

Adaptive low-rank approximations for operator equations: Accuracy control and computational complexity

Markus Bachmayr and Wolfgang Dahmen

ABSTRACT. The challenge of mastering computational tasks of enormous size tends to frequently override questioning the quality of the numerical outcome in terms of *accuracy*. By this we do not mean the accuracy within the discrete setting, which itself may also be far from evident for ill-conditioned problems or when iterative solvers are involved. By *accuracy-controlled computation* we mean the deviation of the numerical approximation from the exact solution of an underlying continuous problem in a relevant metric, which has been the initiating interest in the first place. Can the accuracy of a numerical result be rigorously certified – a question that is particularly important in the context of uncertainty quantification, when many possible sources of uncertainties interact. This is the guiding question throughout this article, which reviews recent developments of low-rank approximation methods for problems in high spatial dimensions. In particular, we highlight the role of *adaptivity* when dealing with such strongly nonlinear methods that integrate in a natural way issues of discrete and continuous accuracy.

1. Introduction

1.1. Background. Quantifiable approximation, recovery, estimation of functions of a very large and even infinite number of variables pose enormous challenges in numerous application contexts of high current interest. The discussion in this article is guided by two sources of high (spatial) dimensionality, namely

- (I) partial differential equations (PDEs) in high-dimensional phase space, and
- (II) families of PDEs depending on a large number of *parameters*, which could arise as design parameters or stem from parametrizing random coefficient fields.

The electronic Schrödinger equation for N particles or Fokker-Planck equations are typical representatives for (I). Both contain a second order diffusion operator as highest-order term, which explains the interest in considering *high-dimensional* diffusion equations on a product domain as a first model class to be studied. While from an analytical point of view this is, in principle, a very well understood problem,

2010 *Mathematics Subject Classification.* Primary 41A46, 41A63, 42C10, 65F08, 65L70, 65J10, 65N12, 65N15.

Key words and phrases. Nonlinear approximation, Tensor formats, low-rank approximation, hard and soft thresholding, high-dimensional diffusion equations, parametric PDEs, approximation classes, a posteriori error bounds, convergence and complexity.

This research was supported in part by the NSF Grant DMS 1720297, and by the SmartState and Williams-Hedberg Foundation.

the focus here is on the spatial dimension d being large, even several hundreds or thousands.

Regarding (II), an intensely studied problem class concerns *parameter dependent* families of operator equations

$$(1.1) \quad R(u; p) = 0 \quad \text{in } \Omega, \quad u|_{\partial\Omega} = 0, \quad p \in \mathcal{P},$$

where this time $\Omega \subset \mathbb{R}^m$ is a “low-dimensional” domain, typically with $m \in \{1, 2, 3\}$, but the operator depends on a parameter p that may range over a high-dimensional (or even infinite-dimensional) parameter domain \mathcal{P} . In an optimal control context, \mathcal{P} can represent a set of design parameters. Another important instance of this type of problems arises when p is actually a *random field* over some probability space modelling highly complex or micro-structured fields, such as permeability in a porous media flow. Expanding such a random field, e.g., as a *Karhunen-Loève expansion*, one arrives at a representation of p in terms of *parameter sequences* $y = (y_1, y_2, \dots) \in Y := [-1, 1]^{\mathcal{I}}$ where in general $\mathcal{I} = \mathbb{N}$. Evaluating $u(y) = u(p(y))$ for many parameter queries, computing quantities of interest of the states $u(y)$, recovering such states from given observations, or estimating the underlying parameters are typical tasks in the context of *uncertainty quantification*.

The common challenge in both problem scenarios (I) and (II) lies in recovering or approximating functions of a large or even infinite number of variables. Classical numerical concepts based on (local) mesh refinements are of very limited use, since they typically suffer from the *curse of dimensionality*, which roughly means that numerical costs grow exponentially with the spatial dimension. The perhaps most promising remedy is to exploit some intrinsic *sparsity* of solutions with respect to *a priori unknown* dictionaries or expansion systems. Technically this amounts to dealing with approximants that are parametrized in a typically rather nonlinear fashion, see Section 1.2.

Such a strategy is indirectly taken up by the following quite common approach to such spatially high-dimensional problems that has been lately attracting significant attention. One starts from a (usually fixed) standard finite difference or finite element discretization, which initially gives rise to a discrete system of equations of enormous size. Numerical tractability is then achieved by solving the (fixed) discrete problem *approximately*, restricting approximants to a low-rank tensor format, see e.g. [16, 41, 45] and the comments in Section 1.2. The choice of the initial spatial resolution and of the tensor ranks is usually based on an educated guess, with little or no provision in the solver for updating according to some target tolerances. So to speak as a tribute to the problem complexity, one is tacitly content with controlling the accuracy of the *discrete approximate* solution with respect to the *exact* solution of the *discrete problem – discrete accuracy* – but *not* with respect to the actual solution of the underlying continuous problem in a problem-relevant metric. In stark contrast, in the sequel *accuracy* or *error* control will always be understood in this latter sense with reference to the solution of the original *continuous* problem. In fact, the predictive power of models like (I) or (II) depends among other things on the ability to quantify this notion of accuracy, which is the central theme in this article.

Corresponding concepts for low-rank or tensor approximations discussed in this article target two central aims: (a) developing methods with a rigorous quantification of accuracy and (b) understanding how the entailed numerical complexity scales with increasing accuracy (with respect to the continuous solution). In fact, it

is (b) that allows one to determine in the end whether a certain solver methodology is actually appropriate or whether numerical efficiency has been compromised at the expense of any meaningful accuracy.

It is important to note that a favorable answer to (b) is tantamount to saying that solutions to the continuous problem are “well approximated by low-rank or tensor expressions”. Such basic approximability properties are discussed in Section 4 for both scenarios (I) and (II) to formulate benchmarks for the performance of solvers. Devising numerical schemes that are able to best exploit these approximability properties requires a proper balancing of both error sources, namely keeping ranks finite and discretizing corresponding tensor factors. Ideally, ranks and low-dimensional discretizations should evolve in a completely intertwined fashion, which is a highly nonlinear process. Insisting on *error-controlled* realization of such processes is the overarching objective of subsequent discussions. A central message is that this is only possible by respecting and exploiting characteristics of the continuous problem, such as intrinsic metrics, which strongly link the discrete and continuous setting. The standard paradigm “first discretize, then analyze” is thus turned around in that computation is “pretended” to take place in the *infinite*-dimensional context. In fact, a convergent iteration formulated in the infinite-dimensional setting is shown to remain convergent when executed only approximately within suitable accuracy tolerances. Discretizations are therefore never fixed beforehand, but adapt at any given stage of the “outer iteration” to such dynamic tolerances.

Therefore, it is important to note that *a priori estimates* alone are not quite sufficient, since they often involve unknown quantities or are valid under assumptions that may not hold or are impossible to check. Instead, computable rigorous *a posteriori* error control is a key element of the proposed approach that must exploit the structure of the continuous problem. This can be carried out most conveniently for those instances of (I) and (II) where, for a proper choice of a Hilbert space V , corresponding *weak formulations*

$$(1.2) \quad a(u, v) = f(v), \quad v \in V,$$

are well-posed. This means that the bilinear form $a(\cdot, \cdot) : V \times V \rightarrow \mathbb{R}$ is symmetric, bounded and coercive, i.e., (1.2) is *V-elliptic*. $a(\cdot, \cdot)$ therefore defines an equivalent scalar product on V , which means that the operator $A : V \rightarrow V'$, defined by $(Aw)(v) = a(w, v)$, $w, v \in V$, is boundedly invertible.

One reason for discussing scenario (I) in comparison to (II) is that the respective *energy spaces* V differ in an essential way, which will be seen to affect the numerical methods in an equally essential way. Nevertheless, the common ground for both (I) and (II) is the ellipticity of $a(\cdot, \cdot)$ in (1.2), or equivalently, that the induced mapping $A : V \rightarrow V'$ is an isomorphism. This means there exist constants $0 < c_a \leq C_a < \infty$ such that

$$(1.3) \quad c_a \|w\|_V \leq \|Aw\|_{V'} \leq C_a \|w\|_V, \quad w \in V.$$

Since $A(u - w) = f - Aw$ this implies in particular, that

$$(1.4) \quad c_a \|u - w\|_V \leq \|f - Aw\|_{V'} \leq C_a \|u - w\|_V, \quad w \in V,$$

i.e., errors in the trial norm are equivalent to the dual norm of corresponding residuals. It is to be stressed that “residual” stands here for the defect in the *infinite-dimensional* setting, not within a fixed finite-dimensional discretization.

Exploiting ellipticity through such error-residual relations is a cornerstone of the proposed approach. Of course, this is in principle also the starting point of adaptive finite element techniques in low dimensions. One then uses duality arguments to derive sharp computable approximations to $\|f - Aw\|_{V'}$. Since these techniques rely crucially on *localization principles*, which are natural in a finite element framework, they are, being fully subject to the curse of dimensionality, in this form *infeasible* in high-dimensional regimes. In particular, any type of mesh would be meaningless in scenario (II), where the number of relevant “activated” parameters may depend on the target accuracy and may not be known beforehand.

Therefore, we will exploit (1.4) in a different way inspired by adaptive wavelet methods [19, 20]. It hinges on first identifying a *Riesz basis* for the energy space V . This allows one to transform (1.2) into an *equivalent infinite-dimensional* system of linear equations where the matrix representation \mathbf{A} of the operator A is now an *isomorphism* from the space of square-summable sequences onto *itself*, rather than mapping V onto a *different*, less smooth space V' . Thus, errors are measured now in the same (Euclidean) norm as residuals. This transformation “preconditions” the problem already on the infinite-dimensional level. An error-controlled approximation of residuals then reduces primarily to an adaptive error-controlled approximate application of the matrix \mathbf{A} within a given low-rank or tensor format, fully intertwining low-rank approximability and spatial sparsity of low-dimensional tensor factors.

Corresponding computational realizations build essentially on recent important developments of tensor calculus, especially for hierarchical tensor formats [13, 14, 31, 35–37, 39, 43, 47, 48]. As previously indicated, a price for rendering high-dimensional problems practically tractable is to employ non-standard parametrization formats for approximants, which naturally complicates numerical processing. The following section attempts to put this into a perspective which is relevant for the remainder of the discussion.

1.2. Nonlinear approximation and parametrization formats. A numerical approximation of a function $u: \Omega \rightarrow \mathbb{R}$, with $\Omega \subset \mathbb{R}^d$, can be regarded as an algorithmic template for a computable substitute that can be parametrized to resemble u . In the most classical scenario, u is supposed to be an element of a Banach space X endowed with a norm $\|\cdot\|_X$, and one looks for increasingly better approximations to u from a *preselected* sequence $X_N \subset X$ of N -dimensional linear subspaces spanned by computationally accessible basis functions $\{\varphi_1, \varphi_2, \dots, \varphi_N\}$ such as polynomials, splines, or finite elements. Thus, for each $N \in \mathbb{N}$, one seeks a computational prescription

$$\Phi_{\text{lin}}(x; N, \mathbf{u}_1, \dots, \mathbf{u}_N) := \sum_{i=1}^N \mathbf{u}_i \varphi_i(x)$$

parametrized by the coefficient vector $\mathbf{u} = (\mathbf{u}_1, \dots, \mathbf{u}_N)^T$ which needs to be adjusted to yield small $\|\Phi_{\text{lin}}(\cdot; N, \mathbf{u}) - u\|_X$. This is also termed *linear* approximation, since the approximant is sought in the linear subspace spanned by a *preselected* set of N basis functions.

When the spaces X_n are nested, $X_n \subset X_{n+1}$, it is possible to construct an *infinite collection* $\{\varphi_1, \varphi_2, \dots\}$ such that each X_N is spanned by the first N elements of this collection and the whole collection forms a *basis* for X in an appropriate

sense, depending on the nature of the space X . For instance, *hierarchical* or *wavelet bases* fall into this category covering, in particular, the finite element setting.

This latter point of view is useful as it offers a convenient unifying framework for *adapting* an approximation procedure to each specific instance of approximants. With additional parameters $\lambda_1, \dots, \lambda_N \in \mathbb{N}$, one may now consider algorithmic templates of the form

$$\Phi_{\text{nonlin}}(x; N, \mathbf{u}_1, \dots, \mathbf{u}_N, \lambda_1, \dots, \lambda_N) := \sum_{i=1}^N \mathbf{u}_i \varphi_{\lambda_i}(x),$$

that allow one to pick those basis functions that are best suited to approximate the target u within a given budget N . Now approximations are generated from the *nonlinear sets*

$$(1.5) \quad \Sigma_N := \bigcup_{\Lambda \subset \mathbb{N}, \#\Lambda=N} \left\{ \sum_{\lambda \in \Lambda} \mathbf{u}_\lambda \varphi_\lambda : \mathbf{u}_\Lambda \in \mathbb{R}^N \right\} \subset X, \quad \mathbf{u}_\Lambda := (\mathbf{u}_\lambda)_{\lambda \in \Lambda}.$$

Since this requires choosing or *activating* N among the infinitely many basis functions, such a process is called *nonlinear approximation*. In practical realizations one typically generates an increasing sequence of activated index sets Λ_n , where the choice of Λ_{n+1} exploits information gained from the preceding stage represented by Λ_n . This form of nonlinear approximation is referred to as *adaptive approximation*. Its rigorous foundation very much relies on being able to quantify the accuracy obtained at a previous stage n . It is tantamount to asking for *certified a posteriori error bounds*, which is a recurrent theme throughout this article. This issue is precisely what requires intertwining the discrete and continuous setting and exploiting, in particular, intrinsic problem metrics.

These concepts have been quite successful and are by now fairly well understood for low spatial dimensions, typically $d \leq 3$. In particular, the performance of such schemes is essentially governed by *Besov regularity*, that is, smoothness in L_p spaces where p is allowed to be less than one, see [27]. In correspondence to the isotropic nature of such classical regularity notions, improved accuracy is achieved by increasing *spatial localization*. Isotropic localization, however, is precisely what causes the curse of dimensionality.

In cases with moderately large d (depending on the problem, for instance, the low two-digit regime), a successful remedy is based on employing *product-type* basis functions $\varphi_\lambda(x) = \varphi_{\lambda_1}^{(1)}(x_1) \cdots \varphi_{\lambda_d}^{(d)}(x_d)$, where $x = (x_1, \dots, x_d)$, each x_i belongs to \mathbb{R}^m for small m , and the $\varphi_j^{(i)}$ form a low-dimensional basis of the type discussed above. Associating with each index λ_i a *scale* denoted by $|\lambda_i|$, typically indicating that $\text{diam}(\text{supp } \varphi_{\lambda_i}^{(i)}) \sim 2^{-|\lambda_i|}$, one considers, for instance, expansions of the form

$$\sum_{|\lambda_1| + \dots + |\lambda_d| \leq L} \mathbf{u}_\lambda \varphi_\lambda(x).$$

In other words, the a priori activated summands never involve simultaneously many fine scale basis functions, and the complexity of such expressions scales like $2^L L^{d-1}$. This is the concept of *sparse grid* or *hyperbolic cross approximation*. One thus gives up on isotropic localization. For such approximations to provide high accuracy one has to demand, however, correspondingly high regularity, which in this context means controlled *mixed partial derivatives* of appropriate order. While in this form

such schemes are still linear, they lend themselves to nonlinear versions in natural ways, see [51].

However, the principle of adaptively activating basis functions from any *pre-elected* basis or dictionary may no longer suffice for large d to warrant affordable computational cost in an accuracy controlled approximation. Instead one has to resort to yet stronger notions of nonlinearity that may allow one to better capture a hidden sparsity related to the problem at hand, thus also asking for a deeper understanding of the underlying continuous model. In other words, the actual “building blocks” for representing the target function in the spirit of classical harmonic analysis, should *no longer be pre-determined* but rather depend on the problem at hand.

A particularly flexible parametrization format that has been lately attracting considerable attention are *deep neural networks* (DNN) of the form

$$(1.6) \quad \Phi_{\text{nn}}(x; (\mathbf{v}_{\ell,i})) := \sigma_1 \left(\sum_{i_1} \mathbf{v}_{1,i_1} \sigma_{2,i_1} \left(\cdots \sigma_{L,i_{L-1}} \left(\sum_{i_L} \mathbf{v}_{L,i_L} x_{i_L} \right) \cdots \right) \right).$$

That is, the mapping taking possibly high dimensional inputs x into an approximation to $u(x)$ is a concatenation of affine maps followed by a componentwise nonlinear map, called “activation function”. This approach is based on the presumption that in many application scenarios target objects are close to what can be covered by such parametrizations. Although existence of efficient approximations of this type has been shown for various cases of interest [49, 56], guaranteeing convergence and error control in the actual computation of such approximations currently remains a wide open problem.

The type of approximation that we focus on here is somewhat “less nonlinear” than (1.6), but of sufficiently strong nonlinearity to make it suitable for a range of problems with very large d . The principle is easiest to explain for $d = 2$: assume that we have product basis functions $\varphi_{i,j}(x) = \varphi_i^{(1)}(x_1) \varphi_j^{(2)}(x_2)$, $i, j \in \mathbb{N}$. We then consider adaptive low-rank approximations of the form

$$(1.7) \quad \begin{aligned} \Phi_{\text{lr}} \left(x; r, N_1, N_2, \mathbf{u}_{\lambda_1, \dots, \lambda_{N_1}}^1, \mathbf{u}_{\mu_1, \dots, \mu_{N_2}}^2 \right) &:= \sum_{k=1}^r \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \mathbf{u}_{k,i}^1 \mathbf{u}_{k,j}^2 \varphi_{\lambda_i}^{(1)}(x_1) \varphi_{\mu_j}^{(2)}(x_2) \\ &= \sum_{k=1}^r \left(\sum_{i=1}^{N_1} \mathbf{u}_{k,i}^1 \varphi_{\lambda_i}^{(1)}(x_1) \right) \left(\sum_{j=1}^{N_2} \mathbf{u}_{k,j}^2 \varphi_{\mu_j}^{(2)}(x_2) \right) =: \sum_{k=1}^r \mathbf{u}_k^1(x_1) \mathbf{u}_k^2(x_2), \end{aligned}$$

with the additional *rank parameter* r .

To be specific, suppose we seek to approximate elements u in a tensor product Hilbert space $H = H_1 \otimes H_2$, where H_i are separable Hilbert spaces spanned by the orthonormal bases $\{\varphi_i^{(q)}\}_{i \in \mathbb{N}}$, $q = 1, 2$, respectively. Ideally, given a target tolerance $\varepsilon > 0$, one would like to adapt the rank r as well as the spatial resolution of “best-suited” modes u_k^1, u_k^2 as functions of x_1, x_2 , respectively, which is obviously a highly nonlinear problem. “Best-suited” means that a given target accuracy can be met with rank r as small as possible. In the present two-dimensional case, this amounts to approximating the *infinite* matrix \mathbf{U} of basis coefficients with respect to the full tensor product basis $\{\varphi_i^{(1)} \otimes \varphi_j^{(2)}\}_{(i,j) \in \mathbb{N}^2}$ with respect to the Frobenius (or Hilbert-Schmidt) norm within the same target tolerance, by a finite *low-rank approximation* $\mathbf{U}^1(\mathbf{U}^2)^\top$, where $\mathbf{U}^1 = (u_{i,k}^1)_{i,k=1}^{N_1,r}$, $\mathbf{U}^2 = (u_{i,k}^2)_{i,k=1}^{N_2,r}$.

Denoting by \mathbf{U}_{N_1, N_2} the finite section of \mathbf{U} corresponding to the truncated basis $\{\varphi_i^{(1)} \otimes \varphi_j^{(2)}\}_{i \leq N_1, j \leq N_2}$, the error $\|\mathbf{U}^1(\mathbf{U}^2)^\top - \mathbf{U}\|_{\text{HS}}$ can be split into two portions $E_1 := \|\mathbf{U}^1(\mathbf{U}^2)^\top - \mathbf{U}_{N_1, N_2}\|_{\text{HS}}$ and $E_2 := \|\mathbf{U}_{N_1, N_2} - \mathbf{U}\|_{\text{HS}}$. The vast majority of studies focuses on controlling only E_1 , essentially ignoring how the overall error depends on the spatial discretizations represented by N_1, N_2 , determining E_2 . In the current situation $d = 2$, given a target accuracy ε_1 for E_1 , the minimal rank r can then be determined by means of the *Singular Value Decomposition* (SVD) of \mathbf{U}_{N_1, N_2} . Observing that in many cases one can achieve $E_1 \leq \varepsilon_1$ with $r \ll N_1, N_2$ indeed signals a substantial complexity reduction in the number of required coefficients $(N_1 + N_2)r$ compared to a full array representation involving $N_1 N_2$ terms. However, this by itself does not say much about the computational cost of approximately solving the original problem within some target tolerance ε , unless the error portions E_1 and E_2 are essentially *balanced*, which actually depends on the (full) *Hilbert-Schmidt decomposition* (HSD) of \mathbf{U} . Such an assessment can, however, not be reached from a linear algebra perspective alone. The central objective of this article is to highlight concepts that allow one to certifiably control the total error $\|\mathbf{U}^1(\mathbf{U}^2)^\top - \mathbf{U}\|_{\text{HS}}$ for large d , which must involve the underlying continuous model.

1.3. Layout. Unfortunately, a straightforward extension of the format (1.7) to $d > 2$ lacks in general stability. Section 2 is therefore devoted to a brief discussion of alternate stable tensor formats and their main properties. In Section 3 we outline a solution strategy for a general class of high-dimensional *elliptic* operator equations covering the two scenarios (I) and (II), addressed above, as special cases. Section 4 addresses the approximability of solutions to high-dimensional elliptic problems by tensor methods. Again, the findings for the two scenarios (I) and (II) turn out to be quite different. Finally, in Section 5 we present techniques for controlling the computational complexity of approximations as well as of numerical schemes based on the general strategy from Section 3, and discuss the conclusions for scenarios (I) and (II).

2. Subspace-based tensor formats and nonlinear approximation

2.1. Tensor formats. This section offers a brief review of approximating elements u in a tensor product

$$(2.1) \quad H := H_1 \otimes \cdots \otimes H_d,$$

of separable Hilbert spaces H_i , endowed with the unique *cross-norm* $\|\cdot\|_H$ associated with an inner product $\langle \cdot, \cdot \rangle_H$ satisfying $\langle \bigotimes_{i=1}^d v_i, \bigotimes_{i=1}^d w_i \rangle_H = \prod_{i=1}^d \langle v_i, w_i \rangle_{H_i}$.

A natural extension of (1.7) to spatial dimensions $d \geq 3$ would be to seek approximations in the form

$$(2.2) \quad u_r(x) = \sum_{k=1}^r u_k^1(x_1) u_k^2(x_2) \cdots u_k^d(x_d),$$

which is often referred to as the *canonical* or *CP-format*, [15, 37]. Again, the issue is to find suitable modes u_k^i that warrant high accuracy at the expense of a low rank r . However, first, the class of rank- r tensors is not closed, and second, best approximations from such classes do in general not exist [26]. Therefore,

representations in this format need to be treated with care and the observed inherent instability suggests looking for more stable tensor fomats [37, 39, 43, 46, 48, 52].

The canonical format is indeed only one special instance of possible tensor representation formats. A particular type of format with certain restrictions on the components are *tensor networks*; for a comprehensive discussion we refer, for instance, to [13]. In what follows we confine the discussion to the important subclass of *tree networks* that heavily draws on the favorable features of the case $d = 2$.

2.2. Tree-based hierarchical formats. To benefit from SVD concepts also when $d > 2$, the key idea is to view a (discrete) tensor \mathbf{u} of order d as a *matrix* by viewing all multi-indices with respect to some subset $\alpha \subset \alpha^* := \{1, \dots, d\}$ as row-indices, while the multi-indices with respect to the remaining variables in $\alpha^c := \alpha^* \setminus \alpha$ form the column indices. An SVD of this matrix gives rise to left and right singular vectors which are now tensors of order $\#\alpha, \#\alpha^c$, respectively. By successively further decomposing the left singular vectors then leads to approximations in terms of tensors of lower and lower order. An underlying successive splitting of groups of variables into smaller ones can be conveniently organized by a so-called (binary) *dimension tree* \mathbb{T}_d whose nodes α are subsets of the *root* $\alpha^* = \{1, \dots, d\}$. \mathbb{T}_d always contains the root α^* . Each node α of cardinality $\#\alpha > 1$ has a unique pair of children $\alpha_1, \alpha_2 \in \mathbb{T}_d$ such that $\alpha = \alpha_1 \cup \alpha_2$, $\alpha_1 \cap \alpha_2 = \emptyset$, $\#\alpha_i < \#\alpha$. The set $L(\mathbb{T}_d)$ of leaves of \mathbb{T}_d is comprised of the nodes α of cardinality one. The corresponding so-called *hierarchical tensor format* was developed by Hackbusch and Kühn [39], see also [48].

It will be important for what follows to formalize these concepts in the context of *infinite-dimensional* tensor product Hilbert spaces H of the form (2.1). To that end, consider again for any $\alpha \in \mathbb{T}_d$ the grouping $H = H_\alpha \otimes H_{\alpha^c}$. The continuous and linear extension of $M_\alpha(u_\alpha \otimes u_{\alpha^c})v := u_\alpha \langle u_{\alpha^c}, v \rangle_{H_{\alpha^c}}$, $v \in H_{\alpha^c}$, to an operator $M_\alpha(u)$ for $u \in H$,

$$(2.3) \quad M_\alpha(u) : H_{\alpha^c} \rightarrow H_\alpha$$

is a Hilbert-Schmidt operator from H_{α^c} to H_α , i.e.,

$$\|M_\alpha(u)\|_{\text{HS}} = \|u\|_H,$$

and hence is compact. In a slight abuse of terminology we refer to $M_\alpha(u)$ as a *matricization* of u . By the spectral theorem for compact operators, there exist orthonormal systems $\{u_k^\alpha\}_{k \in \mathbb{N}}$, $\{u_k^{\alpha^c}\}_{k \in \mathbb{N}}$ (depending on u) of H_α, H_{α^c} , respectively, as well as a sequence of nonnegative numbers σ_k^α tending to zero such that

$$(2.4) \quad u = \sum_{k=1}^{\infty} \sigma_k^\alpha u_k^\alpha \otimes u_k^{\alpha^c}.$$

The σ_k^α are the singular values associated with the matricization $M_\alpha(u)$. If $\sigma_k^\alpha = 0$ for $k > r_\alpha$ we say that u has α -rank (at most) $r_\alpha = r_\alpha(u) \in \mathbb{N}_0 \cup \{\infty\}$ which is the dimension of the range

$$(2.5) \quad U_\alpha = U_\alpha(u) := \overline{\text{range}(M_\alpha(u))}$$

of $M_\alpha(u)$ and hence equals the α^c -rank of u . In general, the subspaces $U_{\alpha,r} := \dim \{u_k^\alpha : k = 1, \dots, r\}$ are *optimal* in the sense that

$$(2.6) \quad \left\| u - \sum_{k=1}^r \langle u, u_k^\alpha \otimes u_k^{\alpha^c} \rangle_H u_k^\alpha \otimes u_k^{\alpha^c} \right\|_H = \min \{ \|u - w\|_H : w \text{ has } \alpha\text{-rank} \leq r \}.$$

If one further decomposes $\alpha = \alpha_1 \cup \alpha_2$ into its children α_1, α_2 , according to \mathbb{T}_d , one has the *nestedness property*

$$(2.7) \quad U_\alpha \subseteq U_{\alpha_1} \otimes U_{\alpha_2},$$

see [37] (cf. Corollary 6.18 and Theorem 6.31 there). Whenever (2.7) holds one can recursively decompose $u \in U_\alpha$ as follows. For any orthonormal bases $\{u_k^\beta\}_{k=1}^{r_\beta}$ of U_β , $\beta \in \mathbb{T}_d$, one can write

$$(2.8) \quad u_k^\alpha = \sum_{k_1=1}^{r_{\alpha_1}} \sum_{k_2=1}^{r_{\alpha_2}} \mathbf{B}^\alpha(k_1, k_2, k) u_{k_1}^{\alpha_1} \otimes u_{k_2}^{\alpha_2}. \quad \alpha \in \mathbb{T}_d \setminus L(\mathbb{T}_d),$$

where $\mathbf{B}^\alpha(k_1, k_2, k) = \langle u_k^\alpha, u_{k_1}^{\alpha_1} \otimes u_{k_2}^{\alpha_2} \rangle$. The tensors \mathbf{B}^α are referred to as *component* or *transfer tensors*. Hence any $u \in H$ can be recovered exactly in terms of its expansion tensor u as

$$(2.9) \quad u = \sum_{k_1=1}^{r_{\alpha_1}^*} \sum_{k_2=1}^{r_{\alpha_2}^*} \mathbf{B}^{\alpha^*}(k_1, k_2) u_{k_1}^{\alpha_1^*} \otimes u_{k_2}^{\alpha_2^*},$$

with $r_{\alpha_1^*} = r_{\alpha_2^*} \in \mathbb{N} \cup \{\infty\}$. A recursive substitution of (2.8) therefore parametrizes a $u \in U_{\alpha_1^*} \otimes U_{\alpha_2^*}$ in terms of the transfer tensors \mathbf{B}^α , $\alpha \in \mathbb{T} \setminus L(\mathbb{T})$ and the *mode frames* $u^\mu \in H_\mu$, $\mu \in \{1, \dots, d\}$. Using the SVD to successively generate the orthonormal systems $(u_k^\alpha)_{k=1}^{r_\alpha}$ is referred to as *hierarchical singular value decomposition HSVD*.

Before proceeding, let us emphasize that such hierarchical decompositions are actually identified by the set of *pairs*

$$(2.10) \quad \mathbb{E} := \{e = \{\alpha, \alpha^c\} : \alpha \in \mathbb{T}_d \setminus \{\alpha^*\}\},$$

called the set of *effective edges*, see [13]. Different dimension trees can give rise to the same \mathbb{E} and hence to the same matricizations. Such trees are in that sense equivalent.

In other words, for $\{\alpha, \alpha^c\} \in \mathbb{E}$ and any \mathbb{T}_d in the equivalence class determined by \mathbb{E} , either α or α^c belongs to \mathbb{T}_d , and this element is the representer of $e = \{\alpha, \alpha^c\}$ denoted by $[e] \in \mathbb{T}_d$. It is easy to see that

$$(2.11) \quad \#\mathbb{E} = 2d - 3 =: E.$$

We fix in what follows an enumeration $\{e_i\}_{i=1}^E$ of the effective edges in \mathbb{E} .

As a consequence, for $u \in H$ and a given \mathbb{E} (and hence all dimension trees in the corresponding equivalence class) we can write for the corresponding matricizations and subspaces $M_i(u) := M_{[e_i]}(u)$, $U_i = U_i(u) = U_{[e_i]}(u)$, respectively. One can then associate with u its \mathbb{E} -rank

$$(2.12) \quad \mathbf{r}_\mathbb{E}(u) \in (\mathbb{N}_0 \cup \{\infty\})^E := (r_i)_{i=1}^E,$$

where $r_i = r_i(u) := \dim U_i(u) = \text{rank}(M_i(u))$, $i = 1, \dots, E$.

In summary, any $u \in H$ can be parametrized by

$$(2.13) \quad \mathbf{p}(u) := (\mathbf{r}_\mathbb{E}(u), (\mathbf{B}^\alpha)_{\alpha \in \mathbb{T}_d \setminus L(\mathbb{T}_d)}, (u^\mu)_{\mu=1}^d).$$

Such a parametrization is obviously *finite* if all transfer tensors have finite ranks, in particular, when all entries in $\mathbf{r}_\mathbb{E}$ are finite. Moreover, the leaf elements u^μ must admit a finite parametrization, e.g. in terms of a truncated orthonormal basis for H_μ . Thus $\#\mathbf{p}(u) \in \mathbb{N} \cup \{\infty\}$ is the total number of parameters needed to represent u in the above hierarchical format.

The appeal of such hierarchical tensor decompositions for high-dimensional approximation lies in the following facts, see also [35, 36].

REMARK 2.1. Suppose for a moment that the maximal number of parameters to determine each leaf mode u^μ is n and that all entries of $r_{\mathbb{E}}(u)$ are bounded by a fixed $r \in \mathbb{N}$. In view of (2.11), we see that the number of parameters to be stored is of the order $O(dr^3 + ndr)$. The total numerical complexity of computing a HSVD under these premises is of the order of $O(dr^4 + dr^2n)$. In such a situation the *representation and computational complexities* of u depend only *linearly* on the spatial dimension d .

Thus, approximating a given $u \in H$ by properly truncated versions u_ε for which $\#\mathfrak{p}(u_\varepsilon)$ is of moderate size may open, at least for certain classes of target functions $u \in H$, a promising avenue to mitigate the curse of dimensionality.

To make this somewhat more precise, note first that the entries of the rank vectors $r_{\mathbb{E}}(u)$ must satisfy in the finite case some *compatibility conditions* in order to comply with (2.7). In fact, it follows from the results just stated that for $i = 1, \dots, E$ one has $r_i(u) \leq r_{[e_i]_1}(u)r_{[e_i]_2}(u)$. For necessary and sufficient conditions on a rank vector $r_{\mathbb{E}}$ we refer to [37, Section 11.2.3]. In what follows we denote by

$$(2.14) \quad \mathfrak{R} = \mathfrak{R}_{\mathbb{E}} \subset (\mathbb{N}_0 \cup \{\infty\})^{\mathbb{E}}$$

the set of all hierarchical rank vectors satisfying the compatibility conditions for nestedness (2.7). For any $r \in \mathfrak{R}_{\mathbb{E}}$ we define then

$$(2.15) \quad H(r) := \{u \in H : r_i(u) \leq r_i \text{ for all } i = 1, \dots, E\}.$$

In what follows we will be concerned with employing elements from such hierarchical tensor classes to approximate solutions to problems of type (I) and (II).

In the previous discussion the dimension tree \mathbb{T}_d , and hence \mathbb{E} , were kept fixed. It is certainly an interesting question of how to adapt \mathbb{T}_d to a given approximand $u \in H$, so as to warrant good approximations at the expense of possibly small ranks $r \in \mathfrak{R}_{\mathbb{E}}$. We will not address this issue in any depth but pause to briefly mention the following cases of interest. Balanced trees arise when successively splitting the nodes α into two children of roughly the same cardinality. The opposite case of a linear tree amounts to always choosing the left child α_1 have cardinality one (i.e., as a leaf node), while collecting the other entries of α in the right child α_2 . The resulting format is termed *Tensor Train* (TT) format [48]. Combining pairs of mode frames and transfer tensors, this gives rise to an explicit (entry-wise) multilinear representation

$$(2.16) \quad u(x_1, \dots, x_d) = \sum_{k_1=1}^{r_1} \cdots \sum_{k_{d-1}=1}^{r_{d-1}} u^1(x_1, k_1) u^2(k_1, x_2, k_2) u^3(k_2, x_3, k_3) \cdots u^d_{k_{d-1}}(x_d)$$

with remaining rank parameters r_1, r_2, \dots, r_{d-1} . This format is also known as *matrix product states* in physics. The representation complexity is now easily seen to be $O(dnr^2)$ which again nourishes hope to defeat the curse of dimensionality for suitably nearly-sparse $u \in H$, see e.g. [13, 48] for more details.

2.3. Complexity reduction. In contrast to the canonical format the above subspace based tensor format warrants now the *existence* of best approximations in $H(r)$. Following [31, 37], we can state this as follows.

THEOREM 2.2. For $u \in H$ and \mathbb{T}_d a dimension tree from the equivalence class of \mathbb{E} , let $r \in \mathfrak{R}_{\mathbb{E}}$ with $0 \leq r_i \leq r_i(u)$ for $i = 1, \dots, E$. Then there exists $v \in H(r)$ such that

$$\|u - v\|_H = \min\{\|u - w\|_H : r_i(w) \leq r_i, i = 1, \dots, E\}.$$

Note that as a consequence of (2.7), an element providing a best approximation to u as in Theorem 2.2 can be written as the result of a projection applied to u ; we make use of this fact in Section 5. While the computation of best approximations is usually not computationally feasible, *near-best approximations* can be obtained at affordable cost: The HSVD produces orthonormal bases $(u_k^i)_{k=1}^{r_i(u)}$ for each U_i , associated to the decreasing sequences $(\sigma_k^{(i)})_{k=1}^{r_i(u)}$ of singular values from (2.4). Truncation to lower ranks $\tilde{r}_i \leq r_i(u)$ for each i gives rise to a projection $\mathfrak{H}_{\tilde{r}}$ into $H(\tilde{r})$. It amounts to truncating the corresponding transfer tensors and mode frames in (2.13), leading to approximations with errors bounded in terms of the quantities

$$(2.17) \quad \tau_{i,r}(u) := \left(\sum_{k>r} |\sigma_k^{(i)}(u)|^2 \right)^{1/2}, \quad t_{\tilde{r}}(u) := \left(\sum_{i=1}^E \tau_{i,\tilde{r}_i}(u)^2 \right)^{\frac{1}{2}},$$

with the following quasi-optimality property shown in [35].

THEOREM 2.3. Let $u \in H$. Then for hierarchical ranks $\tilde{r} = (\tilde{r}_i)_{i=1}^E \in \mathfrak{R}_H$, we have

$$\|u - \mathfrak{H}_{\tilde{r}}(u)\|_H \leq t_{\tilde{r}}(u) \leq \sqrt{2d-3} \inf\{\|u - v\|_H : v \in H(\tilde{r})\}.$$

The operator $\mathfrak{H}_{\tilde{r}}$ in essence truncates the i -ranks or equivalently sets corresponding i -singular values to zero and is therefore sometimes referred to as *hard thresholding*. While \mathfrak{H}_r does not provide truly best rank r -approximations, the inflating constant depends only mildly on the dimension d and it is computationally feasible at affordable cost. For detailed discussions of how to realize $\mathfrak{H}_{\tilde{r}}$ efficiently we refer to [35].

Since $H(r)$ is not a linear space, arithmetic calculations with hierarchical tensors will inevitably increase ranks. To control complexity it is important to approximate an element in $H(r)$ by one of smaller ranks in $H(\tilde{r})$, $\tilde{r} \leq r$, as well as possible. This is often referred to as *recompression*. We defer the discussion of principles of how to properly balance a rank reduction by recompression with the entailed loss of accuracy to a later section.

An alternative type of rank reduction is based on the concept of *soft thresholding* of singular values. It has the advantage of preserving the contraction properties of iterative schemes, as considered in more detail in Section 5.2.3. Soft thresholding as a scalar operation is defined for a given thresholding parameter $\eta > 0$ as

$$(2.18) \quad s_{\eta}(x) := \operatorname{sgn}(x) \max\{|x| - \eta, 0\}, \quad x \in \mathbb{R}.$$

The crucial property of this operation is its *non-expansiveness*, that is,

$$(2.19) \quad |s_{\eta}(x) - s_{\eta}(y)| \leq |x - y|, \quad x, y \in \mathbb{R}.$$

Note that s_{η} can be characterized variationally by

$$(2.20) \quad s_{\eta}(x) := \arg \min_{y \in \mathbb{R}} \left\{ \frac{1}{2} |x - y|^2 + \eta |y| \right\},$$

which can be used to extend this notion to Hilbert-Schmidt operators. To that end, we make use of the well-known fact that for any two Hilbert spaces H, \tilde{H} , and

$v, w \in H \otimes \tilde{H}$ with sequences $\sigma^H(v), \sigma^H(w)$ of singular values associated with the matricizations $M_H(v), M_H(w)$, one has

$$(2.21) \quad \|\sigma^H(v) - \sigma^H(w)\|_{\ell_2} \leq \|v - w\|_{H \otimes \tilde{H}} = \|M_H(v) - M_H(w)\|_{\text{HS}}.$$

We can then define the *nuclear norm* of the matricization $M_H(v)$ of v by

$$(2.22) \quad \|M_H(v)\|_* := \|\sigma^H(v)\|_{\ell_1},$$

and let

$$(2.23) \quad S_\eta(M_H(v)) := \arg \min_{w \in H \otimes \tilde{H}} \left\{ \eta \|M_H(w)\|_* + \frac{1}{2} \|v - w\|_{H \otimes \tilde{H}}^2 \right\}.$$

Returning now to the setting (2.1), with matricizations $M_i(u) = M_{H[e_i]}(u)$, $i = 1, \dots, E$, we denote by M_i and M_i^{-1} the mappings taking a $u \in H$ into its matricization $M_i(u)$ relative to $[e_i] \in \mathbb{T}_d$ and its inverse, respectively, to set

$$(2.24) \quad \mathfrak{S}_{i,\eta}(u) := (M_i^{-1} \circ S_\eta \circ M_i)(u), \quad i = 1, \dots, E,$$

and finally

$$(2.25) \quad \mathfrak{S}_\eta(u) := (\mathfrak{S}_{E,\eta} \circ \dots \circ \mathfrak{S}_{1,\eta})(u).$$

It has been shown in [12] that this inherits the non-expansiveness of the scalar thresholding operator (2.19),

$$(2.26) \quad \|\mathfrak{S}_\eta(u) - \mathfrak{S}_\eta(v)\|_H \leq \|u - v\|_H,$$

see [12, Prop. 3.2].

3. Solution strategies for operator equations

In this section, we introduce the precise formulations of our main model problems and their sequence space representations that render them amenable to low-rank tensor approximations. We then describe a common basic construction principle of numerical solvers and present the basic principles of deriving rigorous complexity bounds.

3.1. Problem classes and representative model scenarios. Let us now consider in more detail the two concrete model problems mentioned in the introduction.

Diffusion problems (I). The first concerns diffusion problems in weak formulations on $V := H_0^1(\Omega)$, where $\Omega = \Omega_1 \times \dots \times \Omega_d$ is a product domain. We wish to find $u \in V$ for given $f \in V'$ such that

$$(3.1) \quad a(u, v) := \int_{\Omega} M \nabla u \cdot \nabla v \, dx = f(v), \quad v \in V.$$

Here we assume for simplicity that $\Omega_i = [0, 1]$ and that $M \in \mathbb{R}^{d \times d}$ is constant and symmetric (with the Poisson problem $M = I$ as a special case). This stationary boundary value problem also constitutes a first step in treating more general evolution or eigenvalue problems on high-dimensional domains. The bilinear form $a(\cdot, \cdot)$ is ensured to be V -elliptic by the assumption that there exist $\gamma, \Gamma > 0$ such that $\gamma \leq \langle M\xi, \xi \rangle \leq \Gamma$ for all $\xi \in \mathbb{R}^d$, $|\xi| = 1$. Defining the operator A by $\langle Au, v \rangle = a(u, v)$, $u, v \in V$, Lax-Milgram's Theorem then says that $A : V \rightarrow V'$ is an *isomorphism*.

It will later be important to note that the energy space V is in this case the *intersection* of tensor products of Hilbert spaces

$$(3.2) \quad V = \bigcap_{j=1}^d L_2(\Omega_1) \otimes \cdots \otimes L_2(\Omega_{j-1}) \otimes H_0^1(\Omega_j) \otimes L_2(\Omega_{j+1}) \otimes \cdots \otimes L_2(\Omega_d).$$

Parametric problems (II). Our second model problem concerns second-order elliptic PDEs with diffusion coefficients depending in an affine manner on scalar parameters $y \in Y := (-1, 1)^d$. Now let $V := H_0^1(\Omega) \otimes L_2(Y, \mu)$, with $\Omega \subset \mathbb{R}^m$ any domain with, e.g., $m \in \{1, 2, 3\}$, and let μ denote the uniform measure on Y . Assuming for simplicity right hand sides $f \in H^{-1}(\Omega)$ independent of $y \in Y$, we consider problems of the form: find $u \in V$ such that for all $v \in V$

$$(3.3) \quad a(u, v) := \int_Y \int_{\Omega} a(y) \nabla u(y) \cdot \nabla v(y) dx d\mu(y) = \int_Y \int_{\Omega} f(v(y)) dx d\mu(y),$$

where $a(y) = \bar{a} + \sum_{i=1}^d y_i \psi_i$. We assume that $\bar{a}, \psi_i \in L_{\infty}(\Omega)$, for $i = 1, \dots, d$, satisfy $\sum_{i=1}^d |\psi_i| \leq \theta \bar{a}$ for some $\theta < 1$, which ensures that the bilinear form $a(\cdot, \cdot)$ is V -elliptic and the induced operator $A : V \rightarrow V'$ is again an isomorphism. In this case the energy space $V = H_0^1(\Omega) \otimes L_2(Y, \mu)$ is a tensor product space endowed with a cross-norm, since

$$L_2(Y, \mu) = \bigotimes_{i=1}^d L_2((-1, 1), \frac{dy_i}{2}).$$

Here the case $d = \infty$ of countably many parameters, which arises, e.g., in the Karhunen-Loève expansion of random fields, is explicitly permitted and the fact that target functions may depend on infinitely many variables is a particular challenge in this scenario.

3.2. Sequence space formulation. The crucial role of rigorous a posteriori error bounds has been emphasized before. Such bounds can be based on the *error-residual relation* (1.4) valid for elliptic problems. Rather than employing duality arguments relying on spatial localization as for problems in low spatial dimensions, which here are infeasible, we will exploit these relations through first transforming the continuous problem into an equivalent one where domain and range of the transformed operator are the same. More precisely, as a first step, to make the considered variational formulations amenable to the concepts discussed in Section 2, we choose coordinates on the underlying Hilbert spaces via suitable isomorphisms to ℓ_2 sequence spaces, endowed with the norm $\|\mathbf{v}\|_{\ell_2(\mathcal{I})}^2 := \sum_{i \in \mathcal{I}} |\mathbf{v}_i|^2$. Linear mappings of this kind are given by *Riesz bases* of these spaces: a family $\Phi := \{\varphi_i\}_{i \in \mathcal{I}}$ in a Hilbert space H is Riesz basis if there exist $c_{\Phi}, C_{\Phi} > 0$ such that

$$(3.4) \quad c_{\Phi} \|\mathbf{v}\|_{\ell_2(\mathcal{I})} \leq \left\| \sum_{i \in \mathcal{I}} \mathbf{v}_i \varphi_i \right\|_H \leq C_{\Phi} \|\mathbf{v}\|_{\ell_2(\mathcal{I})}, \quad \mathbf{v} \in \ell_2(\mathcal{I}),$$

which implies that the corresponding linear mapping $S_{\Phi} : \ell_2(\mathcal{I}) \rightarrow H, \mathbf{v} \mapsto \sum_{i \in \mathcal{I}} \mathbf{v}_i \varphi_i$, is bounded and continuously invertible. Its adjoint is given by $S'_{\Phi} : H' \rightarrow \ell_2(\mathcal{I}), \psi \mapsto (\psi(\varphi_i))_{i \in \mathcal{I}}$. For a detailed description of such Riesz bases in both scenarios (I) and (II), the reader is referred to [6, 10, 11].

Now suppose one has a Riesz basis Φ for the energy space V on which the variational problem: find $u \in V$, such that

$$(3.5) \quad a(u, v) = f(v), \quad v \in V,$$

is posed. Then, for $A : V \rightarrow V'$ induced by (3.5), defining $\mathbf{A} = S'_\Phi A S_\Phi$, $\mathbf{u} = S_\Phi^{-1}u$ and $\mathbf{f} = S'_\Phi f$,

$$(3.6) \quad \mathbf{A}\mathbf{u} = \mathbf{f},$$

is an equivalent reformulation of (3.5) on $\ell_2(\mathcal{I})$. Since A and S_Φ are isomorphisms, $\mathbf{A} : \ell_2(\mathcal{I}) \rightarrow \ell_2(\mathcal{I})$ is a boundedly invertible mapping with the explicit representation $\mathbf{A} = (a(\varphi_\lambda, \varphi_\nu))_{\lambda, \nu \in \mathcal{I}}$ (see e.g. [24]). Moreover, it is not hard to verify that

$$(3.7) \quad c_a c_\Phi^2 \|\mathbf{v}\|_{\ell_2(\mathcal{I})} \leq \|\mathbf{A}\mathbf{v}\|_{\ell_2(\mathcal{I})} \leq C_a C_\Phi^2 \|\mathbf{v}\|_{\ell_2(\mathcal{I})}, \quad \mathbf{v} \in \ell_2(\mathcal{I}).$$

In what follows we often abbreviate $\|\cdot\| := \|\cdot\|_{\ell_2(\mathcal{I})}$ if the index domain \mathcal{I} is clear from the context.

3.3. Iterative solvers with recompression. Assuming that we have a well-conditioned representation $\mathbf{A}\mathbf{u} = \mathbf{f}$ on a product sequence space $\ell_2(\mathcal{I})$ with $\mathcal{I} = \mathcal{I}_1 \times \cdots \times \mathcal{I}_d$, we now aim to iteratively construct an approximation of the solution coefficient sequence \mathbf{u} in hierarchical low-rank format. Due to the high dimensionality of the index set \mathcal{I} , this is only feasible if all steps in such an iteration are performed entirely on low-rank representations, which requires corresponding low-rank representations (or approximations) of \mathbf{A} and \mathbf{f} .

In our present context of elliptic problems, the most straightforward way of obtaining a sequence of approximations converging to \mathbf{u} is a Richardson iteration: As a consequence of (3.7), one can find $\omega > 0$ such that the iterative scheme

$$(3.8) \quad \mathbf{u}^{n+1} = \mathbf{u}^n + \omega(\mathbf{f} - \mathbf{A}\mathbf{u}^n), \quad n = 0, 1, 2, \dots,$$

converges to \mathbf{u} for any \mathbf{u}^0 . Of course, the iteration (3.8) on an infinite-dimensional sequence space cannot be directly realized numerically.

Finite supports: First, one needs to ensure that all \mathbf{u}^n have finitely many nonzero entries. Since \mathbf{f} and each column of \mathbf{A} are generally infinitely supported, this amounts to an appropriate truncation of $\mathbf{A}\mathbf{u}^n - \mathbf{f}$, to arrive at a computable perturbed version

$$(3.9) \quad \mathbf{u}^{n+1} = \mathbf{u}^n - \omega_n \mathbf{r}^n, \quad \mathbf{u}^0 = 0,$$

of (3.8). The most common strategy is to make an educated guess of a fixed $\Lambda \subset \mathcal{I}$ and always use $\mathbf{r}^n = (\mathbf{A}\mathbf{u}^n - \mathbf{f})|_\Lambda$, which enforces $\text{supp } \mathbf{u}^n \subseteq \Lambda$ for all n . However, due to the limited accuracy in the residual approximation, the iteration then only converges to the Galerkin approximation \mathbf{u}_Λ given by $(\mathbf{A}|_{\Lambda \times \Lambda})\mathbf{u}_\Lambda = \mathbf{f}|_\Lambda$. Especially in high-dimensional problems, the appropriate choice (or refinement) of such Λ to achieve a certain target error is typically not obvious. Therefore, discretizations will never be fixed beforehand but will be adaptively updated.

Tensor ranks: A second issue is that, with a basic scheme as in (3.9), the tensor ranks in the representation of \mathbf{u}^n may increase rapidly with respect to n : in the addition of the low-rank representations of two vectors, their ranks are added, whereas the action of an operator in low-rank form leads to a multiplication by its ranks. Although methods using fixed-rank representations of all iterates can be constructed [36, 45], which essentially attempt to obtain an approximation by optimizing each

component in this fixed tensor representation, the high degree of nonlinearity in the resulting problems makes their convergence analysis an extremely delicate problem. In addition, such approaches then still need to be coupled with a procedure for rank adaptation to a given target accuracy.

Thus it is natural to let both the set of activated basis indices and the representation ranks evolve over the course of the iteration and hence to gradually refine both in parallel. The methods of this type that have been studied so far share the construction principle of combining mappings \mathcal{F}_n , providing an error reduction, with mappings \mathcal{R}_n that perform a re-approximation with complexity reduction,

$$(3.10) \quad \mathbf{u}^{n+1} = \mathcal{R}_n(\mathcal{F}_n(\mathbf{u}^n)).$$

Here \mathcal{F}_n can, in principle, be any procedure providing a guaranteed error reduction, such as a fixed-point iteration as in (3.8). Since we want this error reduction to happen with respect to the exact solution \mathbf{u} in $\ell_2(\mathcal{I})$, in general, this necessitates that \mathbf{r}^n , and hence \mathbf{u}^{n+1} , have larger support than \mathbf{u}^n , and that hierarchical ranks need to grow during the iteration. Accordingly, \mathcal{R}_n needs to both eliminate extraneous basis indices and reduce the ranks of iterates.

This entails a compromise between preserving a sufficient error reduction while at the same time preventing too large a growth in the representation costs of the iterates. Ideally, the reduction operation \mathcal{R}_n should be adjusted to the error reduction so as to ensure convergence of the iteration with (up to a multiplicative constant) the *best achievable total computational costs* in terms of the number of operations, that is, to ensure asymptotically optimal complexity. The basic template (3.10) has been used in the construction of the first adaptive wavelet methods with convergence rates [19, 20], where optimality was established. Procedures of the form (3.10) are also a core ingredient in many iterative methods operating on low-rank representations, for instance those proposed in [14–16, 38, 42, 45]. The choices of \mathcal{R}_n employed in these contributions, corresponding to truncation, for instance, to fixed ranks or to variable ranks with ad-hoc tolerances, however, do not ensure a suitable compromise between convergence and complexity.

In Section 5, we consider in detail low-rank solvers with complexity bounds which ensure that this compromise is met for both scenarios (I) and (II). A crucial role is played by a suitably abstracted version of a ‘‘Coarsening Lemma’’ that appeared first in [19] in the context of adaptive wavelet methods. At this point, we next discuss briefly some common aspects of choosing \mathcal{R}_n and \mathcal{F}_n that can be used to ensure convergence of \mathbf{u}^n to \mathbf{u} . By the above choice of reference basis functions, this is equivalent to convergence of the method to the exact solution u in the energy space V . This requires on the one hand the identification of a sequence of finite subsets $\Lambda^n = \Lambda_1^n \times \cdots \times \Lambda_d^n$ such that $\text{supp}(\mathbf{u}^n) \subseteq \Lambda^n$; and on the other hand, for a dimension tree assumed to be given, finding hierarchical ranks and constructing representation coefficients of a hierarchical tensor representation of \mathbf{u}^n .

Two basic ways of constructing \mathcal{F}_n in (3.10) have been considered in the literature: the first is a perturbed iteration (3.9), where for any given index set Λ^n and ranks \mathbf{r}^n of \mathbf{u}^n , one needs to provide a routine that can produce finitely supported \mathbf{r}^n in hierarchical format such that $\|\mathbf{r}^n - (\mathbf{A}\mathbf{u}^n - \mathbf{f})\| \leq \eta$ for any $\eta > 0$. For appropriate choices of η depending on n , this generally requires an enlarged product set $\Lambda^{n+1} = \Lambda_1^{n+1} \times \cdots \times \Lambda_d^{n+1} \supset \Lambda^n$ to satisfy $\text{supp}(\mathbf{r}^n), \text{supp}(\mathbf{u}^{n+1}) \subseteq \Lambda^{n+1}$. Moreover, the tensor representation of \mathbf{u}^{n+1} resulting from performing (3.9) in low-rank format generally has larger ranks \mathbf{r}^{n+1} . This strategy is considered in [6, 9–11].

The second basic construction of \mathcal{F}_n , which is conceptually closer to adaptive finite element methods, uses sequential Galerkin solves: for given Λ^n , define \mathbf{u}^n as the corresponding Galerkin solution, $\mathbf{u}^n := \mathbf{u}_{\Lambda^n}$. Then, find a product set $\Lambda^{n+1} \supseteq \Lambda^n$ such that $\|(\mathbf{A}\mathbf{u}^n - \mathbf{f})|_{\Lambda^{n+1}}\| \geq \tau \|\mathbf{A}\mathbf{u}^n - \mathbf{f}\|$ for a fixed $\tau \in (0, 1)$. This can be achieved by a sufficiently accurate finitely supported approximation \mathbf{r}^{n+1} as in the case of (3.9). The process of approximating the Galerkin solution $\mathbf{u}_{\Lambda^{n+1}}$, e.g., by a Krylov space method, then leads to new hierarchical representation ranks r^{n+1} . This strategy has been analyzed in [2].

Accordingly, \mathcal{R}_n is chosen as a composition of a rank reduction as discussed in Section 2.3, for instance by hard thresholding or soft thresholding of hierarchical singular values, and a subsequent reduction of active basis indices. The concrete realizations of \mathcal{F}_n and \mathcal{R}_n depend on the type of the problem (3.5). They rely on two essential ingredients discussed next.

3.4. Two core ingredients. The main distinction between the two problem scenarios (I) and (II) lies in the structure of the respective energy space V . This has an essential effect on the approximate evaluation of $\mathbf{A}\mathbf{u}^n$. In the case (II) where the space on which the variational problem is posed is a tensor product space endowed with a cross norm, such as $V = H_0^1(\Omega) \otimes (\bigotimes_{i=1}^d L_2((-1, 1), \frac{1}{2}dy_i))$ in (3.3), a choice of Riesz bases $\Phi_0, \Phi_1, \dots, \Phi_d$ in each of the $(d+1)$ factors leads to a *tensor product Riesz basis* Φ of V (indexed, say, by \mathbb{N}^{d+1}), and hence $S_\Phi = S_{\Phi_0} \otimes \dots \otimes S_{\Phi_d}$. As a consequence, low-rank structures in A and u are preserved in their representations $\mathbf{A}: \ell_2(\mathbb{N}^{d+1}) \rightarrow \ell_2(\mathbb{N}^{d+1})$ and $\mathbf{u} \in \ell_2(\mathbb{N}^{d+1})$; for instance, both A in (3.3) and the corresponding \mathbf{A} can be written as sums of $d+1$ Kronecker products. We then aim to find, for instance, a low-rank representation of the coefficient sequence \mathbf{u} as a tensor of order $d+1$ in hierarchical low-rank format.

3.4.1. Finite rank adaptive scaling. In the case (I) of (3.1), the mapping S_Φ can, in general, *not* be chosen to be of Kronecker rank one. To see this, it is instructive to consider the case $A = -\Delta$ on $\Omega = (0, 1)^2$ in (3.1). A simple choice of basis can be derived from the L_2 -normalized eigenfunctions $\tilde{\varphi}_{k_1, k_2}$ of A given by $\tilde{\varphi}_{k_1, k_2}(x_1, x_2) = 2 \sin(\pi k_1 x_1) \sin(\pi k_2 x_2)$ for $k_1, k_2 \in \mathbb{N}$, with corresponding eigenvalues $\pi^2(k_1^2 + k_2^2)$. The functions $\varphi_{k_1, k_2}(x_1, x_2) = 2\pi^{-2}(k_1^2 + k_2^2)^{-1/2} \sin(\pi k_1 x_1) \sin(\pi k_2 x_2)$ are thus an orthonormal basis (that is, a Riesz basis with $c = C = 1$) of $V = H_0^1(\Omega)$. The resulting mapping S_Φ for $\Phi = \{\varphi_{k_1, k_2}\}_{k \in \mathbb{N}^2}$ then cannot be written as a finite sum of Kronecker products due to the presence of the factor $(k_1^2 + k_2^2)^{-1/2}$. This reflects the fact that $H_0^1(\Omega)$ is not endowed with a cross norm, but can rather be characterized as the *intersection* of product spaces $H_0^1(0, 1) \otimes L_2(0, 1) \cap L_2(0, 1) \otimes H_0^1(0, 1)$ with its natural norm.

As shown by (3.2), the situation for larger d and more general Riesz bases is analogous, i.e., Riesz bases for the energy space V do no longer consist of separable functions. As a result the operator \mathbf{A} has in this case *infinite rank*. In order to still facilitate in the end an efficient error-controlled approximate application of the operator representation \mathbf{A} requires a judicious *preconditioning strategy* that realizes a proper balance between rank growth and stability needed to maintain convergence in (3.10).

To that end, a particularly well-suited building block are *tensor product wavelet bases*, constructed from an L_2 -orthonormal wavelet basis $\{\psi_\lambda\}_{\lambda \in \mathbb{V}}$ on a one-dimensional domain $I_1 \subseteq \mathbb{R}$, assuming $\Omega := I_1 \times \dots \times I_1$. Here \mathbb{V} is an appropriate set of scale-space indices, where the level of $\lambda \in \mathbb{V}$ is denoted by $|\lambda|$. Assuming

sufficient regularity of the ψ_λ , with $\mathcal{I} := \times_{i=1}^d \vee$ and $\Psi_\lambda := \psi_{\lambda_1} \otimes \cdots \otimes \psi_{\lambda_d}$, one has that $\Phi := \{\|\Psi_\lambda\|_{H_0^1}^{-1} \Psi_\lambda\}_{\lambda=(\lambda_1, \dots, \lambda_d) \in \mathcal{I}}$ is a Riesz basis of $H_0^1(\Omega)$. Here one has the proportionality $\|\Psi_\lambda\|_{H_0^1}^{-1} \sim \omega_\lambda := (2^{2|\lambda_1|} + \dots + 2^{2|\lambda_d|})^{-1/2}$ uniformly in λ , which leads to a similar lack of separability of S_Φ , see e.g. [19, 51].

Instead, the representation $\mathbf{T} := S'_\Psi A S_\Psi$ of A with respect to the L_2 orthonormal basis Ψ does have a simple low-rank structure. For instance, in the case of the Laplacian, where $\Delta = \partial_1^2 \otimes I + I \otimes \partial_2^2$ in the above example with $d = 2$. However, \mathbf{T} is not bounded on $\ell_2(\mathcal{I})$ and hence cannot be used in (3.8) or its perturbed version (3.9). The low-rank structure is unfortunately lost in the boundedly invertible representation

$$(3.11) \quad \mathbf{A} = \mathbf{S}^{-1} \mathbf{T} \mathbf{S}^{-1}, \quad (\mathbf{S})_{\lambda, \lambda'} = \delta_{\lambda, \lambda'} (2^{2|\lambda_1|} + \dots + 2^{2|\lambda_d|})^{-1/2}, \quad \lambda \in \mathcal{I},$$

after an ideal scaling that takes the topology of V into account and ensures the necessary mapping property of \mathbf{A} relating errors to residuals.

A crucial point for what follows, however, is that in the case $V = H_0^1(\Omega)$, the departure from the original low-rank structure is still controllable in the sense that S_Φ has efficient and explicitly computable low-rank approximations. One class of such approximations is provided by *exponential sum approximations* of the ideal non-separable scaling factors ω_λ . This is exemplified by the following result, paraphrased from [11, Thm. 9].

THEOREM 3.1. *There exist positive sequences $(w_j)_{j \in \mathbb{N}}$, $(t_j)_{j \in \mathbb{N}}$ such that for all $\lambda \in \mathcal{I} = \otimes_{i=1}^d \vee$,*

$$(3.12) \quad \tilde{\omega}_\lambda^{(\infty)} := \sum_{j=1}^{\infty} w_j \prod_{i=1}^d e^{-t_j 2^{2|\lambda_j|}} < \infty \quad \text{satisfies} \quad |\omega_\lambda - \tilde{\omega}_\lambda^{(\infty)}| \leq \frac{\omega_\lambda}{2},$$

and in addition, there exists $C > 0$ such that with

$$(3.13) \quad \tilde{\omega}_\lambda^{(m)} := \sum_{j=1}^m w_j \prod_{i=1}^d e^{-t_j 2^{2|\lambda_i|}},$$

one has

$$(3.14) \quad |\omega_\lambda - \tilde{\omega}_\lambda^{(m)}| \leq \frac{\omega_\lambda}{2}, \quad |\tilde{\omega}_\lambda^{(\infty)} - \tilde{\omega}_\lambda^{(m)}| \leq \eta \omega_\lambda, \quad \lambda \in \Lambda,$$

for any finite $\Lambda \subset \mathcal{I}$, provided that

$$(3.15) \quad m \geq C(1 + |\log \eta| + \max_{\lambda \in \Lambda} |\log \omega_\lambda|).$$

By (3.12), $\tilde{\Phi} := \{\tilde{\omega}_\lambda^{(\infty)} \Psi_\lambda\}_{\lambda \in \mathcal{I}}$ is then also a Riesz basis of $H_0^1(\Omega)$, with constants $c_{\tilde{\Phi}}, C_{\tilde{\Phi}}$ independent of d . For finite subsets of \mathcal{I} , the corresponding mapping $S_{\tilde{\Phi}}$ has efficient low-rank approximations whose Kronecker rank increases logarithmically in the desired error and linearly in

$$\max_{\lambda \in \Lambda} |\log \omega_\lambda| \sim \max_{\substack{j=1, \dots, d \\ \lambda \in \Lambda}} |\lambda_j|.$$

This quantity corresponds roughly to the largest frequency of an activated wavelet basis element. For problems on $H_0^1(\Omega)$, we can thus use the sequence space formulation $\mathbf{A} \mathbf{u} = \mathbf{f}$ with

$$(3.16) \quad \mathbf{A} = S'_{\tilde{\Phi}} A S_{\tilde{\Phi}} =: \tilde{\mathbf{S}}_\infty^{-1} \mathbf{T} \tilde{\mathbf{S}}_\infty^{-1}, \quad \mathbf{u} = S_{\tilde{\Phi}}^{-1} u, \quad \mathbf{f} = S'_{\tilde{\Phi}} f,$$

where $\tilde{\mathbf{S}}_\infty^{-1} = \text{diag}(\tilde{\omega}_\lambda^{(\infty)})_{\lambda \in \mathcal{I}}$. Then, with $\tilde{\Phi}_m := \{\tilde{\omega}_\lambda^{(m)} \Psi_\lambda\}_{\lambda \in \mathcal{I}}$, the infinite matrix $\mathbf{A}^{(m_1, m_2)} := S'_{\tilde{\Phi}_{m_1}} A S_{\tilde{\Phi}_{m_2}}$ provides an explicit low-rank approximation of \mathbf{A} when m_1, m_2 are chosen appropriately according to (3.14). For instance, in (3.1) with $M = I$, using L_2 -orthonormality of the ψ_{λ_i} , we have

$$(3.17) \quad (\mathbf{A}^{(m_1, m_2)})_{\lambda, \mu} = \sum_{j_1=1}^{m_1} \sum_{j_2=1}^{m_2} w_{j_1} w_{j_2} \sum_{k=1}^d e^{-t_{j_1} 2^{2|\lambda_k|} - t_{j_2} 2^{2|\mu_k|}} \langle \psi'_{\lambda_k}, \psi'_{\mu_k} \rangle \prod_{i \neq k} e^{-t_{j_1} 2^{2|\lambda_i|} - t_{j_2} 2^{2|\mu_i|}} \delta_{\lambda_i, \mu_i}.$$

Note that although this gives $dm_1 m_2$ terms of Kronecker rank one, $\mathbf{A}^{(m_1, m_2)}$ can in fact be represented as an operator in hierarchical tensor format with hierarchical ranks bounded by $2m_1 m_2$. One can proceed similarly for \mathbf{f} , and the approximate coefficients of u with respect to the L_2 -basis¹ $\{\Psi_\lambda\}_{\lambda \in \mathcal{I}}$ can be recovered by applying $S_{\tilde{\Phi}_m}$ with sufficiently large m to any finitely supported approximation of \mathbf{u} .

In summary, in problems such as (3.1), one arrives at a coupling between the costs of approximate low-rank representations of \mathbf{A} and the activated finite subsets of basis indices. In terms of the discretizations defined by these subsets, this means that preconditioning the problem in low-rank form to arrive at a well-conditioned matrix equation becomes more costly in terms of the arising ranks as the discretization is refined. How to dynamically adjust the choice of m_1, m_2 for given target tolerances η will be taken up again in Section 5. This issue does not arise in scenario (II), since there V is endowed with a cross-norm.

3.4.2. Contractions and tensor coarsening. The reduction operator \mathcal{R}_n in (3.10) must involve, aside from rank-reduction, a mechanism to control representation complexity in terms of activated basis functions, which is the second core component in all variants of (3.10). Specifically, this requires the identification of a sequence of finite subsets $\Lambda^n = \Lambda_1^n \times \cdots \times \Lambda_d^n$ such that $\text{supp}(\mathbf{u}^n) \subseteq \Lambda^n$ that warrants sufficient accuracy. In purely sparsity-based adaptive methods, as introduced in [19, 20], such a reduction of indices can be done by keeping the entries of largest absolute value. In our context, this is not an option: on one hand, directly sorting the high-dimensional coefficient sequence is not possible for complexity reasons; on the other hand, the tensor decomposition requires the product structure of index sets to be preserved. A possible way around these two restrictions, proposed in [4] and analysed further in [9, 11], is to select the relevant basis indices based on the *one-dimensional sequences*

$$\pi^{(i)}(\mathbf{v}) = (\pi_{\lambda_i}^{(i)}(\mathbf{v}))_{\lambda_i \in \mathcal{I}_i} \in \ell_2(\mathcal{I}_i),$$

where

$$\pi_{\lambda_i}^{(i)}(\mathbf{v}) := \left(\sum_{\lambda_1, \dots, \lambda_{i-1}, \lambda_{i+1}, \dots, \lambda_d} |v_{\lambda_1, \dots, \lambda_{i-1}, \lambda_i, \lambda_{i+1}, \dots, \lambda_d}|^2 \right)^{\frac{1}{2}},$$

¹Schemes for performing SVD-type approximations in H^1 directly using L_2 -basis coefficients have recently been proposed in [1]; however, this leads to weakened quasi-optimality properties.

which were termed *contractions* in [9] and can also be regarded as one-dimensional densities of the tensor. These can be efficiently evaluated using the identity²

$$(3.18) \quad \pi_\lambda^{(i)}(\mathbf{v}) = \left(\sum_k |\mathbf{U}_{\lambda,k}^{(i)}|^2 |\sigma_k^{(i)}|^2 \right)^{\frac{1}{2}}, \quad \lambda \in \mathcal{I}_i,$$

in terms of the *HSVD* mode frames $\mathbf{U}^{(i)}$ and the corresponding sequences of singular values $\sigma^{(i)}$. One easily obtains the following *quasi-optimality* result for coarsening based on these sequences.

PROPOSITION 3.2. *Let $\Lambda(\mathbf{v}; N) = \Lambda^{(1)}(\mathbf{v}; N) \times \cdots \times \Lambda^{(d)}(\mathbf{v}; N) \subset \mathcal{I}$ be the product index set corresponding to the N largest elements of $\{\pi_\lambda^{(i)}(\mathbf{v}) : i = 1, \dots, d, \lambda \in \mathcal{I}_i\}$. Then for any $\hat{\Lambda} = \hat{\Lambda}_1 \times \cdots \times \hat{\Lambda}_d$ with $\sum_{i=1}^d \#\hat{\Lambda}_i \leq N$, with*

$$(3.19) \quad s_N(\mathbf{v}) := \left(\sum_{i=1}^d \sum_{\lambda \in \mathcal{I}_i \setminus \Lambda^{(i)}(\mathbf{v}; N)} |\pi_\lambda^{(i)}(\mathbf{v})|^2 \right)^{\frac{1}{2}}$$

one has

$$\|\mathbf{v} - (\mathbf{v}|_{\Lambda(\mathbf{v}; N)})\| \leq s_N(\mathbf{v}) \leq \sqrt{d} \|\mathbf{v} - (\mathbf{v}|_{\hat{\Lambda}})\|.$$

Note that this result is analogous to Theorem 2.3, and as discussed in further detail in Section 5, the effect of the two types of reduction operations combined in \mathcal{R}_n can be estimated using essentially the same techniques. For finitely supported input sequences \mathbf{w} on \mathcal{I} , in both cases one can implement computational routines with analogous properties:

- **RECOMPRESS**($\mathbf{w}; \eta$), which performs *HSVD* hard thresholding, such that $\mathbf{w}_\eta := \text{RECOMPRESS}(\mathbf{w}; \eta)$ has hierarchical ranks $r_\eta := r_{\mathbb{E}}(\mathbf{w}_\eta)$ with $\max_i r_{\eta, i}(\mathbf{w}_\eta)$ minimal such that $t_{r_\eta}(\mathbf{w}) \leq \eta$, and hence by Theorem 2.3, $\|\mathbf{w} - \mathbf{w}_\eta\|_{\ell_2(\mathcal{I})} \leq \eta$. This routine has the cost of performing the *HSVD* as discussed in Remark 2.1. One has the quasi-optimality property

$$\|\mathbf{w} - \text{COARSEN}(\mathbf{w}; \eta)\| \leq \sqrt{2d-3} \min_{r_{\mathbb{E}}(\hat{\mathbf{w}}) \leq r_\eta} \|\mathbf{w} - \hat{\mathbf{w}}\|.$$

- **COARSEN**($\mathbf{w}; \eta$), which determines the product index set $\Lambda(\mathbf{w}; N_\eta)$ as in Proposition 3.2 with N_η minimal such that $s_{N_\eta}(\mathbf{w}) \leq \eta$, which ensures $\|\mathbf{w} - \text{COARSEN}(\mathbf{w}; \eta)\| \leq \eta$. This requires computing $\pi_\lambda^{(i)}(\mathbf{w})$ from the *HSVD* for all $i = 1, \dots, d$, $\lambda \in \mathcal{I}_i$, and (quasi-)sorting these values. From Proposition 3.2, one has the quasi-optimality property

$$\|\mathbf{w} - \text{COARSEN}(\mathbf{w}; \eta)\| \leq \sqrt{d} \min_{\hat{\Lambda}} \|\mathbf{w} - (\mathbf{w}|_{\hat{\Lambda}})\|,$$

where the minimum is over all $\hat{\Lambda} = \hat{\Lambda}_1 \times \cdots \times \hat{\Lambda}_d$ such that $\sum_i \#\hat{\Lambda}_i \leq N_\eta$.

An important feature of the resulting schemes of the form (3.10), with appropriately adjusted \mathcal{F}_n and \mathcal{R}_n , is their *universality*: they require no information on the approximability of the sought solution \mathbf{u} to produce quasi-optimal solutions. In order to assess the performance of such methods, however, we hypothetically assume \mathbf{u} to have certain approximability properties and then show that the method under consideration will have the expected complexity in such a case without using, however, any a priori knowledge of such approximability properties.

²A technique of using (3.18) for steering adaptivity in high dimensions was also subsequently developed independently in [30].

3.5. Approximation classes. The natural approximability properties to consider are strongly tied to the choice of \mathcal{R}_n , that is, to the routines RECOMPRESS and COARSEN; in our case, these properties are the following intrinsic features of \mathbf{u} :

- (i) The convergence rate of the best approximation in $\ell_2(\mathcal{I})$ of hierarchical rank \mathbf{r} to \mathbf{u} as \mathbf{r} increases, for instance, as $|\mathbf{r}|_\infty \rightarrow \infty$.
- (ii) The asymptotic decay of the decreasing rearrangement of all values $\pi_\lambda^{(i)}(\mathbf{u})$ for $i = 1, \dots, d$ and $\lambda \in \mathcal{I}_i$.

The notion of approximation classes is a classical means to quantify approximation properties. Regarding (i), we follow [9] and consider positive and strictly increasing sequences $\gamma = (\gamma(n))_{n \in \mathbb{N}_0}$ with $\gamma(0) = 1$ and $\gamma(n) \rightarrow \infty$ as $n \rightarrow \infty$. For $\mathbf{v} \in \ell_2(\mathcal{I})$ let

$$\|\mathbf{v}\|_{\mathcal{A}_{\mathcal{H}}(\gamma)} := \sup_{r \in \mathbb{N}_0} \gamma(r) \inf_{\max r_{\mathbb{E}}(\mathbf{w}) \leq r} \|\mathbf{v} - \mathbf{w}\|,$$

where $r_{\mathbb{E}}(\mathbf{w})$ reduces to $\text{rank}(\mathbf{w})$ when $d = 2$ and $\mathcal{I} = \mathcal{I}_1 \times \mathcal{I}_2$, as in the case of separating spatial and parametric variables in the problem class (II). It will be seen in the next section that relevant sequences γ cover a wide range from algebraic to (sub-)exponential rates. This gives rise to the approximation classes

$$(3.20) \quad \mathcal{A}_{\mathcal{H}}(\gamma) := \{\mathbf{v} \in \ell_2(\mathcal{S} \times \mathcal{F}) : \|\mathbf{v}\|_{\mathcal{A}_{\mathcal{H}}(\gamma)} < \infty\}.$$

For $\mathbf{u} \in \mathcal{A}_{\mathcal{H}}(\gamma)$, a hierarchical tensor best approximation is guaranteed to achieve accuracy ε with hierarchical ranks \mathbf{r}_ε whenever

$$(3.21) \quad \max r_\varepsilon \geq \gamma^{-1}(\|\mathbf{u}\|_{\mathcal{A}_{\mathcal{H}}(\gamma)} \varepsilon^{-1}).$$

We will be especially interested in exponential-type decay of low-rank approximation errors, where with $\gamma(r) = e^{cr^\beta}$ for some $c, \beta > 0$, for $\mathbf{v} \in \mathcal{A}_{\mathcal{H}}(\gamma)$ one has

$$\inf_{\max r_{\mathbb{E}}(\mathbf{w}) \leq r} \|\mathbf{v} - \mathbf{w}\| \leq \|\mathbf{v}\|_{\mathcal{A}_{\mathcal{H}}(\gamma)} e^{-cr^\beta}$$

and accordingly

$$\max r_\varepsilon \geq (c^{-1} \ln(\|\mathbf{u}\|_{\mathcal{A}_{\mathcal{H}}(\gamma)} \varepsilon^{-1}))^{1/\beta}.$$

Regarding (ii), for an approximation \mathbf{v} of bounded support to \mathbf{u} , the number of nonzero coefficients $\#\text{supp}_i \mathbf{v}$ 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. \mathcal{A}^s is comprised of all $\mathbf{w} \in \ell_2(\mathcal{J})$ for which the quasi-norm

$$(3.22) \quad \|\mathbf{w}\|_{\mathcal{A}^s(\mathcal{J})} := \sup_{N \in \mathbb{N}_0} (N+1)^s \inf_{\substack{\Lambda \subset \mathcal{J} \\ \#\Lambda \leq N}} \|\mathbf{w} - \mathbf{R}_\Lambda \mathbf{w}\|$$

is finite. The approximation of a sequence $\mathbf{u} \in \mathcal{A}^s$ up to an error ε by a sequence with N_ε nonzero entries is ensured with $N_\varepsilon \geq \|\mathbf{u}\|_{\mathcal{A}^s}^{\frac{1}{s}} \varepsilon^{-\frac{1}{s}}$. The separate sparsity of tensor mode frames for $\mathbf{u} \in \ell_2(\mathcal{I})$ can be quantified by assumptions of the form $\pi^{(i)}(\mathbf{u}) \in \mathcal{A}^s(\mathcal{I}_i)$ for $i = 1, \dots, d$.

Before turning to the complexity analysis, we next give an overview of available approximability results that indicate which approximation classes solutions to both problem types (I) and (II) can be expected to belong to. This will then serve formulating proper benchmarks for the performance of specifications of algorithms.

4. Approximability

In this section, we discuss approximability properties of solutions to problems of both types (I) and (II) with respect to low-rank or hierarchical tensor methods, partly in comparison with best n -term approximation. These theoretical questions concerning the “expressive power” of certain approximation formats serve two purposes. On one hand, they help judging which approximation type is best suited for which problem class. It is clear from the preceding discussions that low-rank or hierarchical tensor formats, due to their “higher level of nonlinearity” may increase the computational burden and also raise coding challenges. This should be justified by a superior performance in terms of the resulting *work-accuracy balance*.

On the other hand, describing best possible approximability, these results provide *benchmarks* for assessing the performance of versions of Algorithm 5.1. This asks for the computational cost required to meet the target accuracy in a given format. It includes two basic questions, namely

- What is the *representation complexity* of a given approximation format for a given problem class, i.e., the number of degrees of freedom required by a given approximation format to achieve the target accuracy?
- What is the corresponding *computational complexity*, that is, can one devise a numerical method that certifiably achieves a given target accuracy at the expense of a computational work that stays as close as possible to the representation complexity?

In this section we address the first question mainly for low-rank and tensor methods. This asks for a *regularity theory* in a broad sense where classical smoothness measures are replaced by quantified approximability. Corresponding notions of *low-rank approximability* would delineate the extent to which tensor methods could ideally mitigate the curse of dimensionality. Such methods exhibit a significantly higher level of nonlinearity than classical best N -term approximation, but are more restricted than DNNs, for which comparable theoretical guarantees are not yet available.

4.1. High-dimensional diffusion. The representation complexity for problems of the type (3.1) is perhaps best understood for the special case that the diffusion matrix M is diagonal [25]. In fact, what matters most is the following structure of the operator A induced by the V -elliptic bilinear form in this case, namely that

$$(4.1) \quad A = \sum_{j=1}^d I \otimes \cdots \otimes I \otimes A_j \otimes I \otimes \cdots \otimes I,$$

has rank d and the “low-dimensional” operators

$$(4.2) \quad A_i : \mathbf{H}_i \rightarrow \mathbf{H}'_i, \quad \langle A_i u, v \rangle = a_i(u, v) : \mathbf{H}_i \times \mathbf{H}_i \rightarrow \mathbb{R},$$

are \mathbf{H}_i -elliptic. The simplest example is $A_j = -\partial_{x_j}(M_j \partial_{x_j})$ and $\mathbf{H}_j = H_0^1(\Omega_j)$, $\Omega_j \subset \mathbb{R}^m$, $j = 1, \dots, d$. As seen below, it is not important that the A_j are second order operators. Again the energy space V has then the form (3.2) with $H_0^1(\Omega_i)$ replaced by \mathbf{H}_i .

While the structure of A may nourish hope that the solution to

$$(4.3) \quad Au = f \quad \text{on } \Omega_1 \times \cdots \times \Omega_d$$

can be accurately approximated by low-rank functions, A^{-1} as a mapping from V' to V has infinite rank which renders the question of quantifiable low-rank approximability less clear. In view of (1.4), it is this mapping property that determines the representation complexity. We review in this section some results from [25] on quantified *low-rank approximability* of solutions to high-dimensional diffusion equations of the form (4.3), (4.1) that do establish rigorous relations between target accuracies and rank growth in a “tamed canonical tensor format”, as explained below. While the canonical format is most convenient for the present purpose, it can be embedded in the hierarchical tensor format as well, see [13]. The term “tamed” accounts for the fact that computing in this format is, in general, prone to instability.

The relevant structural properties of A are conveniently reflected by its *eigen-system*. Denoting by $\{e_{j,k}\}_{k \in \mathbb{N}}$ the set of eigenfunctions of the low-dimensional component operators A_j with eigenvalues $\lambda_{j,k}$, one easily verifies that

$$(4.4) \quad Ae_\nu = \lambda_\nu e_\nu, \quad e_\nu = e_{1,\nu_1} \otimes \cdots \otimes e_{d,\nu_d}, \quad \lambda_\nu = \lambda_{1,\nu_1} + \cdots + \lambda_{d,\nu_d}.$$

This allows one to define the smoothness norm

$$(4.5) \quad \|v\|_s^2 := \sum_{\nu \in \mathbb{N}^d} \lambda_\nu^s |\langle v, e_\nu \rangle|^2.$$

Denoting by \mathbb{H}^t the space of all v for which $\|v\|_t$ is finite, the operator A then acts as an isometry between these scales, $\|A\|_{\mathcal{L}(\mathbb{H}^{t+2}, \mathbb{H}^t)} = 1$, $t \in \mathbb{R}$. While these norms exist for all $t \in \mathbb{R}$, they are equivalent with classical Sobolev norms only for a limited range of smoothness around $t = 1$, depending on the specific problem data. In particular, one always has $V = \mathbb{H}^1$.

Since obviously λ_ν^{-1} is not separable, $A^{-1} : \mathbb{H}^t \rightarrow \mathbb{H}^{t+2}$ has infinite rank. Hence, a bounded rank of data f does not a priori indicate which rank suffices for an approximate solution to warrant a given accuracy.

We will state next in more precise terms what we mean by *low-rank approximability* of a function. Since we will deal with tensors in the CP format, continuing to denote rank-one functions as $(g_1 \otimes \cdots \otimes g_d)(x) := g_1(x_1) \cdots g_d(x_d)$, we have to incorporate means to ensure stability. To that end, we consider for $r \in \mathbb{N}$ *tamed* rank- r tensor classes of the form

$$\mathcal{T}_r(\mathbb{H}^t) := \left\{ g \in \mathbb{H}^t : g = \sum_{k=1}^r \bigotimes_{j=1}^d g_j^{(k)} \right\},$$

endowed with the “norm-like” quantities

$$(4.6) \quad \|g\|_{r,t} := \inf_{\sum_{k=1}^r \bigotimes_{j=1}^d g_j^{(k)} = g} \left\{ \max_{k=1, \dots, r} \{ \|g\|_t, \|g^{(k)}\|_t \} \right\}.$$

Controlling $\|g\|_{r,t}$ is to account for the fact that rank- r representations are in general not unique and to avoid cancellation effects between terms with large norms.

Despite the otherwise inherent instability of the CP format this allows us to introduce *approximation classes* collecting those elements in \mathbb{H}^t that are close to low-rank tensors in a quantifiable way. To that end, consider for any $\zeta > 0$ the *K-functional*

$$(4.7) \quad K_r(v, \delta) := K_r(v, \delta, \mathbb{H}^t, \mathbb{H}^{t+\zeta}) := \inf_{g \in \mathcal{T}_r(\mathbb{H}^{t+\delta})} \left\{ \|v - g\|_t + \delta \|g\|_{r,t+\zeta} \right\}.$$

Then $\delta^{-1}K_r(v, \delta) \leq C$ means that there exists a rank- r tensor g of smoothness $t + \zeta$ for which $\|v - g\|_t \leq \delta C$ and $\|g\|_{r, t + \zeta} \leq C$. To relate the achievable accuracy δ to the rank r , consider as before a strictly increasing sequence of positive numbers $\gamma = (\gamma(n))_{n \in \mathbb{N}}$, called a *growth sequence*. Monotonicity of γ ensures that the inverse γ^{-1} exists. We now collect in the *approximation class* $\mathcal{A}^\gamma = \mathcal{A}^\gamma(\mathbb{H}^t, \mathbb{H}^{t+\zeta})$ all those $v \in \mathbb{H}^t$ for which the quasi-norm

$$(4.8) \quad \|v\|_{\mathcal{A}^\gamma(\mathbb{H}^t, \mathbb{H}^{t+\zeta})} := \sup_{r \in \mathbb{N}_0} \gamma(r) K_r(f, \gamma(r)^{-1}; \mathbb{H}^t, \mathbb{H}^{t+\zeta}),$$

is finite, i.e., those elements in \mathbb{H}^t which can be stably approximated by rank- r tensors in $\mathbb{H}^{t+\zeta}$ with accuracy $O(\gamma(r)^{-1})$. In other words, membership of a v to \mathcal{A}^γ means that approximating a given $v \in \mathcal{A}^\gamma$ by a tensor in $\mathbb{H}^{t+\zeta}$ within accuracy $\varepsilon > 0$, is guaranteed, in analogy to (3.21), by the rank

$$(4.9) \quad r = r(\varepsilon) = \lceil \gamma^{-1}(\|v\|_{\mathcal{A}^\gamma} / \varepsilon) \rceil.$$

Since A^{-1} has infinite rank one cannot expect any low-rank approximability of the data f to be exactly inherited by the solution $u = A^{-1}f$. The following result says that a low-rank approximation order $O(1/\gamma(r))$ for the data f implies a slightly weaker order $O(1/\hat{\gamma}(r))$ for the solution u . The attenuated growth sequence $\hat{\gamma}(r)$ is given as follows. Let $G(x) = x(\ln \gamma(x))^2$, $x > 0$, and define

$$(4.10) \quad \hat{\gamma} := \gamma(G^{-1}(m)).$$

Since G has super-linear growth, G^{-1} has only sublinear growth causing $\hat{\gamma}(n)$ to have a somewhat weaker growth than γ .

THEOREM 4.1. *Assume that the data f belong to $\mathcal{A}^\gamma(\mathbb{H}^t, \mathbb{H}^{t+\zeta}) \subset \mathbb{H}^t$ for some $0 < \zeta \leq 2$. Then, for $\hat{\gamma}$ defined by (4.10), the solution u of (4.3) belongs to $\mathcal{A}^{\hat{\gamma}}(\mathbb{H}^{t+2}, \mathbb{H}^{t+2+\zeta}) \subset \mathbb{H}^{t+2}$ and for each $r \in \mathbb{N}$ there exists a $u_r \in \mathcal{T}_r(\mathbb{H}^{t+2+\zeta})$ such that*

$$(4.11) \quad \|u - u_r\|_{t+2} \leq C \hat{\gamma}(r)^{-1} \|f\|_{\mathcal{A}^\gamma(\mathbb{H}^t, \mathbb{H}^{t+\zeta})},$$

holds for some fixed constant C independent of r .

The relation between γ and $\hat{\gamma}$ is illustrated by taking $\gamma(r) = r^\alpha$, for some positive α . Then $\hat{\gamma}(r) \sim \left(\frac{r}{\alpha^2 \log r}\right)^\alpha$, i.e., up to a log-factor $\hat{\gamma}$, the rank approximation order of the solution exhibits still the same algebraic rate. Instead, for $\gamma(r) = e^{\alpha r}$ one obtains $\hat{\gamma}(r) \sim e^{(\alpha r/C)^{1/3}}$ and thus a more significant relative degradation of low-rank approximability caused by A^{-1} . Of course, when $f \in \mathcal{T}_k(\mathbb{H}^{t+\zeta})$ has a finite tamed rank k , it belongs to $\mathcal{A}^\gamma(\mathbb{H}^t, \mathbb{H}^{t+\zeta})$ for any growth function γ . This explains the high efficiency of low-rank approximations to the solution u in such a case.

The result in Theorem 4.1 is in this form unusual since the approximants $u_r \in \mathcal{T}_r(\mathbb{H}^{t+2+\zeta})$ in (4.11) are not yet determined by finitely many parameters. It remains to identify numerically viable *finitely parametrized* approximants $\bar{u}_r \in \mathcal{T}_r(\mathbb{H}^{t+2+\zeta})$ and bound the complexity of computing them.

We summarize next some related simplified but representative results from [25]. Note first that stable approximability of the data f by somewhat smoother low-rank expressions in $\mathcal{T}_r(\mathbb{H}^{t+\zeta})$ can be viewed as *excess-regularity* of the data that, according to Theorem 4.1, is in this form inherited by the solution. Moreover, a rank-one tensor g belongs to \mathbb{H}^s only if its factors g_i belong to the corresponding

component spaces H_i^s . Therefore, the excess-regularity ζ affects the number of parameters needed to form an ε -accurate rank- r approximant.

More precisely, the following facts are somewhat specialized cases of the results in [25]. For convenience we assume here 1-dimensional component domains $\Omega_i \subset \mathbb{R}$, $\zeta \in (0, 2]$, and $t = -1$. Then, for any $f \in \mathcal{A}^\gamma(H^{-1}, H^{\zeta-1})$ and any $\varepsilon > 0$ there exists a $\bar{u}(\varepsilon) \in \mathcal{T}_{\bar{r}(\varepsilon)}(H^{\zeta+1})$, which is determined by $N(\varepsilon, d) < \infty$ parameters, such that

$$(4.12) \quad \|u - \bar{u}(\varepsilon)\|_1 \leq \varepsilon, \quad \|\bar{u}(\varepsilon)\|_{\bar{r}(\varepsilon), 1} \leq C_1 \|f\|_{\mathcal{A}^\gamma(H^{-1}, H^{\zeta-1})},$$

where

$$(4.13) \quad \bar{r}(\varepsilon) \leq C(\zeta)\gamma^{-1}(\bar{C}\|f\|_{\mathcal{A}^\gamma/\varepsilon})|\ln \varepsilon|^2, \quad N(\varepsilon, d) \leq C_2 d \varepsilon^{-1/\zeta} \bar{r}(\varepsilon)^{1+\frac{1}{\zeta}}.$$

Since in view of (4.9), an ε -accurate approximation of f requires at most a rank $r(\varepsilon) \lesssim \gamma^{-1}(\|f\|_{\mathcal{A}^\gamma/\varepsilon})$, the solution rank $\bar{r}(\varepsilon)$ is only mildly increased by a factor of the order of $|\ln \varepsilon|^2$. Moreover, the total number $N(\varepsilon, d)$ of degrees of freedom exceeds an ‘‘ideal count’’ of $d\varepsilon^{-1/\zeta}\bar{r}(\varepsilon)$, obtained when all tensor factors of $\bar{u}(\varepsilon)$ of smoothness $1 + \zeta$ were known beforehand, by the factor $\bar{r}(\varepsilon)^{\frac{1}{\zeta}}$.

A pivotal constituent behind the above results is again approximation by *exponential sums*. Specifically, one can employ a somewhat simpler version of Theorem 3.1, namely a classical result by Braess and Hackbusch [17, 18] stating that

$$(4.14) \quad \sup_{x \in [1, \infty)} \left| \frac{1}{x} - S_r(x) \right| \leq C e^{-\pi\sqrt{r}}, \quad S_r(x) = \sum_{k=1}^r \omega_{r,k} e^{-\alpha_{r,k} x}.$$

Defining operator exponentials $e^{-A}v$ through the eigensystems (4.4), one can prove the following result that takes the scale-change caused by A^{-1} into proper account.

THEOREM 4.2. *There exists a constant C_0 such that for $t \leq s \leq 1$, $s - t \leq 2$, one has*

$$(4.15) \quad \|A^{-1} - S_r(A)\|_{\mathcal{L}(H^t, H^s)} \leq C_0 e^{-\frac{(2-s+t)\pi}{2}\sqrt{r}}, \quad r \in \mathbb{N}.$$

Of course, this allows one to determine $R(\varepsilon)$ such that for any given rank-one tensor $\tau = \tau_1 \otimes \cdots \otimes \tau_d$ one can approximate $A^{-1}\tau$ with accuracy ε in a desired smoothness scale by the rank- $R(\varepsilon)$ tensor

$$S_{R(\varepsilon)}(A)\tau = \sum_{k=1}^{R(\varepsilon)} \omega_{R,k} \left(\bigotimes_{j=1}^d e^{-\alpha_{R,k} A_j} \tau_j \right), \quad \text{where } R = R(\varepsilon) \sim |\ln \varepsilon|^2,$$

explaining the rank-increase by a factor $|\ln \varepsilon|^2$.

The *representation complexity* bounds (4.13) make use of the implied regularity of the tensor factors. They can be extended to bounds for the *computational complexity* of actually computing an accuracy controlled low-rank approximant by devising a numerical scheme for evaluating exponential sums of the form $S_{R(\varepsilon)}(A)\tau$. Of course, using the eigensystems of A is not a numerically viable option. Instead, approximating exponential factors of the type $e^{-\alpha A_j \tau_j}$ in low dimensions can be based on the Dunford integral representation

$$(4.16) \quad e^{-\alpha A_j} \tau_j = \frac{1}{2\pi i} \int_{\Gamma} e^{-\alpha z} (zI - A_j)^{-1} \tau_j dz,$$

where $\Gamma \subset \mathbb{C}$ is a suitably chosen analytic curve separating the spectrum of A_j from the left complex half plane, and permitting an analytic extension of the integrand to a sufficiently wide region, see [25] for details. Sharp estimates for the

mapping properties of the low-dimensional operator exponentials $e^{-\alpha A_j}$ and using exponential convergence of sinc-quadrature, one can devise for any given target accuracy ε appropriate discretizations of the resolvent problems $(zI - A_j)u_j(z) = \tau_j$ for suitably chosen quadrature points z . Thus, in contrast to earlier work in [32], using such integral representations, the order of first discretizing the operator A_j and then using quadrature, is here reversed in order to exploit mapping properties for error controlled computation. When the component operators A_j are indeed second order $H_0^1(\Omega_j)$ -elliptic and if H^s agrees with $H_0^1(\Omega_j) \cap H^s(\Omega_j)$, $s \in [1, 2]$, one can bound the *number of operations and used degrees of freedom* for computing $\bar{u}(\varepsilon)$ by

$$(4.17) \quad d^{1+\bar{\rho}/\zeta} \varepsilon^{-\bar{\rho}/\zeta} \gamma^{-1} (C \|f\|_{\mathcal{A}^\gamma} / \varepsilon) |\ln \varepsilon|^2 \left(\ln \left(\frac{d}{\varepsilon} \right) \right)^2,$$

where $\bar{\rho}$ is a fixed positive number. This is slightly more pessimistic than (4.13) but nevertheless shows that the curse of dimensionality is clearly avoided for a wide scope of data f .

Exponential sum based schemes are very performant but restricted to the particular operator type (4.1). Nevertheless, in particular, sub-exponential rates are seen to provide relevant benchmarks for high-dimensional diffusion problems with low-rank data. The performance of the more general iterative approach, outlined in Section 3, will therefore be discussed in Section 5 for such benchmarks.

4.2. Multi-parametric problems. As mentioned earlier in Section 3, the uniform ellipticity assumption on the parametric family $a(y)$, $y \in Y$, ensures ellipticity of (3.3) in the tensor product Hilbert space

$$(4.18) \quad V := H_0^1(\Omega) \otimes L_2(Y, \mu) = H_0^1(\Omega) \otimes \left(\bigotimes_{j=1}^d L_2((-1, 1), \frac{dy_j}{2}) \right) = L_2(Y, H_0^1(\Omega), \mu).$$

The mapping $u : y \mapsto u(y) \in H$ is then well-defined and the set $u(Y) \subset H$ of all states that can be attained when traversing the parameter domain, sometimes referred to as *solution manifold*, is compact. Hence the *Kolmogorov n -widths*

$$d_n(u(Y))_V := \inf_{\dim(W)=n} \sup_{y \in Y} \inf_{w \in W} \|u(y) - w\|_V,$$

tend to zero. They quantify linear approximability of $u(Y)$ in the worst case sense with respect to Y viewing u as a mapping into the smaller Bochner space $L_\infty(Y, H_0^1(\Omega))$. A natural accuracy measure for low-rank approximation in V is the mean-squared error

$$(4.19) \quad \delta_n(u, \mu)_V^2 := \inf_{\dim(W)=n} \int_Y \min_{w \in W} \|u(y) - w\|_{H^1(\Omega)}^2 d\mu(y).$$

As exploited earlier, best subspaces realizing $\delta_n(u, \mu)_V$ are obtained with the aid of the Hilbert-Schmidt operator $M_u : L_2(Y, \mu) \rightarrow H_0^1(\Omega)$ defined by

$$M_u \varphi := \int_Y u(y) \varphi(y) d\mu(y),$$

through its Hilbert-Schmidt decomposition

$$(4.20) \quad M_u = \sum_{k=1}^{\infty} \sigma_k v_k \langle \cdot, \phi_k \rangle_{L_2(Y, \mu)},$$

where $\sigma = (\sigma_k)_{k \geq 1} \in \ell_2(\mathbb{N})$ is non-negative and non-increasing. Moreover, the $(v_k)_{k \geq 1}, (\phi_k)_{k \geq 1}$ are orthonormal systems in $H_0^1(\Omega)$ and $L_2(Y, \mu)$, respectively. In particular, the spaces $W_n := \text{span}\{v_k : k = 1, \dots, n\}$ realize $\delta_n(u, \mu)_V$ and

$$(4.21) \quad \delta_n(u, \mu)_V^2 = \sum_{k > n} \sigma_k^2.$$

The largest r for which $\sigma_r > 0$ is the rank of M_u .

Since μ is a probability measure $\delta_n(u, \mu)_V$ is always dominated by the n -widths

$$(4.22) \quad \delta_n(u, \mu)_V \leq d_n(u(Y))_V,$$

so that bounds for d_n provide such for δ_n as well. In this context a crucial fact is *holomorphy* of the map u , see [21–23], even when $d = \infty$, which is, in particular, responsible for decay rates of δ_n, d_n that are robust in d .

Quantifying such rates, however, requires distinguishing two fundamental regimes, namely the case (a) $d < \infty$ possibly very large but fixed, and (b) $d = \infty$. In (a) all parametric components y_i will be viewed as equally important. In scenario (b) uniform ellipticity necessitates a decay of the parametric expansion functions ψ_j in $L_\infty(\Omega)$.

4.2.1. *Finitely many parameters.* It is straightforward to show (for instance, using Taylor expansion as in [5]) that for fixed $d < \infty$, the n -widths of the solution manifold $u(Y)$ satisfy

$$(4.23) \quad d_n(u(Y))_H \leq C e^{-cn^{1/d}},$$

where the constants c, C depends on \bar{a} and r in the uniform ellipticity condition $r \leq \bar{a} - \sum_{i=1}^d |\psi_i| \leq a(y) = \bar{a} + \sum_{i=1}^d y_i \psi_i$.

The rate (4.23) can be slightly improved to $e^{-|\ln \rho| n^{\frac{1}{d-1}}}$ under the assumption

$$(4.24) \quad \sum_{i=1}^d \psi_i = \theta \bar{a} \quad \text{for a } \theta \in (0, 1).$$

More precisely, it is shown in [5] that the partial sums $u_k(y)$ in a Neumann series expansion of $u(y)$ can be represented as

$$(4.25) \quad u_k(y) = \sum_{j=0}^{n(d-1, k)} \phi_{k,j}(y) v_j,$$

for some d -variate polynomials $\phi_{k,j}$ and $v_l \in V$, and therefore have at most rank $n(d-1, k) \leq (k+1)^{d-1}$.

Estimates of the type (4.23) give a first justification for the use of sub-exponential decay rates as benchmark rates. On the other hand, the obtained rates weaken substantially with increasing d . This has led to a careful study of more specialized types of parametric expansions to better understand whether and to what extent low-rank approximation may substantially outperform approximations based on an *a priori* fixed expansion system such as tensor product Taylor or Legendre expansions.

An important model case are parametric expansions with *piecewise constant* coefficients of the form

$$(4.26) \quad a(y) = \bar{a} + \sum_{j=1}^d y_j \theta_j,$$

where $d < \infty$, $\bar{a} = 1$, $\theta_j = b_j \chi_{D_j}$, $b_j \in (0, 1)$, and the subdomains $D_j \subset \Omega$ form a partition of Ω .

Specifically, consider first the spatially one-dimensional case $\Omega = (0, 1)$. As shown in [5, 6], for any $f \in H'$, $u(y)$ has then finite rank at most $2d - 1$. In fact, $u(y)|_{D_i}$ must be spanned for each y by $\chi_I, x\chi_I, F_i\chi_I$ where $F_i'' = f$ on D_i and F_i vanishes at the end points of D_i . Boundary conditions and $d - 1$ continuity conditions at the interior breakpoints leave $2d - 1$ degrees of freedom so that $u(y)$ is contained in a linear space of dimension at most $2d - 1$. Since therefore $d_{2d-1}(u(Y))_V = 0$, $u(y)$ has at most rank $2d - 1$.

For higher spatial dimensions, M_u is finite only in exceptional cases. For piecewise constant diffusion coefficients of the form (4.26) with $\theta_j = \bar{a}\theta\chi_{D_j}$, significantly refined information is derived in [5] by decomposing $H_0^1(\Omega)$ into $\bigoplus_{i=1}^d H_0^1(D_i)$ and its orthogonal complement W with respect to $\langle \cdot, \cdot \rangle_{\bar{a}}$. Employing corresponding orthogonal projections, one can write

$$u(y) = u_W(y) + \sum_{i=1}^d u_i(y),$$

where the $u_i(y)$ are solutions to

$$(1 + \theta y_i) \int_{D_i} \bar{a} \nabla u_i(y) \cdot \nabla v \, dx = \int_{D_i} f v \, dx, \quad v \in H_0^1(D_i), \quad u(y)|_{\Omega \setminus D_i} = 0.$$

Since the u_i have rank one depending only on y_i the portion $\sum_{i=1}^d u_i(y)$ has at most rank d . To determine the rank of u_W , it is written as the harmonic extension of a skeleton component u_Γ . It is then shown that $u_\Gamma(y)$ has a Neumann series representation which ultimately requires estimating the ranks of its partial sums $u_{\Gamma,k}(y)$.

In the case $d = 4$ and D_i being congruent squares, this analysis is carried out showing that in this case the ranks of the partial sums $u_k(y)$ of the Taylor expansion of $u(y)$ are bounded by $8k + 5$, which then leads to a bound $d_n(u(Y))_V \leq Ce^{-cn}$. Hence, for each $n \in \mathbb{N}$ one can find functions $u_k^x \in H_0^1(\Omega)$, $u_k^y \in L_2(Y, \mu)$ such that

$$(4.27) \quad \left\| u - \sum_{k=1}^n u_k^x \otimes u_k^y \right\|_{L_2(Y, H_0^1(\Omega))} \lesssim e^{-cn}.$$

Remarkably, numerical experiments reveal that for partitions into four non-congruent quadrilaterals, corresponding singular values decay at a significantly slower, yet still subexponential rate.

Moreover, numerical experiments support a similar behavior for larger d , with c in the estimate (4.27) depending weakly (approximately as $1/d$) on d . Very similar results are obtained for hierarchical tensor approximations, where also the intermediate ranks for matricizations with respect to the parametric variables enter. While to our knowledge, there are no theoretical bounds on these ranks, the numerical results indicate that the ranks for the spatial matricizations dominate and behave essentially as in the previous case where the parametric variables are not separated.

4.2.2. Infinitely many parameters. The case $d = \infty$, where we consider diffusion coefficients $a(y) = \bar{a} + \sum_{i=1}^{\infty} y_i \psi_i$ parameterized by $y \in Y = (-1, 1)^{\mathbb{N}}$, turns out to be significantly different. A tensor product basis for $L_2(Y, \mu)$ is provided

by the product Legendre polynomials $L_\nu(y) = \prod_{i=1}^{\infty} L_{\nu_j}(y_j)$ for $\nu \in \mathcal{F} := \{\nu \in \mathbb{N}_0^{\mathbb{N}} : \nu_j = 0 \text{ for almost all } j \in \mathbb{N}\}$, giving rise to the orthonormal polynomial basis expansion

$$(4.28) \quad u(y) = \sum_{\nu \in \mathcal{F}} u_\nu L_\nu(y), \quad u_\nu = \int_Y u(y) L_\nu(y) d\mu(y).$$

Also here, a natural first question concerns the separation of spatial and parametric variables as in (4.27), that is, the approximation of u by functions of the form $\sum_{k=1}^n u_k^x \otimes u_k^y$. The smallest possible error $\delta_n(u, \mu)_V = (\sum_{k>n} \sigma_k^2)^{1/2}$ is attained by the truncated Hilbert-Schmidt decomposition of M_u and thus determined by the decay of the singular values $(\sigma_k)_{k \in \mathbb{N}}$ of M_u .

Truncations of the expansion (4.28) provide alternative low-rank approximations, where one takes $u_k^x := u_{\nu_k}$, $u_k^y := L_{\nu_k}$ with $(\nu_k)_{k \in \mathbb{N}}$ chosen such that $\|u_{\nu_k}\|_{H^1(\Omega)} \geq \|u_{\nu_{k+1}}\|_{H^1(\Omega)}$, which yields an approximation error $(\sum_{k>n} \|u_{\nu_k}\|_{H^1(\Omega)}^2)^{1/2}$ with n terms. By the best approximation property of the Hilbert-Schmidt decomposition, clearly $\sum_{k>n} \|u_{\nu_k}\|_{H^1(\Omega)}^2 \geq \sum_{k>n} \sigma_k^2$.

The decay of $\|u_{\nu_k}\|_{H^1(\Omega)}$ as $k \rightarrow \infty$, which depends mainly on the functions ψ_j and their decay as $j \rightarrow \infty$, is well studied, see e.g. [7, 21–23]. One obtains results of the type $(\|u_\nu\|_{H^1(\Omega)})_{\nu \in \mathcal{F}} \in \ell_p(\mathcal{F})$ for some $p \in (0, 2)$, which implies $(\sum_{k>n} \|u_{\nu_k}\|_{H^1(\Omega)}^2)^{1/2} \lesssim n^{-1/p+1/2}$. There are only very limited results on the decay of σ_k in this setting; it depends also, e.g., on the right hand side f .

The following result on $(\sigma_k)_{k \in \mathbb{N}}$ was obtained in [6, §6], adapting arguments from [7, §4.1]: For $\Omega = (0, 1)$ and a certain class of ψ_j such that $u(y)$ is easy to handle analytically, there always exists $f \in H^{-1}(\Omega)$ such that $(\|u_\nu\|_{H^1(\Omega)})_{\nu \in \mathcal{F}} \in \ell_p(\mathcal{F})$ and $(\sigma_k)_{k \in \mathbb{N}} \notin \ell_{p'}(\mathbb{N})$ for any $p' < p$. In other words, with appropriate f , the singular values σ_k have the *same asymptotic decay* as the decreasingly ordered Legendre coefficient norms $\|u_{\nu_k}\|_{H^1(\Omega)}$. This also yields lower bounds on the Kolmogorov widths $d_n(u(Y))_{H^1(\Omega)} \geq \delta_n(u, \mu)_{H^1(\Omega)} \gtrsim n^{-1/p+1/2}$ in these settings. Thus, the performance concerning the number of terms n attainable by any low-rank approximation is, for an adverse choice of data f , asymptotically no better than that of the simpler Legendre expansion (4.28).

As numerical experiments in [6] show, this situation is also observed in generic examples with typical choices of ψ_j and f , where this has not been proved analytically. In summary, in such problems with infinitely many parameters, one can therefore only expect limited gains from optimized low-rank approximations in comparison to Legendre expansions (4.28). Indeed, as explained in more detail in Section 5, corresponding SVD-based computational methods are generally more expensive, since they scale nonlinearly with respect to the expansion ranks.

5. Adaptive error control and computational complexity

A central objective in adaptive methods for high-dimensional problems is the reliable quantification of errors. For scenario (I), one approach to determine rigorous error bounds has been recently proposed in [28] employing the *Hypercircle Principle* in combination with flux-approximations in tensor-train format. However, lower bounds do not seem to be available in this format. Hence, while the involved constant is just one, a quantification of a possible overestimation and its dependence on d seems to remain unclear. Moreover, the approach does not offer any provision for updating the given discretization based on the computed bounds. When

estimating instead residuals directly, as suggested in this article, one issue is the quality of error estimates, in that one needs to ensure that all involved constants are controlled and the resulting bounds remain meaningful for large dimension d . For instance, in naive generalizations of lower-dimensional error estimators, constants in the estimates may depend exponentially on d . A second issue is the numerical cost of their evaluation, which should also remain under control for large d .

The approach that we follow here addresses both issues in conjunction with strategies for updating the approximations. It is based on the proportionality between errors and residuals (1.4) and approximation of these infinite-dimensional residuals. To ensure the quality of error estimates for large d , we rely, in particular, on expressing the underlying problem in terms of suitable basis functions. Achieving near-optimal computational cost of the residual approximation prohibits, of course, ever accessing the underlying full coefficient tensor but requires a somewhat specialized treatment in each of the problem scenarios that we consider.

The resulting methods for error estimation can, in principle, be used in various different constructions of adaptive solvers. In Section 5.2, we combine them with a scheme for which we can subsequently show near-optimal asymptotic computational costs.

5.1. Residual approximations. The accuracy-controlled approximation of residuals $\mathbf{A}\mathbf{v} - \mathbf{f}$ for given finitely represented \mathbf{v} , needed in any realization of (3.10), requires some means of approximating \mathbf{f} in the desired form, and more crucially a scheme for approximately applying the infinite matrix \mathbf{A} . Devising such a routine with provably optimal performance, one faces different obstructions in the two problem scenarios (I), (II).

(I) *Diffusion equations.* Recall that an important requirement on a Riesz basis Φ for the energy space V is that it is a rescaled version of an orthonormal basis Ψ for $L_2(\Omega)$; the latter property is crucial for avoiding constants in the estimates that deteriorate exponentially in d . In this case the (unbounded) representation $\mathbf{T} := S'_\Psi AS_\Psi$ preserves the (formal) low-rank structure of the diffusion operator while the scaled version $\mathbf{A} = S'_\Phi AS_\Phi = \tilde{\mathbf{S}}_\infty^{-1} \mathbf{T} \tilde{\mathbf{S}}_\infty^{-1}$ from (3.16) has infinite rank. A key component of a computable residual approximation in low-rank format thus is to replace the ideal scaling matrix $\tilde{\mathbf{S}}_\infty^{-1}$ by an approximation $\tilde{\mathbf{S}}_m^{-1}$. Here, $m = m(\mathbf{v}, \eta)$ depends through (3.15) on the support of the vector \mathbf{v} , which \mathbf{A} is to be applied to, as well as on the target accuracy η . This ensures that the matrix $\mathbf{A}_{m(\mathbf{v}, \eta)} := \tilde{\mathbf{S}}_m^{-1} \mathbf{T} \tilde{\mathbf{S}}_m^{-1}$ has now finite rank while, on account of Theorem 3.1, the restriction to the activated basis elements is still well-conditioned and thus convergence of (3.10) is still ensured.

The matrix $\mathbf{A}_{m(\mathbf{v}, \eta)}$ is still infinite and its approximate application makes essential use of the *near-sparsity* or *compressibility* of properly rescaled versions of the low-dimensional blocks $\mathbf{T}_{n_i}^{(i)}$ of \mathbf{T} . Exploiting compressibility of such scaled blocks, in turn, is facilitated by the near-separability of the low-rank scaling matrices.

To indicate how this can be used to approximate the infinite-dimensional residual $\mathbf{A}\mathbf{u}^n - \mathbf{f}$, whose ℓ_2 -norm is uniformly proportional to the error $\|\mathbf{u} - \mathbf{u}^n\|_{\ell_2(\mathcal{T})}$, within a suitably updated dynamic accuracy tolerance η_n , we recall the notion of *s*-compressibility* which plays a pivotal role in adaptive wavelet methods, [19, 20]. An infinite matrix $\mathbf{B} \in \mathbb{R}^{\vee \times \vee}$ is called *s*-compressible* if for any $0 < s < s^*$, there exist summable positive sequences $(\alpha_j)_{j \geq 0}$, $(\beta_j)_{j \geq 0}$ and for each $j \geq 0$,

there exists \mathbf{B}_j with at most $\alpha_j 2^j$ nonzero entries per row and column, such that $\|\mathbf{B} - \mathbf{B}_j\| \leq \beta_j 2^{-sj}$.

To approximate $\mathbf{B}\mathbf{v}$ for a given $\mathbf{v} \in \ell_2(\mathcal{V})$ one combines *a priori* knowledge about the matrix \mathbf{B} with *a posteriori* information on the input \mathbf{v} , e.g., by partitioning the support of \mathbf{v} into disjoint subsets Λ_j corresponding to the magnitude of its entries in decreasing order. We then define a \mathbf{v} -dependent approximation $\tilde{\mathbf{B}}_J$ of \mathbf{B} ,

$$(5.1) \quad \mathbf{B}\mathbf{v} \approx \tilde{\mathbf{B}}_J \mathbf{v} := \sum_{j=0}^J \mathbf{B}_{J-j}(\mathbf{v}|_{\Lambda_j}),$$

which satisfies

$$\begin{aligned} \|\mathbf{B}\mathbf{v} - \tilde{\mathbf{B}}_J \mathbf{v}\| &= \left\| \sum_{j=0}^J (\mathbf{B} - \mathbf{B}_{J-j})(\mathbf{v}|_{\Lambda_j}) + \mathbf{B}(\mathbf{v}|_{\Lambda_{J+1}}) \right\| \\ &\leq \sum_{j=0}^J \|\mathbf{B} - \mathbf{B}_{J-j}\| \|\mathbf{v}|_{\Lambda_j}\| + \|\mathbf{B}\| \|\mathbf{v}|_{\Lambda_{J+1}}\|. \end{aligned}$$

Note that for j close to J , \mathbf{B}_{J-j} is a coarse approximation, whereas the most accurate approximations of \mathbf{B} are applied to the large contributions in \mathbf{v} .

For high-dimensional input tensors $\mathbf{v} \in \ell_2(\mathcal{V}^d)$, the crucial step is then to apply approximations of each lower-dimensional component in the tensor decomposition of \mathbf{T} separately in each tensor mode, with accuracies controlled by subdividing the sequences $\pi^{(i)}(\mathbf{v})$ for each mode $i = 1, \dots, d$. As shown in [11], one can then combine this with the approximate rescaling in a routine $\text{APPLY}(\mathbf{v}; \eta) : \mathbf{v} \rightarrow \mathbf{w}_\eta$ with

$$(5.2) \quad \mathbf{w}_\eta := \tilde{\mathbf{S}}_{m(\eta; \mathbf{v})}^{-1} \tilde{\mathbf{T}}_{J(\eta)} \tilde{\mathbf{S}}_{m(\eta; \mathbf{v})}^{-1} \mathbf{v},$$

for exponential sum truncation ranks $m(\eta; \mathbf{v})$ chosen according to (3.15) and compression depth $J(\eta)$, such that $\|\mathbf{A}\mathbf{v} - \mathbf{w}_\eta\| \leq \eta$.

Instead of the energy norm, one can instead measure accuracy in a weaker norm such as the $L_2(\Omega)$ -norm. Since this is a cross-norm one would intuitively expect achieving desired accuracy thresholds at the expense of lower ranks. This is indeed the case as shown in [10], where rigorous error control in $L_2(\Omega)$ is established again in the framework of an equivalent problem formulation in sequence space. The relevant mapping properties are now realized through an *unsymmetric* preconditioning, which requires only a single scaling operation in each application of \mathbf{A} of the form

$$(5.3) \quad \mathbf{A}_{m(\mathbf{v}, \eta)} = \mathbf{S}_{\Phi_{m(\mathbf{v}, \eta)}}^{-2} \mathbf{T},$$

see [10] for further details.

(II) *Parametric PDEs.* As indicated at the beginning of Section 3.4, one can tensorize any Riesz-basis Φ of $H_0^1(\Omega)$ with an orthonormal basis of $L_2(Y, \mu)$. A natural choice for this orthonormal basis are the product Legendre polynomials as in Section 4.2.2, with an orthonormal basis given by $\{L_\nu\}_{\nu \in \mathcal{G}}$, where $\mathcal{G} = \mathbb{N}_0^d$ for $d < \infty$ and by $\mathcal{G} = \mathcal{F}$ for $d = \infty$. The resulting system $\Phi := \{\varphi_\lambda \otimes L_\nu\}_{(\lambda, \nu) \in \mathcal{S} \times \mathcal{G}}$ is then a Riesz basis of $L_2(Y, H_0^1(\Omega))$, which we now use to transform (3.3) into an equivalent system over $\ell_2(\mathcal{S} \times \mathcal{G})$. Defining then with $\psi_0 := \bar{a}$ for $j \in \mathbb{N}_0$

$$(5.4) \quad \mathbf{A}_j := (\langle \psi_j \nabla \varphi_{\lambda'}, \nabla \varphi_\lambda \rangle)_{\lambda, \lambda' \in \mathcal{S}}, \quad \mathbf{f} := (\langle f, \varphi_\lambda \otimes L_\nu \rangle)_{(\lambda, \nu) \in \mathcal{S} \times \mathcal{G}},$$

as well as

$$(5.5) \quad \mathbf{M}_j := \left(\int_Y y_j L_\nu(y) L_{\nu'}(y) d\mu(y) \right)_{\nu, \nu' \in \mathcal{G}}, \quad j \geq 1,$$

and \mathbf{M}_0 as the identity on $\ell_2(\mathcal{G})$, we obtain an equivalent problem

$$(5.6) \quad \mathbf{A}\mathbf{u} = \mathbf{f}, \quad \text{where} \quad \mathbf{A} := \sum_{j=0}^d \mathbf{A}_j \otimes \mathbf{M}_j,$$

on $\ell_2(\mathcal{S} \times \mathcal{G})$. Due to the classical three-term recurrence relations for Legendre polynomials, the matrices \mathbf{M}_j are *bi-diagonal* for $j \geq 1$.

In the case of fixed finite dimension $d < \infty$, the hierarchical tensor format is a natural choice. In fact, it poses significantly less difficulties than in the case of high-dimensional diffusion equations, since the mapping S_Φ is a Kronecker product and the sparse matrices \mathbf{M}_j need no additional compression. The matrices \mathbf{A}_j can be approximated similarly to (5.1), with accuracies controlled by norms of segments of $\|\pi^{(x)}(\mathbf{v})\|$ of the spatial contractions

$$\pi_\lambda^{(x)}(\mathbf{v}) := \left(\sum_{\nu \in \mathcal{G}} |\mathbf{v}_{\lambda, \nu}|^2 \right)^{1/2}, \quad \lambda \in \mathcal{S},$$

for inputs $\mathbf{v} \in \ell_2(\mathcal{S} \times \mathcal{G})$, which can be computed from the *HSVD* as in (3.18). As shown in the next section, corresponding complexity estimates are slightly more favorable than for scenario (I).

In problems with $d = \infty$, residual approximations are more complex due to the additional requirement of *adaptive truncation of dimensionality*: to ensure efficiency with controlled errors, the number of parameters y_j that are activated needs to be adjusted to discretization and low-rank approximation errors. This can be controlled adaptively, taking into account the known decay properties of the parameter expansion functions $\|\psi_j\|_{L_\infty(\Omega)}$, via estimates for the operator truncation error of the type

$$(5.7) \quad e_J := \left\| \sum_{j>J} \mathbf{A}_j \otimes \mathbf{M}_j \right\| \leq CJ^{-S},$$

with some fixed $S > 0$. Such an adaptation has been considered in [33, 34] for generic ψ_j , but as observed in [6], the value of S is strongly influenced by structural features of this expansion. A particularly favorable case are expansions with a multilevel structure, where each j is uniquely associated to a pair $(\ell(j), k(j))$ with inverse mapping $j(\ell, k)$, such that there are $\mathcal{O}(2^{m\ell})$ functions on level ℓ and

$$(5.8) \quad \sum_k |\psi_{j(\ell, k)}| \leq C2^{-\alpha\ell}, \quad \ell \geq 0,$$

for some $\alpha > 0$, with a uniform constant $C > 0$. Here $\ell(j) \in \mathbb{N}_0$ describe the scale and $k(j) \in \mathbb{Z}^m$ (with $\Omega \subset \mathbb{R}^m$) the spatial localization of ψ_j . Such a condition is satisfied in particular when $\text{diam}(\text{supp}(\psi_j)) \sim 2^{-\ell(j)}$, there exists $M > 0$ such that for each ℓ, k , the number $\#\{k' : \text{supp} \psi_{j(\ell, k)} \cap \text{supp} \psi_{j(\ell, k')} \neq \emptyset\} \leq M$, and $\|\psi_j\|_{L_\infty(\Omega)} \sim c_j \sim 2^{-\alpha\ell(j)}$. Expansions with the property (5.8) can be constructed for wide classes of random fields, see [8], where one obtains $S = \alpha/m$.

In principle, for any fixed set of modes x, y_1, \dots, y_J , one can apply the hierarchical format with $J + 1$ tensor modes, which leads to some substantial additional

difficulties in the complexity analysis due to the changing dimensionality. We therefore restrict the following discussion to the basic case of separating spatial indices in \mathcal{S} and parametric indices in \mathcal{F} as in Section 4.2.2, that is, to estimating errors in low-rank approximations to \mathbf{u} of the form

$$(5.9) \quad \bar{\mathbf{u}} = \sum_{k=1}^r \bar{\mathbf{U}}_k^x \otimes \bar{\mathbf{U}}_k^y,$$

with \mathbf{U}_k^x and $\bar{\mathbf{U}}_k^y$ finitely supported on \mathcal{S} and \mathcal{F} , respectively. In view of the results outlined in Section 4.2.2, direct Legendre expansion can be advantageous compared with (5.9) regarding asymptotic costs. Nevertheless, despite the fact that the number of terms r is possibly not significantly smaller than for sparse Legendre expansions, and the total computational cost may in fact be higher, an efficient computation of low-rank approximations with controlled error in $L_2(Y, H_0^1(\Omega))$ and near-optimal basis functions as in (5.9) can still be of interest in many application scenarios, for instance in the case of many parameter queries.

In the case of a low-rank approximation of the residual with simultaneous sparse approximation in the tensor modes, the decomposition of \mathbf{v} is more involved. For $p, q = 0, 1, 2, \dots$ an input \mathbf{v} is then decomposed further into segments $\mathbf{v}_{[p,q]}$ of rank 2^p , where p refers to consecutive groups of 2^p singular values of \mathbf{v} , and q refers to the best 2^q -term approximations of the parametric contractions $\pi^{(y)}(\mathbf{v})$, defined by

$$\pi_\nu^{(y)}(\mathbf{v}) := \left(\sum_{\lambda \in \mathcal{S}} |\mathbf{v}_{\lambda,\nu}|^2 \right)^{1/2}.$$

The routine $\text{APPLY}(\mathbf{A}, \mathbf{v}; \eta) \rightarrow \mathbf{w}_\eta$ then yields finitely supported low-rank approximations

$$(5.10) \quad \mathbf{w}_\eta := \sum_{p,q \geq 0} \sum_{j=1}^{J_{p,q}(\eta)} (\tilde{\mathbf{A}}_{j,p,q} \otimes \mathbf{M}_j) \mathbf{v}_{[p,q]},$$

with $\|\mathbf{w}_\eta - \mathbf{A}\mathbf{v}\| \leq \eta$, where the number of terms $J_{p,q}(\eta)$ is determined in an a posteriori fashion from the norms $\|\mathbf{v}_{[p,q]}\|$ as well as known bounds for the truncation errors $e_{J_{p,q}}$ from (5.7). Moreover, $(\tilde{\mathbf{A}}_{j,p,q} \otimes \mathbf{M}_j) \mathbf{v}_{[p,q]}$ is an error-controlled approximation of $(\mathbf{A}_j \otimes \mathbf{M}_j) \mathbf{v}_{[p,q]}$ obtained as in (5.1). Here, the change in compressibility of \mathbf{A}_j as $j \rightarrow \infty$, which typically is driven by increasingly oscillatory features of the functions ψ_j , crucially determines the resulting computational costs. Also in this respect, a multilevel structure of the parametric expansion, as in (5.8), turns out to be highly advantageous, in that it also leads to improved compressibility of the matrices \mathbf{A}_j . In summary, under this condition one can construct the approximation (5.10) at near-optimal computational costs. For further details, we refer to [6].

REMARK 5.1. A different technique for residual approximation for such problems, using finite element discretizations in space but also relying on contractions $\pi^{(i)}(\mathbf{v})$ similarly to the present approach, was proposed in [30] and shown to provide an upper bound for the error. However, for this method no suitable complexity estimates are available.

5.2. Computational complexity of solvers. Independently of any specific solution algorithm, the residual approximations, discussed above, provide information on numerical errors in high-dimensional problems, where the costs of evaluating these approximations in terms of their quality depend only mildly on the dimensionality of the problem. Such residual approximations can, in particular, be naturally used in iterative solvers as discussed in Section 3.3, which gradually refine simultaneous sparse and low-rank approximations.

As indicated earlier, the residual approximations, as well as many further operations required in iterative solvers, increase the representation complexity of iterates, especially concerning their tensor representation ranks. The computational complexity of the resulting methods is therefore a crucial issue. Ideally, it should be directly related to the convergence of corresponding best approximations; in the case of low-rank approximations, however, one can generally not expect computational costs to remain strictly proportional to the number of degrees of freedom, since the costs of many required procedures such as SVD scale nonlinearly in the ranks.

A basic strategy for controlling approximation complexity, following the template (3.10), is to combine error reduction (based, in particular, on residual approximations) with a suitable complexity reduction. Whereas maintaining convergence to the exact solution is straightforward with the ingredients we have introduced so far, relating the computational complexity to the quality of best approximations is substantially more delicate. We now discuss two basic approaches that rely on hard and soft thresholding with judiciously chosen thresholding parameters.

5.2.1. Hard thresholding. This first approach is based on the observation that if an input v is known to be quantifiably close to an unknown object u then, using properly chosen thresholds, these routines produce a new approximation with a slightly larger error but with near-minimal ranks or representation complexities, respectively.

We will first frame the basic mechanisms underlying the reduction procedures in abstract terms. To this end, let H be a Hilbert space and let $\{\mathcal{S}_r\}_{r \in \mathbb{N}}$ be a family of subsets from which approximations are selected, where the parameter r controls their complexity. For $v \in H$ and $\eta > 0$, let

$$\bar{r}(v, \eta) := \min\{r : (\exists w \in \mathcal{S}_r : \|v - w\| \leq \eta)\},$$

which is the minimal complexity parameter for which the best approximation from these sets attains accuracy η . In addition, assume that there exists a (linear) orthogonal projection $\hat{P}(v, \eta)$ onto $\mathcal{S}_{\bar{r}(v, \eta)}$ such that $\|v - \hat{P}(v, \eta)v\| \leq \eta$.

LEMMA 5.2. *For any r , let $\hat{P}_r: H \rightarrow \mathcal{S}_r$ and $\hat{e}_r: H \rightarrow \mathbb{R}^+$ be such that*

$$\|v - \hat{P}_r(v)\| \leq \hat{e}_r(v),$$

where for some $\kappa > 0$,

$$\hat{e}_r(v) \leq \kappa \inf_{w \in \mathcal{S}_r} \|v - w\|.$$

For any $\alpha, \eta > 0$, and for any $u, v \in H$ with $\|u - v\| \leq \eta$, let \hat{r} be minimal such that $\hat{e}_{\hat{r}}(v) \leq \kappa(1 + \alpha)\eta$. Then, for $\hat{v} := \hat{P}_{\hat{r}}(v) \in \mathcal{S}_{\hat{r}}$ one has

$$(5.11) \quad \|u - \hat{v}\| \leq (1 + \kappa(1 + \alpha))\eta \quad \text{while} \quad \hat{r} \leq \bar{r}(u, \alpha\eta).$$

PROOF. Given $u, v, \eta > 0$ as above, by linearity and orthogonality of $\bar{P}(u, \alpha\eta)$, we have for any fixed $\alpha > 0$

$$(5.12) \quad \begin{aligned} \|v - \bar{P}(u, \alpha\eta)v\|_H &\leq \|(v - u) - \bar{P}(u, \alpha\eta)(v - u)\|_H \\ &+ \|u - \bar{P}(u, \alpha\eta)u\|_H \leq (1 + \alpha)\eta. \end{aligned}$$

On the other hand, we have for any r that

$$\|v - \hat{P}_r(v)\|_H \leq \kappa \inf_{w \in \mathcal{S}_r} \|v - w\|_H,$$

so that by (5.12), we have for $r = \bar{r}(v, \alpha\eta)$

$$\|v - \hat{P}_{\bar{r}(v, \alpha\eta)}(v)\|_H \leq \kappa \inf_{w \in \mathcal{S}_{\bar{r}(v, \alpha\eta)}} \|v - w\|_H \leq \kappa(1 + \alpha)\eta.$$

Since, by definition \hat{r} is the minimal index providing accuracy $\kappa(1 + \alpha)\eta$, we conclude that $\hat{r} \leq \bar{r}(v, \alpha\eta)$ so that the second relation in (5.11) follows, while the first relation holds by the triangle inequality. \square

Lemma 5.2 means roughly the following: suppose one has constructed an explicit finitely parametrized approximation v to an unknown u for which an error bound η is known. In the present context v is the result of an error-controlled reduction step \mathcal{F}_n in (3.10). Then, *re-approximating* the explicitly given v within a judiciously chosen tolerance, which is somewhat larger than the error bound η , provides a new approximation \hat{v} to the unknown u with a slightly worsened accuracy but with near-optimal complexity.

More precisely, the approximation complexity of the computed recompressed approximation \hat{v} in (5.11) can be quantified in terms of the corresponding best approximations of u . For instance, take Σ_r as defined in (1.5), with a fixed basis $\{\varphi_\lambda\}_{\lambda \in \mathbb{N}}$, as the approximation sets, where the complexity parameter is now the number of nonzero coefficients $r \in \mathbb{N}$. Both \bar{P} and \hat{P}_r correspond to retaining the coefficients of largest absolute value, and $\kappa = 1$. Then if $u = \sum_\lambda \mathbf{u}_\lambda \varphi_\lambda$ with $\mathbf{u} \in \mathcal{A}^s$, we have

$$\bar{r}(u, \varepsilon) \leq \|\mathbf{u}\|_{\mathcal{A}^s}^{\frac{1}{s}} \varepsilon^{-\frac{1}{s}},$$

and thus (5.11) gives $\|u - \hat{v}\| \leq \varepsilon$ with $\hat{v} \in \Sigma_{\hat{r}}$, where

$$\hat{r} \leq \left(\frac{2 + \alpha}{\alpha} \right)^{\frac{1}{s}} \|\mathbf{u}\|_{\mathcal{A}^s}^{\frac{1}{s}} \varepsilon^{-\frac{1}{s}}.$$

In this manner, this observation appears first in the context of sparse approximations with respect to a given basis, as in adaptive wavelet schemes [19, 20].

In the context of hierarchical tensor representations, r is the maximal entry in the tuple of hierarchical ranks. In the simplest case $d = 2$, where $H = H_1 \otimes H_2$, we choose \mathcal{S}_r as the elements of rank at most r , and \bar{P}, \hat{P}_r are given by the truncated Hilbert-Schmidt expansion. Then if $\sigma(u) \in \mathcal{A}^s$, in a completely analogous way we obtain $\|u - \hat{v}\| \leq \varepsilon$ with $\hat{v} \in \mathcal{S}_{\hat{r}}$ and

$$\hat{r} = \text{rank}(\hat{v}) \leq \left(\frac{2 + \alpha}{\alpha} \right)^{\frac{1}{s}} \|\sigma(u)\|_{\mathcal{A}^s}^{\frac{1}{s}} \varepsilon^{-\frac{1}{s}}.$$

In this setting, cases where low-rank approximations are particularly attractive are those where u satisfies

$$\inf_{\text{rank}(v) \leq r} \|u - v\| \leq C e^{-cr^\beta}$$

for some $C, c, \beta > 0$, corresponding to exponential-type decay of the singular value sequence $\sigma(u)$. In such cases we instead obtain

$$\hat{r} = \text{rank}(\hat{v}) \leq (c^{-1} \ln(C(2 + \alpha)\alpha^{-1}\varepsilon^{-1}))^{1/\beta} \lesssim |\ln \varepsilon|^{1/\beta}.$$

For higher-order hierarchical tensors, the near-optimal projection \hat{P}_r can be done by the routine RECOMPRESS implementing the HSVD hard thresholding operation where, by Theorem 2.3, $\kappa = \sqrt{2d - 3}$. Under an analogous assumption

$$\inf_{\max r_{\mathbb{E}}(v) \leq r} \|u - v\| \leq Ce^{-cr^\beta},$$

which means that $u \in \mathcal{A}_{\mathcal{H}}((e^{cr^\beta})_{r \in \mathbb{N}_0})$, one then obtains

$$\hat{r} = \max r_{\mathbb{E}}(\hat{v}) \leq \left(c^{-1} \ln(C(1 + (1 + \alpha)\kappa)\alpha^{-1}\varepsilon^{-1}) \right)^{1/\beta}.$$

In the application of Lemma 5.2 to the control of the componentwise sparsity in each tensor mode, as quantified by the contractions $\pi^{(i)}$ in (3.18), r is the total number of nonzero entries $\pi_{\lambda}^{(i)}$ for $i = 1, \dots, d$ and $\lambda \in \mathcal{I}_i$. Proposition 3.2 provides the corresponding quasi-best approximation procedure \hat{P}_r realized by COARSEN. As shown in [9], the two reduction procedures for hierarchical tensor decompositions and componentwise sparsity can be combined to achieve the quasi-optimality property (5.11) simultaneously with respect to both types of approximation.

A corresponding adaptive scheme can be realized as in Algorithm 5.1 using the routines COARSEN and RECOMPRESS described in Section 3.4.2, combined with residual approximations using routines RHS($\mathbf{f}; \eta$), APPLY($\mathbf{A}, \mathbf{v}; \eta$). These latter routines generate, for any given target tolerance $\eta > 0$, error-controlled finite-rank and finitely supported approximations such that

$$(5.13) \quad \|\mathbf{f} - \text{RHS}(\mathbf{f}; \eta)\|_{\ell_2(\mathcal{I})} \leq \eta, \quad \|\mathbf{A}\mathbf{v} - \text{APPLY}(\mathbf{A}, \mathbf{v}; \eta)\|_{\ell_2(\mathcal{I})} \leq \eta.$$

In all cases, the accuracy requirements are ensured by a posteriori bounds, *independently* of any prior assumptions on or knowledge of the given finitely parametrized input.

One step in Algorithm 5.1 takes the basic form

$$(5.14) \quad \bar{\mathbf{u}} \leftarrow \text{COARSEN} \left(\text{RECOMPRESS} \left(\bar{\mathbf{u}} - \omega \mathbf{r}_{\eta_A, \eta_f}(\bar{\mathbf{u}}); \eta_{rc} \right); \eta_{coa} \right),$$

where $\mathbf{r}_{\eta_A, \eta_f}(\bar{\mathbf{u}}) := \text{APPLY}(\mathbf{A}, \bar{\mathbf{u}}; \eta_A) - \text{RHS}(\mathbf{f}; \eta_f)$ with certain tolerances $\eta_A, \eta_f, \eta_{rc}, \eta_{coa}$ of comparable size and geometric decay. The algorithm has an inner loop, where steps (5.14) are applied to achieve an error reduction, corresponding to the mapping \mathcal{F}_n in (3.10). In the outer loop, this is followed by a complexity reduction by RECOMPRESS and COARSEN with sufficiently large tolerances, see step 11. As shown in [9], Lemma 5.2 can be applied to infer joint quasi-optimality of hierarchical ranks and mode frame supports provided that the parameters $\kappa_1, \kappa_2, \kappa_3$, which control the relative size of error and complexity reduction tolerances, are chosen as

$$(5.15) \quad \begin{aligned} \kappa_1 &= \left(1 + (1 + \alpha)(\sqrt{2d - 3} + \sqrt{d} + \sqrt{(2d - 3)d}) \right)^{-1}, \\ \kappa_2 &= \sqrt{2d - 3}(1 + \alpha)\kappa_1, \quad \kappa_3 = \sqrt{d}(\sqrt{2d - 3} + 1)(1 + \alpha)\kappa_1. \end{aligned}$$

REMARK 5.3. It is easy to see that for any given tolerance ε , Algorithm 5.1 terminates after $\mathcal{O}(|\ln \varepsilon|)$ steps, producing a coefficient array \mathbf{u}_ε such that $\|\mathbf{u} - \mathbf{u}_\varepsilon\|_{\ell_2(\mathcal{I})} \leq \varepsilon$. We emphasize that convergence of the algorithm is guaranteed

Algorithm 5.1 $\mathbf{u}_\varepsilon = \text{SOLVE}(\mathbf{A}, \mathbf{f}; \varepsilon)$

```

input   $\begin{cases} \omega > 0 \text{ and } \rho \in (0, 1) \text{ such that } \|\mathbf{I} - \omega \mathbf{A}\| \leq \rho, \\ c_{\mathbf{A}} \geq \|\mathbf{A}^{-1}\|, \varepsilon_0 \geq c_{\mathbf{A}} \|\mathbf{f}\|, \\ \kappa_1, \kappa_2, \kappa_3 \in (0, 1) \text{ with } \kappa_1 + \kappa_2 + \kappa_3 \leq 1, \text{ and } \beta_1 \geq 0, \beta_2 > 0. \end{cases}$ 
output  $\mathbf{u}_\varepsilon$  satisfying  $\|\mathbf{u}_\varepsilon - \mathbf{u}\| \leq \varepsilon$ .
1:  $\mathbf{u}_0 := 0$ 
2:  $k := 0, I := \min\{j : \rho^j(1 + (\omega + \beta_1 + \beta_2)j) \leq \frac{1}{2}\kappa_1\}$ 
3: while  $2^{-k}\varepsilon_0 > \varepsilon$ 
4:    $\mathbf{w}_{k,0} := \mathbf{u}_k, j \leftarrow 0$ 
5:   while  $j < I$ 
6:      $\eta_{k,j} := \rho^{j+1}2^{-k}\varepsilon_0$ 
7:      $\mathbf{r}_{k,j} := \text{APPLY}(\mathbf{w}_{k,j}; \frac{1}{2}\eta_{k,j}) - \text{RHS}(\frac{1}{2}\eta_{k,j})$ 
8:      $\mathbf{w}_{k,j+1} := \text{COARSEN}(\text{RECOMPRESS}(\mathbf{w}_{k,j} - \omega \mathbf{r}_{k,j}; \beta_1 \eta_{k,j}); \beta_2 \eta_{k,j})$ 
9:      $j \leftarrow j + 1$ .
10:  end while
11:   $\mathbf{u}_{k+1} := \text{COARSEN}(\text{RECOMPRESS}(\mathbf{w}_{k,j}; \kappa_2 2^{-(k+1)}\varepsilon_0); \kappa_3 2^{-(k+1)}\varepsilon_0)$ 
12:   $k \leftarrow k + 1$ 
13: end while
14:  $\mathbf{u}_\varepsilon := \mathbf{u}_k$ 

```




FIGURE 1. Illustration of Algorithm 5.1, for $d = 2$ with coefficient arrays determined by sums of outer products of sparse vectors: starting from \mathbf{u}_k , in the inner loop (steps 2–11), ranks are increased and degrees of freedom are activated (shown in red) in $\mathbf{w}_{k,j}$ for $j = 1, \dots, J$. Small coefficients are removed in the re-approximation step of line 11 to obtain \mathbf{u}_{k+1} .

independently of any additional conditions on the unknown solution array \mathbf{u} . The performance of the algorithm, i.e., the numerical cost of computing \mathbf{u}_ε , can, of course, only be quantified under certain assumptions on the solution.

5.2.2. *Complexity bounds in model cases.* Guided by the findings in Section 4, we proceed to formulate the relevant assumptions and *benchmark* properties for the problem scenarios (I) and (II) for which we will quantify the performance of Algorithm 5.1. For a detailed account we refer to [10, 11], in particular regarding the dependence on d of various constants.

In this regard, similar assumptions are suitable for the model problems in scenario (I) and in scenario (II) with $d < \infty$, since in both cases we expect

$$(5.16) \quad \mathbf{u} \in \mathcal{A}_H(\gamma), \quad \gamma(n) = e^{cn^\beta}, \quad \text{for some } c, \beta > 0.$$

For the sparse approximability in each tensor mode, we may assume

$$(5.17) \quad \pi^{(i)}(\mathbf{u}) \in \mathcal{A}^s, \quad i = 1, \dots, d, \quad \text{for some } s > 0.$$

Then for any $\varepsilon > 0$, the result of Algorithm 5.1 after the final complexity reduction step satisfies³, as a consequence of Lemma 5.2,

$$(5.18) \quad \max r_{\mathbb{E}}(\mathbf{u}_\varepsilon) \lesssim (|\ln \varepsilon| + \ln d)^{1/\beta},$$

$$(5.19) \quad \sum_{i=1}^d \#\text{supp}(\pi^{(i)}(\mathbf{u}_\varepsilon)) \lesssim d^{1+s^{-1}} \left(\sum_{i=1}^d \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^s} \right)^{\frac{1}{s}} \varepsilon^{-\frac{1}{s}}.$$

Note that $\#\text{supp}(\pi^{(i)}(\mathbf{u}_\varepsilon))$ is the number of degrees of freedom in tensor mode i . Thus, in other words, ranks and discretization sizes in the result are *quasi-optimal*. Moreover, one has

$$(5.20) \quad \|\mathbf{u}_\varepsilon\|_{\mathcal{A}_{\mathcal{H}}(\gamma_{\mathbf{u}})} \lesssim \sqrt{d} \|\mathbf{u}\|_{\mathcal{A}_{\mathcal{H}}(\gamma_{\mathbf{u}})},$$

$$(5.21) \quad \sum_{i=1}^d \|\pi^{(i)}(\mathbf{u}_\varepsilon)\|_{\mathcal{A}^s} \lesssim d^{1+\max\{1,s\}} \sum_{i=1}^d \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^s}.$$

For the proofs of the estimates (5.18), (5.19), (5.20), (5.21), see [9, 11]. Note that these bounds hold after every execution of the outer loop in Algorithm 5.1, with ε replaced by the accuracy $2^{-k}\varepsilon_0$ of the current iterate \mathbf{u}_k .

On this basis, we obtain estimates for the total computational complexity of the adaptive scheme, which depends mainly on the computational costs of residual approximations in the inner loop. At this point, we need to exploit structural properties of \mathbf{A} . Specifically, we refer to [6, 10, 11] where s^* -compressibility of the tensor components of \mathbf{A} has been established. The value of s^* depends on diffusion coefficients, the choice of the Riesz-basis, and in scenario (II), on the properties of the parametric expansions. Furthermore, we need to make some additional technical assumptions concerning approximability of \mathbf{f} , including $\mathbf{f} \in \mathcal{A}_{\mathcal{H}}(\gamma)$. These properties of the problem data can, in principle, be verified.

In case (I), for the bound of the complexity of $\text{APPLY}(\mathbf{A}, \mathbf{v}; \eta) \rightarrow \mathbf{w}_\eta$, we assume in addition that $u \in H^t(\Omega)$ for some $t > 1$. One can then bound the ranks $\hat{m}(\eta, \mathbf{v})$ required for the approximate rescaling $\tilde{\mathbf{S}}_{m(\mathbf{v}, \eta)}^{-1}$. Then one obtains estimates of the form

$$(5.22) \quad r_i(\mathbf{w}_\eta) \leq (\hat{m}(\eta, \mathbf{v}))^2 R r_i(\mathbf{v}), \quad i = 1, \dots, E,$$

where R bounds the representation ranks of \mathbf{T} (for instance, $R = 2$ in the case $A = -\Delta$). Moreover, whenever $s < s^*$,

$$(5.23) \quad \begin{aligned} \#\text{supp}(\pi^{(i)}(\mathbf{w}_\eta)) &\lesssim \left(\sum_{j=1}^d \|\pi^{(j)}(\mathbf{v})\|_{\mathcal{A}^s} \right)^{1/s} \eta^{-1/s}, \\ \|\pi^{(i)}(\mathbf{w}_\eta)\|_{\mathcal{A}^s} &\lesssim d \|\pi^{(i)}(\mathbf{v})\|_{\mathcal{A}^s}, \end{aligned}$$

³If $\beta \leq 1$, the estimate (5.18) can be quantified more precisely in terms of $\|\mathbf{u}\|_{\mathcal{A}_{\mathcal{H}}(\gamma)}$; see [9].

with hidden constants depending in particular on \mathbf{A} ; for details, see [11, §6.2]. Thus, for a certain approximability range $s < s^*$, the approximate operator application preserves joint mode frame sparsity.

With (5.22), (5.23), and taking into account that the costs of each iteration in the inner loop are dominated by those for the *HSVD* of iterates and residual approximations, we arrive at the bound

$$(5.24) \quad \text{ops}(\mathbf{u}_\varepsilon) \lesssim d^{C_1 \ln d} (1 + |\ln \varepsilon|)^{C_2 \ln d + 2/\beta} \varepsilon^{-\frac{1}{s}}$$

for the total computational costs of Algorithm 5.1, where $C_1, C_2 > 0$ are constants *independent* of d and ε . Similar (and slightly more favorable) estimates hold for (I) with error control in L_2 , see [10].

In scenario (II) with $d < \infty$, the situation is similar, but instead of (5.22) one has the η -independent bound $r_i(\mathbf{w}_\eta) \leq (d+1)r_i(\mathbf{v})$, $i = 1, \dots, E$. As a consequence, the complexity bound improves to

$$(5.25) \quad \text{ops}(\mathbf{u}_\varepsilon) \lesssim d^{C \ln d} |\ln \varepsilon|^{2/\beta} \varepsilon^{-\frac{1}{s}},$$

with $C > 0$ independent of d and ε .

Concerning the dependence on d , the bounds (5.24), (5.25) are in general far from sharp: also in the inner loop of Algorithm 5.1, the ranks of iterates typically remain lower than guaranteed by the available upper bounds. However, the estimates (5.24), (5.25) still show a growth in d that is far below exponential, and thus ensure that low-rank methods can indeed break the curse of dimensionality for these problems also in terms of their total computational cost.

In scenario (II) with $d = \infty$, we again consider low-rank approximation with a single separation between spatial and parametric degrees of freedom as in (5.9). Assuming that the parametric expansion functions ψ_j have multilevel structure according to (5.8), one immediately obtains from the results in [7] that $\pi^{(x)}(\mathbf{u}), \pi^{(y)}(\mathbf{u}) \in \mathcal{A}^s$ for any $s < \alpha/m$. In view of the discussion in 4.2.2, for the singular values $\sigma(\mathbf{u})$ of \mathbf{u} , this implies $\sigma(\mathbf{u}) \in \mathcal{A}^s$ for any $s < \alpha/m$. Unfortunately, one cannot in general expect stronger summability of $\sigma(\mathbf{u})$; that is, in this setting the singular values of \mathbf{u} generally decay only *algebraically*.

In this case, as a consequence of Lemma 5.2, one has

$$(5.26) \quad \text{rank}(\mathbf{u}_\varepsilon) \lesssim \varepsilon^{-\frac{1}{s}} \|\sigma(\mathbf{u})\|_{\mathcal{A}^s}^{\frac{1}{s}}, \quad \|\sigma(\mathbf{u}_\varepsilon)\|_{\mathcal{A}^s} \lesssim \|\sigma(\mathbf{u})\|_{\mathcal{A}^s},$$

as well as estimates for $\#\text{supp } \pi^{(i)}(\mathbf{u}_\varepsilon)$, $i = x, y$, analogous to (5.19), (5.21).

Under certain assumptions on the parametric representation of the diffusion coefficient, the compressibility on \mathbf{A}_j is quantified in [6] based on which $\text{APPLY}(\mathbf{v}; \eta) \rightarrow \mathbf{w}_\eta$, as in (5.10), is shown to exhibit near optimal performance also in this case. Specifically, one has

$$\begin{aligned} \text{rank}(\mathbf{w}_\eta) &\lesssim \eta^{-\frac{1}{s}} \|\sigma(\mathbf{v})\|_{\mathcal{A}^s}^{\frac{1}{s}} (1 + |\log \eta|)^{1/\beta}, \\ \|\sigma(\mathbf{w}_\eta)\|_{\mathcal{A}^s} &\lesssim \|\sigma(\mathbf{v})\|_{\mathcal{A}^s} (1 + |\log \eta|)^{s/\beta}, \end{aligned}$$

as well as

$$\|\pi^{(y)}(\mathbf{w}_\eta)\|_{\mathcal{A}^s} \lesssim \|\pi^{(y)}(\mathbf{v})\|_{\mathcal{A}^s} (1 + |\log \eta|)^{s/\beta}$$

and analogous bounds for $\|\pi^{(x)}(\mathbf{w}_\eta)\|_{\mathcal{A}^s}$. These estimates show a logarithmic degradation in the target tolerances. Nevertheless, they allow one to establish the estimate

$$(5.27) \quad \text{ops}(\mathbf{u}_\varepsilon) \lesssim (1 + |\log \varepsilon|)^\zeta \left(\varepsilon^{-\frac{1}{s}} \|\sigma(\mathbf{u})\|_{\mathcal{A}^s}^{\frac{1}{s}} \right)^2 \sum_{i \in \{x, y\}} \varepsilon^{-\frac{1}{s}} \|\pi^{(i)}(\mathbf{u})\|_{\mathcal{A}^s}^{-\frac{1}{s}},$$

where $\zeta > 0$ depends on s , on $\text{cond}(\mathbf{A})$, and on the parameters in Algorithm 5.1. Further details can be found in [6]. Thus, in this case with algebraic decay of singular values, although the arising ranks are quasi-optimal, the quadratic dependence of the cost of the SVD on these ranks leads to substantially less favorable costs of low-rank approximations. This effect is essentially unavoidable by any low-rank scheme using a separation of spatial and parametric variables.

REMARK 5.4. As outlined in Section 3.3, an alternative to the perturbed Richardson iteration in Algorithm 5.1 consists in solving a sequence of Galerkin discretizations that are successively refined using residual approximations as considered in Section 5.1. As shown in [2], such an approach leads to analogous asymptotic complexity bounds, with potential for quantitative improvements in the numerical costs.

5.2.3. *Soft thresholding.* A second basic approach to control the complexity of iterates in a scheme of the form (3.10) is to choose a reduction operation \mathcal{R}_n which is *non-expansive*, that is, $\|\mathcal{R}_n(\mathbf{u}) - \mathcal{R}_n(\mathbf{v})\| \leq \|\mathbf{u} - \mathbf{v}\|$. As a consequence, if \mathcal{F}_n is a contraction, then $\mathcal{R}_n \circ \mathcal{F}_n$ is still a contraction for each n . This applies, in particular, to the soft thresholding operation, which can be applied entry-wise to sparse expansions, or to the hierarchical singular values of tensors as discussed in Section 2.3.

Compared to the truncation of the *HSVD* by hard thresholding, one has the interesting feature that convergence of the iteration is preserved regardless of the thresholding value. Rather than thresholding with a sufficiently large tolerance whenever a sufficient error reduction has been achieved, one can therefore simply threshold in every iteration step, with parameters that decrease sufficiently slowly.

Let us take a closer look at this in the case $\mathcal{F}_n = \mathcal{F}$, that is, the underlying iteration is stationary. One key observation for establishing quasi-optimality of ranks is that the limit of the iteration with the thresholded mapping $\mathfrak{S}_\alpha \circ \mathcal{F}$ can be related to the thresholded exact solution.

LEMMA 5.5. *Let \mathcal{F} be a contraction with Lipschitz constant $\rho < 1$ and unique fixed point u , let $\alpha > 0$ and let u^α be the unique fixed point of $\mathfrak{S}_\alpha \circ \mathcal{F}$. Then*

$$(1 + \rho)^{-1} \|\mathfrak{S}_\alpha(u) - u\| \leq \|u^\alpha - u\| \leq (1 - \rho)^{-1} \|\mathfrak{S}_\alpha(u) - u\|.$$

Up to mildly dimension-dependent constants (recall that $E = 2d - 3$), the rank reduction by \mathfrak{S}_α produces quasi-optimal ranks:

- (i) In the case of algebraic decay of singular values, that is, $\sigma^{(i)}(u) \in \mathcal{A}^s$ with $s > 0$ for $i = 1, \dots, E$, with $C_{u,s} := \max_i \|\sigma^{(i)}(u)\|_{\mathcal{A}^s}$, we have $\|\mathfrak{S}_\alpha(u) - u\| \leq EC_{u,s}^{1/(2s+1)} \alpha^{2s/(2s+1)}$ with ranks

$$\max \mathbf{r}_{\mathbb{E}}(\mathfrak{S}_\alpha(u)) \leq (C_{u,s} \alpha^{-1})^{2/(2s+1)}.$$

This means that $\|\mathfrak{S}_\alpha(u) - u\| \leq \varepsilon$ with

$$(5.28) \quad \max \mathbf{r}_{\mathbb{E}}(\mathfrak{S}_\alpha(u)) \leq (EC_{u,s})^{1/s} \varepsilon^{-1/s}.$$

- (ii) If $\sigma_k^{(i)}(u) \leq Ce^{-ck^\beta}$ for $k \in \mathbb{N}$ with $C, c, \beta > 0$, then $\|\mathfrak{S}_\alpha(u) - u\| \leq C_1 E(1 + |\log \alpha|)^{1/(2\beta)} \alpha$ with ranks

$$\max r_{\mathbb{E}}(\mathfrak{S}_\alpha(u)) \leq (c^{-1} |\log C \alpha^{-1}|)^{1/\beta} \leq C_2 (1 + |\log \alpha|)^{1/\beta}.$$

This gives $\|\mathfrak{S}_\alpha(u) - u\| \leq \varepsilon$, for $\varepsilon < 1$, with

$$(5.29) \quad \max r_{\mathbb{E}}(\mathfrak{S}_\alpha(u)) \leq C_3 (1 + |\log \varepsilon|)^{1/\beta},$$

where $C_1, C_2, C_3 > 0$ depend on C, c, β .

The link from thresholded approximations to the approximability of exact solutions is provided by the following lemma.

LEMMA 5.6. *Let $u, v \in H$, $\alpha > 0$, and $\varepsilon > 0$ such that $\|u - v\| \leq \varepsilon$.*

- (i) *If $\sigma^{(i)}(u) \in \mathcal{A}^s$ with $s > 0$ for all $i = 1, \dots, E$, then*

$$r_i(\mathfrak{S}_\alpha(v)) \leq \frac{4\varepsilon^2}{\alpha^2} + C_s (\|\sigma^{(i)}(u)\|_{\mathcal{A}^s} \alpha^{-1})^{2/(2s+1)}.$$

- (ii) *If $\sigma_k^{(i)}(u) \leq Ce^{-ck^\beta}$ for $k \in \mathbb{N}$ with $C, c, \beta > 0$, then*

$$r_i(\mathfrak{S}_\alpha(v)) \leq \frac{4\varepsilon^2}{\alpha^2} + (c^{-1} \ln(2C\alpha^{-1}))^{1/\beta}.$$

To obtain an iterative method with quasi-optimal ranks for all iterates, it thus suffices to ensure that the first terms of order ε^2/α^2 in the above estimates remain comparable to the respective second terms by decreasing α sufficiently slowly.

For elliptic problems in well-conditioned representations $\mathbf{A}\mathbf{u} = \mathbf{f}$, this can be realized using a Richardson iteration (3.8). Let $\omega > 0$ be such that $\xi := \|I - \omega\mathbf{A}\| < 1$. The basic iterative method applied to the present problem has the form

$$(5.30) \quad \mathbf{u}^{n+1} = \mathfrak{S}_{\alpha_n}(\mathbf{u}^n - \omega(\mathbf{A}\mathbf{u}^n - \mathbf{f})), \quad n \geq 0,$$

with $\mathbf{u}^0 = 0$ and $\alpha_n \rightarrow 0$ determined (according to [12, Alg. 2]) as follows: set $\alpha_0 = \omega\|\mathbf{f}\|_2/(d-1)$, and for a fixed $\bar{B} > \|\mathbf{A}\|$, take

$$(5.31) \quad \alpha_{n+1} = \begin{cases} \frac{1}{2}\alpha_n, & \text{if } \|\mathbf{u}^{n+1} - \mathbf{u}^n\|_2 \leq \frac{1-\xi}{\xi\bar{B}} \|\mathbf{A}\mathbf{u}^{n+1} - \mathbf{f}\|, \\ \alpha_n, & \text{else.} \end{cases}$$

As shown in [12], the scheme given by (5.30), (5.31), converges linearly and each iterate \mathbf{u}^n satisfies quasi-optimal rank estimates of the form (5.28) or (5.29) provided that the exact solution \mathbf{u} has the corresponding approximability.

This is potentially stronger than the result for hard thresholding, where one cannot rule out that ranks in the inner loop of Algorithm 5.1 between steps 3 and 11 cumulatively increase due to repeated (approximate) application of \mathbf{A} before being reduced to near-optimality in step 11. In contrast, in (5.30), (5.31), the iterates are returned to quasi-optimality after *every single application* of \mathbf{A} . However, in the form given above and in [12], the soft thresholding method still assumes a *fixed discretization* (or an idealized iteration on the full sequence space) and does not yet incorporate adaptive discretizations.

6. Conclusions and outlook

The methods discussed in this article rest on two conceptual pillars. First, they invoke strategies for estimating errors with respect to the underlying continuous problem in highly nonlinear approximations of solutions of high-dimensional PDEs, combining adaptive discretizations with low-rank expansions. Exploiting the mapping properties of the underlying continuous operator is essential. These strategies are based on the approximate *evaluation of residuals* in function spaces and are guaranteed to remain computationally feasible even in very high-dimensional settings as they progress from coarse to fine with certified accuracy at each stage.

Second, they employ basic *complexity reduction mechanisms* for ensuring quasi-optimality of computed approximations. That is, they ensure that their representation complexities remain comparable to those of corresponding best approximations of the same accuracy. This can be achieved, in particular, by the truncation of the hierarchical singular value decomposition up to a judiciously chosen tolerance. An alternate strategy is based on soft thresholding, which has the advantage of preserving the convergence of iterative methods for any thresholding parameter.

The combination of the approximate residual evaluations with such recompression strategies enables the construction of iterative methods that converge to the exact solutions of the continuous problem with near-optimal computational costs. As we have noted, soft thresholding can, in principle, give slightly stronger bounds on the total computational complexity, since the ranks of all iterates are under control, and the bounds are therefore less dependent on the ranks of operators. This approach, however, has not yet been combined with adaptive discretizations.

We have confined the discussion to highlighting the essential conceptual mechanisms. Corresponding findings are illustrated by first numerical experiments in [6, 10, 11]; in particular, comparisons with other methods are given in [10]. While these experiments confirm the near-optimal asymptotic complexity bounds for the resulting methods, much room is certainly left for optimizing corresponding concrete implementations with regard to quantitative practical performance.

Concerning the basic construction of solvers, there is a variety of methods that follow a quite different philosophy in using minimization principles for optimizing tensor decompositions for *fixed discretizations*, for instance ALS [40, 50, 55], DMRG [53, 54], AMEn [29], or Riemannian optimization methods [44]. Comparably little is known, however, on their global convergence properties, let alone the total computational complexity of such methods in relation to the output accuracy. As a possible further direction, they could, however, serve as additional inner iterations for accelerating the convergence of error-controlled methods as considered here.

Here we have concentrated on the application to linear operator equations. For other problem classes, such as eigenvalue problems or time-dependent problems, only rather preliminary results on error-controlled low-rank methods are available, see, e.g., [3, 4]. Many additional challenges in the application of the basic principles discussed here to such problems remain open.

References

- [1] M. Ali and A. Nouy, *Singular value decomposition in Sobolev spaces*, arXiv:1809.11001, 2018.
- [2] M. Ali and K. Urban, *HT-AWGM: A hierarchical Tucker-adaptive wavelet galerkin method for high dimensional elliptic problems*, arXiv preprint arXiv:1805.12016, 2018.

- [3] R. Andreev and C. Tobler, *Multilevel preconditioning and low-rank tensor iteration for space-time simultaneous discretizations of parabolic PDEs*, Numer. Linear Algebra Appl. **22** (2015), no. 2, 317–337, DOI 10.1002/nla.1951. MR3313261
- [4] M. Bachmayr, *Adaptive low-rank wavelet methods and applications to two-electron Schrödinger equations*, Ph.D. thesis, RWTH Aachen, 2012.
- [5] M. Bachmayr and A. Cohen, *Kolmogorov widths and low-rank approximations of parametric elliptic PDEs*, Math. Comp. **86** (2017), no. 304, 701–724, DOI 10.1090/mcom/3132. MR3584545
- [6] M. Bachmayr, A. Cohen, and W. Dahmen, *Parametric PDEs: sparse or low-rank approximations?*, IMA J. Numer. Anal. **38** (2018), no. 4, 1661–1708, DOI 10.1093/imanum/drx052. MR3867379
- [7] M. Bachmayr, A. Cohen, and G. Migliorati, *Sparse polynomial approximation of parametric elliptic PDEs. Part I: Affine coefficients*, ESAIM Math. Model. Numer. Anal. **51** (2017), no. 1, 321–339, DOI 10.1051/m2an/2016045. MR3601010
- [8] M. Bachmayr, A. Cohen, and G. Migliorati, *Representations of Gaussian random fields and approximation of elliptic PDEs with lognormal coefficients*, J. Fourier Anal. Appl. **24** (2018), no. 3, 621–649, DOI 10.1007/s00041-017-9539-5. MR3802287
- [9] M. Bachmayr and W. Dahmen, *Adaptive near-optimal rank tensor approximation for high-dimensional operator equations*, Found. Comput. Math. **15** (2015), no. 4, 839–898, DOI 10.1007/s10208-013-9187-3. MR3371372
- [10] M. Bachmayr and W. Dahmen, *Adaptive low-rank methods for problems on Sobolev spaces with error control in L_2* , ESAIM Math. Model. Numer. Anal. **50** (2016), no. 4, 1107–1136, DOI 10.1051/m2an/2015071. MR3521714
- [11] M. Bachmayr and W. Dahmen, *Adaptive low-rank methods: problems on Sobolev spaces*, SIAM J. Numer. Anal. **54** (2016), no. 2, 744–796, DOI 10.1137/140978223. MR3474850
- [12] M. Bachmayr and R. Schneider, *Iterative methods based on soft thresholding of hierarchical tensors*, Found. Comput. Math. **17** (2017), no. 4, 1037–1083, DOI 10.1007/s10208-016-9314-z. MR3682220
- [13] M. Bachmayr, R. Schneider, and A. Uschmajew, *Tensor networks and hierarchical tensors for the solution of high-dimensional partial differential equations*, Found. Comput. Math. **16** (2016), no. 6, 1423–1472, DOI 10.1007/s10208-016-9317-9. MR3579714
- [14] J. Ballani and L. Grasedyck, *A projection method to solve linear systems in tensor format*, Numer. Linear Algebra Appl. **20** (2013), no. 1, 27–43, DOI 10.1002/nla.1818. MR3007237
- [15] G. Beylkin and M. J. Mohlenkamp, *Numerical operator calculus in higher dimensions*, Proc. Natl. Acad. Sci. USA **99** (2002), no. 16, 10246–10251, DOI 10.1073/pnas.112329799. MR1918798
- [16] M. Billaud-Friess, A. Nouy, and O. Zahm, *A tensor approximation method based on ideal minimal residual formulations for the solution of high-dimensional problems*, ESAIM Math. Model. Numer. Anal. **48** (2014), no. 6, 1777–1806, DOI 10.1051/m2an/2014019. MR3342163
- [17] D. Braess and W. Hackbusch, *Approximation of $1/x$ by exponential sums in $[1, \infty)$* , IMA J. Numer. Anal. **25** (2005), no. 4, 685–697, DOI 10.1093/imanum/dri015. MR2170519
- [18] Dietrich Braess and Wolfgang Hackbusch, *On the efficient computation of high-dimensional integrals and the approximation by exponential sums*, Multiscale, Nonlinear and Adaptive Approximation, R. DeVore and A. Kunoth, Eds., Springer Berlin Heidelberg, 2009.
- [19] A. Cohen, W. Dahmen, and R. DeVore, *Adaptive wavelet methods for elliptic operator equations: convergence rates*, Math. Comp. **70** (2001), no. 233, 27–75, DOI 10.1090/S0025-5718-00-01252-7. MR1803124
- [20] A. Cohen, W. Dahmen, and R. DeVore, *Adaptive wavelet methods. II. Beyond the elliptic case*, Found. Comput. Math. **2** (2002), no. 3, 203–245, DOI 10.1007/s102080010027. MR1907380
- [21] A. Cohen and R. DeVore, *Approximation of high-dimensional parametric PDEs*, Acta Numer. **24** (2015), 1–159, DOI 10.1017/S0962492915000033. MR3349307
- [22] 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), no. 6, 615–646, DOI 10.1007/s10208-010-9072-2. MR2728424
- [23] A. Cohen, R. DeVore, and C. Schwab, *Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE’s*, Anal. Appl. (Singap.) **9** (2011), no. 1, 11–47, DOI 10.1142/S0219530511001728. MR2763359

- [24] W. Dahmen, *Wavelet and multiscale methods for operator equations*, Acta numerica, 1997, Acta Numer., vol. 6, Cambridge Univ. Press, Cambridge, 1997, pp. 55–228, DOI 10.1017/S0962492900002713. MR1489256
- [25] W. Dahmen, R. DeVore, L. Grasedyck, and E. Süli, *Tensor-sparsity of solutions to high-dimensional elliptic partial differential equations*, Found. Comput. Math. **16** (2016), no. 4, 813–874, DOI 10.1007/s10208-015-9265-9. MR3529127
- [26] V. de Silva and L.-H. Lim, *Tensor rank and the ill-posedness of the best low-rank approximation problem*, SIAM J. Matrix Anal. Appl. **30** (2008), no. 3, 1084–1127, DOI 10.1137/06066518X. MR2447444
- [27] R. A. DeVore, *Nonlinear approximation*, Acta numerica, 1998, Acta Numer., vol. 7, Cambridge Univ. Press, Cambridge, 1998, pp. 51–150, DOI 10.1017/S0962492900002816. MR1689432
- [28] S. Dolgov and T. Vejchodský, *Guaranteed a posteriori error bounds for low rank tensor approximate solutions*, (2019).
- [29] S. V. Dolgov and D. V. Savostyanov, *Alternating minimal energy methods for linear systems in higher dimensions*, SIAM J. Sci. Comput. **36** (2014), no. 5, A2248–A2271, DOI 10.1137/140953289. MR3262607
- [30] M. Eigel, M. Pfeffer, and R. Schneider, *Adaptive stochastic Galerkin FEM with hierarchical tensor representations*, Numer. Math. **136** (2017), no. 3, 765–803, DOI 10.1007/s00211-016-0850-x. MR3660302
- [31] A. Falcó and W. Hackbusch, *On minimal subspaces in tensor representations*, Found. Comput. Math. **12** (2012), no. 6, 765–803, DOI 10.1007/s10208-012-9136-6. MR2989473
- [32] I. P. Gavriluk, W. Hackbusch, and B. N. Khoromskij, *\mathcal{H} -matrix approximation for the operator exponential with applications*, Numer. Math. **92** (2002), no. 1, 83–111, DOI 10.1007/s002110100360. MR1917366
- [33] C. J. Gittelsohn, *Adaptive wavelet methods for elliptic partial differential equations with random operators*, Numer. Math. **126** (2014), no. 3, 471–513, DOI 10.1007/s00211-013-0572-2. MR3164144
- [34] C. J. Gittelsohn, *An adaptive stochastic Galerkin method for random elliptic operators*, Math. Comp. **82** (2013), no. 283, 1515–1541, DOI 10.1090/S0025-5718-2013-02654-3. MR3042573
- [35] L. Grasedyck, *Hierarchical singular value decomposition of tensors*, SIAM J. Matrix Anal. Appl. **31** (2009/10), no. 4, 2029–2054, DOI 10.1137/090764189. MR2678955
- [36] L. Grasedyck, D. Kressner, and C. Tobler, *A literature survey of low-rank tensor approximation techniques*, GAMM-Mitt. **36** (2013), no. 1, 53–78, DOI 10.1002/gamm.201310004. MR3095914
- [37] W. Hackbusch, *Tensor spaces and numerical tensor calculus*, Springer Series in Computational Mathematics, vol. 42, Springer, Heidelberg, 2012. MR3236394
- [38] W. Hackbusch, B. N. Khoromskij, and E. E. Tyrtshnikov, *Approximate iterations for structured matrices*, Numer. Math. **109** (2008), no. 3, 365–383, DOI 10.1007/s00211-008-0143-0. MR2399149
- [39] W. Hackbusch and S. Kühn, *A new scheme for the tensor representation*, J. Fourier Anal. Appl. **15** (2009), no. 5, 706–722, DOI 10.1007/s00041-009-9094-9. MR2563780
- [40] S. Holtz, T. Rohwedder, and R. Schneider, *The alternating linear scheme for tensor optimization in the tensor train format*, SIAM J. Sci. Comput. **34** (2012), no. 2, A683–A713, DOI 10.1137/100818893. MR2914300
- [41] B. N. Khoromskij and I. Oseledets, *Quantics-TT collocation approximation of parameter-dependent and stochastic elliptic PDEs*, Comput. Methods Appl. Math. **10** (2010), no. 4, 376–394, DOI 10.2478/cmam-2010-0023. MR2770302
- [42] B. N. Khoromskij and C. Schwab, *Tensor-structured Galerkin approximation of parametric and stochastic elliptic PDEs*, SIAM J. Sci. Comput. **33** (2011), no. 1, 364–385, DOI 10.1137/100785715. MR2783199
- [43] T. G. Kolda and B. W. Bader, *Tensor decompositions and applications*, SIAM Rev. **51** (2009), no. 3, 455–500, DOI 10.1137/07070111X. MR2535056
- [44] D. Kressner, M. Steinlechner, and B. Vandereycken, *Preconditioned low-rank Riemannian optimization for linear systems with tensor product structure*, SIAM J. Sci. Comput. **38** (2016), no. 4, A2018–A2044, DOI 10.1137/15M1032909. MR3519141

- [45] D. Kressner and C. Tobler, *Low-rank tensor Krylov subspace methods for parametrized linear systems*, SIAM J. Matrix Anal. Appl. **32** (2011), no. 4, 1288–1316, DOI 10.1137/100799010. MR2854614
- [46] L. De Lathauwer, B. De Moor, and J. Vandewalle, *A multilinear singular value decomposition*, SIAM J. Matrix Anal. Appl. **21** (2000), no. 4, 1253–1278, DOI 10.1137/S0895479896305696. MR1780272
- [47] I. V. Oseledets, *Tensor-train decomposition*, SIAM J. Sci. Comput. **33** (2011), no. 5, 2295–2317, DOI 10.1137/090752286. MR2837533
- [48] I. V. Oseledets and E. E. Tyrtyshnikov, *Breaking the curse of dimensionality, or how to use SVD in many dimensions*, SIAM J. Sci. Comput. **31** (2009), no. 5, 3744–3759, DOI 10.1137/090748330. MR2556560
- [49] T. Poggio, H. Mhaskar, L. Rosasco, B. Miranda, and Liao Q., *Why and when can deep-but not shallow-networks avoid the curse of dimensionality: A review*, Int. J. Autom. Comput. **14** (2017), no. 5, 503–519.
- [50] T. Rohwedder and A. Uschmajew, *On local convergence of alternating schemes for optimization of convex problems in the tensor train format*, SIAM J. Numer. Anal. **51** (2013), no. 2, 1134–1162, DOI 10.1137/110857520. MR3038114
- [51] C. Schwab and R. Stevenson, *Adaptive wavelet algorithms for elliptic PDE’s on product domains*, Math. Comp. **77** (2008), no. 261, 71–92, DOI 10.1090/S0025-5718-07-02019-4. MR2353944
- [52] L. R. Tucker, *Contributions to mathematical psychology*, ch. The extension of factor analysis to three-dimensional matrices, pp. 109–127, Holt, Rinehart & Winston, New York, 1964.
- [53] G. Vidal, *Efficient classical simulation of slightly entangled quantum computations*, Phys. Rev. Lett. **91** (2003), no. 14, 147902.
- [54] S. R. White, *Density matrix formulation for quantum renormalization groups*, Phys. Rev. Lett. **69** (1992), 2863–2866.
- [55] S. R. White, *Density matrix renormalization group algorithms with a single center site*, Phys. Rev. B **72** (2005), no. 18, 180403.
- [56] D. Yarotsky, *Error bounds for approximations with deep ReLU networks*, Neural Netw. **94** (2017), 103–114.

JOHANNES GUTENBERG-UNIVERSITÄT MAINZ, INSTITUT FÜR MATHEMATIK, STAUDINGERWEG
9 55128 MAINZ, GERMANY

Email address: bachmayr@uni-mainz.de

UNIVERSITY OF SOUTH CAROLINA, MATHEMATICS DEPARTMENT, LECONTE COLLEGE, 1523
GREENE STREET, COLUMBIA, SOUTH CAROLINA 29208

Email address: wolfgang.anton.dahmen@googlemail.com