PREPRINT 454

JULY 2016

igpm | RWTH AACHEN UNIVERSITY

Parametric PDEs: Sparse or low-rank approximations?

Markus Bachmayr[†], Albert Cohen[‡] and Wolfgang Dahmen[†]

Institut für Geometrie und Praktische Mathematik
Templergraben 55, 52062 Aachen, Germany

* This work has been supported by ERC AdG 338977 BREAD, DFG SFB-Transregio 40, DFG Research Group 1779, the Excellence Initiative of the German Federal and State Governments (RWTH Aachen Distinguished Professorship), and DARPA-BAA-15-13.

† Institut für Geometrie und Praktische Mathematik, RWTH Aachen, Templergraben 55, 52056 Aachen, Germany ({bachmayr,dahmen}@igpm.rwth-aachen.de)

‡ Sorbonne Universit´es, UPMC Univ Paris 06, CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, 4 place Jussieu, 75005 Paris, France (cohen@ljll.math.upmc.fr)

Parametric PDEs: Sparse or low-rank approximations?*

Markus Bachmayr†, Albert Cohen‡and Wolfgang Dahmen†

July 14, 2016

Abstract

We consider a class of parametric operator equations where the involved parameters could either be of deterministic or stochastic nature. In both cases we focus on scenarios involving a large number of parameters. Typical strategies for addressing the challenges posed by high dimensionality use low-rank approximations of solutions based on a separation of spatial and parametric variables. One such strategy is based on performing sparse best n-term approximations of the solution map in an *a priori* chosen system of tensor product form in the parametric variables. This approach has been extensively analyzed in the case of tensor product Legendre polynomial bases, for which approximation rates have been established. The objective of this paper is to investigate what can be gained by exploiting further low rank structures, in particular using optimized systems of basis functions obtained by singular value decomposition techniques. On the theoretical side, we show that optimized low-rank expansions can either bring significant or no improvement over sparse polynomial expansions, depending on the type of parametric problem. On the computational side, we analyze an adaptive solver which, at near-optimal computational cost for this type of approximation, exploits low-rank structure as well as sparsity of basis expansions.

MSC 2010: 41A46, 41A63, 42C10, 65D99, 65J10, 65N12, 65N15

Keywords: parameter-dependent PDEs, low-rank approximations, sparse polynomial expansions, a posteriori error estimates, adaptive methods, complexity bounds

1 Introduction

Complex design, optimization, or uncertainty quantification tasks based on parameter dependent families of PDEs arise in virtually all branches of science and engineering. Typical scenarios are models whose physical properties – such as diffusivity, transport velocity or domain geometry – are described by a *finite* number of real parameter values. In certain instances, one may even encounter *infinitely* many parameters of decreasing influence. This occurs for instance in the case of a random stochastic diffusion field represented by an infinite expansion in a given basis. The development and analysis of numerical strategies for capturing the dependence of the PDE on the parameters has been the subject of intensive research efforts in recent years.

*This work has been supported by ERC AdG 338977 BREAD, DFG SFB-Transregio 40, DFG Research Group 1779, the Excellence Initiative of the German Federal and State Governments (RWTH Aachen Distinguished Professorship), and DARPA-BAA-15-13.

†Institut für Geometrie und Praktische Mathematik, RWTH Aachen, Templergraben 55, 52056 Aachen, Germany ({bachmayr,dahmen}@igpm.rwth-aachen.de)

‡Sorbonne Universités, UPMC Univ Paris 06, CNRS, UMR 7598, Laboratoire Jacques-Louis Lions, 4 place Jussieu, 75005 Paris, France (cohen@ljll.math.upmc.fr)

1.1 Problem formulation

The problems that are addressed in this paper have the following general form. Let V be a separable Hilbert space. We consider a parametric operator $A(y)\colon V \to V'$ of the form

$$A(y) := A_0 + \sum_{j\in\mathcal{I}} y_j A_j\,, \quad y \in Y := [-1,1]^{\mathcal{I}}\,, \tag{1.1}$$

where $\mathcal{I} = \{1,\dots,d\}$ or $\mathcal{I} = \mathbb{N}$ in the finite or infinite dimensional case, respectively. In the infinite dimensional case, we require that the above series converges in $\mathcal{L}(V,V')$ for any $y \in Y$. We assume uniform boundedness and ellipticity of $A(y)$ over the parameter domain, that is

$$\langle A(y)v, w\rangle \le R\|v\|_V\|w\|_V \quad \text{and} \quad \langle A(y)v, v\rangle \ge r\|v\|_V^2, \quad v,w \in V,\ y \in Y, \tag{1.2}$$

for some $0 < r \le R < \infty$, which implies in particular that $A(y)$ is boundedly invertible uniformly in $y \in Y$, with

$$\|A(y)\|_{\mathcal{L}(V',V)} \le r^{-1}, \quad y \in Y. \tag{1.3}$$

We also consider parametric data $f\colon Y \to V'$, and for each $y \in Y$, we define $u(y) \in V$ the solution to the equation

$$A(y)\,u(y) = f(y). \tag{1.4}$$

A guiding example is provided by *affinely parametrized* diffusion problems of the form

$$A(y)u := -\operatorname{div}\big(a(y)\nabla u\big) = f, \quad a(y) := \bar{a} + \sum_{j\ge1} y_j\theta_j, \tag{1.5}$$

with homogeneous Dirichlet boundary conditions, posed in the weak sense on a spatial domain D. In this particular case of frequent interest, the data $f \in V'$ is independent of y. The validity of (1.2) is then usually ensured by constraints on the expansion functions θ_j, see (2.1) below. Thus, V is typically a function space defined over some physical domain D, for example the Sobolev space $H^1_0(D)$ in the above case of second order elliptic equations with homogeneous Dirichlet boundary conditions. Therefore the solution may either be viewed as the Hilbert space valued map

$$y \mapsto u(y), \tag{1.6}$$

which acts from Y to V or as the scalar valued map

$$(x,y) \mapsto u(x,y) := u(y)(x), \tag{1.7}$$

where $x \in D$ and $y \in Y$ are refered to as the spatial and parametric variables.

Approximating such solution maps amounts to approximating functions of a large or even infinite number of variables. In applications, one is often interested in specific functionals of the solution. Here we focus on the basic question of approximating the entire solution map in an appropriate norm.

The guiding questions are the following:

(i) Is there a most suitable approximation to cope with the high dimensionality in problems of the form (1.1)? If not, which features of problems (1.1) favor certain approaches over others?

(ii) At what numerical cost can one find these approximations, and how do they depend on particular features of the given problem?

1.2 Sparse and low-rank approximability

Before addressing any concrete numerical schemes, we discuss basic concepts of approximations for the solution map (1.6). The approximations that we consider are of the general form

$$u \approx u_n := \sum_{k=1}^{n} u_k^{\mathrm{x}} \otimes u_k^{\mathrm{y}}, \tag{1.8}$$

where u_k^{x} and u_k^{y} are functions of the spatial and parametric variable respectively. By analogy with a matrix, we say that such a u_n has *rank* at most n.

In other words,

$$u(x,y) \approx u_n(x,y) = \sum_{k=1}^{n} u_k^{\mathrm{x}}(x)\, u_k^{\mathrm{y}}(y), \tag{1.9}$$

or equivalently

$$u(y) \approx u_n(y) = \sum_{k=1}^{n} u_k^{\mathrm{x}}\, u_k^{\mathrm{y}}(y). \tag{1.10}$$

The error between u and u_n can be measured in several ways, the two most common ones being the uniform error $\|u-u_n\|_{L^\infty(Y,V)}$ and the mean-square error $\|u-u_n\|_{L^2(Y,V)}$, where the norms are defined by

$$\|v\|_{L^\infty(Y,V)} := \sup_{y\in Y} \|v(y)\|_V, \quad \|v\|^2_{L^2(Y,V)} := \int_Y \|v(y)\|_V^2 d\mu(y), \tag{1.11}$$

and where μ is a given probability measure over Y. Note that $\|v\|_{L^2(Y,V)} \le \|v\|_{L^\infty(Y,V)}$, and therefore the uniform error dominates the mean-square error.

In the present paper we focus on the mean-square error, with μ being the uniform probability measure. Note that the Bochner space $L^2(Y,V)$ has a tensor product structure $L^2(Y,V) = V \otimes L^2(Y)$ where $L^2(Y) = L^2(Y,\mu)$. The optimal rank n approximation of u in $L^2(Y,V)$ is then given by truncation of the Hilbert-Schmidt decomposition of u interpreted as the operator

$$T_u : v \to \int_Y u(x,y)v(y)d\mu(y), \tag{1.12}$$

acting from $L^2(Y)$ to V. Equivalently the functions $u_1^{\mathrm{x}}, \dots, u_n^{\mathrm{x}}$ are defined as the n first basis vectors in the Karhunen-Loève expansion of the V-valued random variable $u(y)$ when y has the uniform distribution over Y, and $u_n(y)$ is the V-orthogonal projection of $u(y)$ onto

$$V_n := \mathrm{span}\{u_1^{\mathrm{x}}, \dots, u_n^{\mathrm{x}}\} \subset V. \tag{1.13}$$

This particular system of basis functions is a natural benchmark as it minimizes the rank $n = n(\varepsilon)$ required to ensure a mean-square accuracy ε. However, it is unclear how to compute sufficiently good approximations of these basis functions at affordable cost, a point to be taken up again later.

There exist several approaches for deriving computable expansions of the form (1.8). A first category of methods consists in first constructing in an offline stage functions $u_1^{\mathrm{x}}, \dots, u_n^{\mathrm{x}}$ and their span V_n. Then, in an online stage, for any given $y \in Y$, the approximate solution $u_r(y)$ is defined as the Galerkin projection of $u(y)$ on the space V_n. The parametric functions $y \mapsto u_i^{\mathrm{y}}(y)$ are thus implicitly determined by this solution process. Examples of such methods include the reduced basis (RB) and proper orthogonal decomposition (POD) method, see e.g. [31, 38]. These methods benefit from the fact that in many relevant cases the Galerkin projection gives a near-best approximation of $u(y)$ in

V_n. One obstruction is that the construction of spaces V_n which approach the optimal choice described above may require a very large number of offline computations in the case where Y is high-dimensional. For example, in the POD method, one performs approximately a principal component analysis based on samples $u(y^1), \dots, u(y^m)$, where the size of m ensuring an accurate approximation becomes prohibitive in high dimension, see e.g. [26].

A second category of methods is based on defining u_n by approximation of $y \mapsto u(y)$, i.e., the factors u_k^{y} are computed explicitly as well. A low-rank approximation trying to approximately realizing a truncated Hilbert-Schmidt decomposition would be a first example for this category aiming at meeting the above mentioned benchmark. In this case the error caused by truncation should ideally be balanced against the error in approximating the *unknown* basis functions $u_k^{\mathrm{x}}, u_k^{\mathrm{y}}$. These latter approximation problems could be formulated as N-term approximations with respect to fixed bases for the spatial and parametric components, respectively. Overall we thus face a highly nonlinear approximation problem.

A second simpler approach to approximating $y \mapsto u(y)$ is to employ an *a priorily chosen* basis $\{u_1^{\mathrm{y}}, \dots, u_n^{\mathrm{y}}\}$, and compute the u_i^{x} as the corresponding coefficients of this approximation. One prominent example of this approach are orthogonal polynomial expansion methods, see e.g. [17, 18, 33, 42]. In this case, the parametric functions u_i^{y} are picked from the set of *tensorized* Legendre polynomials

$$L_\nu(y) = \prod_{j\geq 1} L_{\nu_j}(y_j), \quad \nu = (\nu_j)_{j\geq 1}, \tag{1.14}$$

with $(L_k)_{k\geq 1}$ the univariate Legendre polynomial sequence normalized in $L^2([-1,1], \frac{dt}{2})$. The functions $(L_\nu)_{\nu\in\mathcal{F}}$ are an orthonormal basis of $L^2(Y)$, where $\mathcal{F}$ is $\mathbb{N}_0^d$ in the case $\mathcal{I} = \{1, \dots, d\}$ or the set of finitely supported sequences of non-negative integers in the case $\mathcal{I} = \mathbb{N}$, that is

$$\mathcal{F} := \{\nu \in \mathbb{N}_0^{\mathbb{N}} : \#\operatorname{supp}\nu < \infty\}. \tag{1.15}$$

One thus has

$$u(y) = \sum_{\nu\in\mathcal{F}} u_\nu L_\nu(y), \quad u_\nu = \int_Y u(y) L_\nu(y) d\mu(y). \tag{1.16}$$

Then, one natural choice for u_n is obtained by restricting the above expansion to the set $\Lambda_n \subset \mathcal{F}$ of indices ν corresponding to the n largest $\|u_\nu\|_V$, since this set minimizes the error $\|u - u_n\|_{L^2(Y,V)}$ among all possible choices of n-term truncations. This strategy for generating sparse polynomial approximations in the context of parametric PDEs was first introduced and analyzed in [11, 12]. In practice, the set Λ_n is not accessible, but provides a benchmark for the performance of algorithms. Note that on the one hand, even if the set Λ_n and the coefficients u_ν were exactly known to us, there is no reason to expect that the resulting u_n achieves an accuracy in $L^2(Y,V)$ comparable to the best rank-n approximations given by the truncation of the Hilbert-Schmidt decomposition.

On the other hand, it is not clear whether best rank-n approximations always achieve a substantially better accuracy in $L^2(Y,V)$ than the best n-term Legendre approximation. In fact, such comparisons may depend strongly on the specific type of the parametric problem. A first central question we would like to address in the present paper is therefore:

(I) *For which types of parametric problems can we expect rank n approximations obtained by Hilbert-Schmidt truncation to be substantially more accurate, regarding the number of terms needed to realize any given target accuracy, than those obtained by best n-term truncation of expansions in a priori chosen bases in the parametric variable such as Legendre*

polynomials?

One may as well go beyond the Hilbert-Schmidt decomposition (1.8) and consider higher-order low-rank tensor representations that correspond to further decompositions of the factors u_k^{y} in (1.8). For simplicity, at this point let us consider this in the finite-dimensional case $d < \infty$ (possibly after truncating the expansion (1.1) for $A(y)$). Introducing an additional tensor decomposition of the factors $\mathbf{u}_k^{\mathrm{y}}$, we obtain the general approximation format

$$u_n = \sum_{k_0=1}^{r_0} u_{k_0}^{\mathrm{x}} \otimes \Big(\sum_{k_1=1}^{r_1} \cdots \sum_{k_d=1}^{r_d} \mathbf{a}_{k_0,k_1,\ldots,k_d} \bigotimes_{j=1}^{d} u_{k_j}^{\mathrm{y},j} \Big), \tag{1.17}$$

where each $u_{k_j}^{\mathrm{y},j}$ is a function of the individual variable y_j. The minimal r_j such that u_n can be represented in the form (1.17) are called *multilinear ranks* of u_n.

In *hierarchical tensor representations* (with the *tensor train format* as a special case), see e.g. [21,24,35], the high-order core tensor $\mathbf{a} = (\mathbf{a}_{k_0,k_1,\ldots,k_d})_{k_0,\ldots,k_d}$ is further decomposed in terms of lower-order tensors, based on matricizations of $\mathbf{a}$. For instance, if

$$\hat{r}_i = \operatorname{rank}\big(\mathbf{a}_{(k_0,\ldots,k_i),(k_{i+1},\ldots k_d)}\big), \tag{1.18}$$

one has a factorized representations of the form

$$\mathbf{a}_{k_0,k_1,\ldots,k_d} = \sum_{\ell_1=1}^{\hat{r}_1} \mathbf{M}^{(1)}_{k_0,k_1,\ell_1} \sum_{\ell_2=1}^{\hat{r}_2} \mathbf{M}^{(2)}_{\ell_1,k_2,\ell_2} \cdots \sum_{\ell_{d-1}=1}^{\hat{r}_{d-1}} \mathbf{M}^{(d-1)}_{\ell_{d-2},k_{d-2},\ell_{d-1}} \mathbf{M}^{(d)}_{\ell_{d-1},k_d} \tag{1.19}$$

in terms of the tensors $\mathbf{M}^{(i)}$, $i = 1,\ldots,d$ of order at most three, and only these low-order tensors need to be stored and manipulated.

This format contains sparse polynomial expansions (1.16) as a special case: let $\nu(k_0)$, $k_0 = 1,\ldots,r_0$ be an enumeration of elements of $\mathbb{N}_0^d$, and choose $\mathbf{a}_{k_0,k_1,\ldots,k_d} = \delta_{\nu(k_0),(k_1,\ldots,k_d)}$, $u_{k_j}^{\mathrm{y},j} = L_{k_j}$. Another noteworthy special case concerns the choices $r_1 = \ldots = r_d = 1$ and $\mathbf{a}_{k_0,1,\ldots,1} = 1$, which is often referred to as the *canonical format*. In this case $n = r_0$ is called the *canonical rank*. Thus, the overall necessary number of terms in such an approximation format required for realizing a target accuracy ε in $L^2(Y,V)$ may be smaller than for the format (1.16), but comes at the expense of an even higher level of structural nonlinearity.

What has been discussed so far, however, does not yet give numerically realizable approximations: the functions u_k^{x} and u_k^{y} in (1.8), $u_{k_0}^{\mathrm{x}}$ and $u_{k_j}^{\mathrm{y},j}$ in (1.17), as well as the coefficients v_ν in (1.16) in turn need to be approximated as well. For instance, one may choose a fixed basis $\{\psi_\lambda\}_{\lambda\in\mathcal{S}}$ of V and expand

$$u_\nu = \sum_{\lambda\in\mathcal{S}} \mathbf{u}_{\lambda,\nu}\psi_\lambda,$$

in (1.16), which leads us to consider full approximations of the form

$$u \approx \sum_{(\lambda,\nu)\in\Lambda} \mathbf{u}_{\lambda,\nu}\, \psi_\lambda \otimes L_\nu. \tag{1.20}$$

with coefficients $\mathbf{u}_{\lambda,\nu} \in \mathbb{R}$ and $\Lambda \subset \mathcal{S}\times\mathcal{F}$. Here one can again ask for best n-term approximations with respect to the basis $\{\psi_\lambda \otimes L_\nu\}_{\lambda\in\mathcal{S},\nu\in\mathcal{F}}$, obtained by minimizing the error over all Λ with $\#\Lambda = n$.

This can be compared to (1.8) using the respective basis expansions for u_k^{x} and u_k^{y}, that is, to approximations of the form

$$u \approx \sum_{k=1}^{n} \Big(\sum_{\lambda \in \Lambda_k^{\mathrm{x}}} \mathbf{u}_{k,\lambda}^{\mathrm{x}} \psi_\lambda \Big) \otimes \Big(\sum_{\nu \in \Lambda_k^{\mathrm{y}}} \mathbf{u}_{k,\nu}^{\mathrm{y}} L_\nu \Big) \tag{1.21}$$

with $\Lambda_k^{\mathrm{x}} \subset \mathcal{S}$, $\Lambda_k^{\mathrm{y}} \subset \mathcal{F}$, $k = 1, \dots, n$.

1.3 Computational complexity

There exist several computational approaches for constructing approximations of the types outlined above, based on either a priori estimates or a posteriori criteria, which will be discussed in more detail later.

Of course, the number of terms n in a separable approximation of the form (1.8) does not yet determine the complexity of this approximation since each of the non a-priorily chosen factors needs to be approximated. This can be addressed as in (1.21) by choosing for each variable a fixed background basis and searching for sparse approximations of each factor with respect to these bases. The resulting approximation is then described by a finite number $n_{\mathrm{dof}} = \sum_{k=1}^{n} (\#\Lambda_k^{\mathrm{x}} + \#\Lambda_k^{\mathrm{y}})$ of degrees of freedom. The smallest possible $n_{\mathrm{dof}}(\varepsilon)$ required for any approximation of accuracy ε in $L^2(Y, V)$ describes the *representation complexity* of a function with respect to the given approximation format. Expanding the arising products, the low-rank representation (1.21) can still be interpreted as a linear combination of prescribed tensor product basis functions, but with an algebraically nonlinear parametrization of coefficients, as opposed to (1.20) where the coefficients enter linearly. Thus the potentially more condensed representation by fewer degrees of freedom comes at the price of handling this nonlinear mapping.

We thus need to address the possibly higher computational cost entailed by a stronger structural nonlinearity, which could well offset the n-term rate: Computing for a given approximation format such an ε-accurate approximation by a given algorithm will require a number $n_{\mathrm{op}} = n_{\mathrm{op}}(\varepsilon)$ of operations that is typically larger than the representation complexity $n_{\mathrm{dof}}(\varepsilon)$.

Here we aim for algorithms that are universal in the sense that they do not require a priori knowledge on the approximability of the solution (e.g., on the decay of coefficients), but adjust to such approximability automatically. This goes hand in hand with a mechanism for obtaining a posteriori error bounds, making use only of the given data, viz., right hand side and operator representation.

A second main question addressed in this paper is therefore:

(II) *For a given parametric problem and approximation format (1.16), (1.8) or (1.17), can one contrive a universal numerical scheme that can achieve any given target accuracy ε, certified a posteriori? Are the approximate solutions close to the minimum required representation complexity $n_{\mathrm{dof}}(\varepsilon)$? And can we relate n_{op} to ε and hence to $n_{\mathrm{dof}}(\varepsilon)$?*

1.4 Relation to previous work

There is a variety of results on the convergence of sparse polynomial expansions (1.16), see, e.g., [4, 11, 12]. Furthermore, some estimates are available that include multilevel spatial discretizations and hence provide upper bounds for the error of best n-term approximation (1.20), see, e.g., [12,13]. Concerning our question (I), there are only few specialized results

on possible gains by low-rank approximations, all concerning the case of finitely many parameters [31], [2].

There are various approaches for generating sparse polynomial expansions, for instance based on collocation [1, 7] or adaptive Taylor expansion [8]. Note that these strategies do not currently yield *a-posteriori* error bounds for the computed solutions, and their performance is thus described by *a-priori* estimates which may not be sharp.

The adaptive method proposed in [14], based on finite element discretization for the spatial variable, yields a posteriori error bounds for the full approximations. However, complexity bounds are given only in terms of the resulting finite element meshes.

Adaptive schemes using wavelet-based spatial discretizations, which yield approximations of the form (1.20), have been studied by Gittelson [19, 20]. In this case, bounds for the complete computational complexity are proven.

Reduced basis and POD methods [26,37,38] correspond to expansions of the form (1.8), where only the spatial basis elements u_k^{x} are explicitly computed. For known variants of these methods, accuracy guarantees in the respective norms require a sufficiently dense sampling of the parameter domain, which becomes prohibitive for large d, and one only obtains a posteriori bounds for the resulting V-error in each given $y \in Y$.

In methods based on higher-order tensor representations, instead of sampling in the parameter domain, one also approximates u_k^{y} as in (1.17), at the price of additional approximability requirements as in (1.19). A variety of schemes have been proposed that operate on fixed discretizations [27, 28, 30, 34], which do not yield information on the discretization error. Based on [14], an adaptive scheme for hierarchical tensor approximation is proposed in [15]. It provides rigorous a posteriori bounds for the approximation error, but is not proven to converge.

1.5 Novelty of the paper and outline

Question (I) posed in §1.2 is addressed in §2, by studying specific examples of parametric problems of the diffusion form (1.5). For a certain class of such problems, we prove that the best n-term Legendre approximation is already near-optimal among all rank-n approximations. In other words, the $L^2(Y, V)$ error achievable by this particular fully separable approximation is bounded, up to a fixed multiplicative constant, by the $L^2(Y, V)$ error achieved by any fully separable approximation of the form (1.17). For other examples, we prove that optimized low rank approximations perform significantly better than best n-term Legendre approximations in terms of faster convergence rates. In summary, regarding n-term approximation there is no universally best strategy.

Question (II) is addressed in sections §3 to §7. A generic algorithm is described in §3 based on the work in [5], which is guaranteed to converge. Furthermore, it yields rigorous a posteriori error bounds, using only information on the problem data. Suitable specifications cover all above mentioned types of approximations (1.8), (1.16), and (1.17). The scheme is formulated in a general sequence space framework, using a discretization of the space $L^2(Y, V)$ through a basis with elements of the form $\psi_\lambda \otimes L_\nu$. Here, $\{\psi_\lambda\}_{\mu\in\mathcal{S}}$ is a given Riesz basis of V (for example, a wavelet basis in the case where V is a Sobolev space) and $\{L_\nu\}_{\nu\in\mathcal{F}}$ is the previously described multivariate Legendre basis. The algorithm performs an iteration in the sequence space $\ell^2(\mathcal{S}\times\mathcal{F})$, which involves at each step specific routines RECOMPRESS and COARSEN aiming at respectively controlling the rank of the current approximation as well as the number of degrees of freedom that describe the different factors involved in this approximation.

We then describe two realizations of this generic algorithm corresponding to two distinct settings. In §4 we apply the algorithm for the generation of approximations with

full separation, as described by (1.17). In this case the RECOMPRESS routine is based on a hierarchical SVD truncation, that is, SVD truncations for the matrices involved in a given hierarchical format. We analyze the performance of the algorithms for classes described by the decay of the corresponding singular value and joint sparsity of the corresponding singular vectors. In this particular realization, it is required that the dimension d of the parametric variable is finite.

§5 to 7 are devoted to the case of *anisotropic* parameter dependence in the diffusion problem (1.5), where we allow the dimension d of the parametric variable to be infinite. In this latter case any numerical scheme involves a truncation of the operator representation (1.1). Being able to retrieve then the representation complexity $n_{\mathrm{dof}}(\varepsilon)$ of the exact solution turns out to depend on how the decay of the truncation error relates to the sparsity parameters governing the representation complexity $n_{\mathrm{dof}}(\varepsilon)$ of the approximation format under consideration. A first major theme in §5 is therefore the interplay between the decay of the corresponding truncation error and the approximability of the solutions. Specifically, we discuss two types of operator parametrizations, namely, in §5.1, Karhunen-Loève type expansions comprising globally supported basis functions with increasing oscillatory behavior, and in §5.2 wavelet-type multilevel expansions with corresponding space-frequency localization. One of our central findings is that in the case of wavelet-type parametrizations one can better exploit the approximability of the solutions in the adaptive scheme, since this has a substantial influence on the compressibility of the operator. It is perhaps worth mentioning that our further algorithmic developments require substantially weaker assumptions on the A_j in (1.1) than the methods in [14,15], which require summability of $(\|A_j\|)_{j\geq 1}$.

In §6 we analyze a specialized version of Algorithm 3.1 producing n-term sparse Legendre expansions. In this version the routine RECOMPRESS is simply the identity, and hence Algorithm 3.1 agrees with the adaptive solver developed and analyzed in [10]. A key ingredient is the adaptive approximation of the operator analyzed in §6.1 (based on matrix compression results in Appendix A) for the operator parametrizations from §5.2. We then establish in §6.2 convergence and complexity rates which significantly improve on earlier results for similar schemes in [19] in that best n-term rates are retrieved for a wider range of solution approximabilities.

In §7 we apply Algorithm 3.1 for the generation of approximations with only space-parameter separation (1.8) and (1.21), assuming again parametrizations of the type considered in §5.2. In particular, in this realization, we allow the dimension d of the parametric variable to be infinite. In this case the RECOMPRESS routine is based on standard SVD truncation. We analyze the performance of the algorithm for classes of solutions described on the one hand by the rate of decay of the singular values which appear in the Hilbert-Schmidt decomposition, and on the other hand by a notion of joint sparsity for the corresponding singular vectors.

For each realization of the basic algorithm, we derive convergence and complexity estimates for suitable benchmark classes motivated by the considerations in §2 and §5 and discuss their optimality in the respective contexts.

2 Low-rank approximability of parametric problems

The goal of this section is to use a particular example of a problem of the type (2.1) in order to identify two basic regimes of parameter-dependent problems which lead to different conclusions when comparing sparse and low-rank approximations.

We shall focus in what follows on *affinely parametrized* diffusion problems of the form

$$A(y)u := -\operatorname{div}\big(a(y)\nabla u\big) = f, \quad a(y) := \bar{a} + \sum_{j\geq 1} y_j\theta_j, \tag{2.1}$$

posed in the weak sense on a domain D with homogeneous Dirichlet boundary conditions, where the expansion of the coefficient satisfies the *uniform ellipticity assumption*.

$$\sum_{j\geq 1}|\theta_j(x)| \leq \bar{a}(x) - \underline{\alpha}, \quad x \in D, \tag{2.2}$$

for some $\underline{\alpha} > 0$. We then have $V = H_0^1(D)$, and the corresponding operators $A_j \colon V \to V'$ for $j \in \{0\} \cup \mathcal{I}$ are defined by

$$\langle A_0 u, v\rangle := \int_D \bar{a}\nabla u \cdot \nabla v \, dx\,, \qquad \langle A_j u, v\rangle := \int_D \theta_j \nabla u \cdot \nabla v \, dx, \quad i \in \mathcal{I},$$

for $u, v \in V$.

For any $\nu = (\nu_j)_{j\geq 1} \in \mathcal{F}$, we define the coefficients

$$t_\nu(y) = \frac{1}{\nu!}\partial^\nu u(y), \quad \nu! := \prod_{j\geq 1}\nu_j!, \tag{2.3}$$

of a Taylor expansion of u at y. Denoting by $e^j = (0, \dots, 0, 1, 0, \dots)$ is the j-th Kroenecker sequence, by differentiating the equation, we find that these coefficients are given by the recursion

$$t_\nu(y) := -A(y)^{-1} \sum_{j\in\operatorname{supp}\nu} A_j t_{\nu-e^j}(y), \tag{2.4}$$

initialized by

$$t_0(y) = A(y)^{-1} f = u(y). \tag{2.5}$$

As simple yet instructive examples, we consider spatially univariate problems with $\Omega =]0,1[$, $\bar{a} = 1$, and

$$\theta_j = b_j \chi_{D_j}, \tag{2.6}$$

where $b_j \in]0,1[$ are constants and the open subintervals D_j of Ω have disjoint closures so that the diffusion coefficient is a strictly positive piecewise constant,

$$a(y) = \bar{a} + \sum_{j\geq 1} y_j b_j \chi_{D_j}.$$

We consider two types of such inclusion systems. As a first scenario, we consider the following.

Example 2.1. Let $d := \#\mathcal{I} < \infty$, $D_j \subset]0,1[$ for $j = 1, \dots, d$ with pairwise disjoint $\overline{D_j}$, and $b_j = \xi$ for some $\xi \in]0,1[$.

For low-rank approximation, we then have the following result.

Proposition 2.2. *In Example 2.1, for any $f \in V'$, one has* $\operatorname{rank}(\mathbf{u}) \leq 4d + 1$.

Proof. This follows by the same arguments as in [2, Example 2.2]: the endpoints of the D_j induce a partition of $]0,1[$ into $2d+1$ intervals. For each such interval I, for any F such that $F'' = f$, we have $u(y)|_I \in \operatorname{span}\{\chi_I, x\,\chi_I, F\,\chi_I\}$. Hence $u(y)$ is contained in a y-independent space of dimension $6d+3$ for all y. In addition, there are $2d+2$ continuity conditions, independent of y, at the interval boundaries, which leaves at most $4d+1$ degrees of freedom. □

We observe on the other hand that the Legendre expansions for this problem involves infinitely many nonzero coefficients, that is, the solution map $y \mapsto u(y)$ is not a polynomial in y. This can be checked for example by considering the Taylor coefficient of order n in a given variable j at the origin, that is,

$$t_{n,j} := t_{ne^j}(0) = \frac{1}{n!}\partial_{y_j}^n u(0). \tag{2.7}$$

As a particular case of (2.4), we have

$$\int_D \bar{a}\nabla t_{n,j} \cdot \nabla v \, dx = -\int_D \theta_j \nabla t_{n-1,j} \cdot \nabla v \, dx. \tag{2.8}$$

Since $t_{0,j} = u(0)$ is not trivial, there is at least one variable j such that $t_{1,j}$ does not vanish on D_j. Then, taking $v = t_{n-1,j}$ in the above recursion shows by contradiction that $t_{n,j}$ does not vanish on D_j, for all values of $n \geq 0$. Thus $y \mapsto u(y)$ cannot be a polynomial.

Low-rank approximations thus give substantially faster convergence than Legendre expansions in this case. Similar results showing substantial advantages of best low-rank approximations have also been obtained for spatially two-dimensional examples of analogous structure in [2].

As a second scenario, we consider a problem with countably many parameters of decreasing influence.

Example 2.3. Let $\mathcal{I} = \mathbb{N}$, and let $D_j \subset]0,1[$ be disjoint with $|D_j| > 0$ for all j. In addition, let $(b_j)_{j\geq 1} \in \ell^q(\mathbb{N})$ for some $q > 0$.

As an immediate consequence of the results in [4, §4.1], one has the following.

Proposition 2.4. *In Example 2.3, for all right hand sides $f \in V'$, one has $(\|u_\nu\|_V)_{\nu\in\mathcal{F}} \in \ell^p(\mathcal{F})$ for $p = \frac{2q}{2+q}$, and there exists $f \in V'$ such that $(\|u_\nu\|_V)_{\nu\in\mathcal{F}} \notin \ell^{p'}(\mathcal{F})$ for $0 < p' < p$.*

If σ_n are the singular values of u, then for the decreasing rearrangement $(u_n^*)_{n\geq 1}$ of $(\|u_\nu\|_V)_{\nu\in\mathcal{F}}$ we clearly have $u_n^* \geq \sigma_n$. As the following new result shows by similar arguments as in [4, §4.1], in general the singular values actually do not have faster decay in this situation than the ordered norms of the Legendre coefficients.

Proposition 2.5. *In Example 2.3, if $(b_j) \notin \ell^{q'}(\mathbb{N})$ for any $0 < q' < q$, then there exists an $f \in V'$ such that the singular values of u are not in $\ell^{p'}(\mathbb{N})$ for $0 < p' < p = \frac{2q}{2+q}$.*

Proof. We first observe that the singular values of $u = \sum_{\nu\in\mathcal{F}} u_\nu \otimes L_\nu$ are bounded from below by those of $\tilde{u} = \sum_{j\geq 1} u_{e^j} \otimes L_{e^j}$, with e^j denoting the j-th Kronecker sequence. This follows from the fact that $\tilde{u} = (\mathrm{I} \otimes \tilde{P})u$, where $\tilde{P}$ is the projector onto $\overline{\mathrm{span}}\{L_{e^j}\}_{j\geq 1}$.

For u_{e^j}, one has by Rodrigues' formula the explicit representation

$$u_{e^j} = \frac{\sqrt{3}}{2}\int_Y t_j(y)\,(1-y_j^2)\,d\mu(y) \tag{2.9}$$

in terms of the first-order derivatives $t_{e^j}(y) = \partial_{y_j} u(y)$.

Let h_j be the symmetric hat functions with support D_j. We now choose

$$f = -\sum_{j\geq 1} c_j h_j'',$$

where $\sum_{j\geq 1} c_j^2/|D_j| < \infty$, which yields $f \in V'$ and

$$t_0(y) = \sum_{j\geq 1} (1 + b_j y_j)^{-1} c_j h_j.$$

By (2.4),

$$t_{e^j}(y) = -(1+b_j y_j)^{-2} b_j c_j h_j$$

and as a consequence of (2.9),

$$u_{e^j} = -M_j b_j c_j h_j, \qquad M_j := \frac{\sqrt{3}}{2}\int_{-1}^{1} \frac{1-y^2}{(1+b_j y)^2}\frac{dy}{2} \geq \frac{1}{4\sqrt{3}(1-\max_j b_j)} =: M_0.$$

We thus obtain

$$\langle u_{e^i}, u_{e^j}\rangle_V = 0, \quad i \neq j,$$

as well as

$$\|u_{e^j}\|_V \geq M_0 b_j c_j \|h_j\|_V = \frac{2M_0 b_j c_j}{\sqrt{|D_j|}}.$$

Since (b_j) is precisely in $\ell^q(\mathbb{N})$, by choosing $c_j = b_j^{q/2}\sqrt{|D_j|}$, which guarantees in particular that $(c_j/\sqrt{|D_j|})_{j\geq 1} \in \ell^2(\mathbb{N})$ as required, we arrive at the statement. □

The above result shows that from an asymptotic point of view, in Example 2.3, there is in general nothing to be gained by low-rank approximation: there always exist right hand sides f such that the singular values have *exactly the same* asymptotic decay as the ordered norms of Legendre coefficients.

Numerical tests as in Example 5.2 indicate that this also holds true for problems with different types of parametrization and more general f.

Remark 2.6. *The conclusion of Proposition 2.5 reveals that, in the case of Example 2.3 and if $(b_j)_{j\geq 1} \notin \ell^{q'}$ for all $0 < q' < q$, then any separable approximation of the form* (1.8) *satisfies*

$$\|u - u_n\|_{L^2(Y,V)} \geq c_r n^{-r}, \quad n \geq 1, \tag{2.10}$$

for some $c_r > 0$, whenever $r > \frac{1}{q}$. In turn, we also have

$$\|u - u_n\|_{L^\infty(Y,V)} \geq c_r n^{-r}, \quad n \geq 1. \tag{2.11}$$

This implies that the Kolmogorov n-width

$$d_n(\mathcal{M})_V = \inf_{\dim(E)\leq n} \max_{v\in\mathcal{M}} \operatorname{dist}(v,E)_V, \tag{2.12}$$

of the solution manifold $\mathcal{M} := \{u(y) \,:\, y \in Y\}$ satisfies a similar lower bound

$$d_n(\mathcal{M})_V \geq c_r n^{-r}, \quad n \geq 1. \tag{2.13}$$

*While upper bounds for $d_n(\mathcal{M})_V$ in parametric PDEs are typically proved by exhibiting a particular separable approximation and studying its convergence in $L^\infty(Y,V)$, see [**?**, 2], lower bounds are generally out of reach and the ones given above constitute a notable exception.*

Remark 2.7. *One arrives at analogous observations in similar higher-dimensional settings. The construction of Example 2.3 immediately carries over to spatial domains with $m > 1$ when the definition of f is based on higher-dimensional hat functions. Examples similar to Example 2.1 with $m = 2$ have been considered in [2], where the sequence of singular values is no longer compactly supported but still decays exponentially.*

3 A generic algorithm

In this section, we follow the approach developed in [5], by first reformulating the general equation (1.4) in a sequence space, and then introducing a generic resolution algorithm based on this equivalent formulation.

We first notice that (1.4) may also be written as,

$$Au = f, \tag{3.1}$$

where A is elliptic and boundedly invertible from $L^2(Y,V)$ to $L^2(Y,V')$ and can be defined in a weak sense by

$$\langle Au, v\rangle := \int_Y \langle A(y)u(y), v(y)\rangle d\mu(y), \quad u, v \in L^2(Y,V). \tag{3.2}$$

We assume that $f \in L^2(Y,V')$, so that there exists a unique solution $u \in L^2(Y,V)$.

Given a Riesz basis $\{\psi_\lambda\}_{\lambda\in\mathcal{S}}$ of V, we tensorize it with the orthonormal basis $\{L_\nu\}_{\nu\in\mathcal{F}}$ of $L^2(Y)$. The resulting system $\{\psi_\lambda \otimes L_\nu\}_{(\lambda,\nu)\in\mathcal{S}\times\mathcal{F}}$ is a Riesz basis of $L_2(Y,V)$, which we now use to discretize (3.1). For this purpose, we define the matrices

$$\mathbf{A}_j := \big(\langle A_j\psi_{\lambda'}, \psi_\lambda\rangle\big)_{\lambda,\lambda'\in\mathcal{S}} \quad \text{and} \quad \mathbf{M}_j = \left(\int_Y y_j L_\nu(y) L_{\nu'}(y)\, d\mu(y)\right)_{\nu,\nu'\in\mathcal{F}}, \tag{3.3}$$

where $\mathbf{M}_0$ is set to be the identity matrix, and the right hand side column vector

$$\mathbf{f} := \big(\langle f, \psi_\lambda \otimes L_\nu\rangle\big)_{(\lambda,\nu)\in\mathcal{S}\times\mathcal{F}}. \tag{3.4}$$

We thus obtain an equivalent problem

$$\mathbf{A}\mathbf{u} = \mathbf{f} \tag{3.5}$$

on $\ell^2(\mathcal{S}\times\mathcal{F})$ where

$$\mathbf{A} := \sum_{j\geq 0} \mathbf{A}_j \otimes \mathbf{M}_j \tag{3.6}$$

and $\mathbf{u} = \big(u_{\lambda,\nu}\big)_{(\mu,\nu)\in\mathcal{S}\times\mathcal{F}}$ is the coordinate vector of u in the basis $\{\psi_\mu \otimes L_\nu\}_{(\mu,\nu)\in\mathcal{S}\times\mathcal{F}}$.

Regarding $\nu \in \mathcal{F}$ as the column index of the infinite matrix $\mathbf{u} = (\mathbf{u}_{\mu,\nu})_{\mu\in\mathcal{S},\nu\in\mathcal{F}}$, we denote by $\mathbf{u}_\nu$ the columns of $\mathbf{u}$, which are precisely the basis representations of the Legendre coefficients $u_\nu \in V$.

In what follows we always denote by $\|\cdot\|$ the ℓ^2-norm on the respective index set which could be $\mathcal{S}$, $\mathcal{F}$ or $\mathcal{S}\times\mathcal{F}$, or the corresponding operator norm when this is clear from the context. Since $\{\psi_\mu\}_{\mu\in\mathcal{S}}$ is a Riesz basis for V we have $\|u_\nu\|_V \sim \|\mathbf{u}_\nu\|$ uniformly in $\nu \in \mathcal{F}$ which together with boundedness and ellipticity of A implies that $\mathbf{A}$ is bounded and elliptic on $\ell^2(\mathcal{S}\times\mathcal{F})$ and we have

$$\|\mathbf{u}\| \sim \|\mathbf{A}\mathbf{u}\| \sim \|Au\|_{L^2(Y,V')} \sim \|u\|_{L^2(Y,V)} \tag{3.7}$$

with uniform constants. On account of (3.7), solving (3.5) approximately up to some target accuracy is equivalent to solving (3.5) in ℓ^2 to essentially the same accuracy.

As a further consequence, one can find a fixed positive ω such that $\|\mathbf{I} - \omega\mathbf{A}\| \leq \rho < 1$, ensuring that a simple Richardson iteration converges with a fixed error reduction rate per step. This serves as the conceptual starting point for the adaptive low-rank approximation scheme introduced in [5] as given in Algorithm 3.1.

Algorithm 3.1 $\mathbf{u}_\varepsilon = \textsc{solve}(\mathbf{A}, \mathbf{f}; \varepsilon)$

input $\omega > 0$ and $\rho \in (0,1)$ such that $\|\mathrm{I} - \omega\mathbf{A}\| \le \rho$, $\lambda_{\mathbf{A}} \le \|\mathbf{A}^{-1}\|^{-1}$,
$\kappa_1, \kappa_2, \kappa_3 \in (0,1)$ with $\kappa_1 + \kappa_2 + \kappa_3 \le 1$, and $\beta \ge 0$.
output $\mathbf{u}_\varepsilon$ satisfying $\|\mathbf{u}_\varepsilon - \mathbf{u}\| \le \varepsilon$.

$\mathbf{u}_0 := 0$, $\delta := \lambda_{\mathbf{A}}^{-1}\|\mathbf{f}\|$
$k := 0$, $J := \min\{j\colon \rho^j(1 + (\omega+\beta)j) \le \frac{1}{2}\kappa_1\}$
while $\frac{1}{2^k}\delta > \varepsilon$
    $\mathbf{w}_0 := \mathbf{u}_k$, $j \leftarrow 0$
    repeat
        $\eta_j := \rho^{j+1}\frac{1}{2^k}\delta$
        $\mathbf{r}_j := \textsc{apply}(\mathbf{w}_j; \frac{1}{2}\eta_j) - \textsc{rhs}(\frac{1}{2}\eta_j)$
        $\mathbf{w}_{j+1} := \textsc{recompress}(\mathbf{w}_j - \omega\mathbf{r}_j; \beta\eta_j)$
        $j \leftarrow j + 1$.
    until $(j \ge J \quad \vee \quad \lambda_{\mathbf{A}}^{-1}\rho\|\mathbf{r}_{j-1}\| + (\lambda_{\mathbf{A}}^{-1}\rho + \omega + \beta)\eta_{j-1} \le \frac{1}{2^{k+1}}\kappa_1\delta)$
    $\mathbf{u}_{k+1} := \textsc{coarsen}\big(\textsc{recompress}(\mathbf{w}_j; \frac{1}{2^{k+1}}\kappa_2\delta); \frac{1}{2^{k+1}}\kappa_3\delta\big)$
    $k \leftarrow k + 1$
end while
$\mathbf{u}_\varepsilon := \mathbf{u}_k$

This basic algorithmic template can be used to produce various types of sparse and low-rank approximations, with appropriate choices of the subroutines APPLY, RHS, COARSEN, and RECOMPRESS.

The procedures COARSEN and RECOMPRESS are independent of the considered $\mathbf{A}$ and $\mathbf{f}$, and satisfy

$$\|\textsc{coarsen}(\mathbf{v};\eta) - \mathbf{v}\| \le \eta, \quad \|\textsc{recompress}(\mathbf{v};\eta) - \mathbf{v}\| \le \eta, \tag{3.8}$$

for any $\eta \ge 0$ and any compactly supported $\mathbf{v} \in \ell^2(\mathcal{S} \times \mathcal{F})$. Here COARSEN is intended to reduce the support of the sequence $\mathbf{v}$, whereas RECOMPRESS reduces the rank of $\mathbf{v}$ in a low-rank tensor representation. The particular realizations of these routines depend on the dimensionality of the problem and on the type of approximation. We shall use the constructions given in [5].

The routines APPLY and RHS are assumed to satisfy, for compactly supported $\mathbf{v}$ and any $\eta > 0$, the requirements

$$\|\textsc{apply}(\mathbf{v};\eta) - \mathbf{A}\mathbf{v}\| \le \eta, \quad \|\textsc{rhs}(\eta) - \mathbf{f}\| \le \eta. \tag{3.9}$$

Their construction not only depends on the type of approximation, but also on the specific problem under consideration. These two routines are indeed the main driver of adaptivity in Algorithm 3.1, and a major part of what follows concerns the construction of APPLY in different scenarios.

It hinges on the compression of matrices by exploiting their near-sparsity in certain basis representations. We use the following notion introduced in [9]: A bi-infinite matrix $\mathbf{B}$ is called *s^*-compressible* if there exist matrices $\mathbf{B}_n$ with $\alpha_n 2^n$ entries per row and column and such that

$$\|\mathbf{B} - \mathbf{B}_n\| \le \beta_n 2^{-sn}, \quad \text{for } \ 0 < s < s^*, \tag{3.10}$$

and where the sequences $\boldsymbol{\alpha} = (\alpha_n)_{n\in\mathbb{N}}$ and $\boldsymbol{\beta} = (\beta_n)_{n\in\mathbb{N}}$ are summable. Here we always assume $\mathbf{B}_0 = 0$. Furthermore, a bi-infinite matrix $\mathbf{B}$ that has at most k entries in each row and column is called *k-sparse*.

Remark 3.1. *As shown in [5], regardless of the specifications of the routines* APPLY, RHS, COARSEN, RECOMPRESS*, Algorithm 3.1 terminates after finitely many steps and its output* $\mathbf{u}_\varepsilon$ *satisfies* $\|\mathbf{u}-\mathbf{u}_\varepsilon\|\le\varepsilon$.

At this point, we record for later usage a particular feature of $\mathbf{A}$ that arises as a consequence of our choice of tensor product orthogonal polynomials for the parameter-dependence: The approximate application of $\mathbf{A}$ is facilitated by the fact that the matrices $\mathbf{M}_j$ are *bidiagonal*. That is, in view of the three-term recurrence relation

$$tL_n(t)=p_{n+1}L_{n+1}(t)+p_nL_{n-1}(t),\quad L_{-1}\equiv 0, \tag{3.11}$$

where

$$p_0=0,\qquad p_n=\frac{1}{\sqrt{4-n^{-2}}},\quad n>0, \tag{3.12}$$

one has $\int_U y_j\,L_\nu(y)L_\mu(y)\,d\mu(y)=0$ whenever $j\notin\operatorname{supp}\nu\cup\operatorname{supp}\mu$, providing

$$(\mathbf{M}_j)_{\nu,\nu'}=p_{\nu_j}\delta_{\nu+e^j,\nu'}+p_{\nu_j-1}\delta_{\nu-e^j,\nu'} \tag{3.13}$$

with the Kronecker sequence $(e_i^j)_{i\in\mathcal{I}}:=(\delta_{i,j})_{i\in\mathcal{I}}\in\mathcal{F}$.

4 Hierarchical tensor representations in the case of finitely many parameters

We begin by considering the setting

$$\mathcal{I}=\{1,\dots,d\}. \tag{4.1}$$

then $\mathcal{F}=\mathbb{N}_0^d$ and $\mathbf{u}\in\ell^2(\mathcal{S}\times\mathbb{N}_0\times\cdots\times\mathbb{N}_0)$.

Here we are interested in the case that all coordinates in $\mathcal{I}$ have comparable influence. As illustrated in §2, a direct sparse Legendre expansion of $\mathbf{u}$ over $\mathcal{S}\times\mathcal{F}$ will then in general be infeasible already for moderately large d. However, one may as well exploit Cartesian product structure in $\mathcal{F}$, regarding $\mathbf{u}$ as a higher-order tensor, and using corresponding hierarchical low-rank representations. As we shall detail in what follows, the results of [5] can be adapted to this problem in a rather straightforward manner.

It will be convenient to introduce a numbering of tensor modes as follows: $\mathcal{G}_{\mathrm{x}}:=\mathcal{S}$, $\mathcal{G}_1:=\mathbb{N}_0$, …, $\mathcal{G}_d:=\mathbb{N}_0$. We additionally introduce the notation

$$\hat{\mathcal{I}}:=\{\mathrm{x}\}\cup\mathcal{I}.$$

The representations of higher-order tensors which we consider are built on the Hilbert-Schmidt case via *matricizations*: for each nonempty $M\subset\hat{\mathcal{I}}$, $\mathbf{u}$ induces a compact operator $T_{\mathbf{u}}^{(M)}\colon\ell^2(\times_{i\in\hat{\mathcal{I}}\setminus M}\mathcal{G}_i)\to\ell^2(\times_{i\in M}\mathcal{G}_i)$.

In terms of the left singular vector $\{\mathbf{U}_k^{(i)}\}_{k\in\mathbb{N}}$ of $T_{\mathbf{u}}^{(\{i\})}$, $i\in\hat{\mathcal{I}}$, we obtain the *HOSVD representation* [32] in the *Tucker format* [40, 41] as in (1.17),

$$\mathbf{u}=\sum_{1\le k_i\le r_i\colon i\in\hat{\mathcal{I}}}\mathbf{a}_{\mathsf{k}}\bigotimes_{i\in\hat{\mathcal{I}}}\mathbf{U}_{k_i}^{(i)}, \tag{4.2}$$

where $\mathbf{a}=(\mathbf{a}_{\mathsf{k}})_{\mathsf{k}\in\mathbb{N}^{d+1}}$ is referred to as *core tensor* and $(r_{\mathrm{x}},r_1,\dots,r_d)$ as the *multilinear ranks* of $\mathbf{u}$.

The *hierarchical tensor format* [25], on which the variant of our scheme described in this section is based, can be interpreted as a further decomposition of $\mathbf{a}$ into tensors of order at most three. This decomposition is obtained using further matricizations of the tensor according to a recursive decomposition of the set of modes $\hat{\mathcal{I}}$ into a binary tree, which we denote by $\mathcal{D}$. For each $\alpha \in \mathcal{D}$, the rank of the corresponding matricization $T_{\mathbf{u}}^{(\alpha)}$ is denoted by $\operatorname{rank}_\alpha(\mathbf{u})$, where $\operatorname{rank}_{\hat{\mathcal{I}}}(\mathbf{u}) = 1$ for all $\mathbf{u} \neq 0$, and we set

$$\operatorname{rank}(\mathbf{u}) := \big(\operatorname{rank}_\alpha(\mathbf{u})\big)_{\alpha \in \mathcal{D}\setminus\hat{\mathcal{I}}}. \tag{4.3}$$

The hierarchical format can offer substantially more favorable complexity characteristics for large d than (4.2). The left singular vectors of the involved matricizations yield a *hierarchical singular value decomposition* [21]. We refer also to [16, 22, 24, 25, 29] for detailed expositions regarding the finitely supported case (see also [35, 36] for the related *tensor train* representation), and to [5] for analogous results for tensors in sequence spaces, with notation analogous to the present paper.

The *contractions*

$$\pi^{(i)}(\mathbf{v}) = \big(\pi^{(i)}_{\nu_i}(\mathbf{v})\big)_{\nu_i \in \mathcal{G}_i}, \quad \pi^{(i)}_\mu(\mathbf{v}) = \Big(\sum_{\nu\colon \nu_i = \mu} |\mathbf{v}_\nu|^2\Big)^{1/2}, \quad i \in \hat{\mathcal{I}}, \tag{4.4}$$

as introduced in [5], can be evaluated efficiently (without any d-dimensional summations) due to the relation

$$\pi^{(i)}_\mu(\mathbf{v}) = \Big(\sum_k |\mathbf{U}^{(i)}_{k,\mu}|^2 |\sigma^{(i)}_k|^2\Big)^{1/2}, \tag{4.5}$$

where $\sigma^{(i)}_k$ are the mode-i singular values of $\mathbf{v}$. As in our previous notation, we abbreviate $\operatorname{supp}_i \mathbf{v} := \operatorname{supp}\big(\pi^{(i)}(\mathbf{v})\big)$, $i \in \hat{\mathcal{I}}$.

4.1 Adaptive scheme

In the present case, we consider Algorithm 3.1 with the routines RECOMPRESS and COARSEN for the hierarchical format as given in [5, Rem. 15]. RECOMPRESS is based on a truncation of a hierarchical singular value decomposition up to a prescribed accuracy $\eta > 0$, which can be ensured based on the ℓ^2-norm of omitted singular values of matricizations. We denote this operation by $\hat{\mathrm{P}}_\eta$. It satisfies the quasi-optimality property [21]

$$\|\mathbf{v} - \hat{\mathrm{P}}_\eta(\mathbf{v})\| \leq \sqrt{2d-3}\inf\big\{\|\mathbf{v} - \mathbf{w}\|\colon \operatorname{rank}(\mathbf{w}) \leq \operatorname{rank}(\hat{\mathrm{P}}_\eta(\mathbf{v}))\big\}, \tag{4.6}$$

with the inequality between ranks understood componentwise.

COARSEN retains the degrees of freedom for each mode that correspond to the largest contractions (4.4). Let $(\mu^*_{i,k})_{k\in\mathbb{N}}$ be such that $(\pi^{(i)}_{\mu^*_{i,k}}(\mathbf{v}))_{k\in\mathbb{N}}$ is nonincreasing. Denote for $\Lambda \subset \mathcal{S}\times\mathcal{F}$ by $\mathrm{R}_\Lambda\,\mathbf{v}$ the array obtained by retaining all entries of $\mathbf{v}$ corresponding to indices in Λ, while replacing all others by zero. Given $\eta > 0$, we define the product set

$$\Lambda(\eta) = \bigtimes_{i\in\hat{\mathcal{I}}} \{\mu^*_{i,k}\colon k \leq N_i\},$$

where N_i, $i \in \hat{\mathcal{I}}$ are chosen to such that $\sum_{i\in\hat{\mathcal{I}}} N_i$ is minimal subject to the condition

$$\Big(\sum_{i\in\hat{\mathcal{I}}}\sum_{k>N_i} |\pi^{(i)}_{\mu^*_{i,k}}(\mathbf{v})|^2\Big)^{1/2} \leq \eta. \tag{4.7}$$

Noting that the left side in (4.7) is an upper bound for $\|\mathbf{v}-\mathrm{R}_{\Lambda(\eta)}\,\mathbf{v}\|$, we define COARSEN as a numerical realization of $\hat{\mathrm{C}}_\eta\,\mathbf{v} := \mathrm{R}_{\Lambda(\eta)}\,\mathbf{v}$, for which one has an analogous quasi-optimality property as in (4.6) with constant $\sqrt{d}$.

Furthermore, $\mathbf{A}$ as defined in (3.6) is here a *finite* sum of Kronecker product operators, which considerably simplifies the construction of the corresponding routine APPLY. More specifically, $\mathbf{A}$ is a sum of $d+1$ Kronecker product terms $\mathbf{A}_j \otimes \mathbf{M}_j$, $j = 0,\dots,d$. The action of $\mathbf{A}$ can thus increase each hierarchical rank of its argument at most by a factor of $d+1$. Consequently, APPLY can be obtained following the generic construction given in [5], provided that the operators $\mathbf{A}_j$ and $\mathbf{M}_j$ acting on each mode have the required compressibility properties. Recall that by (3.13), the infinite matrices $\mathbf{M}_j$ are bidiagonal, and hence do not require any further approximation. To use the construction of [5], we thus only need that the operators $\mathbf{A}_0,\dots,\mathbf{A}_d$ acting on the spatial variables are s^*-compressible.

Remark 4.1. *In contrast to the case considered in [6], here the Hilbert space $\mathcal{H} = V \otimes L^2(Y)$ on which the problem is posed is endowed with a cross norm. As a consequence, the isomorphism that takes $v \in \mathcal{H}$ to its coefficients $\mathbf{v} \in \ell^2(\mathcal{S}\times\mathcal{F})$ with respect to the tensor product basis is of Kronecker rank one. The original low-rank structure* (1.1) *of $A(y)$ is therefore preserved in the ℓ^2-representation* (3.6) *of the problem.*

4.2 Approximability of solutions

We illustrate next the approximability in hierarchical tensor format of a problem with finitely many parameters. In our test problem, the coefficients are piecewise constant and vary according to a spatially two-dimensional "checkerboard" geometry. That is, we consider a partition of $D :=]0,1[^2$ into congruent square subdomains D_j, $j = 1,\dots,d$. For $d = 16$, this is illustrated in Figure 1.

The low-rank approximability of such problems with respect to space-parameter separation has been studied in [2]. For the case $d = 4$ (that is, a 2×2-checkerboard), it is shown there that for each $n \in \mathbb{N}$ one can find u_k^{x}, u_k^{y} for $k = 1,\dots,n$ such that for some $c > 0$,

$$\Big\| u - \sum_{k=1}^{n} u_k^{\mathrm{x}} \otimes u_k^{\mathrm{y}} \Big\|_{L^2(Y,V)} \lesssim e^{-cn}.$$

Numerical tests indicate that an analogous estimate can be achieved also for geometries of the type shown in Figure 1 with $d = 9, 16, 25, \dots$, where c has a moderate dependence on d. Note also that for a hierarchical tensor representation, the ranks of further matricizations enter as well. We are not aware of any bounds for these additional ranks. The numerically observed decay of the corresponding singular values for different values of d (using a linear dimension tree) are shown in Figure 2.

Remark 4.2. *As we have noted for the spatially one-dimensional case in Example 2.1 in §2, for the separation between spatial and parametric variables for that case one always obtains* fixed *finite ranks that grow linearly in the number of parameters d. Note, however, that the approximation ranks corresponding to further separations among the parametric variables may then still not be uniformly bounded; see e.g. [28, Prop. 2.5] for an analysis of a simple example.*

4.3 Convergence analysis

Our complexity results aim at the following type of statements: given a certain approximability of the solution, the algorithm recovers the corresponding convergence rates without their explicit knowledge.

		$\cdots$	D_{16}
		$\iddots$	$\vdots$
$\vdots$	$\iddots$		
D_1	$\cdots$		

Figure 1: Example geometry of piecewise constant coefficients, $d = 16$.

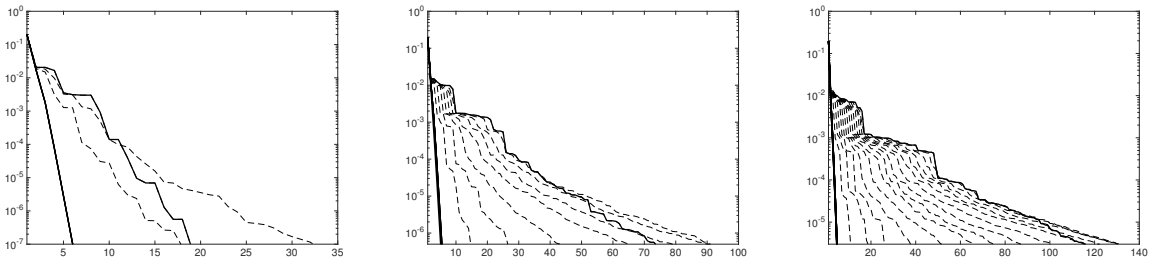

$d = 4$ $d = 9$ $d = 16$

Figure 2: Hierarchical singular values of $\mathbf{u}$, where D has $\sqrt{d} \times \sqrt{d}$-checkerboard geometry as in Figure 1. *Solid lines:* singular values of matricizations $T_{\mathbf{u}}^{(\{i\})}$ associated to $i \in \hat{\mathcal{I}}$, *dashed lines:* singular values of further matricizations in the hierarchical representation. The horizontal axes show the numbers of the decreasingly ordered singular values.

To describe these approximability properties, we now recall the definition of approximation classes to quantify the convergence of hierarchical low-rank approximations from [5]. Let $\gamma = \big(\gamma(n)\big)_{n\in\mathbb{N}_0}$ be positive and strictly increasing with $\gamma(0) = 1$ and $\gamma(n) \to \infty$ as $n \to \infty$, for $\mathbf{v} \in \ell^2(\mathcal{S}\times\mathcal{F})$ let

$$|\mathbf{v}|_{\mathcal{A}_\mathcal{H}(\gamma)} \coloneqq \sup_{r\in\mathbb{N}_0} \gamma(r) \inf_{|\mathrm{rank}(\mathbf{w})|_\infty \le r} \|\mathbf{v}-\mathbf{w}\|$$

as well as

$$\mathcal{A}_\mathcal{H}(\gamma) \coloneqq \big\{\mathbf{v}\in\ell^2(\mathcal{S}\times\mathcal{F}) : |\mathbf{v}|_{\mathcal{A}_\mathcal{H}(\gamma)} < \infty\big\}, \quad \|\mathbf{v}\|_{\mathcal{A}_\mathcal{H}(\gamma)} \coloneqq \|\mathbf{v}\| + |\mathbf{v}|_{\mathcal{A}_\mathcal{H}(\gamma)}.$$

We restrict our considerations to γ that satisfy

$$\rho_\gamma \coloneqq \sup_{n\in\mathbb{N}} \gamma(n)/\gamma(n-1) < \infty\,,$$

which corresponds to a restriction to at most exponential growth.

For an approximation $\mathbf{v}$ of bounded support to $\mathbf{u}$, the number of nonzero coefficients $\#\operatorname{supp}_i \mathbf{v}$ required in each tensor mode to achieve a certain accuracy depends on the best n-term approximability of the sequences $\pi^{(i)}(\mathbf{u})$.

This approximability by sparse sequences is quantified by the classical *approximation classes* $\mathcal{A}^s = \mathcal{A}^s(\mathcal{J})$, where $s > 0$ and $\mathcal{J}$ is a countable index set, comprised of all $\mathbf{w} \in \ell_2(\mathcal{J})$ for which the quasi-norm

$$\|\mathbf{w}\|_{\mathcal{A}^s(\mathcal{J})} \coloneqq \sup_{N\in\mathbb{N}_0} (N+1)^s \inf_{\substack{\Lambda\subset\mathcal{J}\\ \#\Lambda\le N}} \|\mathbf{w} - \mathrm{R}_\Lambda\,\mathbf{w}\| \tag{4.8}$$

is finite.

In particular, if $\pi^{(i)}(\mathbf{u}) \in \mathcal{A}^s(\mathcal{G}_i)$, these sequences can be approximated within accuracy η by finitely supported sequences with $\mathcal{O}(\eta^{-1/s})$ nonzero entries. In what follows, we do not explicitly specify the index set in the spaces $\mathcal{A}^s$ when this is clear from the context.

The observations of Section 4.2 lead us to the following benchmark assumptions for analyzing the complexity of the algorithm.

Assumptions 4.3. *For the hierarchical tensor approximation in the case* (4.1) *of d parametric variables, we assume the following:*

(i) $\pi^{(i)}(\mathbf{u}), \pi^{(i)}(\mathbf{f}) \in \mathcal{A}^s(\mathcal{G}_i)$, $i \in \hat{\mathcal{I}}$, *for an* $s > 0$.

(ii) $\mathbf{u}, \mathbf{f} \in \mathcal{A}_\mathcal{H}(\gamma)$, *where* $\gamma(n) \coloneqq e^{\bar{c}n^{1/\bar{b}}}$ *with* $\bar{b}, \bar{c} > 0$.

(iii) *The* $\mathbf{A}_j$, $j \in \hat{\mathcal{I}}$, *are* s^**-compressible for an* $s^* > s$, *and hence there exist matrices* $\mathbf{A}_{j,n}$ *with* $\alpha_{j,n}2^n$ *entries per row and column and such that* $\|\mathbf{A}_j - \mathbf{A}_{j,n}\| \le \beta_{j,n}2^{-sn}$, *and where the sequences* $\boldsymbol{\alpha}_j = (\alpha_{j,n})_{n\in\mathbb{N}}$ *and* $\boldsymbol{\beta}_j = (\beta_{j,n})_{n\in\mathbb{N}}$ *are summable.*

(iv) *The routine* RHS *satisfies, for sufficiently small* $\eta > 0$ *and* $\mathbf{f}_\eta \coloneqq \mathrm{RHS}(\eta)$,

$$\#\operatorname{supp}_i(\mathbf{f}_\eta) \lesssim \eta^{-\frac{1}{s}}\|\pi^{(i)}(\mathbf{f})\|_{\mathcal{A}^s}, \quad \|\pi^{(i)}(\mathbf{f}_\eta)\|_{\mathcal{A}^s} \lesssim \|\pi^{(i)}(\mathbf{f})\|_{\mathcal{A}^s},$$
$$|\mathrm{rank}(\mathbf{f}_\eta)|_\infty \lesssim \big(d^{-1}\ln(\|\mathbf{f}\|_{\mathcal{A}_\mathcal{H}(\gamma)}/\eta)\big)^b,$$

with hidden constants that do not depend on d, and there exists $C > 0$ independent of d such that the required number of operations is bounded by

$$C\Big(d|\mathrm{rank}(\mathbf{f}_\eta)|_\infty^3 + |\mathrm{rank}(\mathbf{f}_\eta)|_\infty \sum_{i\in\hat{\mathcal{I}}} \#\operatorname{supp}_i(\mathbf{f}_\eta)\Big).$$

We will use the above assumptions as a reference point for the scaling with respect to ε of the computational complexity. In order to also compare different parametric dimensionalities d in the complexity bounds, we additionally need a specific reference family of d-dependent problems. We introduce the following model class motivated by the considerations of Section 4.2.

Assumptions 4.4. *For the quantities in Assumptions 4.3, in addition let the following hold:*

(i) $\|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^s}$, $\|\pi^{(i)}(\mathbf{f})\|_{\mathcal{A}^s}$, $i \in \hat{\mathcal{I}}$, *and* $\|\mathbf{u}\|_{\mathcal{A}_{\mathcal{H}}(\gamma)}$, $\|\mathbf{f}\|_{\mathcal{A}_{\mathcal{H}}(\gamma)}$ *as well as* $\bar{c}^{-1}$ *grow at most polynomially in* d.

(ii) $\bar{b}$ *and* $\|\mathbf{A}_j\|$, $\|\boldsymbol{\alpha}_j\|_{\ell^1}$, $\|\boldsymbol{\beta}_j\|_{\ell^1}$ *for* $j \in \hat{\mathcal{I}}$ *are bounded independently of* d.

It needs to be emphasized that Algorithm 3.1 does not require any knowledge on the approximability of $\mathbf{u}$ stated in Assumptions 4.3 and 4.4; these merely describe a model case for complexity bounds. Recall from Remark 3.1 that Algorithm 3.1 produces $\mathbf{u}_\varepsilon$ satisfying $\|\mathbf{u} - \mathbf{u}_\varepsilon\| \leq \varepsilon$ in finitely many steps.

Theorem 4.5. *Let Assumptions 4.3 hold, let* $\alpha > 0$ *and let* $\kappa_1, \kappa_2, \kappa_3$ *in Algorithm 3.1 be chosen as*

$$\kappa_1 = \big(1 + (1+\alpha)(\sqrt{d} + \sqrt{2d-3} + \sqrt{d(2d-3)})\big)^{-1},$$
$$\kappa_2 = \sqrt{2d-3}(1+\alpha)\kappa_1, \quad \kappa_3 = \sqrt{d}(1+\sqrt{2d-3})(1+\alpha)\kappa_1.$$

Then for each $\varepsilon > 0$ *with* $\varepsilon < \varepsilon_0$, *the approximation* $\mathbf{u}_\varepsilon$ *produced by Algorithm 3.1 satisfies*

$$|\mathrm{rank}(\mathbf{u}_\varepsilon)|_\infty \leq \big(\bar{c}^{-1} \ln\big[2(\alpha\kappa_1)^{-1}\rho_\gamma \|\mathbf{u}\|_{\mathcal{A}_{\mathcal{H}}(\gamma)} \varepsilon^{-1}\big]\big)^{\bar{b}} \lesssim \left(\frac{|\ln\varepsilon| + \ln d}{\bar{c}}\right)^{\bar{b}}, \tag{4.9}$$

as well as

$$\sum_{i\in\hat{\mathcal{I}}} \#\,\mathrm{supp}_i(\mathbf{u}_\varepsilon) \lesssim d^{1+\frac{1}{s}} \Big(\sum_{i\in\hat{\mathcal{I}}} \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^s}\Big)^{\frac{1}{s}} \varepsilon^{-\frac{1}{s}}. \tag{4.10}$$

Let in addition Assumptions 4.4 hold, then there exist $c, C > 0$ *such that the number of required operations is bounded by*

$$C d^{c\ln d} |\ln\varepsilon|^{2\bar{b}} \varepsilon^{-\frac{1}{s}}, \tag{4.11}$$

where c *depends on* α, ρ, ω, s, *and* C *may additionally depend on* $\mathbf{u}$ *and* $\mathbf{f}$.

Proof. The validity of (4.9) and (4.10) follows by [5, Thm. 7], which can be immediately applied to the result of line 11 in Algorithm 3.1. Concerning (4.11), we can apply [5, Thm. 8] (with $R_i = d$ and uniform constants $\hat{C}^{(i)}_{\hat{\alpha}}$, $\hat{C}^{(i)}_{\hat{\beta}}$ and $\hat{C}^{(i)}_{\tilde{\mathbf{A}}}$ in the notation used there) to obtain, for $\mathbf{w}_\eta \coloneqq \textsc{apply}(\mathbf{v};\eta)$,

$$\#\,\mathrm{supp}_i(\mathbf{w}_\eta) \lesssim d^{1+s^{-1}} \eta^{-\frac{1}{s}} \Big(\sum_{j\in\mathcal{I}_{\mathbf{x}}} \|\pi^{(i)}(\mathbf{v})\|_{\mathcal{A}^s}\Big)^{\frac{1}{s}}, \quad \|\pi^{(i)}(\mathbf{w}_\eta)\|_{\mathcal{A}^s} \lesssim d^{1+s} \|\pi^{(i)}(\mathbf{v})\|_{\mathcal{A}^s},$$

as well as $\mathrm{rank}(\mathbf{w}_\eta) \leq (d+1)\,\mathrm{rank}(\mathbf{v})$. With these estimates, (4.11) follows exactly as in [5, Thm. 9]. □

5 Anisotropic dependence on infinitely many parameters

We now turn to the case $\mathcal{I} = \mathbb{N}$, that is, problems involving countably many parameters $(y_j)_{j\geq 1}$ that have decreasing influence as j increases. Here we consider problems of the type (1.5),

$$-\operatorname{div}\big(a(y)\nabla u(y)\big) = f, \quad a(y) = \bar{a} + \sum_{j\geq 1} y_j \theta_j, \tag{5.1}$$

for $u \in L^2(Y, V)$, with $V = H^1_0(D)$ on a domain $D \subset \mathbb{R}^m$ and $Y = [-1, 1]^{\mathbb{N}}$, under the uniform ellipticity assumption (2.2) on a.

In this section we address two principal interrelated issues. On the one hand, the performance of the previously discussed different approximation formats (realized by variants of Algorithm 3.1) for the present scenario; on the other hand, the influence of the particular type of parametric expansion system $(\theta_j)_{j\in\mathbb{N}}$ in (5.1) on the performance of the approximation method.

Any numerical scheme necessarily involves a truncation of the series for $\mathbf{A}$. Defining for each nonnegative integer M the corresponding truncation error

$$e_M := \Big\| \sum_{j>M} \mathbf{A}_j \otimes \mathbf{M}_j \Big\| \tag{5.2}$$

of replacing $\mathbf{A}$ by $\sum_{j=1}^{M} \mathbf{A}_j \otimes \mathbf{M}_j$, where $e_0 = \|\mathbf{A}\|$, the decay of e_M describes the approximability of $\mathbf{A}$. We will be concerned with algebraic rates

$$e_M \leq C M^{-S}, \quad M \in \mathbb{N}, \tag{5.3}$$

where $C, S > 0$ are fixed constants. Note that in particular, our further developments do not require summability of $(\|\theta_j\|_{L^\infty})_{j\geq 1}$ as assumed, e.g., in [14, 15].

In principle, the results of the previous section concerning a full separation of variables based on hierarchical tensor formats could be applied with any finite truncation dimension d. However, assuming (5.3), a total error of order ε requires $d(\varepsilon) \sim \varepsilon^{-1/S}$. As a consequence, due to the d-dependent quasi-optimality (4.6) of the hierarchical SVD truncation, we can only obtain a highly suboptimal complexity bound in (4.11).

Concerning low-rank decompositions, we therefore concentrate here on a more basic case, namely a separation of spatial and parametric variables in a Hilbert-Schmidt decomposition: The sequence $\mathbf{u}$ defines a Hilbert-Schmidt operator $\ell^2(\mathcal{F}) \to \ell^2(\mathcal{S})$ with singular value decomposition

$$\mathbf{u} = \sum_{k=1}^{\infty} \sigma_k \mathbf{U}^{(\mathrm{x})}_k \otimes \mathbf{U}^{(\mathrm{y})}_k, \tag{5.4}$$

where $\sigma_k \geq 0$, $\{\mathbf{U}^{(\mathrm{x})}_k\}$, $\{\mathbf{U}^{(\mathrm{y})}_k\}$ are orthonormal in $\ell^2(\mathcal{S})$ and $\ell^2(\mathcal{F})$, respectively, and

$$\Big\| \mathbf{u} - \sum_{k=1}^{r} \sigma_k \mathbf{U}^{(\mathrm{x})}_k \otimes \mathbf{U}^{(\mathrm{y})}_k \Big\|^2 = \sum_{k>r} \sigma_k^2 = \min_{\operatorname{rank}(\mathbf{w})\leq r} \|\mathbf{u} - \mathbf{w}\|^2. \tag{5.5}$$

Since this separation also occurs in any hierarchical representation, the resulting Hilbert-Schmidt rank provides a lower bound for the maximum hierarchical ranks that one can obtain in a hierarchical format involving further matricizations.

To understand the joint approximability of the infinite vectors $\mathbf{U}^{(i)}_k$, $i = \mathrm{x}, \mathrm{y}$, we consider the particular contractions defined, for $\mathbf{v} \in \ell^2(\mathcal{S} \times \mathcal{F})$, by

$$\pi^{(\mathrm{x})}(\mathbf{v}) := \Big(\Big(\sum_{\nu\in\mathcal{F}} |\mathbf{v}_{\lambda,\nu}|^2 \Big)^{1/2} \Big)_{\lambda\in\mathcal{S}}, \quad \pi^{(\mathrm{y})}(\mathbf{v}) := \Big(\Big(\sum_{\lambda\in\mathcal{S}} |\mathbf{v}_{\lambda,\nu}|^2 \Big)^{1/2} \Big)_{\nu\in\mathcal{F}}. \tag{5.6}$$

Note that $\pi_\nu^{(\mathrm{y})}(\mathbf{u})$ is uniformly proportional to the norm of the corresponding Legendre coefficient of u, that is, $\pi_\nu^{(\mathrm{y})}(\mathbf{u}) \sim \|u_\nu\|_V$.

Let $(\lambda_k^*)_{k\in\mathbb{N}}$ and $(\nu_k^*)_{k\in\mathbb{N}}$ be such that $(\pi_{\lambda_k^*}^{(\mathrm{x})}(\mathbf{v}))_{k\in\mathbb{N}}$ and $(\pi_{\nu_k^*}^{(\mathrm{y})}(\mathbf{v}))_{k\in\mathbb{N}}$ are nonincreasing, respectively. Then the singular values $\sigma_k(\mathbf{v})$ of $\mathbf{v}$ satisfy

$$\sigma_k(\mathbf{v}) \le \pi_{\lambda_k^*}^{(\mathrm{x})}(\mathbf{v}),\, \pi_{\nu_k^*}^{(\mathrm{y})}(\mathbf{v}), \quad k \in \mathbb{N}. \tag{5.7}$$

In view of our results for Example 2.3 (and the further numerical experiments of Example 5.2), we expect algebraic decay of singular values, which we quantify in terms of classes $\mathcal{A}_{\mathcal{H}}(\gamma)$ specialized to tensors of order two and to the specific sequence $\gamma(k) := (1+k)^{\bar{s}}$. This yields the approximation classes

$$\Sigma^{\bar{s}} := \Big\{\mathbf{v} \in \ell^2(\mathcal{S}\times\mathcal{F}) : \sup_{k\in\mathbb{N}}(1+k)^{\bar{s}}\Big(\sum_{j>k}\sigma_k(\mathbf{v})^2\Big)^{1/2} =: \|\mathbf{v}\|_{\Sigma^{\bar{s}}} < \infty\Big\}.$$

The approximate sparsity of the sequences $\pi^{(\mathrm{x})}(\mathbf{v})$, $\pi^{(\mathrm{y})}(\mathbf{v})$ is measured in terms of the largest $s_{\mathrm{x}}, s_{\mathrm{y}} > 0$ such that $\pi^{(\mathrm{x})}(\mathbf{v}) \in \mathcal{A}^{s_{\mathrm{x}}}(\mathcal{S})$, $\pi^{(\mathrm{y})}(\mathbf{v}) \in \mathcal{A}^{s_{\mathrm{y}}}(\mathcal{F})$ according to (4.8).

In the light of the findings for Example 2.3 in §2, showing that sparse Legendre expansions can be nearly optimal in the n-term sense for infinite parametric expansions with terms of decreasing influence, we include in the competition a method for constructing a direct sparse approximation in a tensor product basis $\{\psi_\lambda \otimes L_\nu\}_{\lambda\in\mathcal{S},\nu\in\mathcal{F}}$ of $L^2(Y,V)$. This variant of Algorithm 3.1 is similar to the scheme proposed in [20], following the approach of [9, 10].

All adaptive strategies that we consider are in essence driven by rigorous a posteriori bounds in terms of residuals and involve dynamically adapted applications of *compressed* versions of the spatial operator components $\mathbf{A}_j$. This leads in the end to a complete convergence and complexity analysis for the respective fully discrete method.

However, the efficiency of approximately evaluating residuals depends on the compressibility of the $\mathbf{A}_j$ (which is affected by the oscillatory nature of the expansion system $(\theta_j)_{j\in\mathbb{N}}$) and on the decay of the truncation errors (5.3). This is where the second issue of the influence of the parametric expansion on the performance of adaptive methods comes into play. In this regard we compare below two representative models of expansion systems: on the one hand, globally supported θ_j as they arise in Karhunen-Loève expansions for stationary processes, and on the other hand, multilevel- or wavelet type expansions of a. We show that in the former case the decay (5.3) is typically too slow to be able to fully exploit the approximability properties of the solution, a fact that has already been observed in [20]. In contrast, as shown in §5.2, a multilevel parametrization leads to near-optimal complexity bounds for the considered model classes. In this context, recall from [3] that for lognormal diffusion coefficients, in typical model examples of Matérn covariances both parametrization types are possible.

We proceed next considering the two scenarios that serve as a motivation for the assumptions that we shall make in the complexity analysis of the adaptive schemes. In particular, they demonstrate the fundamental impact of the choice of θ_j on the achievable computational complexity.

5.1 Global-type parameter expansions

To first shed some light on certain fundamental issues in approximating $\mathbf{A}$, we use the example of a simple operator which is representative of a certain class of parametrizations of the coefficient a.

We consider the following spatially one-dimensional setting with $D =]0,1[$: for a monotonically decreasing positive sequence $(c_j)_{j\in\mathbb{N}}$ with $\sum_{j\geq 1} c_j \leq \frac{1}{2}$, take $\theta_j = c_j \cos(j\pi\cdot)$, so that

$$a(y) = 1 + \sum_{j\geq 1} y_j c_j \cos(j\pi\cdot). \tag{5.8}$$

We consider this as a representative model case for the setting of globally supported θ_j that become increasingly oscillatory as $j \to \infty$.

Let us now discuss the approximability of the operator, and how it relates to the approximability of the solution. With $\alpha > \frac{1}{2}$, consider $c_j := \frac{1}{2\zeta(\frac{1}{2}+\alpha)} j^{-(\frac{1}{2}+\alpha)}$, where $t \mapsto \zeta(t)$ is the standard zeta function. Then $\|\mathbf{A}_j\| \sim c_j$,

$$e_M \lesssim M^{-\alpha+\frac{1}{2}},$$

and therefore $S = \alpha - \frac{1}{2}$. This corresponds precisely to the example considered in [20].

These considerations do not depend on a particular choice of basis of $V = H^1_0(D)$. For a complete sparse approximation of $\mathbf{A}$, however, this choice is relevant, since it determines the s^*-compressibility properties of the $\mathbf{A}_j$. In the present case (5.8) obtain a particularly simple representation using the orthonormal basis $\{\phi_k\}_{k\geq 1}$ with $\phi_k(x) := \sqrt{2}(k\pi)^{-1}\sin(k\pi x)$ of V. Then not only the $\mathbf{M}_j$, but also the $\mathbf{A}_j$ are bidiagonal, since $\mathbf{A}_0 = \mathrm{I}$ and for $j \geq 1$,

$$A_{j,kl} = c_j \int_0^1 \cos(j\pi x)\, \phi_k'(x)\, \phi_l'(x)\, dx = \frac{c_j}{2}\big(\delta_{k+l,j} + \delta_{l-k,j} + \delta_{k-l,j}\big).$$

As a consequence, each $\mathbf{A}_j$ is indeed exactly sparse and does not need to be further approximated. In other words, in this case the approximability of $\mathbf{A}$ is determined completely by e_M. In case that $\mathbf{A}_j$ need to be compressed as well, one obtains a further reduction of the compressibility of $\mathbf{A}$ relative to S, see Remark 6.3.

We also have $(\|\theta_j\|)_{j\in\mathbb{N}} \in \ell^p(\mathbb{N})$ for any $p > \frac{1}{\frac{1}{2}+\alpha}$, and as a consequence of the results in [12],

$$\pi^{(\mathrm{y})}(\mathbf{u}) \in \mathcal{A}^s \quad \text{for any } s < \alpha.$$

Hence we can take s_{y} to be any value less than α. In summary, in the present example one has

$$\alpha - \frac{1}{2} = S < s_{\mathrm{y}} < \alpha.$$

Remark 5.1. *As the above example demonstrates, for parametrizations of such type the operator approximability is in general strictly weaker than the relevant sparse approximation rates of the solution. Thus, in an adaptive scheme based on approximate residual evaluations one has to expect that the overall complexity is determined by the rate S in (5.3), and the sparsity of the solution cannot be fully exploited since $S < s_{\mathrm{y}}$. This leads us to consider below an alternate parametrization type.*

5.2 Wavelet-type expansions

Let $\Xi = \{\xi_\mu\}_{\mu\in\Lambda}$ be a system of compactly supported multilevel basis functions with $\mathrm{diam}(\mathrm{supp}(\xi_\mu)) \sim 2^{-|\mu|}$ and $\|\xi_\mu\|_{L^\infty(D)} = 1$. With $(\mu_j)_{j\geq 1}$ an enumeration of Λ by increasing level and some fixed $\alpha > 0$, we consider

$$\theta_j = c_{\mu_j}\xi_{\mu_j}, \quad \text{where } c_{\mu_j} = 2^{-\alpha|\mu_j|}. \tag{5.9}$$

To simplify notation, we also set $c_{\mu_0} \coloneqq 1$, $\xi_{\mu_0} \coloneqq \bar{a}$, and $|\mu_0| \coloneqq 0$. Note that for what follows, it would in fact suffice to assume $c_\mu \sim 2^{-\alpha|\mu|}$, with a constant that is uniform over Λ, but we assume equality to simplify the exposition. If the ξ_μ are sufficiently smooth, we then have $a(y) \in B^\alpha_{\infty,\infty}(D)$, that is, for any $y \in Y$,

$$a(y) \in C^\alpha(D) \quad \text{for } \alpha \notin \mathbb{N}. \tag{5.10}$$

Let us first compare the resulting rates in the spatially one-dimensional case $m = 1$. As shown in [4], here we have $\pi^{(\mathrm{y})}(\mathbf{u}) \in \mathcal{A}^s$ for any $s > 0$ with

$$s < s^*_{\mathrm{y}} \coloneqq \alpha,$$

where one may have $\|\pi^{(\mathrm{y})}(\mathbf{u})\|_{\mathcal{A}^s} \to \infty$ as $s \to s^*_{\mathrm{y}}$. Here we also obtain

$$e_M \lesssim M^{-\alpha},$$

that is, $S = \alpha$. In summary,

$$s_{\mathrm{y}} < s^*_{\mathrm{y}} = S,$$

which in the light of Remark 5.1 is a more favorable situation.

Regarding s^*_{x}, for sufficiently regular f and sufficiently regular wavelets ψ_λ, we also have $\pi^{(\mathrm{x})}(\mathbf{u}) \in \mathcal{A}^s$ for any

$$s < s^*_{\mathrm{x}} \coloneqq \alpha.$$

This can be seen as follows: Let $0 < s < \alpha$, and let ψ_λ be sufficiently smooth to form a Riesz basis of $H^{1+s}(D)$. Then

$$\sum_{\lambda\in\mathcal{S}} 2^{2s|\lambda|}|\pi^{(\mathrm{x})}_\lambda(\mathbf{u})|^2 = \sum_{\lambda\in\mathcal{S}}\sum_{\nu\in\mathcal{F}} 2^{2s|\lambda|}|\mathbf{u}_{\lambda\nu}|^2 \sim \int_Y \|u(y)\|^2_{H^{1+s}(D)}\, d\mu(y).$$

By (5.10) and [23, Thm. 9.1.16] we have $\|u(y)\|_{H^{1+s}} \lesssim \|f\|_{H^{-1+s}}$ uniformly in y for any $s < \alpha$ (where uniformity in y can be seen by inspection of the proof, see also [23, Thm. 9.1.8]).

When $m > 1$, we obtain in the same manner

$$s^*_{\mathrm{x}} = s^*_{\mathrm{y}} = S = \frac{\alpha}{m}, \tag{5.11}$$

where for the result concerning s^*_{x} we additionally need to assume a sufficiently smooth domain D; otherwise, one may have $s^*_{\mathrm{x}} < \alpha/m$.

For low-rank approximation, a crucial question is how s_{x} and s_{y} relate to the largest $\bar{s}$ such that $(\sigma_k(\mathbf{u}))_{k\in\mathbb{N}} \in \Sigma^{\bar{s}}$. In view of (5.7), for this $\bar{s}$ one always has $\bar{s} \geq s_{\mathrm{x}}, s_{\mathrm{y}}$. In addition, we are interested in the performance of best n-term approximation of the form (1.20), that is, the largest s such that $\mathbf{u} \in \mathcal{A}^s(\mathcal{S}\times\mathcal{F})$. We are not aware of results that give sharp statements on $\bar{s}$ and s for the present example with parametrization as in (5.9). The following representative numerical example gives an indication of what one may expect. x

Example 5.2. We consider $m = 1$ with $D =]0, 1[$, $\bar{a} = 1$, $f = 1$ and

$$\theta_j(x) = c_\alpha 2^{-\alpha\ell} h(2^\ell x - k), \quad j = 2^\ell + k$$

for $\ell \geq 0$ and $k = 0, \ldots, 2^\ell - 1$, where $h(x) = (1 - |2x - 1|)_+$ and c_α is chosen so as to ensure uniform ellipticity. In other words, the parameter is expanded in a Schauder hat

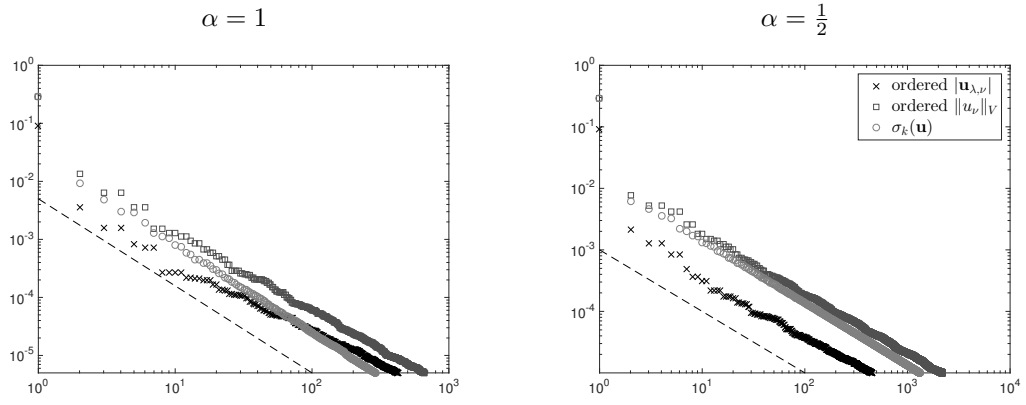

Figure 3: Absolute values of tensor product expansion coefficients $|\mathbf{u}_{\lambda,\nu}|$, Legendre coefficient norms $\|u_\nu\|_V$ (proportional to $\pi_\nu^{(\mathrm{y})}(\mathbf{u})$) and singular values $\sigma_k(\mathbf{u})$, for $\alpha = 1$ and $\alpha = \frac{1}{2}$ in Example 5.2. The dashed lines show the known asymptotic decay rates of $\|u_\nu\|_V$ according to [4].

function basis. The resulting observed decay[1] of $|\mathbf{u}_{\lambda,\nu}|$, of $\|u_\nu\|_V$ (which are proportional to $\pi_\nu^{(\mathrm{y})}(\mathbf{u})$), and of $\sigma_k(\mathbf{u})$ is shown in Figure 3. Note that in both cases $\alpha = 1$ and $\alpha = \frac{1}{2}$, the $\sigma_k(\mathbf{u})$ clearly decay at exactly the limiting rate $s_{\mathrm{y}}^* + \frac{1}{2}$ that is theoretically guaranteed for the Legendre coefficient norms $\|u_\nu\|_V$; the $\|u_\nu\|_V$ themselves approach this rate only fairly late. Thus, we can at best expect $\bar{s} = s_{\mathrm{y}}^* = \alpha$ here. Note also that the decay of $|\mathbf{u}_{\lambda,\nu}|$ is closely aligned to that of $\|u_\nu\|_V$.

In contrast to the setting of section 5.1, in the present setting (5.9), where the θ_j have multilevel structure, the e_M have sufficient decay to match the approximability of the solution. The construction of a routine APPLY that can take full advantage of this approximability, however, depends in an essential way on the compressibility of the matrices $\mathbf{A}_j$, which in turn depends on the basis that is used for V. This is considered in more detail in the next section.

6 Spatial-parametric sparse approximation

In this section we consider a version of Algorithm 3.1 that produces n-term approximations to $u \in L^2(Y, V)$ in terms of the wavelet-Legendre tensor product basis $\{\psi_\lambda \otimes L_\nu\}_{\lambda\in\mathcal{S},\nu\in\mathcal{F}}$. That is, the approximation that we seek in this case is of the form

$$u \approx u_n = \sum_{(\lambda,\nu)\in\Lambda_n} \mathbf{u}_{\lambda\nu}\psi_\lambda \otimes L_\nu, \tag{6.1}$$

where we aim to identify Λ_n that yields an error close to that of the best n-term approximation in this basis.

In this case, COARSEN performs a standard coarsening operation on a sequence, and RECOMPRESS$(\mathbf{v};\eta) \coloneqq \mathbf{v}$ for any η. The scheme thus reduces to the adaptive method of [10], which has been considered for this particular type of approximation of parametric PDEs also in [20]. The key ingredient that remains to be described is the adaptive application of $\mathbf{A}$ to representations of the form (6.1).

[1]Here we choose ψ_λ to be piecewise cubic L^2-orthonormal multiwavelets rescaled to be Riesz bases of H^1.

6.1 Adaptive operator application

In view of Remark 5.1 and the preceding discussion we concentrate in what follows on wavelet-type parametrizations as in (5.9). The compressibility of the corresponding matrices $\mathbf{A}_j$ is analyzed in Appendix A. These findings justify to proceed under the following hypothesis (see (A.15), (A.16)).

Assumptions 6.1. *Let $\{\theta_j\}_{j\geq 1}$ satisfy (5.9). There exist a $\tau > \frac{\alpha}{m}$ and matrices $\mathbf{A}_{n,j}$, $n \in \mathbb{N}_0$, with the following properties:*

(i) *the number of nonvanishing entries in each column of $\mathbf{A}_{n,j}$ does not exceed a uniform constant multiple of $\big(1+|\mu_j|^q\big)2^n$, for some $q \geq 1$;*

(ii) *one has*

$$\|\mathbf{A}_j - \mathbf{A}_{j,n}\| \lesssim c_{\mu_j} 2^{-\tau n}, \quad n \in \mathbb{N}, \tag{6.2}$$

where the hidden constant is independent of j, n and c_{μ_j} is defined as in (5.9).

Specifically, it is shown in Appendix A that the above assumptions can be realized for arbitrarily large τ by choosing the functions ξ_μ and the spatial wavelets sufficiently smooth, the latter having sufficiently many vanishing moments.

As we shall see, making use of a multilevel structure in the parametrization that leads to Assumptions 6.1, one can obtain substantially better compressibility of $\mathbf{A}$ than under the generic assumptions used in [20]. We show that under such a hypothesis the matrix $\mathbf{A}$ is s^*-compressible where $s^* < \alpha/m$ comes as close to α/m as one wishes when τ is suitably large. Consequently, the n-term approximability of $\mathbf{u}$ can be fully exploited.

Proposition 6.2. *Let Assumption 6.1 be valid. Then $\mathbf{A}$ is s^*-compressible for any*

$$s^* < \frac{\alpha}{m}\frac{2\tau}{1+2\tau}. \tag{6.3}$$

Proof. We construct approximations $\mathbf{A}_\mathbf{n}$ of $\mathbf{A}$ by choosing sequences $\mathbf{n} = (n_j)_{j\geq 0}$ of bounded support and defining $\mathbf{A}_\mathbf{n}\colon \ell^2(\mathcal{S}\times\mathcal{F}) \to \ell^2(\mathcal{S}\times\mathcal{F})$ by

$$\mathbf{A}_\mathbf{n} := \sum_{j\geq 0} \mathbf{A}_{j,n_j} \otimes \mathbf{M}_j. \tag{6.4}$$

Our aim is to find such $\mathbf{n}^J$ such that the corresponding $\mathbf{A}^J := \mathbf{A}_{\mathbf{n}^J}$ satisfy

$$\|\mathbf{A} - \mathbf{A}^J\| \lesssim J^{-2}2^{-s^*J}, \quad J \in \mathbb{N}, \tag{6.5}$$

with s^* as in the assertion, and such that $\mathbf{A}^J$ is $J^{-2}2^J$-sparse, i.e., the number of nonzero entries in the each column of $\mathbf{A}^J$ does not exceed a fixed constant multiple of $J^{-2}2^J$.

We take $L \in \mathbb{N}$ be arbitrary but fixed. Recall that we assume μ_j to be ordered by increasing level, that is, $|\mu_{j+1}| \geq |\mu_j|$. We now consider $(n_j)_{j\geq 0}$ such that $n_j = 0$ for $j > M_L := \lceil L^{2m/\alpha}2^{mL}\rceil$. Since then $e_{M_L} \lesssim L^{-2}2^{-\alpha L}$ (see (5.3), (5.11)), we obtain

$$\|\mathbf{A} - \mathbf{A}_\mathbf{n}\| \lesssim \Big\|\sum_{j=0}^{M_L} \big(\mathbf{A}_j - \mathbf{A}_{j,n_j}\big)\otimes \mathbf{M}_j\Big\| + L^{-2}2^{-\alpha L}. \tag{6.6}$$

Within each level $\ell \geq 0$, i.e., for each μ with $|\mu| = \ell$, there are only finitely many μ' with $|\mu'| = \ell$ such that $\operatorname{supp}\xi_\mu \cap \operatorname{supp}\xi_{\mu'} \neq \emptyset$. Since the images of $\mathbf{A}_j$ corresponding to ξ_{μ_j} with disjoint support are orthogonal, we obtain

$$\Big\|\sum_{j=0}^{M_L} \big(\mathbf{A}_j - \mathbf{A}_{j,n_j}\big)\otimes \mathbf{M}_j\Big\| \lesssim \sum_{\ell=0}^{\lceil L+\frac{2}{\alpha}\log_2 L\rceil} \Big(\sum_{j\colon |\mu_j|=\ell} \|\mathbf{A}_j - \mathbf{A}_{j,n_j}\|^2\Big)^{\frac{1}{2}}, \tag{6.7}$$

where the constant depends on the maximum number of wavelets of overlapping support on each level. Taking

$$n_j = n_\ell = \Big\lceil \frac{m\ell}{2\tau} + \frac{\alpha}{\tau}\Big(L + \frac{2}{\alpha}\log_2 L - \ell\Big) + \frac{1}{\tau}\log_2(1+\ell)^2 \Big\rceil$$

for μ_j of level ℓ and recalling that for such j we have $|c_{\mu_j}| = 2^{-\alpha\ell}$ gives

$$\sum_{\ell=0}^{\lceil L+\frac{2}{\alpha}\log_2 L\rceil} \Big(\sum_{j\colon |\mu_j|=\ell} \|\mathbf{A}_j - \mathbf{A}_{j,n_j}\|^2\Big)^{\frac{1}{2}} \lesssim \sum_{\ell=0}^{\lceil L+\frac{2}{\alpha}\log_2 L\rceil} 2^{\frac{m}{2}\ell} 2^{-\alpha\ell} 2^{-\tau n_\ell} \lesssim L^{-2}2^{-\alpha L}. \tag{6.8}$$

The resulting $\mathbf{A_n}$ is N_L-sparse with

$$\begin{aligned} N_L &\lesssim \sum_{\ell=0}^{\lceil L+\frac{2}{\alpha}\log_2 L\rceil} (1+\ell^q) 2^{m\ell} 2^{n_\ell} \lesssim L^{\frac{2}{\alpha}} 2^{\frac{\alpha}{\tau}L} \sum_{\ell=0}^{\lceil L+\frac{2}{\alpha}\log_2 L\rceil} (1+\ell)^{q+\frac{2}{\tau}} 2^{(1+\frac{1}{2\tau})m\ell - \frac{\alpha}{\tau}\ell} \\ &\lesssim L^{q+\frac{2(1+m)}{\alpha}} 2^{\frac{1+2\tau}{2\tau}mL}, \end{aligned} \tag{6.9}$$

where we have used $\tau > \alpha/m$.

We now fix $s^* > 0$ with $s^* < t := \frac{\alpha}{m}\frac{2\tau}{1+2\tau}$ and take $J := \lceil \frac{t}{s^*}\frac{1+2\tau}{2\tau} mL\rceil = \lceil \frac{\alpha}{s^*}L\rceil$ and $\mathbf{n}^J := \mathbf{n}$. Since then $N_L \lesssim J^{q+\frac{2(1+m)}{\alpha}} 2^{\frac{s^*}{t}J}$ we see that $N_L \lesssim J^{-2}2^J$ with a constant that depends on α, m and increases when s^* approaches t. It immediately follows from (6.8) that

$$\|\mathbf{A} - \mathbf{A}^J\| \lesssim J^{-2}2^{-s^*J} \tag{6.10}$$

with a constant depending on m. Thus $\mathbf{A}$ is s^*-compressible for any $s^* < t$. □

Remark 6.3. *In [20], where the case of globally supported θ_j as in §5.1 is considered, compressibility with only $s^* = \frac{1}{2}(\alpha - \frac{1}{2})$ for $\mathbf{A}$ with $m = 1$ is obtained when taking the compression of the individual $\mathbf{A}_j$ into account. The example given in [20] is completely analogous to the case $\theta_j \sim j^{-(\alpha+\frac{1}{2})}\cos(j\pi\cdot)$ considered in §5.1. In this case, one finds that Assumptions 6.1 are replaced by $\|\mathbf{A}_j - \mathbf{A}_{j,n}\| \lesssim j^{-(\alpha+\frac{1}{2})}2^{-\gamma n}$ with $\mathcal{O}(j(1+\log_2 j)2^n)$ entries per row and column. We comment further in Remark A.3 on how this leads to the limitation to $s^* = \frac{1}{2}(\alpha - \frac{1}{2})$.*

6.2 Complexity

Let APPLY be constructed according to Proposition 6.2 with s^* as in (6.3), and let $\mathbf{u} \in \mathcal{A}^s$ for an $s < s^*$. Then $\mathbf{f} \in \mathcal{A}^s$ and we can construct RHS satisfying

$$\#\operatorname{supp}(\mathrm{RHS}(\eta)) \lesssim \eta^{-\frac{1}{s}}\|\mathbf{f}\|_{\mathcal{A}^s}^{\frac{1}{s}}, \quad \|\mathrm{RHS}(\eta)\|_{\mathcal{A}^s} \lesssim \|\mathbf{f}\|_{\mathcal{A}^s}.$$

By the results in [10] we obtain the following complexity bound.

Theorem 6.4. *For any given $\varepsilon > 0$, the approximation $\mathbf{u}_\varepsilon$ produced by the above modification of Algorithm 3.1 operating on approximations of the form (6.1) satisfies $\|\mathbf{u} - \mathbf{u}_\varepsilon\| \le \varepsilon$, and if $\mathbf{u} \in \mathcal{A}^s$, we have*

$$\#\operatorname{supp}(\mathbf{u}_\varepsilon) \lesssim \varepsilon^{-\frac{1}{s}}\|\mathbf{u}\|_{\mathcal{A}^s}^{\frac{1}{s}}, \quad \|\mathbf{u}_\varepsilon\|_{\mathcal{A}^s} \lesssim \|\mathbf{u}\|_{\mathcal{A}^s},$$

and the number of operations is bounded up to a multiplicative constant by $1 + \varepsilon^{-\frac{1}{s}}\|\mathbf{u}\|_{\mathcal{A}^s}^{\frac{1}{s}}$.

7 Low-rank approximation

We now turn to an adaptive method for finding low-rank approximations based on the Hilbert-Schmidt decomposition (5.4),

$$\mathbf{u} = \sum_{k=1}^{\infty} \sigma_k \mathbf{U}_k^{(\mathrm{x})} \otimes \mathbf{U}_k^{(\mathrm{y})}\,.$$

On the one hand, a fast decay of the singular values σ_k means that relatively low ranks suffice to realize a given target accuracy. At the same time the generally infinitely supported mode frames $\{\mathbf{U}_k^{(\mathrm{x})}\}$, $\{\mathbf{U}_k^{(\mathrm{y})}\}$ need to be approximated by finitely supported sequences that ideally should be best n-term approximations. Hence, in order to minimize computational complexity one should intertwine adaptivity in rank and in the basis expansions. As in the scheme considered in §4, this is done by iteratively improving low-rank expansions of varying ranks, while at the same time identifying finitely supported approximations in $\ell^2(\mathcal{S})$ and $\ell^2(\mathcal{F})$, both based on approximate residual evaluations.

For the low-rank approximation, the routines RECOMPRESS and COARSEN used in Algorithm 3.1 are based on the specialization to tensors of order two of the routines described in the previous section. RECOMPRESS$(\mathbf{v};\eta)$ is a numerical realization of $\hat{\mathrm{P}}_\eta(\mathbf{v})$, which we define as the operator producing the best low-rank approximation of $\mathbf{v}$ with error at most η with respect to $\|\cdot\|$, obtained by truncating the singular value decomposition of its argument.

The routine COARSEN$(\mathbf{v};\eta)$ is constructed as in §4 based on the contractions $\pi^{(\mathrm{x})}(\mathbf{v})$, $\pi^{(\mathrm{y})}(\mathbf{v})$ defined as in (5.6). The following result differs from [5, Theorem 7], which is formulated for general hierarchical tensors, in that we now consider differing sparsity classes for the contractions $\pi^{(i)}$, $i = \mathrm{x},\mathrm{y}$. In view of the preceding discussion, it is reasonable to assume possibly different but algebraic decay for both contractions.

Theorem 7.1. *Let $\mathbf{u},\mathbf{v} \in \ell^2(\mathcal{S}\times\mathcal{F})$ with $\|\mathbf{u}-\mathbf{v}\| \le \eta$. Then for*

$$\mathbf{w}_\eta := \hat{\mathrm{C}}_{2^{3/2}(1+\alpha)\eta}\big(\hat{\mathrm{P}}_{(1+\alpha)\eta}(\mathbf{v})\big)\,, \tag{7.1}$$

we have

$$\|\mathbf{u}-\mathbf{w}_\eta\| \le \big(2+\alpha+2^{3/2}(1+\alpha)\big)\eta. \tag{7.2}$$

Moreover, when $\mathbf{u}\in\mathcal{A}(\gamma)$, $\pi^{(i)}(\mathbf{u})\in\mathcal{A}^{s_i}$, $i=\mathrm{x},\mathrm{y}$, we have

$$|\mathrm{rank}(\mathbf{w}_\eta)|_\infty \le \gamma^{-1}\big(\rho_\gamma\|\mathbf{u}\|_{\mathcal{A}(\gamma)}/(\alpha\eta)\big)\,, \qquad \|\mathbf{w}_\eta\|_{\mathcal{A}(\gamma)} \le C_1\|\mathbf{u}\|_{\mathcal{A}(\gamma)}, \tag{7.3}$$

with $C_1 = 2(1+\alpha^{-1})$ and

$$\begin{aligned} \#\,\mathrm{supp}_{\mathrm{x}}(\mathbf{w}_\eta) + \#\,\mathrm{supp}_{\mathrm{y}}(\mathbf{w}_\eta) &\le 2 + \left(\frac{2\|\pi^{(\mathrm{x})}(\mathbf{u})\|_{\mathcal{A}^{s_\mathrm{x}}}}{\alpha\eta}\right)^{\frac{1}{s_\mathrm{x}}} + \left(\frac{2\|\pi^{(\mathrm{y})}(\mathbf{u})\|_{\mathcal{A}^{s_\mathrm{y}}}}{\alpha\eta}\right)^{\frac{1}{s_\mathrm{y}}} \\ \|\pi^{(i)}(\mathbf{w}_\eta)\|_{\mathcal{A}^{s_i}} &\le C_2\|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}}, \quad i=\mathrm{x},\mathrm{y}, \end{aligned} \tag{7.4}$$

where C_2 depends on α and s_i, $i=\mathrm{x},\mathrm{y}$.

The estimates (7.1), (7.2) have been already shown in [5]. The only deviation concerns the stability estimate (7.4), which we prove in Appendix B.

To apply Algorithm 3.1 it remains to specify the approximate application of $\mathbf{A}$ by the procedure APPLY to representations of the form (5.4). As part of this procedure, we shall also use a modified routine COARSEN$_\mathrm{y}$ which operates only on the second tensor mode and leaves supp$_\mathrm{x}$ unchanged. For this routine, we shall only use the simpler statement that for any $\mathbf{v}\in\ell^2(\mathcal{S}\times\mathcal{F})$ with $\pi^{(\mathrm{y})}(\mathbf{v})\in\mathcal{A}^{s_\mathrm{y}}(\mathcal{F})$, $\mathbf{v}_\mathrm{y} := \mathrm{COARSEN}_\mathrm{y}(\mathbf{v};\eta)$ satisfies

$$\#\,\mathrm{supp}_{\mathrm{y}}(\mathbf{v}_\mathrm{y}) \lesssim \eta^{-\frac{1}{s_\mathrm{y}}}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_\mathrm{y}}}^{\frac{1}{s_\mathrm{y}}}, \quad \|\pi^{(\mathrm{y})}(\mathbf{v}_\mathrm{y})\|_{\mathcal{A}^{s_\mathrm{y}}} \lesssim \|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_\mathrm{y}}}.$$

7.1 Adaptive operator application

We now describe a specification of the more generic routine used in [5] that is tailored to exploit anisotropy in the parametrizations of parametric operators. For given any given $\eta > 0$ and finitely supported $\mathbf{v}$ we aim to construct $\mathbf{w}_\eta$ such that $\|\mathbf{A}\mathbf{v} - \mathbf{w}_\eta\| \leq \eta$. We follow here the general strategy of combining a priori knowledge on $\mathbf{A}$ with a posteriori information on $\mathbf{v}$, which is given in terms of a suitable decomposition of $\mathbf{v}$.

To that end, we first apply a preprocessing step to the finitely supported input $\mathbf{v}$ that consists of applications of RECOMPRESS and COARSEN_y. We choose for a given $\eta > 0$ the tolerances of order η in such a way that the resulting $\mathbf{v}_\eta$ satisfies

$$\|\mathbf{v} - \mathbf{v}_\eta\| \leq \frac{\eta}{2\|\mathbf{A}\|}. \tag{7.5}$$

As a consequence, for any positive s_y, $\bar{s}$ we have

$$\operatorname{rank}(\mathbf{v}_\eta) \lesssim \eta^{-\frac{1}{\bar{s}}} \|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}}, \quad \|\mathbf{v}_\eta\|_{\Sigma^{\bar{s}}} \lesssim \|\mathbf{v}\|_{\Sigma^{\bar{s}}}, \tag{7.6}$$

and

$$\#\operatorname{supp}_\text{y}(\mathbf{v}_\eta) \lesssim \eta^{-\frac{1}{s_\text{y}}} \|\pi^{(\text{y})}(\mathbf{v})\|_{\mathcal{A}^{s_\text{y}}}^{\frac{1}{s_\text{y}}}, \quad \|\pi^{(\text{y})}(\mathbf{v}_\eta)\|_{\mathcal{A}^{s_\text{y}}} \lesssim \|\pi^{(\text{y})}(\mathbf{v})\|_{\mathcal{A}^{s_\text{y}}}. \tag{7.7}$$

We then have the SVD of $\mathbf{v}_\eta$ at hand,

$$\mathbf{v}_\eta = \sum_{k=1}^{K} \sigma_k \mathbf{U}_k^{(\text{x})} \otimes \mathbf{U}_k^{(\text{y})}, \tag{7.8}$$

and set $K_p = \{2^p, \ldots, \min\{K, 2^{p+1} - 1\}\}$, for $p = 0, 1, \ldots$, $p \leq \log_2 K$. Furthermore, for $q = 0, 1, \ldots$, let $\hat{\Lambda}_q^{(\text{y})}$ be the support of the best 2^q-term approximation of $\pi^{(\text{y})}(\mathbf{v}_\eta)$. We set $\Lambda_0^{(\text{y})} := \hat{\Lambda}_0^{(\text{y})}$ and $\Lambda_q^{(\text{y})} := \hat{\Lambda}_q^{(\text{y})} \setminus \hat{\Lambda}_{q-1}^{(\text{y})}$ for $q > 0$. With this subdivision of $\operatorname{supp}_\text{y}(\mathbf{v}_\eta)$, we now define

$$\mathbf{v}_{[p,q]} := \mathrm{R}_{\mathcal{S} \times \Lambda_q^{(\text{y})}} \sum_{k \in K_p} \sigma_k \mathbf{U}_k^{(\text{x})} \otimes \mathbf{U}_k^{(\text{y})} = \sum_{k \in K_p} \sigma_k \mathbf{U}_k^{(\text{x})} \otimes \mathrm{R}_{\Lambda_q^{(\text{y})}} \mathbf{U}_k^{(\text{y})}, \tag{7.9}$$

and obtain

$$\mathbf{A}\mathbf{v}_\eta = \sum_{p,q \geq 0} \sum_{j=0}^{\infty} (\mathbf{A}_j \otimes \mathbf{M}_j) \mathbf{v}_{[p,q]} = \sum_{p,q \geq 0} \sum_{j=0}^{\infty} \sum_{k \in K_p} \sigma_k \left(\mathbf{A}_j \mathbf{U}_k^{(\text{x})}\right) \otimes \left(\mathbf{M}_j \, \mathrm{R}_{\Lambda_q^{(\text{y})}} \mathbf{U}_k^{(\text{y})}\right). \tag{7.10}$$

To construct an approximation $\mathbf{w}_\eta$ of $\mathbf{A}\mathbf{v}_\eta$ based on this decomposition, we truncate the summations over j for each p, q at some index $M_{p,q} \in \mathbb{N}$, to be determined later, and then replace the remaining terms $\mathbf{A}_j$ by compressed versions, again depending on the respective p, q. With e_M defined for nonnegative integer M as in (5.2), for any given choice of $M_{p,q}$ we have

$$\Big\| \mathbf{A}\mathbf{v} - \sum_{p,q \geq 0} \sum_{j=0}^{M_{p,q}} (\mathbf{A}_j \otimes \mathbf{M}_j) \mathbf{v}_{[p,q]} \Big\| \leq \sum_{p,q \geq 0} e_{M_{p,q}} \|\mathbf{v}_{[p,q]}\|. \tag{7.11}$$

We now choose the $M_{p,q}$ such that

$$\sum_{p,q \geq 0} e_{M_{p,q}} \|\mathbf{v}_{[p,q]}\| \leq \frac{\eta}{4}. \tag{7.12}$$

This can be done by choosing positive weights $\alpha_{p,q}$ such that $\sum_{p,q} \alpha_{p,q} = 1$, computing $\|\mathbf{v}_{[p,q]}\|$, and adjusting the $M_{p,q}$ so as to guarantee that

$$e_{M_{p,q}} \|\mathbf{v}_{[p,q]}\| \leq \eta_{p,q} := \alpha_{p,q}\eta/4. \tag{7.13}$$

We will give an a priori choice for $M_{p,q}$ below, but one may as well use e.g. the Greedy scheme proposed in [19] for selecting these values.

Next, in order to realize an approximate application of the (generally) infinite matrices $\mathbf{A}_j$ to $\mathbf{U}_k^{(\mathrm{x})}$ in (7.10) we replace $\mathbf{A}_j \mathbf{v}_{[p,q]}$ by by an approximation $\tilde{\mathbf{A}}_{j,p,q} \mathbf{v}_{[p,q]}$ using (3.10) so as to satisfy

$$\left\| \sum_{j=0}^{M_{p,q}} (\mathbf{A}_j - \tilde{\mathbf{A}}_{j,p,q}) \otimes \mathbf{M}_j \, \mathbf{v}_{[p,q]} \right\| \leq \eta_{p,q}. \tag{7.14}$$

The approximate operators $\tilde{\mathbf{A}}_{j,p,q}$ will be specified later. The sought approximation of $\mathbf{A}\mathbf{v}$ can now be obtained as

$$\mathbf{w}_\eta := \sum_{p,q \geq 0} \sum_{j=0}^{M_{p,q}} (\tilde{\mathbf{A}}_{j,p,q} \otimes \mathbf{M}_j) \mathbf{v}_{[p,q]}, \tag{7.15}$$

which by the above construction satisfies the (computable) error bound

$$\|\mathbf{A}\mathbf{v}_\eta - \mathbf{w}_\eta\| \leq \sum_{p,q \geq 0} \left(e_{M_{p,q}} \|\mathbf{v}_{[p,q]}\| + \eta_{p,q} \right) \leq \eta/2 \,, \tag{7.16}$$

so that

$$\|\mathbf{A}\mathbf{v} - \mathbf{w}_\eta\| \leq \eta. \tag{7.17}$$

In summary, the above adaptive approximation of $\mathbf{A}$ to a given finitely supported $\mathbf{v}$ involves the following steps:

APPLY : $\mathbf{v} \to \mathbf{w}_\eta$, with $\mathbf{v}$ given by its SVD

(S1): compute $\mathbf{v}_\eta \coloneqq \mathrm{COARSEN}_{\mathrm{y}}(\mathrm{RECOMPRESS}(\mathbf{v}; \eta/4\|\mathbf{A}\|); \eta/4\|\mathbf{A}\|)$ and (quasi-)sort[2] the entries of $\pi^{(\mathrm{y})}(\mathbf{v}_\eta)$ to obtain the sets $\Lambda_q^{(\mathrm{y})}$;

(S2): compute the quantities $\|\mathbf{v}_{p,q}\|$ and determine the truncation values $M_{p,q} = M_{p,q}(\eta)$;

(S3): compute the quantities $\left(\pi_\nu^{(\mathrm{x})}(\mathbf{v}_{[p,q]})\right)_{\nu \in \mathcal{S}}$ and use these to obtain the compressed matrices $\tilde{\mathbf{A}}_{j,p,q}$, using (7.9) in the assembly step (7.15).

7.2 Complexity analysis

To quantify the complexity of computing $\mathbf{w}_\eta$ in (7.15) we need to specify the properties of the operator $A(y)$ as well as the sparsity properties of the input. In view of our preceding discussion, in the scenario of primary interest, the singular values of the solution $\mathbf{u}$ as well as the best n-term approximations of the contractions $\pi^{(i)}(\mathbf{u})$, $i \in \{\mathrm{x}, \mathrm{y}\}$, exhibit algebraic decay rates. As before, these rates are denoted by $\bar{s}$ and s_{x}, s_{y}, respectively.

As indicated earlier, the complexity of the above scheme depends, in particular, on the operator approximability by truncation. We adhere to the natural assumption that

[2] As usual, to warrant a linear scaling, instead of exact ordering it suffices to perform a quasi-sorting into buckets according to some fixed exponential decay.

$e_M \leq CM^{-S}$ for some positive S, see (5.3). Since, as explained in §5.2, we have $S > s_i$, $i \in \{\mathrm{x}, \mathrm{y}\}$, in the expansion model (5.9), we confine the subsequent discussion to this setting, where $S = \frac{\alpha}{m}$. We gather next the properties upon which the complexity analysis will be based.

Assumptions 7.2. *The solution* $\mathbf{u}$ *to* (3.5) *and the matrix* $\mathbf{A}$ *have the following properties:*

(i) *One has* $\pi^{(i)}(\mathbf{u}), \pi^{(i)}(\mathbf{f}) \in \mathcal{A}^{s_i}$, $i = \mathrm{x}, \mathrm{y}$, *with* $s_\mathrm{x}, s_\mathrm{y} > 0$.

(ii) $\mathbf{u}, \mathbf{f} \in \Sigma^{\bar{s}}$ *for some* $\bar{s} \geq s_\mathrm{x}, s_\mathrm{y}$.

(iii) *There exists a constant* C *such that* $e_M \leq CM^{-S}$, $M \in \mathbb{N}$, *where* e_M *is defined by* (5.2) *and*

$$S \geq \bar{s}, s_\mathrm{y}. \tag{7.18}$$

(iv) *The representations* $\mathbf{A}_j$, $j \in \mathbb{N}$, *satisfy Assumptions 6.1 where* τ *satisfies*

$$\frac{2\tau}{1+2\tau}\frac{\alpha}{m} = \frac{2\tau}{1+2\tau}S > s_\mathrm{x}. \tag{7.19}$$

(v) *The routine* RHS *satisfies, for sufficiently small* $\eta > 0$ *and* $\mathbf{f}_\eta \coloneqq \mathrm{RHS}(\eta)$,

$$\begin{gathered}\#\operatorname{supp}_i(\mathbf{f}_\eta) \lesssim \eta^{-\frac{1}{s_i}}\|\pi^{(i)}(\mathbf{f})\|_{\mathcal{A}^{s_i}}, \quad \|\pi^{(i)}(\mathbf{f}_\eta)\|_{\mathcal{A}^{s_i}} \lesssim \|\pi^{(i)}(\mathbf{f})\|_{\mathcal{A}^{s_i}}, \quad i \in \{\mathrm{x}, \mathrm{y}\},\\ \operatorname{rank}(\mathbf{f}_\eta) \lesssim \eta^{-\frac{1}{\bar{s}}}\|\mathbf{f}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}}, \quad \|\mathbf{f}_\eta\|_{\Sigma^{\bar{s}}} \lesssim \|\mathbf{f}\|_{\Sigma^{\bar{s}}},\end{gathered}$$

and requires $\mathcal{O}\big(\eta^{-\frac{1}{\bar{s}}-\frac{1}{\min\{s_\mathrm{x}, s_\mathrm{y}\}}}\big)$ *operations.*

The main result of this section states that up to a logarithmic factor the sparsity properties of the input are preserved by the output of APPLY.

Theorem 7.3. *Suppose that the properties listed under Assumptions 7.2 hold. Then, given any finitely supported input* $\mathbf{v} \in \ell^2(\mathcal{S} \times \mathcal{F})$, *the output* $\mathbf{w}_\eta$ *produced by the procedure* APPLY, *based on the steps* (S1)–(S3), *satisfies*

$$\|\mathbf{A}\mathbf{v} - \mathbf{w}_\eta\| \leq \eta. \tag{7.20}$$

Moreover, with some $b \leq 2 + \frac{4}{s_\mathrm{x}}$ *one has*

$$\operatorname{rank}(\mathbf{w}_\eta) \lesssim \eta^{-\frac{1}{\bar{s}}}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}}(1+|\log\eta|)^b, \quad \|\mathbf{w}_\eta\|_{\Sigma^{\bar{s}}} \lesssim \|\mathbf{v}\|_{\Sigma^{\bar{s}}}(1+|\log\eta|)^{\bar{s}b}, \tag{7.21}$$

and

$$\begin{gathered}\#\operatorname{supp}_\mathrm{y}(\mathbf{w}_\eta) \lesssim \eta^{-\frac{1}{s_\mathrm{y}}}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_\mathrm{y}}}^{\frac{1}{s_\mathrm{y}}}(1+|\log\eta|)^b,\\ \|\pi^{(\mathrm{y})}(\mathbf{w}_\eta)\|_{\mathcal{A}^{s_\mathrm{y}}} \lesssim \|\mathbf{v}\|_{\mathcal{A}^{s_\mathrm{y}}}(1+|\log\eta|)^{s_\mathrm{y}b},\end{gathered} \tag{7.22}$$

as well as

$$\begin{gathered}\#\operatorname{supp}_\mathrm{x}(\mathbf{w}_\eta) \lesssim \eta^{-\frac{1}{s_\mathrm{x}}}\|\pi^{(\mathrm{x})}(\mathbf{v})\|_{\mathcal{A}^{s_\mathrm{x}}}^{\frac{1}{s_\mathrm{x}}}(1+|\log\eta|)^b,\\ \|\pi^{(\mathrm{x})}(\mathbf{w}_\eta)\|_{\mathcal{A}^{s_\mathrm{x}}} \lesssim \|\pi^{(\mathrm{x})}(\mathbf{v})\|_{\mathcal{A}^{s_\mathrm{x}}}(1+|\log\eta|)^{s_\mathrm{x}b},\end{gathered} \tag{7.23}$$

where the constants depend also on s_i, $|\log\|\pi^{(i)}(\mathbf{v})\|_{\mathcal{A}^{s_i}}|$, $i \in \{\mathrm{x}, \mathrm{y}\}$, *and on* τ *in Assumption 7.2.*

Proof. The error bound (7.20) is implied by the construction. As for the remaining claims, to assess the complexity of computing $\mathbf{w}_\eta$, given by (7.15), we estimate first $M_{p,q} = M_{p,q}(\eta)$ in terms of η. To obtain a priori bounds for the $M_{p,q}$, we use Assumptions 7.2(i) and (ii) to conclude that

$$\|\mathbf{v}_{[p,q]}\| \le 2^{-s_{\mathrm{y}} q}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}\,, \qquad \|\mathbf{v}_{[p,q]}\| \le 2^{-\bar{s}p}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}\,. \tag{7.24}$$

Then Assumption 7.2(iii) and (7.24) yield the sufficient conditions

$$M_{p,q} = M_{p,q}(\eta) \ge \left(\frac{4C\min\{2^{-\bar{s}p}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}, 2^{-s_{\mathrm{y}} q}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}\}}{\alpha_{p,q}\,\eta}\right)^{\frac{1}{S}}. \tag{7.25}$$

From (7.15) and the decomposition (7.10) we see that

$$\mathrm{rank}(\mathbf{w}_\eta) \le \sum_{p,q\ge 0} M_{p,q}2^p\,, \quad \#\,\mathrm{supp}_{\mathrm{y}}(\mathbf{w}_\eta) \le \sum_{p,q\ge 0} 3M_{p,q}2^q\,. \tag{7.26}$$

Note that the factor of 3 in the bound for $\#\,\mathrm{supp}_{\mathrm{y}}(\mathbf{w}_\eta)$ results from the bidiagonal form of the matrices $\mathbf{M}_j$; that is, the action of each of these matrices can add at most twice the number of nonzero entries in the preimage sequence, in addition to the existing ones.

The following lemma provides bounds for the right hand sides in (7.26).

Lemma 7.4. *For any fixed constant $a > 1$ choose*

$$\alpha_{p,q} = c\,\big((1+p)(1+q)\big)^{-a}, \quad c := \Big(\sum_{p,q\ge 0}\big((1+p)(1+q)\big)^{-a}\Big)^{-1}, \tag{7.27}$$

as weights in (7.25). Then for $S \ge \bar{s}$ one has

$$\sum_{p,q} 2^p M_{p,q} \lesssim \eta^{-\frac{1}{S}}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{S}}\big(1+\log_2 \#\,\mathrm{supp}_{\mathrm{y}}(\mathbf{v})\big)^{1+\frac{a}{S}} \times \big(1+\log_2\mathrm{rank}(\mathbf{v}_\eta)\big)^{1+\frac{a}{S}}\big(\mathrm{rank}(\mathbf{v}_\eta)\big)^{1-\frac{\bar{s}}{S}}, \tag{7.28}$$

where the constant depends on $a, S, \bar{s}$, on c in (7.27), and on C in Assumptions 7.2(iii).

Similarly, for $S \ge s_{\mathrm{y}}$ one has

$$\sum_{p,q} 2^q M_{p,q} \lesssim \eta^{-\frac{1}{S}}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}^{\frac{1}{S}}\big(1+\log_2\mathrm{rank}(\mathbf{v}_\eta)\big)^{1+\frac{a}{S}} \times \big(1+\log_2 \#\,\mathrm{supp}_{\mathrm{y}}(\mathbf{v}_\eta)\big)^{1+\frac{a}{S}}\big(\#\,\mathrm{supp}_{\mathrm{y}}(\mathbf{v}_\eta)\big)^{1-\frac{s_{\mathrm{y}}}{S}} \tag{7.29}$$

with similar dependencies of the constants as before, but with $\bar{s}$ replaced by s_{y}.

Proof. Bounding $M_{p,q} \lesssim \eta^{-\frac{1}{S}}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{S}}(1+q)^{\frac{a}{S}}(1+p)^{\frac{a}{S}}2^{-\frac{\bar{s}p}{S}}$, we derive

$$\sum_{p,q} 2^p M_{p,q} \lesssim \eta^{-\frac{1}{S}}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{S}}\big(1+\log_2 \#\,\mathrm{supp}_{\mathrm{y}}(\mathbf{v}_\eta)\big)^{1+\frac{a}{S}}\sum_{p}(1+p)^{\frac{a}{S}}2^{p\left(1-\frac{\bar{s}}{S}\right)}, \tag{7.30}$$

which gives (7.28), where the constant depends on $a, S, \bar{s}$ and c, C from (7.25).

To bound $\sum_{p,q} 2^q M_{p,q}$ we use $M_{p,q} \lesssim \eta^{-\frac{1}{S}}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}^{\frac{1}{S}}(1+p)^{\frac{a}{S}}(1+q)^{\frac{a}{S}}2^{-\frac{s_{\mathrm{y}} q}{S}}$ and obtain

$$\sum_{p,q} 2^q M_{p,q} \lesssim \eta^{-\frac{1}{S}}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}^{\frac{1}{S}}\big(1+\log_2\mathrm{rank}(\mathbf{v}_\eta)\big)^{1+\frac{a}{S}}\sum_{q}(1+q)^{\frac{a}{S}}2^{q\left(1-\frac{s_{\mathrm{y}}}{S}\right)}$$

which yields (7.29). □

We proceed estimating the various sparsity norms of $\mathbf{w}_\eta$. We first address rank growth and parametric sparsity, which are independent of the specific choice of $\tilde{\mathbf{A}}_{j,p,q}$. Using (7.26) and (7.28) in Lemma 7.4 together with (7.6) and (7.7), for $S \geq \bar{s}$ we obtain

$$\begin{aligned}\operatorname{rank}(\mathbf{w}_\eta) &\lesssim \eta^{-\frac{1}{S}}(1+|\log\eta|)^{2\left(1+\frac{a}{S}\right)}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{S}}\,\eta^{-\frac{1}{\bar{s}}\left(1-\frac{\bar{s}}{S}\right)}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}\left(1-\frac{\bar{s}}{S}\right)},\\ &= \eta^{-\frac{1}{\bar{s}}}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}}(1+|\log\eta|)^{2\left(1+\frac{a}{S}\right)}, \end{aligned} \tag{7.31}$$

where the constant depends also on $|\log\|\pi^{(i)}(\mathbf{v})\|_{\mathcal{A}^{s_i}}|$, $i\in\{\mathrm{x},\mathrm{y}\}$. Now suppose that N_η is an upper bound for $\operatorname{rank}(\mathbf{w}_\eta)$. To simplify the exposition, let us assume without loss of generality that $\eta\in(0,1)$. Then, by definition, one has

$$\begin{aligned}\|\mathbf{w}_\eta\|_{\Sigma^{\bar{s}}} &= \sup_{N\leq N_\eta} N^{\bar{s}} \inf_{\operatorname{rank}(\mathbf{w})\leq N}\|\mathbf{w}_\eta-\mathbf{w}\| \leq \sup_{B\in[1,\eta^{-1}]} N_{B\eta}^{\bar{s}}\|\mathbf{w}_\eta-\mathbf{w}_{B\eta}\|\\ &\leq \sup_{B\in[1,\eta^{-1}]} N_{B\eta}^{\bar{s}}\big(\|\mathbf{w}_\eta-\mathbf{A}\mathbf{v}_\eta\|+\|\mathbf{A}\mathbf{v}_\eta-\mathbf{w}_{B\eta}\|\big) \leq \sup_{B\in[1,\eta^{-1}]} 2B\eta N_{B\eta}^{\bar{s}}.\end{aligned}$$

Now we can invoke for each $B\in[1,\eta^{-1}]$ the upper bound for $\operatorname{rank}(\mathbf{v}_\eta)$ given by (7.31), and observe that the resulting bound is maximized for $B=\eta^{-1}$ when $S\geq\bar{s}$. This gives

$$\|\mathbf{w}_\eta\|_{\Sigma^{\bar{s}}} \lesssim \|\mathbf{v}\|_{\Sigma^{\bar{s}}}(1+|\log\eta|)^{2\bar{s}\left(1+\frac{a}{S}\right)}, \tag{7.32}$$

which confirms (7.21).

Similarly, using the second estimate in (7.26) and (7.29) in Lemma 7.4 and invoking (7.7) yields, for $S\geq s_{\mathrm{y}}$,

$$\#\operatorname{supp}_{\mathrm{y}}(\mathbf{w}_\eta) \lesssim \eta^{-\frac{1}{S}}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}^{\frac{1}{S}}(1+|\log_2\eta|)^{2+\frac{2a}{S}}\left(\frac{\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}}{\eta}\right)^{\frac{1}{s_{\mathrm{y}}}\left(1-\frac{s_{\mathrm{y}}}{S}\right)}. \tag{7.33}$$

By the same argument as before one obtains

$$\#\operatorname{supp}_{\mathrm{y}}(\mathbf{w}_\eta) \lesssim \eta^{-\frac{1}{s_{\mathrm{y}}}}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}^{\frac{1}{s_{\mathrm{y}}}}(1+|\log_2\eta|)^{2+\frac{2a}{S}}. \tag{7.34}$$

We can then continue as above, denoting by M_η an upper bound for $\#\operatorname{supp}_{\mathrm{y}}(\mathbf{w}_\eta)$, to argue

$$\begin{aligned}\|\pi^{(\mathrm{y})}(\mathbf{w}_\eta)\|_{\mathcal{A}^{s_{\mathrm{y}}}} &\leq \sup_{B\in[1,\eta^{-1}]} M_{B\eta}^{s_{\mathrm{y}}}\big(\|\pi^{(\mathrm{y})}(\mathbf{w}_\eta)-\pi^{(\mathrm{y})}(\mathbf{A}\mathbf{v}_\eta)\|+\|\pi^{(\mathrm{y})}(\mathbf{w}_{B\eta})-\pi^{(\mathrm{y})}(\mathbf{A}\mathbf{v}_\eta)\|\big)\\ &\leq \sup_{B\in[1,\eta^{-1}]} M_{B\eta}^{s_{\mathrm{y}}}\big(\|\mathbf{w}_\eta-\mathbf{A}\mathbf{v}\|+\|\mathbf{w}_{B\eta}-\mathbf{A}\mathbf{v}\|\big)\\ &\leq \sup_{B\in[1,\eta^{-1}]} 2B\eta M_{B\eta}^{s_{\mathrm{y}}}.\end{aligned}$$

Thus we obtain

$$\|\pi^{(\mathrm{y})}(\mathbf{w}_\eta)\|_{\mathcal{A}^{s_{\mathrm{y}}}} \lesssim \|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}(1+|\log_2\eta|)^{2s_{\mathrm{y}}\left(1+\frac{a}{S}\right)}, \tag{7.35}$$

which together with (7.34) shows (7.22).

We now turn to estimating $\#\operatorname{supp}_{\mathrm{x}}(\mathbf{w}_\eta)$ and $\|\pi^{(\mathrm{x})}(\mathbf{w}_\eta)\|_{\mathcal{A}^{s_{\mathrm{x}}}}$. To this end, we specify suitable compressed matrices $\tilde{\mathbf{A}}_{j,p,q}$ in (7.14). Denoting by $\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})_\ell$ the best ℓ-term approximation of $\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})$, we set $\Lambda_{p,q,0}=\operatorname{supp}(\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})_1)$ and

$$\Lambda_{p,q,n} := \operatorname{supp}\left(\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})_{2^n}\right)\setminus\operatorname{supp}\left(\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})_{2^{n-1}}\right),\quad n\in\mathbb{N}.$$

Note that

$$\|\mathrm{R}_{\Lambda_{p,q,n}\times\mathcal{F}}\,\mathbf{v}_{[p,q]}\| \leq \|\mathrm{R}_{\Lambda_{p,q,n}}\,\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})\| \leq 2^{-s_{\mathrm{x}}n}\|\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})\|_{\mathcal{A}^{s_{\mathrm{x}}}}.$$

To proceed we employ the following convenient reformulation of Proposition 6.2.

Remark 7.5. *Let $M \in \mathbb{N}$ and any $s^* < \frac{2\tau}{1+2\tau} S$. Then for any $J \in \mathbb{N}$ we can find $\mathbf{A}_j^J$, $j \geq 0$, such that*

$$\Big\| \sum_{j=0}^{M} \big(\mathbf{A}_j - \mathbf{A}_j^J\big) \otimes \mathbf{M}_j \Big\| \leq \beta_J 2^{-s^* J},$$

and the following holds: for each $\lambda \in \mathcal{S}$, for the sum of the number of corresponding nonzero column entries of the $\mathbf{A}_j^J$ we have the bound

$$\sum_{j=0}^{M} \#\operatorname{supp}\big(\mathbf{A}_{j,\lambda'\lambda}^J\big)_{\lambda' \in \mathcal{S}} \leq \alpha_J 2^J. \tag{7.36}$$

Here $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are positive summable sequences.

For a suitable nonnegative integer $N = N_{j,p,q,\eta}$, let $\tilde{\mathbf{A}}_{j,p,q} \coloneqq \sum_{n=0}^{N} \mathbf{A}_j^{N-n} \mathrm{R}_{\Lambda_{p,q,n}}$ and

$$\mathbf{w}_{p,q} \coloneqq \sum_{j=0}^{M_{p,q}} (\tilde{\mathbf{A}}_{j,p,q} \otimes \mathbf{M}_j) \mathbf{v}_{[p,q]}. \tag{7.37}$$

Then

$$\Big\| \mathbf{w}_{p,q} - \sum_{j=0}^{M_{p,q}} \big(\mathbf{A}_j \otimes \mathbf{M}_j\big) \mathbf{v}_{[p,q]} \Big\| = \Big\| \sum_{j=0}^{M_{p,q}} \sum_{n=0}^{N} \big((\mathbf{A}_j^{N-n} - \mathbf{A}_j)\, \mathrm{R}_{\Lambda_{p,q,n}} \otimes \mathbf{M}_j \big) \mathbf{v}_{[p,q]} \Big\|.$$

Using Remark 7.5 with $s^* = s_{\mathrm{x}}$, the right side can be estimated by

$$\begin{aligned} \sum_{n=0}^{N} \beta_{N-n} 2^{-s_{\mathrm{x}}(N-n)} 2^{-s_{\mathrm{x}} n} \|\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})\|_{\mathcal{A}^{s_{\mathrm{x}}}} + 2\|\mathbf{A}\| \sum_{n>N} 2^{-s_{\mathrm{x}} n} \|\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})\|_{\mathcal{A}^{s_{\mathrm{x}}}} \\ \lesssim 2^{-s_{\mathrm{x}} N} \|\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})\|_{\mathcal{A}^{s_{\mathrm{x}}}}, \end{aligned}$$

where the constant depends on s_{x}, $\|\mathbf{A}\|$, and $\|\boldsymbol{\beta}\|_{\ell^1}$. By (7.36), we obtain

$$\#\operatorname{supp}_{\mathrm{x}}(\mathbf{w}_{p,q}) \lesssim \sum_{n=0}^{N} 2^n \alpha_{N-n} 2^{N-n} \lesssim 2^N. \tag{7.38}$$

If we now choose the smallest N such that (7.14) holds, i.e., $2^{-s_{\mathrm{x}} N} \|\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})\|_{\mathcal{A}^{s_{\mathrm{x}}}} \lesssim \eta_{p,q}$, we obtain

$$\#\operatorname{supp}_{\mathrm{x}}(\mathbf{w}_{p,q}) \lesssim \eta_{p,q}^{-\frac{1}{s_{\mathrm{x}}}} \|\pi^{(\mathrm{x})}(\mathbf{v}_{[p,q]})\|_{\mathcal{A}^{s_{\mathrm{x}}}}^{\frac{1}{s_{\mathrm{x}}}} \lesssim \eta_{p,q}^{-\frac{1}{s_{\mathrm{x}}}} \|\pi^{(\mathrm{x})}(\mathbf{v}_{\eta})\|_{\mathcal{A}^{s_{\mathrm{x}}}}^{\frac{1}{s_{\mathrm{x}}}}.$$

Keeping the definition of $\eta_{p,q} = \alpha_{p,q}\eta$ and (7.6), (7.7) in mind, summing over p, q gives (7.23) with $b = 2\big(1 + \frac{a}{s_{\mathrm{x}}}\big) > 2\big(1 + \frac{a}{S}\big)$, where the bound on $\|\pi^{(\mathrm{y})}(\mathbf{w}_{\eta})\|_{\mathcal{A}^{s_{\mathrm{y}}}}$ follows as in (7.32) and (7.35). □

Remark 7.6. *Note that in Assumptions 7.2, we state that $S \geq \bar{s}, s_{\mathrm{y}}$ and $S > s_{\mathrm{x}}$. While other cases can in principle be considered in the same manner, the convergence rate S of the operator truncation then limits the achievable efficiency: if $S < \bar{s}$, for instance, it is easy to see that in general one can only obtain* $\operatorname{rank}(\mathbf{w}_{\eta}) \sim \mathcal{O}(\eta^{-1/S})$.

Proposition 7.7. *Under the assumptions of Theorem 7.3, let* $\mathbf{v}$ *be given by its SVD with* $r := \operatorname{rank}(\mathbf{v})$ *and* $n_i := \#\operatorname{supp}_i(\mathbf{v})$ *for* $i \in \{\mathrm{x},\mathrm{y}\}$. *Then for the number of operations* $\operatorname{ops}(\mathbf{w}_\eta)$ *required to obtain* $\mathbf{w}_\eta$, *one has*

$$\operatorname{ops}(\mathbf{w}_\eta) \lesssim (n_\mathrm{x}+n_\mathrm{y})r^2 + \Big((1+|\log\eta|)^{\frac{2a}{s_\mathrm{x}}}\eta^{-\frac{1}{s_\mathrm{x}}}\|\pi^{(\mathrm{x})}(\mathbf{v})\|_{\mathcal{A}^{s_\mathrm{x}}}^{\frac{1}{s_\mathrm{x}}} + (1+|\log\eta|)^{\frac{2a}{S}}\eta^{-\frac{1}{s_\mathrm{y}}}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_\mathrm{y}}}^{\frac{1}{s_\mathrm{y}}}\Big)\eta^{-\frac{1}{\bar{s}}}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}}. \quad (7.39)$$

For the proof, we refer to Appendix B.

Theorem 7.8. *For any* $\varepsilon > 0$, *the approximation* $\mathbf{u}_\varepsilon$ *produced by Algorithm 3.1, specified as above for approximations of the form* (1.21) *based on Hilbert-Schmidt decomposition* (5.4), *satisfies* $\|\mathbf{u}-\mathbf{u}_\varepsilon\| \le \varepsilon$.

Moreover, if Assumptions 7.2 hold, then

$$\operatorname{rank}(\mathbf{u}_\varepsilon) \lesssim \varepsilon^{-\frac{1}{\bar{s}}}\|\mathbf{u}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}}, \quad \|\mathbf{u}_\varepsilon\|_{\Sigma^{\bar{s}}} \lesssim \|\mathbf{u}\|_{\Sigma^{\bar{s}}} \quad (7.40)$$

and

$$\sum_{i\in\{\mathrm{x},\mathrm{y}\}} \#\operatorname{supp}_i(\mathbf{u}_\varepsilon) \lesssim \sum_{i\in\{\mathrm{x},\mathrm{y}\}} \varepsilon^{-\frac{1}{s_i}}\|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}}^{-\frac{1}{s_i}}, \quad \|\pi^{(i)}(\mathbf{u}_\varepsilon)\|_{\mathcal{A}^{s_i}} \lesssim \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}}. \quad (7.41)$$

The number of operations $\operatorname{ops}(\mathbf{u}_\varepsilon)$ *required to produce* ε *then satisfies*

$$\operatorname{ops}(\mathbf{u}_\varepsilon) \lesssim 1 + (1+|\log\varepsilon|)^\zeta\Big(\varepsilon^{-\frac{1}{\bar{s}}}\|\mathbf{u}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}}\Big)^2 \sum_{i\in\{\mathrm{x},\mathrm{y}\}} \varepsilon^{-\frac{1}{s_i}}\|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}}^{-\frac{1}{s_i}}, \quad (7.42)$$

where $\zeta > 0$ *depends on* s_x, $\operatorname{cond}(\mathbf{A})$, *and on the choice of* κ_1, β *in Algorithm 3.1. The constants in* (7.40), (7.41), *and* (7.42) *may also depend on* S, $\bar{s}$, s_y, *and on the further parameters of Algorithm 3.1.*

Proof. We follow the general strategy of the proofs as in [5] and in Theorem 4.5, combining the properties of the complexity reduction procedures COARSEN and RECOMPRESS with the specific adaptive operator application that we have constructed for the present problem.

The bound (7.40) and (7.41) follow from Theorem 7.1 applied to the result of line 11 in Algorithm 3.1. Note that here, the number J of inner iterations depends only on $\operatorname{cond}(\mathbf{A})$ (via ρ, ω) and on the choice of κ_1 and β. With the complexity estimates for APPLY from Theorem 7.3 and Proposition 7.7 at hand, we obtain (7.42). □

Remark 7.9. *As can be seen from the proofs Theorem 7.3 and Proposition 7.7, the numerical cost for the approximate operator application is dominated by the cost of performing orthogonalizations of the input. In particular, this leads to a quadratic dependence on the approximation ranks. The number of subsequent operations required to construct the low-rank representation of the output, however, remains proportional to the respective number of degrees of freedom.*

8 Summary and conclusions

In this work, we have studied the approximation of the solution map $Y \ni y \mapsto u(y) \in V$ in $L^2(Y,V)$ for parametric diffusion problems, where the parameter domain Y is of high or infinite dimensionality. We have considered approximations based on sparse expansions

in terms of tensor product Legendre polynomials in y, low-rank approximations based on separation of spatial and parametric variables, and higher-order tensor decompositions using further hierarchical low-rank approximation among the parametric variables. Each of these approximations can be regarded as an expansion in terms of a fixed tensor product reference basis, with the degrees of freedom entering in different nonlinear ways.

The central aim is to investigate the performance of adaptive algorithms for each type of approximation that require as input only information on the parametric operator and right hand side, and that produce rigorous and computable a posteriori error bounds. These goals are achieved, in a unified manner for all considered types of approximations, by Algorithm 3.1. Such algorithms are necessarily based on the approximate evaluation of residuals. They are also intrusive, in that they do not treat the underlying parametrized problem as a black box; however, we are not aware of any non-intrusive method with comparable properties.

Although the resulting schemes do not use a priori information on the convergence of the respective approximations of the solution map, they still produce approximations of near-optimal complexity (e.g., with respect to the number of terms or tensor ranks). The question of also guaranteeing a near-optimal operation count for constructing these approximations is more delicate: this computational complexity depends on the costs of approximating the residual, and thus on the approximability properties of the operator. In the case of low-rank approximations, due to the required orthogonalizations, the number of operations also scales at least quadratically with respect to the arising tensor ranks.

Especially keeping the latter point in mind, there is no single type of approximation that is most favorable in all of the representative model scenarios that we have considered. In the case of finitely many parameters of comparable influence, hierarchical tensor representations of u turn out to be advantageous: We can show near-optimal computational complexity on certain natural approximability classes (as in Assumptions 4.3, 4.4) for the adaptive scheme based on the method in [5].

The situation turns out to be different in the case of infinitely many parameters of decreasing influence. We have proven in §2, for a certain class of such problems, that the norms of Legendre coefficients of u have the same asymptotic decay as the singular values in its Hilbert-Schmidt decomposition. In other words, the ranks in a corresponding low-rank approximation need to increase at the same rate as the number of terms in a sparse Legendre expansion as we accuracy is increased. The numerical tests given in Figure 3 of §5 indicate that this holds true also for substantially more general problems. As a consequence, even with the careful residual evaluation given in §7, which can preserve near-optimal ranks, due to the nonlinear scaling with respect to the ranks the computational complexity of finding low-rank approximations scales worse than a direct sparse expansion as considered in §6. This conclusion remains true also for hierarchical tensor decompositions involving the same separation between spatial and parametric variables.

It is interesting to note that the results in Figure 3 also indicate that best n-term approximations with respect to a suitable fixed spatial-parametric tensor product basis have the same decay as the Legendre coefficient norms. For a fully discrete approximation one can consequently obtain the same convergence rate with respect to the number of terms as for a semi-discrete parametric Legendre expansion, which is better than the currently available estimates.

For both schemes in §7 and §5, we have seen that whether the residual can be evaluated at a cost that matches the approximability of the solution depends on the type of parameter-dependence in the diffusion coefficient. As the simple example given in §5.1 shows, in the case of diffusion coefficients expanded in terms of increasingly oscillatory

functions of global support, the complexity of the methods is in general dominated by the residual evaluation. However, in the case of diffusion coefficients whose parametrization has a multilevel structure, we have demonstrated that one can come arbitrarily close to fully exploiting the approximability of u.

References

[1] I. Babuška, F. Nobile, and R. Tempone, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM Rev., 52 (2010), pp. 317–355.

[2] M. Bachmayr and A. Cohen, *Kolmogorov widths and low-rank approximations of parametric elliptic PDEs*, Math. Comp., (2016). In press.

[3] M. Bachmayr, A. Cohen, and G. Migliorati, *Representations of gaussian random fields and approximation of elliptic pdes with lognormal coefficients.* arXiv:1603.05559.

[4] ——, *Sparse polynomial approximation of parametric elliptic PDEs. Part I: affine coefficients.* arXiv:1509.07045 [math.NA], 2015.

[5] M. Bachmayr and W. Dahmen, *Adaptive near-optimal rank tensor approximation for high-dimensional operator equations*, 2015.

[6] ——, *Adaptive low-rank methods: Problems on Sobolev spaces*, 2016.

[7] J. Beck, R. Tempone, F. Nobile, and L. Tamellini, *On the optimal polynomial approximation of stochastic PDEs by Galerkin and collocation methods*, Math. Models Methods Appl. Sci., 22 (2012), pp. 1250023, 33.

[8] A. Chkifa, A. Cohen, and C. Schwab, *High-dimensional adaptive sparse polynomial interpolation and applications to parametric PDEs*, Found. Comput. Math., 14 (2014), pp. 601–633.

[9] A. Cohen, W. Dahmen, and R. DeVore, *Adaptive wavelet methods for elliptic operator equations: Convergence rates*, Mathematics of Computation, 70 (2001), pp. 27–75.

[10] A. Cohen, W. Dahmen, and R. DeVore, *Adaptive wavelet methods II – beyond the elliptic case*, Foundations of Computational Mathematics, 2 (2002), pp. 203–245.

[11] 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), pp. 615–646.

[12] A. Cohen, R. Devore, and C. Schwab, *Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE's*, Anal. Appl. (Singap.), 9 (2011), pp. 11–47.

[13] D. Dũng, *Linear collective collocation and Galerkin methods for parametric and stochastic elliptic PDEs*, ArXiv e-prints, (2015).

[14] M. Eigel, C. J. Gittelson, C. Schwab, and E. Zander, *Adaptive stochastic Galerkin FEM*, Comput. Methods Appl. Mech. Engrg., 270 (2014), pp. 247–269.

[15] M. Eigel, M. Pfeffer, and R. Schneider, *Adaptive stochastic Galerkin FEM with hierarchical tensor representions.* Preprint 2153, WIAS Berlin, 2015.

[16] A. Falcó and W. Hackbusch, *On minimal subspaces in tensor representations*, Foundations of Computational Mathematics, 12 (2012), pp. 765–803.

[17] R. Ghanem and P. D. Spanos, *Polynomial chaos in stochastic finite elements*, J. Appl. Mech., 57 (1990), pp. 197–202.

[18] R. G. Ghanem and P. D. Spanos, *Stochastic Finite Elements: A Spectral Approach*, Dover, second ed., 2007.

[19] C. Gittelson, *An adaptive stochastic galerkin method for random elliptic operators*, Mathematics of Computation, 82 (2013), pp. 1515–1541.

[20] C. J. Gittelson, *Adaptive wavelet methods for elliptic partial differential equations with random operators*, Numer. Math., 126 (2014), pp. 471–513.

[21] L. Grasedyck, *Hierarchical singular value decomposition of tensors*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 2029–2054.

[22] L. Grasedyck, D. Kressner, and C. Tobler, *A literature survey of low-rank tensor approximation techniques*, GAMM-Mitteilungen, 36 (2013), pp. 53–78.

[23] W. Hackbusch, *Elliptic Differential Equations: Theory and Numerical Treatment*, vol. 18 of Springer Series in Computational Mathematics, Springer-Verlag Berlin Heidelberg, 1992.

[24] ——, *Tensor Spaces and Numerical Tensor Calculus*, vol. 42 of Springer Series in Computational Mathematics, Springer-Verlag Berlin Heidelberg, 2012.

[25] W. Hackbusch and S. Kühn, *A new scheme for the tensor representation*, Journal of Fourier Analysis and Applications, 15 (2009), pp. 706–722.

[26] M. Kahlbacher and S. Volkwein, *Galerkin proper orthogonal decomposition methods for parameter dependent elliptic systems*, Discuss. Math. Differ. Incl. Control Optim., 27 (2007), pp. 95–117.

[27] B. N. Khoromskij and I. Oseledets, *Quantics-TT collocation approximation of parameter-dependent and stochastic elliptic PDEs*, Comput. Methods Appl. Math., 10 (2010), pp. 376–394.

[28] B. N. Khoromskij and C. Schwab, *Tensor-structured Galerkin approximation of parametric and stochastic elliptic PDEs*, SIAM J. Sci. Comput., 33 (2011), pp. 364–385.

[29] T. G. Kolda and B. W. Bader, *Tensor decompositions and applications*, SIAM Review, 51 (2009), pp. 455–500.

[30] D. Kressner and C. Tobler, *Low-rank tensor Krylov subspace methods for parametrized linear systems*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 1288–1316.

[31] T. Lassila, A. Manzoni, A. Quarteroni, and G. Rozza, *Generalized reduced basis methods and n-width estimates for the approximation of the solution manifold of parametric PDEs*, in Analysis and numerics of partial differential equations, vol. 4 of Springer INdAM Ser., 2013, pp. 307–329.

[32] L. D. Lathauwer, B. D. Moor, and J. Vandewalle, *A multilinear singular value decomposition*, SIAM Journal on Matrix Analysis and Applications, 21 (2000), pp. 1253–1278.

[33] O. P. Le Maître and O. M. Knio, *Spectral Methods for Uncertainty Quantification*, Springer, New York, 2010.

[34] H. G. Matthies and E. Zander, *Solving stochastic systems with low-rank tensor compression*, Linear Algebra and its Applications, 436 (2012), pp. 3819–3838.

[35] I. Oseledets and E. Tyrtyshnikov, *Breaking the curse of dimensionality, or how to use SVD in many dimensions*, SIAM Journal on Scientific Computing, 31 (2009), pp. 3744–3759.

[36] I. V. Oseledets, *Tensor-train decomposition*, SIAM Journal on Scientific Computing, 33 (2011), pp. 2295–2317.

[37] G. Rozza, *Separated Representations and PGD-based Model Reduction*, Springer, Vienna, 2014, ch. Fundamentals of reduced basis method for problems governed by parametrized PDEs and applications, pp. 153–227.

[38] G. Rozza, D. B. P. Huynh, and A. T. Patera, *Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: application to transport and continuum mechanics*, Arch. Comput. Methods Eng., 15 (2008), pp. 229–275.

[39] R. Stevenson, *On the compressibility of operators in wavelet coordinates*, SIAM Journal on Mathematical Analysis, 35 (2004), pp. 1110–1132.

[40] L. R. Tucker, *Contributions to Mathematical Psychology*, Holt, Rinehart & Winston, New York, 1964, ch. The extension of factor analysis to three-dimensional matrices, pp. 109–127.

[41] ——, *Some mathematical notes on three-mode factor analysis*, Psychometrika, 31 (1966), pp. 279–311.

[42] D. Xiu, *Numerical Methods for Stochastic Computations. A Spectral Method Approach*, Princeton University Press, Princeton, NJ, 2010.

A Compressibility of parametric operators

The approximate application of the operator $\mathbf{A}$ in Algorithm 3.1 must involve, in particular, an approximate application of the spatial components $\mathbf{A}_j$. With the exception of very particular situations (such as the model case considered in §5.1), the infinite matrices $\mathbf{A}_j$ are not sparse, but contain infinitely many nonzero entries in each column. Their approximation hinges on the compressibility of these operator representations as defined in Assumptions 6.1.

These are closely related to s^*-compressibility of $\mathbf{A}_j$ as in (3.10), which here means that there exist matrices $\mathbf{A}_{j,n}$ with $\alpha_{j,n}2^n$ entries per row and column and such that

$$\|\mathbf{A}_j - \mathbf{A}_{j,n}\| \leq \beta_{j,n}2^{-sn}, \quad \text{for } 0 < s < s^*, \tag{A.1}$$

and where $\boldsymbol{\alpha}_j, \boldsymbol{\beta}_j \in \ell^1$. This is known to hold for each fixed j when employing a piecewise polynomial wavelet-type Riesz basis $\{\psi_\lambda\}_{\lambda\in\mathcal{S}}$ for V, see e.g. [9,39]. However, when insisting

on the same compressibility bound s^* for all $\mathbf{A}_j$, the quantities $\|\boldsymbol{\alpha}_j\|_{\ell^1}$ and $\|\boldsymbol{\beta}_j\|_{\ell^1}$ can in general not be expected to both remain uniformly bounded in j when the θ_j become increasingly oscillatory as in the model cases considered in §5. This dependence is reflected in Assumptions 6.1.

In the light of the discussion in §5 we confine ourselves to operators $\mathbf{A}_j$ arising from multilevel representations of the parameter of the form (5.9). To obtain this compression, we use a wavelet basis $\{\psi_\lambda\}_{\lambda\in\mathcal{S}}$ in the spatial variable.

To understand the basic mechanism, recall that the compressibility of the $\mathbf{A}_j$ is governed by the modulus of its entries $\langle\theta_j\nabla\psi_\lambda,\nabla\psi_{\lambda'}\rangle$, where $\langle\cdot,\cdot\rangle$ denotes the L^2-inner product. Specifically, recall e.g. from [9] that compression strategies for wavelet representations of an elliptic second order operator with diffusion field $c\in L^\infty(D)$ are based on bounds of the type

$$|\langle c\nabla\psi_\lambda,\nabla\psi_{\lambda'}\rangle|\lesssim\|c\|_{W^{b-m/2}(L^\infty(D))}2^{-||\lambda|-|\lambda'||b},\tag{A.2}$$

where m is the dimensionality of the spatial domain, and where $b>m/2$ depends on the smoothness of the diffusion coefficient c and of the wavelets ψ_λ. From this one derives the compression order

$$s^*=\frac{b}{m}-\frac{1}{2}.\tag{A.3}$$

Specifically, for piecewise polynomial wavelets with kth order vanishing moments and for $c\in H^{k+1}(D)$, whenever the support of the higher-level wavelet is essentially disjoint from the singular support of the lower-level wavelet, one has $b=\frac{m}{2}+k+1$. In those cases one formally gets $s^*=(k+1)/m$. A subtle analysis of the remaining cases where the singular support of the lower-level wavelet is overlapped, which prevents the highest possible order of vanishing moments from being applicable, shows that an overall compression rate $s^*>k/m$ can be obtained, which is the highest possible n-term convergence rate, see [39].

However, in our case the overall compression rate is also limited by the decay of the operator truncation error (5.2). In view of Proposition 6.2, the objective here is rather to have a compression rate for the individual components $\mathbf{A}_j$ that is as high as possible, so that one approaches the limiting value imposed by (5.2).

We now summarize the conditions on the multilevel parametric expansion functions and the spatial wavelet basis under which we will verify Assumptions 6.1. To simplify notation, let $S_\lambda:=\operatorname{supp}\psi_\lambda$.

Assumptions A.1. *For some $\gamma>0$,*

$$\xi_\mu\nabla\psi_{\lambda'}\in H^\gamma(S_\lambda),\quad \mu\in\Lambda,\ \lambda,\lambda'\in\mathcal{S},\tag{A.4}$$

and the ψ_λ have vanishing moments of order k with $k>\gamma-1$.

Note that the $\nabla\psi_\lambda$ then have vanishing moments of order $k+1>\gamma$. If $|\lambda|,|\mu|\le|\lambda'|$, using (A.4) we obtain the standard estimate

$$|\langle\xi_\mu\nabla\psi_\lambda,\nabla\psi_{\lambda'}\rangle|\le\inf_{P\in\Pi^m_{k+1}}\|\xi_\mu\nabla\psi_\lambda-P\|_{L^2(S_{\lambda'})}\|\psi_{\lambda'}\|_{L^2}\lesssim 2^{-|\lambda'|\gamma}|\xi_\mu\nabla\psi_\lambda|_{H^\gamma(S_{\lambda'})}.\tag{A.5}$$

Combining this with $|\xi_\mu\nabla\psi_\lambda|_{H^\gamma(S_{\lambda'})}\lesssim 2^{-\frac{m}{2}||\lambda|-|\lambda'||}2^{\gamma\max\{|\mu|,|\lambda|\}}$, we obtain

$$|\langle\xi_\mu\nabla\psi_\lambda,\nabla\psi_{\lambda'}\rangle|\lesssim 2^{-(\gamma+\frac{m}{2})||\lambda|-|\lambda'||}2^{\gamma(|\mu|-|\lambda|)_+}.\tag{A.6}$$

Note that the requirement (A.4) could be weakened along the lines of [39] to *piecewise* smoothness, in which case combinations of wavelets with overlapping singular supports

need to be considered separately. Since this is not essential for our purposes, to keep the exposition accessible we do not consider this in further detail.

The consequences of the estimate (A.6) depend on the relations between $|\mu|$, $|\lambda|$, and $|\lambda'|$. We distinguish the three following cases:

If $|\mu| \leq |\lambda|, |\lambda'|$, we obtain an estimate analogous to the standard case (A.2),

$$|\langle \xi_\mu \nabla \psi_\lambda, \nabla \psi_{\lambda'} \rangle| \lesssim 2^{-(\gamma + \frac{m}{2})||\lambda| - |\lambda'||}. \tag{A.7}$$

If $|\lambda| \leq |\mu| < |\lambda'|$, we obtain the modified estimate

$$|\langle \xi_\mu \nabla \psi_\lambda, \nabla \psi_{\lambda'} \rangle| \lesssim 2^{-\gamma(|\lambda'| - |\mu|)} 2^{-\frac{m}{2}||\lambda| - |\lambda'||}. \tag{A.8}$$

Note that for each fixed μ and fixed levels $|\lambda|, |\lambda'|$, there exist in this case $\mathcal{O}(2^{m(|\lambda'| - |\mu|)})$ entries that may be nonzero.

Finally, if $|\lambda|, |\lambda'| \leq |\mu|$, then for each μ, there exist $|\mu|$ indices λ such that the corresponding supports overlap, and in turn there exist $\mathcal{O}(|\mu|^2)$ pairs of λ, λ' that may give a nonvanishing entry. These entries satisfy

$$|\langle \xi_\mu \nabla \psi_\lambda, \nabla \psi_{\lambda'} \rangle| \lesssim 2^{-m|\mu|} 2^{\frac{m}{2}(|\lambda| + |\lambda'|)}. \tag{A.9}$$

Note that we do not assume any vanishing moments for ξ_μ. Hence in general not much can be gained by discarding further entries in this third case.

Our strategy for dealing with the increasingly oscillatory nature of ξ_μ as $|\mu| \to \infty$ is to retain a common compression rate s^* in (A.1) uniformly in μ without losing the decay induced by the factors c_μ, which is the result required in Assumptions 6.1. To take increasingly oscillatory behavior into account, we need to retain additional entries of the $\mathbf{A}_j$ in the cases (A.8) and (A.9). This results in the j-dependent number of nonzero entries in each row and column of the compressed operators $\mathbf{A}_{j,n}$, which is of order $\mathcal{O}((1 + |\mu_j|^q) 2^n)$.

Let $a_{\mu_j, \lambda, \lambda'}$ denote the entries of $\mathbf{A}_j$, that is,

$$a_{\mu_j, \lambda, \lambda'} = c_{\mu_j} \langle \xi_{\mu_j} \nabla \psi_\lambda, \nabla \psi_{\lambda'} \rangle.$$

Proposition A.2. *Under Assumptions A.1, the conditions in Assumptions 6.1 are satisfied, with $\tau := \gamma/m$ and $q := \max\{1, \tau^{-1}\}$, for $\mathbf{A}_{j,n}$ obtained by retaining only those entries from $\mathbf{A}_j = (a_{\mu_j, \lambda, \lambda'})_{\lambda, \lambda' \in \mathcal{S}}$ for which*

$$d_{\mu_j}(\lambda, \lambda') := \max\big\{|\lambda|, |\lambda'|\big\} - \max\big\{|\mu_j|, \min\{|\lambda|, |\lambda'|\}\big\} \leq \frac{n}{m} + \frac{\log_2(1 + |\mu_j|)}{\gamma}.$$

Proof. For $j \in \mathbb{N}$, we set $\mu := \mu_j$. In a first step, for $N > 0$, we obtain a compressed version $\mathbf{A}_j^N$ of $\mathbf{A}_j$ as follows: for the column λ, retain only those entries with row index λ' such that $d_\mu(\lambda, \lambda') \leq N$. Note that by symmetry of d_μ in its two arguments and that of $\mathbf{A}_j$, the approximation $\mathbf{A}_j^N$ is also symmetric. We now show that for

$$N = N_n := \frac{n}{m} + \frac{\log_2(1 + |\mu|)}{\gamma}, \tag{A.10}$$

we arrive at Assumptions 6.1. We use the standard weighted Schur Lemma, which in the present symmetric case yields that

$$\omega_\lambda^{-1} \sum_{\lambda' \colon d_\mu(\lambda, \lambda') > N} \omega_{\lambda'} |a_{\mu, \lambda, \lambda'}| \leq B, \quad \lambda \in \mathcal{S}, \qquad \text{implies} \quad \|\mathbf{A}_j - \mathbf{A}_j^N\| \leq B. \tag{A.11}$$

Note that $d_\mu(\lambda,\lambda') > 0$ implies that $|\lambda| > |\mu|$ or $|\lambda'| > |\mu|$. Thus, as a particular consequence of (A.6), if $d_\mu(\lambda,\lambda') > 0$ we have

$$|a_{\mu,\lambda,\lambda'}| \lesssim c_\mu 2^{-\gamma d_\mu(\lambda,\lambda')} 2^{-\frac{m}{2}||\lambda|-|\lambda'||}. \tag{A.12}$$

With the usual choice $\omega_\lambda := 2^{-\frac{m}{2}|\lambda|}$, and setting

$$I(\lambda;N) := \{\lambda' \colon d_\mu(\lambda,\lambda') > N\},$$

we obtain

$$\omega_\lambda^{-1} \sum_{\lambda' \in I(\lambda;N)} \omega_{\lambda'} |a_{\mu,\lambda,\lambda'}| \lesssim c_\mu \sum_{\lambda' \in I(\lambda;N)} 2^{-m(|\lambda'|-|\lambda|)_+} 2^{-\gamma d_\mu(\lambda,\lambda')}.$$

We now decompose $I(\lambda;N) = I_1 \cup I_2 \cup I_3 \cup I_4$, where

$$\begin{aligned} I_1 &:= \{\lambda' \in I(\lambda;N) \colon |\lambda'| \le |\mu| < |\lambda|\}, & I_2 &:= \{\lambda' \in I(\lambda;N) \colon |\lambda| \le |\mu| < |\lambda'|\}, \\ I_3 &:= \{\lambda' \in I(\lambda;N) \colon |\mu| < |\lambda'| \le |\lambda|\}, & I_4 &:= \{\lambda' \in I(\lambda;N) \colon |\mu| < |\lambda| < |\lambda'|\}. \end{aligned}$$

Since $\#(I_1) \lesssim 1 + |\mu|$,

$$\sum_{\lambda' \in I_1} 2^{-m(|\lambda'|-|\lambda|)_+} 2^{-\gamma d_\mu(\lambda,\lambda')} \lesssim (1+|\mu|) 2^{-\gamma N}. \tag{A.13}$$

Likewise, we obtain the estimates

$$\begin{aligned} \sum_{\lambda' \in I_2} 2^{-m(|\lambda'|-|\lambda|)_+} 2^{-\gamma d_\mu(\lambda,\lambda')} &\lesssim \sum_{\ell=|\mu|+N}^{\infty} \sum_{\substack{\lambda' \in I_2 \\ |\lambda'|=\ell}} 2^{-m(|\lambda'|-|\lambda|)} 2^{-\gamma\ell} \\ &\lesssim \sum_{\ell=|\mu|+N}^{\infty} 2^{-\gamma\ell} \big(2^{m(\ell-|\mu|)} 2^{-m(\ell-|\lambda|)}\big) \\ &\lesssim 2^{-\gamma N} \end{aligned}$$

and

$$\sum_{\lambda' \in I_3} 2^{-m(|\lambda'|-|\lambda|)_+} 2^{-\gamma d_\mu(\lambda,\lambda')} \lesssim \sum_{\ell=|\mu|}^{|\lambda|-N} 2^{-\gamma(|\lambda|-\ell)} \lesssim 2^{-\gamma N},$$

as well as

$$\begin{aligned} \sum_{\lambda' \in I_4} 2^{-m(|\lambda'|-|\lambda|)_+} 2^{-\gamma d_\mu(\lambda,\lambda')} &\lesssim \sum_{\ell=|\lambda|+N}^{\infty} 2^{-\gamma\ell} \sum_{\substack{\lambda' \in I_4 \\ |\lambda'|=\ell}} 2^{-m(\ell-|\lambda|)} \\ &\lesssim \sum_{\ell=|\lambda|+N}^{\infty} 2^{-\gamma\ell} \big(2^{m(\ell-|\lambda|)} 2^{-m(\ell-|\lambda|)}\big) \\ &\lesssim 2^{-\gamma N}. \end{aligned}$$

Note that $\#(I_3) \lesssim N$ and $\#(I_2), \#(I_4) \lesssim 2^{mN}$. Except for (A.13), the constants in these bounds are independent of μ. In summary, we thus obtain

$$\|\mathbf{A}_j - \mathbf{A}_j^N\| \lesssim c_\mu (1+|\mu|)\, 2^{-\gamma N} \tag{A.14}$$

with a uniform constant. As pointed out above, each column of $\mathbf{A}_j^N$ has at most $\mathcal{O}(|\mu| + 2^{mN})$ entries. With $\tau = \gamma/m$ and N_n as in (A.10), the estimate (A.14) takes the desired form

$$\|\mathbf{A}_j - \mathbf{A}_j^{N_n}\| \lesssim c_\mu 2^{-\tau n}, \tag{A.15}$$

where the number of nonzero entries can be bounded further by

$$|\mu| + 2^{mN_n} \lesssim |\mu| + 2^n(1+|\mu|)^{\frac{m}{\gamma}} \lesssim \big(1 + |\mu|^{\max\{1,m/\gamma\}}\big)2^n, \tag{A.16}$$

and thus Assumptions 6.1 are valid. □

Relation (A.15), (A.16) show that the resulting compression rate is limited by the smoothness of the expansion functions ξ_μ and the spatial wavelets ψ_λ and by the number of vanishing moments of the ψ_λ, expressed by the value γ. As Proposition 6.2 shows, with increasing γ the rate of compressibility of the complete operator $\mathbf{A}$ approaches the limiting value determined by the decay of its tail (5.2).

Remark A.3. *Proposition A.2 yields, as we have also noted in Remark 6.3, a compressibility result for multilevel-type parametrizations that is substantially more favorable than what can in general be obtained for globally supported θ_j. In the case $\theta_j \sim j^{-(\alpha+\frac{1}{2})}\cos(j\pi\cdot)$ on $D =]0,1[$ considered in §5.1, in place of* (A.6) *we obtain the analogous bound*

$$|\langle \theta_j \psi'_\lambda, \psi'_{\lambda'}\rangle| \lesssim j^{-(\alpha+\frac{1}{2})} 2^{-(\gamma+\frac{1}{2})||\lambda|-|\lambda'||} 2^{\gamma(\log_2 j - |\lambda|)_+}.$$

One may thus proceed as in the proof of Proposition A.2, with $|\mu|$ replaced by $\log_2 j$, to obtain $\mathbf{A}_{j,n}$ such that

$$\|\mathbf{A}_j - \mathbf{A}_{j,n}\| \lesssim j^{-(\alpha+\frac{1}{2})} 2^{-\gamma n}.$$

However, among the pairs of indices (λ,λ') with $|\lambda| \le \log_2 j$, we are eventually left with $\mathcal{O}(j(1+\log_2 j)2^n)$ entries per row and column.

Using these bounds to obtain a compressibility result for $\mathbf{A}$ as in Proposition 6.2, in this case we have, for $\mathbf{A_n}$ as defined in (6.4)*, the simpler estimate*

$$\|\mathbf{A} - \mathbf{A_n}\| \lesssim \sum_{j=0}^{M} \big\|\mathbf{A}_j - \mathbf{A}_{j,n_j}\big\| + M^{-(\alpha-\frac{1}{2})} \lesssim 2^{-\gamma n_0} + \sum_{j=1}^{M} j^{-(\alpha+\frac{1}{2})} 2^{-\gamma n_j} + M^{-(\alpha-\frac{1}{2})}.$$

Choosing n_j appropriately to ensure that the right hand side is of order $M^{-(\alpha-\frac{1}{2})}$ and summing the resulting total numbers of nonzero entries, as in [20] one arrives at the limiting value $s^* = \frac{1}{2}(\alpha - \frac{1}{2})$ for the compressibility of $\mathbf{A}$.

B Proofs of auxiliary results

Proof of Theorem 7.1. The estimates (7.1), (7.2) are obtained exactly as in [5]. To prove (7.4) we follow the lines of the argument in [5], and adopt the notation used there, let $N \in \mathbb{N}$ be the minimal integer such that

$$\|\mathbf{u} - \bar{C}_{\mathbf{u},N}\mathbf{u}\| \le \alpha\eta. \tag{B.1}$$

Then

$$\begin{aligned}
\alpha\eta &< \|\mathbf{u} - \bar{C}_{\mathbf{u},N-1}\mathbf{u}\| && \text{(B.2)}\\
&\le \inf_{\#\Lambda_{\mathrm{x}}+\#\Lambda_{\mathrm{y}}\le N-1} \Big\{ \|\pi^{(\mathrm{x})}(\mathbf{u}) - \mathrm{R}_{\Lambda_{\mathrm{x}}}\pi^{(\mathrm{x})}(\mathbf{u})\| + \|\pi^{(\mathrm{y})}(\mathbf{u}) - \mathrm{R}_{\Lambda_{\mathrm{y}}}\pi^{(\mathrm{y})}(\mathbf{u})\| \Big\}\\
&\le (\#\Lambda_{\mathrm{x}})^{-s_{\mathrm{x}}} \|\pi^{(\mathrm{x})}(\mathbf{u})\|_{\mathcal{A}^{s_{\mathrm{x}}}} + (\#\Lambda_{\mathrm{y}})^{-s_{\mathrm{y}}} \|\pi^{(\mathrm{y})}(\mathbf{u})\|_{\mathcal{A}^{s_{\mathrm{y}}}}. && \text{(B.3)}
\end{aligned}$$

Abbreviating $n_i := \#\Lambda_i$, $i = \mathrm{x}, \mathrm{y}$, to obtain a good upper for bound $N = N(\eta)$ from (B.1), we would like to find the minimal $n_\mathrm{x} + n_\mathrm{y}$ such that

$$\alpha\eta \le (n_\mathrm{x})^{-s_\mathrm{x}} \|\pi^{(\mathrm{x})}(\mathbf{u})\|_{\mathcal{A}^{s_\mathrm{x}}} + (n_\mathrm{y})^{-s_\mathrm{y}} \|\pi^{(\mathrm{y})}(\mathbf{u})\|_{\mathcal{A}^{s_\mathrm{y}}}, \tag{B.4}$$

to conclude that $N(\eta) \le n_\mathrm{x} + n_\mathrm{y}$. Equilibrating the upper bound yields a pair $n_\mathrm{x}, n_\mathrm{y}$ given by

$$n_i = n_i(\eta) := \left\lceil \left(2\|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}(\gamma_i)} / \alpha\eta \right)^{1/s_i} \right\rceil, \quad i = \mathrm{x}, \mathrm{y}, \tag{B.5}$$

This yields

$$\#\operatorname{supp}_\mathrm{x} \mathbf{w}_\eta + \#\operatorname{supp}_\mathrm{y} \mathbf{w}_\eta \le 2 + \left(\frac{2\|\pi^{(\mathrm{x})}(\mathbf{u})\|_{\mathcal{A}^{s_\mathrm{x}}}}{\alpha\eta} \right)^{1/s_\mathrm{x}} + \left(\frac{2\|\pi^{(\mathrm{y})}(\mathbf{u})\|_{\mathcal{A}^{s_\mathrm{y}}}}{\alpha\eta} \right)^{1/s_\mathrm{y}}, \tag{B.6}$$

which is the first inequality in (7.4).

Regarding the second inequality in (7.4), note first that

$$N \le B_i n_i, \quad i = \mathrm{x}, \mathrm{y}, \tag{B.7}$$

where B_i depend only on $s_\mathrm{x}, s_\mathrm{y}$. To bound $\|\pi^{(i)}(\mathbf{w}_\eta)\|_{\mathcal{A}^{s_i}}$ we only need to estimate

$$\sup_n n^{s_i} \inf_{\#\operatorname{supp}\hat{\mathbf{w}} \le n} \|\hat{\mathbf{w}} - \pi^{(i)}(\mathbf{w}_\eta)\|, \quad i = \mathrm{x}, \mathrm{y},$$

for $n \le \#\operatorname{supp}_i \mathbf{w}_\eta \le N$. To that end, denoting by $\hat{\mathbf{u}}_n^{(i)}$ a best n-term approximation to $\pi^{(i)}(\mathbf{u})$ and using (7.3), we obtain

$$\begin{aligned}
\inf_{\#\operatorname{supp}\hat{\mathbf{w}} \le n} \|\hat{\mathbf{w}} - \pi^{(i)}(\mathbf{w}_\eta)\| &\le \|\pi^{(i)}(\mathbf{w}_\eta) - \pi^{(i)}(\mathbf{u})\| + \|\pi^{(i)}(\mathbf{u}) - \hat{\mathbf{u}}_n^{(i)}\| \\
&\le \|\mathbf{w}_\eta - \mathbf{u}\| + n^{-s_i} \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}} \\
&\le C(\alpha)\eta + n^{-s_i} \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}} \\
&\le \frac{2C(\alpha)}{\alpha} n_i^{-s_i} \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}} + n^{-s_i} \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}},
\end{aligned}$$

where we have used (B.5) and where $C(\alpha) := \big(2 + \alpha + 2^{3/2}(1+\alpha)\big)$. Hence

$$\begin{aligned}
&n^{s_i} \inf_{\#\operatorname{supp}\hat{\mathbf{w}} \le n} \|\hat{\mathbf{w}} - \pi^{(i)}(\mathbf{w}_\eta)\| \\
&\qquad \le \left(1 + \frac{2C(\alpha)}{\alpha} \left(\frac{n}{n_i} \right)^{s_i} \right) \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}} \le \left(1 + \frac{2C(\alpha) B_i^{s_i}}{\alpha} \right) \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^{s_i}},
\end{aligned}$$

which completes the proof. □

Proof of Proposition 7.7. As we assume $\mathbf{v}$ to be given in SVD form, RECOMPRESS in step (S1) of the procedure APPLY takes only $\mathcal{O}(r)$ operations. Since it preserves the SVD form, the subsequent COARSEN using quasi-sorting takes $\mathcal{O}(r(n_\mathrm{x} + n_\mathrm{y}))$ operations (with the computation of the contractions as the dominating contribution).

In computing the quantities $\|\mathbf{v}_{p,q}\|$ and $\pi_\nu^{(\mathrm{x})}(\mathbf{v}_{[p,q]})$ in steps (S2) and (S3), we need to take into account that the vectors $\mathrm{R}_{\Lambda_q^{(\mathrm{y})}} \mathbf{U}_k^{(\mathrm{y})}$, $k \in K_p$, need no longer be orthonormal.

To this end, let $\mathbf{V} \in \mathbb{R}^{2^q \times 2^p}$ denote the matrix with columns $\mathbf{V}_k := \sigma_k \mathrm{R}_{\Lambda_q^{(\mathrm{y})}} \mathbf{U}_k^{(\mathrm{y})}$, and let

$$\hat{\mathbf{u}}_\nu = (\mathbf{U}_{k,\nu}^{(\mathrm{x})})_{k \in K_p} \in \mathbb{R}^{2^p}.$$

If $q \geq p$, we compute the Gramian $\mathbf{V}^T\mathbf{V}$, which takes $\mathcal{O}(2^{2p+q})$ operations. We then directly obtain $\|\mathbf{v}_{[p,q]}\|^2 = \mathrm{tr}(\mathbf{V}^T\mathbf{V})$. Moreover, for each given ν we can evaluate

$$|\pi_\nu^{(\mathrm{x})}(\mathbf{v}_{[p,q]})|^2 = \hat{\mathbf{u}}_\nu^T(\mathbf{V}^T\mathbf{V})\hat{\mathbf{u}}_\nu$$

using $\mathcal{O}(2^{2p})$ operations. If $p > q$, we first factorize $\mathbf{V}^T = \mathbf{Q}\mathbf{R}$, where $\mathbf{Q} \in \mathbb{R}^{2^p \times 2^q}$ has orthonormal columns and $\mathbf{R} \in \mathbb{R}^{2^q \times 2^q}$. This takes $\mathcal{O}(2^{p+2q})$ operations. In addition, we form $\mathbf{R}\mathbf{R}^T$ using $\mathcal{O}(2^{3q})$ operations. We then have $\|\mathbf{v}_{[p,q]}\|^2 = \mathrm{tr}(\mathbf{R}\mathbf{R}^T)$ and for each ν, we can evaluate $\hat{\mathbf{u}}_\nu^T\mathbf{Q}$ and subsequently $|\pi_\nu^{(\mathrm{x})}(\mathbf{v}_{[p,q]})|^2 = (\hat{\mathbf{u}}_\nu^T\mathbf{Q})(\mathbf{R}\mathbf{R}^T)(\hat{\mathbf{u}}_\nu^T\mathbf{Q})^T$ using $\mathcal{O}(2^{p+q} + 2^{2q})$ operations.

Altogether, abbreviating $r_\eta \coloneqq \mathrm{rank}(\mathbf{v}_\eta)$ and $n_{\eta,\mathrm{y}} \coloneqq \#\,\mathrm{supp}_{\mathrm{y}}(\mathbf{v}_\eta)$, the computational work required for obtaining $\|\mathbf{v}_{[p,q]}\|$ and $|\pi_\nu^{(\mathrm{x})}(\mathbf{v}_{[p,q]})|$ is of order

$$(n_{\mathrm{x}} + n_{\eta,\mathrm{y}})r_\eta^2 \leq (n_{\mathrm{x}} + n_{\mathrm{y}})r^2. \tag{B.8}$$

With these values at hand, it remains to assemble $\mathbf{w}_\eta$ in the form (7.15), which amounts to building each $\mathbf{w}_{p,q}$ as in (7.37). The action of the bidiagonal matrices $\mathbf{M}_j$, on the one hand, for each p, q and j requires 2^{p+q} operations, and the total costs for assembling the y-components of the result are therefore bounded up to a constant by

$$\begin{aligned}\sum_{p,q\geq 0} 2^{p+q}M_{p,q} &\lesssim \eta^{-\frac{1}{\bar{S}}}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{S}}}(1+\log_2(n_{\eta,\mathrm{y}}))^{\frac{a}{\bar{S}}}n_{\eta,\mathrm{y}}(1+\log_2(r_\eta))^{\frac{a}{\bar{S}}}(r_\eta)^{1-\frac{\bar{s}}{\bar{S}}}\\ &\lesssim \eta^{-\frac{1}{\bar{s}}}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}}(1+|\log\eta|)^{\frac{2a}{\bar{S}}}\eta^{-\frac{1}{s_{\mathrm{y}}}}\|\pi^{(\mathrm{y})}(\mathbf{v})\|_{\mathcal{A}^{s_{\mathrm{y}}}}^{\frac{1}{s_{\mathrm{y}}}},\end{aligned}$$

where the estimate on the right is obtained as in (7.30) and (7.31). Assembling the x-components requires the action of the approximate operators $\tilde{\mathbf{A}}_{p,q,j}$. By our construction, the combined action of $\tilde{\mathbf{A}}_{p,q,j}$, $j = 1, \ldots, M_{p,q}$, on a single vector $\mathbf{U}_k^{(\mathrm{x})}$, $k \in K_p$, takes a number of operations proportional to the resulting $\#\,\mathrm{supp}_{\mathrm{x}}(\mathbf{w}_{p,q})$. Consequently, the total number of operations for the x-components is bounded up to a constant by

$$\begin{aligned}\sum_{p,q\geq 0} 2^p \#\,\mathrm{supp}_{\mathrm{x}}(\mathbf{w}_{p,q}) &\lesssim \sum_{p,q\geq 0} 2^p(1+p)^{\frac{a}{s_{\mathrm{x}}}}(1+q)^{\frac{a}{s_{\mathrm{x}}}}\eta^{-\frac{1}{s_{\mathrm{x}}}}\|\pi^{(\mathrm{x})}(\mathbf{v}_\eta)\|_{\mathcal{A}^{s_{\mathrm{x}}}}^{\frac{1}{s_{\mathrm{x}}}}\\ &\lesssim r_\eta(1+|\log\eta|)^{\frac{2a}{s_{\mathrm{x}}}}\eta^{-\frac{1}{s_{\mathrm{x}}}}\|\pi^{(\mathrm{x})}(\mathbf{v}_\eta)\|_{\mathcal{A}^{s_{\mathrm{x}}}}^{\frac{1}{s_{\mathrm{x}}}}\\ &\lesssim \eta^{-\frac{1}{\bar{s}}}\|\mathbf{v}\|_{\Sigma^{\bar{s}}}^{\frac{1}{\bar{s}}}(1+|\log\eta|)^{\frac{2a}{s_{\mathrm{x}}}}\eta^{-\frac{1}{s_{\mathrm{x}}}}\|\pi^{(\mathrm{x})}(\mathbf{v}_\eta)\|_{\mathcal{A}^{s_{\mathrm{x}}}}^{\frac{1}{s_{\mathrm{x}}}}. \qquad \square\end{aligned}$$