymbol{k} \in \mathscr{K}} f_{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}), \tag{2.5}$$

where $\mathscr{K} \subset \mathbb{N}_0^d$ is a prescribed set of multi-indices defining the truncated PC basis and $f_{\boldsymbol{k}}$ are deterministic coefficients. In this work we restrict ourselves to the case of uniform probabilitiy measures, leading to constant densities, ρ_i, and a basis corresponding to the tensorization of translated and rescaled Legendre polynomials ψ_k^i. Note, however, that the methodology may be naturally extended to any tensorized orthonormal basis. In the following, we shall omit, when not necessary, references to specific dimensions $i = 1, \cdots, d$.

2.1.2 Non-Intrusive Spectral Projection

To determine the coefficients in (2.5), we consider so-called non-intrusive methods where the coefficients of a quantity of interest F are determined from a finite ensemble of deterministic model evaluations (or realizations) for different values of $\boldsymbol{\xi} \in \Xi$ [24]. We focus on the Non-Intrusive Spectral Projection (NISP) [33,25] approach, which consists of requiring the approximation error $F(\boldsymbol{\xi}) - \sum_{\boldsymbol{k} \in \mathscr{K}} f_{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi})$ to be orthogonal to the span of the polynomial basis. Using the orthonormal character of the basis, this leads to

$$\left\langle F - \sum_{\boldsymbol{k}' \in \mathscr{K}} f_{\boldsymbol{k}'} \Psi_{\boldsymbol{k}'}, \Psi_{\boldsymbol{k}} \right\rangle = 0 \Rightarrow f_{\boldsymbol{k}} = \langle F, \Psi_{\boldsymbol{k}} \rangle = \int_{\Xi} F(\boldsymbol{\xi}) \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi} \quad \forall \, \boldsymbol{k} \in \mathscr{K}. \tag{2.6}$$

Classically, the integrals in the right-hand side of (2.6) are approximated using an N_q point quadrature rule of the form

$$f_{\boldsymbol{k}} = \int_{\Xi} F(\boldsymbol{\xi}) \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi} \approx \sum_{q=1}^{N_q} F(\boldsymbol{\xi}^{(q)}) \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}^{(q)}) w^{(q)}, \tag{2.7}$$

where the $\{\boldsymbol{\xi}^{(q)}, q = 1, \cdots, N_q\}$ is the set of quadrature points having associated weights $w^{(q)}$. Equation (2.7) shows that the main computational burden of NISP is to compute N_q model evaluations for $F(\boldsymbol{\xi}^{(q)})$, so the complexity of the method is $\mathscr{O}(N_q)$. In traditional tensor-product quadratures, N_q scales exponentially with the number of dimensions d (the so-called "Curse of Dimensionality") and is intractable for all but very low dimensional problems [24]. This issue has motivated the use of sparse grid methods [20], which significantly reduce the computation burden by reducing the number of model evaluations to construct the PC approximation of F.

2.2 Sparse Pseudo-Spectral Projection

2.2.1 Direct Spectral Projection

The direct spectral projection approach is based on a straightforward tensor-product construction in which all integrands are approximately identically using all available nodes and weights. This is briefly outlined below.

Consider $\mathrm{Q}_1, \mathrm{Q}_2, \dots$ a sequence of 1-D quadrature formulas having increasing polynomial exactness, denoted p_l for Q_l, that is

$$\int G(\xi) \rho(\xi) d\xi = \mathrm{Q}_l G = \sum_{q=1}^{N(l)} G\left(\xi^{(q,l)}\right) w_{1D}^{(q,l)}, \quad \forall G \in \pi_{p_l},$$

where π_{p_l} the set of polynomials of degree less or equal to p_l and $N(l)$ is the number of points in the formula Q_l. We call l the level of the formula Q_l. Prescribing the multi-index $\boldsymbol{l} = (l_1 \cdots l_d) \in \mathbb{N}_+^d$, the full-tensor-product (FT) quadrature of an integrable function, G, can be written as follows:

$$\begin{aligned} \mathrm{Q}_{\boldsymbol{l}}^{FT} G &= \left(\mathrm{Q}_{l_1} \otimes \cdots \otimes \mathrm{Q}_{l_d}\right) G = \sum_{q_1=1}^{N(l_1)} \cdots \sum_{q_d=1}^{N(l_d)} G\left(\xi_1^{(q_1,l_1)}, \cdots, \xi_d^{(q_d,l_d)}\right) w_{1D}^{(q_1,l_1)} \cdots w_{1D}^{(q_d,l_d)} \\ &= \sum_{q=1}^{N(\boldsymbol{l})} G(\boldsymbol{\xi}^{(q,\boldsymbol{l})}) w_{\boldsymbol{l}}^{(q,\boldsymbol{l})} \end{aligned} \tag{2.8}$$

where $N(\boldsymbol{l}) = \prod_{i=1}^d N(l_i)$. (Note that G generically denotes both univariate and multivariate functions.) The (FT) quadrature rule can be used in (2.7) to compute the NISP coefficients $f_{\boldsymbol{k}}$, using $G \equiv F\Psi_{\boldsymbol{k}}$, $\boldsymbol{k} \in \mathscr{K}$. We call using the same quadrature rule to compute all the coefficients $f_{\boldsymbol{k}\in\mathscr{K}}$, the Direct Spectral Projection (DSP).

Note that in the case of uniform probability measures, the use of *nested sequences*, for instance based on the Clenshaw-Curtis [3,43], Fejer, [11] and Gauss-Patterson-Kronrod [23,31] rules, is often preferred. The complexity of $\mathrm{Q}_{\boldsymbol{l}}^{FT}$ for such nested sequences is $N(\boldsymbol{l})$, and so it increases exponentially with d. Sparse grids mitigate this complexity by first introducing the 1-D difference formulas between two successive level, $\Delta_l^{\mathrm{Q}} = Q_l - Q_{l-1}$, $\Delta \mathrm{Q}_1 = \mathrm{Q}_1$, such that (2.8) can be recast as

$$\begin{aligned} \mathrm{Q}_{\boldsymbol{l}}^{FT} G &= \left(\mathrm{Q}_{l_1} \otimes \cdots \otimes \mathrm{Q}_{l_d}\right) G \\ &= \left(\left(\sum_{i_1=1}^{l_1} \Delta_{i_1}^{\mathrm{Q}}\right) \otimes \cdots \otimes \left(\sum_{i_d=1}^{l_d} \Delta_{i_d}^{\mathrm{Q}}\right)\right) G = \sum_{i_1=1}^{l_1} \cdots \sum_{i_d=1}^{l_d} \left(\Delta_{i_1}^{\mathrm{Q}} \otimes \cdots \otimes \Delta_{i_d}^{\mathrm{Q}}\right) G \\ &= \sum_{\boldsymbol{i} \in \mathscr{L}_{\boldsymbol{l}}^{FT}} \left(\Delta_{i_1}^{\mathrm{Q}} \otimes \cdots \otimes \Delta_{i_d}^{\mathrm{Q}}\right) G = \sum_{\boldsymbol{i} \in \mathscr{L}_{\boldsymbol{l}}^{FT}} \boldsymbol{\Delta}_{\boldsymbol{i}}^{\mathrm{Q}} G, \end{aligned} \tag{2.9}$$

where $\mathscr{L}_{\boldsymbol{l}}^{FT} = \{\boldsymbol{i} \in \mathbb{N}_+^d, i_j \le l_j \text{ for } j = 1, \cdots, d\}$ is the (FT) multi-index set of tensorizations of 1D difference formulas, i.e. of Δ_l's. (We shall refer to these simply as "tensorizations".) The sparse quadrature rule $\mathrm{Q}_{\mathscr{L}}$ is finally constructed by considering the summation over a subset $\mathscr{L}$ of tensorized quadrature differences:

$$\mathrm{Q}_{\mathscr{L}} G = \sum_{\boldsymbol{i} \in \mathscr{L}} \boldsymbol{\Delta}_{\boldsymbol{i}}^{\mathrm{Q}} G, \quad \mathscr{L} \subset \mathscr{L}_{\boldsymbol{l}}^{FT}. \tag{2.10}$$

The set of tensorizations $\mathscr{L}$ must be admissible in the sense that the following condition holds [13,12]

$$\forall\, \boldsymbol{i} = (i_1 \cdots i_d) \in \mathscr{L} : i_{1 \le j \le d} > 1 \Rightarrow \boldsymbol{i} - \hat{\mathrm{e}}_j \in \mathscr{L},$$

where $\{\hat{\mathrm{e}}_j, j = 1, \cdots, d\}$ is the canonical unit vectors of $\mathbb{N}_0^d$. This admissibility condition is necessary to preserve the telescopic property of the sum of quadrature differences. In the following, we denote by $\mathscr{G}(\mathscr{L})$ the set of nodes in the sparse grid,

$$\mathscr{G}(\mathscr{L}) \doteq \bigcup_{\boldsymbol{i} \in \mathscr{L}} \left\{ \left(\xi^{(q_1,i_1)}, \cdots, \xi^{(q_d,i_d)}\right), 1 \le q_j \le N(i_j), 1 \le j \le d \right\}, \quad N(\mathscr{L}) = |\mathscr{G}(\mathscr{L})|, \tag{2.11}$$

where $|\cdot|$ is the cardinality of a set.

A quadrature rule $\mathrm{Q}_{\mathscr{L}}$ is said to be *sufficiently exact* with respect to a polynomial multi-index set $\mathscr{K}$ when for any couple $(\boldsymbol{k}, \boldsymbol{k}') \in \mathscr{K} \times \mathscr{K}$, the basis orthonormality conditions is recovered using the discrete inner product:

$$\langle U, V \rangle_{\mathrm{Q}_{\mathscr{L}}} \doteq \mathrm{Q}_{\mathscr{L}}(UV) \tag{2.12}$$

In other words, $\mathrm{Q}_{\mathscr{L}}$ is sufficiently exact if for any $(\boldsymbol{k}, \boldsymbol{k}') \in \mathscr{K} \times \mathscr{K}$, we have:

$$\langle \Psi_{\boldsymbol{k}}, \Psi_{\boldsymbol{k}'} \rangle_{\mathrm{Q}_{\mathscr{L}}} = \begin{cases} 1 \text{ when } \boldsymbol{k} = \boldsymbol{k}' \\ 0 \text{ otherwise} \end{cases}$$

This immediately implies that when U and V belong to the space spanned by the basis, their discrete and continuous inner products coincide. If $\mathrm{Q}_{\mathscr{L}}$ is not sufficiently exact, severe and detrimental "internal aliasing" effects will ensue when performing DSP [44,4]. (By internal aliasing, we refer to errors in numerical inner product between elements lying in the span of the basis, whereas external aliasing occurs as a result of truncation.) In the case of FT quadratures, the largest set, $\mathscr{K}^*(\boldsymbol{l})$, for which internal aliasing errors vanish can be easily determined from the degrees of polynomial exactness of the 1-D sequence, using the so-called half-accuracy set. Specifically, $\mathscr{K}^*(\boldsymbol{l})$ is given by

$$\mathscr{K}^*(\boldsymbol{l}) = \left\{ \boldsymbol{k} \in \mathbb{N}_0^d : k_i \le p_{l_i}/2,\ i = 1, \cdots, d \right\}. \tag{2.13}$$

We shall denote by $\mathrm{P}_{\boldsymbol{l}}^{FT} F$ the (FT) projection operator of $F \in L_2(\Xi, \rho)$ onto the span of $\{\Psi_{\boldsymbol{k}}, \boldsymbol{k} \in \mathscr{K}^*(\boldsymbol{l})\}$, namely

$$\mathrm{P}_{\boldsymbol{l}}^{FT} F \doteq \sum_{\boldsymbol{k} \in \mathscr{K}^*(\boldsymbol{l})} f_{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}), \quad f_{\boldsymbol{k}} = \mathrm{Q}_{\boldsymbol{l}}^{FT}(F \Psi_{\boldsymbol{k}}). \tag{2.14}$$

An immediate consequence of using the sparse quadrature formula with $\mathscr{L} \subset \mathscr{L}_{\boldsymbol{l}}^{FT}$, is that $\mathrm{Q}_{\mathscr{L}}$ is not sufficiently exact with respect to $\mathscr{K}^*(\boldsymbol{l})$, as the sparse formula is not able to exactly integrate the product of high-order monomials. We can still define largest sets $\mathscr{K}^*(\mathscr{L})$ for which the sparse quadrature is sufficiently exact, but these sets are usually non-unique leading to ambiguous maximal internal aliasing-free projection spaces. In any case, we observe that the use of sparse quadrature in place of FT ones, while reducing the complexity ($|\mathscr{G}(\mathscr{L})| \ll \prod_{j=1}^{d} N(l_j)$), leads also to a significant reduction of the polynomial space allowing for internal aliasing-free DSP ($|\mathscr{K}^*(\mathscr{L})| \ll |\mathscr{K}^*(\boldsymbol{l})|$). The sparse pseudo-spectral projection (PSP) method presented below allows one to consider polynomial spaces larger than defined $\mathscr{K}^*(\mathscr{L})$, for the same sparse grid $\mathscr{G}(\mathscr{L})$ and without introducing internal aliasing.

2.2.2 *Sparse Pseudo-Spectral Projection*

The sparse PSP method relies on the definition of a sequence of *1-D* projection operators [5,4]. Let $\{\mathrm{P}_{l \geq 1}\}$ denote the sequence of 1-D projection operators associated to the sequence $\{\mathrm{Q}_{l \geq 1}\}$ of 1-D quadrature, where

$$\mathrm{P}_l : F(\xi) \mapsto \mathrm{P}_l F(\xi) \equiv \sum_{k=0}^{p_l/2} f_k \psi_k(\xi) \in \pi_{p_l/2}, \quad f_k = \sum_{q=1}^{N(l)} f(\xi^{(q)}) \psi_k(\xi^{(q,l)}) w_{1D}^{(q,l)}. \tag{2.15}$$

Note that $\mathrm{P}_{l \geq 1}$ is free of internal aliasing, owing to the dependence with l of the projection spaces, which ensures that the quadrature is sufficiently exact. Consider for $l \geq 1$ the difference of successive 1-D projection operators,

$$\Delta_{l=1}^{\mathrm{P}} = \mathrm{P}_1, \quad \Delta_{l>1}^{\mathrm{P}} = \mathrm{P}_l - \mathrm{P}_{l-1},$$

such that $\mathrm{P}_l = \sum_{i=1}^{l} \Delta \mathrm{P}_i$. The FT projection operator in (2.14) can be recast in terms of a sum of tensorized difference projection operators as follows

$$\begin{aligned} \mathrm{P}_{\boldsymbol{l}}^{FT} F &= \left(\mathrm{P}_{l_1} \otimes \cdots \otimes \mathrm{P}_{l_d}\right) F \\ &= \left(\left(\sum_{i_1=1}^{l_1} \Delta_{i_1}^{\mathrm{P}} \right) \otimes \cdots \otimes \left(\sum_{i_d=1}^{l_d} \Delta_{i_d}^{\mathrm{P}} \right) \right) F = \sum_{i_1=1}^{l_1} \cdots \sum_{i_d=1}^{l_d} \left(\Delta_{i_1}^{\mathrm{P}} \otimes \cdots \otimes \Delta_{i_d}^{\mathrm{P}}\right) F \\ &= \sum_{\boldsymbol{i} \in \mathscr{L}_{\boldsymbol{l}}^{FT}} \left(\Delta_{i_1}^{\mathrm{P}} \otimes \cdots \otimes \Delta_{i_d}^{\mathrm{P}}\right) F. \end{aligned} \tag{2.16}$$

The sparse PSP operator $\mathrm{P}_{\mathscr{L}}$ is finally obtained by considering a summation over an admissible subset $\mathscr{L} \subset \mathscr{L}_{\boldsymbol{l}}^{FT}$ of tensorized difference projection operators. This results in

$$\mathrm{P}_{\mathscr{L}} F = \sum_{\boldsymbol{i} \in \mathscr{L}} \boldsymbol{\Delta}_{\boldsymbol{i}}^{\mathrm{P}} F, \quad \boldsymbol{\Delta}_{\boldsymbol{i}}^{\mathrm{P}} F \doteq \left(\Delta_{i_1}^{\mathrm{P}} \otimes \cdots \otimes \Delta_{i_d}^{\mathrm{P}}\right) F. \tag{2.17}$$

The key point here is that the sparse PSP operator in (2.17) involves a telescopic sum of differences in the projection of F onto subspaces of increasing dimensions, where each individual (FT) tensorized quadrature projection is evaluated without internal aliasing. As a result, the sparse PSP in (2.17) is free of internal aliasing. Further, $\mathrm{P}_{\mathscr{L}} F$ belongs to the span of $\{\Psi_{\boldsymbol{k}}, \boldsymbol{k} \in \mathscr{K}(\mathscr{L})\}$, where $\mathscr{K}(\mathscr{L}) = \bigcup_{\boldsymbol{i} \in \mathscr{L}} \mathscr{K}^*(\boldsymbol{i})$. We have $\mathscr{K}(\mathscr{L}) \supseteq \mathscr{K}^*(\mathscr{L})$ and the projection space of the sparse PSP is usually significantly larger than that of the DSP, while the two methods have the same complexity relying both on the same sparse grid $\mathscr{G}(\mathscr{L})$. Note, however, that the global sparse quadrature rule $\mathrm{Q}_{\mathscr{L}}$ is generally not sufficiently accurate with respect to $\mathscr{K}(\mathscr{L})$. The inclusion of the sets $\mathscr{K}(\mathscr{L})$ and $\mathscr{K}^*(\mathscr{L})$ is illustrated in Fig. 2.1.

2.3 Adaptive Pseudo-Spectral Projection

The PSP method does not, in itself, guarantee a low projection error if the projection space is not rich enough (external aliasing). To address this issue, we adapt the method proposed in [13] to the context of the sparse PSP. The adaptive PSP method was first introduced in [4,44] and subsequently used in [41,18] in the context of Bayesian inference and experimental design. The algorithm is explained in detail in [44,4,18] and is only briefly outlined below.

We begin by defining two distinct multi-index sets, the set $\mathscr{O}$ of old multi-indices and the set $\mathscr{A}$ of active multi-indices, such that $\mathscr{L} = \mathscr{O} \cup \mathscr{A}$ is admissible, and $\mathscr{A}$ is $\mathscr{O}$-admissible, that is

$$\forall \boldsymbol{l} \in \mathscr{A}, \quad l_i > 1 \Rightarrow \boldsymbol{l} - \hat{\mathrm{e}}_i \in \mathscr{O}, \quad \text{for } i = 1, \cdots, d.$$

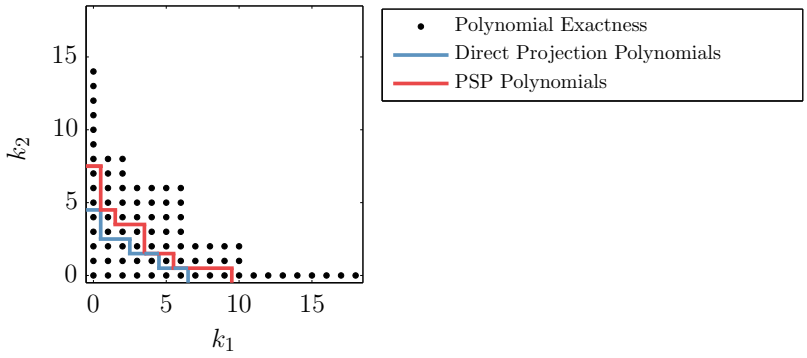

Fig. 2.1 Comparison of polynomials multi-index sets for DSP and PSP using a quadrature rule $\mathrm{Q}_{\mathscr{L}}$ in $d = 2$ dimensions. Points are multi-indices (k_1, k_2) such $G \in \pi_{k_1} \times \pi_{k_2}$ is exactly integrated by $\mathrm{Q}_{\mathscr{L}}$. Multi-indices of the DSP set $\mathscr{K}^*(\mathscr{L})$ are below the blue line, while elements of the PSP set $\mathscr{K}(\mathscr{L})$ are below the red-line. In this illustrative example, $\mathscr{L}$ is the region bounded by the indices (0,4), (1,2), (2,2), (3,1), (4,1) and (6,0). The Gauss-Kronrod-Patterson rule is used to determine the depicted boundaries of $\mathscr{K}(\mathscr{L})$ and $\mathscr{K}^*(\mathscr{L})$.

The set $\mathscr{A}$ contains the multi-indices available for adaption. Associated with each $\boldsymbol{l} \in \mathscr{A}$ is a refinement indicator $\varepsilon(\boldsymbol{l}) \geq 0$. In view of (2.17), a natural measure for the contribution of a tensorization $\boldsymbol{l} \in \mathscr{L}$ to the projection is the $L_2(\Xi, \rho)$ norm of the associated tensor-product-delta-operator; we take

$$\varepsilon(\boldsymbol{l}) \doteq \left\| \left(\Delta_{l_1}^{\mathrm{P}} \otimes \cdots \otimes \Delta_{l_d}^{\mathrm{P}} \right) F \right\|_{L_2(\Xi,\rho)}. \tag{2.18}$$

The indicator $\varepsilon(\boldsymbol{l})$ is then the L_2-norm of the projection surplus associated to $\boldsymbol{l}$. At each iteration of the adaptive algorithm, the *critical* multi-index $\boldsymbol{l}^* \in \mathscr{A}$ with the highest indicator $\varepsilon(\boldsymbol{l})$ is selected. $\boldsymbol{l}^*$ is removed from $\mathscr{A}$ and added to $\mathscr{O}$. Then, each of the d multi-indices of the forward neighborhood of $\boldsymbol{l}^*$,

$$\mathscr{F}(\boldsymbol{l}^*) \doteq \{\boldsymbol{l}_i^+ = \boldsymbol{l}^* + \hat{\mathbf{e}}_i, i = 1, \cdots, d\}, \tag{2.19}$$

is added to $\mathscr{A}$ provided that $\mathscr{L} \cup \{\boldsymbol{l}_i^+\}$ remains $\mathscr{O}$-admissible. Note that none of the forward neighbors of $\boldsymbol{l}^*$ may be admissible to complete $\mathscr{A}$. When new multi-indices are added to $\mathscr{A}$, the sparse grid $\mathscr{G}(\mathscr{O} \cup \mathscr{A})$ and projection $\mathrm{P}_{\mathscr{O}\cup\mathscr{A}}F$ are updated (involving model evaluations at the new grid points only), and the indicators $\varepsilon(\boldsymbol{l})$ of the new multi-indices are computed. The adaptive procedure can then be repeated.

Also associated with the adaptive procedure is a convergence estimator, $\eta(\mathscr{A})$, defined as

$$\eta^2(\mathscr{A}) = \sum_{\boldsymbol{l} \in \mathscr{A}} \varepsilon^2(\boldsymbol{l}). \tag{2.20}$$

As discussed in [4], $\eta(\mathscr{A})$ is usually an acceptable surrogate for the projection error. Therefore, the adaptivity continues until $\eta < \texttt{Tol}$ or $\mathscr{A} = \emptyset$, where $\texttt{Tol} > 0$ is a user defined error tolerance.

3 Methodology

The aPSP algorithm introduced above does not inherently distinguish between the individual directions, ξ_i. In many situations, however, we would be interested in a finer control of the projection error along certain directions. This is the case for instance when the set of input variables is made of subsets of variables having different origin and nature (uncertainty sources). To remain general, we then consider a partition of $\boldsymbol{\xi}$ into two independent sets of variables $\boldsymbol{\xi}_{\mathrm{u}}$, and $\boldsymbol{\xi}_{\mathrm{p}}$ with respective ranges Ξ_{u} and Ξ_{p}, densities ρ_{u} and ρ_{p}, such that

$$\Xi = \Xi_{\mathrm{u}} \times \Xi_{\mathrm{p}}, \quad \rho(\boldsymbol{\xi}) = \rho_{\mathrm{u}}(\boldsymbol{\xi}_{\mathrm{u}}) \times \rho_{\mathrm{p}}(\boldsymbol{\xi}_{\mathrm{p}}).$$

We shall denote d_{u} and d_{p} the dimensions of Ξ_{u} and Ξ_{p}, and we seek to construct in an adaptive way a PCE $F(\boldsymbol{\xi} = (\boldsymbol{\xi}_{\mathrm{p}}, \boldsymbol{\xi}_{\mathrm{u}})) \approx \sum_{\boldsymbol{k} \in \mathscr{K}} f_{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi})$ of some quantity of interest F, while somehow controlling the errors along $\boldsymbol{\xi}_{\mathrm{u}}$ and $\boldsymbol{\xi}_{\mathrm{p}}$ independently.

3.1 Nested Projection

In the Nested Projection (NP), the aPSP method of Section 2.3 is applied in a nested fashion, allowing for independent control of the convergence in $\boldsymbol{\xi}_{\mathrm{u}}$ and $\boldsymbol{\xi}_{\mathrm{p}}$.

3.1.1 Nested PC expansion

For $F(\boldsymbol{\xi}_\mathrm{p},\boldsymbol{\xi}_\mathrm{u}) \in L_2(\Xi_\mathrm{p},\rho_\mathrm{p}) \otimes L_2(\Xi_\mathrm{u},\rho_\mathrm{u})$, the nested expansion consists in considering first a *truncated* spectral expansion of $F(\boldsymbol{\xi}_\mathrm{p},\cdot)$ with respect to $\boldsymbol{\xi}_\mathrm{u}$, referred to hereafter as the inner expansion:

$$F(\boldsymbol{\xi}_\mathrm{p},\boldsymbol{\xi}_\mathrm{u}) \approx \sum_{\boldsymbol{k}\in\mathscr{K}_\mathrm{u}} f^{\boldsymbol{k}}(\boldsymbol{\xi}_\mathrm{p})\Psi_{\boldsymbol{k}}(\boldsymbol{\xi}_\mathrm{u}), \qquad \boldsymbol{\xi}_\mathrm{p} \text{ fixed.} \tag{3.1}$$

We then perform a PC expansion in $\boldsymbol{\xi}_\mathrm{p}$ (outer expansion) of the coefficients $f^{\boldsymbol{k}}(\boldsymbol{\xi}_\mathrm{p})$ in Equation (3.1), namely

$$f^{\boldsymbol{k}}(\boldsymbol{\xi}_\mathrm{p}) \approx \sum_{\boldsymbol{j}\in\mathscr{K}_\mathrm{p}} c_{\boldsymbol{j}}^{\boldsymbol{k}}\Psi_{\boldsymbol{j}}(\boldsymbol{\xi}_\mathrm{p}). \tag{3.2}$$

Substituting Equation (3.2) into Equation (3.1), we obtain the global PC expansion of $F(\boldsymbol{\xi}_\mathrm{p},\boldsymbol{\xi}_\mathrm{u})$:

$$F(\boldsymbol{\xi}_\mathrm{p},\boldsymbol{\xi}_\mathrm{u}) \approx \sum_{\boldsymbol{k}\in\mathscr{K}_\mathrm{u}}\sum_{\boldsymbol{j}\in\mathscr{K}_\mathrm{p}} c_{\boldsymbol{j}}^{\boldsymbol{k}}\Psi_{\boldsymbol{j}}(\boldsymbol{\xi}_\mathrm{p})\Psi_{\boldsymbol{k}}(\boldsymbol{\xi}_\mathrm{u}). \tag{3.3}$$

Note that multi-indices in the two sets $\mathscr{K}_\mathrm{u}$ and $\mathscr{K}_\mathrm{p}$ have different sizes in general. We pursue this idea of nested expansion in the context of the PSP method, as outlined below.

3.1.2 Inner projection

Let $\mathscr{G}(\mathscr{L}_\mathrm{p}) = \{\boldsymbol{\xi}_\mathrm{p}^{(i)}, i = 1,\cdots,N(\mathscr{L}_\mathrm{p})\}$ be an admissible sparse grid in Ξ_p, see (2.11). We call $\mathscr{G}(\mathscr{L}_\mathrm{p})$ the outer sparse grid. At each node, $\boldsymbol{\xi}_\mathrm{p}^{(i)}$, of the outer grid the aPSP is employed to approximate $G^{(i)}(\boldsymbol{\xi}_\mathrm{u}) \doteq F(\boldsymbol{\xi}_\mathrm{p}^{(i)},\boldsymbol{\xi}_\mathrm{u})$, thus providing:

$$G^{(i)}(\boldsymbol{\xi}_\mathrm{u}) \approx \sum_{\boldsymbol{k}\in\mathscr{K}_\mathrm{u}^{(i)}} g_{\boldsymbol{k}}^{(i)}\Psi_{\boldsymbol{k}}(\boldsymbol{\xi}_\mathrm{u}). \tag{3.4}$$

Note that this construction is purely local at $\boldsymbol{\xi}_\mathrm{p}^{(i)}$ and involves evaluations of F at a fixed value $\boldsymbol{\xi}_\mathrm{p} = \boldsymbol{\xi}_\mathrm{p}^{(i)} \in \Xi_\mathrm{p}$. The adaptation of the inner sparse grid uses refinement indicators ε_{in} based on $G^{(i)}(\boldsymbol{\xi}_\mathrm{u})$ in Equation (2.18). The inner adaptation is carried out until a prescribed termination criterion $\mathtt{Tol}_\mathrm{u}$ is reached. The same criterion $\mathtt{Tol}_\mathrm{u}$ is used for all nodes of the outer sparse grid. Therefore, different outer nodes $\boldsymbol{\xi}_\mathrm{p}^{(i)}$ may use different inner sparse grids to support the aPSP of their local functions $G^{(i)}$. We denote $\mathscr{G}_\mathrm{u}^{(i)} = \mathscr{G}(\mathscr{L}_\mathrm{u}^{(i)})$ the resulting inner sparse grid at node $\boldsymbol{\xi}_\mathrm{p}^{(i)}$, and $\mathscr{K}_\mathrm{u}^{(i)} \doteq \mathscr{K}(\mathscr{L}_\mathrm{u}^{(i)})$ the multi-index set of polynomial tensorizations in the approximation of $G^{(i)}$.

3.1.3 Outer projection

For the outer sparse grid, we again rely on the aPSP scheme, but with refinement indicators ε_{out} based on $H(\boldsymbol{\xi}_\mathrm{p}) \equiv \|F(\boldsymbol{\xi}_\mathrm{p},\cdot)\|_{L_2(\Xi_\mathrm{u},\rho_\mathrm{u})}$. The function $H(\boldsymbol{\xi}_\mathrm{p})$ is unknown but needs only be evaluated at the nodes $\boldsymbol{\xi}_\mathrm{p}^{(i)}$ of $\mathscr{G}(\mathscr{L}_\mathrm{p})$, where we can take as a surrogate the norm of the approximation of $G^{(i)}(\boldsymbol{\xi}_\mathrm{u})$. This leads to

$$H(\boldsymbol{\xi}_\mathrm{p}^{(i)}) = \|G^{(i)}\|_{L_2(\Xi_\mathrm{u},\rho_\mathrm{u})} \approx \sqrt{\sum_{\boldsymbol{k}\in\mathscr{K}_\mathrm{u}^{(i)}} (g_{\boldsymbol{k}}^{(i)})^2}.$$

The adaptation of the outer sparse grid is carried out until the termination criterion $\mathtt{Tol}_\mathrm{p}$ is reached. Since the inner projections are performed locally, the points $\boldsymbol{\xi}_\mathrm{p}^{(i)}$ support approximations of $G^{(i)}$ defined over generally different polynomial subspaces of $L_2(\Xi_\mathrm{u},\rho_u)$, as prescribed by $\mathscr{K}_\mathrm{u}^{(i)}$. A "global" polynomial subspace of $L_2(\Xi_\mathrm{u},\rho_\mathrm{u})$ can be defined as the union of all the $\mathscr{K}_\mathrm{u}^{(i)}$,

$$\mathscr{K}_\mathrm{u}(\mathscr{L}_\mathrm{p}) \doteq \bigcup_{i=1}^{N(\mathscr{L}_\mathrm{p})} \mathscr{K}_\mathrm{u}^{(i)}, \tag{3.5}$$

and extending the inner approximations as follows

$$G^{(i)}(\boldsymbol{\xi}_\mathrm{u}) \approx \sum_{\boldsymbol{k}\in\mathscr{K}_\mathrm{u}(\mathscr{L}_\mathrm{p})} \tilde{g}_{\boldsymbol{k}}^{(i)}\Psi_{\boldsymbol{k}}(\boldsymbol{\xi}_\mathrm{u}), \quad \tilde{g}_{\boldsymbol{k}}^{(i)} = \begin{cases} g_{\boldsymbol{k}}^{(i)} & \boldsymbol{k}\in\mathscr{K}_\mathrm{u}^{(i)}, \\ 0 & \text{otherwise.} \end{cases} \tag{3.6}$$

The definition of $\mathscr{K}_\mathrm{u}(\mathscr{L}_\mathrm{p})$ enables us to define the global polynomial expansion of F, simply by performing the PSP of the set of inner PC coefficients $\{\tilde{g}_{\boldsymbol{k}}^{(i)}, \boldsymbol{k}\in\mathscr{K}_\mathrm{u}(\mathscr{L}_\mathrm{p})\}$, defined at each node $\boldsymbol{\xi}_\mathrm{p}^{(i)}$ of the outer sparse grid. This operation

results in a vector of coefficients, $c_{\boldsymbol{k}}^{\boldsymbol{j}}$, with $\boldsymbol{k} \in \mathscr{K}_{\mathtt{u}}(\mathscr{L}_{\mathtt{p}})$ and $\boldsymbol{j} \in \mathscr{K}_{\mathtt{p}}(\mathscr{L}_{\mathtt{p}})$, where $\mathscr{K}_{\mathtt{p}}$ denotes the multi-index set of the projection subspace associated with the H-adapted outer sparse grid. This leads to the global representation:

$$F(\boldsymbol{\xi}_{\mathtt{p}}, \boldsymbol{\xi}_{\mathtt{u}}) \simeq \sum_{\boldsymbol{j} \in \mathscr{K}(\mathscr{L}_{\mathtt{p}})} \sum_{\boldsymbol{k} \in \mathscr{K}_{\mathtt{u}}(\mathscr{L}_{\mathtt{p}})} c_{\boldsymbol{k}}^{\boldsymbol{j}} \Psi_{\boldsymbol{j}}(\boldsymbol{\xi}_{\mathtt{p}}) \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}_{\mathtt{u}}). \tag{3.7}$$

Note that by construction, we end up with an expansion based on a full summation over the multi-indices in $\mathscr{K}_{\mathtt{p}}(\mathscr{L}_{\mathtt{p}})$ and in $\mathscr{K}_{\mathtt{u}}(\mathscr{L}_{\mathtt{p}})$, in other words with a structure that is reminiscent of fully tensorized constructions. As previously discussed, this is achieved by setting some (non computed) inner expansion coefficients to zero, see Equation (3.6). Though this could potentially result in errors, it is anticipated that these coefficients would primarily affect higher-order modes having small amplitudes only. Nonetheless, we will examine this issue in Section 4 based on results obtained for an idealized test problem.

Though the fully tensored nature of the $\boldsymbol{\xi}_{\mathtt{p}}$ and $\boldsymbol{\xi}_{\mathtt{u}}$ spaces is expected to be computationally demanding, an advantage of the nested projection is that it provides a clear handle on convergence in the two spaces separately. Note, however, this does not readily translate into estimates of convergence for the global representation in Equation (3.7). Finally, we also note that the outer and inner projections may be inverted in the nested projection depending on the nature of the problem, and that the approach could conceptually be extended to any number of parameter spaces.

3.2 Product Space Projection

In this alternative approach, we focus on expanding $F(\boldsymbol{\xi})$, with $\boldsymbol{\xi} = (\boldsymbol{\xi}_{\mathtt{p}}, \boldsymbol{\xi}_{\mathtt{u}})$. However, applying directly the adapted PSP method to $F(\boldsymbol{\xi})$ would treat all parameters equally. This equal treatment may raise potential issues. For example, F may exhibit larger variations with $\boldsymbol{\xi}_{\mathtt{p}}$ than with $\boldsymbol{\xi}_{\mathtt{u}}$, and consequently the adaptation may ignore essential variations with $\boldsymbol{\xi}_{\mathtt{u}}$. Irrespective of whether such issue arises, it is desirable to incorporate means to separately assess resolution along distinct dimensions, as well as criteria that enable us to tune adaptation along these dimensions.

3.2.1 Theoretical considerations

Recall that the aPSP involves two distinct sets of tensorizations, $\mathscr{O}$ and $\mathscr{A}$, which defines the projection of F through

$$\mathrm{P}_{\mathscr{O} \cup \mathscr{A}} F = \sum_{\boldsymbol{l} \in \mathscr{O}} \boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F + \sum_{\boldsymbol{l} \in \mathscr{A}} \boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F,$$

where $\boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F$ is called the projection surplus associated with the tensorization $\boldsymbol{l}$. Therefore, the second sum in the equation above represents the global projection surplus associated with the active set $\mathscr{A}$ of tensorizations available for adaptation. This global projection surplus, denoted $\delta \mathrm{P}_{\mathscr{A}} F$, belongs to the polynomial subspace defined by the multi-index set $\mathscr{K}(\mathscr{A})$, i.e.

$$\delta \mathrm{P}_{\mathscr{A}} F \equiv \sum_{\boldsymbol{l} \in \mathscr{A}} \boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F = \sum_{\boldsymbol{k} \in \mathscr{K}(\mathscr{A})} \delta f_{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}), \quad \mathscr{K}(\mathscr{A}) = \bigcup_{\boldsymbol{l} \in \mathscr{A}} \mathscr{K}^*(\boldsymbol{l}).$$

We are now in position to measure the relative contributions of $\boldsymbol{\xi}_{\mathtt{p}}$ and $\boldsymbol{\xi}_{\mathtt{u}}$, by means of the orthogonal Sobol decomposition of $\delta \mathrm{P}_{\mathscr{A}} F$ (see Appendix A), that is formally

$$\delta \mathrm{P}_{\mathscr{A}} F(\boldsymbol{\xi}) = (\delta \mathrm{P}_{\mathscr{A}} F)_{\emptyset} + (\delta \mathrm{P}_{\mathscr{A}} F)_{\mathtt{p}}(\boldsymbol{\xi}_{\mathtt{p}}) + (\delta \mathrm{P}_{\mathscr{A}} F)_{\mathtt{u}}(\boldsymbol{\xi}_{\mathtt{u}}) + (\delta \mathrm{P}_{\mathscr{A}} F)_{\mathtt{p} \cup \mathtt{u}}(\boldsymbol{\xi}_{\mathtt{p}}, \boldsymbol{\xi}_{\mathtt{u}}).$$

The squared norm of the global projection surplus is then equal to the squared norms of the functions in the decomposition of $\delta \mathrm{P}_{\mathscr{A}} F$; we have

$$\| \delta \mathrm{P}_{\mathscr{A}} F \|_{L_2(\Xi, \rho)}^2 = V_{\emptyset}^{\mathscr{A}} + V_{\mathtt{p}}^{\mathscr{A}} + V_{\mathtt{u}}^{\mathscr{A}} + V_{\mathtt{p} \cup \mathtt{u}}^{\mathscr{A}}, \tag{3.8}$$

where $V_{\emptyset}^{\mathscr{A}} = (\delta \mathrm{P}_{\mathscr{A}} F)_{\emptyset}^2$, while other terms can be easily computed form the PC expansion of $\delta \mathrm{P}_{\mathscr{A}} F$, as explained in Appendix A.2. Further, $V_{\emptyset}^{\mathscr{A}}$ is the squared surplus of the projection on the mean mode and it is usually small compared to the other contributions to the norm of the global surplus. The partial variance $V_{\mathtt{p}}^{\mathscr{A}}$ (resp. $V_{\mathtt{u}}^{\mathscr{A}}$) then measures the available projection surplus along directions in $\mathtt{p}$ (resp. in $\mathtt{u}$) only, whereas $V_{\mathtt{p} \cup \mathtt{u}}^{\mathscr{A}}$ is the measure along the mixed directions. The partial estimates can be used to tune the adaptation strategy and to gauge accuracy control along the $\mathtt{p}$ and $\mathtt{u}$ directions. For instance, if $V_{\mathtt{u}}^{\mathscr{A}}$ becomes less than a certain tolerance value, we can decide to prevent the adaptive process to further enrich the sparse grid along the hyperplane $\Xi_{\mathtt{u}}$ of Ξ.

3.2.2 Construction of directional indicators

Though feasible, the decomposition of the different contributions to the (squared) norm of the available projection surplus in (3.8) induces some computational overheads, as it requires first the PC expansion of $\delta \mathrm{P}_{\mathscr{A}} F$ which involves a summation of all the projection surpluses $\boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F$ for $\boldsymbol{l} \in \mathscr{A}$. We would prefer to derive indicators of the available surpluses along the $\mathtt{p}$, $\mathtt{u}$ and mixed directions, through a summation of local indicators associated with each $\boldsymbol{l} \in \mathscr{A}$. A first simplification is obtained by performing *locally* the Sobol decomposition of each of the projection surpluses,

$$\boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F = \sum_{\boldsymbol{k} \in \mathscr{K}^*(\boldsymbol{l})} \delta f_{\boldsymbol{k}}^{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}) \Rightarrow \|\Delta_{\boldsymbol{l}}^{\mathrm{P}} F\|^2_{L_2(\Xi,\rho)} = V_{\emptyset}^{\boldsymbol{l}} + V_{\mathrm{p}}^{\boldsymbol{l}} + V_{\mathrm{u}}^{\boldsymbol{l}} + V_{\mathrm{p}\cup\mathrm{u}}^{\boldsymbol{l}},$$

and derive upper-bounds for the decomposition of the available surplus using

$$V_{\emptyset}^{\mathscr{A}} \leq \sum_{\boldsymbol{l}\in\mathscr{A}} V_{\emptyset}^{\boldsymbol{l}}, \quad V_{\mathrm{p}}^{\mathscr{A}} \leq \sum_{\boldsymbol{l}\in\mathscr{A}} V_{\mathrm{p}}^{\boldsymbol{l}}, \quad V_{\mathrm{u}}^{\mathscr{A}} \leq \sum_{\boldsymbol{l}\in\mathscr{A}} V_{\mathrm{u}}^{\boldsymbol{l}}, \quad V_{\mathrm{p}\cup\mathrm{u}}^{\mathscr{A}} \leq \sum_{\boldsymbol{l}\in\mathscr{A}} V_{\mathrm{p}\cup\mathrm{u}}^{\boldsymbol{l}}.$$

We shall consider in the following the case where the polynomial exactness of the 1-D quadrature formulas at level $l = 1$ is $p_1 \leq 1$. This is a common situation in sparse grid methods that for efficiency generally use a low number of 1-D points at the first level, typically $N(1) = 1$ (the first tensorization $\boldsymbol{l} = (1, \cdots, 1)$ involves $N(1)^d$ points, so a low value is mandatory for large d). In this case, the 1-D projection operators $\mathrm{P}_1 F \in \pi_0$ (see equation 2.15), corresponding to the projection onto the space of constant functions. The following observation allows for a simple definition of the surplus decomposition: the surplus $\boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F$ having $l_i = 1$ is a function independent of the variables ξ_i. Further, for multi-indices $\boldsymbol{l} = (\boldsymbol{l}_{\mathrm{p}}, \boldsymbol{l}_{\mathrm{u}})$, where $\boldsymbol{l}_{\mathrm{p}}$ and $\boldsymbol{l}_{\mathrm{u}}$ are sub-multi-indices related to directions $\mathtt{p}$ and $\mathtt{u}$ respectively, the projection surplus $\boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F$ is a function of a subset of variables in $\boldsymbol{\xi}_{\mathrm{p}}$ (resp. $\boldsymbol{\xi}_{\mathrm{u}}$) only, if $|\boldsymbol{l}_{\mathrm{p}}| > d_{\mathrm{p}}$ and $|\boldsymbol{l}_{\mathrm{u}}| = d_{\mathrm{u}}$ (resp. $|\boldsymbol{l}_{\mathrm{p}}| = d_{\mathrm{p}}$ and $|\boldsymbol{l}_{\mathrm{u}}| > d_{\mathrm{u}}$). That is, at least one index of $\boldsymbol{l}_{\mathrm{p}}$ is greater than 1 and all indices of $\boldsymbol{l}_{\mathrm{u}}$ equal 1. Since by construction the multi-index $(1, \cdots, 1) \notin \mathscr{A}$, this suggests the partition of $\mathscr{A}$ into three disjoint subsets

$$\mathscr{A} = \mathscr{A}_{\mathrm{p}} \cup \mathscr{A}_{\mathrm{u}} \cup \mathscr{A}_{\mathrm{p,u}},$$

where

$$\mathscr{A}_{\mathrm{p}} \doteq \{\boldsymbol{l} = (\boldsymbol{l}_{\mathrm{p}}, \boldsymbol{l}_{\mathrm{u}}) \in \mathscr{A}, |\boldsymbol{l}_{\mathrm{u}}| = d_{\mathrm{u}}\}, \quad \mathscr{A}_{\mathrm{u}} \doteq \{\boldsymbol{l} = (\boldsymbol{l}_{\mathrm{p}}, \boldsymbol{l}_{\mathrm{u}}) \in \mathscr{A}, |\boldsymbol{l}_{\mathrm{p}}| = d_{\mathrm{p}}\}, \tag{3.9}$$

and $\mathscr{A}_{\mathrm{p,u}} \doteq \mathscr{A} \setminus (A_{\mathrm{p}} \cup \mathscr{A}_{\mathrm{u}})$.

The estimator η of the available surplus norm is broken accordingly, resulting in

$$\eta^2(\mathscr{A}) = \sum_{\boldsymbol{l}\in\mathscr{A}} \varepsilon^2(\boldsymbol{l}) = \eta_{\mathrm{p}}^2 + \eta_{\mathrm{u}}^2 + \eta_{\mathrm{p,u}}^2, \quad \eta_{\bullet}^2 = \sum_{\boldsymbol{l}\in\mathscr{A}_{\bullet}} \varepsilon^2(\boldsymbol{l}). \tag{3.10}$$

Note that having $\boldsymbol{l} \in \mathscr{A}_{\mathrm{p,u}}$ does not imply that it is associated to a projection surplus depending necessarily on a set of variables belonging to both $\boldsymbol{\xi}_{\mathrm{p}}$ and $\boldsymbol{\xi}_{\mathrm{u}}$, nor that its Sobol decomposition yields $(\boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F)_{\mathrm{p}} = (\boldsymbol{\Delta}_{\boldsymbol{l}}^{\mathrm{P}} F)_{\mathrm{u}} = 0$, because of external aliasing and quadrature error. Nonetheless, the indicators η_{p}, η_{u} and $\eta_{\mathrm{p,u}}$ can effectively be used to construct directional stopping criteria in the adaptive procedure. The main advantage of using the proposed indicators, compared to performing the actual Sobol decomposition of the global projection surplus in (3.8), is that the update of the η's during the adaptive process is immediate, because the indicators $\varepsilon(\boldsymbol{l})$ are easy to compute (locally) and do not change when new tensorizations are introduced in the sparse grid.

3.2.3 Adaptive strategies

The directional adaptive methods proposed below follow the usual incremental completion of the tensorizations set, adding to $\mathscr{O}$ the critical tensorization $\boldsymbol{l}^*$ with the highest indicator ε, and completing $\mathscr{A} \setminus \{\boldsymbol{l}^*\}$ with elements of the forward neighbors of $\boldsymbol{l}^*$ maintaining the $\mathscr{O}$ admissibility. The process is repeated until $\mathscr{A} = \emptyset$ (exhaustion of available tensorization) or $\eta^2 < \mathtt{Tol}$, where $\mathtt{Tol}$ is a user defined global tolerance. Here, we further restrict the set of admissible tensorizations for the completion of $\mathscr{A}$, whenever a convergence is achieved along the $\mathtt{p}$ or $\mathtt{u}$ directions. To this end, we first introduce two convergence criteria associated to the $\mathtt{p}$ and $\mathtt{u}$ dimensions,

$$C_{\mathrm{p}}(\eta_{\mathrm{p}}, \eta_{\mathrm{p,u}}) < \mathtt{Tol}_{\mathrm{p}}^2, \quad C_{\mathrm{u}}(\eta_{\mathrm{u}}, \eta_{\mathrm{p,u}}) < \mathtt{Tol}_{\mathrm{u}}^2,$$

where C_{p} and C_{u} are appropriate combinations of the directional and mixed indicators, such that $C_{\bullet}(\eta_{\bullet}, \eta_{\mathrm{p,u}}) \leq \eta^2$ for $\bullet = \mathrm{p}, \mathrm{u}$. For instance, one may consider

$$C_{\bullet}(\eta_{\bullet}, \eta_{\mathrm{p,u}}) = \eta_{\bullet}^2 + \alpha\, \eta_{\mathrm{p,u}}^2, \quad 0 \leq \alpha \leq 1. \tag{3.11}$$

Note that the tolerances $\mathtt{Tol}_{\mathrm{p}}$ and $\mathtt{Tol}_{\mathrm{u}}$, and the parameter α should generally be carefully selected, particularly when the estimates C_{p} and C_{u} have disparate magnitudes.

If none of the directional criteria is satisfied, we apply the classical method using all forward neighbors $\mathscr{F}(\boldsymbol{l}^*) = \mathscr{F}_{\mathtt{p}}(\boldsymbol{l}^*) \cup \mathscr{F}_{\mathtt{u}}(\boldsymbol{l}^*)$, with $\mathscr{F}_{\mathtt{p}}(\boldsymbol{l}^*) \doteq \{(\boldsymbol{l}^*_{\mathtt{p}} + \hat{\mathtt{e}}_i, \boldsymbol{l}^*_{\mathtt{u}}), i = 1, \cdots, d_{\mathtt{p}}\}$ the p-forward neighbors of $\boldsymbol{l}^*$ and similarly for their u−counterpart $\mathscr{F}_{\mathtt{u}}(\boldsymbol{l}^*)$. On the contrary, when the p (resp. u) directional criterion is satisfied, we only consider $\tilde{\mathscr{F}}_{\mathtt{p}} \subseteq \mathscr{F}_{\mathtt{p}}$ (resp. $\tilde{\mathscr{F}}_{\mathtt{u}} \subseteq \mathscr{F}_{\mathtt{u}}$) forward neighbors. Different reduction strategies can be conceived. In the computational examples below, we investigate the following two alternatives:

T1: Restrict the inclusion of the forward neighbors in converged directions to only tensorizations of levels lower or equal to the highest levels reached in these directions. Specifically, denoting $\bullet = \mathtt{p}$ or u, the restricted forward neighbors are defined by

$$\tilde{\mathscr{F}}_{\bullet}(\boldsymbol{l}^*) \equiv \begin{cases} \{\boldsymbol{l} \in \mathscr{F}_{\bullet}(\boldsymbol{l}^*), \boldsymbol{l}_{\bullet} \leq \tilde{\boldsymbol{l}}_{\bullet}, \} & C_{\bullet} < \mathtt{Tol}_{\bullet}, \\ \mathscr{F}_{\bullet}(\boldsymbol{l}^*) & \text{otherwise}, \end{cases}$$

where $\tilde{\boldsymbol{l}}_{\bullet}$ is defined component-wise as the largest $\bullet$-multi-index in $\mathscr{O} \cup \mathscr{A}$, i.e.

$$(\tilde{\boldsymbol{l}}_{\bullet})_i = \max_{\boldsymbol{l}_{\bullet} \in \mathscr{O} \cup \mathscr{A}} (\boldsymbol{l}_{\bullet})_i, \quad 1 \leq i \leq d_{\bullet}.$$

Likewise, the inequality between two multi-indices, $\boldsymbol{l} \leq \boldsymbol{l}'$, is understood component-wise, that is for multi-indices with length d

$$\boldsymbol{l} \leq \boldsymbol{l}' \Leftrightarrow l_i \leq l'_i \text{ for } 1 \leq i \leq d.$$

Note that in general, strategy T1 leads to the inclusion of additional forwards neighbors with either mixed $\mathtt{p}-\mathtt{u}$ dependence such that $\boldsymbol{l}^*_{\bullet} \leq \tilde{\boldsymbol{l}}_{\bullet}$, or dependence only in directions other than $\bullet$.

T2: The second strategy is more radical in that no tensorizations along converged directions are considered for the completion of $\mathscr{A}$. That is,

$$\tilde{F}_{\bullet}(\boldsymbol{l}^*) \equiv \begin{cases} \emptyset & C_{\bullet} < \mathtt{Tol}_{\bullet}, \\ F_{\bullet}(\boldsymbol{l}^*) & \text{otherwise}, \end{cases}.$$

The adaptation proceeds classically by completing $\mathscr{A} \setminus \{\boldsymbol{l}^*\}$ with the elements of $\tilde{F}_{\mathtt{p}}(\boldsymbol{l}^*) \cup \tilde{F}_{\mathtt{u}}(\boldsymbol{l}^*)$ maintaining the $\mathscr{O}$ admissibility.

The two strategies T1 and T2 are schematically illustrated in Fig. 3.1 in the case of an achieved directional convergence, where it is assumed $C_{\mathtt{p}} = \eta_{\mathtt{p}} < \mathtt{Tol}_{\mathtt{p}}$. Notice that in Fig. 3.1 (b) the tensorization $(3,2)$ is added, but not in Fig. 3.1 (e), where tensorization $(2,3)$ *is* added despite the halting line because it belongs to the u-forward neighborhood of $\boldsymbol{l}^*$. Note also that in Fig. 3.1 (a) and (d), $\boldsymbol{l}^* = \{3,1\}$ for which no forward neighbor is added because of $\mathscr{O}$-admissibility ($(3,2)$ missing) and because the p direction has converged ($\tilde{\boldsymbol{l}}_{\mathtt{p}} = (3)$). However, when $\boldsymbol{l}^*$ *is* removed from $\mathscr{A}$ we have $\mathscr{A}_{\mathtt{p}} = \emptyset$, making $\eta_{\mathtt{p}} = C_{\mathtt{p}} = 0$. This is not a concern since this only occurs after convergence in that dimension.

Adaptation continues until either *both* $\eta_{\mathtt{p}} \leq \mathtt{Tol}_{\mathtt{p}}$ and $\eta_{\mathtt{u}} \leq \mathtt{Tol}_{\mathtt{u}}$ or, until the global $\eta \leq \mathtt{Tol}$. Unlike the nested adaptation, the product-space adaption has both independent convergence metrics and a global convergence metric. Also, whereas extension of the nested adaptation approach to greater than two spaces is not recommended, the product-space construction trivially lends itself to any number of spaces or groupings of dimensions.

Note that the present aPSP scheme with directionally tuned refinement naturally exploits sparse-tensorization in the product space, and is consequently expected to be more efficient than the nested algorithm. This is achieved at the potential expense of a coarser control of the resolution properties along the p and u directions. The effect of this compromise is examined in the analysis below, together with a comparative analysis of the performance of the two algorithms.

4 Test Problem

4.1 Test Problem Definition

In order to examine the behavior of the methods presented in Section 3, we made use of the two-dimensional $(d = 2)$ test function,

$$F(\boldsymbol{\xi}) = F(\xi_1, \xi_2) = \left(1 + \frac{1/3}{2\xi_1 + \xi_2 + 7/2}\right) \exp\left[-\left(\frac{1}{2}\left(\xi_2 - \frac{1}{5}\right)(\xi_2 + 1)\right)^2\right], \quad \boldsymbol{\xi} \in [-1,1]^2, \tag{4.1}$$

allowing us to rapidly evaluate performance and to conduct error analyses. The function F (Fig. 4.1) was designed so as to exhibit a mixed dependence in ξ_1 and ξ_2 and a greater dependence on ξ_1 as measured by the Sobol sensitivity indices of F which are $S_{\{1\}} = 0.531$, $S_{\{2\}} = 0.199$, and $S_{\{1,2\}} = 0.270$ (see Appendix A). Because the nested method adapts differently the inner and outer directions, we shall call (1,2)-Nested method the case where the inner-loop acts on ξ_2 (that is $\mathtt{u} = \{2\}$) and the outer-loop acts on ξ_1 ($\mathtt{p} = \{1\}$), whereas for the (2,1)-Nested method the roles of these directions

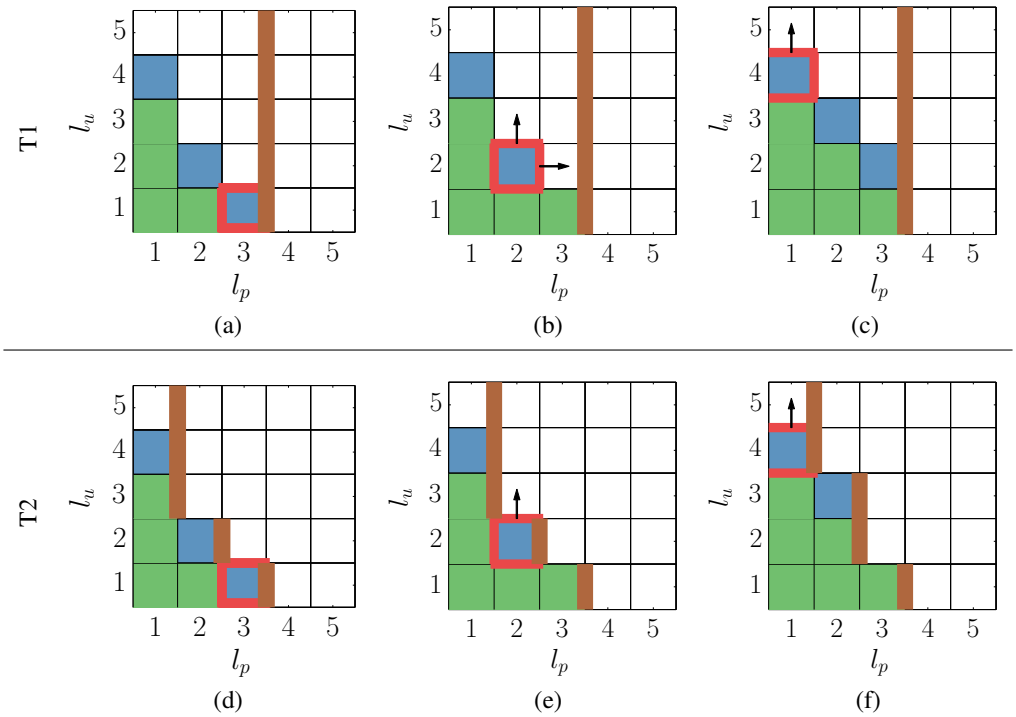

Fig. 3.1 Schematic demonstration of strategy T1 (top plots, a-c) and T2 (bottom plots, d-f) for the case $\eta_{\mathrm{p}}^2 < \mathtt{Tol}_{\mathrm{p}}$. Shown in blue is $\mathscr{A}$, and in green is $\mathscr{O}$. The red box denotes $\boldsymbol{l}^*$ and the thick-brown line represents the halting-condition imposed by the two strategies.

are exchanged. To illustrate the quantities of interest driving the adaptation in the (1,2)-Nested method, Fig. 4.1 also depicts for a set of values $\xi_1^{(i)}$ the cross-sections $F(\xi_1^{(i)},\xi_2) = G^{(i)}(\xi_2)$ used for the inner-loop adaptation, as well as the outer-loop quantity of interest $H(\xi_1) = \|F(\xi_1,\cdot)\|_{L_2([-1,1])}$.

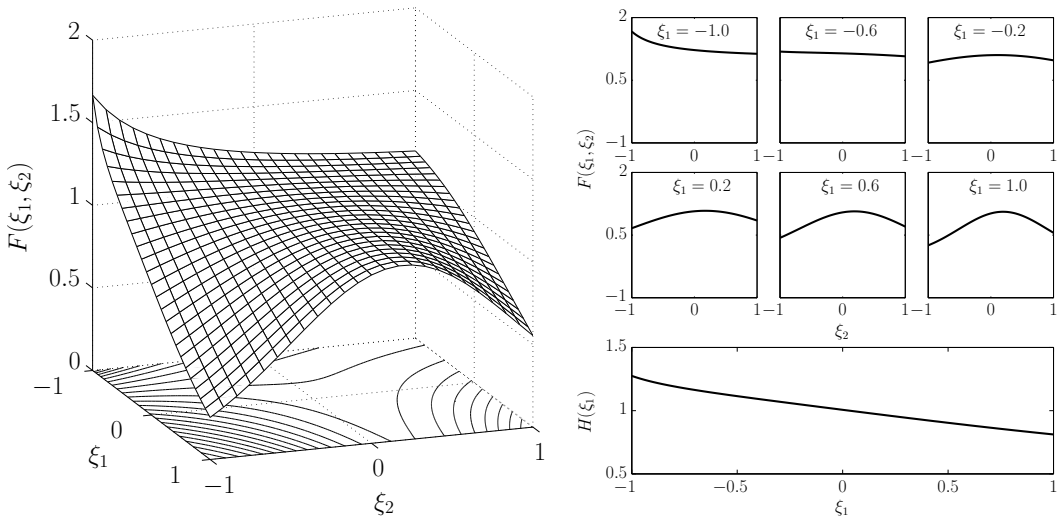

(a) Function F (b) Functions for the (1,2)-Nested methods.

Fig. 4.1 The plot in (a) shows the 2D surface $F(\xi_1,\xi_2)$. The plots in (b) depict the functions used in the (1,2)-Nested method: cross-sections $G^{(i)}(\xi_2)$ for selected values of $\xi_1^{(i)}$, as indicated (four top plots), and the function $H(\xi_1)$ for outer-loop adaptation (bottom).

4.2 Error Analysis

Let us denote here $\widetilde{F}(\boldsymbol{\xi})$ the PCE surrogate of F obtained with one of the adaptive projection methods, and $Z(\boldsymbol{\xi})$ the error function:

$$Z(\boldsymbol{\xi}) \doteq F(\boldsymbol{\xi}) - \widetilde{F}(\boldsymbol{\xi}) = F(\boldsymbol{\xi}) - \sum_{\boldsymbol{k}\in\mathscr{K}} \widetilde{f}_{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}). \tag{4.2}$$

Introducing the exact (but unknown) PC expansion of F, Z can be split into two orthogonal contributions:

$$F(\boldsymbol{\xi}) = \sum_{\boldsymbol{k}\in\mathbb{N}_0^d} f_{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}) \quad \Rightarrow \quad Z(\boldsymbol{\xi}) = \underbrace{\sum_{\boldsymbol{k}\in\mathscr{K}} \left(f_{\boldsymbol{k}} - \widetilde{f}_{\boldsymbol{k}}\right) \Psi_{\boldsymbol{k}}(\boldsymbol{\xi})}_{\text{Internal aliasing}} + \underbrace{\sum_{\boldsymbol{k}\in\mathbb{N}_0^d\setminus\mathscr{K}} f_k \Psi_k(\boldsymbol{\xi})}_{\text{Projection error}}.$$

Another decomposition, more relevant to the task of error control, consists of orthogonal Sobol decomposition (see Appendix A) of Z according to:

$$Z(\boldsymbol{\xi}) = Z_{\emptyset} + Z_{\{1\}}(\xi_1) + Z_{\{2\}}(\xi_2) + Z_{\{1,2\}}(\xi_1,\xi_2).$$

The *relative* error norm, ζ, can then be decomposed into

$$\zeta^2 \doteq \frac{\|Z\|_{L_2}^2}{\|F\|_{L_2}^2} = \zeta_{\emptyset}^2 + \zeta_{\{1\}}^2 + \zeta_{\{2\}}^2 + \zeta_{\{1,2\}}^2, \tag{4.3}$$

where $\zeta_\bullet = \|Z_\bullet\|_{L_2}/\|F\|_{L_2}$. Thus, $\zeta_\emptyset$ measures the relative error on the mean, $\zeta_{\{1\}}$ (resp. $\zeta_{\{2\}}$) measure the relative error in direction ξ_1 (resp. ξ_2) only, and $\zeta_{\{1,2\}}$ measures the relative error in mixed directions. Note that in higher dimension, this decomposition of the error norm can be extended to arbitrary partitions of the dimensions.

Following the discussion in Appendix A.2, the relative errors can be easily computed from the PCE of Z and F. Since for the present example an explicit PCE for F in (4.1) is not available, we instead construct a highly accurate projection of F on a 90^{th} total-order polynomial basis using full-tensored Gauss-Legendre rule having a total of 100^2 quadrature points.

4.3 Behavior of the adaptive methods

In this section, we analyze and contrast the performance of the different adaptive schemes developed above. To this end, computations are performed for the test problem, using the tolerances reported in Table 4.1. Recall that the nested projection uses one tolerance for each direction, $\texttt{Tol}_{\{1\}}$ and $\texttt{Tol}_{\{2\}}$, that serve as stopping criteria for the inner or outer loops depending if the (1,2) or (2,1)-Nested method is considered. The aPSP uses a global tolerance $\texttt{Tol}$ and the aPSP-T1 and aPSP-T2 methods (see Section 3.2.3) have in addition two directional tolerances $\texttt{Tol}_1$ and $\texttt{Tol}_2$. When $\texttt{Tol}_1 = \texttt{Tol}_2 = 0$, one recovers the original aPSP method and is then used as reference in the presented results.

	Nested	aPSP (reference)	aPSP-T1, 2
$\texttt{Tol}$	–	10^{-5}	10^{-5}
$\texttt{Tol}_1$	10^{-7}	0	10^{-7}
$\texttt{Tol}_2$	10^{-9}	0	10^{-9}

Table 4.1 Tolerance values for different adaptive methods used for the test problem.

4.3.1 Adaptive Product Space Projection with directional criteria

Due to the low-dimensionality of the test problem, the product-space adaptation was performed for the tolerance values in Table 4.1 with $\alpha = 0$ in (3.11). Figure 4.2 depicts the squared projection surplus norms, η^2, and their Sobol decompositions into pure and mixed contributions, during the adaptive process for the aPSP and aPSP-T1 methods. Results for aPSP-T2 are not shown as they lead to results similar to those for aPSP-T1. The evolutions of the adaptation criteria are reported, as a function of the number of evaluations of F required by the two methods. The plots show that the two methods follow a similar path but that aPSP-T1 terminates earlier as the dimensional tolerances are met. In addition, it is seen that $\eta_{\{1\}}$ in the aPSP-T1 method becomes zero after few adaptive iterations, for $\mathscr{A}_{\{1\}} = \emptyset$, However, this occurs after $\eta_{\{1\}}$ drops below the specified tolerance $\texttt{Tol}_1$, indicated by a dashed line in Fig. 4.2(b), as we would have expected. Note that for the two methods, $\eta_{\{1,2\}}^2 \gg \eta_{\{1\}}, \eta_{\{2\}}$ indicating that much of the remaining surplus is due to mixed terms.

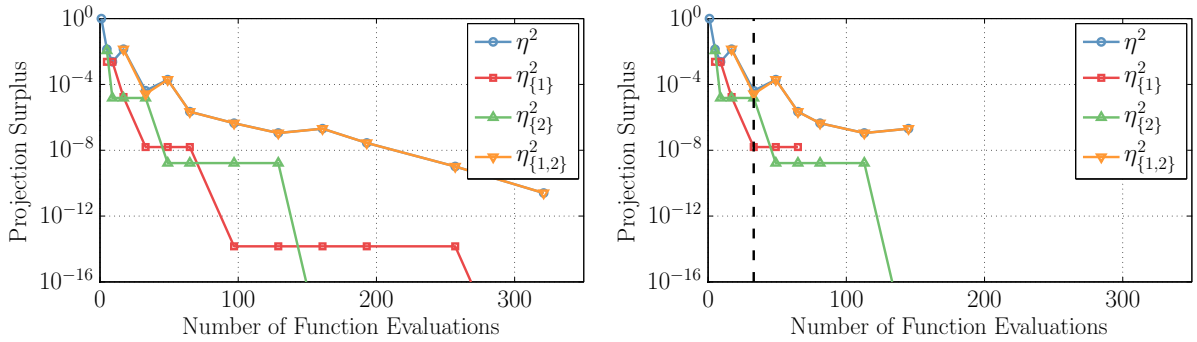

(a) aPSP (b) aPSP-T1

Fig. 4.2 Relative projection surplus η^2 and their Sobol decompositions versus the number of function evaluations for the test function in Equation (4.1). Plotted in (a) are the results for aPSP (reference). Plot (b) shows the results of the aPSP-T1. The dashed line in (b) represents the point at which the aPSP-T1 converge in direction $\{1\}$.

The results thus confirm that method aPSP-T1 is able to control the adaptivity for separate sets of directions, and to control termination of the refinement along the corresponding dimensions.

To further illustrate the behavior of the directionally-tuned adaption, we plot in Fig. 4.3 the pseudo-spectral projection multi-index sets, $\mathscr{L}$, the selection indicators $\varepsilon(\boldsymbol{k})\ \forall\ \boldsymbol{k} \in \mathscr{L}$, and the adaptive paths for (a) the aPSP, (b) the aPSP-T1 and (c) the aPSP-T2 methods. The number inside each box is the iteration at which each index is added. The results indicate that, as expected for the present settings, the reference aPSP method does include more mixed multi-indices ($\boldsymbol{l} = (l_1, l_2)$ with $l_1, l_2 > 1$) compared to the other two methods, corresponding to a richer projection space. Indeed, for aPSP-T1 and aPSP-T2 the adaptation halts refining sooner than aPSP along direction $\{1\}$, which is associated to a larger $\texttt{Tol}_1$ value, while still refining along direction $\{2\}$ as for aPSP. The plots also highlight the effect of selecting strategy T1 or T2, the latter being more aggressive in the sense that it results into a coarser multi-index set $\mathscr{L}$ with fewer mixed tensorizations than for aPSP-T2.

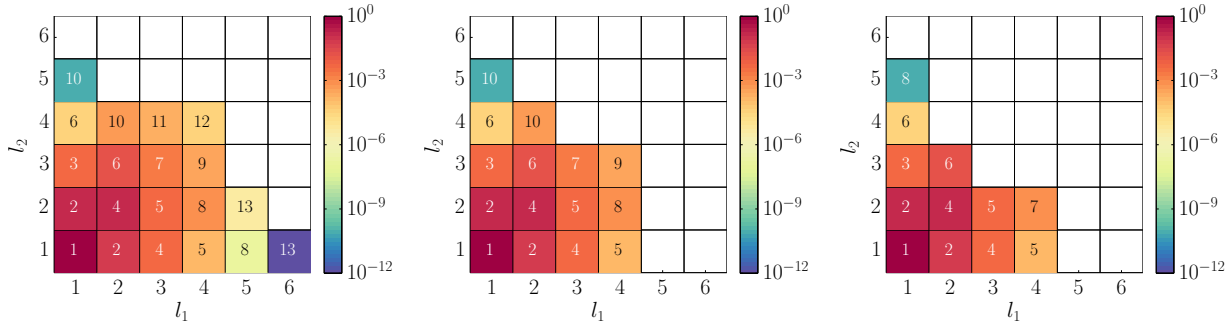

(a) Reference (b) Method T1 (c) Method T2

Fig. 4.3 Multi-index sets for the pseudo-spectral projection using (a) aPSP, (b) aPSP-T1 and (c) aPSP-T2 methods. Also shown are the color-coded values of the norm of the projection surpluses $\varepsilon(\boldsymbol{l})$ while the iteration number at which that multi-index was added to $\mathscr{L}$.

4.3.2 Nested Projection

We now focus on the behavior of the nested adaptation with the tolerance values reported in Table 4.1. The adaptation was performed for both the (1,2) and (2,1)-Nested algorithms. Figure 4.4 presents the final grids of points where the function has been evaluated, for the reference aPSP method and the two nested projection alternatives. The results indicate that, in contrast to aPSP, the nested adaptations exhibit a *nearly* fully-tensorized grids of points between the two spaces and a larger number of function evaluations as a consequence. The plots also illustrate how at different outer-direction grid points, the inner-space grids are not necessarily identical.

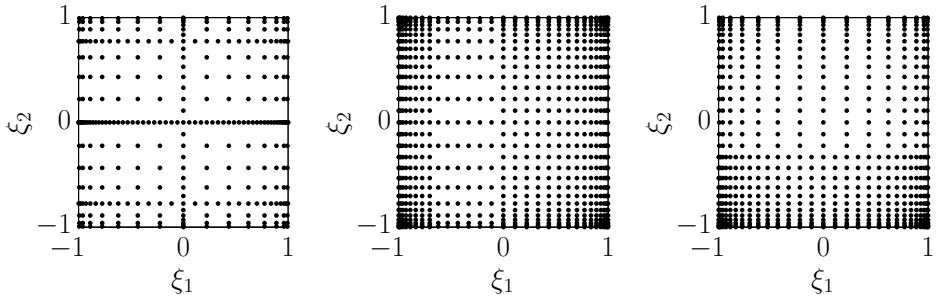

(a) Reference (b) (1,2)-Nested grid (c) (2,1)-Nested grid

Fig. 4.4 Comparison of the final grids $\mathscr{G}(\mathscr{L})$ for (a) aPSP method, (b) the (1,2)-Nested adaptation, and (c) (2,1)-Nested adaptation.

4.4 Error and Performance Analysis

Figure 4.5 shows the global error estimates, ζ^2, plotted against the number of function evaluations, for: (a) the (1,2)-Nested method, (b) the (2,1)-Nested Method, (c) the aPSP method, and (d) the aPSP-T1 method. Results for aPSP-T2 are similar to those of aPSP-T1, and are consequently omitted. Plotted along with ζ^2 is its decomposition based on (4.3).

The results in Fig. 4.5(a) and 4.5(b) illustrate an intrinsic feature of the nested construction, namely that the error associated with inner dimension is reduced immediately. This can be appreciated by the fact that $\zeta^2_{\{2\}}$ in Fig. 4.5(a) and the $\zeta^2_{\{1\}}$ in Fig. 4.5(b) decrease rapidly. Meanwhile, the global error closely follows the error in the outer dimension. Note that for the present example, the mixed-term error $\zeta_{\{1,2\}}$ has a dominant contribution to the global error in both the nested, aPSP and aPSP-T1 methods. This observation, however, need not hold in general, particularly in situations when F has essentially an additive form: $F(\boldsymbol{\xi}_\mathrm{p},\boldsymbol{\xi}_\mathrm{u}) \approx F_\mathrm{p}(\boldsymbol{\xi}_\mathrm{p}) + F_\mathrm{u}(\boldsymbol{\xi}_\mathrm{u})$.

The performance of the aPSP-T1 can be assessed by contrasting the results in Figs. 4.5(c) and 4.5(d). Specifically, the results indicate that, consistent with the behavior of the η^2 indicators in Fig. 4.2, the aPSP terminates with a lower global error ζ^2; however, this comes at the cost of more function evaluations. The results also illustrate the effectiveness of aPSP-T1 in guiding and halting the refinement in direction $\{1\}$, with a small impact on the global error. This feature will prove especially useful in a higher dimensional applications, where tighter control on directional refinement can be substantially advantageous.

Finally, contrasting the results in Figs. 4.5(a) and (b) with those in Figs. 4.5(c) and (d), the following general conclusion can be drawn from the test example. First, all adaptive methods exhibit a similar complexity, in the sense that the global errors have similar decay rates with the number of function evaluations. Second, the nested method offers an efficient means to reduce at a minimal cost the error along the inner direction. Third, the aPSP-T1,2 methods offer flexibility to balance and halt the adaptation along specific directions.

Finally, we examine the impact of setting to zero the α_i^k coefficients in Equation (3.6) in the nested construction at outer grid points where these coefficients are not defined. To this end, we performed a computation where the missing coefficients were actually computed by means of an inner projection based on the union of the inner grids at all outer points. In other words, the nested grid was completed to a full-tensorization grid, requiring additional evaluations of the function. For the present example and other numerical tests (not shown) it was found that the resulting improvement in the global error was marginal while the completion of the grid could require a significant number of functions evaluations, with an overall degradation of the efficiency. This point is illustrated for the present example in Fig. 4.6.

5 High-Dimensional Problem

This section illustrates the behavior of the methods outlined above for a high-dimensional combustion chemistry problem that has motivated their development. The problem consists in the optimal design of shock-tube ignition experiments, and thus makes use of global surrogates that describe the dependence of key experimental observables (i.e. QoIs) on both uncertain rate parameters (to be calibrated by yet unperformed experiments) and design variables (to be optimally selected based on prior information).

The shock tube experiment presently considered focuses on methane ignition. The reaction model is based on the GRI-Mech 3.0 [37] mechanism, involving 57 neutral species and 325 neutral species reactions. The model is augmented with ion-chemistry mechanism developed by Prager *et. al* [32], including 11 charged species and 67 ionic reactions.

Since we are primarily interested in characterizing the performance of adaptive constructions, we focus on a single experimental observable, namely the peak electron concentration that is achieved during the ignition event. The uncertain parameters concern the rate constants of 22 reactions that were deemed especially relevant based on a reaction pathway analysis. Details of this analysis are beyond the scope of the present work, and will be described elsewhere [21]. The ranges of the uncertain reaction rates are specified in terms of factors, as summarized in Table 5.1(a). Based on these

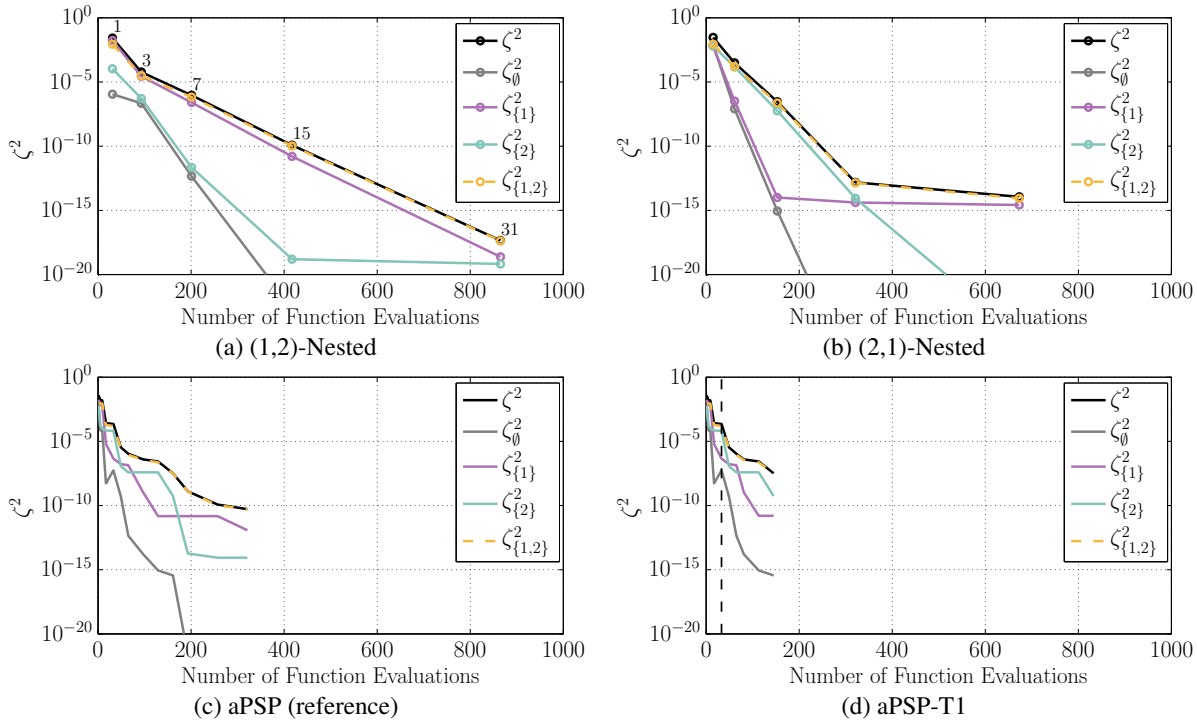

Fig. 4.5 ζ^2 and its Sobol decomposition versus number of function evaluations, for (a) the (1,2)-Nested, (b) the (2,1)-Nested, (c) the aPSP, and (d) aPSP-T1 methods.

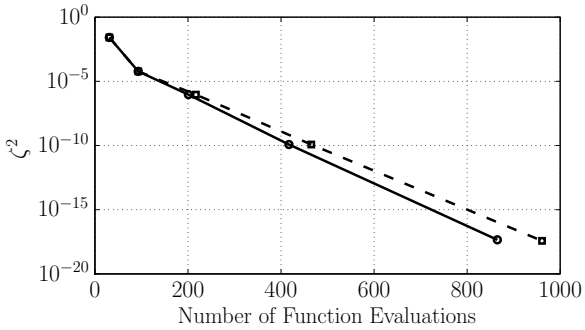

Fig. 4.6 Comparison of the global error ζ^2 versus number of function evaluations, for the nested method (solid line) and its completion to full-tenzorization (dotted line).

factors, the uncertain rates are parametrized using independent random variables that are uniformly distributed in the interval $[-1,1]$. This leads to a 22-dimensional germ $\boldsymbol{\xi}_{\mathrm{u}}$. We assume that the initial pressure is atmospheric, and focus on three-design variables, namely the initial temperature, and the initial concentrations of CH_4 and O_2. Thus, the design-variable vector, $\boldsymbol{\xi}_{\mathrm{p}}$, is three-dimensional, and the ranges of its components are specified in Table 5.1(b). We associate with $\boldsymbol{\xi}_{\mathrm{p}}$ a constant weight ρ_{p}. Once design variables are selected, the evolution of the system for a specific realization of the germ $\boldsymbol{\xi}_{\mathrm{u}}$ is obtained using an extended version of the TChem software [35,34]. The problem is then set in terms of the germ $\boldsymbol{\xi} = (\boldsymbol{\xi}_{\mathrm{u}}, \boldsymbol{\xi}_{\mathrm{p}}) \in [-1,1]^{22+3}$ with density $\rho_{\boldsymbol{\xi}} = 2^{-25}$.

We applied the nested, aPSP and aPSP-T1 methods using the the tolerance values listed in Table 5.2. These values were set so as to obtain a suitably accurate approximation of the selected QoI. Also note that in the present case, d_u is significantly larger than d_p; consequently, we used lower tolerance values for the u directions than for the p directions.

We initially ran the aPSP-reference and found that the QoI exhibits a large sensitivity to the design variables $\boldsymbol{\xi}_{\mathrm{p}}$. This can be appreciated from Fig. 5.1, which provides the first-order and total sensitivity indices, estimated from the PCE after the algorithm has converged. The results clearly illustrate the dominant contribution of the design variables to the overall variance of the peak electron concentration. The results also indicate that for the present case there appears to be only a small contribution in the mixed term, specifically accounting for about 4% of the total variance; not surprisingly, the computed first-order and total sensitivity indices for the individual parameters have close values.

We examined (not shown) the behavior of the aPSP and aPSP-T1 methods during the adaptation. The aPSP-T1 method followed a similar path to the aPSP, converging early in the p direction while the p directions are subsequently

ID	Reaction	Uncertain Range
326	$CH + O \Leftrightarrow HCO^+ + E^-$	$\pm 50\%$
328	$HCO^+ + H_2O \Leftrightarrow H_3O^+ + CO$	$\pm 50\%$
330	$H_3O^+ + E^- \Leftrightarrow H_2O + H$	$\pm 25\%$
331	$H_3O^+ + E^- \Leftrightarrow OH + H + H$	$\pm 25\%$
332	$H_3O^+ + E^- \Leftrightarrow H_2 + OH$	$\pm 25\%$
333	$H_3O^+ + E^- \Leftrightarrow O + H_2 + H$	$\pm 25\%$
336	$HCO^+ + CH_3 \Leftrightarrow C_2H_3O^+ + H$	$2^{\pm}$
338	$H_3O^+ + CH_2CO \Leftrightarrow C_2H_3O^+ + H_2O$	$\pm 25\%$
340	$C_2H_3O^+ + O \Leftrightarrow HCO^+ + CH_2O$	$\pm 50\%$
349	$OH^+ + O \Leftrightarrow HO_2 + E^-$	$\pm 50\%$
350	$OH^- + H \Leftrightarrow H_2O + E^-$	$2^{\pm}$
353	$OH^- + CH_3 \Leftrightarrow CH_3OH + E^-$	$2^{\pm}$
361	$O^- + H_2O \Leftrightarrow OH^- + OH$	$2^{\pm}$
366	$O^- + H_2 \Leftrightarrow H_2O + E^-$	$2^{\pm}$
369	$O^- + CO \Leftrightarrow CO_2 + E^-$	$2^{\pm}$
006	$O + CH \Leftrightarrow H + CO$	$\pm 25\%$
049	$H + CH \Leftrightarrow C + H_2$	$1.5^{\pm}$
091	$OH + CH \Leftrightarrow H + HCO$	$\pm 25\%$
093	$OH + CH_2 \Leftrightarrow CH + H_2O$	$\pm 25\%$
125	$CH + O_2 \Leftrightarrow O + HCO$	$\pm 50\%$
126	$CH + H_2 \Leftrightarrow H + CH_2$	$\pm 50\%$
127	$CH + H_2O \Leftrightarrow H + CH_2O$	$3^{\pm}$

(a) Uncertain Parameters ($\boldsymbol{\xi}_u$-space)

Design Parameter	Interval
Temperature (K)	$[2250, 2750]$
CH_4 Mole Fraction (%)	$[0.5, 2.0]$
O_2 Mole Fraction (%)	$[0.2, 0.75]$

(b) Design Variables ($\boldsymbol{\xi}_p$-space)

Table 5.1 (a) Uncertainty in rate parameters, and (b) ranges of design variables, for the CH_4+O_2 shock-tube ignition problem. The uncertainty in the rate parameters are specified either in terms of a percentage are relative, or as a multiplicative factor. In all cases, the distributions are represented in terms of canonical random variables uniformly distributed over $[-1, 1]$.

	Nested	aPSP (reference)	aPSP-T1
`Tol`	–	2.75×10^{-3}	2.75×10^{-3}
$\texttt{Tol}_\texttt{p}$	3.50×10^{-3}	0	3.50×10^{-3}
$\texttt{Tol}_\texttt{u}$	2.50×10^{-3}	0	2.50×10^{-3}

Table 5.2 Tolerance values for the Nested-(p,u), aPSP (reference), and aPSP-T1 methods with $\alpha = 0.1$ in Equation (3.11).

refined. We also examined whether the similar adaptation paths were affected by the selected tolerance values, namely by rerunning aPSP-T1 with larger values of $\texttt{Tol}_\texttt{p}$, thus forcing the refinement along the p directions to terminate earlier. However, this had little impact on the grid refinement path, even after the specified $\texttt{Tol}_\texttt{p}$ value was reached. This is not surprising, in light of the fact that the contribution of the mixed terms is small.

The projection surplus, η^2, and its Sobol decomposition are plotted in Fig. 5.2. For both the aPSP and the aPSP-T1 adaptations, η_p^2 immediately and rapidly decays as the sparse grid is adapted, whereas the mixed direction surplus is clearly dominant. The results also indicate that behavior of the projection surplus and its Sobol components are very similar for both algorithms. This is consistent with our earlier observation that the adaptation of the sparse grids evolve similar for both algorithms, though aPSP-T1 evidently terminates earlier.

Similarly to what we observed with the two-dimensional test function in Section 4, the dimensionally guided adaptation affected the adaptive refinement path only slightly. Both aPSP and aPSP-T1 exhibit similar decay of the error indicators with the number of model realizations, and as further discussed below similar behavior for errors estimates obtained from independent samples. However, with the selected tolerances, aPSP-T1 terminated with 3411 model evaluations, whereas aPSP required 5651 realizations to reach termination. This illustrates the flexibility of directional criteria in tuning adaptation (and as further illustrated below resolution), along specific dimensions.

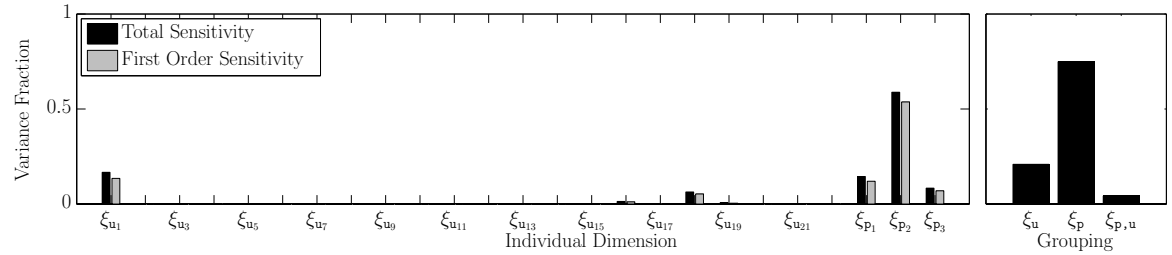

Fig. 5.1 Sensitivity of the peak electron concentration in the shock-tube ignition problem. Plotted on the left are the first-order and and total sensitivity indices for individual parameters. The right plot depicts the sensitivity indices for the $\boldsymbol{\xi}_{\mathtt{u}}$ and $\boldsymbol{\xi}_{\mathtt{p}}$ vectors, as well as the contribution of the mixed term.

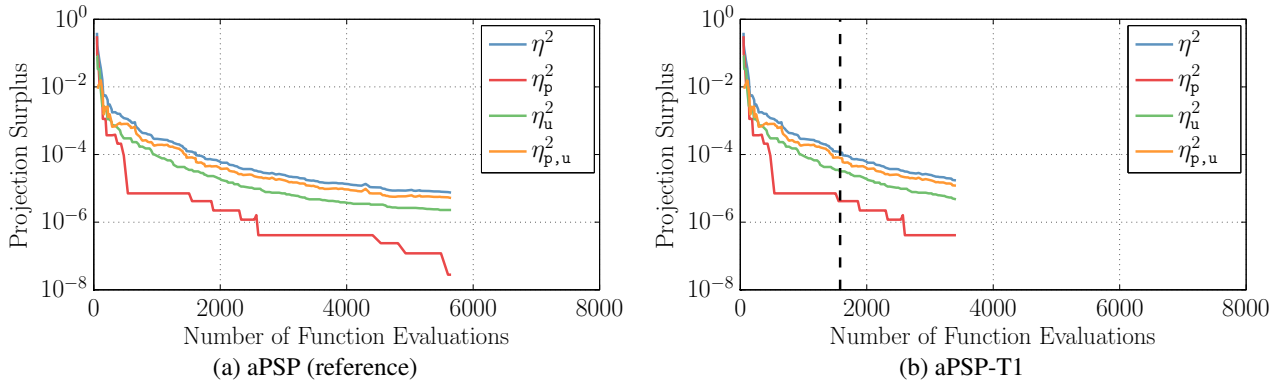

Fig. 5.2 Projection surplus, η, and its Sobol decomposition for (a) aPSP, and (b) aPSP-T1 methods. Also plotted in (b) is a dashed line representing the point at which the adaptation converged in the $\mathtt{p}$ directions.

We performed the nested adaptation with the 22 uncertain parameters ($\boldsymbol{\xi}_{\mathtt{u}}$) as the inner dimensions and the 3 design paramaters ($\boldsymbol{\xi}_{\mathtt{p}}$) as the outer dimensions. The adaptation was allowed to proceed until both $\mathtt{Tol}_{\mathtt{u}}$ and $\mathtt{Tol}_{\mathtt{p}}$ values were met (recall that the nested construction does not involve a global tolerance estimate). As discussed earlier, different grid locations in the $\xi_{\mathtt{p}}$ space admit different adaptations into $\boldsymbol{\xi}_{\mathtt{u}}$. The dependence of the number of inner grid points at each point in $\boldsymbol{\xi}_{\mathtt{p}}$ domain is visualized in Fig. 5.3. Notice that, towards some boundaries of $\boldsymbol{\xi}_{\mathtt{p}}$ domain, inner grids are generated with many more points than in other areas. Despite the flexibility afforded in locally adapting at each outer point, at termination the nested method required a total 76,411 model evaluations. This is over one order of magnitude larger than what was required by the aPSP methods to achieve a comparable error.

To validate the results we generated a Monte-Carlo ensemble consisting of 4096 realizations. An additional 8192 simulations were performed in order to obtain estimates of the Sobol decomposition of the relative error, ζ^2, between the Monte-Carlo sample and the sparse adaptive algorithm. The ζ^2 values, and their Sobol decompositions thus estimated, are presented in Fig. 5.4.

The results of Fig. 5.4(a) demonstrate how the MC estimates of the error in the inner directions are quickly reduced, whereas the global error is predominantly dictated by the refinement in outer dimensions. This is in stark contrast to the behavior seen for the aPSP methods, where the first-order contributions $\zeta_{\mathtt{p}}$ and $\zeta_{\mathtt{u}}$ *both* drop rapidly at early stages of the adaptation while the decay of the global error follows that the mixed term. The results also indicate that the error control afforded by the nested method comes at the expense of a large computational cost, namely requiring more than an order-of-magnitude increase in the number of model realizations. Figure 5.4(c) demonstrates the effectiveness of the aPSP-T1 method in terminating adaptation along $\mathtt{u}$ directions. The results also illustrate the flexibility of specifying directional criteria, which for the presently selected values result in termination of aPSP-T1 with significantly smaller number of model evaluations than aPSP. Finally, we note that for all three methods, the error estimates obtained based on the MC sample drop to very small values as the iterations converge to their specified tolerances. This provides confidence in the validity of all three methods and in the suitability of the resulting surrogates.

6 Conclusion

Two approaches were developed for obtaining a global surrogate of QoIs involving parameters belonging to different spaces. The first approach relied on a nested construction where the two spaces were adapted separately. Specifically, at each realization of the outer variables, an adaptive Smolyak pseudo-spectral projection was performed in the inner space. The local PC coefficients at the outer grid points were then projected in the outer space.

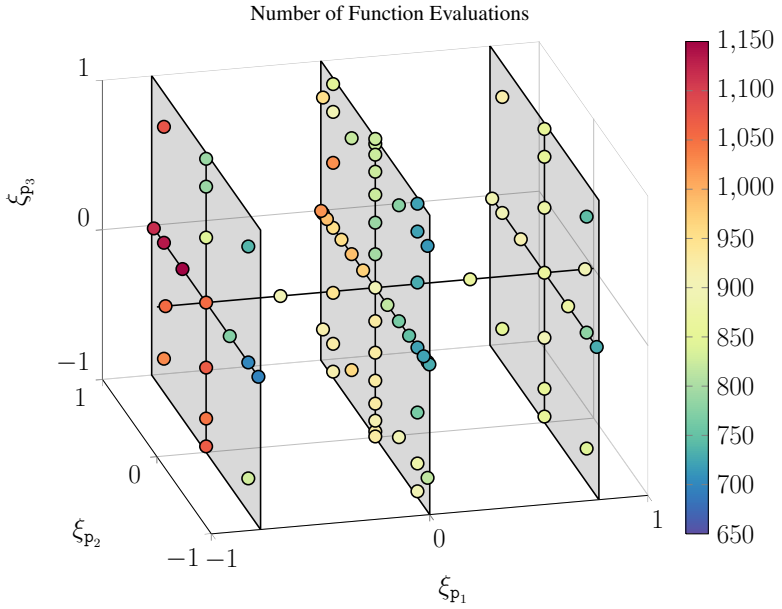

Fig. 5.3 Number of model evaluations required at each realization of the design variables ($\boldsymbol{\xi}_{\mathtt{p}}$), needed to calculate the local $\boldsymbol{\xi}_{\mathtt{u}}$ projection. Shaded $\mathtt{p}_2 - \mathtt{p}_3$ planes are for reference.

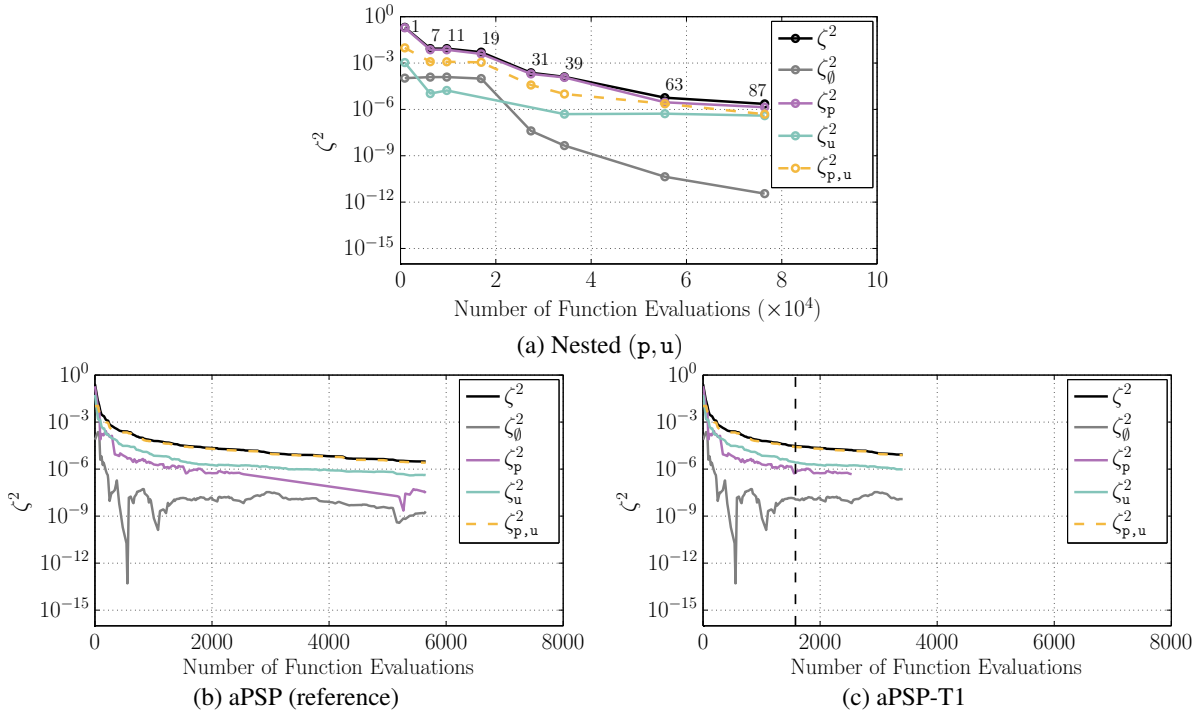

Fig. 5.4 ζ^2 values for the (a) Nested $(\mathtt{p},\mathtt{u})$, (b) aPSP, and (c) aPSP-T1 methods. In (a), the legend next to individual points refer to the number of outer ($\boldsymbol{\xi}_{\mathtt{p}}$) realizations. The dotted vertical line in (c) represents the point at which $\mathtt{Tol}_{\mathtt{p}}$ is reached. Note that the scale of the x–axis differs in (a) from (b,c).

The second approach was based on considering a suitable product space combining all parameters, and on performing an adaptive pseudo-spectral projection in this product space. In addition, we developed a decomposition of the convergence indicator η, which enabled us to tune adaptivity along individual dimensions. Two versions of the adaptive algorithm were developed for this purpose, allowing to terminate refinement along specific directions when specified directional tolerances are met.

The validity of the construction was analyzed, including a Sobol-decomposition of the error between the constructed surrogate-based estimates and the true model. We analyzed both approaches in light of applications on a simple test

problem, as well as detailed simulations of a shock-tube ignition problem involving 3 design variables and 22 uncertain rate parameters.

Computational tests indicated that the nested approach affords great flexibility in controlling the error in individual sets of parameters, first reducing the error of the inner expansion locally at each point of the outer variables sparse grid, and then reducing the error of the outer expansion and consequently the global representation. However, this flexibility, which is gained from the tensorized construction of the nested approach, comes at the expense of larger computational burden. In contrast, the adaptive pseudo-spectral projection refinement in the product-space approach enabled the construction of suitable surrogates at a smaller cost compared to the nested algorithm. The tests also reveal that directional tuning algorithm can provide an effective means of resolution control, affording appreciable savings over unconstrained refinement.

Both the nested and product-space approaches offer many avenues for further improvement. Specifically, the nested approach is not limited to performing pseudo-spectral projections in both the inner and outer space. The incorporation of other methods for constructing surrogates, such as collocation or regression, would potentially constitute avenues for improvement. In particular, the use of regularized solvers such as compressed sensing [7,8] appear to be promising, namely for reducing the required number of model evaluations. The product-space approach with the decomposed error indicators also offers promising areas for further extensions. For instance, the tolerances used for tuning adaptivity could be dynamically adjusted, to guide refinement towards specific hyperplanes of the parameter space. Another possibility would be to incorporate cost functions into the sparse grid refinement algorithm, so as to optimally exploit computational resources. Such extensions will be considered in future work.

A Global Sensitivity Analysis

A.1 Sensitivity indices

Following notation of Section 2, consider $F : \boldsymbol{\xi} \in \Xi \mapsto F(\Xi) \in L_2(\Xi, \rho)$, where $\boldsymbol{\xi} = (\xi_1 \cdots \xi_d)$ is a vector of d independent real-valued random variable with joined density ρ. Let $\mathscr{D} = \{1, \cdots, d\}$ and for $\mathfrak{u} \in \mathscr{D}$ denote $|\mathfrak{u}| = \mathrm{Card}(\mathfrak{u})$ and $\mathfrak{u}_\sim = \mathscr{D} \setminus \mathfrak{u}$, such that $\mathfrak{u} \cup \mathfrak{u}_\sim = \mathscr{D}$, $\mathfrak{u} \cap \mathfrak{u}_\sim = \emptyset$. Given $\mathfrak{u} \in \mathscr{D}$ we denote $\boldsymbol{\xi}_\mathfrak{u}$ the sub-vector of $\boldsymbol{\xi}$ with components $(\xi_{\mathfrak{u}_1} \cdots \xi_{\mathfrak{u}_{|\mathfrak{u}|}})$, so $\boldsymbol{\xi} = (\boldsymbol{\xi}_\mathfrak{u} \boldsymbol{\xi}_{\mathfrak{u}_\sim})$. Under the stated assumptions, the function $F(\boldsymbol{\xi})$ has a *unique* orthogonal decomposition of the form [15]

$$F(\boldsymbol{\xi}) = \sum_{\mathfrak{u} \in \mathscr{D}} f_\mathfrak{u}(\boldsymbol{\xi}_\mathfrak{u}), \quad \langle f_\mathfrak{u}, f_\mathfrak{v} \rangle = 0 \text{ if } \mathfrak{u} \neq \mathfrak{v}. \tag{A.1}$$

The functions $f_\mathfrak{u}$ can be recursively expressed as [39]

$$f_\mathfrak{u}(\boldsymbol{\xi}_\mathfrak{u}) = \mathbb{E}\{F \mid \boldsymbol{\xi}_\mathfrak{u}\} - \sum_{\substack{\mathfrak{v} \in \mathscr{D} \\ \mathfrak{v} \subsetneq \mathfrak{u}}} f_\mathfrak{v}(\boldsymbol{\xi}_\mathfrak{v}),$$

where $\mathbb{E}\{F \mid \boldsymbol{\xi}_\mathfrak{u}\}$ is the conditional expectation of F given $\boldsymbol{\xi}_\mathfrak{u}$, namely

$$\mathbb{E}\{F \mid \boldsymbol{\xi}_\mathfrak{u}\} = \int F(\boldsymbol{\xi}_\mathfrak{u} \boldsymbol{\xi}_{\mathfrak{u}_\sim}) \rho(\boldsymbol{\xi}_{\mathfrak{u}_\sim}) d\boldsymbol{\xi}_\sim.$$

The decomposition (A.1) being orthogonal, the variance of F, $\mathbb{V}\{F\}$, is decomposed into

$$\mathbb{V}\{F\} = \sum_{\substack{\mathfrak{u} \in \mathscr{D} \\ \mathfrak{u} \neq \emptyset}} V_\mathfrak{u}, \quad V_\mathfrak{u} = \mathbb{V}\{f_\mathfrak{u}\}. \tag{A.2}$$

The partial variance $V_\mathfrak{u}$ measures the contribution to $\mathbb{V}\{F\}$ of the interactions between the variables $\boldsymbol{\xi}_\mathfrak{u}$. Since there are 2^d such partial variances, the sensitivity analysis is usually reduced to a simpler characterization, based on first and total-order sensitivity indices associated to individual variables ξ_i or group of variables $\boldsymbol{\xi}_\mathfrak{u}$. The first-order, $\mathtt{S}_\mathfrak{u}$, and total-order, $\mathtt{T}_\mathfrak{u}$, sensitivity indices associated to $\boldsymbol{\xi}_\mathfrak{u}$ are given by [16]

$$\mathbb{V}\{F\}\mathtt{S}_\mathfrak{u} = \sum_{\substack{\mathfrak{v} \in \mathscr{D} \\ \mathfrak{v} \subseteq \mathfrak{u}}} V_\mathfrak{v} = \mathbb{V}\{\mathbb{E}\{F \mid \boldsymbol{\xi}_\mathfrak{u}\}\}, \quad \mathbb{V}\{F\}\mathtt{T}_\mathfrak{u} = \sum_{\substack{\mathfrak{v} \in \mathscr{D} \\ \mathfrak{v} \cap \mathfrak{u} \neq \emptyset}} V_\mathfrak{v} = \mathbb{V}\{F\} - \mathbb{V}\{\mathbb{E}\{F \mid \boldsymbol{\xi}_{\mathfrak{u}_\sim}\}\}. \tag{A.3}$$

The first-order index $\mathtt{S}_\mathfrak{u}$ is then the fraction of variance that arises due to the individual variables in $\boldsymbol{\xi}_\mathfrak{u}$ and their mutual interactions, only; the total-order sensitivity index $\mathtt{T}_\mathfrak{u}$ is the fraction of the variance arising due to the variables in $\boldsymbol{\xi}_\mathfrak{u}$, their mutual interactions and all their interactions with all other variables in $\boldsymbol{\xi}_{\mathfrak{u}_\sim}$. Clearly, $\mathtt{S}_\mathfrak{u} \leq \mathtt{T}_\mathfrak{u}$, and $\mathtt{T}_\mathfrak{u} = 1 - \mathtt{S}_{\mathfrak{u}_\sim}$ [16].

A.2 Sensitivity indices of PC expansion

The partial variances and sensitivity indices of F can be easily computed from the PC expansion of F in the form (2.5)

$$F(\boldsymbol{\xi}) \approx \sum_{\boldsymbol{k} \in \mathscr{K}} f_{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}),$$

where $\mathscr{K}$ is the multi-index set of polynomial tensorizations as discussed in section 2.1 (see (2.4)). Owing to the linear structure and polynomial form of the PC expansion, one can easily partition $\mathscr{K}$ into distinct subsets $\mathscr{K}_\mathfrak{u}$ contributing to the PC expansion of $f_\mathfrak{u}(\boldsymbol{\xi}_\mathfrak{u})$ [42,6]. Specifically, the PC approximation of $f_\mathfrak{u}$ is

$$f_\mathfrak{u}(\boldsymbol{\xi}_\mathfrak{u}) \approx \sum_{\boldsymbol{k} \in \mathscr{K}_\mathfrak{u}} f_{\boldsymbol{k}} \Psi_{\boldsymbol{k}}(\boldsymbol{\xi}), \quad \mathscr{K}_\mathfrak{u} \doteq \{\boldsymbol{k} \in \mathscr{K}; k_{1 \leq i \leq d} > 0 \text{ if } i \in \mathfrak{u}, k_{1 \leq i \leq d} = 0 \text{ if } i \notin \mathfrak{u}\}.$$

Observe that $\mathscr{K}_\emptyset = \{(0\cdots 0)\}$. Then, $\forall \mathfrak{u} \in \mathscr{D}, \mathfrak{u} \neq \emptyset$, we have $V_\mathfrak{u} \approx \sum_{\boldsymbol{k}\in\mathscr{K}_\mathfrak{u}} f_{\boldsymbol{k}}^2$, while the variance of F is approximated by

$$\mathbb{V}\{F\} \approx \sum_{\boldsymbol{k}\in\mathscr{K}\setminus\mathscr{K}_\emptyset} f_{\boldsymbol{k}}^2.$$

The approximations for the first and total-order sensitivity indices $\mathtt{S}_\mathfrak{u}$ and $\mathtt{T}_\mathfrak{u}$ can be easily derived through (A.3), by taking the corresponding unions of subsets $\mathscr{K}_\mathfrak{v}$. For instance, in the case of singleton $\mathfrak{u} = \{i\}$, we have

$$\mathtt{S}_{\{i\}} \approx \frac{1}{\mathbb{V}\{F\}} \sum_{\boldsymbol{k}\in\mathscr{K}^{\mathtt{S}}_{\{i\}}} f_{\boldsymbol{k}}^2, \qquad \mathscr{K}^{\mathtt{S}}_{\{i\}} \doteq \{\boldsymbol{k}\in\mathscr{K}; k_i > 0, k_{j\neq i} = 0\}, \tag{A.4}$$

$$\mathtt{T}_{\{i\}} \approx \frac{1}{\mathbb{V}\{F\}} \sum_{\boldsymbol{k}\in\mathscr{K}^{\mathtt{T}}_{\{i\}}} f_{\boldsymbol{k}}^2, \qquad \mathscr{K}^{\mathtt{T}}_{\{i\}} \doteq \{\boldsymbol{k}\in\mathscr{K}; k_i > 0\}. \tag{A.5}$$

Acknowledgements This work was supported by the US Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, under Award Number DE-SC0008789. The authors wish to express their gratitude to Dr. Cosmin Safta for providing a pre-release version of TChem that enables the simulation of adiabatic combustion at constant volume.

References

1. Alexanderian, A., Winokur, J., Sraj, I., Srinivasan, A., Iskandarani, M., Thacker, W., Knio, O.: Global sensitivity analysis in an ocean general circulation model: a sparse spectral projection approach. Computational Geosciences **16**(3), 1–22 (2012)
2. Chaloner, K., Verdinelli, I.: Bayesian experimental design: A review. Statistical Science **10**(3), pp. 273–304 (1995). URL http://www.jstor.org/stable/2246015
3. Clenshaw, C.W., Curtis, A.R.: A method for numerical integration on an automatic computer. Numerische Mathematik **2**(1), 197–205 (1960)
4. Conrad, P.R., Marzouk, Y.M.: Adaptive smolyak pseudospectral approximations. Submitted: CoRR **abs/1209.1406** (2012)
5. Constantine, P., Eldred, M., Phipps, E.: Sparse pseudospectral approximation method. Computer Methods in Applied Mechanics and Engineering (2012)
6. Crestaux, T., Le Maître, O., Martinez, J.: Polynomial chaos expansion for sensitivity analysis. Reliability Engineering & System Safety **94**(7), 1161–1172 (2009)
7. Donoho, D.L., Elad, M.: Optimally sparse representation in general (nonorthogonal) dictionaries via 1 minimization. Proceedings of the National Academy of Sciences **100**(5), 2197–2202 (2003)
8. Doostan, A., Owhadi, H.: A non-adapted sparse approximation of pdes with stochastic inputs. Journal of Computational Physics **230**(8), 3015–3034 (2011)
9. Eldred, M., Swiler, L.P., Tang, G.: Mixed aleatory-epistemic uncertainty quantification with stochastic expansions and optimization-based interval estimation. Reliability Engineering & System Safety **96**(9), 1092–1113 (2011)
10. Eldred, M.S.: Design under uncertainty employing stochastic expansion methods. International Journal for Uncertainty Quantification **1**(2) (2011)
11. Fejér, L.: Mechanische quadraturen mit positiven cotesschen zahlen. Mathematische Zeitschrift **37**(1), 287–309 (1933)
12. Gerstner, T., Griebel, M.: Numerical integration using sparse grids. Numerical Algorithms pp. 209–232 (1998)
13. Gerstner, T., Griebel, M.: Dimension-adaptive tensor-product quadrature. Computing **71**(1), 65–87 (2003)
14. Ghanem, R.G., Spanos, P.D.: Stochastic Finite Elements: A Spectral Approach. Courier Dover Publications (2003)
15. Hoeffding, W.: A class of statistics with asymptotically normal distribution. The annals of Mathematical Statistics **19**, 293–325 (1948)
16. Homma, T., Saltelli, A.: Importance measures in global sensitivity analysis of nonlinear models. Reliability Engineering & System Safety **52**(1), 1–17 (1996)
17. Hosder, S., Walters, R.W., Balch, M.: Efficient sampling for non-intrusive polynomial chaos applications with multiple uncertain input variables. In: Proceedings of the 48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, AIAA paper, vol. 1939 (2007)
18. Huan, X., Marzouk, Y.M.: Simulation-based optimal bayesian experimental design for nonlinear systems. Journal of Computational Physics (2012)
19. Isukapalli, S., Balakrishnan, S., Georgopoulos, P.: Computationally efficient uncertainty propagation and reduction using the stochastic response surface method. In: Decision and Control, 2004. CDC. 43rd IEEE Conference on, vol. 2, pp. 2237–2243. IEEE (2004)
20. Keese, A., Matthies, H.: Numerical methods and Smolyak quadrature for nonlinear stochastic partial differential equations. Tech. rep., Institute of Scientific Computing TU Braunschweig Brunswick (2003)
21. Kim, D., Rizzi, F., Cheng, K.W., Han, J., Bisetti, F., Knio, O.: Manuscript in preparation (2014)
22. Knio, O., Le Maître, O.: Uncertainty propagation in cfd using polynomial chaos decomposition. Fluid Dynamics Research **38**(9), 616–640 (2006)
23. Kronrod, A.S.: Nodes and weights of quadrature formulas: sixteen-place tables. Consultants Bureau New York (1965)
24. Le Maître, O., Knio, O.: Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics. Scientific Computation. Springer (2010)
25. Le Maître, O., Reagan, M., Najm, H., Ghanem, R., Knio, O.: A stochastic projection method for fluid flow. II. Random process. Journal of Computational Phyics **181**, 9–44 (2002)
26. Li, R., Ghanem, R.: Adaptive polynomial chaos expansions applied to statistics of extremes in nonlinear random vibration. Probabilistic engineering mechanics **13**(2), 125–136 (1998)
27. Lindley, D.V.: On a measure of the information provided by an experiment. The Annals of Mathematical Statistics pp. 986–1005 (1956)
28. Marzouk, Y., Najm, H.: Dimensionality reduction and polynomial chaos acceleration of bayesian inference in inverse problems. Journal of Computational Physics **228**(6), 1862–1902 (2009)
29. Marzouk, Y., Najm, H., Rahn, L.A.: Stochastic spectral methods for efficient bayesian solution of inverse problems. Journal of Computational Physics **224**(2), 560–586 (2007). URL http://dx.doi.org/10.1016/j.jcp.2006.10.010
30. Najm, H.: Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics. Ann. Rev. Fluid Mech. **41**, 35–52 (2009)
31. Patterson, T.: The optimum addition of points to quadrature formulae. Math. Comput **22**, 847–856 (1968)

32. Prager, J., Riedel, U., Warnatz, J.: Modeling ion chemistry and charged species diffusion in lean methane–oxygen flames. Proceedings of the Combustion Institute **31**(1), 1129–1137 (2007)
33. Reagan, M.T., Najm, H.N., Ghanem, R.G., Knio, O.M.: Uncertainty quantification in reacting-flow simulations through non-intrusive spectral projection. Combustion and Flame **132**(3), 545–555 (2003)
34. Safta, C.: private communication (2013)
35. Safta, C., Najm, H., Knio, O.: TChem. http://www.sandia.gov/tchem/
36. Sebastiani, P., Wynn, H.P.: Maximum entropy sampling and optimal bayesian experimental design. Journal of the Royal Statistical Society: Series B (Statistical Methodology) **62**(1), 145–157 (2000)
37. Smith, G.P., Golden, D.M., Frenklach, M., Moriarty, N.W., Eiteneer, B., Goldenberg, M., Bowman, C.T., Hanson, R.K., Song, S., Jr., W.C.G., Lissianski, V.V., Qin, Z.: Gri-mech 3.0. http://www.me.berkeley.edu/gri_mech/ (2012)
38. Smolyak, S.: Quadrature and interpolation formulas for tensor products of certain classes of functions. Dokl. Akad. Nauk SSSR **4**(240-243), 123 (1963)
39. Sobol, I.M.: Sensitivity estimates for nonlinear mathematical models. Math. Model. Comput. Exp. **1**, 407–414 (1993)
40. Sraj, I., Iskandarani, M., Carlisle Thacker, W., Srinivasan, A., Knio, O.: Drag parameter estimation using gradients and hessian from a polynomial chaos model surrogate. Monthly Weather Review **142**(2), 933–941 (2014)
41. Sraj, I., Iskandarani, M., Srinivasan, A., Thacker, W.C., Winokur, J., Alexanderian, A., Lee, C.Y., Chen, S.S., Knio, O.M.: Bayesian inference of drag parameters using axbt data from typhoon fanapi. Monthly Weather Review **141**(7), 2347–2367 (2013)
42. Sudret, B.: Global sensitivity analysis using polynomial chaos expansions. Reliability Engineering & System Safety **93**(7), 964–979 (2008)
43. Waldvogel, J.: Fast construction of the fejér and clenshaw–curtis quadrature rules. BIT Numerical Mathematics **46**(1), 195–202 (2006)
44. Winokur, J., Conrad, P., Sraj, I., Knio, O., Srinivasan, A., Thacker, W., Marzouk, Y., Iskandarani, M.: A priori testing of sparse adaptive polynomial chaos expansions using an ocean general circulation model database. Computational Geosciences **17**(6), 899–911 (2013). DOI 10.1007/s10596-013-9361-3. URL http://dx.doi.org/10.1007/s10596-013-9361-3