

SPARSE ADAPTIVE TAYLOR APPROXIMATION ALGORITHMS FOR PARAMETRIC AND STOCHASTIC ELLIPTIC PDES *

ABDELLAH CHKIFA^{1,2}, ALBERT COHEN^{1,2}, RONALD DEVORE³
AND CHRISTOPH SCHWAB⁴

Abstract. The numerical approximation of parametric partial differential equations is a computational challenge, in particular when the number of involved parameter is large. This paper considers a model class of second order, linear, parametric, elliptic PDEs on a bounded domain D with diffusion coefficients depending on the parameters in an affine manner. For such models, it was shown in [9, 10] that under very weak assumptions on the diffusion coefficients, the entire family of solutions to such equations can be simultaneously approximated in the Hilbert space $V = H_0^1(D)$ by multivariate sparse polynomials in the parameter vector y with a controlled number N of terms. The convergence rate in terms of N does not depend on the number of parameters in V , which may be arbitrarily large or countably infinite, thereby breaking the curse of dimensionality. However, these approximation results do not describe the concrete construction of these polynomial expansions, and should therefore rather be viewed as benchmark for the convergence analysis of numerical methods. The present paper presents an adaptive numerical algorithm for constructing a sequence of sparse polynomials that is proved to converge toward the solution with the optimal benchmark rate. Numerical experiments are presented in large parameter dimension, which confirm the effectiveness of the adaptive approach.

Mathematics Subject Classification. 65N35, 65L10, 35J25.

Received June 25, 2011. Revised February 27, 2012.
Published online November 29, 2012.

Keywords and phrases. Parametric and stochastic PDE's, sparse polynomial approximation, high dimensional problems, adaptive algorithms.

* *This research was supported by the Office of Naval Research Contracts ONR-N00014-08-1-1113, ONR N00014-09-1-0107, the AFOSR Contract FA95500910500, the ARO/DoD Contracts W911NF-05-1-0227 and W911NF-07-1-0185, the National Science Foundation Grant DMS 0915231; the excellency chair of the Foundation "Science Mathématiques de Paris" awarded to Ronald DeVore in 2009. This publication is based on work supported by Award No. KUS-C1-016-04, made by King Abdullah University of Science and Technology (KAUST). This research is also supported by the Swiss National Science Foundation under Grant SNF 200021-120290/1 and by the European Research Council under grant ERC AdG247277. CS acknowledges hospitality by the Hausdorff Institute for Mathematics, Bonn, Germany.*

¹ UPMC Univ. Paris 06, UMR 7598, Laboratoire Jacques-Louis Lions, 75005 Paris, France. chkifa@ann.jussieu.fr; cohen@ann.jussieu.fr

² CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, 75005 Paris, France.

³ Department of Mathematics, Texas A&M University, College Station, 77843 TX, USA. rdevore@math.tamu.edu

⁴ Seminar for Applied Mathematics, ETH Zürich, 8092 Zürich, Switzerland. schwab@math.ethz.ch

1. INTRODUCTION

We consider parametric partial differential equations of the general form

$$\mathcal{D}(u, y) = 0$$

where $u \mapsto \mathcal{D}(u, y)$ is a partial differential operator that depends on a vector of parameters y , and therefore so does the solution $u = u(y)$. Parametric problems of this type arise in modeling complex systems in various contexts:

- Stochastic modelling: the parameters y are realizations of random variables which reflects the fact that the diffusion coefficient is not known exactly and is therefore modelled as a random field. The user is interested in the resulting statistical properties of the solution u . This is the point of view adopted for example in [1, 2, 13, 15, 19, 22, 23]. Here, the parameter y is generally infinite dimensional and any numerically viable approximation must necessarily address the issue of dimensional reduction, in addition to that of discretization.
- Deterministic modelling: the parameters y are known or controlled by the user, who is interested in studying the dependence of u with respect to these parameters for various purposes (for example, optimizing an output of the equation with respect to y). This is the point of view adopted for example in [5, 20].

In both of these settings, the main computational challenge is to *simultaneously* solve the entire parametric family of equations up to a prescribed accuracy ε , with reasonable computation cost. This task is particularly difficult when the number of involved parameters is finite, but large, and more so if there are even countably many parameters, *i.e.* when the dimension of the parameter domain is infinite; this is the problem analyzed in the present paper.

To avoid unnecessary technicalities, we present our algorithms for the model parametric elliptic boundary value problem

$$-\operatorname{div}(a\nabla u) = f \text{ in } D \subset \mathbb{R}^m, \quad u = 0 \text{ on } \partial D, \quad (1.1)$$

where $f \in H^{-1}(D)$ and

$$a(x, y) := \bar{a}(x) + \sum_{j=1}^d y_j \psi_j(x),$$

for $y = (y_j)_{j=1, \dots, d} \in U := [-1, 1]^d$. Here, \bar{a} and the ψ_j are functions in $L^\infty(D)$.

We are interested in algorithms that can be applied for a large number d of parameters and whose performance is *independent of d* . It is also of mathematical interest to allow the possibility of a countably infinite number, $d = \infty$, of parameters. This corresponds to an infinite parameter sequence $y = (y_j)_{j \geq 1} \in U := [-1, 1]^{\mathbb{N}}$.

Note that we may view the finite dimensional setting as a particular case of the latter one with $\psi_j = 0$ for $j > d$. Having this in mind, and since we are interested in deriving numerical results that are robust with respect to the parameter dimension d , we place ourselves in the infinite dimensional setting

$$a(x, y) := \bar{a}(x) + \sum_{j \geq 1} y_j \psi_j(x),$$

throughout this paper and understand sums and products over $j \geq 1$ to be either over the finite range $j = 1, \dots, d$ or over all $j \in \mathbb{N}$. Here and throughout, we denote by $\mathbb{N} = \{1, 2, 3, \dots\}$ the set of natural numbers and by $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. In the stochastic context, an infinite series $\bar{a} + \sum_{j \geq 1} y_j \psi_j$ could result, for example, from a Karh unen-Lo eve expansion of a random field $a(x, \omega)$, see, for example, [25].

We assume at a minimum that the sequence $(\|\psi_j\|_{L^\infty(D)})_{j \geq 1}$ is bounded. By rearranging indices, if necessary, we may assume without loss of generality that $\|\psi_j\|_{L^\infty(D)}$, $j = 1, 2, \dots$, is a non-increasing sequence. We work under the (minimal) *uniform ellipticity assumption* **UEA**(r, R) for some $0 < r \leq R < \infty$:

$$0 < r \leq \bar{a}(x) + \sum_{j \geq 1} y_j \psi_j(x) \leq R, \quad x \in D, \quad y \in U. \quad (1.2)$$

Observe that the lower inequality in this assumption is equivalent to

$$\sum_{j \geq 1} |\psi_j(x)| \leq \bar{a}(x) - r, \quad x \in D. \tag{1.3}$$

Assumption **UEA**(r, R) ensures existence and uniqueness of the solution $u(y)$ in $V = H_0^1(D)$, for all $y \in U$, with the *a priori* estimate

$$\|u(y)\|_V \leq C_0 := \frac{\|f\|_{V^*}}{r},$$

where $\|v\|_V := \|\nabla v\|_{L^2(D)}$. We also introduce the average energy norm

$$\|v\|_{\bar{a}} := \left(\int_D \bar{a} |\nabla v|^2 \right)^{1/2},$$

which, under **UEA**(r, R), is equivalent to the $H_0^1(D)$ norm $\|\cdot\|_V$ in the sense that

$$\sqrt{r}\|v\|_V \leq \|v\|_{\bar{a}} \leq \sqrt{R}\|v\|_V, \quad v \in V.$$

We are interested in approximating the map $y \mapsto u(y)$ by multivariate polynomials in y with coefficients in V . One way to obtain such an approximation is by truncating in some way the Taylor expansion of u at $y = 0$. This expansion has the general form $\sum_{\nu \in \mathcal{F}} t_\nu y^\nu$, with the notation

$$y^\nu := \prod_{j \geq 1} y_j^{\nu_j} \quad \text{and} \quad t_\nu := \frac{1}{\nu!} \partial^\nu u(0) \in V \quad \text{with} \quad \nu! := \prod_{j \geq 1} \nu_j! \quad \text{and} \quad 0! := 1,$$

for $\nu = (\nu_j)_{j \geq 1} \in \mathcal{F}$. In the finite dimensional setting $d < \infty$, the index set \mathcal{F} coincides with $\mathbb{N}_0^d = \{0, 1, \dots\}^d$. In the infinite dimensional setting, \mathcal{F} denotes the (countable) set of all sequences of nonnegative integers which are *finitely supported* (i.e. those sequence for which only finitely many terms are nonzero).

We are thus interested in the summability properties in V of partial sums of the Taylor series $\sum_{\nu \in \mathcal{F}} t_\nu y^\nu$. The following result was proved in [10] (Thm. 1.2).

Theorem 1.1. *If $(\|\psi_j\|_{L^\infty(D)})_{j \geq 1} \in \ell^p(\mathbb{N})$ for some $0 < p < 1$ and if **UEA**(r, R) holds, then*

$$u(y) = \sum_{\nu \in \mathcal{F}} t_\nu y^\nu, \tag{1.4}$$

in the sense of unconditional convergence in $L^\infty(U, V)$ and $(\|t_\nu\|_V)_{\nu \in \mathcal{F}} \in \ell^p(\mathcal{F})$.

Analogous results for other types of partial differential equations are available in [17, 18]. Theorems such as 1.1 have important implications concerning the efficient numerical approximation of the parametric solution map $y \mapsto u(y)$ which, in turn, opens a perspective for novel computational approaches in parameter identification problems [24]. To describe these, for any sequence $(a_\nu)_{\nu \in \mathcal{F}}$ of real numbers indexed by the countable set \mathcal{F} and any $k \geq 1$, we define the sets $\Lambda_k^* := \Lambda_k^*((a_\nu)_{\nu \in \mathcal{F}})$ of k largest elements in absolute value. The sets Λ_k^* are generally not unique because of possible ties in the size of the $|a_\nu|$. However, if $(a_\nu)_{\nu \in \mathcal{F}} \in \ell^p(\mathcal{F})$ and if Λ_k^* is any of these sets, then for any $q > p$

$$\left(\sum_{\nu \notin \Lambda_k^*} |a_\nu|^q \right)^{1/q} \leq \|(a_\nu)\|_{\ell^p(\mathcal{F})} (k+1)^{-\frac{1}{p} + \frac{1}{q}}. \tag{1.5}$$

This elementary fact is proved (see [11]) by introducing the decreasing rearrangement $(a_j^*)_{j>0}$ of the sequence $(|a_\nu|)_{\nu \in \mathcal{F}}$ and by combining the two observations

$$\left(\sum_{\nu \notin \Lambda_k^*} |a_\nu|^q \right)^{1/q} = \left(\sum_{j>k} (a_j^*)^q \right)^{1/q} \leq \left(\sum_{j>k} (a_{k+1}^*)^{q-p} (a_j^*)^p \right)^{1/q} \leq (a_{k+1}^*)^{1-p/q} \| (a_\nu) \|_{\ell_p(\mathcal{F})}^{p/q},$$

and

$$(k+1)(a_{k+1}^*)^p \leq \sum_{j \leq k+1} (a_j^*)^p \leq \| (a_\nu) \|_{\ell_p(\mathcal{F})}^p.$$

Working under the assumptions of the above theorem, for each $k = 1, 2, \dots$, we denote by $\Lambda_k^* \subset \mathcal{F}$ the set of indices $\nu \in \mathcal{F}$ corresponding to the k largest values of the $\|t_\nu\|_V$, with ties broken in an arbitrary (but consistent) way. We then have

$$\sup_{y \in U} \|u(y) - \sum_{\nu \in \Lambda_k^*} t_\nu y^\nu\|_V \leq \sum_{\nu \notin \Lambda_k^*} \|t_\nu\|_V \leq \| \|t_\nu\|_V \|_{\ell_p(\mathcal{F})} (k+1)^{-s}, \quad s := \frac{1}{p} - 1. \tag{1.6}$$

The sparse polynomials $\sum_{\nu \in \Lambda_k^*} t_\nu y^\nu$ therefore provide a simultaneous approximation of the family $\{u(y) ; y \in U\}$ at the cost of computing k functions $t_\nu \in V$. Quite remarkably the rate k^{-s} and the constant in (1.6) is independent of the number of parameters y_j which may be countably infinite. Thus, (1.6) implies that one can in principle overcome the curse of dimensionality in the approximation of $u(y)$.

In computation, however, the sets Λ_k^* in (1.6) are not known to us and to find them we would ostensibly have to compute all of the t_ν , which is infeasible. To obtain computable sequences of index sets, we shall not insist on optimality: we shall say that a nested sequence $(\Lambda_n)_{n \geq 0}$ of finite subsets $\Lambda_n \subset \mathcal{F}$ is near optimal if there is a constant $C \geq 1$ such that for s as in (1.6) and for every $n \geq 0$

$$\sum_{\nu \notin \Lambda_n} \|t_\nu\|_V \leq C \| \|t_\nu\|_V \|_{\ell_p(\mathcal{F})} (\#\Lambda_n)^{-s}. \tag{1.7}$$

The goal of the present paper is to give a concrete algorithm that adaptively builds a near optimal sequence $(\Lambda_n)_{n \geq 0}$ and corresponding Taylor coefficients $(t_\nu)_{\nu \in \Lambda_n}$ at a cost that scales linearly in $\#\Lambda_n$. We should point out that a similar program was developed when solving a *single* PDE by either adaptive wavelet methods [7, 8, 14] or by adaptive finite element methods [4, 12, 21, 26]. In these papers, it was proved that certain iterative refinement algorithms based on *a posteriori* analysis generate adaptive wavelet sets or adaptive meshes such that the approximate solution converges with the optimal rate allowed by the exact solution. A common point between these algorithms and the one that we study in this paper is the use of a *bulk chasing procedure* in order to build the set Λ_{n+1} from the set Λ_n . However, our present setting is significantly different, since the index sets Λ_n are picked from the infinite dimensional lattice \mathcal{F} and the coefficients associated to each $\nu \in \Lambda_n$ are functions in V instead of numbers.

Our paper is organized as follows. In Section 2, we show that the sets Λ_n may be picked from a restricted class called *monotone sets* while retaining the optimal rate.

We show in Section 3 how the Taylor coefficients associated to a monotone set may be recursively computed by solving one elliptic boundary value problem at a time per coefficient, and we establish a useful estimate for the energy of the Taylor coefficients outside a monotone set.

A first adaptive algorithm is proposed in Section 4 and proved to converge with the optimal rate in the sense that the sets Λ_n generated by the algorithm satisfy (1.7). A defect of this algorithm is that the bulk chasing procedure at step n requires the computation of the $\|t_\nu\|_{\bar{a}}$ for ν in a neighbourhood of Λ_n which has infinite cardinality and is therefore not practical.

We remedy this defect in Section 5 by introducing a second algorithm which operates the bulk search on a finite set, and which is also proved to converge with the optimal rate s in (1.6).

We study in Section 6 the additional error which is induced on the approximation of the map $y \mapsto u(y)$ by the spatial discretization when solving the boundary value problems that give the Taylor coefficients, for example by a finite element method on D . We prove that the additional error introduced by the finite element discretization of the coefficients is independent of the number of computed Taylor coefficients.

Finally, numerical experiments are presented in Section 7, for finite but high dimensional test cases ($y \in [-1, 1]^d$ with d up to 255), and using finite element for the spatial discretization. We test the adaptive bulk search strategy, and compare it with non-adaptive strategies based on *a priori* choices of the sets Λ_n . These experiments confirm that, without any information based on *a priori* analysis, the adaptive approach produces near-optimal sets of active indices in terms of convergence rate. We also propose alternate adaptive strategies which are computationally much cheaper than the bulk search and exhibit, in our numerical examples, the same convergence rate as the approximations obtained by bulk search yet without complete theoretical justification. In the practically relevant case where the goal of computation is to compute an average in y of the solution (corresponding to an expectation of the random solution) we show that the results based on our adaptive algorithm strongly outperform those using the Monte-Carlo method

2. MONOTONE SETS

In this section, we give a finer description of the approximation properties of the Taylor series by introducing the notion of *monotonicity*. This notion is based on the following ordering of \mathcal{F} : for $\mu, \nu \in \mathcal{F}$, $\mu \leq \nu$ if and only if $\mu_j \leq \nu_j$ for all $j \geq 1$. We will also say that $\mu < \nu$ if and only if $\mu \leq \nu$ and $\mu_j < \nu_j$ for at least one value of j .

Definition 2.1. A sequence $(a_\nu)_{\nu \in \mathcal{F}}$ of nonnegative real numbers is said to be *monotone decreasing* if and only if for all $\mu, \nu \in \mathcal{F}$

$$\mu \leq \nu \Rightarrow a_\nu \leq a_\mu.$$

A non empty set $\Lambda \subset \mathcal{F}$ is called *monotone* if and only if $\nu \in \Lambda$ and $\mu \leq \nu \Rightarrow \mu \in \Lambda$. For a monotone set $\Lambda \subset \mathcal{F}$, we define its margin $\mathcal{M} = \mathcal{M}(\Lambda)$ as follows:

$$\mathcal{M}(\Lambda) := \{\nu \notin \Lambda ; \exists j > 0 : \nu - e_j \in \Lambda\}, \tag{2.1}$$

where $e_j \in \mathcal{F}$ is the Kronecker sequence: $(e_j)_i = \delta_{ij}$ for $i, j \in \mathbb{N}$.

Notice that the margin $\mathcal{M}(\Lambda)$ is an infinite set even when Λ is finite since there are infinitely many variables. In the finite dimensional setting $d < \infty$, the margin is a finite set. Any nonempty monotone set contains the null index $(0, 0, \dots)$, which we will denote in what follows with slight abuse of notation by 0 . Intersections and unions of monotone sets are also monotone. Also, note that $\Lambda \cup \mathcal{M}(\Lambda)$ is a monotone set.

For any $\nu \in \mathcal{F}$, we let $|\nu| := \sum_{i \geq 1} \nu_i$. We say that ν is maximal in a set $\Lambda \subset \mathcal{F}$ if and only if there exists no $\mu > \nu$ in Λ . If $\Lambda \subset \mathcal{F}$ satisfies $N := N(\Lambda) := \max_{\nu \in \Lambda} |\nu| < \infty$, then any $\nu \in \Lambda$ for which $|\nu| = N$ is a maximal element. In particular, any finite set Λ has at least one maximal element. If Λ is monotone and if ν is maximal in Λ , then $\Lambda - \{\nu\}$ is monotone.

Remark 2.2. If $(a_\nu)_{\nu \in \mathcal{F}}$ is a monotone sequence, the set $\Lambda_k^* = \Lambda_k^*((a_\nu)_{\nu \in \mathcal{F}})$ of indices corresponding to the k -largest a_ν in absolute value is always a monotone set whenever it is unique: it is then equivalently given by $\Lambda_k^* = \{\nu \in \mathcal{F} : a_\nu \geq \eta\}$ for some threshold η which depends on k . In the case of non-uniqueness, there exists at least one realization of a Λ_k^* which is monotone. We refer to such a set as a monotone realization of Λ_k^* . Such a realization may be constructed as follows: consider the largest threshold η such that the monotone set $\{\nu \in \mathcal{F} : a_\nu \geq \eta\}$ has more than k elements, and trim this set by removing iteratively a maximal ν until it has exactly k elements.

Remark 2.3. We localize the notion of monotone sequences and monotone sets as follows: if $\mathcal{F}_0 \subset \mathcal{F}$ is any subset, we say that the sequence $(a_\nu)_{\nu \in \mathcal{F}}$ is monotone on \mathcal{F}_0 (or that $(a_\nu)_{\nu \in \mathcal{F}_0}$ is monotone) if and only if

$$\mu, \nu \in \mathcal{F}_0 \text{ and } \mu \leq \nu \Rightarrow a_\nu \leq a_\mu.$$

Clearly, a monotone sequence is monotone on any set \mathcal{F}_0 . Likewise we say that a subset $\mathcal{F}_1 \subset \mathcal{F}_0$ is monotone in \mathcal{F}_0 if and only if

$$\nu \in \mathcal{F}_1, \mu \in \mathcal{F}_0 \text{ and } \mu \leq \nu \Rightarrow \mu \in \mathcal{F}_1.$$

In the case where \mathcal{F}_0 is monotone, this is equivalent to saying that \mathcal{F}_1 is monotone. If (a_ν) is monotone on \mathcal{F}_0 , a set of indices corresponding to the k -largest a_ν in absolute value with $\nu \in \mathcal{F}_0$ is monotone in \mathcal{F}_0 whenever it is unique. If it is not unique, there exists at least one realization of such a set which is monotone. This set may be obtained by the same trimming procedure as in Remark 2.2.

The *monotone majorant* of a bounded sequence $(a_\nu)_{\nu \in \mathcal{F}}$ is the sequence

$$\mathbf{a}_\nu := \max_{\mu \geq \nu} |a_\mu|, \quad \nu \in \mathcal{F}.$$

We define $\ell_m^p(\mathcal{F})$ as the set of all sequences which have their monotone majorant in $\ell^p(\mathcal{F})$. Clearly, $\ell_m^p(\mathcal{F})$ is a linear space with respect to addition of sequences and scalar multiplication. We equip this space with the norm

$$\|(a_\nu)\|_{\ell_m^p(\mathcal{F})} := \|(\mathbf{a}_\nu)\|_{\ell^p(\mathcal{F})}.$$

Now, if $(a_\nu)_{\nu \in \mathcal{F}} \in \ell_m^p(\mathcal{F})$, $0 < p < 1$, and Λ_k is any monotone realization of $\Lambda_k^*((\mathbf{a}_\nu)_{\nu \in \mathcal{F}})$, then the sets Λ_k are monotone and satisfy

$$\sum_{\nu \notin \Lambda_k} |a_\nu| \leq \sum_{\nu \notin \Lambda_k} \mathbf{a}_\nu \leq \|(a_\nu)\|_{\ell_m^p(\mathcal{F})} (k+1)^{-s}, \quad s := \frac{1}{p} - 1. \tag{2.2}$$

Theorem 2.4. Under the assumptions of Theorem 1.1, the sequence $(\|t_\nu\|_V)_{\nu \in \mathcal{F}}$ belongs to $\ell_m^p(\mathcal{F})$.

Proof. The result will follow from the estimates that underpin the Proof of Theorem 1.1 given in [10]. A sequence $\rho = (\rho_j)_{j \geq 1}$ is said to be admissible of order δ if

$$\sum_{j \geq 1} \rho_j |\psi_j(x)| \leq \bar{a}(x) - \delta, \quad x \in D.$$

We denote by \mathcal{A}_δ the set of all δ -admissible sequences ρ for which $\rho_j \geq 1$, for all j . It was shown in [10] that for any $0 < \delta < r$

$$\|t_\nu\|_V \leq \frac{\|f\|_{V^*}}{\delta} \inf_{\rho \in \mathcal{A}_\delta} \rho^{-\nu}. \tag{2.3}$$

In particular, taking $\delta = \frac{r}{2}$,

$$\|t_\nu\|_V \leq b_\nu := 2C_0 \inf_{\rho \in \mathcal{A}_{\frac{r}{2}}} \rho^{-\nu}. \tag{2.4}$$

It was moreover shown that under the assumptions of Theorem 1.1 the sequence $(b_\nu)_{\nu \in \mathcal{F}}$ belongs to $\ell^p(\mathcal{F})$ which thus leads to the Proof of Theorem 1.1. We now observe that the sequence b_ν is monotone because for any $\rho \in \mathcal{A}_{\frac{r}{2}}$

$$\mu \leq \nu \Rightarrow \rho^{-\nu} \leq \rho^{-\mu},$$

and thus

$$\mu \leq \nu \Rightarrow b_\nu \leq b_\mu.$$

Therefore, if (\mathbf{a}_ν) denotes the monotone majorant of the sequence $(\|t_\nu\|_V)$, we also find that

$$\mathbf{a}_\nu \leq b_\nu.$$

It follows that $\|(\|t_\nu\|_V)\|_{\ell_m^p(\mathcal{F})} \leq \|(\mathbf{a}_\nu)\|_{\ell^p(\mathcal{F})} < \infty$. □

3. RECURSIVE ESTIMATES

It was shown in [9] that the Taylor coefficients satisfy a recursion relation obtained by differentiating the variational formulation

$$\int_D a(x, y) \nabla u(x, y) \nabla v(x) dx = \int_D f(x) v(x) dx \quad v \in V,$$

at $y = 0$. Namely, we obtain by induction that $t_\nu \in V$ is the solution to the elliptic boundary value problem given in weak form by (see Eqs. (4.6) and (4.10) of [9])

$$\int_D \bar{a} \nabla t_\nu \nabla v = - \sum_{j \text{ s.t. } \nu_j \neq 0} \int_D \psi_j \nabla t_{\nu - e_j} \nabla v, \quad v \in V. \tag{3.1}$$

This recurrence allows one to compute all Taylor coefficients from the first coefficient $t_0 = u(0)$ corresponding to $\nu = 0$ which satisfies

$$\int_D \bar{a} \nabla t_0 \nabla v = \int_D f v, \quad v \in V. \tag{3.2}$$

In practice, these boundary value problems can only be solved approximately by space discretization, for example by the finite element method. We shall deal with this issue in §6 and assume for the moment that they can be solved exactly. For any monotone set of indices Λ , the recursion (3.1) determines the Taylor coefficients $\{t_\nu \in V : \nu \in \Lambda\}$ uniquely; determining them requires the successive numerical solution of the “nominal” elliptic problems (3.2) with $\#\Lambda$ many right hand sides. In particular, then, for computing numerical approximations of the $(t_\nu)_{\nu \in \Lambda}$, a discretized single, parameter-independent “nominal” elliptic problem (3.2) in the domain D must be solved with $\#\Lambda$ many load cases. Since our adaptive algorithms will be based on the norms $\|t_\nu\|_{\bar{a}}$ of Taylor coefficients through the recursion (3.1), we introduce the abbreviated notation

$$\bar{t}_\nu := \|t_\nu\|_{\bar{a}}, \quad \nu \in \mathcal{F},$$

and the following quantities for bounding for the right hand side of (3.1)

$$d_{\mu, j} := \int_D |\psi_j| |\nabla t_\mu|^2, \quad \mu \in \mathcal{F}, j \geq 1.$$

We observe that **UEA**(r, R) implies for almost every $x \in D$

$$\sum_{j \geq 1} |\psi_j(x)| \leq \gamma \bar{a}(x) \tag{3.3}$$

with

$$\gamma = 1 - \frac{r}{R} < 1. \tag{3.4}$$

It follows that for any $\mu \in \mathcal{F}$, we have

$$\sum_{j \geq 1} d_{\mu, j} \leq \gamma \bar{t}_\mu^2. \tag{3.5}$$

Lemma 3.1. *Under assumption **UEA** (r, R), we have for any $\nu \in \mathcal{F}$,*

$$\bar{t}_\nu^2 \leq \alpha \sum_{j \text{ s.t. } \nu_j \neq 0} d_{\nu - e_j, j}, \tag{3.6}$$

with

$$\alpha := \frac{R}{R + r} < 1. \tag{3.7}$$

Proof. Taking $v = t_\nu$ in (3.1), we find that

$$\bar{t}_\nu^2 = - \sum_{j \text{ s.t. } \nu_j \neq 0} \int_D \psi_j \nabla t_{\nu - e_j} \nabla t_\nu, \quad (3.8)$$

and therefore

$$\bar{t}_\nu^2 \leq \frac{1}{2} \sum_{j \text{ s.t. } \nu_j \neq 0} \int_D |\psi_j| |\nabla t_{\nu - e_j}|^2 + \frac{1}{2} \sum_{j \text{ s.t. } \nu_j \neq 0} \int_D |\psi_j| |\nabla t_\nu|^2. \quad (3.9)$$

Using (3.3) in the second term of (3.9) gives

$$(1 - \gamma/2) \bar{t}_\nu^2 \leq \frac{1}{2} \sum_{j \text{ s.t. } \nu_j \neq 0} \int_D |\psi_j| |\nabla t_{\nu - e_j}|^2,$$

from which we derive (3.6). \square

For any set $\Lambda \subset \mathcal{F}$, we introduce

$$e(\Lambda) := \sum_{\nu \in \Lambda} \bar{t}_\nu^2, \quad \sigma(\Lambda) := \sum_{\nu \in \mathcal{F} \setminus \Lambda} \bar{t}_\nu^2, \quad (3.10)$$

which is a measure of the energy of the Taylor coefficients on Λ and the energy on its complement respectively. Our next lemma shows that if Λ is a monotone set, the energy outside Λ is controlled by the energy on the margin $\mathcal{M} = \mathcal{M}(\Lambda)$.

Lemma 3.2. *Under assumption (UEA)(r, R), we have for any monotone set Λ and its margin \mathcal{M} ,*

$$\sigma(\Lambda) \leq \frac{1}{1 - \delta} e(\mathcal{M}), \quad (3.11)$$

with

$$\delta = \frac{R - r}{R + r} < 1. \quad (3.12)$$

Proof. We first note that

$$\sigma(\Lambda) = e(\mathcal{M}) + \sigma(\tilde{\Lambda}), \quad (3.13)$$

where we have set $\tilde{\Lambda} := \Lambda \cup \mathcal{M}$. According to Lemma 3.1, we may write

$$\sigma(\tilde{\Lambda}) \leq \alpha \sum_{\nu \in \mathcal{F} \setminus \tilde{\Lambda}} \left(\sum_{j \text{ s.t. } \nu_j \neq 0} d_{\nu - e_j, j} \right) \leq A + B, \quad (3.14)$$

where

$$A := \alpha \sum_{\nu \in \mathcal{F} \setminus \tilde{\Lambda}} \left(\sum_{j \text{ s.t. } \nu - e_j \in \mathcal{F} \setminus \tilde{\Lambda}} d_{\nu - e_j, j} \right) = \alpha \sum_{\mu \in \mathcal{F} \setminus \tilde{\Lambda}} \left(\sum_{j \text{ s.t. } \mu + e_j \in \mathcal{F} \setminus \tilde{\Lambda}} d_{\mu, j} \right),$$

and

$$B := \alpha \sum_{\nu \in \mathcal{F} \setminus \tilde{\Lambda}} \left(\sum_{j \text{ s.t. } \nu - e_j \in \tilde{\Lambda}} d_{\nu - e_j, j} \right) = \alpha \sum_{\mu \in \mathcal{M}} \left(\sum_{j \text{ s.t. } \mu + e_j \in \mathcal{F} \setminus \tilde{\Lambda}} d_{\mu, j} \right).$$

In this splitting, we have used the fact that if $\nu \in \mathcal{F} \setminus \tilde{\Lambda}$ and $\nu_j \neq 0$, we have either $\nu - e_j \in \mathcal{F} \setminus \tilde{\Lambda}$ or $\nu - e_j \in \mathcal{M}$. Using (3.5), we may control the first term A by

$$A \leq \alpha\gamma \sum_{\mu \in \mathcal{F} \setminus \tilde{\Lambda}} \bar{t}_\mu^2 = \alpha\gamma\sigma(\tilde{\Lambda}),$$

and by the same argument we obtain

$$B \leq \alpha\gamma e(\mathcal{M}).$$

Combining these estimates with (3.14), it follows that

$$(1 - \alpha\gamma)\sigma(\tilde{\Lambda}) \leq \alpha\gamma e(\mathcal{M}),$$

and thus by (3.13)

$$\sigma(\Lambda) \leq \left(1 + \frac{\alpha\gamma}{1 - \alpha\gamma}\right) e(\mathcal{M}),$$

which gives the final result. □

4. A BULK CHASING ALGORITHM

In this section, we introduce the notion of bulk chasing and show how this idea can be used to build an adaptive algorithm for generating a near optimal sequence of sets (Λ_n) in the sense of (1.7). We shall see that this algorithm is not numerically feasible but nevertheless will guide us in the construction of more practical algorithms in the sections that follow.

To begin the discussion, let us assume that we have a finite monotone set Λ with margin $\mathcal{M} = \mathcal{M}(\Lambda)$, for which we have already computed t_ν , $\nu \in \Lambda$. From this knowledge we can directly compute certain t_ν , $\nu \in \mathcal{M}$ from the recurrence (3.1). Namely, if $\mathcal{I}_1(\mathcal{M})$ is the set of $\nu \in \mathcal{M}$ such that each $\nu - e_j \in \Lambda$ whenever $\nu_j \geq 1$, then we can compute t_ν for all $\nu \in \mathcal{I}_1(\mathcal{M})$ since we already know each of the $t_{\nu - e_j}$ that occur in (3.1). We can then repeat this process and compute t_ν for any $\nu \in \mathcal{I}_2(\mathcal{M})$ where $\mathcal{I}_2(\mathcal{M})$ is the set of all $\nu \in \mathcal{M} \setminus \mathcal{I}_1(\mathcal{M})$ such that $\nu - e_j \in \Lambda \cup \mathcal{I}_1(\mathcal{M})$ whenever $\nu_j > 0$. Continuing in this way, we can compute all of the $t_\nu \in \mathcal{M}$. Notice that one only needs a finite number of the sets $\mathcal{I}_j(\mathcal{M})$ to exhaust \mathcal{M} since Λ is finite.

For a fixed $0 < \theta < 1$, we consider the following (not yet practical) algorithm:

Algorithm 1. Define $\Lambda_0 := \{0\}$ and compute $t_0 := u(0)$ and $\bar{t}_0 := \|t_0\|_{\bar{a}}$. For $n = 0, 1, \dots$ do the following:

- given that Λ_n has been defined and $(t_\nu)_{\nu \in \Lambda_n}$ have been computed, we define $\mathcal{M}_n = \mathcal{M}(\Lambda_n)$ and compute t_ν , $\nu \in \mathcal{M}_n$, by using the recursion (3.1);
- compute \bar{t}_ν for $\nu \in \mathcal{M}_n$ and their monotone envelope inside \mathcal{M}_n defined by

$$\bar{\mathbf{t}}_\nu(\mathcal{M}_n) := \sup\{\bar{t}_\mu : \mu \geq \nu \text{ and } \mu \in \mathcal{M}_n\};$$

- enrich Λ_n by the indices $\nu \in \mathcal{M}_n$ corresponding to the k largest $\bar{\mathbf{t}}_\nu(\mathcal{M}_n)$ for $k = 1, 2, \dots$, until the obtained monotone set $\Lambda_{n+1} \subset \Lambda_n \cup \mathcal{M}_n$ satisfies $e(\Lambda_{n+1} \cap \mathcal{M}_n) \geq \theta e(\mathcal{M}_n)$;
- go to step $n + 1$.

Remark 4.1. Since its construction is based on sorting the quantities $\bar{\mathbf{t}}_\nu(\mathcal{M}_n)$, similar to the trimming procedure of Remark 2.2, the new set Λ_{n+1} can be made monotone. Another option would be to define Λ_{n+1} as the smallest monotone set such that $\Lambda_n \subset \Lambda_{n+1} \subset \Lambda_n \cup \mathcal{M}_n$ and such that $e(\Lambda_{n+1} \cap \mathcal{M}_n) \geq \theta e(\mathcal{M}_n)$. The Proof of Theorem 4.2 below also applies in this case. However, we do not currently know how to efficiently implement the search for such a minimal monotone set.

Algorithm 1 is not satisfactory for several reasons. A first problem is that we can only solve the boundary value problems (3.1) approximately, for example using a finite element discretization. We analyze the additional error induced by this discretization in Section 6. Another problem is that, in our infinite dimensional setting the margin \mathcal{M}_n has infinite cardinality, and therefore there are infinitely many \bar{t}_ν to be computed which requires in principle solving infinitely many boundary value problems for the corresponding t_ν . Although this problem does not occur in the finite dimensional setting $d < \infty$, it is still reflected by the fact that the size of \mathcal{M}_n is potentially much larger than that of Λ_n as d gets large and therefore solving the boundary value problems for all $\nu \in \mathcal{M}_n$ becomes the main source of computational complexity. We propose a solution to this problem in the next section Section 5.

For the present, we remain with the above algorithm and prove its optimality. We first establish the following result on the decay of the energy of Taylor coefficients.

Theorem 4.2. *Under the assumptions of Theorem 1.1, we have $(\bar{t}_\nu)_{\nu \in \mathcal{F}} \in \ell_m^p(\mathcal{F})$ and the sets Λ_n satisfy*

$$\sigma(\Lambda_n) \leq C_1 \|(\bar{t}_\nu)_{\nu \in \mathcal{F}}\|_{\ell_m^p(\mathcal{F})}^2 (\#\Lambda_n)^{-2t}, \quad t := \frac{1}{p} - \frac{1}{2}, \quad (4.1)$$

where C_1 only depends on (r, R, θ, t) .

Proof. We have shown in Theorem 2.4 that $(\|t_\nu\|_V)_{\nu \in \mathcal{F}} \in \ell_m^p$. Since $\bar{t}_\nu \leq \sqrt{R}\|t_\nu\|_V$ for all $\nu \in \mathcal{F}$, it follows that $(\bar{t}_\nu)_{\nu \in \mathcal{F}} \in \ell_m^p(\mathcal{F})$ as stated.

To prove (4.1), we first control the cardinality of the updated set $\Lambda_{n+1} \cap \mathcal{M}_n$. Recall that Λ_{n+1} was obtained by a iterative enrichment of Λ_n according to the largest $\bar{\mathbf{t}}_\nu(\mathcal{M}_n)$. Let $\bar{\Lambda}_{n+1} \subset \Lambda_{n+1}$ be the set corresponding to the last step of this enrichment before Λ_{n+1} has been reached. We define $\mathcal{S}_n := \bar{\Lambda}_{n+1} \cap \mathcal{M}_n$, so that

$$\#\Lambda_{n+1} \cap \mathcal{M}_n = 1 + \#\mathcal{S}_n,$$

and

$$e(\mathcal{S}_n) < \theta e(\mathcal{M}_n).$$

We thus have

$$\begin{aligned} (1 - \theta)e(\mathcal{M}_n) &\leq e(\mathcal{M}_n) - e(\mathcal{S}_n) \\ &= \sum_{\nu \in \mathcal{M}_n \setminus \mathcal{S}_n} |\bar{t}_\nu|^2 \\ &\leq \sum_{\nu \in \mathcal{M}_n \setminus \mathcal{S}_n} |\bar{\mathbf{t}}_\nu(\mathcal{M}_n)|^2 \\ &\leq \|(\bar{t}_\nu)_{\nu \in \mathcal{F}}\|_{\ell_m^p(\mathcal{F})}^2 (\#\mathcal{S}_n + 1)^{-2t}, \end{aligned}$$

where for the last line we have used (1.5) with $q = 2$, applied to the sequence $(\bar{\mathbf{t}}_\nu(\mathcal{M}_n))_{\nu \in \mathcal{M}_n}$ and the fact that $\bar{\mathbf{t}}_\nu(\mathcal{M}_n) \leq \bar{\mathbf{t}}_\nu$ where $(\bar{\mathbf{t}}_\nu)_{\nu \in \mathcal{F}}$ is the monotone envelope of $(\bar{t}_\nu)_{\nu \in \mathcal{F}}$. It follows that

$$\#\Lambda_{n+1} \cap \mathcal{M}_n \leq (1 - \theta)^{-1/2t} \|(\bar{t}_\nu)_{\nu \in \mathcal{F}}\|_{\ell_m^p(\mathcal{F})}^{1/t} e(\mathcal{M}_n)^{-1/2t},$$

and therefore, using (3.11),

$$\#\Lambda_{n+1} \cap \mathcal{M}_n \leq [(1 - \delta)(1 - \theta)]^{-1/2t} \|(\bar{t}_\nu)_{\nu \in \mathcal{F}}\|_{\ell_m^p(\mathcal{F})}^{1/t} \sigma(\Lambda_n)^{-1/2t}. \quad (4.2)$$

We next observe that $\sigma(\Lambda_n)$ decreases geometrically. Indeed, we can write

$$\sigma(\Lambda_{n+1}) = \sigma(\Lambda_n) - e(\Lambda_{n+1} \cap \mathcal{M}_n) \leq \sigma(\Lambda_n) - \theta e(\mathcal{M}_n).$$

Since by (3.11), we have $e(\mathcal{M}_n) \geq (1 - \delta)\sigma(\Lambda_n)$, we obtain

$$\sigma(\Lambda_{n+1}) \leq \kappa \sigma(\Lambda_n), \quad (4.3)$$

with $\kappa := 1 - \theta(1 - \delta) = 1 - \frac{2\theta r}{R+r} < 1$. Iterating (4.3), we obtain

$$\sigma(\Lambda_k) \geq \kappa^{n-k} \sigma(\Lambda_n).$$

Combining this with (4.2), we may now estimate the global cardinality of $\#(\Lambda_n)$ by

$$\begin{aligned} \#(\Lambda_n) &\leq \#(\Lambda_0) + \sum_{k=0}^{n-1} \#(\Lambda_{k+1} \cap \mathcal{M}_k) \\ &\leq 1 + [(1 - \delta)(1 - \theta)]^{-1/2t} \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^{1/t} \sum_{k=0}^{n-1} \sigma(\Lambda_k)^{-1/2t} \\ &\leq 1 + [(1 - \delta)(1 - \theta)]^{-1/2t} \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^{1/t} \sigma(\Lambda_n)^{-1/2t} \sum_{k=0}^{n-1} \kappa^{(n-k)/2t} \\ &\leq 1 + C \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^{1/t} \sigma(\Lambda_n)^{-1/2t} \end{aligned}$$

where $C := \left(\frac{\kappa}{(1-\delta)(1-\theta)}\right)^{1/2t} (1 - \kappa^{1/2t})^{-1}$. This last inequality can be rewritten as

$$\sigma(\Lambda_n) \leq C^{2t} \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^2 (\#(\Lambda_n) - 1)^{-2t}.$$

If $\#(\Lambda_n) > 1$, we have established (4.1) with $C_1 := (2C)^{2t}$. If $\#(\Lambda_n) = 1$ then

$$\sigma(\Lambda_n) \leq \|(\bar{t}_\nu)\|_{\ell^2(\mathcal{F})}^2 \leq \|(\bar{t}_\nu)\|_{\ell^p(\mathcal{F})}^2 \leq \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^2,$$

and (4.1) also holds, with the constant C_1 only depending on r, R, θ and t . \square

As a corollary, we also obtain the optimal decay of the ℓ^1 tail of the \bar{t}_ν and therefore of the error between u and its partial Taylor sum.

Corollary 4.3. *Under the assumptions of Theorem 1.1, we have $(\bar{t}_\nu)_{\nu \in \mathcal{F}} \in \ell_m^p(\mathcal{F})$ and the sets Λ_n satisfy*

$$\sum_{\nu \in \mathcal{F} \setminus \Lambda_n} \bar{t}_\nu \leq C_2 \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})} (\#(\Lambda_n))^{-s}, \quad s := \frac{1}{p} - 1, \quad (4.4)$$

where $C_2 := 1 + \sqrt{C_1}$ with C_1 the constant in (4.1). Consequently we have

$$\sup_{y \in U} \|u(y) - \sum_{\nu \in \Lambda_n} t_\nu y^\nu\|_V \leq \sum_{\nu \notin \Lambda_n} \|t_\nu\|_V \leq \frac{1}{\sqrt{r}} \sum_{\nu \notin \Lambda_n} \bar{t}_\nu \leq \frac{C_1}{\sqrt{r}} \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})} (\#(\Lambda_n))^{-s}. \quad (4.5)$$

Proof. Let $m := \#(\Lambda_n) \geq 1$ and consider the set $\Lambda_m^* \subset \mathcal{F}$ corresponding to the m largest \bar{t}_ν , we have

$$\begin{aligned} \sum_{\nu \notin \Lambda_n} \bar{t}_\nu &= \sum_{\nu \notin \Lambda_m^*} \bar{t}_\nu + \sum_{\nu \in \Lambda_m^* \setminus \Lambda_n} \bar{t}_\nu - \sum_{\nu \in \Lambda_n \setminus \Lambda_m^*} \bar{t}_\nu \\ &\leq \sum_{\nu \notin \Lambda_m^*} \bar{t}_\nu + \sum_{\nu \in \Lambda_m^* \setminus \Lambda_n} \bar{t}_\nu \\ &\leq \|(\bar{t}_\nu)\|_{\ell^p(\mathcal{F})} (m+1)^{-s} + m^{1/2} e (\Lambda_m^* \setminus \Lambda_n)^{1/2} \\ &\leq (\|(\bar{t}_\nu)\|_{\ell^p(\mathcal{F})} + \sqrt{C_1} \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}) (\#(\Lambda_n))^{-s} \\ &\leq (1 + \sqrt{C_1}) \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})} (\#(\Lambda_n))^{-s}, \end{aligned}$$

where we have used both (1.5) and (4.1). This establishes (4.4), which implies (4.5) since $\|\cdot\|_V \leq \frac{1}{\sqrt{r}} \|\cdot\|_{\bar{a}}$. \square

5. A SECOND ALGORITHM

We now want to modify Algorithm 1 in order to restrict the computation of the t_ν to a finite subset of \mathcal{M}_n . In the modified algorithm, we set a target accuracy $\varepsilon > 0$ and design the procedure in such a way that the algorithm terminates when $\sigma(\Lambda_n) \leq \varepsilon$. As explained in the next section, an additional error occurs due to finite element spatial discretization, and therefore a relevant choice for ε is an estimated value of this additional error, such as given by standard residual-based finite element error estimator.

In order to restrict the margin \mathcal{M}_n to a finite subset, we introduce a procedure SPARSE that has the following properties: if Λ is a finite monotone set and \mathcal{M} is its infinite margin, and if $(\bar{t}_\nu)_{\nu \in \Lambda}$ are known, then for any $\eta > 0$,

$$\mathcal{N} := \text{SPARSE}(\Lambda, (\bar{t}_\nu)_{\nu \in \Lambda}, \eta),$$

is a finite subset of \mathcal{M} which is monotone in \mathcal{M} and such that $e(\mathcal{M} \setminus \mathcal{N}) \leq \eta$.

There are several possible concrete realizations of this procedure. Here is a simple one. We define

$$\bar{\psi}_j := \frac{\psi_j}{\bar{a}},$$

and choose an integer $J > 0$ large enough such that

$$\left\| \sum_{j>J} |\bar{\psi}_j| \right\|_{L^\infty(D)} \leq \left(\frac{\alpha e(\Lambda)}{1 - \alpha \gamma} \right)^{-1} \eta, \quad (5.1)$$

where α and γ are given by (3.4) and (3.7), and we define

$$\mathcal{N} := \text{SPARSE}(\Lambda, (\bar{t}_\nu)_{\nu \in \Lambda}, \eta) := \{\nu \in \mathcal{M} ; \nu - e_j \in \Lambda \Rightarrow j \leq J\}.$$

Clearly \mathcal{N} is finite with $\#\mathcal{N} \leq J\#\Lambda$.

Lemma 5.1. *With the above definition of \mathcal{N} , one has*

$$e(\mathcal{M} \setminus \mathcal{N}) = \sum_{\nu \in \mathcal{M} \setminus \mathcal{N}} \bar{t}_\nu^2 \leq \eta. \quad (5.2)$$

Proof. We proceed in a similar way to the Proof of Lemma 3.2, by first writing

$$e(\mathcal{M} \setminus \mathcal{N}) \leq \alpha \sum_{\nu \in \mathcal{M} \setminus \mathcal{N}} \left(\sum_{j \text{ s.t. } \nu_j \neq 0} d_{\nu - e_j, j} \right) \leq A + B, \quad (5.3)$$

where now

$$A := \alpha \sum_{\nu \in \mathcal{M} \setminus \mathcal{N}} \left(\sum_{j \text{ s.t. } \nu - e_j \in \mathcal{M} \setminus \mathcal{N}} d_{\nu - e_j, j} \right) = \alpha \sum_{\mu \in \mathcal{M} \setminus \mathcal{N}} \left(\sum_{j \text{ s.t. } \mu + e_j \in \mathcal{M} \setminus \mathcal{N}} d_{\mu, j} \right),$$

and

$$B := \alpha \sum_{\nu \in \mathcal{M} \setminus \mathcal{N}} \left(\sum_{j \text{ s.t. } \nu - e_j \notin \mathcal{M} \setminus \mathcal{N}} d_{\nu - e_j, j} \right) = \alpha \sum_{\mu \in \Lambda \cup \mathcal{N}} \left(\sum_{j \text{ s.t. } \mu + e_j \in \mathcal{M} \setminus \mathcal{N}} d_{\mu, j} \right).$$

In this splitting, we have used the fact that if $\nu \in \mathcal{M} \setminus \mathcal{N}$ and $\nu_j \neq 0$, we have either $\nu - e_j \in \mathcal{M} \setminus \mathcal{N}$ or $\nu - e_j \in \Lambda \cup \mathcal{N}$. Using (3.5), we can bound A by

$$A \leq \alpha \gamma \sum_{\mu \in \mathcal{M} \setminus \mathcal{N}} \bar{t}_\mu^2 = \alpha \gamma e(\mathcal{M} \setminus \mathcal{N}).$$

To bound B , we first claim that for any $\mu \in \Lambda \cup \mathcal{N}$ such that $\mu + e_j \in \mathcal{M} \setminus \mathcal{N}$, we must have $\mu \in \Lambda$ and $j > J$. Indeed, since $\mu + e_j \in \mathcal{M} \setminus \mathcal{N}$, the definition of \mathcal{N} guarantees that $\mu + e_j = \tilde{\nu} + e_k$ for some $\tilde{\nu} \in \Lambda$ and $k > J$. If $j = k$ we have our claim. If $j \neq k$ then necessarily $\tilde{\nu} - e_j \in \Lambda$ because of the monotonicity of Λ and therefore μ can be written as the sum of $\tilde{\nu} - e_j \in \Lambda$ and e_k , which means that μ is not in \mathcal{N} . Thus, we have verified our claim. From the claim, it follows that the only j 's that may contribute in the summation inside B are such that $j > J$ and $\nu - e_j \in \Lambda$. Hence,

$$\begin{aligned} B &\leq \alpha \sum_{\mu \in \Lambda} \sum_{j > J} d_{\mu,j} \\ &= \alpha \sum_{\mu \in \Lambda_D} \int \left(\sum_{j > J} |\psi_j| \right) |\nabla t_\nu|^2 \\ &= \alpha \sum_{\mu \in \Lambda_D} \int \left(\sum_{j > J} |\bar{\psi}_j| \right) \bar{a} |\nabla t_\nu|^2 \\ &\leq \alpha \left\| \sum_{j > J} |\bar{\psi}_j| \right\|_{L^\infty} e(\Lambda) \leq (1 - \alpha\gamma) \eta. \end{aligned}$$

Combining the bounds for A and B with (5.3), we obtain

$$e(\mathcal{M} \setminus \mathcal{N}) \leq \frac{B}{1 - \alpha\gamma} \leq \eta,$$

as desired. \square

For a fixed $0 < \theta < 1$ and target accuracy $\varepsilon > 0$, we now consider the following algorithm:

Algorithm 2. Define $\Lambda_0 := \{0\}$, compute $t_0 := u(0)$ and set $\eta_0 := \bar{t}_0 = \|t_0\|_{\bar{a}}$; for $n = 0, 1, \dots$

- Given Λ_n and $(t_\nu)_{\nu \in \Lambda_n}$, define $\mathcal{M}_n := \mathcal{M}(\Lambda_n)$.
- For $j = 0, 1, \dots$
 - Define $\eta_{n,j} := 2^{-j} \eta_n$ and $\mathcal{M}_{n,j} := \text{SPARSE}(\Lambda_n, (\bar{t}_\nu)_{\nu \in \Lambda_n}, \eta_{n,j})$;
 - compute \bar{t}_ν for $\nu \in \mathcal{M}_{n,j}$ and compute $e(\mathcal{M}_{n,j})$;
 - if $e(\mathcal{M}_{n,j}) + \eta_{n,j} \leq (1 - \delta)\varepsilon$, with δ as in (3.12), then terminate the Algorithm and output the set $\Lambda(\varepsilon) := \Lambda_n$;
 - else if $e(\mathcal{M}_{n,j}) < \frac{4-2\theta}{1-\theta} \eta_{n,j}$, then go directly to step $j + 1$;
 - else if $e(\mathcal{M}_{n,j}) \geq \frac{4-2\theta}{1-\theta} \eta_{n,j}$, then terminate the inner loop in j , and define $\eta_{n+1} := \eta_{n,j}$. Compute t_ν for $\nu \in \mathcal{M}_{n,j}$ using (3.1) and define

$$\bar{\mathbf{t}}_\nu(\mathcal{M}_{\mathbf{n},j}) := \sup\{\bar{\mathbf{t}}_\mu : \mu \geq \nu \text{ and } \mu \in \mathcal{M}_{\mathbf{n},j}\}.$$

Enrich Λ_n by the indices $\nu \in \mathcal{M}_{n,j}$ corresponding to the k largest $\bar{\mathbf{t}}_\nu(\mathcal{M}_{\mathbf{n},j})$ for $k = 1, 2, \dots$, until the obtained monotone set $\Lambda_{n+1} \subset \Lambda_n \cup \mathcal{M}_{n,j}$ satisfies

$$e(\Lambda_{n+1} \cap \mathcal{M}_{n,j}) \geq \theta e(\mathcal{M}_{n,j}) + (2 - \theta) \eta_{n,j}.$$

- Go to step $n + 1$.

Theorem 5.2. Under the assumptions of Theorem 1.1 with $p \leq 1$, Algorithm 2 terminates for a finite value n^* and outputs a set $\Lambda := \Lambda(\varepsilon) := \Lambda_{n^*}$ which satisfies

$$\sigma(\Lambda) \leq \varepsilon. \tag{5.4}$$

Moreover, one has

$$\sigma(\Lambda) \leq \bar{C}_1 \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^2 (\#\Lambda)^{-2t}, \quad t := \frac{1}{p} - \frac{1}{2}, \quad (5.5)$$

and

$$\sup_{y \in U} \|u(y) - \sum_{\nu \in \Lambda} t_\nu y^\nu\|_V \leq \bar{C}_2 \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})} (\#\Lambda)^{-s}, \quad s := \frac{1}{p} - 1, \quad (5.6)$$

where the constants \bar{C}_1 and $\bar{C}_2 := \frac{1+\sqrt{\bar{C}_1}}{\sqrt{r}}$ only depend on (r, R, θ, t) .

Proof. We first claim that for each n the inner j loop terminates for some j . To see this, we note that this loop advances j only when $e(\mathcal{M}_{n,j}) < \frac{4-2\theta}{1-\theta}\eta_{n,j}$ and $e(\mathcal{M}_{n,j}) > (1-\delta)\varepsilon - \eta_{n,j}$. This cannot happen when j is large because the $\eta_{n,j}$ tend to zero with increasing j .

Next, we need to check that for the output j of the inner loop, we are able to determine the index set Λ_{n+1} satisfying the bulk criterion. To see this, we note that for the output j of the inner loop, we must have $e(\mathcal{M}_{n,j}) > \frac{4-2\theta}{1-\theta}\eta_{n,j}$ and therefore

$$(1-\theta)e(\mathcal{M}_{n,j}) - (2-\theta)\eta_{n,j} > (2-\theta)\eta_{n,j} > 0.$$

or equivalently

$$e(\mathcal{M}_{n,j}) > \theta e(\mathcal{M}_{n,j}) + (2-\theta)\eta_{n,j}.$$

This shows that Λ_{n+1} exists since the choice $\Lambda_n \cup \mathcal{M}_{n,j}$ satisfies the bulk criterion. Since $e(\mathcal{M}_n \setminus \mathcal{M}_{n,j}) \leq \eta_{n,j}$, it follows that

$$e(\Lambda_{n+1} \cap \mathcal{M}_{n,j}) \geq \theta e(\mathcal{M}_{n,j}) + (2-\theta)\eta_{n,j} \geq \theta e(\mathcal{M}_n) + (2-2\theta)\eta_{n,j},$$

and thus since $\theta < 1$

$$e(\Lambda_{n+1} \cap \mathcal{M}_{n,j}) \geq \theta e(\mathcal{M}_n). \quad (5.7)$$

Next we have to check that the outer n loop terminates. To see this we note that from (5.7) and the same reasoning as in the Proof of Theorem 4.2, the contraction property

$$\sigma(\Lambda_{n+1}) \leq \kappa \sigma(\Lambda_n), \quad (5.8)$$

holds whenever the set Λ_{n+1} is created by the algorithm and $0 < \kappa < 1$ is the constant of (4.3). Obviously, this contraction property shows that for any given $c > 0$, if the algorithm did not terminate for n large enough we must have $e(\mathcal{M}_n) \leq c$ because $e(\mathcal{M}_n) \leq \sigma(\Lambda_n)$. Here we take $c = \frac{(4-2\theta)(1-\delta)\varepsilon}{5-3\theta}$. If such a value of n has been reached by the algorithm and if Λ_n is not selected during the inner loop, this means that for the last value of j in this inner loop, we have

$$\eta_{n,j} \leq \frac{1-\theta}{4-2\theta} e(\mathcal{M}_{n,j}),$$

and thus

$$e(\mathcal{M}_{n,j}) + \eta_{n,j} \leq \frac{5-3\theta}{4-2\theta} e(\mathcal{M}_{n,j}) \leq \frac{5-3\theta}{4-2\theta} e(\mathcal{M}_n) \leq (1-\delta)\varepsilon.$$

Therefore, the algorithm terminates.

Let n^* be the terminal value of n and let j^* be the terminal value of j for the inner j loop applied to this value $n = n^*$. We have $e(\mathcal{M}_{n^*,j^*}) + \eta_{n^*,j^*} \leq (1-\delta)\varepsilon$ and thus

$$\sigma(\Lambda_{n^*}) \leq (1-\delta)^{-1} e(\mathcal{M}_{n^*}) \leq (1-\delta)^{-1} (e(\mathcal{M}_{n^*,j^*}) + \eta_{n^*,j^*}) \leq \varepsilon.$$

Therefore, (5.4) holds for the generated set $\Lambda = \Lambda_{n^*}$.

We shall next prove (5.5) by using an argument similar to that used in the Proof of Theorem 4.2. For any $n < n^*$, we bound the cardinality of the update set $\Lambda_{n+1} \cap \mathcal{M}_{n,j}$ where $j = j(n)$ is the terminal value of j for this n . By the same argument as in the Proof of Theorem 4.2, we obtain

$$(1 - \theta)e(\mathcal{M}_{n,j}) - (2 - \theta)\eta_{n,j} \leq \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^2 \#(\Lambda_{n+1} \cap \mathcal{M}_{n,j})^{-2t}.$$

We next observe that

$$\begin{aligned} (1 - \theta)e(\mathcal{M}_{n,j}) - (2 - \theta)\eta_{n,j} &= \frac{1}{3}(1 - \theta)e(\mathcal{M}_{n,j}) + \frac{2}{3}(1 - \theta)e(\mathcal{M}_{n,j}) - (2 - \theta)\eta_{n,j} \\ &\geq \frac{1}{3}(1 - \theta)e(\mathcal{M}_{n,j}) + \frac{2}{3}(4 - 2\theta)\eta_{n,j} - (2 - \theta)\eta_{n,j} \\ &= \frac{1}{3}(1 - \theta)e(\mathcal{M}_{n,j}) + \frac{1}{3}(2 - \theta)\eta_{n,j} \\ &\geq \frac{1 - \theta}{3}(e(\mathcal{M}_{n,j}) + \eta_{n,j}) \\ &\geq \frac{1 - \theta}{3}e(\mathcal{M}_n) \geq (1 - \delta) \left(\frac{1 - \theta}{3}\right) \sigma(\Lambda_n) =: \bar{C}_2^{-1} \sigma(\Lambda_n). \end{aligned}$$

Here the last inequality follows from Lemma 3.2 and the next to last inequality follows from the fact that $e(\mathcal{M}_n \setminus \mathcal{M}_{n,j}) \leq \eta_{n,j}$ by virtue of (5.2). Therefore

$$\begin{aligned} \#(\Lambda) &= 1 + \sum_{n=0}^{n^*-1} \#(\Lambda_{n+1} \cap \mathcal{M}_{n,j(n)}) \\ &\leq 1 + \bar{C}_2^{1/t} \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^{1/t} \sum_{n=0}^{n^*-1} \sigma(\Lambda_n)^{-1/2t} \\ &\leq 1 + \bar{C}_2^{1/t} \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^{1/t} \sigma(\Lambda)^{-1/2t} \sum_{k=0}^{n^*-1} \kappa^{(n^*-k)/2t} \\ &\leq 1 + C \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^{1/t} \sigma(\Lambda)^{-1/2t}, \end{aligned}$$

where $C := \bar{C}_2^t (1 - \kappa^{1/2t})^{-1}$. This last inequality can be rewritten as

$$\sigma(\Lambda) \leq \bar{C}_1 \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^2 (\#(\Lambda))^{-2t},$$

in the same way we argued to complete the Proof of Theorem 4.2. Hence, we have proved (5.5). Finally, (5.6) follows from (5.5) in the same way we have derived Corollary 4.3 from Theorem 4.2. \square

Although Algorithm 2 algorithm meets the benchmark of the optimal rate (1.7) under the minimal assumptions of Theorem 1.1, a closer inspection shows that it is not completely optimal from a computational point of view. Indeed, consider the number $B = B(\varepsilon) = B_{n^*}$ of boundary value problems which have actually been solved in order to compute the functions t_ν for ν in the final set $\Lambda = \Lambda(\varepsilon) = \Lambda_{n^*}$. Ideally, we would hope that this number is not much larger than the cardinality of Λ , so that we may actually retrieve the convergence estimates (5.5) and (5.6) in terms of B instead of $\#(\Lambda)$.

However, the number B involves the size of the restricted margin which is produced by the procedure SPARSE, and which might in principle be substantially larger than the set that is finally selected by the bulk search. Retrieving the same convergence rate in terms of B would actually require that when the accuracy η prescribed in SPARSE is of the same order as the current accuracy $e(\Lambda)$, then the cardinality of the produced set \mathcal{N} should be bounded by the optimal rate

$$\#(\mathcal{N}) \leq C \|(\bar{t}_\nu)\|_{\ell_m^p(\mathcal{F})}^{1/t} \eta^{-1/2t}. \tag{5.9}$$

A brief inspection seems to indicate that only a lower rate is achieved by our SPARSE procedure: on the one hand we know that

$$\#(\mathcal{N}) \leq J\#(A),$$

and that the set A has its cardinality optimally controlled by $\eta^{-1/2t}$, and on the other hand the number J that ensures (5.1) is of the order $\eta^{-1/s}$ where $s = \frac{1}{p} - 1 = t - \frac{1}{2}$. Therefore $\eta^{-1/2t}$ in (5.9) is *a priori* replaced by the non-optimal rate $\eta^{-1/(2t^2-t)}$

In order to remedy this defect, one would need to design more elaborate realizations of SPARSE in order to obtain a set \mathcal{N} of smaller, hopefully optimal, cardinality. One option that could lead to such a SPARSE procedure would be to make use of the available *a priori bounds* on the $\|t_\nu\|_V$ such as such as (2.3) and (2.4) in order to control the energy outside of the set \mathcal{N} . We do not embark in this direction here. Another option for lowering the CPU cost, which appears to work quite well in practice yet without a complete theoretical justification, will be proposed in Section 7.

6. SPACE DISCRETIZATION

The boundary value problems that recursively give the coefficients t_ν cannot be solved exactly. Instead, we would use a Galerkin method in a finite dimensional space $V_h \subset V$, typically a finite element space although this is not crucial in the present analysis which would also apply to spectral or wavelet discretization. We shall show in this section that it is possible to choose the *same* space V_h to approximate *all* t_ν and still retain the performance of Algorithms 1 and 2.

For the purpose of simplicity, we consider here the situation where the same \hat{E} spatial discretization is used for all ν . Yet, the analysis in Section 8 of [9] reveals that substantial computational gain may be expected if the spatial discretization is allowed to vary with ν (typically, coarser discretizations should be used for the computation of smaller Taylor coefficients). The possibility of adaptively choosing the approximation space parameter h depending on ν should also be explored but requires a more involved analysis. A future objective is therefore to design a solution algorithm that adaptively monitors the spatial resolution as new coefficients are being computed.

We define the finite element approximation $u_h(y) \in V_h$ as the solution to

$$\int_D a(x, y) \nabla u_h(y) \nabla v_h = \int_D f v_h \quad \forall v_h \in V_h. \tag{6.1}$$

By assumption **UEA**(r, R), for any closed subspace $V_h \subset V$ the Finite Element approximation is uniquely defined and the analysis in [9, 10] and all results of the present paper also apply to the discretized problem.

In particular, for every $h > 0$, the finite element approximation $u_h(y) \in V_h$ can be represented as a convergent Taylor expansion about $y = 0$ *i.e.*, $u_h(y) = \sum_{\nu \in \mathcal{F}} t_{\nu,h} y^\nu$, where

$$t_{\nu,h} := \frac{1}{\nu!} \partial^\nu u_h(0) \in V_h.$$

Moreover, the $\|t_{\nu,h}\|_V$ can be estimated by the same bound $b_\nu := 2C_0 \inf_{\rho \in \mathcal{A}_r} \rho^{-\nu}$ as the $\|t_\nu\|_V$, leading to a result similar to Theorem 2.4.

Theorem 6.1. *Under the assumptions of Theorem 1.1, the sequence $(\|t_{\nu,h}\|_V)_{\nu \in \mathcal{F}}$ belongs to $\ell_m^p(\mathcal{F})$. Moreover $\|(\|t_{\nu,h}\|_V)\|_{\ell^p(\mathcal{F})}$ is bounded independent of h .*

The coefficients $t_{\nu,h}$ can be computed recursively by solving linear systems corresponding to the space-discretized boundary value problems

$$\int_D \bar{a} \nabla t_{\nu,h} \nabla v_h = - \sum_{j \text{ s.t. } \nu_j \neq 0} \int_D \psi_j \nabla t_{\nu - e_j, h} \nabla v_h \quad \forall v_h \in V_h. \tag{6.2}$$

For the approximate Taylor coefficients, we introduce once more their energies as $\bar{t}_{\nu,h} := \|t_{\nu,h}\|_{\bar{a}}$. We may define $e_h(\Lambda)$ and $\sigma_h(\Lambda)$ and apply Algorithms 1 or 2 in a similar way, by simply replacing t_ν and \bar{t}_ν by $t_{\nu,h}$ and $\bar{t}_{\nu,h}$. For these algorithms, we obtain convergence results by the exact same approach as without space discretization.

Theorem 6.2. *Under the assumptions of Theorem 1.1 with $p < 1$, the application of each of the Algorithms 1 or 2 in the space discretized setting yields a sequence of sets (Λ_n) such that*

$$\sigma_h(\Lambda_n) \leq C_1 \|(\bar{t}_{\nu,h})\|_{\ell_m^p(\mathcal{F})} (\#\Lambda_n)^{-2t}, \quad t := \frac{1}{p} - \frac{1}{2}, \tag{6.3}$$

and

$$\sum_{\nu \notin \Lambda_n} e_{\nu,h} \leq C_2 \|(\bar{t}_{\nu,h})\|_{\ell_m^p(\mathcal{F})} (\#\Lambda_n)^{-s}, \quad s := \frac{1}{p} - 1, \tag{6.4}$$

where C_1 and C_2 are as in the continuous setting (depending on r, R, θ and on t , but being independent of h).

Consequently we have

$$\sup_{y \in U} \|u_h(y) - \sum_{\nu \in \Lambda_n} t_{\nu,h} y^\nu\|_V \leq \sum_{\nu \notin \Lambda_n} \|t_{\nu,h}\|_V \leq \frac{C_1}{\sqrt{r}} \|(\bar{t}_{\nu,h})\|_{\ell_m^p(\mathcal{F})} (\#\Lambda_n)^{-s}. \tag{6.5}$$

We finally quantify the space discretization error. The well-known theory of finite elements tells us that the rate of convergence of $\|u(y) - u_h(y)\|_V$ in terms of the decay of h is controlled by the smoothness of u in the scale of the H^s Sobolev space and the order of the finite element spaces V_h which are employed. For example, when using Lagrange finite elements of order k , we have

$$\sup_{y \in U} \|u(y) - u_h(h)\|_V \leq C_3 h^r \sup_{y \in U} \|u(y)\|_{H^{1+r}(D)}, \tag{6.6}$$

for all $r \leq k$. This leads to the following result.

Corollary 6.3. *Under the assumptions of Theorem 1.1 and assuming that $\sup_{y \in U} |u(y)|_{H^{1+r}} < \infty$ and that we use Lagrange finite elements of order $k \geq r$, then applying Algorithms 1 or 2 in the space discretized setting, we obtain*

$$\sup_{y \in U} \|u(y) - \sum_{\nu \in \Lambda_n} t_{\nu,h} y^\nu\|_V \leq C_3 h^r \sup_{y \in U} \|u(y)\|_{H^{1+r}(D)} + \frac{C_1}{\sqrt{r}} \|(\bar{t}_{\nu,h})\|_{\ell_m^p(\mathcal{F})} (\#\Lambda_n)^{-s}. \tag{6.7}$$

The largest value of r for which $\sup_{y \in U} |u(y)|_{H^{1+r}} < \infty$ is determined by:

- the smoothness of the right hand side f ;
- the smoothness of the diffusion coefficient a ;
- the smoothness of the boundary of D .

As an example, for $f \in L^2(D)$ and coefficients $a(x, y)$ which satisfy

$$\sup_{y \in U} \|a(\cdot, y)\|_{W^{1,\infty}(D)} < \infty, \tag{6.8}$$

we note that (1.1) implies that the solution $u(y)$ satisfies the Poisson equation

$$-\Delta u(y) = \frac{1}{a} [f - \nabla a \cdot \nabla u(y)] \quad \text{in } D, \quad u(y)|_{\partial D} = 0. \tag{6.9}$$

Therefore, for every $y \in U$ the solution $u(y)$ belongs to the space

$$W = \{v \in V : \Delta v \in L^2(D)\}.$$

If the domain D is convex, it is well known $W = H^2(D) \cap H_0^1(D)$. For more general Lipschitz domains, it is also known that $W = H^{1+r}(D) \cap H_0^1(D)$ for some $\frac{1}{2} \leq r \leq 1$. We refer to [16] for a general treatment of elliptic problems on non-smooth domains.

In the numerical experiment that follow, we actually deal with coefficients $a(x, y)$ which are piecewise constant on a partition of $D = [0, 1]^2$ into fixed sub-squares independent of y . Such coefficients obviously do not satisfy (6.8), however regularity results are also known in this setting and give that the solution $u(y)$ belong to $H^{1+r}(D) \cap H_0^1(D)$ for some $0 < r < \frac{1}{2}$ that depends on the maximal contrast R/r , see for example [3].

7. NUMERICAL EXPERIMENTS

7.1. Numerical tests and strategies

In this section, we study the numerical performance of the algorithm proposed in this paper. In this algorithm, the choice of Λ_n is made adaptively and it is based on a bulk search procedure. In particular, we want to compare this choice of Λ_n , with non-adaptive choices. We also study some variants using other adaptive strategies.

The algorithms that we analyzed in the previous sections were formulated regardless of whether the dimension of y is finite or infinite. In the present numerical tests, we use a parameter vector $y = (y_j)_{j=1, \dots, d}$, of dimension d that ranges up to 255. More precisely, we consider the following problem on the unit square $D := [0, 1] \times [0, 1]$:

$$-\operatorname{div}(a \nabla u) = f \text{ in } D, \quad u = 0 \text{ on } \partial D,$$

where for illustration purposes we take $f(x) = f(x_1, x_2) := x_1 x_2$. As to the diffusion coefficient $a(x, y)$ and the choice of the ψ_j we consider two different test cases.

Test 1. Characteristic functions

We partition D into $d = 64$ squares D_j of equal shape and consider a diffusion coefficient that is piecewise constant on each subdomain:

$$a(x, y) = \bar{a}(x) + \sum_{j=1}^{64} y_j \psi_j(x), \quad \bar{a} = 1 \text{ and } \psi_j = \alpha_j \chi_{D_j}. \quad (7.1)$$

Since in this case the ψ_j have disjoint supports, the uniform ellipticity assumption simply means that the weights $\alpha_j = \|\psi_j\|_{L^\infty(D)}$ are all strictly less than 1. To study the consistency of the numerical results with our theory, we also need that the sequence α_j has some decay, since in the case of an infinite sequence we require that $(\|\psi_j\|_{L^\infty(D)})_{j \geq 1}$ is summable. In our numerical test we take in (7.1)

$$\alpha_j = \frac{0.9}{j^3}. \quad (7.2)$$

The uniform ellipticity assumption **UEA**(r, R) therefore holds with $r = 0.1$ and $R = 1.9$. This test is not physically realistic: it corresponds to a diffusion which is uncorrelated between the different subdomains and with variability that strongly differs between subdomains labelled by small and large values of j . Its main purpose is to compare adaptive strategies with a non-adaptive one based on the estimates (2.3) and (2.4) for $\|t_\nu\|_V$, since these estimates can be explicitly computed for this test case as we explain further.

Test 2. Haar wavelets

For the same domain $D = [0, 1]^2$, we consider the bidimensional, isotropic Haar wavelet basis

$$h_{l,k}^i(x) = h^i(2^l x - k), \quad l \in \mathbb{N}, \quad k = (k_1, k_2) \in \{0, \dots, 2^l - 1\}^2, \quad i = 1, 2, 3,$$

where the generating wavelets h^i are defined by $h^1(x_1, x_2) := \varphi(x_1)h(x_2)$, $h^2(x_1, x_2) := h(x_1)\varphi(x_2)$ and $h^3(x_1, x_2) := h(x_1)h(x_2)$, with $\varphi := \chi_{[0,1]}$ and $h := \chi_{[0,1/2[} - \chi_{[1/2,1[}$. Any function in $L^2(D)$ has a unique

expansion into the orthogonal basis composed of all the above wavelets together with χ_D . We refer to [6] for a detailed treatment of wavelets. We consider parametric diffusion coefficients of the form

$$a(x, y) := \bar{a}(x) + \sum_{l=0}^L \beta_l \sum_{i=1}^3 \sum_{k \in \{0, \dots, 2^l - 1\}^2} y_{l,k,i} h_{l,k}^i(x), \quad \bar{a} = 1, \quad (7.3)$$

where L is a fixed integer representing the finest scale level, $(\beta_l)_{l=0, \dots, L}$ a positive sequence and $y_{l,k,i}$ are the parametric variables ranging over $[-1, 1]$.

With the above normalization $\|h_{l,k}^i\|_{L^\infty(D)} = 1$, it is known that the wavelet coefficients of a Hölder-continuous function in $C^{0,\gamma}(\bar{D})$ with $0 < \gamma \leq 1$ decay like $\mathcal{O}(2^{-\gamma l})$ as the scale level l grows. The rate of decay of the sequence β_l therefore reflects the amount smoothness (or correlation in the stochastic context) in a . We consider the general form

$$\beta_l := c 2^{-\gamma l} \quad c := 0.3 \frac{1 - 2^{-\gamma}}{1 - 2^{-(L+1)\gamma}} = 0.3 \frac{2^\gamma - 1}{2^{2\gamma} - 2^{-L\gamma}}, \quad (7.4)$$

which ensures that the uniform ellipticity assumption **UEA**(r, R) holds with $r = 0.1$ and $R = 1.9$. In the numerical tests, we consider the two particular values

$$\gamma = 0.5 \quad \text{and} \quad \gamma = 3,$$

in order to compare the effect of low smoothness (short range correlation) and of high smoothness (long range correlation) on the behaviour of our algorithms. Using the relabelling

$$\psi_j := \beta_l h_{l,k}^i \quad \text{and} \quad y_j := y_{l,k,i}, \quad \text{when} \quad j = 2^{2l} + 3(2^l k_1 + k_2) + i - 1,$$

we may rewrite the above expansion (7.3) in the form $a(x, y) := \bar{a}(x) + \sum_{j=1}^d y_j \psi_j(x)$ adopted in this paper, with

$$d := 2^{2(L+1)} - 1.$$

In order to study the robustness of the method to the dimensionality, we consider different values $L = 1, 2, 3$ for the maximal scale, which corresponds to taking $d = 15, 63, 255$. Note that after this relabelling, the sequence $(\|\psi_j\|_{L^\infty(D)})_{j \geq 1}$ decays like $\mathcal{O}(j^{-\gamma/2})$, and therefore like $\mathcal{O}(j^{-1/4})$ and $\mathcal{O}(j^{-3/2})$ for $\gamma = 0.5$ and $\gamma = 3$ respectively.

As in Section 6, we use one fixed finite element space for the spatial discretization of all active Taylor coefficients. Therefore, for the different strategies of building the coefficient sets Λ_n , we actually study the decay of the Taylor expansion error *for the finite element solution*

$$\sup_{y \in U} \|u_h(y) - \sum_{\nu \in \Lambda_n} t_{\nu,h} y^\nu\|_V, \quad (7.5)$$

as $\#(\Lambda_n)$ grows, bearing in mind that the finite element discretization induces an additional discretization error $\sup_{y \in U} \|u(y) - u_h(y)\|_V$ which can be bounded according to (6.6).

For the generation of the sets Λ_n of “active” Taylor coefficients, we compare three non-adaptive strategies that are based on *a priori* choices of the sets Λ_n , and three adaptive strategies (in particular Algorithm 1) that exploit the results of earlier computations. For the sake of notational simplicity, we describe these algorithms without the additional finite element discretization, therefore using the notation t_ν instead of $t_{\nu,h}$. Their adaptation in the finite element setting is of course straightforward.

Non-adaptive strategies:

- **Algorithm QN:** for $n \geq 0$, we take $\Lambda_n = \{\nu \in \mathcal{F}, \text{ s.t. } \max(\nu_j) \leq n\}$. Therefore $\text{Span}\{y \mapsto y^\nu : \nu \in \Lambda_n\}$ is the space \mathbb{Q}_n of polynomials of degree at most n in each variable. The dimension of this space $\#\Lambda_n = (n+1)^d$ grows exponentially with the dimension d of y , reflecting the curse of dimensionality.

- **Algorithm PN:** for $n \geq 0$, we take $\Lambda_n = \{\nu \in \mathcal{F}, \text{ s.t } |\nu| \leq n\}$. Therefore $\text{Span}\{y \mapsto y^\nu : \nu \in \Lambda_n\}$ is the space \mathbb{P}_n of polynomials of total degree at most n in d variables. The dimension of this space $\#\Lambda_n = \binom{n+d}{n}$, although smaller than that of \mathbb{Q}_n by an order $d!$ still grows exponentially with d .
- **Algorithm LE (largest estimates):** as explained in the Proof of Theorem 2.4, there are available estimates such as (2.3) and (2.4) for $\|t_\nu\|_V$. It is therefore natural to choose for Λ_n the set of indices corresponding to the n largest of these estimates, for example the n largest

$$e_\nu := \frac{\|f\|_{V^*}}{\delta} \inf_{\rho \in \mathcal{A}_\delta} \rho^{-\nu},$$

for some given $0 < \delta < r$. As already explained, such sets are monotone by construction. In practice, it is not always simple to compute the exact value of the infimum in the above definition of e_ν . However, this problem has a simple solution in the case where the ψ_j 's have disjoint supports since all ρ_j can be optimized separately. For the diffusion coefficient considered in Test 1 this easily leads to the solution

$$\rho_j^* = \frac{1 - \delta}{\alpha_j},$$

and therefore

$$e_\nu := \frac{\|f\|_{V^*}}{\delta} \prod_{j \geq 1} \left(\frac{\alpha_j}{1 - \delta} \right)^{\nu_j}.$$

The set Λ_n may also be viewed as the set of those ν such that e_ν exceeds a certain threshold $t = t(n) > 0$ that decreases with n , and are therefore of the form

$$\Lambda_n := \left\{ \nu : \sum_{j \geq 1} a_j \nu_j \leq \Theta(n) \right\} \quad \text{with } a_j := -\log \left(\frac{\alpha_j}{1 - \delta} \right) \quad \text{and } \Theta(n) = -\log \left(\frac{t(n)\delta}{\|f\|_{V^*}} \right).$$

Note that if all α_j (and therefore a_j) were equal, then this would give the same *a priori* choice as the previously described Algorithm PN. In our case, the weight factors a_j decrease with j , resulting in some anisotropy in the sets Λ_n : higher polynomial degrees are expected for small values of j which represent the most “active” variables. A similar choice of polynomial space was studied in [23] for collocation methods.

In our numerical tests, we have used the value $\delta = \frac{r}{2} = 0.05$. One could use an even sharper *a priori* estimate on the $\|t_\nu\|_V$ by taking the infimum of e_ν also over all $\delta \in]0, r[$. This leads to a similar estimate but now with δ depending on ν according to $\delta = \min(r, \frac{1}{1+|\nu|})$.

In contrast, the exact value of e_ν is not easily computable for the diffusion coefficient considered in Test 2, and we need to rely on less sharp estimates for the $\|t_\nu\|_V$, based on specific choices for the sequence ρ as explained further.

Adaptive strategies:

- **Algorithm BS (bulk search):** this is simply Algorithm 1 based on the bulk search procedure as proposed in Section 4, and applied under the given finite element discretization. Since we work in finite dimension, it is in theory possible to apply this algorithm without the need to restrict the margin as it is done in Algorithm 2. In our numerical test, we have used the value $\theta = 0.2$ for the bulk parameter (we observed that the error curves are almost identical when θ ranges in $[0.05, 0.95]$).
- **Algorithm LN (largest neighbor):** although Algorithm 1 may be performed in the finite dimensional context, the size of the current margin \mathcal{M}_n relative to the size of the current set Λ_n becomes a source of computational slow-down as d grows. An alternate strategy is to only consider the *reduced margin*

$$\mathcal{I}_1(\Lambda_n) := \{\nu \notin \Lambda_n : \nu_j \neq 0 \Rightarrow \nu - e_j \in \Lambda_n\},$$

which are the indices in \mathcal{M}_n for which the Taylor coefficients t_ν can directly be computed from those indexed by A_n . We then define

$$A_{n+1} = A_n \cup \{\nu^*\},$$

where

$$\nu^* := \operatorname{Argmax}_{\nu \in \mathcal{I}_1(A_n)} \bar{t}_\nu.$$

The intuition for considering such a strategy is that if the sequence $(\bar{t}_\nu)_{\nu \in \mathcal{F}}$ were monotone, then this would select the \bar{t}_ν in decreasing order. The potential pay-off is that the reduced margin is much smaller than \mathcal{M}_n (in particular, it is easy to check that at most d boundary value problems need to be solved at each iteration). As we shall see, this strategy gives excellent results, although we have no proof similar to Algorithm 1 that it performs optimally in the sense of convergence rates.

- **Algorithm LNE (largest neighbor estimate):** in order to save further computational cost, we can use majorants of \bar{t}_ν in order to decide on the new set A_{n+1} . From (3.7), one straightforward upper estimate for \bar{t}_ν is

$$\bar{t}_\nu \leq N_\nu := \left(\alpha \sum_{j \text{ s.t. } \nu_j \neq 0} \|\bar{\psi}_j\|_{L^\infty(D)} \bar{t}_{\nu - e_j}^2 \right)^{\frac{1}{2}}. \tag{7.6}$$

One can then construct the new set A_{n+1} as in the previous Algorithm LN, by using N_ν instead of \bar{t}_ν . The saving comes from the fact that computing N_ν is much cheaper than computing t_ν .

7.2. Numerical results for Test 1

We have compared the various strategies using 4 choices of finite element spaces based on uniform triangulations of D obtained by splitting each element of a square mesh into two triangles: (i) 8×8 squares and \mathbb{P}_1 finite elements ($\dim(V_h) = 49$), (ii) 16×16 squares and \mathbb{P}_1 finite elements ($\dim(V_h) = 225$) (iii) 16×16 squares and \mathbb{P}_2 finite elements ($\dim(V_h) = 961$), (iv) 32×32 squares and \mathbb{P}_1 finite elements ($\dim(V_h) = 961$). We display in Figure 1 the error curves for the six strategies described above for the generation of the sets A_n . These error curves represent the supremum error (7.5) (estimated by taking the supremum over a random choice of 100 values of y) as a function of $\#(A_n)$. Note that for certain strategies, such as PN, QN and BS, the number $\#(A_n)$ does not grow by 1 at each iteration and therefore only takes a few integer values. In such cases, we obtain all intermediate values for the error curves by filling the intermediates indices in $A_{n+1} \setminus A_n$ by lexicographic order.

We also indicate for each choice of finite element space an estimate of the FE error $\sup_{y \in U} \|u(y) - u_h(y)\|_V$. This estimate is done by replacing $u(y)$ by a finite element solution on a very fine mesh obtained from 256×256 squares and taking the supremum over the same random choice of 100 values of y .

We record three major observations about the error curves.

- First, not much difference in the error curves is observed as we modify the spatial discretization, once it is finer than 8×8 . In fact, a closer inspection also show that the sets A_n selected by the adaptive algorithms change very little as we modify the spatial discretization. This suggest that the same sets and error curves would be obtained if there were no spatial discretization at all, *i.e.* if we were computing the t_ν by exactly solving the boundary value problems (3.1). In particular, the portion of the error curves which is below the value of the finite element error is still relevant to us, since this portion does not seem to change as this error is diminished.
- Second, we observe that the adaptive strategies BS and LN outperform all non adaptive strategies. They give almost identical error curves, which indicates that the LN strategy is preferable since it has lower computational cost. In contrast, a loss in performance is observed if we instead use LNE. As to the non-adaptive strategies, LE outperforms PN and QN which do not produce any anisotropy in the coefficient sets. It is interesting to note that with 100 coefficients, the Taylor approximation error of the adaptive strategies is dominated by the finite element error, while it is still above it with 10^4 coefficients when using PN and QN.

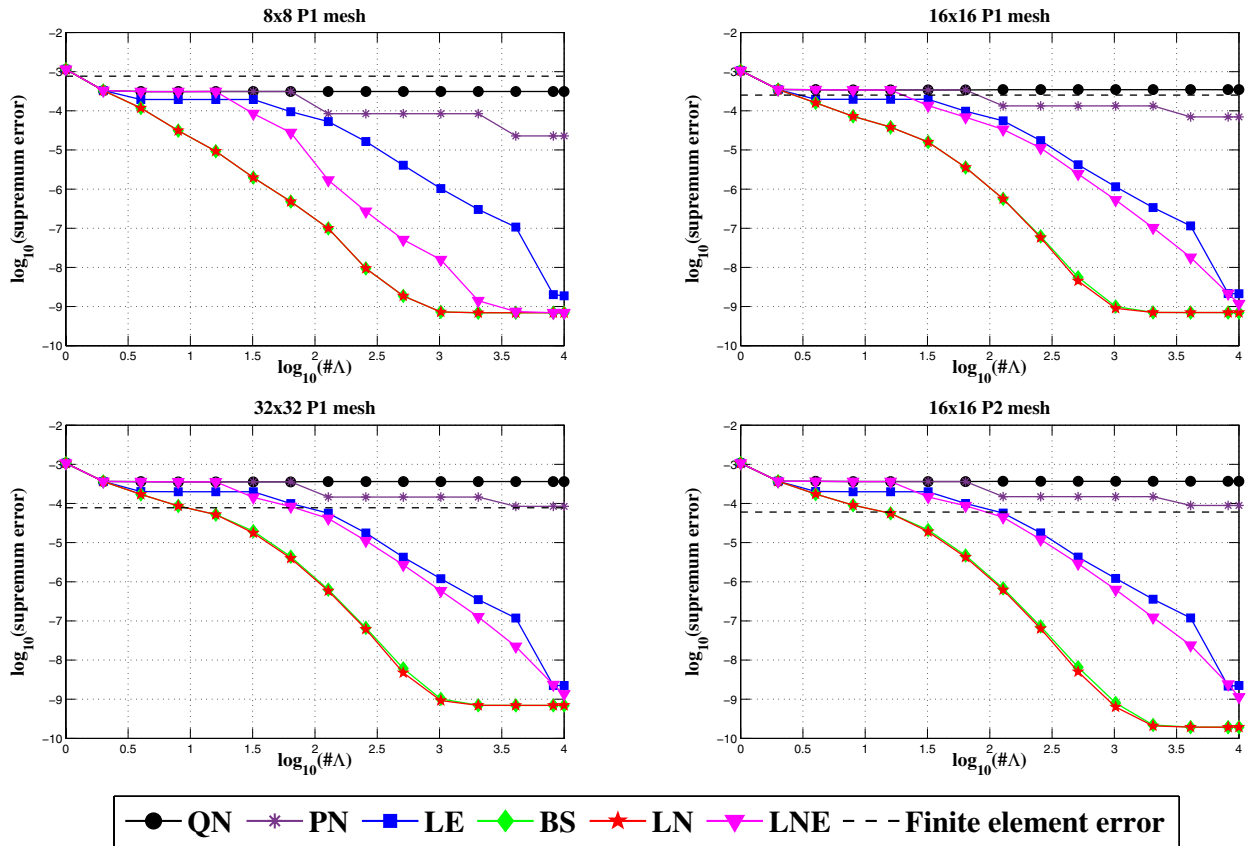


FIGURE 1. Comparison the different strategies for finite element spaces (i) (upper left), (ii) (upper right), (iii) (lower left) and (iv) (lower right).

- Finally, we observe a stagnation of order 10^{-9} in the supremum error. We interpret this by the fact that our algorithm computes once and for all the Taylor coefficients and that small numerical error resulting from linear system inversion accumulate in such computations. In turn the computed Taylor development converges towards a limit which slightly differs from $u_h(y)$.

In order to obtain a fair comparison between the different algorithms, we also show in Figure 2 their error curves in terms of the total number of boundary value problems which have been solved, and which is a better reflection of the CPU time (here we only consider the spatial discretization by 16×16 squares \mathbb{P}_1 finite elements). For non-adaptive strategies and for LNE, this number is the same as $\#(\Lambda)$, but it exceeds it moderately for LN and more strongly for BS. In this new comparison, we observe that the algorithm LN gives the best performance, followed by LNE and LE.

Remark 7.1. Since we have observed that the error curves and selected adaptive sets do not depend much on the finite element space discretization, an interesting perspective for gaining CPU time is to first use a coarse grid finite element space to find the adaptive coefficients sets Λ_n . One may then use a finer grid for the computation of the coefficients in such sets, therefore avoiding the overhead caused by solving more boundary value problems than $\#(\Lambda_n)$ with the fine discretization. We may also use the coarse grid error curves to estimate the number of Taylor coefficients that we need to compute with the fine discretization in order to reach a prescribed accuracy.




FIGURE 2. Comparison of the different strategies in term of total number of solved bvp.

Remark 7.2. Our analysis shows that we can set a stopping criterion for our adaptive algorithm based on the accuracy of the Taylor approximation to $u_h(y)$: the algorithm terminates at some step n such that

$$\sup_{y \in U} \|u_h(y) - \sum_{\nu \in \Lambda_n} t_{\nu,h} y^\nu\|_V \leq \varepsilon,$$

where $\varepsilon > 0$ is a prescribed tolerance. A natural choice is to choose ε of the same order as the finite element error

$$\sup_{y \in U} \|u_h(y) - u(y)\|_V.$$

While this last quantity is not exactly known to us, it can be bounded by above according to *a priori* estimate (6.6) based on our knowledge of the maximal Sobolev smoothness of $u(y)$, or estimated in a finer way based on *a posteriori* analysis.

Remark 7.3. In all three adaptive approaches, the specific choice of numbering coordinates y_j might influence the selection of the approximations once ties in certain quantities occur. In the present numerical experiments, the 64 coordinates were enumerated in lexicographic order according to the location of the support of the ψ_j in D . We performed the same experiments with several random reshufflings of the indexation (so that the most significant parameter y_j does not appear as first coordinate) which rendered indistinguishable results from the ones reported here; although this finding is, to some extent, implementation dependent, it strongly suggests that the presented algorithms will perform well also for more general parameter dependences, where the most significant coordinate appears only in high dimension.

We also have investigated the convergence of the mean value solution $\bar{u} = \mathbb{E}(u)$ when the y_j are i.i.d. random variables which are uniformly distributed in $[-1, 1]$. Given a Taylor approximation $u_\Lambda(y) := \sum_{\nu \in \Lambda} t_\nu y^\nu$ computed for a certain set Λ by one of the proposed strategies, this mean value may thus be approximated by

$$\bar{u}_\Lambda := \sum_{\nu \in \Lambda} t_\nu \mathbb{E}(y^\nu),$$




FIGURE 3. Comparison of the different strategies with Monte Carlo method.

with

$$\mathbb{E}(y^\nu) = \prod_{j \geq 1} \mathbb{E}(y_j^{\nu_j}) = \prod_{j \geq 1} \left(\int_{-1}^1 t^{\nu_j} \frac{dt}{2} \right) = \prod_{j \geq 1} \frac{1 + (-1)^{\nu_j}}{2 + 2\nu_j}.$$

We are ensured that the difference between the averages \bar{u} and \bar{u}_Λ does not exceed the supremum error in y between $u(y)$ and $u_\Lambda(y)$ which was previously estimated for the various methods. Since we do not know the exact value of \bar{u} for the computation of the error, we replace it by the value \bar{u}_Λ obtained with BS algorithm when $\#(\Lambda) = 10\,000$, which is thus accurate up to an error of order 10^{-10} . This allows us to make the comparison between performance of the various strategies for approximating \bar{u} by the error curves in terms of the number of coefficients. In addition we may compare this with the accuracy of the Monte-Carlo method, which consists in computing the empirical average

$$\bar{u}_n := \frac{1}{n} \sum_{i=1}^n u(y^i),$$

where y^1, \dots, y^n are independent random draws of the vector y . Since the MC method requires solving n boundary value problems, we compare its performance to the previous methods when the total number of solved boundary value problem is n , as n varies. The results are displayed in Figure 3. For the MC method, we display the average of the error curves for 6 independant realizations in order to illustrate the expected error $\mathbb{E}(\|\bar{u} - \bar{u}_n\|_V)$ rather than the error $\|\bar{u} - \bar{u}_n\|_V$ for a particular realization (which is more oscillatory). The $n^{-1/2}$ rate of decay of the MC method is clearly outperformed by the Taylor approximation methods based on the adaptive selection of Λ , which is rather striking in view of the large dimension $d = 64$. Note however, that in contrast to the Taylor approximation method, the MC approach allows us to solve all boundary value problems in parallel, however with a different stiffness matrix for each problem.

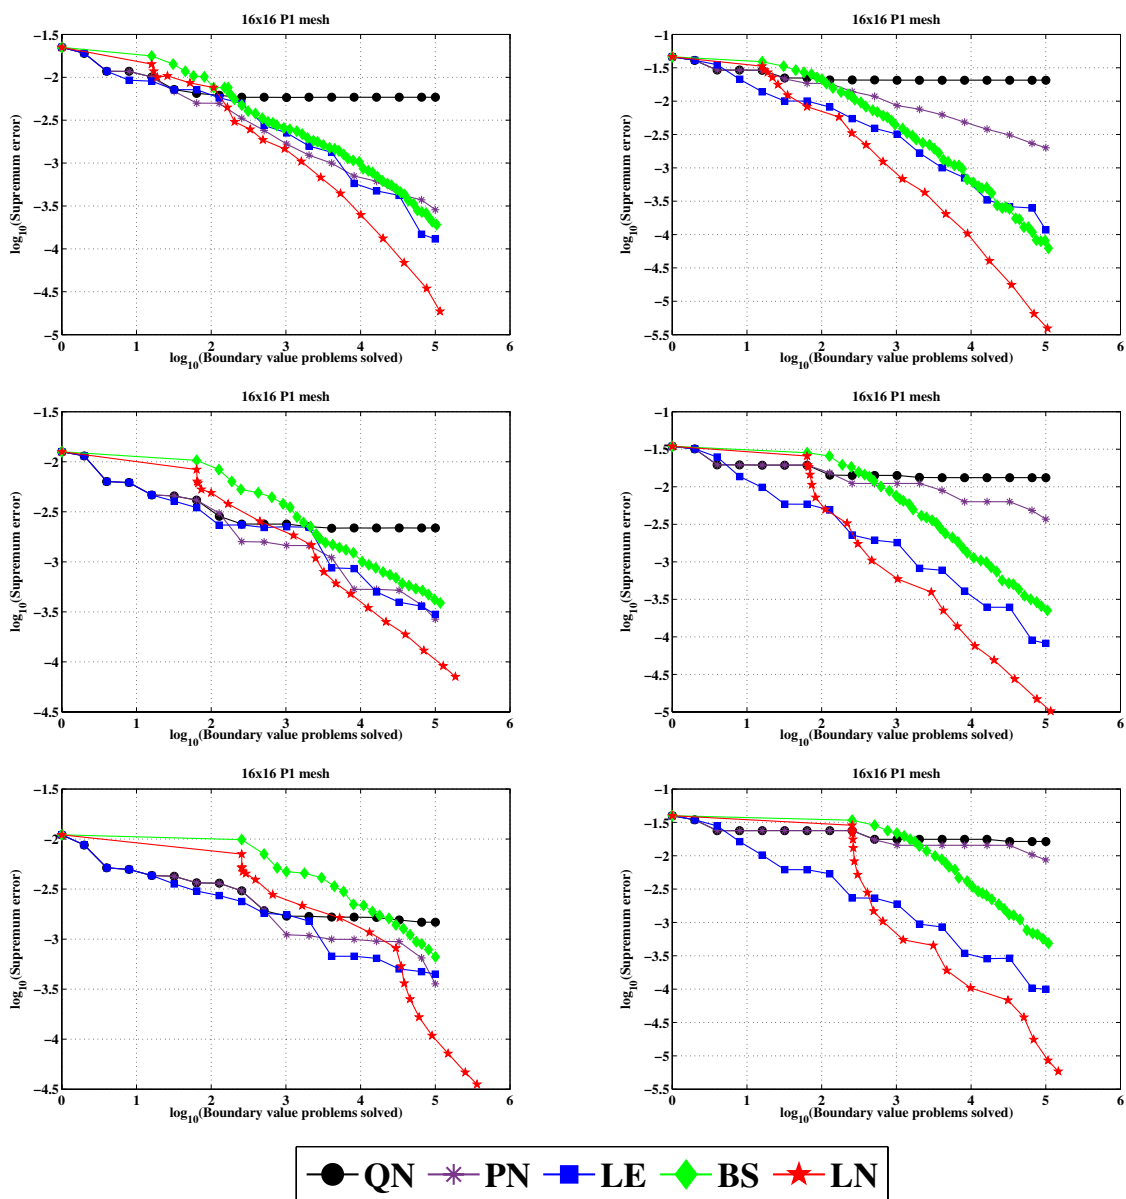


FIGURE 4. Comparison the different strategies for for $\gamma = 0.5$ (left) and $\gamma = 3$ (right), and different dimensionalities $d = 15$ (up), $d = 63$ (middle) and $d = 255$ (bottom).

7.3. Numerical results for Test 2

For this test, we only display the comparison between Algorithms PN, QN, LE, BS and LN, since we observed that the LNE algorithm does not perform as good as LN, similar to Test 1.

In contrast to Test 1, we cannot use the *a priori* bound

$$e_\nu := \frac{\|f\|_{V^*}}{\delta} \inf_{\rho \in \mathcal{A}_\delta} \rho^{-\nu},$$




FIGURE 5. Comparison the index sets Λ with $\#(\Lambda) = 10000$ projected on the components (1, 5) for Algorithms QN (upper left), PN (upper right), LE (lower left) and LN (lower right).

with $\delta = r/2 = 0.05$ to derive the choice of the active indices in Algorithm LE, since this optimization problem has no simply computable solution. We therefore rely on sub-optimal bounds of the form

$$e_\nu \leq \tilde{e}_\nu := \frac{\|f\|_{V^*}}{\delta} \rho^{-\nu},$$

where $\rho = \rho(\nu)$ is a particular sequence in \mathcal{A}_δ , which is chosen depending on ν , in contrast to Test 1. The best results were obtained with the following choice for $\rho(\nu)$:

$$\rho_j(\nu) = \frac{A(\nu)}{\|\psi_j\|_{L^\infty(D)}},$$

if $\nu_j \neq 0$ and $\rho_j(\nu) = 0$ if $\nu_j = 0$, where $A(\nu)$ is the largest positive number such that the result $\rho(\nu)$ is in \mathcal{A}_δ . This number is thus defined in such a way that

$$\left\| \sum_{j \text{ s.t. } \nu_j \neq 0} \rho_j(\nu) |\psi_j| \right\|_{L^\infty(D)} = 1 - \delta = 0.95.$$

The sets Λ_n chosen in Algorithm LE now correspond to the n largest \tilde{e}_ν .

Figure 4 displays the error curves in terms of the number of solved boundary value problems, similar to Figure 2 for Test 1, for the two values $\gamma = 0.5$ and $\gamma = 3$, and for different maximal scale level $L = 1, 2, 3$, corresponding to dimension $d = 15, 63, 255$. We record several observations about the error curves:

- In the low smoothness/short range correlation case $\gamma = 0.5$, Algorithms BS, LE and LN do not perform significantly better than Algorithm PN which corresponds to the standard choice of polynomials of fixed total degree. This can be explained by the fact that, in this case, all $\|\psi_j\|_{L^\infty(D)}$ have roughly the same range of magnitude so that all variables y_j are equally important. In turn, in this example the active index sets selected by BS, LE and LN are not highly anisotropic, and perform similar than those of PN. In contrast,

the high smoothness/long range correlation case $\gamma = 3$ highly benefits from an anisotropic selection (higher degree tends to be allocated to the variables associated to coarse scale wavelets), and in turn Algorithms BS, LE and LN significantly outperform PN and QN.

- In the low smoothness/correlation case $\gamma = 0.5$, all algorithms are subject to the curse of dimensionality in the sense that the error curves deteriorate as d increases. This includes the adaptive algorithms, which is not in contradiction with our theoretical results. Indeed, this value of γ corresponds to a decay in $\mathcal{O}(j^{-1/4})$ for the sequence $(\|\psi_j\|_{L^\infty(D)})_{j \geq 1}$, which is therefore not ℓ^p -summable for any exponent $p < 1$. In contrast, the high smoothness/ long range correlation case $\gamma = 3$ is not subject to the curse of dimensionality when Algorithms LE and LN are being used. Note that this value of γ corresponds to a decay in $\mathcal{O}(j^{-3/2})$ for the sequence $(\|\psi_j\|_{L^\infty(D)})_{j \geq 1}$, which is therefore ℓ^p summable for $p > 2/3$.
- For both values of γ , Algorithms LN and LE gives the best performances, and Algorithm BS is subject to the curse of dimensionality due to the cost of solving boundary value problems for all indices in the margin \mathcal{M}_n .
- The error curves for Algorithm LN and BS start decreasing only after a certain number of boundary value problems has been solved. This is simply due to the fact that at the very first step, d boundary value problems have to be solved, corresponding to the cardinality of the margin \mathcal{M}_0 of $\Lambda_0 = \{0\}$ (which at this stage is the same as the reduced margin).
- Generally speaking, Algorithm LN seems to be the one that performs best.

In order to have an idea of the geometry of the coefficients sets Λ produced by the different strategies, consider the projection on two variables $j = 1, 5$, associated to the first and second scale level $l = 0, 1$, *i.e.* the sets

$$\{(\nu_1, \nu_5) : \nu \in \Lambda\}.$$

We compare these sets in Figure 5, when $\#(\Lambda) = 10\,000$ for the strategies QN, PN, LE and LN in the case $\gamma = 3$ and $d = 63$. As expected, the sets obtained for the non-adaptive choices QN and PN do not reach a high degree due to the curse of dimensionality: when $d = 63$ the dimension of the spaces \mathbb{Q}_1 and \mathbb{P}_3 clearly exceeds 10 000 and therefore no degree higher than 1 and 3 can be reached for any variable when using these two methods respectively. In contrast, the adaptive strategies capture the anisotropic feature of the problem and reach a higher polynomial degrees in the most active variable y_1 . We did not plot the sets generated by BS which are quite similar to LN. Note however that their geometry differs from that of the set generated by LE based on the *a priori* estimates \tilde{e}_ν .

8. CONCLUSION

In this paper, we have introduced an adaptive algorithm for sparse Taylor approximation of parametric and stochastic PDEs. This algorithm has several remarkable features, in particular:

- (i) It builds the polynomial expansion by a solving a sequence of boundary value problems (3.1) which all have the same stiffness matrix.
- (ii) From a theoretical point of view, its convergence with respect to the polynomial dimension can be proved to be near optimal.

This second property is reflected in the numerical tests. In contrast to other approaches for selecting the active index sets, the algorithm does not uses any information based on *a priori* analysis, and yet performs at least as good as when using such *a priori* choices.

It is worth mentioning that this computational approach can be applied verbatim to other models, such as a parabolic equation of the form

$$\partial_t u - \operatorname{div}(a \nabla u) = f, \quad t \in [0, T], \quad x \in D, \quad u(x, 0) = u_0(x), \quad u(x, t) = 0, \quad x \in \partial D,$$

with a of the same parametric form as in this paper and the functions $\psi_j = \psi_j(x, t)$ satisfying a similar condition as (UEA). In that case, the solution space is $V = L^2([0, T], H_0^1(D))$. However, let us stress that our approach is strongly tied to the affine structure of a with respect to the parameter vector y , in contrast to other methods such as collocation.

REFERENCES

- [1] I. Babuska, R. Tempone and G.E. Zouraris, Galerkin finite element approximations of stochastic elliptic partial differential equations. *SIAM J. Numer. Anal.* **42** (2004) 800–825.
- [2] I. Babuška, F. Nobile and R. Tempone, A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM J. Numer. Anal.* **45** (2007) 1005–1034.
- [3] C. Bernardi and R. Verfürth, Adaptive finite element methods for elliptic equations with non-smooth coefficients. *Numer. Math.* **85** (2000) 579–608.
- [4] P. Binev, W. Dahmen, and R. DeVore, Adaptive finite element methods with convergence rates. *Numer. Math.* **97** (2004) 219–268.
- [5] A. Buffa, Y. Maday, A.T. Patera, C. Prudhomme and G. Turinici, *A priori* convergence of the greedy algorithm for the parameterized reduced basis. *ESAIM: M2AN* **3** (2012) 595–603.
- [6] A. Cohen, *Numerical analysis of wavelet methods*. Elsevier, Amsterdam (2003).
- [7] A. Cohen, W. Dahmen and R. DeVore, Adaptive wavelet methods for elliptic operator equations – Convergence rates. *Math. Comput.* **70** (2000) 27–75.
- [8] A. Cohen, W. Dahmen and R. DeVore, Adaptive wavelet methods for operator equations – Beyond the elliptic case. *J. FoCM* **2** (2002) 203–245.
- [9] A. Cohen, R. DeVore and C. Schwab, Convergence rates of best N -term Galerkin approximations for a class of elliptic sPDEs. *Found. Comput. Math.* **10** (2010) 615–646.
- [10] A. Cohen, R. DeVore and C. Schwab, Analytic regularity and polynomial approximation of parametric and stochastic PDE's. *Anal. Appl.* **9** (2011) 11–47.
- [11] R. DeVore, Nonlinear approximation. *Acta Numer.* **7** (1998) 51–150.
- [12] W. Dörfler, A convergent adaptive algorithm for Poisson's equation. *SIAM J. Numer. Anal.* **33** (1996) 1106–1124.
- [13] Ph. Frauenfelder, Ch. Schwab and R.A. Todor, Finite elements for elliptic problems with stochastic coefficients. *Comput. Methods Appl. Mech. Eng.* **194** (2005) 205–228.
- [14] T. Gantumur, H. Harbrecht and R. Stevenson, An optimal adaptive wavelet method without coarsening of the iterands. *Math. Comput.* **76** (2007) 615–629.
- [15] R. Ghanem and P. Spanos, Spectral techniques for stochastic finite elements. *Arch. Comput. Methods Eng.* **4** (1997) 63–100.
- [16] P. Grisvard, *Elliptic problems on non-smooth domains*. Pitman (1983).
- [17] V.H. Hoang and Ch. Schwab, *Sparse tensor Galerkin discretizations for parametric and random parabolic PDEs I: Analytic regularity and gpc-approximation*. Report 2010-11, Seminar for Applied Mathematics, ETH Zürich (in review).
- [18] V.H. Hoang and Ch. Schwab, *Analytic regularity and gpc approximation for parametric and random 2nd order hyperbolic PDEs*. Report 2010-19, Seminar for Applied Mathematics, ETH Zürich (to appear in *Anal. Appl.* (2011)).
- [19] M. Kleiber and T.D. Hien, *The stochastic finite element methods*. John Wiley & Sons, Chichester (1992).
- [20] R. Milani, A. Quarteroni and G. Rozza, Reduced basis methods in linear elasticity with many parameters. *Comput. Methods Appl. Mech. Eng.* **197** (2008) 4812–4829.
- [21] P. Morin, R.H. Nochetto and K.G. Siebert, Data oscillation and convergence of adaptive FEM. *SIAM J. Numer. Anal.* **38** (2000) 466–488.
- [22] F. Nobile, R. Tempone and C.G. Webster, A sparse grid stochastic collocation method for elliptic partial differential equations with random input data. *SIAM J. Numer. Anal.* **46** (2008) 2309–2345.
- [23] F. Nobile, R. Tempone and C.G. Webster, An anisotropic sparse grid stochastic collocation method for elliptic partial differential equations with random input data. *SIAM J. Numer. Anal.* **46** (2008) 2411–2442.
- [24] Ch. Schwab and A.M. Stuart *Sparse deterministic approximation of Bayesian inverse problems*. Report 2011-16, Seminar for Applied Mathematics, ETH Zürich (to appear in *Inverse Probl.*).
- [25] Ch. Schwab and R.A. Todor, Karhunen-Loève, Approximation of random fields by generalized fast multipole methods. *J. Comput. Phys.* **217** (2000) 100–122.
- [26] R. Stevenson, Optimality of a standard adaptive finite element method. *Found. Comput. Math.* **7** (2007) 245–269.