# Low-rank tensor methods for partial differential equations 

Markus Bachmayr<br>Institut für Geometrie und Praktische Mathematik, RWTH Aachen University, Templergraben 55, 52056 Aachen, Germany<br>E-mail: bachmayr@igpm.rwth-aachen.de

Low-rank tensor representations can provide highly compressed approximations of functions. These concepts, which essentially amount to generalizations of classical techniques of separation of variables, have proved to be particularly fruitful for functions of many variables. We focus here on problems where the target function is given only implicitly as the solution of a partial differential equation. A first natural question is under which conditions we should expect such solutions to be efficiently approximated in low-rank form. Due to the highly nonlinear nature of the resulting low-rank approximations, a crucial second question is at what expense such approximations can be computed in practice. This article surveys basic construction principles of numerical methods based on low-rank representations as well as the analysis of their convergence and computational complexity.

2020 Mathematics Subject Classification:
Primary 41A46, 41A63, 65D40
Secondary 65F55, 65J10, 65M12, 65N12, 65N25, 65Y20

## CONTENTS

1 Introduction ..... 2
2 Subspace tensor formats ..... 14
3 Approximations of functions and operators in low-rank tensor formats ..... 39
4 Approximate inverses and preconditioning of elliptic operators ..... 55
5 Solving operator equations in low-rank format ..... 65
6 Eigenvalue problems ..... 98
7 Time-dependent problems ..... 101
References ..... 107

## 1. Introduction

A matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$ of rank $r$ can be represented in low-rank format $\mathbf{M}=\mathbf{A B}{ }^{\top}$ with matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times r}$, which when $r \ll n$ substantially reduces the required storage from $n^{2}$ to $2 n r$. Moreover, the low-rank format also leads to a reduction of the computational costs of many linear algebra operations, such as multiplication by matrices from the left or right. This basic idea can be generalized to tensors of order $d>2$, represented - in the real-valued case that we focus on here - by elements of $\mathbb{R}^{n_{1} \times \ldots \times n_{d}}$ with $n_{1}, \ldots, n_{d} \in \mathbb{N}$. This leads to low-rank formats for higher-order tensors, which can be used to significantly reduce the dependence on $d$ of storage requirements and computational complexity of numerical methods. One can thus hope to avoid the curse of dimensionality, that is, the exponential scaling with respect to $d$ obtained with standard methods designed for lower-dimensional problems.

A main motivation for the development of numerical methods based on low-rank tensor representations is the application to partial differential equations posed on high-dimensional domains, which frequently arise in the context of quantum physics or probabilistic models. Closely related are methods for model order reduction of multi-parametric problems. In these settings, coefficient arrays of basis expansions of solutions can be regarded as high-order tensors. Low-rank approximations of these high-dimensional coefficient arrays lead to nonlinear approximation methods, that is, approximations are not chosen from linear spaces and are parametrized in a nonlinear manner. As above in the parametrization of $\mathbf{M}$ by the factors $\mathbf{A}$ and B, this nonlinearity is of multiplicative type.

Certain types of low-rank tensor representations have proved to be particularly suited to the construction of numerical methods for problems in very high dimensions, such as hierarchical tensors, which are also known as tree tensor networks in physics. Since we consider approximations in such formats for functions that are given only implicitly as solutions of partial differential equations, not only the solution itself but also the problem data (such as domains, right-hand sides, boundary or initial conditions) need to have efficient low-rank approximations. A central issue in the construction of solvers is the adaptation of rank parameters in tensor representations: these ranks needs to be chosen sufficiently large to ensure convergence, but should also be kept as small as possible due to their strong influence on the computational costs of operating on low-rank representations.

The focus of this article is in particular on how low-rank tensor formats interact with discretizations of partial differential equations in the construction of numerical solvers. The remainder of this section gives a conceptual overview; further details and references are given in the respective later sections.

### 1.1. Notation

Vectors, matrices and tensors are denoted by bold-face letters. On the one hand, we consider tensors on finite index sets, that is, elements of $\mathbb{R}^{n_{1} \times \cdots \times n_{d}}$. These
are tensors of order $d$, where each dimension in the index set is called a mode of the tensor, with mode $i$ having mode size $n_{i}$ for $i=1, \ldots, d$. On the other hand, an important role is played by tensors with square-summable entries on countable $d$-dimensional index sets, which are sequences in $\ell_{2}\left(\mathbb{N}^{d}\right)$. In both cases, in order to avoid too many subscripts, for the entry of a tensor $\mathbf{t}$ of order $d$ at index $\left(v_{1}, \ldots, v_{d}\right) \in \mathbb{N}^{d}$, we write $\mathbf{t}\left[v_{1}, \ldots, v_{d}\right]$.

For the unit cube in $d$ dimensions, which frequently appears in high-dimensional model problems, we use the abbreviation $\square_{d}=(0,1)^{d}$. We write supp $f$ for the support of the function $f$ and use the same notation for the supports of sequences: for $\mathbf{v} \in \mathbb{R}^{I}$ with $I$ a countable index set,

$$
\operatorname{supp} \mathbf{v}=\{v \in I: \mathbf{v}[v] \neq 0\}
$$

We write cond(A) for the spectral norm condition number of the mapping defined by $\mathbf{A}$.

For scalar expressions $A$ and $B$, we write $A \lesssim B$ to denote $A \leq C B$ with a constant $C>0$ that does not depend on quantities that explicitly appear in $A$ or $B$; $A \gtrsim B$ for $B \lesssim A$; and $A \approx B$ for $A \lesssim B$ and $A \gtrsim B$. Inequalities between tuples (such as tuples of rank parameters) are always to be understood componentwise.

### 1.2. Classes of model problems

Let us first consider several typical model cases of high-dimensional partial differential equations (PDEs). A first problem of frequent interest is the high-dimensional diffusion problem of the basic form

$$
\begin{gather*}
\partial_{t} u-\operatorname{div}_{x}\left(A \operatorname{grad}_{x} u+b u\right)=f \text { in }(0, T] \times \Omega, \\
u=0 \text { on }(0, T] \times \partial \Omega, \quad u(0, x)=u_{0}(x) \text { for } x \in \Omega . \tag{1.1}
\end{gather*}
$$

Here we consider a bounded domain $\Omega$ with homogeneous Dirichlet boundary conditions for simplicity, but problems of this type can also be considered on unbounded domains such as $\Omega=\mathbb{R}^{d}$. The sought solution $u$ as well as the source term $f$ and initial condition $u_{0}$ are scalar functions, whereas $A$ and $b$ take values in $\mathbb{R}^{d \times d}$ and $\mathbb{R}^{d}$, respectively. Equations of this type arise, for instance, as Kolmogorov equations describing the time evolutions of probability densities of stochastic processes.

A second source of high-dimensional problems is given by quantum-physical models, where the dimensionality of the problem scales with the number of particles. In classical quantum mechanics, the time evolution of quantum systems is modelled by the time-dependent Schrödinger equation for the complex-valued wavefunction $u$. For a given potential $V$ and initial state $u_{0}$, it takes the form

$$
\begin{equation*}
\mathrm{i} \partial_{t} u=-\Delta u+V u \text { on } \mathbb{R}^{d},\left.\quad u\right|_{t=0}=u_{0} . \tag{1.2}
\end{equation*}
$$

The eigenpairs of the Hamiltonian operator $-\Delta+V$ on the right-hand side correspond to the quantum states for the total energy of the system; the associated Schrödinger
eigenvalue problem reads

$$
\begin{equation*}
-\Delta u+V u=\lambda u \text { on } \mathbb{R}^{d} . \tag{1.3}
\end{equation*}
$$

One example that is of central importance in quantum chemistry is the electronic Schrödinger eigenvalue problem, which provides the (non-relativistic) standard model of electronic structure of molecules for given positions of atomic nuclei. In this case we have $d=3 N$ for a system of $N$ electrons, and $V$ is a sum of Coulomb potentials.

We are also frequently interested in stationary states of (1.1), which are described by second-order elliptic equations. A frequently considered model problem, with the simplification $b=0$, reads

$$
\begin{equation*}
-\operatorname{div}(A \operatorname{grad} u)=f \text { in } \Omega, \quad u=0 \text { on } \partial \Omega \tag{1.4}
\end{equation*}
$$

Such problems of elliptic type also arise as parts of algorithms for the parabolic problem (1.1) and the Schrödinger problems (1.2), (1.3) (with $A=\mathrm{I}$ and $\Omega=\mathbb{R}^{d}$ ) and are thus a natural starting point for the development of numerical methods.

The above model problems have in common that they involve differential operators on high-dimensional domains. A different type of high-dimensionality originates from problem data that depend on parameters, for instance some parameter $y$ from a set $Y$. In the example of (1.4), with parameter-dependent diffusion tensor $A_{y}$ and associated solutions $u_{y}$, this leads to the mapping $y \mapsto u_{y}$ defined for all $y \in Y$ by

$$
\begin{equation*}
-\operatorname{div}\left(A_{y} \operatorname{grad} u_{y}\right)=f \text { in } \Omega, \quad u=0 \text { on } \partial \Omega . \tag{1.5}
\end{equation*}
$$

Instead of evaluating $u_{y}$ independently for each required $y$, in the context of model order reduction or uncertainty quantification, we are frequently interested in economical approximations of the higher-dimensional function $u(x, y)=u_{y}(x)$ for $(x, y) \in \Omega \times Y$.

After treating the basics of low-rank tensor representations in Section 2, in the following sections we first focus on elliptic problems. Their solution is of interest in its own right, but also provides important building blocks for more involved problems. In particular, the norm in which solutions are to be approximated is important: in (1.4) and (1.5), we typically require approximations in the norm of the Sobolev space $H^{1}(\Omega)$, which is equivalent to the energy norm associated to the problem. How this needs to be taken into account in the construction of low-rank solvers is treated in Section 4. After treating low-rank solvers for elliptic problems in Section 5, we return to eigenvalue problems and time-dependent problems in Sections 6 and 7.

Concerning the application of low-rank approximations to solutions of PDE problems as above, there are several fundamental questions: Can we have algorithms with guaranteed convergence to the exact solution, that is, with the ability to produce approximations of the solution of the original PDE problem with any desired accuracy with respect to the appropriate norm? Can these algorithms be
ensured to be numerically stable? Can the costs of numerically computing low-rank approximations be quantified, and can we ensure that these costs are optimal in an appropriate sense? For problems in high dimensions $d$, we aim to avoid the curse of dimensionality, that is, an exponential dependence of the costs on $d$. This imposes requirements on both the approximability of solutions and the costs of numerical algorithms.

### 1.3. Separation of variables

A basic instance of low-rank methods is provided by the classical techniques of separation of variables in partial differential equations, where higher-dimensional problems are decomposed into simpler ones. As an example, consider the eigenvalue problem for the negative Laplacian

$$
-\Delta=-\sum_{i=1}^{d} \partial_{x_{i}}^{2}
$$

on $\square_{d}$ with homogeneous Dirichlet boundary values on $\partial \square_{d}$. From the solutions of the one-dimensional eigenvalue problem for the negative second derivative on $(0,1)$ with homogeneous Dirichlet boundary conditions,

$$
-\partial_{x}^{2} \sin (\ell \pi x)=(\ell \pi)^{2} \sin (\ell \pi x), \quad \ell \in \mathbb{N}
$$

it is easy to see that for each $k=\left(k_{1}, \ldots, k_{d}\right) \in \mathbb{N}^{d}$, the functions

$$
\Phi_{k}\left(x_{1}, \ldots, x_{d}\right)=\prod_{i=1}^{d} \sin \left(k_{i} \pi x_{i}\right)
$$

satisfy $\left.\Phi_{k}\right|_{\partial \square_{d}}=0$ and

$$
-\Delta \Phi_{k}=\pi^{2}\left(k_{1}^{2}+\cdots+k_{d}^{2}\right) \Phi_{k}
$$

Thus $\Phi_{k}$ is an eigenfunction of $-\Delta$ on $\square_{d}$, and since $\operatorname{span}\left\{\Phi_{k}\right\}_{k \in \mathbb{N}^{d}}$ is dense in $L_{2}\left(\square_{d}\right)$, all eigenfunctions in $L_{2}\left(\square_{d}\right)$ are of this form.

By a similar use of the structure of the Laplacian, certain multidimensional initial value problems can also be reduced to spatially one-dimensional equations. Let

$$
g\left(x_{1}, \ldots, x_{d}\right)=\prod_{i=1}^{d} g_{i}\left(x_{i}\right)
$$

with

$$
g_{i}(x)=\sum_{k=1}^{\infty} c_{k}^{(i)} \sin (k \pi x)
$$

Then the solution of the heat equation $\partial_{t} u-\Delta u=0$ on $\mathbb{R}^{+} \times \square_{d}$ with homogeneous Dirichlet boundary conditions and initial condition $u(0, \cdot)=g$ on $\square_{d}$ has the product
form

$$
u\left(t, x_{1}, \ldots, x_{d}\right)=\prod_{i=1}^{d} u_{i}\left(t, x_{i}\right), \quad u_{i}(t, x)=\sum_{k=1}^{\infty} c_{k}^{(i)} e^{-(k \pi)^{2} t} \sin (k \pi x)
$$

where $u_{i}$ solve the spatially univariate problems $\partial_{t} u_{i}-\partial_{x}^{2} u_{i}=0, u_{i}(0, \cdot)=g_{i}$.
This is an instance of the following more general structure: let $H_{1}, \ldots, H_{d}$ be separable Hilbert spaces and consider the tensor product Hilbert space $H=$ $\otimes_{i=1}^{d} H_{i}$ and an operator $A: H \rightarrow H$ of the Laplacian-type form

$$
\begin{equation*}
A=A_{1} \otimes \mathrm{I} \otimes \cdots \otimes \mathrm{I}+\mathrm{I} \otimes A_{2} \otimes \mathrm{I} \otimes \cdots \otimes \mathrm{I}+\cdots+\mathrm{I} \otimes \cdots \otimes \mathrm{I} \otimes A_{d} \tag{1.6}
\end{equation*}
$$

where for $i=1, \ldots, d, A_{i}: H_{i} \rightarrow H_{i}$ are self-adjoint and positive definite, and where I denotes the identity mapping. Assume that the normalized eigenfunctions $\left\{\phi_{k}^{(i)}\right\}_{k \in \mathbb{N}}$ of $A_{i}$ are an orthonormal basis of $H_{i}$ for $i=1, \ldots, d$. Then, for each $k=\left(k_{1}, \ldots, k_{d}\right)$,

$$
\Phi_{k}=\bigotimes_{i=1}^{d} \phi_{k_{i}}^{(i)}
$$

is an eigenfunction of $A$. Since $\left\{\Phi_{k}\right\}_{k \in \mathbb{N}^{d}}$ is an orthonormal basis of $H$, every eigenfunction of $A$ is of this form. As a further consequence, for $g \in H$ of the form $g=\bigotimes_{i=1}^{d} g_{i}$ with $g_{i} \in H_{i}$, we also have that the solution of the initial value problem $u^{\prime}+A u=0, u(0)=g$ takes the form

$$
\begin{equation*}
u(t)=\bigotimes_{i=1}^{d} e^{-t A_{i}} g_{i} \tag{1.7}
\end{equation*}
$$

where $e^{-t A_{i}} g_{i}$ solves an analogous evolution problem for the operator $A_{i}$, with initial value $g_{i}$ on $H_{i}$ for each $i$. Thus, for $A$ and $f$ of this particular structure, both the eigenvalue problem and the associated homogeneous evolution problem can be reduced to problems on the component spaces $H_{1}, \ldots, H_{d}$.

However, such simple product structures no longer apply when considering slightly different problems, such as non-homogeneous evolution equations $u^{\prime}+A u=$ $f$ or associated stationary problems $A u=f$. In these cases there is generally no longer a straightforward reduction to lower-dimensional problems. In some cases we can still find semi-analytical approximations that yield a reduction: given $A$ of the Laplace-like form (1.6), for $u$ solving $A u=f=\bigotimes_{i=1}^{d} f_{i}$ we have efficient exponential sum approximations of the form

$$
\begin{equation*}
u=A^{-1} f \approx \sum_{k=1}^{r} \omega_{k} \bigotimes_{i=1}^{d} e^{-\alpha_{k} A_{i}} f_{i} \tag{1.8}
\end{equation*}
$$

It can be shown, as detailed in Section 4.2, that there exist coefficients $\omega_{k}, \alpha_{k}>0$ for $k=1, \ldots, r$ such that we can achieve approximation error $\varepsilon$ for $u$ in suitable norms with $r \approx|\log \varepsilon|^{p}$, with some $p \geq 1$ depending in particular on the chosen
norm. Note that the same construction can still be applied to $f$ given by a sum of elementary tensor products.

Approximations as in (1.8) indicate that sums of separable terms can provide a more general framework for reducing high- to lower-dimensional problems. The above example, however, still depends quite strongly on the particular structure of the problem. This motivates more generally applicable computational methods based on low-rank tensor representations. In such approaches, solutions $u$ of quite generic problems are approximated by more adaptable structures for separating coordinates which, however, still offer the advantage of reducing the computational effort.

### 1.4. Low-rank tensor representations

The notion of tensors can be approached in different ways. For a detailed treatment of different definitions, we refer to Lim (2021). Although tensors can be defined in a completely abstract manner as elements of tensor products of vector spaces, they can also be regarded as describing multilinear mappings on Cartesian products of vector spaces; with respect to bases of these vector spaces, tensors can be described by their coefficient arrays, which need to satisfy certain invariance properties with respect to changes of bases. Note that in physics, tensor fields (that is, tensorvalued mappings on manifolds) are also often referred to simply as tensors, but these geometric objects do not play a role in what follows.

The discrete analogue of separable functions, as considered in the previous subsection, is given by tensors of rank one, that is, elementary tensors of the form $\mathbf{t}=\mathbf{v}^{(1)} \otimes \mathbf{v}^{(2)} \otimes \cdots \otimes \mathbf{v}^{(d)}$ with $\mathbf{v}^{(i)} \in \mathbb{R}^{n_{i}}$ for $i=1, \ldots, d$. The entries of $\mathbf{t}$ are then given by

$$
\mathbf{t}\left[j_{1}, \ldots, j_{d}\right]=\prod_{i=1}^{d} \mathbf{v}^{(i)}\left[j_{i}\right] .
$$

A natural generalization is the canonical tensor format: clearly, each $\mathbf{t} \in$ $\mathbb{R}^{n_{1} \times \cdots \times n_{d}}$ can be represented as a sum of elementary tensors in the form

$$
\begin{equation*}
\mathbf{t}=\sum_{k=1}^{r} \bigotimes_{i=1}^{d} \mathbf{v}_{k}^{(i)}, \tag{1.9}
\end{equation*}
$$

with $\mathbf{v}_{k}^{(i)} \in \mathbb{R}^{n_{i}}$ for $k=1, \ldots, r, i=1, \ldots, d$, where the smallest possible $r$ in such a representation is called the tensor rank of $\mathbf{t}$. This tensor format, also called the canonical polyadic (CP) format, appeared already in Hitchcock $(1927,1928)$.

The use of representations (1.9) in iterative methods for solving high-dimensional partial differential equations was proposed in the seminal papers by Beylkin and Mohlenkamp (2002, 2005). An important requirement for such algorithms is a computational procedure for approximating a given tensor $\mathbf{t}^{*}$ by a tensor of lower rank, for instance the lowest rank permissible for a certain accuracy. In other words,
for some prescribed rank $r$, we aim to minimize $\left\|\mathbf{t}^{*}-\mathbf{a}\right\|_{2}$ over all tensors a of the same mode sizes as $\mathbf{t}^{*}$ with rank at most $r$.

As demonstrated by de Silva and Lim (2008), this problem of best approximation of tensor rank $r$ is generally ill-posed for representations of the form (1.9). This is illustrated by the following example (de Silva and Lim 2008, Prop. 4.6): let $\mathbf{x}_{i}, \mathbf{y}_{i}$ be linearly independent vectors in $\mathbb{R}^{n_{i}}$ for $i=1,2,3$, and define the tensors

$$
\begin{equation*}
\mathbf{t}_{n}=n\left(\mathbf{x}_{1}+\frac{1}{n} \mathbf{y}_{1}\right) \otimes\left(\mathbf{x}_{2}+\frac{1}{n} \mathbf{y}_{2}\right) \otimes\left(\mathbf{x}_{3}+\frac{1}{n} \mathbf{y}_{3}\right)-n \mathbf{x}_{1} \otimes \mathbf{x}_{2} \otimes \mathbf{x}_{3}, \tag{1.10}
\end{equation*}
$$

which are of rank two. Then the limit of the tensors $\mathbf{t}_{n}$ as $n \rightarrow \infty$ is

$$
\mathbf{t}^{*}=\mathbf{y}_{1} \otimes \mathbf{x}_{2} \otimes \mathbf{x}_{3}+\mathbf{x}_{1} \otimes \mathbf{y}_{2} \otimes \mathbf{x}_{3}+\mathbf{x}_{1} \otimes \mathbf{x}_{2} \otimes \mathbf{y}_{3}
$$

which can be verified to have rank three.
This shows that, unlike the matrix rank for $d=2$, for orders $d>2$, the tensor rank is not lower semicontinuous, and in particular, the infimum

$$
\inf \left\{\left\|\mathbf{t}^{*}-\mathbf{a}\right\|_{2}: \mathbf{a} \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}} \text { has tensor rank at most } 2\right\}=0
$$

is not attained. The properties of the tensor rank in the higher-order case are different from those of the matrix rank in other respects too. For instance, as shown by Hillar and Lim (2013), determining the tensor rank of a given tensor with $d>2$ is generally an NP-hard problem.

The ill-posedness of the best tensor rank- $r$ approximation problem poses a fundamental difficulty in the construction of reliable numerical approximation procedures, even when we are content with near-best approximations. These issues can be circumvented by different types of low-rank representations that can be characterized in terms of matrix ranks. The most classical instance of such a representation is the Tucker format, considered in the case $d=3$ by Tucker (1964, 1966). In this particular case, for $\mathbf{t} \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$, a representation in Tucker format takes the form

$$
\begin{equation*}
\mathbf{t}=\sum_{k_{1}=1}^{r_{1}} \sum_{k_{2}=1}^{r_{2}} \sum_{k_{3}=1}^{r_{3}} \mathbf{a}\left[k_{1}, k_{2}, k_{3}\right] \mathbf{U}_{k_{1}}^{(1)} \otimes \mathbf{U}_{k_{2}}^{(2)} \otimes \mathbf{U}_{k_{3}}^{(3)} \tag{1.11}
\end{equation*}
$$

with the core tensor $\mathbf{a} \in \mathbb{R}^{r_{1} \times r_{2} \times r_{3}}$ and vectors $\mathbf{U}_{k}^{(i)} \in \mathbb{R}^{n_{i}}$ for $k=1, \ldots, r_{i}$ and $i=1,2,3$.

The entrywise smallest possible tuple ( $r_{1}, r_{2}, r_{3}$ ) for representing $\mathbf{t}$ in the form (1.11) is called the multilinear rank of $\mathbf{t}$. It can be characterized in terms of matrix ranks: for the multilinear rank $\left(r_{1}, r_{2}, r_{3}\right)$ of $\mathbf{t}$, with the matricizations

$$
\begin{aligned}
& \mathbf{M}_{1}=\left(\mathbf{t}\left[v_{1}, v_{2}, v_{3}\right]\right)_{v_{1} \in\left\{1, \ldots, n_{1}\right\},\left(v_{2}, v_{3}\right) \in\left\{1, \ldots, n_{2}\right\} \times\left\{1, \ldots, n_{3}\right\},}, \\
& \mathbf{M}_{2}=\left(\mathbf{t}\left[v_{1}, v_{2}, v_{3}\right]\right)_{v_{2} \in\left\{1, \ldots, n_{2}\right\},\left(v_{1}, v_{3}\right) \in\left\{1, \ldots, n_{1}\right\} \times\left\{1, \ldots, n_{3}\right\},}, \\
& \mathbf{M}_{3}=\left(\mathbf{t}\left[v_{1}, v_{2}, v_{3}\right]\right)_{v_{3} \in\left\{1, \ldots, n_{3}\right\},\left(v_{1}, v_{2}\right) \in\left\{1, \ldots, n_{1}\right\} \times\left\{1, \ldots, n_{2}\right\}},
\end{aligned}
$$

of $\mathbf{t}$, we have

$$
r_{1}=\operatorname{rank} \mathbf{M}_{1}, \quad r_{2}=\operatorname{rank} \mathbf{M}_{2}, \quad r_{3}=\operatorname{rank} \mathbf{M}_{3} .
$$

In other words, we obtain the most efficient representations of the form (1.11) by choosing $\left\{\mathbf{U}_{k}^{(i)}\right\}_{k=1, \ldots, r_{i}}$ as a basis of the subspace range $\left(\mathbf{M}_{i}\right) \subset \mathbb{R}^{n_{i}}$ for $i=1,2,3$.

For the Tucker format, the existence of best approximations of prescribed rank is guaranteed (see Section 2.7). Moreover, standard linear algebra procedures for matrices can be used for the computation of near-best approximations by the higher-order singular value decomposition, which can serve as the basis of reliable computational methods. However, this format by itself is not suitable for avoiding the curse of dimensionality, since in general the number of entries $\prod_{i=1}^{d} r_{i}$ of the core tensor a still scales exponentially with respect to $d$.

This problem is addressed by more general subspace tensor formats that use further decompositions of a into lower-order tensors, which amounts to choosing subspaces not only for single tensor modes as in the Tucker format, but for groups of modes. Instances of such formats that are suitable for problems in very high dimensions are hierarchical tensors, also known as tree tensor networks in physics, and the special case of tensor trains, which are also known as matrix product states. For tensors that satisfy further constraints on additional rank parameters, representations in these tensor formats can achieve favourable scaling with respect to $d$. Note that there are two cases in which all formats considered below coincide: they reduce to standard representations of low-rank matrices for order $d=2$, and at arbitrary order they contain elementary tensors as special cases.

In general, a tensor format defines a multilinear representation mapping $\tau$ that for certain fixed parameters (such as tensor order) maps a corresponding number $N$ of representation component tensors $\mathbf{X}_{1}, \ldots, \mathbf{X}_{N}$ to the represented tensor,

$$
\mathbf{t}=\tau(\mathrm{X}), \quad \mathrm{X}=\left(\mathbf{X}_{1}, \ldots, \mathbf{X}_{N}\right)
$$

In the simplest case of low-rank matrices in $\mathbb{R}^{n_{1} \times n_{2}}$ with factors $\mathbf{X}_{1} \in \mathbb{R}^{n_{1} \times r}$ and $\mathbf{X}_{2} \in \mathbb{R}^{n_{2} \times r}$, the usual representation mapping is

$$
\begin{equation*}
\tau\left(\mathbf{X}_{1}, \mathbf{X}_{2}\right)=\mathbf{X}_{1} \mathbf{X}_{2}^{\top} \tag{1.12}
\end{equation*}
$$

for any rank parameter $r$. In the case of the Tucker format in $d$ dimensions, the arguments of the representation mapping are $\mathbf{a} \in \mathbb{R}^{r_{1} \times \cdots \times r_{d}}$ and the matrices $\mathbf{U}^{(i)} \in \mathbb{R}^{n_{i} \times r_{i}}$ with columns $\mathbf{U}_{k}^{(i)}, k=1, \ldots, r_{i}$ as above for $i=1, \ldots, d$.

These representation mappings are not injective: for instance, for any given $\mathbf{X}_{1}$, $\mathbf{X}_{2}$ as in (1.12) and $\mathbf{G} \in \operatorname{GL}(r, \mathbb{R})$, we have $\tau\left(\mathbf{X}_{1}, \mathbf{X}_{2}\right)=\tau\left(\mathbf{X}_{1} \mathbf{G}^{-1}, \mathbf{X}_{2} \mathbf{G}^{\top}\right)$. A central feature of tensor formats is that they are closed under vector space operations, that is, under scaling and addition of tensors. For example, the addition of two low-rank matrices can again be written in the form (1.12) with summed rank parameters by a simple concatenation of components, as detailed in Section 2.2.

### 1.5. Low-rank approximations of functions

Low-rank tensors provide a natural tool for approximating functions on highdimensional product domains. We exemplify this by the approximation of elements of $L_{2}\left(\square_{d}\right)$ by tensor product basis expansions.

Let $\left\{\varphi_{\nu}\right\}_{\nu \in \mathbb{N}}$ be an orthonormal basis of $L_{2}(0,1)$. Then, for any $f \in L_{2}\left(\square_{d}\right)$, we have the basis expansion

$$
f=\sum_{\left(v_{1}, \ldots, v_{d}\right) \in \mathbb{N}^{d}} \mathbf{f}\left[v_{1}, \ldots, v_{d}\right] \bigotimes_{i=1}^{d} \varphi_{v_{i}}
$$

where

$$
\mathbf{f}\left[v_{1}, \ldots, v_{d}\right]=\int_{\square_{d}} f \bigotimes_{i=1}^{d} \varphi_{v_{i}} \mathrm{~d} x
$$

Note that the mapping $f \mapsto \mathbf{f}$ defined in this manner is an isometry between $L_{2}\left(\square_{d}\right)$ and $\ell_{2}\left(\mathbb{N}^{d}\right)$. Due to the tensor product structure of this mapping, if $f$ has low-rank structure, it is inherited by the basis coefficients $\mathbf{f}$ : in the simplest case, if $f=\bigotimes_{i=1}^{d} f_{i}$ with $f_{i} \in L_{2}(0,1)$ for $i=1, \ldots, d$, then

$$
\mathbf{f}=\bigotimes_{i=1}^{d} \mathbf{f}_{i}, \quad \text { where } \mathbf{f}_{i}[v]=\int_{0}^{1} f_{i} \varphi_{v} \mathrm{~d} x, v \in \mathbb{N} .
$$

This hints at the usefulness of considering tensor representations not only for tensors of finite mode sizes but also on sequence spaces defined on high-dimensional index sets such as $\ell_{2}\left(\mathbb{N}^{d}\right)$. In this setting, infinite basis expansions are covered in a natural way, and low-rank approximations can be studied independently of any particular choice of finitely many summands in the basis expansion.

The reduction to a problem on $\ell_{2}\left(\mathbb{N}^{d}\right)$ via a basis representation is also convenient for computational purposes, since standard algorithms for rank reduction in subspace tensor formats provide error estimates in this norm. However, an additional difficulty that is important for applications to partial differential equations is that for Sobolev spaces $H^{s}\left(\square_{d}\right)$, the isomorphism to $\ell_{2}\left(\mathbb{N}^{d}\right)$ provided by suitable bases is not of tensor product form, and thus does not preserve low-rank structures. As considered in detail in Section 4, this issue can be overcome by what amounts to a discretization-dependent low-rank approximation of this isomorphism.

Tensor approximations are of interest not only for problems of high dimensionality: although it is natural to assign tensor modes to independent variables in an approximation problem, entirely different parametrizations of functions by tensors can be useful for particular purposes. For instance, in the context of occupation number representations of symmetric or antisymmetric wavefunctions in quantum physics, each tensor mode corresponds to a lower-dimensional basis function. Another example is that of methods based on reinterpreting basis expansion coefficients of lower-dimensional functions as high-order tensors. In this case, by a suitable encoding of indices, tensor modes correspond to scales in the approximation problem. In particular, this leads to efficient generic approximations for problems with singularities. These examples are considered in detail in Section 3. In such a context, classical notions of dimensionality need no longer be the determining factor for the computational complexity of low-rank approximations.

### 1.6. Low-rank methods as an instance of nonlinear approximation schemes

An important feature of low-rank tensor formats is that they are closed under addition, but addition of two tensor representations generally leads to an increase of the associated rank parameters. In the simplest case of matrices, the result of adding two matrices of ranks $r_{1}$ and $r_{2}$ may have rank up to $r_{1}+r_{2}$. In lowrank methods, approximations of functions are thus not sought from linear spaces (spanned, for instance, by a fixed set of basis functions) but from sets with nonlinear structure, where the representation complexity of elements increases under vector space operations.

This difficulty arises in an analogous manner in sparse nonlinear approximations with dictionaries; see for instance the survey by DeVore (1998). Here, we choose a fixed dictionary of functions $\left\{\psi_{\nu}\right\}_{\nu \in \mathbb{N}}$, such as a Riesz basis or frame of the relevant function space $H$. For a given approximand $f \in H$ and error tolerance $\varepsilon>0$, we choose a suitable finite subset $\Lambda_{\varepsilon}(f) \subset \mathbb{N}$ such that $f$ has an approximation $f_{\varepsilon} \in \operatorname{span}\left\{\psi_{v}: v \in \Lambda_{\varepsilon}(f)\right\}$ with $\left\|f-f_{\varepsilon}\right\|_{H} \leq \varepsilon$. This notion of sparse approximation in some sense involves a lesser degree of nonlinearity than the low-rank approximations considered above, since $f_{\varepsilon}$ is parametrized linearly by its expansion coefficients. However, we have a nonlinear dependence on $f$ via the choice of non-zero entries of the coefficient sequence, and when adding sparse approximations of two different functions with $n_{1}$ and $n_{2}$ non-zero coefficients, respectively, the result may also have up to $n_{1}+n_{2}$ non-zero coefficients. Hence, in contrast to approximations from fixed linear subspaces, the complexity of representations generally increases under vector space operations.

A natural benchmark for sparse approximations is provided by best $n$-term approximation, where in the above setting, for given $f \in H$ and for each budget of $n \in \mathbb{N}$ terms, we consider the error achievable by a best choice of $\Lambda_{n}(f) \subset \mathbb{N}$ with $\# \Lambda_{n}(f) \leq n$, that is, the best $n$-term errors

$$
\min _{\substack{\Lambda_{n}(f) \subset \mathbb{N} \\ \# \Lambda_{n}(f) \leq n}} \min \left\{\left\|f-f_{n}\right\|_{H}: f_{n} \in \operatorname{span}\left\{\psi_{v}: v \in \Lambda_{n}(f)\right\}\right\} .
$$

Ideally, numerical methods for computing sparse approximations should yield errors proportional to the best $n$-term errors when using $n$ coefficients or (as a much stronger requirement) when using a multiple of $n$ elementary operations. Such questions also arise in the context of low-rank methods when we need to identify suitable lower-dimensional basis expansions in each tensor mode. This combination of sparse and low-rank approximations is discussed in further detail in Sections 3.2 and 5.3.

Another important instance of nonlinear approximations is neural networks, where inputs are transformed by several layers of alternating linear mappings with entrywise nonlinear transformations such as $\operatorname{ReLU}(x)=\max \{x, 0\}$. Although neural networks can in principle approximate many classes of functions very efficiently, either matching or surpassing the performance of classical approximation
methods, finding such approximations as solutions of regression problems by machine learning methods can generally be quite a problematic task. In this regard, a distinctive feature of subspace tensor formats is that the generation of near-best low-rank approximations can be reduced to the computation of singular value decompositions, for which reliable numerical algorithms are available.

### 1.7. Basic algorithmic concepts

Standard methods for solving partial differential equations eventually amount to solving systems of equations for grid values or coefficients with respect to some basis. In the case of low-rank approximations, we are facing the additional difficulty that these basis coefficients are represented indirectly in terms of components of a tensor representation. Moreover, the rank parameters that are required for a certain approximation quality are typically not known a priori. Since the process of generating approximate solutions in low-rank format is not straightforward, the computational complexity of low-rank solvers is just as important as the basic approximability of solutions when considering the viability of such approximations.

With the exception of semi-analytical explicit low-rank approximations as in (1.8) that exist for problems with certain special structures (see Section 4.2), methods for computing low-rank approximations are generally iterative: starting from a certain initial approximation $\tau\left(\mathrm{X}^{0}\right)$ of a solution in low-rank format with representation $\mathrm{X}^{0}$, each step from $\tau\left(\mathrm{X}^{n}\right)$ to $\tau\left(\mathrm{X}^{n+1}\right)$ is performed by operating only on the representation components $X^{n}$.

Concerning the construction of such iterations, one can broadly distinguish two basic approaches. One common strategy is to realize standard iterative solvers, such as Krylov subspace methods, in tensor format. Perhaps the simplest example is Richardson iteration for solving systems of linear equations $\mathbf{A x}=\mathbf{b}$ for positive definite $\mathbf{A}$. With starting vector $\mathbf{x}^{0}$ and a damping parameter $\omega>0$, the iteration reads

$$
\begin{equation*}
\mathbf{x}^{n+1}=\mathbf{x}^{n}-\omega\left(\mathbf{A} \mathbf{x}^{n}-\mathbf{b}\right) . \tag{1.13}
\end{equation*}
$$

To realize this iteration in low-rank form, the problem data $\mathbf{A}$ and $\mathbf{b}$ also need to be represented in a compatible format. For instance, let $d=2$ with

$$
\mathbf{x}^{n}=\sum_{k=1}^{r} \mathbf{v}_{k}^{(1)} \otimes \mathbf{v}_{k}^{(2)} \quad \text { and } \quad \mathbf{A}=\sum_{\ell=1}^{R} \mathbf{A}_{\ell}^{(1)} \otimes \mathbf{A}_{\ell}^{(2)}
$$

Then

$$
\begin{equation*}
\mathbf{A} \mathbf{x}^{n}=\sum_{\ell=1}^{R} \sum_{k=1}^{r} \mathbf{A}_{\ell}^{(1)} \mathbf{v}_{k}^{(1)} \otimes \mathbf{A}_{\ell}^{(2)} \mathbf{v}_{k}^{(2)} \tag{1.14}
\end{equation*}
$$

that is, we still have a representation in the same low-rank format, but with ranks multiplied by (at most) $R$. Except in special cases such as $R=1$, without further
precautions, the rank parameters of the representations $X^{n}$ of the iterates $\mathbf{x}^{n}$ in (1.13) will thus increase exponentially with respect to $n$.

The iteration (1.13) thus needs to be combined with a rank reduction: let $\mathrm{X}^{n+1}=\mathcal{F}\left(\mathrm{X}^{n}\right)$ be the iteration on the low-rank representations of the $\mathbf{x}^{n}$ defined by (1.13); we then replace this with $X^{n+1}=(\mathcal{R} \circ \mathcal{F})\left(X^{n}\right)$, where $\mathcal{R}$ is a suitable rank reduction procedure, such as a truncated singular value decomposition or one of its higher-order generalizations for subspace tensor formats. In devising criteria for performing this rank reduction, we need to balance computational costs against preserving the convergence of the iteration. When always truncating to some fixed rank parameter, the costs in each step are controlled, but the convergence of the iteration is difficult to guarantee; conversely, when truncating up to a prescribed error, it becomes a non-trivial problem to estimate the resulting ranks and hence the computational costs.

The second strategy is to exploit the particular structure of the tensor format and to directly optimize the components of the representation. Again assuming $\mathbf{A}$ to be positive definite, we can rewrite the original system of equations as a optimization problem: find a low-rank representation X such that

$$
\frac{1}{2}\langle\mathbf{A} \tau(\mathrm{X}), \tau(\mathrm{X})\rangle-\langle\mathbf{b}, \tau(\mathrm{X})\rangle
$$

is minimized. Variational formulations of this kind can be combined with alternating minimization procedures that operate on one or two components in $X$ at a time. Solvers can also be constructed to use the Riemannian manifold structure of fixed-rank tensors.

There exist many different realizations of the two approaches described above, as well as methods that combine elements of both. An overview of some prominent examples is given in Section 5.1. In the remainder of Section 5 we then focus on methods that achieve guarantees on convergence and computational complexity by striking a balance between error reduction and controlled ranks.

### 1.8. Notes on the literature

There is now a large body of literature on low-rank tensor approximations. For a survey on classical tensor representations including canonical and Tucker formats, we refer to Kolda and Bader (2009). Tensors in the context of computational applications have also been considered in surveys by Grasedyck, Kressner and Tobler (2013), Hackbusch (2014), Khoromskij (2015), Bachmayr, Schneider and Uschmajew (2016) and Bachmayr and Dahmen (2020), as well as in the monographs by Hackbusch (2019) and Khoromskij (2018). The canonical tensor rank and its geometric and algebraic aspects are considered in a monograph by Landsberg (2012) and in a survey by Lim (2021).

In the present article, at the expense of developments that are aimed at other applications, such as machine learning, we consider low-rank methods specifically from the perspective of the numerical solution of partial differential equations.

Although we aim at a general overview of such methods, we focus on the use of subspace tensor formats in schemes with rigorous deterministic performance guarantees.

Concerning applications of tensor representations in signal processing and machine learning not covered by this article, we refer to the surveys by Signoretto, Tran Dinh, De Lathauwer and Suykens (2014), Cichocki et al. (2015), Sidiropoulos et al. (2017) and Janzamin et al. (2019), as well as to Stoudenmire and Schwab (2016), Chen, Batselier, Suykens and Wong (2017), Novikov, Trofimov and Oseledets (2018) and Michel and Nouy (2022). Concerning methods for solving PDEs using tensor representations with random sampling, which we do not consider in further detail, see for example Richter, Sallandt and Nüsken (2021), Oster, Sallandt and Schneider (2022) and Fackeldey, Oster, Sallandt and Schneider (2022).

## 2. Subspace tensor formats

In this section we consider in detail the particular class of low-rank representations of higher-order tensors that provide the foundation for all numerical methods that follow. We first introduce the basic notions of tensor products of abstract vector spaces. To speak about the asymptotic low-rank approximability of functions and construct numerical methods with discretization-independent performance, we need to consider tensor products of infinite-dimensional spaces, where we restrict ourselves to Hilbert spaces.

### 2.1. Tensor product Hilbert spaces

For continuous functions $f$ and $g$, in analogy to the outer product of vectors, we have $(f \otimes g)(x, y)=f(x) g(y)$ for $x$ and $y$ from the respective domains. However, for Hilbert spaces such as $L_{2}$-spaces for which point evaluation is not defined, a more generally applicable construction of tensor products is required. Here we first define the algebraic tensor product of vector spaces, which contains all finite linear combinations of elementary tensors, and subsequently consider closures in suitable norms to obtain topological tensor products. Treatments of the basic notions considered in this section can also be found in Light and Cheney (1985), Hackbusch (2019) and Khoromskij (2018). For the more general case of Banach spaces, we refer to Defant and Floret (1993), Ryan (2002) and Hackbusch (2019).

For arbitrary vector spaces $V, W$, the algebraic tensor product $V \otimes_{a} W$ provides an abstract notion of finite sums of elementary tensor products of the form

$$
\begin{equation*}
\sum_{k=1}^{r} \lambda_{k} v_{k} \otimes w_{k} \tag{2.1}
\end{equation*}
$$

where $r \in \mathbb{N}$ and $\lambda_{k} \in \mathbb{R}, v_{k} \in V, w_{k} \in W$ for $k=1, \ldots, r$. To this end, we define

$$
T_{a}(V, W)=\bigcup_{n=1}^{\infty}(\mathbb{R} \times V \times W)^{n}
$$

with the natural addition operation by concatenation, and we identify expressions

$$
\begin{equation*}
t=\left(\left(\lambda_{k}, v_{k}, w_{k}\right)\right)_{k=1, \ldots, r} \in T_{a}(V, W) \tag{2.2}
\end{equation*}
$$

with functionals on all bilinear forms $b$ on $V \times W$ by

$$
t(b)=\sum_{k=1}^{r} \lambda_{k} b\left(v_{k}, w_{k}\right)
$$

We then define an equivalence relation $\sim_{a}$ on $T_{a}(V, W)$ by treating two expressions $t, \tilde{t} \in T_{a}(V, W)$ as equivalent if and only if they give rise to the same such functional, that is, if $t(b)=\tilde{t}(b)$ for all bilinear forms $b$ on $V \times W$.

Definition 2.1. The algebraic tensor product $V \otimes_{a} W$ of two vector spaces is defined as the quotient vector space $T_{a}(V, W) / \sim_{a}$.

For the equivalence class in $V \otimes_{a} W$ represented by $t$ as in (2.2), we use the notation (2.1). We have the following basic properties; see Light and Cheney (1985, Sec. 1) and Hackbusch (2019, Sec. 3.2).

## Proposition 2.2.

(i) For $\alpha, \beta, \gamma, \delta \in \mathbb{R}, v_{1}, v_{2} \in V$ and $w_{1}, w_{2} \in W$,

$$
\left(\alpha v_{1}+\beta v_{2}\right) \otimes\left(\gamma w_{1}+\delta w_{2}\right)=\alpha \gamma v_{1} \otimes w_{1}+\alpha \delta v_{1} \otimes w_{2}+\beta \gamma v_{2} \otimes w_{1}+\beta \delta v_{2} \otimes w_{2}
$$

(ii) If $\left\{v_{1}, \ldots, v_{r_{1}}\right\} \subset V$ and $\left\{w_{1}, \ldots, w_{r_{2}}\right\} \subset W$ with $r_{1}, r_{2} \in \mathbb{N}$ are linearly independent, then $\left\{v_{i} \otimes w_{j}: i=1, \ldots, r_{1}, j=1, \ldots, r_{2}\right\} \subset V \otimes_{a} W$ is linearly independent.

The construction of algebraic tensor products can be iterated and is associative up to isomorphism, that is, $U \otimes_{a}\left(V \otimes_{a} W\right) \simeq\left(U \otimes_{a} V\right) \otimes_{a} W$ for vector spaces $U, V, W$; we thus simply write $U \otimes_{a} V \otimes_{a} W$.

Topological tensor products are obtained by taking closures of algebraic tensor products in specific norms. For the purposes of this article, it will be sufficient to consider topological tensor products of Hilbert spaces. In contrast to the case of more general Banach spaces, we then have a single canonical choice of tensor product norm. Let $H_{1}, \ldots, H_{d}$ be Hilbert spaces. For elementary tensors $v=$ $\otimes_{i=1}^{d} v_{i}, w=\bigotimes_{i=1}^{d} w_{i} \in H_{1} \otimes_{a} \cdots \otimes_{a} H_{d}$, we define the inner product

$$
\begin{equation*}
\langle v, w\rangle_{H}=\prod_{i=1}^{d}\left\langle v_{i}, w_{i}\right\rangle_{H_{i}}, \tag{2.3}
\end{equation*}
$$

extend this to sums of elementary tensors by bilinearity, and define the tensor product Hilbert space $H=H_{1} \otimes \cdots \otimes H_{d}$ as the closure of $H_{1} \otimes_{a} \cdots \otimes_{a} H_{d}$ with respect to the induced norm. By definition, the norm induced by (2.3) is a
cross-norm, that is,

$$
\begin{equation*}
\left\|v_{1} \otimes \cdots \otimes v_{d}\right\|_{H}=\prod_{i=1}^{d}\left\|v_{i}\right\|_{H_{i}} \quad \text { for } v_{i} \in H_{i}, i=1, \ldots, d \tag{2.4}
\end{equation*}
$$

Unless otherwise stated, for tensor product Hilbert spaces, we assume this canonical choice of inner product (2.3) and induced norm in what follows.

For tensor product Hilbert spaces $H=\bigotimes_{i=1}^{d} H_{i}$ and $G=\bigotimes_{i=1}^{d} G_{i}$ and bounded linear operators $A_{i}: H_{i} \rightarrow G_{i}$ for $i=1, \ldots, d$, we define the tensor product operator $A_{1} \otimes \cdots \otimes A_{d}$ by its action on elementary tensor products as

$$
\begin{equation*}
\left(A_{1} \otimes \cdots \otimes A_{d}\right)\left(v_{1} \otimes \cdots \otimes v_{d}\right)=\left(A_{1} v_{1}\right) \otimes \cdots \otimes\left(A_{d} v_{d}\right) \tag{2.5}
\end{equation*}
$$

This definition is uniquely extended to finite sums of elementary tensors by linearity and to infinite sums by the unique extension of bounded operators defined on dense subsets. The following identity can be obtained immediately from Light and Cheney (1985, Lem. 1.30).

Theorem 2.3. For $A_{1}, \ldots, A_{d}$ as in (2.5),

$$
\left\|A_{1} \otimes \cdots \otimes A_{d}\right\|_{H \rightarrow G}=\prod_{i=1}^{d}\left\|A_{i}\right\|_{H_{i} \rightarrow G_{i}}
$$

When each of the mappings $A_{1}, \ldots, A_{d}$ in (2.5) is an isomorphism, the corresponding tensor product operator $A_{1} \otimes \cdots \otimes A_{d}$ is called a tensor space isomorphism, and we have $\left(A_{1} \otimes \cdots \otimes A_{d}\right)^{-1}=A_{1}^{-1} \otimes \cdots \otimes A_{d}^{-1}$.
Remark 2.4. For countable sets $\mathcal{I}_{1}, \mathcal{I}_{2}$, we have $\ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)=\ell_{2}\left(\mathcal{I}_{1}\right) \otimes \ell_{2}\left(\mathcal{I}_{2}\right)$. In other words, the norm $\|\cdot\|_{\ell_{2}\left(\mathcal{I}_{1}\right) \otimes \ell_{2}\left(\mathcal{I}_{2}\right)}$ defined by (2.4) coincides with the norm of $\ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)$ and $\operatorname{span}\left\{\mathbf{v}_{1} \otimes \mathbf{v}_{2}: \mathbf{v}_{1} \in \ell_{2}\left(\mathcal{I}_{1}\right), \mathbf{v}_{2} \in \ell_{2}\left(\mathcal{I}_{2}\right)\right\}$ is dense in $\ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)$. Similarly, for measurable sets $\Omega_{1}, \Omega_{2}$, we have $L_{2}\left(\Omega_{1} \times \Omega_{2}\right)=L_{2}\left(\Omega_{1}\right) \otimes L_{2}\left(\Omega_{2}\right)$; for a proof, see Light and Cheney (1985, Thm 1.39).

It will suffice to restrict ourselves to separable Hilbert spaces in what follows. Any such space, by any choice of orthonormal basis, is isometrically isomorphic to $\ell_{2}(\mathbb{N})$. In general, such an isomorphism does not respect tensor product structures. These are preserved, however, in the case of tensor product orthonormal bases, as formulated in the following theorem; for a proof, see Weidmann (1980, Thm 3.12).

Theorem 2.5. Let $H_{1}, \ldots, H_{d}$ be separable Hilbert spaces, and for $i=1, \ldots, d$, let $\left\{\varphi_{\nu}^{(i)}\right\}_{\nu \in \mathbb{N}}$ be an orthonormal basis of $H_{i}$. Then

$$
\left\{\bigotimes_{i=1}^{d} \varphi_{v_{i}}^{(i)}\right\}\left(v_{1}, \ldots, v_{d}\right) \in \mathbb{N}^{d}
$$

is an orthonormal basis of $\bigotimes_{i=1}^{d} H_{i}$.

Remark 2.6. With the basis isomorphisms

$$
F_{i}: \ell_{2}(\mathbb{N}) \rightarrow H_{i}, \quad \mathbf{v} \mapsto \sum_{v \in \mathbb{N}} \mathbf{v}[v] \varphi_{v}^{(i)}
$$

as a consequence of the cross-norm property (2.4), the isomorphism associated to the tensor product basis

$$
F: \ell_{2}\left(\mathbb{N}^{d}\right) \rightarrow \bigotimes_{i=1}^{d} H_{i}, \quad \mathbf{v} \mapsto \sum_{v \in \mathbb{N}^{d}} \mathbf{v}[v] \bigotimes_{i=1}^{d} \varphi_{v_{i}}^{(i)}
$$

can be written in tensor product form as $F=\bigotimes_{i=1}^{d} F_{i}$ and is thus a tensor space isomorphism.

For countable sets $\mathcal{I}_{1}, \mathcal{I}_{2}$, elements of $\ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)=\ell_{2}\left(\mathcal{I}_{1}\right) \otimes \ell_{2}\left(\mathcal{I}_{2}\right)$ can be identified with bounded linear operators from $\ell_{2}\left(\mathcal{I}_{2}\right)$ to $\ell_{2}\left(\mathcal{I}_{1}\right)$ by interpreting their entries as those of an infinite matrix: the operator associated to $\mathbf{u} \in \ell_{2}\left(\mathcal{I}_{1}\right) \otimes \ell_{2}\left(\mathcal{I}_{2}\right)$ is given by

$$
\begin{equation*}
\ell_{2}\left(\mathcal{I}_{2}\right) \ni \mathbf{v} \mapsto\left(\sum_{v \in \mathcal{I}_{2}} \mathbf{u}\left[v^{\prime}, v\right] \mathbf{v}[v]\right)_{v^{\prime} \in \mathcal{I}_{1}} \tag{2.6}
\end{equation*}
$$

where the right-hand side is in $\ell_{2}\left(\mathcal{I}_{1}\right)$ as a consequence of the Cauchy-Schwarz inequality.

Definition 2.7. For countable sets $\mathcal{I}_{1}, \mathcal{I}_{2}$ and $\mathbf{u} \in \ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)$, we define

$$
\operatorname{rank}(\mathbf{u})=\operatorname{dim} \operatorname{range}(\mathbf{u}) \in \mathbb{N}_{0} \cup\{\infty\}
$$

where we identify $\mathbf{u}$ with the operator given by (2.6).

For $\mathbf{u} \in \ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)$, the operator defined by (2.6) belongs to the Hilbert-Schmidt class, where the Hilbert-Schmidt inner product is simply the inner product in $\ell_{2}\left(\mathcal{I}_{1} \times\right.$ $\mathcal{I}_{2}$ ). In particular, Hilbert-Schmidt operators are compact. As a consequence, we obtain the following adaptation of a classical result that goes back to Schmidt (1907).

Theorem 2.8 (Hilbert-Schmidt decomposition). For each $\mathbf{u} \in \ell_{2}\left(\mathbb{N}^{2}\right)$, there exist orthonormal systems $\left\{\mathbf{U}_{k}^{(1)}\right\}_{k \in \mathbb{N}}$ and $\left\{\mathbf{U}_{k}^{(2)}\right\}_{k \in \mathbb{N}}$ in $\ell_{2}(\mathbb{N})$ as well as a nonincreasing, non-negative sequence $\left(\sigma_{k}\right)_{k \in \mathbb{N}} \in \ell_{2}(\mathbb{N})$ such that

$$
\begin{equation*}
\mathbf{u}=\sum_{k=1}^{\infty} \sigma_{k} \mathbf{U}_{k}^{(1)} \otimes \mathbf{U}_{k}^{(2)} \tag{2.7}
\end{equation*}
$$

with convergence in $\ell_{2}\left(\mathbb{N}^{2}\right)$. Moreover, for any $r \in \mathbb{N}$, we have the best approximation property

$$
\begin{align*}
\left\|\mathbf{u}-\sum_{k=1}^{r} \sigma_{k} \mathbf{U}_{k}^{(1)} \otimes \mathbf{U}_{k}^{(2)}\right\|_{\ell_{2}\left(\mathbb{N}^{2}\right)} & =\left(\sum_{k=r+1}^{\infty} \sigma_{k}^{2}\right)^{1 / 2} \\
& =\min \left\{\|\mathbf{u}-\mathbf{v}\|_{\ell_{2}\left(\mathbb{N}^{2}\right)}: \operatorname{rank}(\mathbf{v}) \leq r\right\} \tag{2.8}
\end{align*}
$$

and the best approximation is unique precisely when $\sigma_{r}>\sigma_{r+1}$.
The expansion (2.7) is also called the singular value decomposition (SVD) of the operator induced by $\mathbf{u}$. In this context, the expansion is identified with the mapping on $\ell_{2}(\mathbb{N})$ given by

$$
\sum_{k=1}^{\infty} \sigma_{k}\left\langle\mathbf{U}_{k}^{(2)}, \cdot\right\rangle \mathbf{U}_{k}^{(1)}
$$

The orthonormal systems $\left\{\mathbf{U}_{k}^{(1)}\right\}_{k \in \mathbb{N}}$ and $\left\{\mathbf{U}_{k}^{(2)}\right\}_{k \in \mathbb{N}}$ are called the left and right singular vectors, respectively, and $\left(\sigma_{k}\right)_{k \in \mathbb{N}}$ is the sequence of singular values.

Remark 2.9. Let us note the following immediate consequences of Theorem 2.8.
(i) The same statement applies to arbitrary countable index sets $\mathcal{I}_{1}, \mathcal{I}_{2}$, in which case it yields orthonormal systems in $\ell_{2}\left(\mathcal{I}_{1}\right)$ and $\ell_{2}\left(\mathcal{I}_{2}\right)$. In particular, in this section we also use it with $\mathcal{I}_{1}=\mathbb{N}^{d_{1}}$ and $\mathcal{I}_{2}=\mathbb{N}^{d_{2}}$ with $d_{1}, d_{2} \in \mathbb{N}$.
(ii) When supp $\mathbf{u} \subseteq \mathcal{I}_{1} \times \mathcal{I}_{2}$ with finite $\mathcal{I}_{1}, \mathcal{I}_{2}$, which corresponds to the case of the singular value decomposition of a matrix in $\mathbb{R}^{\mathcal{I}_{1} \times \mathcal{I}_{2}}$ representing a linear mapping from $\mathbb{R}^{\mathcal{I}_{2}}$ to $\mathbb{R}^{\mathcal{I}_{1}}$, the sum in (2.7) has at most $\min \left\{\# \mathcal{I}_{1}, \# \mathcal{I}_{2}\right\}$ summands. In this case Theorem 2.8 is also known as the Eckart-Young theorem (Eckart and Young 1936).

### 2.2. Operations on low-rank matrices

Before turning to higher-order tensors, it may be instructive to first take a closer look at low-rank matrices from the computational point of view, that is, at the procedures for performing typical operations such as addition or singular value decomposition in low-rank form.

There are different possibilities for representing matrices in low-rank form that leave different types of redundancy in the representation. As mentioned in Section 1.4, matrices of rank at most $r$ in $\mathbb{R}^{m \times n}$ can be represented by $\mathbf{A} \in \mathbb{R}^{m \times r}$ and $\mathbf{B} \in \mathbb{R}^{n \times r}$ via the representation mapping

$$
\tau(\mathbf{A}, \mathbf{B})=\mathbf{A} \mathbf{B}^{\top}
$$

where $\tau(\mathbf{A}, \mathbf{B})=\tau\left(\mathbf{A G}^{-1}, \mathbf{B G}^{\top}\right)$ for any $\mathbf{G} \in \mathrm{GL}(r, \mathbb{R})$. In addition, we can demand that either $\mathbf{A}$ or $\mathbf{B}$ have orthonormal columns, in which case $\tau(\mathbf{A}, \mathbf{B})=\tau(\mathbf{A G}, \mathbf{B G})$ for any $\mathbf{G} \in \mathrm{O}(r)$. An alternative representation is via

$$
\begin{equation*}
\tau(\mathbf{U}, \mathbf{S}, \mathbf{V})=\mathbf{U S V}^{\top} \tag{2.9}
\end{equation*}
$$

where $\mathbf{U}, \mathbf{V}$ are required to have orthonormal columns and $\mathbf{S} \in \mathbb{R}^{r \times r}$. In this case $\tau(\mathbf{U}, \mathbf{S}, \mathbf{V})=\tau\left(\mathbf{U} \mathbf{G}_{1}, \mathbf{G}_{1}^{\top} \mathbf{S G}_{2}, \mathbf{V G}_{2}\right)$ for all $\mathbf{G}_{1}, \mathbf{G}_{2} \in \mathrm{O}(r)$. Here $\mathbf{S}$ is diagonal precisely when the columns of $\mathbf{U}$ and $\mathbf{V}$ are left and right singular vectors of $\tau(\mathbf{U}, \mathbf{S}, \mathbf{V})$, respectively, in which case the diagonal entries of $\mathbf{S}$ are (up to sign) the corresponding singular values.

In the following description of operations, we focus on representations of the form (2.9) with $\mathbf{S}$ not necessarily diagonal, since this case is conceptually closest to higher-order tensor representations.

### 2.2.1. Orthogonalization

Let $\tilde{\mathbf{U}} \in \mathbb{R}^{m \times r}, \tilde{\mathbf{V}} \in \mathbb{R}^{n \times r}$ and $\tilde{\mathbf{S}} \in \mathbb{R}^{r \times r}$ be given. Then a representation satisfying the orthogonality requirement can be found as follows. Perform QR decompositions $\tilde{\mathbf{U}}=\mathbf{U R}_{1}, \tilde{\mathbf{V}}=\mathbf{V R}_{2}$ to obtain $\mathbf{U} \in \mathbb{R}^{m \times r}$ and $\mathbf{V} \in \mathbb{R}^{n \times r}$ with orthonormal columns, as well as (triangular) $\mathbf{R}_{1}, \mathbf{R}_{2} \in \mathbb{R}^{r \times r}$; then set $\mathbf{S}=\mathbf{R}_{1} \tilde{\mathbf{S}} \mathbf{R}_{2}^{\top}$ to obtain $\tau(\mathbf{U}, \mathbf{S}, \mathbf{V})=\tau(\tilde{\mathbf{U}}, \tilde{\mathbf{S}}, \tilde{\mathbf{V}})$, where $\mathbf{U}, \mathbf{V}$ have orthonormal columns. Note that in this form the Hilbert-Schmidt and spectral norms of $\tau(\mathbf{U}, \mathbf{S}, \mathbf{V})$ can be obtained as the respective norms of the matrix $\mathbf{S}$, which in this case is generally not diagonal.

Remark 2.10. To perform orthogonalization as outlined above, performing the QR decompositions requires $O\left((m+n) r^{2}\right)$ operations, and computing $\mathbf{S}$ in addition requires $O\left(r^{3}\right)$ operations. These QR decompositions can be computed by a numerically stable direct method, for instance by applying at most $r$ orthogonal Householder reflectors.

### 2.2.2. Addition of low-rank matrices

For $\lambda_{1}, \lambda_{2} \in \mathbb{R}$ and two given representations $\tau\left(\mathbf{U}_{1}, \mathbf{S}_{1}, \mathbf{V}_{1}\right)$ and $\tau\left(\mathbf{U}_{2}, \mathbf{S}_{2}, \mathbf{V}_{2}\right)$ of compatible sizes with ranks $r_{1}$ and $r_{2}$, respectively, to obtain a low-rank representation of $\lambda_{1} \tau\left(\mathbf{U}_{1}, \mathbf{S}_{1}, \mathbf{V}_{1}\right)+\lambda_{2} \tau\left(\mathbf{U}_{2}, \mathbf{S}_{2}, \mathbf{V}_{2}\right)$ we assemble the block matrices

$$
\tilde{\mathbf{U}}=\left[\begin{array}{ll}
\mathbf{U}_{1} & \mathbf{U}_{2}
\end{array}\right], \quad \tilde{\mathbf{S}}=\left[\begin{array}{ll}
\lambda_{1} \mathbf{S}_{1} & \\
& \lambda_{2} \mathbf{S}_{2}
\end{array}\right], \quad \tilde{\mathbf{V}}=\left[\begin{array}{ll}
\mathbf{V}_{1} & \mathbf{V}_{2}
\end{array}\right]
$$

corresponding to a representation with rank at most $r_{1}+r_{2}$, and subsequently orthogonalize with results $\mathbf{U}, \mathbf{S}, \mathbf{V}$ as above to arrive at $\tau(\mathbf{U}, \mathbf{S}, \mathbf{V})=\lambda_{1} \tau\left(\mathbf{U}_{1}, \mathbf{S}_{1}, \mathbf{V}_{1}\right)+$ $\lambda_{2} \tau\left(\mathbf{U}_{2}, \mathbf{S}_{2}, \mathbf{V}_{2}\right)$.

### 2.2.3. Singular value decomposition

To obtain the SVD form of a given representation $\tau(\mathbf{U}, \mathbf{S}, \mathbf{V})$ with $\mathbf{U}, \mathbf{V}$ having orthonormal columns, we compute an SVD of the matrix $\mathbf{S}$,

$$
\mathbf{S}=\mathbf{Q} \mathbf{\Sigma} \mathbf{P}^{\top}
$$

where $\mathbf{Q}, \mathbf{P} \in \mathrm{O}(r)$ and $\boldsymbol{\Sigma}$ is diagonal. Then, as summarized for later use in the following simple observation, the sought SVD form is given by $\tau(\mathbf{U Q}, \mathbf{\Sigma}, \mathbf{V P})$.

Proposition 2.11. Let $\mathbf{M}=\mathbf{U S V}^{\top}$, where $\mathbf{U} \in \mathbb{R}^{m \times r}$ and $\mathbf{V} \in \mathbb{R}^{n \times r}$ have orthonormal columns and $\mathbf{S} \in \mathbb{R}^{r \times r}$. Let $\mathbf{S}=\mathbf{Q} \boldsymbol{\Sigma} \mathbf{P}^{\top}$ be a singular value decomposition of $\mathbf{S}$; then $\mathbf{M}=\hat{\mathbf{U}} \mathbf{\Sigma} \hat{\mathbf{V}}^{\top}$ with $\hat{\mathbf{U}}=\mathbf{U Q}, \hat{\mathbf{V}}=\mathbf{V P}$ is a singular value decomposition of $\mathbf{M}$.

Remark 2.12 (computation of the SVD). Truncation of the SVD yields solutions of best low-rank approximation problems by Theorem 2.8, and the computation of SVDs as outlined above is thus of central importance. Note that the SVD of a rectangular matrix $\mathbf{A}$ can be obtained from the eigendecompositions of the symmetric matrices $\mathbf{A}^{\top} \mathbf{A}$ and $\mathbf{A} \mathbf{A}^{\top}$, which necessarily require an iterative scheme. These latter eigenvalue problems can be solved by transformation to tridiagonal form by a direct method, followed by a QR iteration for the resulting symmetric tridiagonal eigenvalue problems that has been shown by Wilkinson (1968) to be globally at least linearly convergent; see also Parlett (1998, Sec. 8.10). This implies that we obtain an algorithm for approximating the SVD with guaranteed convergence and complexity bounds.

When performed with machine precision, the approach via symmetric eigenvalue problems has the disadvantage that a loss of relative precision can occur for the smallest singular values. This is avoided by the specialized numerical SVD methods of Golub and Kahan (1965) and Golub and Reinsch (1970). There are also variants of the SVD algorithm with potentially improved quantitative performance based on divide-and-conquer strategies (Gu and Eisenstat 1995), but these can occasionally be observed to fail to converge in practical computations.

For each fixed relative precision, the work required for the SVD of an $r \times r$-matrix scales as $O\left(r^{3}\right)$, and thus the above procedure for the SVD of low-rank matrices requires

$$
O\left(r^{3}+(m+n) r^{2}\right)
$$

operations in total, which is of the same order as the costs of orthogonalization, despite the different type of numerical algorithms.

### 2.3. Matricizations of tensors and product notation

We now introduce basic notions that play a role in the definitions of the different tensor formats that we consider below. We consider these for tensors of order $d$ with countable index sets, that is, for tensors in $\ell_{2}\left(\mathbb{N}^{d}\right)$. Considering tensors in infinite-dimensional spaces is essential for discretization-independent notions of low-rank approximability of functions and for formulating algorithms that are robust under discretization refinement. Note that the corresponding definitions contain tensors of finite size as special cases with support on finite subsets of $\mathbb{N}^{d}$, and practical computations are carried out on such finite subsets as described above for the matrix case.

In what follows, for the set of all modes of a tensor of order $d$, we use the abbreviation

$$
\alpha^{*}=\{1, \ldots, d\}
$$

For a multi-index $v \in \mathbb{N}^{d}$ and $\alpha \subseteq \alpha^{*}$, we write $v_{\alpha}=\left(v_{i}\right)_{i \in \alpha}$ for its sub-index for the modes in $\alpha$ and set

$$
\alpha^{\mathrm{c}}=\alpha^{*} \backslash \alpha
$$

Definition 2.13. For non-empty $\alpha \subseteq \alpha^{*}$, the $\alpha$-matricization of $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ (also referred to as unfolding or flattening of $\mathbf{v}$ ) is given by

$$
\begin{equation*}
\operatorname{mat}_{\alpha}(\mathbf{v})=(\mathbf{v}[v])_{v_{\alpha} \in \mathbb{N}^{\# \alpha}, v_{\alpha} \in \mathbb{N}^{d-\# \alpha}} \in \ell_{2}\left(\mathbb{N}^{\# \alpha}\right) \otimes \ell_{2}\left(\mathbb{N}^{d-\# \alpha}\right) \tag{2.10}
\end{equation*}
$$

For matricizations with respect to single coordinates $i=1, \ldots, d$, we also write mat $_{i}(\mathbf{v})$ in place of $\operatorname{mat}_{\{i\}}(\mathbf{v})$. The associated $\alpha$-rank of $\mathbf{v}$ is defined by

$$
\operatorname{rank}_{\alpha}(\mathbf{v})=\operatorname{rank}_{\operatorname{mat}}^{\alpha}(\mathbf{v}) \in \mathbb{N}_{0} \cup\{\infty\}
$$

with the abbreviation $\operatorname{rank}_{i}(\mathbf{v})=\operatorname{rank}_{\{i\}}(\mathbf{v})$ for $i=1, \ldots, d$.
For $\alpha \subseteq \alpha^{*}$, the bijective linear mapping mat ${ }_{\alpha}$ provides an identification of tensors in $\ell_{2}\left(\mathbb{N}^{d}\right)$ with infinite matrices, where indices of modes in $\alpha$ are used as row indices and the remaining ones as column indices. As in (2.6), these infinite matrices $\operatorname{mat}_{\alpha}(\mathbf{v})$ are identified with linear operators $\operatorname{mat}_{\alpha}(\mathbf{v}): \ell_{2}\left(\mathbb{N}^{d-\# \alpha}\right) \rightarrow$ $\ell_{2}\left(\mathbb{N}^{\# \alpha}\right)$. Note that $\operatorname{mat}_{\alpha}(\mathbf{v})^{\top}=\operatorname{mat}_{\alpha^{c}}(\mathbf{v})$ and thus

$$
\operatorname{rank}_{\alpha^{\mathrm{c}}}(\mathbf{v})=\operatorname{rank}_{\operatorname{mat}}^{\alpha^{\mathrm{c}}}(\mathbf{v})=\operatorname{rank}_{\operatorname{mat}}^{\alpha}(\mathbf{v})^{\top}=\operatorname{rank}_{\operatorname{mat}}^{\alpha}(\mathbf{v})=\operatorname{rank}_{\alpha}(\mathbf{v}) .
$$

Remark 2.14. For general separable Hilbert spaces $H_{i}, i=1, \ldots, d$, an analogous notion of matricizations of elements of $H=\bigotimes_{i=1}^{d} H_{i}$ can be introduced; the $\alpha$-matricization is then an element of $\bigotimes_{i \in \alpha} H_{i} \otimes \bigotimes_{i \in \alpha^{\mathrm{c}}} H_{i}$, identified with an operator from $\bigotimes_{i \in \alpha} H_{i}$ to $\bigotimes_{i \in \alpha^{\mathrm{c}}} H_{i}$. The corresponding $\alpha$-ranks are invariant under tensor space isomorphisms. In view of Remark 2.6 , by choosing a suitable orthonormal basis for each space $H_{i}$, it thus suffices to consider tensors in $\ell_{2}\left(\mathbb{N}^{d}\right)=\bigotimes_{i=1}^{d} \ell_{2}(\mathbb{N})$.
Definition 2.15. Let $\mathbf{t} \in \ell_{2}\left(\mathbb{N}^{d}\right)$, let $\mathbf{U}: \ell_{2}(\mathbb{N}) \rightarrow \ell_{2}(\mathbb{N})$ be linear and bounded, and let $i \in\{1, \ldots, d\}$. The ( $i$ th) mode product of $\mathbf{t}$ and $\mathbf{U}$ is defined by

$$
\mathbf{t} \times_{i} \mathbf{U}=\operatorname{mat}_{i}^{-1}\left(\mathbf{U} \operatorname{mat}_{i}(\mathbf{t})\right)
$$

or entrywise, with the identification (2.6),

$$
\left(\mathbf{t} \times_{i} \mathbf{U}\right)\left[v_{1}, \ldots, v_{d}\right]=\sum_{v_{i}^{\prime} \in \mathbb{N}} \mathbf{U}\left[v_{i}, v_{i}^{\prime}\right] \mathbf{t}\left[v_{1}, \ldots, v_{i-1}, v_{i}^{\prime}, v_{i+1}, \ldots, v_{d}\right]
$$

for all $v_{1}, \ldots, v_{d} \in \mathbb{N}$. In addition, for non-empty $\alpha \subset \alpha^{*}$ and $\mathbf{W}: \ell_{2}\left(\mathbb{N}^{\# \alpha}\right) \rightarrow$ $\ell_{2}\left(\mathbb{N}^{\# \alpha}\right)$, we introduce the $\alpha$-mode product

$$
\mathbf{t} \times_{\alpha} \mathbf{W}=\operatorname{mat}_{\alpha}^{-1}\left(\mathbf{W} \operatorname{mat}_{\alpha}(\mathbf{t})\right)
$$

acting on the modes in $\alpha$.

### 2.4. Tucker format

We have already briefly considered the Tucker tensor format in Section 1.4. In the general definition of the multilinear rank of a tensor in $\ell_{2}\left(\mathbb{N}^{d}\right)$, for $i=1, \ldots, d$, we rely on mode-wise matricizations

$$
\begin{equation*}
\operatorname{mat}_{i}(\mathbf{v})=(\mathbf{v}[v])_{v_{i} \in \mathbb{N},\left(v_{1}, \ldots, v_{i-1}, v_{i+1}, \ldots, v_{d}\right) \in \mathbb{N}^{d-1}} \tag{2.11}
\end{equation*}
$$

as special cases of Definition 2.13.
Definition 2.16. The multilinear rank of a tensor $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ is the tuple

$$
\operatorname{rank}_{\mathrm{ML}}(\mathbf{v})=\left(\operatorname{rank}_{i}(\mathbf{v})\right)_{i=1, \ldots, d} \in\left(\mathbb{N}_{0} \cup\{\infty\}\right)^{d}
$$

For each $i \in\{1, \ldots, d\}$, $\operatorname{since}^{\operatorname{rank}_{i}(\mathbf{v})=\operatorname{dim} \operatorname{range}^{\operatorname{mat}}(\mathbf{v}) \text {, we can choose an }}$ orthonormal system $\left\{\mathbf{U}_{k}^{(i)}\right\}_{k=1, \ldots, \operatorname{rank}_{i}(\mathbf{v})}$ in $\ell_{2}(\mathbb{N})$ that is an orthonormal basis of $\overline{\text { range } \operatorname{mat}_{i}(\mathbf{v})}$, referred to as orthonormal mode frame. Let $\mathbf{U}^{(i)}$ be the matrix with columns $\mathbf{U}_{k}^{(i)}$. In terms of the mode product, the representation in Tucker format of $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ generalizing (1.11) can then be written in the form

$$
\begin{equation*}
\mathbf{v}=\mathbf{a} \times_{1} \mathbf{U}^{(1)} \times_{2} \mathbf{U}^{(2)} \times_{3} \cdots \times_{d} \mathbf{U}^{(d)} \tag{2.12}
\end{equation*}
$$

where the core tensor $\mathbf{a}$ of size $\operatorname{rank}_{1}(\mathbf{v}) \times \cdots \times \operatorname{rank}_{d}(\mathbf{v})$ satisfies

$$
\mathbf{a}\left[k_{1}, \ldots, k_{d}\right]=\left\langle\mathbf{v}, \mathbf{U}_{k_{1}}^{(1)} \otimes \cdots \otimes \mathbf{U}_{k_{d}}^{(d)}\right\rangle
$$

Note that since an isometric isomorphism on $\ell_{2}$ is applied in each mode, we have

$$
\begin{equation*}
\|\mathbf{v}\|_{\ell_{2}}=\|\mathbf{a}\|_{\ell_{2}} . \tag{2.13}
\end{equation*}
$$

Remark 2.17 (storage costs). Let $\operatorname{rank}_{\mathrm{ML}}(\mathbf{v})=\left(r_{i}\right)_{i=1, \ldots, d}$ with $r_{i} \in \mathbb{N}, i=$ $1, \ldots, d$, and $\operatorname{supp} \mathbf{v} \subseteq\left\{1, \ldots, n_{1}\right\} \times \cdots \times\left\{1, \ldots, n_{d}\right\}$. Then the number of storage locations used for the corresponding representation of $\mathbf{v}$ in Tucker format is

$$
\prod_{i=1}^{d} r_{i}+\sum_{i=1}^{d} r_{i} n_{i} .
$$

The multilinear rank refers to the minimum number of basis vectors required in each mode frame. Tensors can also be given in the form (2.12) in a potentially redundant representation with non-orthogonal mode frames. In this case we refer to the corresponding rank parameters (that is, the mode sizes of the core tensor) as representation ranks.

Remark 2.18 (basic operations). For the operations considered in Section 2.2 for low-rank matrices, we have the following analogous results for the Tucker format.
(i) Addition of tensors in Tucker format can be performed analogously to the matrix case described in Section 2.2.2 by concatenation of mode frames and diagonal stacking of core tensors, so that the multilinear rank of the sum is at most the componentwise sum of multilinear ranks.
(ii) Orthogonalization. Given $\tilde{\mathbf{a}} \times_{1} \tilde{\mathbf{U}}^{(1)} \cdots \times_{d} \tilde{\mathbf{U}}^{(d)}$, where $\tilde{\mathbf{U}}^{(1)}, \ldots, \tilde{\mathbf{U}}^{(d)}$ have non-orthogonal columns, one can orthonormalize in a similar way to the matrix case in Section 2.2.1: perform a $Q R$ decomposition $\tilde{\mathbf{U}}^{(1)}=\mathbf{U}^{(1)} \mathbf{R}_{1}$ and replace the core tensor by $\tilde{\mathbf{a}} \times_{1} \mathbf{R}_{1}$; carry out the same steps for $i=2, \ldots, d$. The total computational costs scale as

$$
O\left(\sum_{i=1}^{d} r_{i}^{2} \prod_{s \neq i} r_{s}+\sum_{i=1}^{d} r_{i}^{2} n_{i}\right)
$$

(iii) Higher-order singular value decomposition. In addition, the mode frames can be chosen to be left singular vectors of the respective matricizations. The resulting representation has been termed higher-order singular value decomposition (HOSVD) in De Lathauwer, De Moor and Vandewalle (2000); its properties, which to some extent parallel those of the SVD of matrices, are considered in Section 2.7. To compute the HOSVD numerically, assume that a representation $\mathbf{v}=\mathbf{a} \times{ }_{1} \mathbf{U}^{(1)} \cdots \times_{d} \mathbf{U}^{(d)}$ with orthonormal mode frames is given. Perform a $Q R$ factorization $\operatorname{mat}_{1}(\mathbf{a})^{\top}=\mathbf{V}_{1} \mathbf{R}_{1}$ with $\mathbf{V}_{1}$ having orthonormal columns and $\mathbf{R}_{1} \in \mathbb{R}^{r_{1} \times r_{1}}$, and compute the $\operatorname{SVD} \mathbf{R}_{1}^{\top}=\mathbf{Q} \Sigma \mathbf{P}^{\top}$. By Proposition 2.11, the columns of $\hat{\mathbf{U}}^{(1)}=\mathbf{U}^{(1)} \mathbf{Q}$ are then left singular vectors of $\operatorname{mat}_{i}(\mathbf{v})$. After replacing the core tensor by $\mathbf{a} \times \mathbf{Q}^{\top}$, the same can be carried out for the further tensor modes, which eventually yields a decomposition $\hat{\mathbf{a}} \times_{1} \hat{\mathbf{U}}^{(1)} \cdots \times_{d} \hat{\mathbf{U}}^{(d)}$, where the columns of each $\hat{\mathbf{U}}^{(i)}$ are left singular vectors of the respective matricization. In view of Remark 2.12, the total number of operations for this procedure scales as

$$
O\left(\sum_{i=1}^{d} r_{i}^{2} \prod_{s \neq i} r_{s}+\sum_{i=1}^{d} r_{i}^{2}\left(n_{i}+r_{i}\right)\right) .
$$

Note that for $d>2$, in contrast to the classical SVD, â is in general a fully populated tensor without diagonal structure.

In summary, the Tucker format can be used to express a given tensor, separately in each mode, with respect to adapted basis vectors given by the mode frames. However, the number of entries of the core tensor still scales exponentially with respect to the number of modes.

### 2.5. Tensor trains or matrix product states

The tensor train (TT) format, introduced in Oseledets (2009b), Oseledets and Tyrtyshnikov (2009a,b) and Oseledets (2011b), is based on a different set of matricizations such that an application to very high tensor orders becomes feasible. In
quantum physics these tensor representations are known as matrix product states. In this context they were introduced in Vidal (2003) for the purpose of approximating wavefunctions; for an overview of corresponding applications, see Schollwöck (2011).

Definition 2.19. The tensor train rank of a tensor $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ is given by the vector

$$
\operatorname{rank}_{\mathrm{TT}}(\mathbf{v})=\left(\operatorname{rank}_{\{1, \ldots, i\}}(\mathbf{v})\right)_{i=1, \ldots, d-1}
$$

We now again deduce from this notion of rank a corresponding tensor representation. To this end, we use that for $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$,

$$
\begin{equation*}
\overline{\text { range } \operatorname{mat}_{\{1, \ldots, i-1, i\}}(\mathbf{v})} \subseteq \overline{\operatorname{range~mat}_{\{1, \ldots, i-1\}}(\mathbf{v})} \otimes \ell_{2}(\mathbb{N}) \tag{2.14}
\end{equation*}
$$

This can be deduced from Hackbusch (2019, Thm 6.31).
To simplify notation, we abbreviate $r_{i}(\mathbf{v})=\operatorname{rank}_{\{1, \ldots, i\}}(\mathbf{v})$. By the definition of $r_{i}(\mathbf{v})$, for $i=1, \ldots, d-1$ there exist orthonormal bases $\left\{\mathbf{U}_{k}^{\{1, \ldots, i\}}\right\}_{k=1, \ldots, r_{i}(\mathbf{v})}$ of $\overline{\text { range } \operatorname{mat}_{\{1, \ldots, i\}}(\mathbf{v})} \subseteq \ell_{2}\left(\mathbb{N}^{i}\right)$. In particular,

$$
\mathbf{v}\left[j_{1}, \ldots, j_{d}\right]=\sum_{k_{d-1}=1}^{r_{d-1}(\mathbf{v})} \mathbf{U}_{k_{d-1}}^{\{1, \ldots, d-1\}}\left[j_{1}, \ldots, j_{d-1}\right] \mathbf{V}_{d}\left[k_{d-1}, j_{d}\right]
$$

with the coefficients

$$
\mathbf{V}_{d}\left[k_{d-1}, j_{d}\right]=\sum_{\left(j_{1}, \ldots, j_{d-1}\right) \in \mathbb{N}^{d-1}} \mathbf{U}_{k_{d-1}}^{\{1, \ldots, d-1\}}\left[j_{1}, \ldots, j_{d-1}\right] \mathbf{v}\left[j_{1}, \ldots, j_{d-1}, j_{d}\right] .
$$

We now proceed recursively: by (2.14), for $1<i<d$ we obtain a third-order coefficient tensor $\mathbf{V}_{i}$ such that

$$
\mathbf{U}_{k}^{\{1, \ldots, i\}}\left[j_{1}, \ldots, j_{i}\right]=\sum_{k_{i-1}=1}^{r_{i-1}(\mathbf{v})} \mathbf{U}_{k_{i-1}}^{\{1, \ldots, i-1\}}\left[j_{1}, \ldots, j_{i-1}\right] \mathbf{V}_{i}\left[k_{i-1}, j_{i}, k\right]
$$

Setting $\mathbf{V}_{1}\left[j_{1}, k\right]=\mathbf{U}_{k}^{\{1\}}\left[j_{1}\right]$, we arrive at the tensor train representation

$$
\begin{equation*}
\mathbf{v}\left[j_{1}, \ldots, j_{d}\right]=\sum_{k_{1}=1}^{r_{1}(\mathbf{v})} \cdots \sum_{k_{d-1}=1}^{r_{d-1}(\mathbf{v})} \mathbf{V}_{1}\left[j_{1}, k_{1}\right] \mathbf{V}_{2}\left[k_{1}, j_{2}, k_{2}\right] \cdots \mathbf{V}_{d}\left[k_{d-1}, j_{d}\right] \tag{2.15}
\end{equation*}
$$

Remark 2.20 (storage costs). Let $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ such that $\operatorname{rank}_{\mathrm{TT}}(\mathbf{v})=\left(r_{i}\right)_{i=1, \ldots, d-1}$ with $r_{i} \in \mathbb{N}, i=1, \ldots, d-1$, and $\operatorname{supp} \mathbf{v} \subseteq\left\{1, \ldots, n_{1}\right\} \times \cdots \times\left\{1, \ldots, n_{d}\right\}$, with $n_{1}, \ldots, n_{d} \in \mathbb{N}$. Then the number of entries to be stored for the representation (2.15) is

$$
\sum_{i=1}^{d} r_{i-1} n_{i} r_{i}
$$

where we set $r_{0}=r_{d}=1$.

Although the above bound has only a linear explicit dependence on $d$, note that the scaling with respect to $d$ also depends on the ranks, which can in principle still scale exponentially with respect to $d$. Suitable rank bounds are thus crucial for actually avoiding the curse of dimensionality.

As a shorthand notation for explicitly specifying tensor train representations, individual cores $\mathbf{V} \in \mathbb{R}^{r_{1} \times n \times r_{2}}$ with $\mathbf{V}\left[k_{1}, j, k_{2}\right]=\mathbf{a}_{k_{1}, k_{2}}[j]$ for certain vectors $\mathbf{a}_{k_{1}, k_{2}}$ can be written in the block form

$$
\mathbf{V}=\left[\begin{array}{ccc}
\mathbf{a}_{1,1} & \cdots & \mathbf{a}_{1, r_{2}}  \tag{2.16}\\
\vdots & \ddots & \vdots \\
\mathbf{a}_{r_{1}, 1} & \cdots & \mathbf{a}_{r_{1}, r_{2}}
\end{array}\right] .
$$

This can be combined with the following special notion of product, introduced by De Launey and Seberry (1994), to obtain a compact notation.

Definition 2.21 (strong Kronecker product). For $\mathbf{V}_{1} \in \mathbb{R}^{r_{1} \times n_{1} \times r_{2}}$ and $\mathbf{V}_{2} \in$ $\mathbb{R}^{r_{2} \times n_{2} \times r_{3}}$, we define

$$
\left(\mathbf{V}_{1} \bowtie \mathbf{V}_{2}\right)\left[k_{1}, j_{1}, j_{2}, k_{3}\right]=\sum_{k_{2}=1}^{r_{2}} \mathbf{V}_{1}\left[k_{1}, j_{1}, k_{2}\right] \mathbf{V}_{2}\left[k_{2}, j_{2}, k_{3}\right]
$$

This product thus acts like standard matrix multiplication on the block form (2.16), where individual entries are combined by the tensor product. As a specific example,

$$
\left[\begin{array}{ll}
\mathbf{a}_{11} & \mathbf{a}_{12} \\
\mathbf{a}_{21} & \mathbf{a}_{22}
\end{array}\right] \bowtie\left[\begin{array}{ll}
\mathbf{b}_{11} & \mathbf{b}_{12} \\
\mathbf{b}_{21} & \mathbf{b}_{22}
\end{array}\right]=\left[\begin{array}{ll}
\mathbf{a}_{11} \otimes \mathbf{b}_{11}+\mathbf{a}_{12} \otimes \mathbf{b}_{21} & \mathbf{a}_{11} \otimes \mathbf{b}_{12}+\mathbf{a}_{12} \otimes \mathbf{b}_{22} \\
\mathbf{a}_{21} \otimes \mathbf{b}_{11}+\mathbf{a}_{22} \otimes \mathbf{b}_{21} & \mathbf{a}_{21} \otimes \mathbf{b}_{12}+\mathbf{a}_{22} \otimes \mathbf{b}_{22}
\end{array}\right],
$$

where the product of two rank-two blocks of tensor order one leads to a new rank-two block of order two. With this notation, we write (2.15) concisely as

$$
\begin{equation*}
\mathbf{v}=\mathbf{V}_{1} \bowtie \mathbf{V}_{2} \bowtie \cdots \bowtie \mathbf{V}_{K} \tag{2.17}
\end{equation*}
$$

### 2.6. Hierarchical tensors and tree tensor networks

We now consider hierarchical tensors as introduced in the context of numerical analysis by Hackbusch and Kühn (2009); see also Hackbusch (2019) and Falcó, Hackbusch and Nouy (2021). In the physics literature, these are a special case of tree tensor networks; see Shi, Duan and Vidal (2006) and Murg, Verstraete, Legeza and Noack (2010).
Definition 2.22. A set $\mathbb{T} \subset 2^{\{1, \ldots, d\}}$ is called a binary dimension tree (for dimension $d$ ) if the following conditions hold.
(i) $\{1, \ldots, d\} \in \mathbb{T}$ and $\emptyset \notin \mathbb{T}$.
(ii) For each $\alpha, \beta \in \mathbb{T}$ with $\alpha \neq \beta$, either $\alpha \subset \beta$ or $\beta \subset \alpha$ or $\alpha \cap \beta=\emptyset$.
(iii) For each $\alpha \in \mathbb{T}$ with $\# \alpha>1$, there exists precisely one subset $\left\{\beta_{1}, \beta_{2}\right\} \subseteq \mathbb{T}$ with $\beta_{1} \cap \beta_{2}=\emptyset$ such that $\alpha=\beta_{1} \cup \beta_{2}$.

As a consequence of this definition, if $\mathbb{T}$ is a binary dimension tree, then for all $i \in\{1, \ldots, d\}$ we have $\{i\} \in \mathbb{T}$; these singleton elements are referred to as leaves of the tree.

As in Bachmayr and Schneider (2017), we introduce a set of effective edges $\mathbb{E}$ of $\mathbb{T}$ as the pairs

$$
\mathbb{E}=\left\{\left\{\alpha, \alpha^{\mathrm{c}}\right\}: \alpha \in \mathbb{T} \backslash\left\{\alpha^{*}\right\}\right\}
$$

for which we have $\# \mathbb{E}=2 d-3$. For each $e \in \mathbb{E}$, we define the representer [ $e$ ] as the $\alpha \in e$ such that $\alpha \in \mathbb{T}$; if this element is not unique, we make an arbitrary choice of $[e]$ among the elements of $e$.

Definition 2.23. Let $\mathbb{T}$ be a dimension tree with associated effective edges $\mathbb{E}$. The hierarchical rank of a tensor $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ is given by

$$
\operatorname{rank}_{\mathbb{E}}(\mathbf{v})=\left(\operatorname{rank}_{[e]}(\mathbf{v})\right)_{e \in \mathbb{E}} .
$$

Remark 2.24. While for many purposes the elements of the dimension tree $\mathbb{T}$ are the decisive quantities, the ranks are indexed by the effective edges $\mathbb{E}$. These correspond to the different matricizations associated to $\mathbb{T}$, where those that can be obtained from each other by transposition are treated as equal (since these necessarily have the same rank). In general, there are several different $\mathbb{T}$ that lead to the same set of matricizations $\mathbb{E}$. For example, for each of the five different $\mathbb{T}$ for $d=4$ shown in Figure 2.1, we have the same effective edges

$$
\begin{aligned}
\mathbb{E}= & \{\{\{1\},\{2,3,4\}\},\{\{2\},\{1,3,4\}\}, \\
& \{\{3\},\{1,2,4\}\},\{\{4\},\{1,2,3\}\},\{\{1,2\},\{3,4\}\}\} .
\end{aligned}
$$

These $\mathbb{T}$ are thus in principle equivalent in that they lead to tensor representations with the same ranks; the difference lies in the placement of the root element $\{1,2,3,4\}$ in the tree. Intuitively, when looking only at the connectivity of the tree between modes (rather than its labels), the different $\mathbb{T}$ arise by placing the root element on different effective edges.

Similarly to the tensor train format, the notion of hierarchical rank leads to a representation of high-order tensors in terms of lower-order tensors, which can be obtained as follows. Let $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ be given. For each $\alpha \in \mathbb{T} \backslash\left\{\alpha^{*}\right\}$, there exists an orthonormal system $\left\{\mathbf{U}_{k}^{\alpha}\right\}_{k=1, \ldots, \operatorname{rank}_{\alpha}(\mathbf{v})}$ in $\ell_{2}\left(\mathbb{N}^{\# \alpha}\right)$ that is an orthonormal basis of range $\operatorname{mat}_{\alpha}(\mathbf{v})$. Such an orthonormal system is again called an orthonormal $\left(\alpha\right.$-)mode frame. In what follows, for each $\alpha$, we write $\mathbf{U}^{\alpha}$ for the (potentially infinite) matrix that has columns $\mathbf{U}_{k}^{\alpha}$ for $k=1, \ldots, \operatorname{rank}_{\alpha}(\mathbf{v})$.

In a first step, for $\mathbf{v}$ as above, we have the decomposition

$\{3\} \quad\{4\}$



$\{3\} \quad\{4\}$

$\{1,2,3,4\}$

\{1,2
$\{1\} \quad\{2\}$

Figure 2.1. Examples of dimension trees $\mathbb{T}$ with $d=4$ that lead to the same effective edges $\mathbb{E}$.
where $\alpha_{1}^{*}, \alpha_{2}^{*}$ are the children of the root element $\alpha^{*}$ and where

$$
\mathbf{b}^{\alpha^{*}}\left[\ell_{1}, \ell_{2}\right]=\left\langle\mathbf{v}, \mathbf{U}_{\ell_{1}}^{\alpha_{1}^{*}} \otimes \mathbf{U}_{\ell_{2}}^{\alpha_{2}^{*}}\right\rangle ;
$$

note that in this particular case $\operatorname{rank}_{\alpha_{1}^{*}}(\mathbf{v})=\operatorname{rank}_{\alpha_{2}^{*}}(\mathbf{v})$, since the children of the root share (up to transposition) the same matricization. For $\alpha \in \mathbb{T} \backslash\left\{\alpha^{*}\right\}$ with $\# \alpha>1$, for the children $\alpha_{1}, \alpha_{2} \in \mathbb{T}$ of $\alpha$, we have

$$
\begin{equation*}
\overline{\text { range } \operatorname{mat}_{\alpha}(\mathbf{v})} \subseteq \overline{\text { range mat } \alpha_{\alpha_{1}}(\mathbf{v})} \otimes \overline{\text { range mat } \alpha_{\alpha_{2}}(\mathbf{v})}, \tag{2.19}
\end{equation*}
$$

which can be shown analogously to (2.14). As a consequence

$$
\begin{equation*}
\mathbf{U}_{k}^{\alpha}=\sum_{\ell_{1}=1}^{\operatorname{rank}_{\alpha_{1}}(\mathbf{v})} \sum_{\ell_{2}=1}^{\mathrm{rank}_{\alpha_{2}}(\mathbf{v})} \mathbf{b}^{\alpha}\left[k, \ell_{1}, \ell_{2}\right] \mathbf{U}_{\ell_{1}}^{\alpha_{1}} \otimes \mathbf{U}_{\ell_{2}}^{\alpha_{2}}, \quad k=1, \ldots, \operatorname{rank}_{\alpha}(\mathbf{v}) \tag{2.20}
\end{equation*}
$$

with the transfer tensors $\mathbf{b}^{\alpha}$ of size $\operatorname{rank}_{\alpha}(\mathbf{v}) \times \operatorname{rank}_{\alpha_{1}}(\mathbf{v}) \times \operatorname{rank}_{\alpha_{2}}(\mathbf{v})$ given by

$$
\mathbf{b}^{\alpha}\left[k, \ell_{1}, \ell_{2}\right]=\left\langle\mathbf{U}_{k}^{\alpha}, \mathbf{U}_{\ell_{1}}^{\alpha_{1}} \otimes \mathbf{U}_{\ell_{2}}^{\alpha_{2}}\right\rangle
$$

The nestedness property (2.19) implies the restriction

$$
\operatorname{rank}_{\alpha}(\mathbf{v}) \leq \operatorname{rank}_{\alpha_{1}}(\mathbf{v}) \operatorname{rank}_{\alpha_{2}}(\mathbf{v})
$$

We let $\operatorname{Ranks}(\mathbb{E})$ denote the set of feasible rank vectors in $\left(\mathbb{N}_{0} \cup\{\infty\}\right)^{\mathbb{E}}$ of hierarchical tensors.

In summary, applying (2.20) recursively, from (2.18) we obtain a representation of $\mathbf{v}$ by the transfer tensors $\mathbf{b}^{\alpha}$ for all $\alpha \in \mathbb{T}$ with $\# \alpha>1$ and the mode frames $\left\{\mathbf{U}_{k}^{\{i\}}\right\}_{k=1, \ldots, \operatorname{rank}_{i}(\mathbf{v})}$ for $i=1, \ldots, d$. All operations on hierarchical tensors are then performed exclusively on these representation components.

Remark 2.25. One can also specify tensors in this form with non-orthogonal mode frames in a potentially redundant representation. In this case we refer to the mode sizes of the transfer tensors (which do not necessarily reflect the true hierarchical ranks) as representation ranks.

For $d=4$, in the case of the balanced binary tree

$$
\mathbb{T}=\{\{1,2,3,4\},\{1,2\},\{3,4\},\{1\},\{2\},\{3\},\{4\}\}
$$

the decomposition can be written as

$$
\begin{equation*}
\mathbf{v}=\sum_{k_{1}=1}^{\mathrm{rank}_{1}(\mathbf{v})} \cdots \sum_{k_{4}=1}^{\mathrm{rank}_{4}(\mathbf{v})} \mathbf{a}\left[k_{1}, \ldots, k_{4}\right] \mathbf{U}_{k_{1}}^{\{1\}} \otimes \cdots \otimes \mathbf{U}_{k_{4}}^{\{4\}} \tag{2.21}
\end{equation*}
$$

where, with $r_{12}=\operatorname{rank}_{\{1,2\}}(\mathbf{v}), r_{34}=\operatorname{rank}_{\{3,4\}}(\mathbf{v})$,

$$
\mathbf{a}\left[k_{1}, \ldots, k_{4}\right]=\sum_{\ell_{1}=1}^{r_{12}} \sum_{\ell_{2}=1}^{r_{34}} \mathbf{b}^{\{1,2,3,4\}}\left[\ell_{1}, \ell_{2}\right] \mathbf{b}^{\{1,2\}}\left[\ell_{1}, k_{1}, k_{2}\right] \mathbf{b}^{\{3,4\}}\left[\ell_{2}, k_{3}, k_{4}\right] .
$$

Since it can be interpreted as a direct generalization of the Tucker format - with additional constraints on ranks of matricizations and a resulting further decomposition of the core tensor - this tensor format is also referred to in the literature as the hierarchical Tucker format.

Remark 2.26 (storage costs). The number of coefficients that need to be stored for a tensor in hierarchical format of size $n_{1} \times \cdots \times n_{d}$, with the abbreviations $r_{\alpha}=\operatorname{rank}_{\alpha}(\mathbf{v})$ and $r_{i}=\operatorname{rank}_{i}(\mathbf{v})$, is

$$
\sum_{\substack{\alpha \in \mathbb{T} \\ \# \alpha>1}} r_{\alpha} r_{\alpha_{1}} r_{\alpha_{2}}+\sum_{i=1}^{d} r_{i} n_{i}
$$

where $\alpha_{1}, \alpha_{2}$ denote the children of each $\alpha$.
Remark 2.27 (relation to the tensor train format). The tensor train format has the effective edges

$$
\begin{equation*}
\mathbb{E}=\{\{\{1, \ldots, i\},\{i+1, \ldots, d\}\}: i=1, \ldots, d-1\}, \tag{2.22}
\end{equation*}
$$

which do not directly correspond to a dimension tree according to Definition 2.22, due to missing leaf elements. From a practical point of view, this means that pairs of transfer tensors and mode frames that are kept separate in the hierarchical format are contracted to a single component in the tensor train format. When keeping this separation, the linear dimension tree

$$
\{\{1, \ldots, d\},\{1\},\{2, \ldots, d\},\{2\},\{3, \ldots, d\}, \ldots,\{d-1, d\},\{d-1\},\{d\}\}
$$

yields a hierarchical tensor representation that is analogous to the tensor train format. By constraining the mode frames to be identity mappings, the tensor train
format can alternatively be seen as a special case of this representation (Uschmajew and Vandereycken 2013). For a comparison of the respective notions of ranks, see Grasedyck and Hackbusch (2011).

Remark 2.28 (orthogonalization). A representation of a finitely supported $\mathbf{v} \in$ $\ell_{2}\left(\mathbb{N}^{d}\right)$ in hierarchical format by non-orthonormal mode frames $\tilde{\mathbf{U}}^{\{i\}}$ for $i=1, \ldots, d$ and transfer tensors $\tilde{\mathbf{b}}^{\alpha}$ for $\alpha \subseteq \alpha^{*}$ with $\# \alpha>1$ can be transformed such that all $\alpha$ mode frames are orthonormal; see Grasedyck (2009/10) and Hackbusch and Kühn (2009). First, for each leaf $\{i\}$ with parent $\beta_{i} \in \mathbb{T}$, perform a QR decomposition $\tilde{\mathbf{U}}^{\{i\}}=\mathbf{U}^{\{i\}} \mathbf{R}_{i}$. With $\tilde{\mathbf{U}}^{\{i\}}$ attached to mode $m$ of the third-order tensor $\tilde{\mathbf{b}}^{\beta_{i}}$, replace $\tilde{\mathbf{b}}^{\beta_{i}}$ by $\tilde{\mathbf{b}}^{\beta_{i}} \times_{m} \mathbf{R}_{i}$. For $\alpha \in \mathbb{T} \backslash\left\{\alpha^{*}\right\}$ with parent $\beta$ and children $\alpha_{1}, \alpha_{2}$ such that $\mathbf{U}^{\alpha_{1}}$ and $\mathbf{U}^{\alpha_{2}}$ are already orthonormalized, perform a $Q R$ factorization $\operatorname{mat}_{\{2,3\}}\left(\tilde{\mathbf{b}}^{\alpha}\right)=\mathbf{C}_{\alpha} \mathbf{R}_{\alpha}$; with $\tilde{\mathbf{b}}^{\alpha}$ attached to mode $m$ of $\tilde{\mathbf{b}}^{\beta}$, replace $\tilde{\mathbf{b}}^{\beta}$ by $\tilde{\mathbf{b}}^{\beta} \times{ }_{m} \mathbf{R}_{\alpha}$, and set $\mathbf{b}^{\alpha}=\operatorname{mat}_{\{2,3\}}^{-1}\left(\mathbf{C}_{\alpha}\right)$.

After performing this recursively for all elements of $\mathbb{T} \backslash\left\{\alpha^{*}\right\}$, the computed $\mathbf{b}^{\alpha}$ and $\mathbf{U}^{\{i\}}$ yield a hierarchical tensor representation of $\mathbf{v}$ such that all associated mode frames $\mathbf{U}^{\alpha}$ for $\alpha \in \mathbb{T} \backslash\left\{\alpha^{*}\right\}$ are orthonormal, that is, $\left(\mathbf{U}^{\alpha}\right)^{\top} \mathbf{U}^{\alpha}=\mathbf{I}$. To assess the total computational costs of this procedure, we again assume finite mode sizes $n_{1}, \ldots, n_{d}$. We assume the input representation to be given with matricization ranks $\left(r_{\alpha}\right)_{\alpha \in \mathbb{T}}$. This means that for $\alpha \in \mathbb{T}$ with children $\alpha_{1}, \alpha_{2}$, the transfer tensor $\tilde{\mathbf{b}}^{\alpha}$ has size $r_{\alpha} \times r_{\alpha_{1}} \times r_{\alpha_{2}}$. In addition, we set $r_{\alpha^{*}}=1$. The number of operations then scales as

$$
\begin{equation*}
O\left(\sum_{\substack{\alpha \in \mathbb{T} \\ \# \alpha>1}}\left(r_{\alpha}^{2} r_{\alpha_{1}} r_{\alpha_{2}}+r_{\alpha} r_{\alpha_{1}}^{2} r_{\alpha_{2}}+r_{\alpha} r_{\alpha_{1}} r_{\alpha_{2}}^{2}\right)+\sum_{i=1}^{d} r_{\{i\}}^{2} n_{i}\right) \tag{2.23}
\end{equation*}
$$

Here, orthogonalization is performed with respect to the root element $\alpha^{*}$. Hence, as a consequence of (2.18), we have

$$
\|\mathbf{v}\|_{\ell_{2}}=\left\|\mathbf{b}^{\alpha^{*}}\right\|_{\ell_{2}}
$$

We can also orthogonalize towards different edges in the dimension tree, which amounts to switching between alternative dimension trees that have the same effective edges $\mathbb{E}$ (as illustrated in Figure 2.1); see Bachmayr and Schneider (2017) for details. Tensor train representations can be orthogonalized in a similar fashion, where we distinguish between left and right orthogonality; see Oseledets (2011b) and Holtz, Rohwedder and Schneider (2012b).

Addition of tensors in hierarchical or tensor train format can be done by concatenation of representations similarly to the case of low-rank matrices as in Section 2.2.2 and Tucker tensors in Remark 2.18(i). The hierarchical or tensor train ranks of such sums are then again bounded componentwise by the sums of the ranks of summands.

Proposition 2.29. With $\mathbf{v}_{k}^{(i)} \in \ell_{2}(\mathbb{N})$ for $k=1, \ldots, r$ and $i=1, \ldots, d$, let

$$
\mathbf{v}=\sum_{k=1}^{r} \mathbf{v}_{k}^{(1)} \otimes \cdots \otimes \mathbf{v}_{k}^{(d)}
$$

Then we have the componentwise estimate $\operatorname{rank}_{\mathbb{E}}(\mathbf{v}) \leq r$.
Proof. Each elementary tensor product $\mathbf{v}_{k}=\mathbf{v}_{k}^{(1)} \otimes \cdots \otimes \mathbf{v}_{k}^{(d)}$ can be represented in hierarchical format with $\operatorname{rank}_{\mathbb{E}}\left(\mathbf{v}_{k}\right) \equiv 1$ for any $\mathbb{T}$, where all transfer tensors have size $1 \times 1 \times 1$. The statement thus follows by addition of these representations.

Remark 2.30. Tree tensor networks can be considered based on trees of more general arities than the binary ones used Definition 2.22. These arities, however, enter exponentially into the storage costs (with the Tucker format as an extreme case), and thus binary trees are typically the most favourable choice.

In certain application scenarios, it can also be of interest to consider tensor networks without tree structure, where components are connected according to a more general graph with cycles. A simple example is given by tensor chains (which are also known as cyclic MPS in physics) with three components in the form

$$
\mathbf{t}\left[j_{1}, j_{2}, j_{3}\right]=\sum_{k_{1}=1}^{r_{1}} \sum_{k_{2}=1}^{r_{2}} \sum_{k_{3}=1}^{r_{3}} \mathbf{U}\left[k_{1}, j_{1}, k_{2}\right] \mathbf{V}\left[k_{2}, j_{2}, k_{3}\right] \mathbf{W}\left[k_{3}, j_{3}, k_{1}\right] .
$$

However, tensor networks with cycles lead to computational issues similar to those described for the canonical format in Section 1.4 (Landsberg, Qi and Ye 2012), and the algorithmic framework for subspace formats based on matricizations is not applicable to such networks. For a detailed discussion, see Hackbusch (2019, Sec. 12.4) and Hackbusch (2014, Sec. 9.4).

### 2.7. Hierarchical singular value decomposition and low-rank approximations

For tensor order $d=2$, by Theorem 2.8, low-rank best approximations in $\ell_{2}\left(\mathbb{N}^{2}\right)$ (and thus in any separable tensor product Hilbert space) can be obtained by truncation of the SVD. As noted in Remark 2.12, for finite matrices, these best approximations can be computed by a polynomial-time algorithm that is actually not much more costly than other standard matrix factorizations.

As we saw in Section 1.4, in the case of the canonical tensor format (1.9), lowrank best approximations generally do not exist. This is different in subspace tensor formats, where we have the following result.
Theorem 2.31. For any $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$, any dimension tree $\mathbb{T}$ and any $r=\left(r_{e}\right)_{e \in \mathbb{E}} \in$ $(\mathbb{N} \cup\{\infty\})^{\mathbb{E}}$, there exists a best approximation with respect to the $\ell_{2}$-norm in $\left\{\mathbf{w} \in \ell_{2}\left(\mathbb{N}^{d}\right): \operatorname{rank}_{\mathbb{E}}(\mathbf{w}) \leq r\right\}$.
Proof. Let $S_{\mathrm{r}}=\left\{\mathbf{w} \in \ell_{2}\left(\mathbb{N}^{d}\right): \operatorname{rank}_{\mathbb{E}}(\mathbf{w}) \leq r\right\}$. Then

$$
S_{\mathrm{r}}=\bigcap_{e \in \mathbb{E}}\left\{\mathbf{w} \in \ell_{2}\left(\mathbb{N}^{d}\right): \operatorname{rank}_{[e]}(\mathbf{w}) \leq r_{e}\right\}
$$

and for each $e$, the set $\left\{\mathbf{w} \in \ell_{2}\left(\mathbb{N}^{d}\right): \operatorname{rank}_{[e]}(\mathbf{w}) \leq r_{e}\right\}$ with constraint on $\operatorname{rank}_{[e]}(\mathbf{w})=\operatorname{rank}^{\operatorname{mat}}{ }_{[e]}(\mathbf{w})$ is weakly sequentially closed. This can be deduced from the lower semicontinuity of the rank of finite-dimensional matrices; see for example Uschmajew (2013, Lem. 6.5) or Hackbusch (2014, Sec. 6.4). Thus $S_{\mathrm{r}}$ is weakly sequentially closed. Since $\ell_{2}\left(\mathbb{N}^{d}\right)$ is reflexive, this allows us to obtain the weak limit of a minimizing sequence in $S_{\mathrm{r}}$, which is a minimizer by weak lower semicontinuity of the norm.

Related results, including the case of the Tucker format, have been obtained in Uschmajew (2010). In the case of general Banach spaces, the existence of best low-rank approximations has been treated by Falcó and Hackbusch (2012) using the notion of minimal subspaces.

In the case of subspace tensor formats for $d>2$ considered in this section, there exist normal forms of tensor representations that share certain features of the SVD of matrices, such as the higher-order SVD discussed in Remark 2.18 in the case of the Tucker format. For general hierarchical tensors, we have an analogous notion developed in Grasedyck (2009/10) and Hackbusch and Kühn (2009); see also Hackbusch (2019, Sec. 11.3.3).

Definition 2.32. A hierarchical tensor representation of $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ with dimension tree $\mathbb{T}$ is a hierarchical singular value decomposition (HSVD) if, for all $\alpha \in \mathbb{T} \backslash\left\{\alpha^{*}\right\}$, the mode frame $\left\{\mathbf{U}_{k}^{\alpha}\right\}_{k=1, \ldots, \operatorname{rank}_{\alpha}(\mathbf{v})}$ is made up of left singular vectors of $\operatorname{mat}_{\alpha}(\mathbf{v})$, ordered such that the corresponding sequence of singular values, denoted by $\sigma^{\alpha}(\mathbf{v})=\left(\sigma_{k}^{\alpha}(\mathbf{v})\right)_{k=1, \ldots, \mathrm{rank}_{\alpha}(\mathbf{v})}$, is non-increasing.

In the case of the tensor train representation, the analogous form in terms of left singular vectors of matricizations is referred to as tensor train SVD (TT-SVD), introduced in Oseledets (2011b). In physics, an equivalent representation is known as the Vidal representation of matrix product states (Vidal 2003).

Remark 2.33 (existence and uniqueness). We can always choose the mode frames for each $\alpha \in \mathbb{T} \backslash\left\{\alpha^{*}\right\}$ as left singular vectors of the $\alpha$-matricization, whose existence is ensured by Theorem 2.8. Thus a hierarchical singular value decomposition exists for every $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$. It is unique precisely when the sequences $\sigma^{\alpha}(\mathbf{v})$ are strictly decreasing for all $\alpha$.

A procedure for numerically transforming a hierarchical tensor representation to HSVD form was introduced by Grasedyck (2009/10). Starting from a representation with orthonormalized mode frames, the matrices transforming these mode frames can be extracted from eigenvalue decompositions of certain recursively assembled Gramians associated to the nodes in $\mathbb{T}$. The computational costs are of the same order (2.23) as the orthogonalization, which for $\mathbf{v}$ of mode sizes $n_{1} \times \cdots \times n_{d}$ is bounded by

$$
\begin{equation*}
O\left(d \max _{\alpha \in \mathbb{T}} \operatorname{rank}_{\alpha}^{4}(\mathbf{v})+\sum_{i=1}^{d} \operatorname{rank}_{i}(\mathbf{v})^{2} n_{i}\right) . \tag{2.24}
\end{equation*}
$$

For the details of this procedure, we refer to Grasedyck (2009/10) and Hackbusch (2019, Sec. 11.4.2). Oseledets (2011b) has introduced a procedure for computing TT-SVD representations of tensor trains following similar lines to the HOSVD in Remark 2.18(iii).

### 2.7.1. Truncated hierarchical SVD

We now consider rank truncation of HSVD representations in analogy to the truncated SVD of matrices in Theorem 2.8.

Let $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ be given in HSVD form with orthonormal mode frames $\mathbf{U}^{\alpha}$ for $\alpha \in \mathbb{T} \backslash\left\{\alpha^{*}\right\}$. For each $\alpha$ and $r$, we write $\mathbf{U}^{\alpha, r}$ for $\mathbf{U}^{\alpha}$ truncated to its first $r$ columns. For each $e \in \mathbb{E}$, we define the orthogonal projector

$$
\mathrm{P}_{\mathbf{v}, r}^{e} \mathbf{v}=\mathbf{v} \times_{[e]} \mathbf{U}^{[e], r}\left(\mathbf{U}^{[e], r}\right)^{\top}
$$

We now introduce a notion of level of nodes in $\mathbb{T}$ as the distance to the root element $\alpha^{*}$, so that $\alpha^{*}$ has level zero and the children of $\alpha^{*}$ have level one. We next choose an ordering $e(1), e(2), \ldots, e(\# \mathbb{E})$ of $\mathbb{E}$ such that $[e(1)],[e(2)], \ldots$ have non-decreasing level. With this ordering, for $r=\left(r_{e}\right)_{e \in \mathbb{E}} \in \operatorname{Ranks}(\mathbb{E})$ we define the sequence of projections

$$
\begin{equation*}
\mathrm{P}_{\mathbf{v}, \mathrm{r}}=\mathrm{P}_{\left.\mathbf{v}, r_{e(\# \mathbb{E}}\right)}^{e(\# \mathbb{E})} \cdots \mathrm{P}_{\mathbf{v}, r_{e(2)}}^{e(2)} \mathrm{P}_{\mathbf{v}, r_{e(1)}}^{e(1)} \tag{2.25}
\end{equation*}
$$

which projects onto truncated mode frames starting from the root of $\mathbb{T}$. This sequence of projections realizes a truncation to hierarchical rank at most $r$.
Lemma 2.34. Let $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ and $r \in \operatorname{Ranks}(\mathbb{E})$ with $r \leq \operatorname{rank}_{\mathbb{T}}(\mathbf{v})$. Then $\operatorname{rank}_{\mathbb{E}}\left(\mathrm{P}_{\mathbf{v}, \mathrm{r}} \mathbf{v}\right) \leq \mathrm{r}$.

Proof. As a consequence of the ordering in (2.25), for each $e=\left\{\alpha, \alpha^{\mathrm{c}}\right\} \in \mathbb{E}$, once the projection $\mathrm{P}_{\mathbf{v}, r_{e}}^{e}$ has been applied, all projections that follow in the ordering act only on modes that are subsets of either $\alpha$ or $\alpha^{\mathrm{c}}$ and thus act from either the left or right on $\operatorname{mat}_{\alpha}(\mathbf{v})$. The subsequent projections can thus only decrease $\operatorname{rank}_{\alpha}(\mathbf{v})$.

Note that the ordering of the projections in (2.25) is essential: without the ordering from root to leaves, applying the projections may indeed increase the ranks beyond the prescribed $r$; see Grasedyck (2009/10).

Remark 2.35 (applying the truncation). It is easy to see that applying $\mathrm{P}_{\mathbf{v}, \mathrm{r}}$ to an HSVD representation amounts to selecting a subset of basis vectors in each mode frame. Since the mode frames are ordered by non-increasing singular values, this is accomplished by simply truncating all transfer tensors and mode frames in the HSVD to match the hierarchical rank r. A completely analogous procedure can be performed for the tensor train format (Oseledets 2011b).

Truncating the HSVD does not provide the best approximation with the given ranks from Theorem 2.31. However, we instead obtain a quasi-best approximation: the truncation error can be estimated by a fixed multiple of the best approximation
error. The following result to this effect was obtained for hierarchical tensors in Grasedyck (2009/10).

Theorem 2.36. Let $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ and $r=\left(r_{e}\right)_{e \in \mathbb{E}} \in \operatorname{Ranks}(\mathbb{E})$ with $r \leq \operatorname{rank}_{\mathbb{E}}(\mathbf{v})$. Then

$$
\begin{aligned}
\left\|\mathbf{v}-\mathrm{P}_{\mathbf{v}, \mathrm{r}} \mathbf{v}\right\|_{\ell_{2}} & \leq\left(\sum_{e \in \mathbb{E}} \sum_{k>r_{e}}\left(\sigma_{k}^{[e]}(\mathbf{v})\right)^{2}\right)^{1 / 2} \\
& \leq \sqrt{\# \mathbb{E}} \min \left\{\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}}: \mathbf{w} \in \ell_{2}\left(\mathbb{N}^{d}\right), \operatorname{rank}_{\mathbb{E}}(\mathbf{w}) \leq \mathrm{r}\right\}
\end{aligned}
$$

Proof. Let $P$ be an orthogonal projection and $Q$ a bounded operator on $\ell_{2}\left(\mathbb{N}^{d}\right)$, respectively. Then, since range $(I-P) \perp$ range $(P)$ and $\|P\|_{\ell_{2} \rightarrow \ell_{2}}=1$,

$$
\begin{aligned}
\|\mathbf{v}-P Q \mathbf{v}\|_{\ell_{2}}^{2} & =\|(I-P) \mathbf{v}+P(I-Q) \mathbf{v}\|_{\ell_{2}}^{2} \\
& =\|(I-P) \mathbf{v}\|_{\ell_{2}}^{2}+\|P(I-Q) \mathbf{v}\|_{\ell_{2}}^{2} \\
& \leq\|(I-P) \mathbf{v}\|_{\ell_{2}}^{2}+\|(I-Q) \mathbf{v}\|_{\ell_{2}}^{2} .
\end{aligned}
$$

Applying this inductively, starting from

$$
P=\mathrm{P}_{\mathbf{v}, r_{e(\# \mathrm{E})}}^{e(\# \mathbb{E})} \quad \text { and } \quad Q=\mathrm{P}_{\left.\mathbf{v}, r_{e(\# \mathbb{E}}-1\right)}^{e(\# \mathbb{E}-1)} \cdots \mathrm{P}_{\mathbf{v}, r_{e(1)}}^{e(1)},
$$

we arrive at

$$
\left\|\mathbf{v}-\mathrm{P}_{\mathbf{v}, \mathbf{v}}\right\|_{\ell_{2}}^{2} \leq \sum_{e \in \mathbb{E}}\left\|\mathbf{v}-\mathrm{P}_{\mathbf{v}, r_{e}}^{e} \mathbf{v}\right\|_{\ell_{2}}^{2}
$$

By Theorem 2.8, for each $e \in \mathbb{E}$,

$$
\left\|\mathbf{v}-\mathrm{P}_{\mathbf{v}, r_{e}}^{e} \mathbf{v}\right\|_{\ell_{2}}^{2}=\sum_{k>r_{e}}\left(\sigma_{k}^{[e]}(\mathbf{v})\right)^{2}=\min \left\{\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}}^{2}: \operatorname{rank}_{[e]}(\mathbf{w}) \leq r_{e}\right\} .
$$

Since $\min \left\{\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}}: \operatorname{rank}_{[e]}(\mathbf{w}) \leq r_{e}\right\} \leq \min \left\{\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}}: \operatorname{rank}_{\mathbb{E}}(\mathbf{w}) \leq r\right\}$, the statement follows.

Remark 2.37. Theorem 2.36 applies to all subspace tensor formats discussed so far.
(i) In the case of the Tucker format, which corresponds to the non-binary dimension tree $\mathbb{T}=\{\{1, \ldots, d\},\{1\}, \ldots,\{d\}\}$ with $\# \mathbb{E}=d$, Theorem 2.36 reduces to the error bounds for truncated higher-order singular value decompositions obtained by De Lathauwer et al. (2000).
(ii) For matricizations (2.22) of the tensor train format, we have \#E $=d-1$. In this case Theorem 2.36 yields the result for the tensor train format obtained by Oseledets (2011b).

Remark 2.38. The rank truncation described here is based on transforming the entire representation to HSVD form and then performing the truncation. Variants that sequentially transform single matricizations into SVD form and directly
perform the corresponding truncation have been investigated in Kühn (2012) (see also Hackbusch 2019, Sec. 11.4.2.2); the procedure for truncating the TT-SVD proposed by Oseledets (2011b) is also of this type. An analogous method for general tree tensor networks has been analysed by Ceruti, Lubich and Sulz (2023). Due to the decreased ranks of intermediate results, in practice this sequential truncation can reduce the total computational costs. A potential disadvantage is that error tolerances need to be allocated to matricizations before all singular values are known.

### 2.7.2. Soft thresholding of hierarchical tensor representations

We now turn to an alternative strategy for rank reduction of hierarchical tensors with the additional feature of non-expansiveness, which can be exploited in the construction of iterative methods.

Let $\mathbf{v} \in \ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)$, where $\mathcal{I}_{1}, \mathcal{I}_{2}$ are countable index sets, have the singular value decomposition

$$
\begin{equation*}
\mathbf{v}=\sum_{k=1}^{\infty} \sigma_{k} \mathbf{U}_{k}^{(1)} \otimes \mathbf{U}_{k}^{(2)} \tag{2.26}
\end{equation*}
$$

according to Theorem 2.8. For $\delta \geq 0$ and $x \in \mathbb{R}$, let $s_{\delta}(x)=\operatorname{sgn}(x) \max \{|x|-\delta, 0\}$, which satisfies

$$
\begin{equation*}
s_{\delta}(x)=\underset{y \in \mathbb{R}}{\arg \min }\left\{\frac{1}{2}|x-y|^{2}+\delta|y|\right\} . \tag{2.27}
\end{equation*}
$$

We define the nonlinear soft thresholding operator $S_{\delta}$ on $\ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)$ in terms of the singular value decomposition (2.26) by

$$
\begin{equation*}
S_{\delta}(\mathbf{v})=\sum_{k=1}^{\infty} s_{\delta}\left(\sigma_{k}\right) \mathbf{U}_{k}^{(1)} \otimes \mathbf{U}_{k}^{(2)} \tag{2.28}
\end{equation*}
$$

Lemma 2.39. For each $\delta \geq 0$, the operator $S_{\delta}$ is non-expansive, that is,

$$
\left\|S_{\delta}(\mathbf{v})-S_{\delta}(\mathbf{w})\right\|_{\ell_{2}} \leq\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}} \quad \text { for all } \mathbf{v}, \mathbf{w} \in \ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)
$$

Proof. We have the characterization

$$
\begin{equation*}
S_{\delta}(\mathbf{v})=\underset{\mathbf{w} \in \ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)}{\arg \min }\left\{\frac{1}{2}\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}}^{2}+\delta\|\sigma(\mathbf{w})\|_{\ell_{1}}\right\} \tag{2.29}
\end{equation*}
$$

where $\sigma(\mathbf{w})$ denotes the sequence of singular values of $\mathbf{w}$. This is shown in Bachmayr and Schneider (2017, Lem. 2.3) by a reduction to the scalar case (2.27); an alternative proof for $\# \mathcal{I}_{1}, \# \mathcal{I}_{2}<\infty$ based on subgradient characterizations is given in Cai, Candès and Shen (2010). The right-hand side (2.29) is the definition of the proximity operator on $\ell_{2}\left(\mathcal{I}_{1} \times \mathcal{I}_{2}\right)$ of the functional $\|\sigma(\cdot)\|_{\ell_{1}}$. As shown by Moreau (1965), every proximity operator is non-expansive.

Based on (2.28), we now define a soft thresholding operation for hierarchical tensors in $\ell_{2}\left(\mathbb{N}^{d}\right)$ that preserves the non-expansiveness. For a given dimension tree $\mathbb{T}$ and associated effective edges $\mathbb{E}$, for each $e \in \mathbb{E}$, we set

$$
S_{\delta}^{e}=\operatorname{mat}_{[e]}^{-1} \circ S_{\delta} \circ \operatorname{mat}_{[e]}
$$

We now choose an enumeration $e(1), \ldots, e(\# \mathbb{E})$ of the elements of $\mathbb{E}$ and define the hierarchical tensor soft thresholding operator

$$
\begin{equation*}
\mathrm{ST}_{\delta}=S_{\delta}^{e(\# \mathbb{E})} \circ \cdots \circ S_{\delta}^{e(1)} \tag{2.30}
\end{equation*}
$$

of Bachmayr and Schneider (2017). By repeated application of Lemma 2.39, we immediately obtain non-expansiveness of $\mathrm{ST}_{\delta}$ on $\ell_{2}\left(\mathbb{N}^{d}\right)$.

Proposition 2.40. For any $\delta \geq 0$,

$$
\left\|\mathrm{ST}_{\delta}(\mathbf{v})-\mathrm{ST}_{\delta}(\mathbf{w})\right\|_{\ell_{2}} \leq\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}} \quad \text { for all } \mathbf{v}, \mathbf{w} \in \ell_{2}\left(\mathbb{N}^{d}\right) .
$$

Remark 2.41. For any dimension tree, by a suitable enumeration of $\mathbb{E}$, the action of $\mathrm{ST}_{\delta}$ can be implemented with the same asymptotic complexity (2.24) as the computation of the HSVD; see Bachmayr and Schneider (2017, Alg. 1).

The following characterization of the $\delta$-dependence of the error incurred by application of $\mathrm{ST}_{\delta}$ in terms of singular values of matricizations is proved in Bachmayr and Schneider (2017, Lem. 3.4).

Theorem 2.42. For $\delta>0, e \in \mathbb{E}$ and $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$, with

$$
r_{\delta}^{e}(\mathbf{v})=\#\left\{k \in \mathbb{N}: \sigma_{k}^{[e]}(\mathbf{v})>\delta\right\}, \quad D_{\delta}^{e}(\mathbf{v})=\sqrt{\delta^{2} r_{\delta}^{e}(\mathbf{v})+\sum_{k>r_{\delta}^{e}(\mathbf{v})}\left|\sigma_{k}^{[e]}(\mathbf{v})\right|^{2}},
$$

we have the bounds

$$
\max _{e \in \mathbb{E}} D_{\delta}^{e}(\mathbf{v}) \leq\left\|\mathrm{ST}_{\delta}(\mathbf{v})-\mathbf{v}\right\|_{\ell_{2}} \leq \sum_{e \in \mathbb{E}} D_{\delta}^{e}(\mathbf{v}) .
$$

Moreover, $D_{\delta}^{e}(\mathbf{v}) \rightarrow 0$ as $\delta \rightarrow 0$ for each $e \in \mathbb{E}$.

### 2.8. Matrices and operators in tensor formats

When applying operators on $\ell_{2}\left(\mathbb{N}^{d}\right)$ to subspace representations, we again need to represent the results in the same format. For this to be efficiently realizable we require a representation of operators that is compatible with the considered low-rank format. In the case $d=2$ corresponding to low-rank matrices, as outlined in (1.14), sums of tensor product operators can be applied in this format and are in this sense compatible: we obtain a representation of the same form, generally with increased ranks.

In the case of the Tucker format, corresponding representations of operators on $\ell_{2}\left(\mathbb{N}^{d}\right)$ are of the analogous form

$$
\begin{equation*}
\sum_{\ell_{1}=1}^{R_{1}} \cdots \sum_{\ell_{d}=1}^{R_{d}} \mathbf{c}\left[\ell_{1}, \ldots, \ell_{d}\right] \mathbf{A}_{\ell_{1}}^{(i)} \otimes \cdots \otimes \mathbf{A}_{\ell_{d}}^{(i)} \tag{2.31}
\end{equation*}
$$

with component operators $\mathbf{A}_{\ell}^{(i)}, \ell=1, \ldots, R_{i}$, for $i=1, \ldots, d$. Applying this representation to a Tucker representation of multilinear ranks $\left(r_{1}, \ldots, r_{d}\right)$ then leads to a (potentially redundant) representation of the result with representation ranks $\left(R_{1} r_{1}, \ldots, R_{d} r_{d}\right)$.

The case of hierarchical tensors is similar: we have the same basic form as in (2.31), but the core tensor $\mathbf{c}$ needs to be decomposed in terms of transfer tensors as in (2.21) according to the given dimension tree $\mathbb{T}$. In the particular case of the tensor train representation, the shorthand notation in terms of the block form (2.16) can also be applied to operators, for instance

$$
\left[\begin{array}{ll}
\mathbf{I} & \mathbf{A}_{1}
\end{array}\right] \bowtie\left[\begin{array}{cc}
\mathbf{I} & \mathbf{A}_{2}  \tag{2.32}\\
0 & \mathbf{I}
\end{array}\right] \bowtie\left[\begin{array}{c}
\mathbf{A}_{\mathbf{3}} \\
\mathbf{I}
\end{array}\right]=\mathbf{A}_{1} \otimes \mathbf{I} \otimes \mathbf{I}+\mathbf{I} \otimes \mathbf{A}_{2} \otimes \mathbf{I}+\mathbf{I} \otimes \mathbf{I} \otimes \mathbf{A}_{3} .
$$

Applying this representation to a tensor train of representation rank (2, 2),

$$
\left[\begin{array}{ll}
\mathbf{u}_{1} & \mathbf{u}_{2}
\end{array}\right] \bowtie\left[\begin{array}{ll}
\mathbf{v}_{1,1} & \mathbf{v}_{1,2} \\
\mathbf{v}_{2,1} & \mathbf{v}_{2,2}
\end{array}\right] \bowtie\left[\begin{array}{l}
\mathbf{w}_{1} \\
\mathbf{w}_{2}
\end{array}\right],
$$

yields the result with representation rank $(4,4)$,

$$
\left[\begin{array}{llll}
\mathbf{u}_{1} & \mathbf{u}_{2} & \mathbf{A}_{1} \mathbf{u}_{1} & \mathbf{A}_{1} \mathbf{u}_{2}
\end{array}\right] \bowtie\left[\begin{array}{cccc}
\mathbf{v}_{1,1} & \mathbf{v}_{1,2} & \mathbf{A}_{2} \mathbf{v}_{1,1} & \mathbf{A}_{2} \mathbf{v}_{1,2} \\
\mathbf{v}_{2,1} & \mathbf{v}_{2,2} & \mathbf{A}_{2} \mathbf{v}_{2,1} & \mathbf{A}_{2} \mathbf{v}_{2,2} \\
0 & 0 & \mathbf{v}_{1,1} & \mathbf{v}_{1,2} \\
0 & 0 & \mathbf{v}_{2,1} & \mathbf{v}_{2,2}
\end{array}\right] \bowtie\left[\begin{array}{c}
\mathbf{A}_{3} \mathbf{w}_{1} \\
\mathbf{A}_{3} \mathbf{w}_{2} \\
\mathbf{w}_{1} \\
\mathbf{w}_{2}
\end{array}\right] .
$$

### 2.9. Condition numbers of tensor representations

The subspace tensor formats considered in this section avoid some of the difficulties associated with the canonical format that are outlined in Section 1.4. However, under certain conditions these representations can still be sensitive to round-off errors. This can happen in particular when the components are in non-orthogonalized form with redundancies that lead to large cancellations, which can occur in intermediate results in numerical methods.

Example 2.43. In order to illustrate how this issue can arise in tensor decompositions, let us consider the tensor $\mathbf{v} \in \bigotimes_{\ell=1}^{d} \mathbb{R}^{2}$ with all entries equal to one, which can be represented in TT format with rank one in the form

$$
\mathbf{v}=\binom{1}{1} \otimes \cdots \otimes\binom{1}{1}
$$

However, one can also write $\mathbf{v}$ in a redundant rank-two representation

$$
\mathbf{v}=\mathbf{X}_{1} \bowtie \cdots \bowtie \mathbf{X}_{d}
$$

with, for any $R>0$,

$$
\begin{align*}
& \mathbf{X}_{1}=\left[\left(1+R^{-d}\right)\binom{R}{R} \quad-\binom{R}{R}\right], \\
& \mathbf{X}_{2}=\cdots=\mathbf{X}_{d-1}=\left[\begin{array}{ll}
\binom{R}{R} & \binom{0}{0} \\
\binom{0}{0} & \binom{R}{R}
\end{array}\right], \quad \mathbf{X}_{d}=\left[\begin{array}{l}
\binom{R}{R} \\
\binom{R}{R}
\end{array}\right] . \tag{2.33}
\end{align*}
$$

We now define $\mathbf{v}_{\boldsymbol{\varepsilon}}$ by replacing precisely one of the cores $\mathbf{X}_{i}$ with $1<i<d$ by

$$
\mathbf{X}_{i, \varepsilon}=\left[\begin{array}{cc}
(1+\varepsilon)\binom{R}{R} & \binom{0}{0} \\
\binom{0}{0} & \binom{R}{R}
\end{array}\right],
$$

so that $\mathbf{v}_{\varepsilon}$ has all entries equal to $1+\left(R^{d}+1\right) \varepsilon$. For the relative error in the represented tensor, we thus obtain

$$
\frac{\left\|\mathbf{v}-\mathbf{v}_{\varepsilon}\right\|_{2}}{\|\mathbf{v}\|_{2}}=\left(R^{d}+1\right) \varepsilon
$$

Even when $R$ is of moderate size, a significant loss of precision can thus occur for large $d$.

Redundancies similar to those in Example 2.43 can arise in particular when applying operators in low-rank representation. Although one can subsequently apply orthogonalization to remove the redundancies, in general these orthogonalizations will already lead to large numerical errors. This effect can be quantified by the following notion introduced in Bachmayr and Kazeev (2020).

Definition 2.44. Let $\tau: \mathcal{C} \rightarrow \ell_{2}\left(\mathbb{N}^{d}\right)$ be the representation mapping of a tensor format, where $\mathcal{C}$ is the set of admissible tuples of components, and let $\mathrm{X}=$ $\left(\mathbf{X}_{1}, \ldots, \mathbf{X}_{N}\right) \in \mathcal{C}$. For $n=1, \ldots, N$, we define the representation condition numbers of $X$ by

$$
\begin{align*}
\operatorname{rcond}_{n}(\mathrm{X})= & \lim _{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \sup \left\{\frac{\|\tau(\mathrm{X})-\tau(\mathrm{Y})\|_{\ell_{2}\left(\mathbb{N}^{d}\right)}}{\left.\|\tau(\mathrm{X})\|_{\ell_{2}(\mathbb{N} d}\right)}\right. \\
\mathrm{Y}=\left(\mathbf{X}_{1}, \ldots, \mathbf{X}_{n-1},\right. & \left.\mathbf{Y}_{n}, \mathbf{X}_{n+1}, \ldots, \mathbf{X}_{N}\right), \\
& \left.\left\|\mathbf{X}_{n}-\mathbf{Y}_{n}\right\|_{\ell_{2}} \leq \varepsilon\left\|\mathbf{X}_{n}\right\|_{\ell_{2}}\right\} . \tag{2.34}
\end{align*}
$$

For tensor trains and more general hierarchical tensors, the quantities defined in (2.34) can be computed directly from norms of components; see Bachmayr and Kazeev (2020, Prop. 1). By multilinearity of $\tau$ for these tensor formats, if $\mathbf{x}=\tau(\mathrm{X})$ and $\mathbf{y}=\tau(\mathrm{Y})$ with $\mathrm{X}=\left(\mathbf{X}_{1}, \ldots, \mathbf{X}_{N}\right)$ and $\mathrm{Y}=\left(\mathbf{Y}_{1}, \ldots, \mathbf{Y}_{N}\right)$ such that

$$
\left\|\mathbf{X}_{n}-\mathbf{Y}_{n}\right\|_{\ell_{2}} \leq \varepsilon\left\|\mathbf{X}_{n}\right\|_{\ell_{2}} \text { for } n=1, \ldots, N
$$

we obtain

$$
\frac{\|\mathbf{x}-\mathbf{y}\|_{\ell_{2}}}{\|\mathbf{x}\|_{\ell_{2}}} \leq \sum_{n=1}^{N} \operatorname{rcond}_{n}(\mathrm{X}) \varepsilon+O\left(\varepsilon^{2}\right)
$$

Definition 2.45. With the notation of Definition 2.44, let A be an operator representation acting on elements of $\mathcal{C}$ and let $A \bullet X$ be the result of applying $A$ to $X \in \mathcal{C}$. For $n=1, \ldots, N$, we then define the operator representation condition number of A as

$$
\begin{equation*}
\operatorname{oprcond}_{n}(\mathrm{~A})=\sup _{\mathrm{X} \in \mathcal{C}} \frac{\operatorname{rcond}_{n}(\mathrm{~A} \bullet \mathrm{X})}{\operatorname{rcond}_{n}(\mathrm{X})} \tag{2.35}
\end{equation*}
$$

Hence operator representations that are ill-conditioned in the sense of large condition numbers (2.35) can lead to a correspondingly large deterioration of the representation condition numbers of the tensors that they are acting on, for instance by introducing cancellations in the representation as in Example 2.43. This can occur in certain representations of differential operators considered in Sections 3.4 and 4.5.

### 2.10. Manifold structure

The set of matrices of a given size $m \times n$ of fixed rank $r \in \mathbb{N}$ is a smooth embedded submanifold of $\mathbb{R}^{m \times n}$ (Helmke and Shayman 1995). This means that techniques of Riemannian optimization can be applied for solving minimization problems on fixed-rank matrices (Absil, Mahony and Sepulchre 2008), considered in further detail in Section 5.1.3. The structure of such fixed-rank matrix manifolds and their tangent spaces is also utilized in dynamical low-rank approximation (Koch and Lubich 2007a) of initial value problems; see Section 7.1.

Analogous manifold properties hold for Tucker tensors, tensor trains and hierarchical tensors with fixed ranks. This has been established for tensors with finite mode sizes for Tucker tensors in Koch and Lubich (2010), for tensor trains in Holtz et al. (2012b) and, using a different approach, for hierarchical tensors and tensor trains in Uschmajew and Vandereycken (2013); see also Arnold and Jahnke (2014). These constructions can be transferred to the tensor product Hilbert space setting of $\ell_{2}\left(\mathbb{N}^{d}\right)$; see Uschmajew (2013). The Tucker format in Banach spaces is considered in Falcó, Hackbusch and Nouy (2019). For a general discussion of manifold properties in the context of tensor networks, we refer to Bachmayr et al. (2016, Sec. 3.8).

## 3. Approximations of functions and operators in low-rank tensor formats

When using tensor formats in approximation problems, a natural approach is to directly assign tensor modes to coordinates in the given problem. In this case entries of a tensor representing a function typically correspond to values on a product grid or to coefficients in a basis expansion. However, tensors can also be used for encoding functions in completely different ways, where the modes are associated to basis functions or to scales in the problem. We now consider some common ways of representing functions by tensors and corresponding representations of operators. As a particular example, we consider representations of the Laplacian as the most basic example of an elliptic operator and the simplest corresponding boundary value problem.

Example 3.1. On the bounded domain $\Omega$, the classical formulation of the Poisson problem reads: for a given $f \in C(\Omega)$, find $u \in C^{2}(\Omega) \cap C(\bar{\Omega})$ such that

$$
\begin{equation*}
-\Delta u=f \text { in } \Omega, \quad u=0 \text { on } \partial \Omega \tag{3.1}
\end{equation*}
$$

The corresponding weak formulation is: for $f \in H^{-1}(\Omega)$, find $u \in H_{0}^{1}(\Omega)$ such that

$$
\begin{equation*}
\int_{\Omega} \nabla u \cdot \nabla v \mathrm{~d} x=f(v) \quad \text { for all } v \in H_{0}^{1}(\Omega) \tag{3.2}
\end{equation*}
$$

We first consider grid-based discretizations by finite differences or finite elements in Section 3.1. In Section 3.2 we then turn to representations in terms of Riesz bases, which play a central role in the following sections in obtaining well-conditioned problems and in adaptively refining discretizations. More specialized uses of lowrank tensors are considered in Section 3.3 on occupation number representations and Section 3.4 on tensorized representations, where tensor modes correspond to basis functions and to scales in the approximation problem, respectively.

### 3.1. Uniform grids

The most common way of representing function approximations in low-rank tensor methods is by point values on a regular tensor product grid. As an example, let us consider a uniform product grid in $\Omega=\square_{d}$ with $N$ points in each coordinate direction. With the grid spacing $h=1 /(N+1)$, the set of grid points in $\Omega$ is given by

$$
\Omega_{h}=\left\{x_{j_{1}, \ldots, j_{d}}: j \in\{1, \ldots, N\}^{d}\right\}, \quad \text { where } x_{j_{1}, \ldots, j_{d}}=\left(j_{1} h, \ldots, j_{d} h\right) .
$$

### 3.1.1. Finite differences

In finite difference discretizations, the original solution $u$ of the $\operatorname{PDE}$ on $\Omega$ is replaced by a grid function $u_{h}: \Omega_{h} \rightarrow \mathbb{R}$ given by the coefficient tensor $\mathbf{u}_{h} \in$ $\mathbb{R}^{N \times \cdots \times N}$ with entries

$$
\mathbf{u}_{h}\left[j_{1}, \ldots, j_{d}\right]=u_{h}\left(x_{j_{1}, \ldots, j_{d}}\right) \approx u\left(x_{j_{1}, \ldots, j_{d}}\right)
$$

Using low-rank tensor approximations of $\mathbf{u}_{h}$ in finite difference methods requires suitable representations of the corresponding discretization matrices. In Example 3.1 with $\Omega=\square_{d}$ and $f \in C(\Omega)$, we obtain a linear system of equations $\mathbf{A}_{\mathrm{FD}} \mathbf{u}_{h}=\mathbf{f}_{h}$, where $\mathbf{f}_{h}[j]=f\left(x_{j}\right)$ for $j \in\{1, \ldots, N\}^{d}$ and

$$
\begin{align*}
\mathbf{A}_{\mathrm{FD}}= & h^{-2} \boldsymbol{\Delta}_{2} \otimes \mathbf{I} \otimes \mathbf{I} \otimes \cdots \otimes \mathbf{I}+\mathbf{I} \otimes h^{-2} \boldsymbol{\Delta}_{2} \otimes \mathbf{I} \otimes \cdots \otimes \mathbf{I} \\
& +\cdots+\mathbf{I} \otimes \mathbf{I} \otimes \cdots \otimes \mathbf{I} \otimes h^{-2} \boldsymbol{\Delta}_{2}, \tag{3.3}
\end{align*}
$$

where the $\mathbf{I}$ are $N \times N$ identity matrices and where we set

$$
\boldsymbol{\Delta}_{2}=\left(\begin{array}{ccccc}
2 & -1 & & &  \tag{3.4}\\
-1 & 2 & -1 & & \\
& \ddots & \ddots & \ddots & \\
& & \ddots & \ddots & -1 \\
& & & -1 & 2
\end{array}\right)
$$

Although the right-hand side of (3.3) has $d$ summands, the matrix $\mathbf{A}_{\mathrm{FD}}$ can be written in the operator versions of the Tucker, tensor train or hierarchical tensor formats with all ranks equal to two. In particular, as observed by Kazeev and Khoromskij (2012), in the tensor train format we have, analogously to (2.32),

$$
\mathbf{A}_{\mathrm{FD}}=\left[\begin{array}{ll}
\mathbf{I} & h^{-2} \boldsymbol{\Delta}_{2}
\end{array}\right] \bowtie\left[\begin{array}{cc}
\mathbf{I} & h^{-2} \boldsymbol{\Delta}_{2} \\
0 & \mathbf{I}
\end{array}\right]^{\bowtie(d-2)} \bowtie\left[\begin{array}{c}
h^{-2} \boldsymbol{\Delta}_{2} \\
\mathbf{I}
\end{array}\right],
$$

where for a core $\mathbf{X}$ we write $\mathbf{X}^{\bowtie n}$ for the $n$-fold strong Kronecker product of $\mathbf{X}$ with itself. It is easy to see that for the condition number with respect to the spectral norm, we have $\operatorname{cond}\left(\mathbf{A}_{\mathrm{FD}}\right)=\operatorname{cond}\left(\boldsymbol{\Delta}_{2}\right)$ independently of $d$, where $\operatorname{cond}\left(\boldsymbol{\Delta}_{2}\right)=O\left(h^{-2}\right)$.

We thus obtain a discretization with simple structure that serves as a benchmark problem in many works, where it is typically considered with rather coarse grids of sizes $N=2^{5}$ to $N=2^{7}$ that can be treated without preconditioning. For large grid sizes, the use of a preconditioner in low-rank format becomes necessary due to the large condition number of $\mathbf{A}_{\mathrm{FD}}$. However, an important limitation of such finite difference discretizations is that they offer no mechanism for obtaining a $d$-robust computable error bound for the discretization error. This is especially important in the high-dimensional setting, since the scaling with respect to $d$ of the quantities arising in a priori error estimates is typically unknown. Moreover, finite differences are less suitable for problems with non-smooth data (such as $f \notin C(\Omega)$ ).

### 3.1.2. Finite elements

An alternative discretization that also uses the values on the grid $\Omega_{h}$, but that is suitable for problems with low regularity and offers more advanced means for error control, is provided by lowest-order multilinear finite elements. In this case, with the univariate piecewise linear hat functions $\varphi_{j}(x)=\max \{1-|x-j h| / h, 0\}$, $j \in\{1, \ldots, N\}$, with $h=1 /(N+1)$ as above, we consider a Galerkin discretization
of (3.2): find $u_{h} \in V_{h}$ such that

$$
\begin{equation*}
\int_{\Omega} \nabla u_{h} \cdot \nabla v_{h} \mathrm{~d} x=f\left(v_{h}\right) \quad \text { for all } v_{h} \in V_{h} \tag{3.5}
\end{equation*}
$$

with the subspace

$$
V_{h}=\operatorname{span}\left\{\bigotimes_{i=1}^{d} \varphi_{j_{i}}: j \in\{1, \ldots, N\}^{d}\right\} \subset H_{0}^{1}(\Omega)
$$

The coefficients $\mathbf{u}_{h} \in \mathbb{R}^{N \times \cdots \times N}$ such that

$$
u_{h}=\sum_{j \in\{1, \ldots, N\}^{d}} \mathbf{u}_{h}[j] \bigotimes_{i=1}^{d} \varphi_{j_{i}}
$$

then again correspond to the values of $u_{h}$ at the grid points. The system of equations resulting from (3.5) is $\mathbf{A}_{\mathrm{FE}} \mathbf{u}_{h}=\mathbf{f}_{h}$ with

$$
\mathbf{f}_{h}[j]=f\left(\bigotimes_{i=1}^{d} \varphi_{j_{i}}\right)
$$

and

$$
\begin{align*}
\mathbf{A}_{\mathrm{FE}}= & \mathbf{A}_{1} \otimes \mathbf{M}_{1} \otimes \mathbf{M}_{1} \otimes \cdots \otimes \mathbf{M}_{1}+\mathbf{M}_{1} \otimes \mathbf{A}_{1} \otimes \mathbf{M}_{1} \otimes \cdots \otimes \mathbf{M}_{1} \\
& +\cdots+\mathbf{M}_{1} \otimes \mathbf{M}_{1} \otimes \cdots \otimes \mathbf{M}_{1} \otimes \mathbf{A}_{1} \tag{3.6}
\end{align*}
$$

where

$$
\mathbf{A}_{1}=\left(\int_{0}^{1} \varphi_{i}^{\prime} \varphi_{j}^{\prime} \mathrm{d} x\right)_{i, j=1, \ldots, N}=h^{-1} \boldsymbol{\Delta}_{2}
$$

and

$$
\mathbf{M}_{1}=\left(\int_{0}^{1} \varphi_{i} \varphi_{j} \mathrm{~d} x\right)_{i, j=1, \ldots, N}=\frac{h}{6}\left(\begin{array}{ccccc}
4 & 1 & & & \\
1 & 4 & 1 & & \\
& \ddots & \ddots & \ddots & \\
& & \ddots & \ddots & 1 \\
& & & 1 & 4
\end{array}\right)
$$

Concerning representation ranks, $\mathbf{A}_{\mathrm{FE}}$ can be represented analogously to $\mathbf{A}_{\mathrm{FD}}$ with all ranks equal to two in the formats considered.

Remark 3.2. In contrast to the finite difference case, since $\operatorname{cond}\left(\mathbf{M}_{1}\right)>1$, it is easy to see that the condition number of $\mathbf{A}_{\mathrm{FE}}$ deteriorates exponentially with increasing $d$. This means in particular that iterative methods applied to $\mathbf{A}_{\mathrm{FE}}$ become extremely inefficient for large $d$. Note that once the discrete problem is represented in the basis of multilinear hat functions, this problem is not easy to circumvent, since any application of $\mathbf{M}_{1} \otimes \cdots \otimes \mathbf{M}_{1}$ or its inverse may incur a relative error increasing exponentially with respect to $d$.

### 3.2. Basis representations

Rather than working with grid values, in many approximation problems it is more advantageous to consider expansions in terms of bases of the relevant function spaces that can serve as dictionaries for nonlinear sparse approximation as outlined in Section 1.6.

Example 3.3. In parametric and stochastic problems, we are often interested in approximation in spaces $L_{2}\left(\mathbb{R}^{\mathcal{N}}, \mu\right)$ with $\mu=\bigotimes_{i \in \mathcal{N}} \mu_{i}$, where $\mathcal{N}$ can be finite or countable and $\mu_{i}$ is a probability measure on $\mathbb{R}$ for each $i \in \mathcal{N}$. To approximate smooth functions in such spaces, a natural choice of basis is given by tensor products of orthonormal polynomials $\left(P_{v}^{(i)}\right)_{v \in \mathbb{N}_{0}}$ in $L_{2}\left(\mathbb{R}, \mu_{i}\right)$ for $i \in \mathcal{N}$, so that each $u \in L_{2}\left(\mathbb{R}^{\mathcal{N}}, \mu\right)$ has an expansion of the form

$$
u=\sum_{v \in \mathcal{F}} \mathbf{u}[v] \bigotimes_{i \in \mathcal{N}} P_{v_{i}}^{(i)}, \quad \text { where } \mathbf{u}[v]=\int_{\mathbb{R}^{\mathcal{N}}} u \bigotimes_{i \in \mathcal{N}} P_{v_{i}}^{(i)} \mathrm{d} \mu, v \in \mathcal{F},
$$

with the index set

$$
\mathcal{F}=\left\{v \in \mathbb{N}_{0}^{\mathcal{N}}: v_{i} \neq 0 \text { for finitely many } i \in \mathcal{N}\right\}
$$

The coefficient sequences $\mathbf{u}$, restricted to suitable finite Cartesian product subsets of $\mathcal{F}$, can be approximated in low-rank formats. For a detailed treatment of such product polynomial expansions, we refer to Schwab and Gittelson (2011) and Cohen and DeVore (2015).

In many scenarios, the notion of orthonormal bases is too restrictive, and we instead work with bases satisfying the following weaker requirements.

Definition 3.4. Let $H$ be a separable Hilbert space, let $\mathcal{I}$ be a countable set and let $\gamma_{v} \in H$ for $v \in \mathcal{I}$. Then $\left\{\gamma_{v}\right\}_{v \in \mathcal{I}}$ is a Riesz sequence if there exist $c, C>0$ such that for all $\mathbf{v} \in \ell_{2}(\mathcal{I})$,

$$
\begin{equation*}
c\|\mathbf{v}\|_{\ell_{2}(\mathcal{I})} \leq\left\|\sum_{v \in \mathcal{I}} \mathbf{v}[v] \gamma_{v}\right\|_{H} \leq C\|\mathbf{v}\|_{\ell_{2}(\mathcal{I})} \tag{3.7}
\end{equation*}
$$

and a Riesz basis of $H$ if in addition $\operatorname{span}\left\{\gamma_{v}: v \in \mathcal{I}\right\}$ is dense in $H$.
In other words, in addition to the density requirement, for a Riesz basis the mapping taking coefficient sequences to represented elements in the Hilbert space is an isomorphism; when $c=C=1$, the Riesz basis is in fact an orthonormal basis, where this mapping is an isometric isomorphism.

When approximating an arbitrary $u=\sum_{v \in \mathcal{I}} \mathbf{u}[v] \gamma_{v}$ with coefficient sequence $\mathbf{u} \in \ell_{2}(\mathcal{I})$ by an expansion $\tilde{u}=\sum_{v \in \Lambda} \mathbf{u}[v] \gamma_{v}$ with $\Lambda \subset \mathcal{I}$, by (3.7) we have in particular

$$
\begin{equation*}
c\|\mathbf{u}-\tilde{\mathbf{u}}\|_{\ell_{2}} \leq\|u-\tilde{u}\|_{H} \leq C\|\mathbf{u}-\tilde{\mathbf{u}}\|_{\ell_{2}} \tag{3.8}
\end{equation*}
$$

which means that up to the fixed constants $c, C$, approximation in $H$-norm is equivalent to approximation of coefficient sequences in $\ell_{2}(\mathcal{I})$. This is particularly useful
in the context of low-rank methods, since the methods discussed in Section 2.7 for rank truncation of tensors based on the HSVD yield error bounds in $\ell_{2}$-norm.

For Riesz bases in tensor product Hilbert spaces, we have the following analogue of Theorem 2.5 for orthonormal bases.
Proposition 3.5. For $i=1, \ldots, d$, let $\left\{\gamma_{v}^{(i)}\right\}_{\nu \in \mathbb{N}}$ be a Riesz basis of the separable Hilbert space $H_{i}$ with

$$
c_{i}\|\mathbf{v}\|_{\ell_{2}(\mathbb{N})} \leq\left\|\sum_{v \in \mathbb{N}} \mathbf{v}[v] \gamma_{v}^{(i)}\right\|_{H_{i}} \leq C_{i}\|\mathbf{v}\|_{\ell_{2}(\mathbb{N})} \quad \text { for all } \mathbf{v} \in \ell_{2}(\mathbb{N})
$$

Then $\left\{\gamma_{v_{1}}^{(1)} \otimes \cdots \otimes \gamma_{v_{d}}^{(d)}\right\}_{v \in \mathbb{N}^{d}}$ is a Riesz basis of $H=\bigotimes_{i=1}^{d} H_{i}$, and for all $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$,

$$
\left(\prod_{i=1}^{d} c_{i}\right)\|\mathbf{v}\|_{\ell_{2}\left(\mathbb{N}^{d}\right)} \leq\left\|\sum_{v \in \mathbb{N}^{d}} \mathbf{v}[v] \bigotimes_{i=1}^{d} \gamma_{v_{i}}^{(i)}\right\|_{H} \leq\left(\prod_{i=1}^{d} C_{i}\right)\|\mathbf{v}\|_{\ell_{2}\left(\mathbb{N}^{d}\right)}
$$

The natural spaces for treating variational formulations of linear elliptic PDEs are the Hilbertian Sobolev spaces $H^{k}(\Omega)$ with $k \in \mathbb{N}$ on the given domain $\Omega$, in particular $H^{1}(\Omega)$ in the case of second-order problems as in Example 3.1. In particular, the weak formulation of a problem of this type with homogeneous Dirichlet boundary conditions as in (3.2) can be written as an operator equation $A u=f$ for $u \in H_{0}^{1}(\Omega)$, where $f \in H^{-1}(\Omega)$, and $A: H_{0}^{1}(\Omega) \rightarrow H^{-1}(\Omega)$ is defined by

$$
\begin{equation*}
\langle A v, w\rangle=\int_{\Omega} \nabla v \cdot \nabla w \mathrm{~d} x, \quad v, w \in H_{0}^{1}(\Omega) \tag{3.9}
\end{equation*}
$$

Note that by the Lax-Milgram theorem, $A$ is an isomorphism.
With the aid of a Riesz basis of $H_{0}^{1}(\Omega)$, such operator equations can be reduced to problems on $\ell_{2}$-sequence spaces; see Cohen, Dahmen and DeVore (2001) and Dahmen (1997).

Proposition 3.6. Let $V$ be a separable Hilbert space, let $A: V \rightarrow V^{\prime}$ be an isomorphism satisfying $\|A\|_{V \rightarrow V^{\prime}} \leq C_{A}$ and $\left\|A^{-1}\right\|_{V^{\prime} \rightarrow V} \leq c_{A}^{-1}$ with $c_{A}, C_{A}>0$, and let $\left\{\gamma_{\nu}\right\}_{\nu \in \mathcal{I}}$ be a Riesz basis of $V$ such that with $c_{\Gamma}, C_{\Gamma}>0$,

$$
c_{\Gamma}\|\mathbf{v}\|_{\ell_{2}(\mathcal{I})} \leq\left\|\sum_{v \in \mathcal{I}} \mathbf{v}[v] \gamma_{v}\right\|_{H} \leq C_{\Gamma}\|\mathbf{v}\|_{\ell_{2}(\mathcal{I})} \quad \text { for all } \mathbf{v} \in \ell_{2}(\mathcal{I}) .
$$

In addition, let the bi-infinite matrix $\mathbf{A} \in \mathbb{R}^{\mathcal{I} \times \mathcal{I}}$ and the sequence $\mathbf{f} \in \mathbb{R}^{\mathcal{I}}$ be defined by

$$
\begin{equation*}
\mathbf{A}\left[v, v^{\prime}\right]=\left\langle A \gamma_{v^{\prime}}, \gamma_{v}\right\rangle, \quad \mathbf{f}[v]=f\left(\gamma_{v}\right), \quad \text { for } v, v^{\prime} \in \mathcal{I} . \tag{3.10}
\end{equation*}
$$

Then $\mathbf{A}$ is an isomorphism on $\ell_{2}(\mathcal{I})$ with

$$
\begin{equation*}
c_{A} c_{\Gamma}^{2}\|\mathbf{v}\|_{\ell_{2}(\mathcal{I})} \leq\|\mathbf{A v}\|_{\ell_{2}(\mathcal{I})} \leq C_{A} C_{\Gamma}^{2}\|\mathbf{v}\|_{\ell_{2}(\mathcal{I})} \quad \text { for all } \mathbf{v} \in \ell_{2}(\mathcal{I}) \tag{3.11}
\end{equation*}
$$

Moreover, $\mathbf{f} \in \ell_{2}(\mathcal{I})$ and the unique solution $u \in V$ of $A u=f$ is given by $u=\sum_{v \in \mathcal{I}} \mathbf{u}[v] \gamma_{v}$, where $\mathbf{u} \in \ell_{2}(\mathcal{I})$ is the unique solution of $\mathbf{A u}=\mathbf{f}$.

Proof. With the isomorphism $S_{\Gamma}: \ell_{2}(\mathcal{I}) \rightarrow V, \mathbf{v} \mapsto \sum_{v \in \mathcal{I}} \mathbf{v}[v] \gamma_{v}$, we have $\mathbf{A}=S_{\Gamma}^{\prime} A S_{\Gamma}$ and $\mathbf{f}=S_{\Gamma}^{\prime} f$. Since $c_{\Gamma}\|\mathbf{v}\|_{\ell_{2}} \leq\left\|S_{\Gamma} \mathbf{v}\right\|_{V} \leq C_{\Gamma}\|\mathbf{v}\|_{\ell_{2}}$ for all $\mathbf{v} \in \ell_{2}(\mathcal{I})$ and $S_{\Gamma}^{\prime}: V^{\prime} \rightarrow \ell_{2}(\mathcal{I})$ is also an isomorphism with $c_{\Gamma}\|g\|_{V^{\prime}} \leq\left\|S_{\Gamma}^{\prime} g\right\|_{\ell_{2}} \leq C_{\Gamma}\|g\|_{V^{\prime}}$ for all $g \in V^{\prime}$, the statements follow.
Remark 3.7. The spaces $H^{1}(\Omega)$ and $H_{0}^{1}(\Omega)$ for $\Omega=\square_{d}$ cannot be characterized as tensor product Hilbert spaces, which means that no simple tensor product constructions of Riesz bases as in Proposition 3.5 are applicable. Note that

$$
\bigotimes_{i=1}^{d} H^{1}(0,1)=H_{\mathrm{mix}}^{1}\left(\square_{d}\right),
$$

where the cross-norm defined by (2.4) reads

$$
\|v\|_{H_{\text {mix }}^{1}}^{2}=\sum_{\substack{\alpha \in \mathbb{N}_{0}^{d} \\ \max _{i} \alpha_{i} \leq 1}} \int_{\square_{d}}\left|\partial^{\alpha} v\right|^{2} \mathrm{~d} x
$$

In contrast, $H^{1}\left(\square_{d}\right)$ can be obtained by taking the closure of $H_{\text {mix }}^{1}\left(\square_{d}\right)$ with respect to the standard $H^{1}$-norm

$$
\begin{equation*}
\|v\|_{H^{1}}^{2}=\sum_{\substack{\alpha \in \mathbb{N}_{0}^{d} \\ \alpha_{1}+\cdots+\alpha_{d} \leq 1}} \int_{\square_{d}}\left|\partial^{\alpha} v\right|^{2} \mathrm{~d} x \tag{3.12}
\end{equation*}
$$

and $H_{\text {mix }}^{1} \subsetneq H^{1}$. The norm (3.12) can also be characterized as the canonical norm on an intersection of tensor product Hilbert spaces, that is,

$$
H^{1}\left(\square_{d}\right)=\bigcap_{i=1}^{d} \bigotimes_{j=1}^{d} X_{i j}, \quad X_{i j}= \begin{cases}H^{1}(0,1) & i=j, \\ L_{2}(0,1) & i \neq j,\end{cases}
$$

and the analogous characterization holds for the subspace $H_{0}^{1}\left(\square_{d}\right)$ with $H^{1}(0,1)$ replaced by $H_{0}^{1}(0,1)$. A detailed treatment of related questions is given in Ali and Nouy (2020a).

In order to still obtain Riesz bases of $H^{1}\left(\square_{d}\right)$ with some degree of product structure, one can consider tensor product bases of $L_{2}\left(\square_{d}\right)$ that become Riesz bases of $H^{1}\left(\square_{d}\right)$ by appropriate normalization.

Example 3.8. Let $\phi_{j}(x)=\sqrt{2} \sin (j \pi x)$ for $j \in \mathbb{N}, x \in(0,1)$. Then $\left\{\phi_{j}\right\}_{j \in \mathbb{N}}$ is an orthonormal basis of $L_{2}(0,1)$ and by Theorem 2.5 , setting $\Phi_{j}=\bigotimes_{i=1}^{d} \phi_{j_{i}}$ for $j \in \mathbb{N}^{d},\left\{\Phi_{j}\right\}_{j \in \mathbb{N}^{d}}$ is an orthonormal basis of $L_{2}\left(\square_{d}\right)$. We easily verify that

$$
\begin{equation*}
\left\{\left\|\Phi_{j}\right\|_{H^{1}}^{-1} \Phi_{j}\right\}_{j \in \mathbb{N}^{d}}, \quad \text { where }\left\|\Phi_{j}\right\|_{H^{1}}=\sqrt{1+\pi^{2}\left(j_{1}^{2}+\cdots+j_{d}^{2}\right)}, \tag{3.13}
\end{equation*}
$$

is a Riesz basis (in fact, even an orthonormal basis) of $H_{0}^{1}\left(\square_{d}\right)$ with the norm (3.12). Here the tensor product functions $\Phi_{j}$ are modified by the factor $\left\|\Phi_{j}\right\|_{H^{1}}^{-1}$, which
is not separable with respect to $j_{1}, \ldots, j_{d}$. This lack of separability subsequently needs to be dealt with in any low-rank approximation method using the basis (3.13).

One construction of Riesz bases that is more suitable for locally refined approximations of potentially less regular functions is that of wavelets. A classical example of such a basis is the Haar wavelet: we define the scaling function $\varphi$ and the mother wavelet $\psi$ by

$$
\varphi=\chi_{[0,1]}, \quad \psi=\chi_{[0,1 / 2]}-\chi_{(1 / 2,1)}
$$

and, for $j \in \mathbb{N}_{0}$ and $k=0, \ldots, 2^{j}-1$, we introduce the notation

$$
\begin{equation*}
\varphi_{j, k}(x)=2^{j / 2} \varphi\left(2^{j} x-k\right), \quad \psi_{j, k}(x)=2^{j / 2} \psi\left(2^{j} x-k\right) \tag{3.14}
\end{equation*}
$$

Then, for each $J \in \mathbb{N}$,

$$
\begin{align*}
\operatorname{span}\left(\left\{\varphi_{0,0}\right\}\right. & \left.\cup\left\{\psi_{j, k}: j \in\{0, \ldots, J-1\}, k \in\left\{0, \ldots, 2^{j}-1\right\}\right\}\right) \\
= & \operatorname{span}\left\{\varphi_{J, k}: k \in\left\{0, \ldots, 2^{J}-1\right\}\right\} \tag{3.15}
\end{align*}
$$

Since the union over $J$ of the spaces of piecewise constant functions on the right is dense in $L_{2}(0,1)$, and by the orthonormality properties of the functions defined in (3.14), we obtain that

$$
\left\{\varphi_{0,0}\right\} \cup\left\{\psi_{j, k}: j \in \mathbb{N}_{0}, k \in\left\{0, \ldots, 2^{j}-1\right\}\right\}
$$

is an orthonormal basis of $L_{2}(0,1)$. Due to their lack of regularity, Haar wavelets are not suitable for the construction of Riesz bases of $H^{1}$-spaces.

Based on more general hierarchies of nested spaces, referred to as multiresolution analyses, wavelet bases of any desired regularity can be constructed. Here we give a brief overview, and refer to the survey by Dahmen (1997) and the monographs by Daubechies (1992) and Cohen (2003) for further details.

The wavelets constructed by Daubechies (1988) provide orthonormal bases of $L_{2}(\mathbb{R})$ with compactly supported scaling function and mother wavelet, and they can be constructed to have arbitrarily high Hölder and Sobolev regularity (at the price of increasing support size). The Haar wavelet arises as the case of lowest order of this construction. The wavelets in this family of higher regularity, however, are fractals that are given only in terms of refinement relations, and are thus difficult to use in numerical methods.

The biorthogonal wavelets constructed by Cohen, Daubechies and Feauveau (1992) are also compactly supported and can be constructed to have arbitrarily high regularity, but are obtained as linear combinations of spline functions and are thus easy to handle numerically. However, the price to pay is that these wavelets yield non-orthonormal Riesz bases of $L_{2}(\mathbb{R})$. The use of such wavelets on bounded intervals requires modifications to basis elements that touch the boundary. Such adaptations to wavelet Riesz bases of $L_{2}(0,1)$ that are suitable for numerical methods have been constructed by Dahmen, Kunoth and Urban (1999) and Primbs (2010).

For such wavelet bases, we use the following customary generic notation: we write wavelet bases as $\left\{\psi_{\nu}\right\}_{v \in \mathrm{~V}}$, where $\vee$ denotes a suitable countable index set such that the functions $\left\{\psi_{v}\right\}_{\nu \in V}$ comprise both the scaling functions on the lowest level and the wavelets that form the basis. In addition, each $v$ includes a level $j$, denoted by $|v|$, and a translation parameter $k$ in (3.14). Without loss of generality, for the lowest level in the basis, we assume $\min _{v \in \mathrm{~V}}|v|=0$. The boundary adaptations mentioned above then lead to wavelet bases that in particular have the following properties: with some $t>0$, for all $v \in \vee$,

$$
\begin{equation*}
\psi_{v} \in H^{1+t}(0,1), \quad \text { diam supp } \psi_{v} \lesssim 2^{-|v|} \tag{3.16}
\end{equation*}
$$

and there exist $c_{0}, c_{1}>0$ and $C_{0}, C_{1}>0$ such that for all $\mathbf{v} \in \ell_{2}(\vee)$,

$$
\begin{equation*}
c_{0}\|\mathbf{v}\|_{\ell_{2}(v)} \leq\left\|\sum_{v \in \mathrm{~V}} \mathbf{v}[v] \psi_{v}\right\|_{L^{2}(0,1)} \leq C_{0}\|\mathbf{v}\|_{\ell_{2}(v)} \tag{3.17}
\end{equation*}
$$

as well as

$$
\begin{equation*}
c_{1}\|\mathbf{v}\|_{\ell_{2}(v)} \leq\left\|\sum_{v \in \mathrm{~V}} \mathbf{v}[v]\right\| \psi_{v}\left\|_{H^{1}(0,1)}^{-1} \psi_{v}\right\|_{H^{1}(0,1)} \leq C_{1}\|\mathbf{v}\|_{\ell_{2}(v)} \tag{3.18}
\end{equation*}
$$

Remark 3.9. In constructing Riesz bases of $H^{1}\left(\square_{d}\right)$ using properties (3.17), (3.18), we face a similar issue to that observed for finite elements in Remark 3.2: as shown by Dijkema, Schwab and Stevenson (2009), for the $H^{1}$-normalized tensor product wavelets

$$
\begin{equation*}
\Psi_{v}=\left\|\bigotimes_{i=1}^{d} \psi_{v_{i}}\right\|_{H^{1}\left(\square_{d}\right)}^{-1} \bigotimes_{i=1}^{d} \psi_{v_{i}}, \quad v \in \mathrm{~V}^{d}, \tag{3.19}
\end{equation*}
$$

we obtain the bounds

$$
\begin{equation*}
c_{1} c_{0}^{d-1}\|\mathbf{v}\|_{\ell_{2}\left(\mathrm{~V}^{d}\right)} \leq\left\|\sum_{v \in \mathrm{~V}^{d}} \mathbf{v}[v] \Psi_{\nu}\right\|_{H^{1}\left(\square_{d}\right)} \leq C_{1} C_{0}^{d-1}\|\mathbf{v}\|_{\ell_{2}\left(\mathrm{~V}^{d}\right)} \tag{3.20}
\end{equation*}
$$

for all $\mathbf{v} \in \ell_{2}\left(\mathrm{~V}^{d}\right)$. Similarly to Remark 3.2, one can show that whenever $c_{0}<C_{0}$, the bounds in (3.20) and in the equivalence of errors (3.8) deteriorate exponentially with respect to $d$. In general, this equally affects the condition numbers of operator representations (3.11).

The above observation shows that to obtain a wavelet Riesz basis of $H^{1}\left(\square_{d}\right)$ with $d$-robust constants we essentially need to start from an orthonormal wavelet basis of $L_{2}(0,1)$ satisfying (3.17) with $c_{0}=C_{0}=1$. By the wavelet constructions mentioned above, we do not obtain the regularity and support properties (3.16) at the same time as orthonormality in $L_{2}(0,1)$ when using wavelets that are at the same time piecewise polynomials. However, this is achieved by certain spline multiwavelets, which are generated by dilations and translations of more than one mother wavelet, constructed by Donovan, Geronimo and Hardin (1996, 1999). For such multiwavelets we keep the above notation, where each index in $\vee$ now also
encodes the respective mother wavelet. The Riesz basis $\left\{\Psi_{\nu}\right\}_{\nu \in \mathrm{V} d}$ of $H^{1}\left(\square_{d}\right)$ obtained from these spline multiwavelets according to (3.19) then satisfies (3.20) with $c_{0}=C_{0}=1$, and thus with constants independent of $d$.

Remark 3.10. If, in addition to (3.16), (3.17) with $c_{0}=C_{0}=1$, and (3.18), we have $\psi_{v} \in H_{0}^{1}(0,1)$ for all $v \in \mathrm{~V}$, then with

$$
\Psi_{v}=\left(\sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right)^{-1 / 2} \bigotimes_{i=1}^{d} \psi_{v_{i}}
$$

$\left\{\Psi_{\nu}\right\}_{\nu \in \mathrm{V}^{d}}$ is a Riesz basis of $H_{0}^{1}\left(\square_{d}\right)$.
For the isomorphism $A: H_{0}^{1}\left(\square_{d}\right) \rightarrow H^{-1}\left(\square_{d}\right)$ defined by the Laplacian in (3.9), let us now consider basis representations as in (3.10) with respect to such a Riesz basis. By (3.11) and (3.20), for $\mathbf{A}: \ell_{2}\left(\mathrm{~V}^{d}\right) \rightarrow \ell_{2}\left(\mathrm{~V}^{d}\right)$ defined by

$$
\begin{equation*}
\mathbf{A}\left[v, v^{\prime}\right]=\left\langle A \Psi_{v^{\prime}}, \Psi_{v}\right\rangle=\int_{\square_{d}} \nabla \Psi_{v^{\prime}} \cdot \nabla \Psi_{v} \mathrm{~d} x, \quad v, v^{\prime} \in \vee^{d} \tag{3.21}
\end{equation*}
$$

we then have $c_{1}\|\mathbf{v}\|_{\ell_{2}} \leq\|\mathbf{A v}\|_{\ell_{2}} \leq C_{1}\|\mathbf{v}\|_{\ell_{2}}$ with $c_{1}, C_{2}$ from (3.18).
Moreover, we have the decomposition $\mathbf{A}=\mathbf{D T D}$ with the infinite diagonal matrix

$$
\mathbf{D}\left[v, v^{\prime}\right]=\left(\sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right)^{-1 / 2} \delta_{v, v^{\prime}}, \quad v, v^{\prime} \in \mathrm{V}^{d}
$$

and

$$
\mathbf{T}\left[v, v^{\prime}\right]=\sum_{i=1}^{d} \int_{0}^{1} \psi_{v_{i}}^{\prime} \psi_{v_{i}^{\prime}}^{\prime} \mathrm{d} x \prod_{j \neq i} \int_{0}^{1} \psi_{v_{j}} \psi_{v_{j}^{\prime}} \mathrm{d} x=\sum_{i=1}^{d} \int_{0}^{1} \psi_{v_{i}}^{\prime} \psi_{v_{i}^{\prime}}^{\prime} \mathrm{d} x \prod_{j \neq i} \delta_{v_{j}, v_{j}^{\prime}},
$$

so that with $\mathbf{T}_{1}\left[v, v^{\prime}\right]=\int_{0}^{1} \psi_{v}^{\prime} \psi_{v^{\prime}}^{\prime} \mathrm{d} x$ for $v, v^{\prime} \in \mathrm{V}$,

$$
\begin{equation*}
\mathbf{T}=\mathbf{T}_{1} \otimes \mathbf{I} \otimes \cdots \otimes \mathbf{I}+\cdots+\mathbf{I} \otimes \cdots \otimes \mathbf{I} \otimes \mathbf{T}_{1} \tag{3.22}
\end{equation*}
$$

Here we have used the $L_{2}$-orthonormality of $\left\{\psi_{\nu}\right\}_{\nu \in \mathrm{V}}$. Note that $\mathbf{D}$ can be regarded as a two-sided diagonal preconditioner for the operator $\mathbf{T}$, which is unbounded on $\ell_{2}\left(\vee^{d}\right)$. For the operator equation $A u=f$ with $f \in H^{-1}\left(\square_{d}\right)$, the basis representation $\mathbf{A u}=\mathbf{f}$ can be written in the form

$$
\begin{equation*}
\mathbf{D T D u}=\mathbf{D g}, \quad \text { where } \mathbf{g}=\left(f\left(\bigotimes_{i=1}^{d} \psi_{v_{i}}\right)\right)_{v \in \mathrm{~V}^{d}} \tag{3.23}
\end{equation*}
$$

As already observed in Section 3.1.1, operators of the form (3.22) can be represented with all ranks equal to two in Tucker, tensor train or hierarchical tensor formats. However, similarly to Example 3.8, D does not have an explicit low-rank
representation, which reflects the lack of tensor structure of $H_{0}^{1}\left(\square_{d}\right)$ noted in Remark 3.7. Note that since $\left\|\psi_{\nu}\right\|_{H^{1}(0,1)} \approx 2^{|\nu|}$ for the wavelets under consideration,

$$
\begin{equation*}
\mathbf{D}[v, v]=\left(\sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right)^{-1 / 2}=\left(\sum_{i=1}^{d} 2^{2\left|v_{i}\right|}\right)^{-1 / 2} \approx 2^{-\max _{i}\left|v_{i}\right|} \tag{3.24}
\end{equation*}
$$

uniformly for all $v \in \mathrm{~V}^{d}$. Approaches for obtaining efficient low-rank approximations of $\mathbf{D}$ are considered in Section 4.

Remark 3.11. For $\left\{\Psi_{\nu}\right\}_{\nu \in \mathrm{V}^{d}}$ constructed from $L_{2}$-orthonormal univariate wavelets as above, by (3.8), for any $u, \tilde{u} \in H_{0}^{1}\left(\square_{d}\right)$ with

$$
u=\sum_{\nu \in \mathrm{V}^{d}} \mathbf{u}[v] \Psi_{\nu}, \quad \tilde{u}=\sum_{v \in \mathrm{~V}^{d}} \tilde{\mathbf{u}}[v] \Psi_{\nu}
$$

we have in particular

$$
c_{1}\|\mathbf{u}-\tilde{\mathbf{u}}\|_{\ell_{2}} \leq\|u-\tilde{u}\|_{H^{1}} \leq C_{1}\|\mathbf{u}-\tilde{\mathbf{u}}\|_{\ell_{2}}
$$

This means that rank truncations of coefficient tensors can be performed by standard HSVD-based procedures with controlled $H^{1}$-error. Without the use of a Riesz basis, error control in $H^{1}$ could in principle also be achieved by different means, for instance direct optimization without the aid of the SVD (which is in general difficult to do reliably), or by alternative approaches considered in Ali and Nouy (2020a,b) which, however, involve additional restrictions.

The preconditioning of operators provided by Riesz basis representations also yields a proportionality between solution error in $H^{1}$ and residual: with $A$ and $\mathbf{A}$ as in (3.21), if $u$ is the solution of $A u=f \in H^{-1}\left(\square_{d}\right)$ with basis representation $\mathbf{A u}=\mathbf{f}$, then

$$
\|u-\tilde{u}\|_{H^{1}} \approx\|\mathbf{A}(\mathbf{u}-\tilde{\mathbf{u}})\|_{\ell_{2}}=\|\mathbf{f}-\mathbf{A} \tilde{\mathbf{u}}\|_{\ell_{2}}
$$

with constants independent of $d$. This means that the total approximation error in $H^{1}$ with respect to exact solutions of the considered PDE can be estimated from finitely supported approximations of the residuals $\mathbf{f}-\mathbf{A} \tilde{\mathbf{u}}$, which generally have infinite support in $\vee^{d}$.

Remark 3.12. Adaptively refined Riesz basis expansions are used in combination with hierarchical tensor representations in infinite-dimensional Hilbert spaces in Bachmayr (2012a), Bachmayr and Dahmen (2015) and subsequently in Bachmayr and Dahmen (2016b) and Ali and Urban (2020). A similar formalism focused on continuous variables rather than basis expansions has been developed independently in Bigoni, Engsig-Karup and Marzouk (2016) based on the tensor train format, termed functional tensor train representation. This is combined with heuristic adaptivity based on spectral basis functions in Gorodetsky, Karaman and Marzouk (2019); see also Dolgov, Kressner and Strössner (2021).

### 3.3. Occupation number representations

In the Schrödinger equations (1.2) and (1.3), solutions describing systems with $N$ particles are functions depending on $N$ particle coordinates. We additionally need to ensure symmetry of solutions under exchange of particle coordinates when considering bosonic particles (due to their indistinguishability), or antisymmetry in case of fermionic particles (corresponding to the Pauli exclusion principle).

One case of particular importance in chemistry is the electronic Schrödinger equation, as a special case of the eigenvalue problem (1.3) for a Hamiltonian operator $H$ of the basic form

$$
\begin{equation*}
H=-\frac{1}{2} \Delta+\sum_{i} V_{1}\left(x_{i}\right)+\sum_{i<j} V_{2}\left(x_{i}, x_{j}\right) \tag{3.25}
\end{equation*}
$$

where $x_{i}, i=1, \ldots, N$ are the coordinates of each particle and $V_{2}$ is a symmetric function. Hamiltonians of this structure arise frequently in quantum physics. In the case of the electronic Schrödinger case, $V_{2}$ takes the form of a Coulomb interaction between electrons, and $V_{1}$ denotes external Coulomb potentials of atomic nuclei. The associated eigenfunctions, referred to as wavefunctions, describe the quantum states of electrons for given nuclear positions. Unless the model under consideration has certain special features (such as the presence of magnetic fields), the wavefunctions are real-valued. Since electrons are fermions, they are required to be antisymmetric under exchange of each pair of electron coordinates. Here each electron has coordinates in $\Omega=\mathbb{R}^{3} \times \mathbb{Z}_{2}$, where the binary degree of freedom in $\mathbb{Z}_{2}$ accounts for electron spin.

The hierarchical tensor format is not well suited to the direct representation of antisymmetric functions (Hackbusch 2018). However, this issue can be circumvented by an indirect parametrization provided by the formalism of second quantization that is widely used in quantum physics and chemistry. The starting point for this formalism is an orthonormal basis $\left\{\phi_{\nu}\right\}_{\nu \in \mathbb{N}}$ of $L_{2}(\Omega)$. These basis functions are also referred to as orbitals. From these functions, we form antisymmetrized tensor products, called Slater determinants, for all $v_{1}, \ldots, v_{N} \in \mathbb{N}$ with $v_{1}<v_{2}<\cdots<v_{N}$,

$$
\begin{aligned}
\Phi_{\nu_{1}, \ldots, v_{N}}\left(x_{1}, \ldots, x_{N}\right) & =\frac{1}{\sqrt{N!}} \sum_{\pi \in S_{N}} \operatorname{sgn}(\pi) \prod_{i=1}^{N} \phi_{v_{i}}\left(x_{\pi(i)}\right) \\
& =\frac{1}{\sqrt{N!}} \operatorname{det}\left(\phi_{\nu_{i}}\left(x_{j}\right)\right)_{i, j=1, \ldots, N}
\end{aligned}
$$

for $x_{1}, \ldots, x_{N} \in \Omega$, which provide an orthonormal basis of the subspace $\mathcal{V}_{N}$ of antisymmetric functions in $L_{2}\left(\Omega^{N}\right)$.

We now restrict to a subset of $K \geq N$ orbitals $\phi_{1}, \ldots, \phi_{K}$. These finitely many orbitals may also be optimized for the problem under consideration; see for instance Krumnow, Veis, Legeza and Eisert (2016). The corresponding Slater determinants
are an orthonormal basis of the subspace

$$
\mathcal{V}_{N}^{K}=\operatorname{span}\left\{\Phi_{v}: v \in\{1, \ldots, K\}^{N}, v_{1}<\cdots<v_{N}\right\} \subset \mathcal{V}_{N}
$$

We next introduce an injective mapping from $\mathcal{V}_{N}^{K}$ to the space

$$
\mathcal{F}^{K}=\bigotimes_{i=1}^{K} \mathbb{R}^{2}
$$

To this end, let $\mathbf{e}_{0}=(1,0)^{\top}, \mathbf{e}_{1}=(0,1)^{\top} \in \mathbb{R}^{2}$. For each single Slater determinant $\Phi_{v} \in \mathcal{V}_{N}^{K}$ and scalar coefficient $c_{\nu}$, we define

$$
c_{\nu} \Phi_{v} \mapsto c_{v} \bigotimes_{i=1}^{K} \mathbf{e}_{b_{i}(v)}, \quad \text { where } b_{i}(v)= \begin{cases}1 & i \in\left\{v_{1}, \ldots, v_{d}\right\} \\ 0 & \text { otherwise }\end{cases}
$$

and extend this by linearity to a mapping from $\mathcal{V}_{N}^{K}$ to $\mathcal{F}^{K}$. We obtain an identification of linear combinations of Slater determinants with tensors in $\mathcal{F}^{K}$ that encode the occupations of orbitals in each Slater determinant. This representation is thus also called occupation number representation. It has the advantage that the elements of $\mathcal{F}^{K}$, referred to as occupation numbers, do not need to satisfy any antisymmetry requirement.

Operators on $\mathcal{F}^{K}$ can be expressed in terms of the annihilation operators $\mathbf{a}_{i}$ and creation operators $\mathbf{a}_{i}^{*}$ for $i=1, \ldots, K$, which are given by

$$
\mathbf{a}_{i}=\mathbf{S}^{\otimes(i-1)} \otimes \mathbf{J} \otimes \mathbf{I}^{\otimes(K-i)}, \quad \mathbf{a}_{i}^{*}=\mathbf{S}^{\otimes(i-1)} \otimes \mathbf{J}^{\top} \otimes \mathbf{I}^{\otimes(K-i)}
$$

in terms of the components

$$
\mathbf{I}=\left(\begin{array}{ll}
1 & 0  \tag{3.26}\\
0 & 1
\end{array}\right), \quad \mathbf{J}=\left(\begin{array}{ll}
0 & 1 \\
0 & 0
\end{array}\right), \quad \mathbf{S}=\left(\begin{array}{cc}
1 & 0 \\
0 & -1
\end{array}\right)
$$

In the occupation number representation, antisymmetry of Slater determinants is reflected by the anticommutation relations satisfied by annihilation and creation operators: for $i, j \in\{1, \ldots, K\}$,

$$
\begin{equation*}
\mathbf{a}_{i} \mathbf{a}_{j}^{*}+\mathbf{a}_{j}^{*} \mathbf{a}_{i}=\delta_{i j}, \quad \mathbf{a}_{i}^{*} \mathbf{a}_{j}^{*}+\mathbf{a}_{j}^{*} \mathbf{a}_{i}^{*}=\mathbf{a}_{i} \mathbf{a}_{j}+\mathbf{a}_{j} \mathbf{a}_{i}=0 \tag{3.27}
\end{equation*}
$$

Making use of the antisymmetry properties of Slater determinants, which lead to the so-called Slater-Condon rules (Helgaker, Jørgensen and Olsen 2000, Chap. 1), for $H$ as in (3.25) we obtain the representation

$$
\begin{equation*}
\mathbf{H}=\sum_{i, j=1}^{K} t_{i j} \mathbf{a}_{i}^{*} \mathbf{a}_{j}+\sum_{i, j, k, l=1}^{K} v_{i j k l} \mathbf{a}_{i}^{*} \mathbf{a}_{j}^{*} \mathbf{a}_{k} \mathbf{a}_{l} \tag{3.28}
\end{equation*}
$$

with the coefficients

$$
\begin{aligned}
t_{i j} & =\int \varphi_{i}(x)\left\{-\frac{1}{2} \Delta+V_{1}(x)\right\} \varphi_{j}(x) \mathrm{d} x \\
v_{i j k l} & =\int \varphi_{i}(x) \varphi_{j}(y) V_{2}(x, y) \varphi_{k}(x) \varphi_{l}(y) \mathrm{d} x \mathrm{~d} y
\end{aligned}
$$

Note that the number of particles is not fixed a priori for elements of $\mathcal{F}^{K}$ and thus needs to be added as a constraint in this formulation. The occupation numbers $\mathbf{x} \in \mathcal{F}^{K}$ that represent wavefunctions with $N$ particles are precisely the eigenfunctions with eigenvalue $N$ of the particle number operator

$$
\begin{equation*}
\mathbf{P}=\sum_{i=1}^{K} \mathbf{a}_{i}^{*} \mathbf{a}_{i} \tag{3.29}
\end{equation*}
$$

that is, they lie in the subspace

$$
\mathcal{F}_{N}^{K}=\left\{\mathbf{x} \in \mathcal{F}^{K}: \mathbf{P} \mathbf{x}=N \mathbf{x}\right\}
$$

The Ritz-Galerkin approximation in $\mathcal{V}_{N}^{K}$ of the original eigenvalue problem for $H$ is then equivalent to

$$
\begin{equation*}
\mathbf{H} \mathbf{x}=\lambda \mathbf{x}, \quad \text { where } \mathbf{x} \in \mathcal{F}_{N}^{K}, \mathbf{x} \neq 0 \tag{3.30}
\end{equation*}
$$

and in particular, the approximation in $\mathcal{V}_{N}^{K}$ of the lowest eigenvalue of $H$ is characterized by

$$
\min _{0 \neq \mathbf{x} \in \mathcal{F}_{N}^{K}} \frac{\langle\mathbf{H} \mathbf{x}, \mathbf{x}\rangle}{\langle\mathbf{x}, \mathbf{x}\rangle} .
$$

The problem (3.30) in occupation number representation is directly amenable to approximations in hierarchical tensor format, in this context typically referred to as tree tensor networks. A frequent choice is tensor trains, in this context known as matrix product states, corresponding to representing or approximating the sought $\mathbf{x} \in \mathcal{F}^{K}$ in the form

$$
\begin{equation*}
\mathbf{x}=\mathbf{X}_{1} \bowtie \cdots \bowtie \mathbf{X}_{K} \tag{3.31}
\end{equation*}
$$

as in (2.17). For a more detailed survey of this approach, which has been successfully applied in quantum chemistry and many other applications of quantum physics, we refer to Szalay et al. (2015). Note that in this formulation the dimensionality $N$ of the original problem enters only indirectly, and the main parameters determining the computational complexity are the ranks and the number of orbitals $K$.

Remark 3.13. It is clear that $\mathbf{H}$ can be represented with ranks $O\left(K^{4}\right)$, since each operator $\mathbf{a}_{i}^{*} \mathbf{a}_{j}^{*} \mathbf{a}_{k} \mathbf{a}_{l}$ is of rank one. However, this is not sharp: as observed in Crosswhite and Bacon (2008) and Keller, Dolfi, Troyer and Reiher (2015), H can generically be represented with ranks $O\left(K^{2}\right)$. For a detailed mathematical treatment, see Bachmayr, Götte and Pfeffer (2022); further rank reductions are possible, for instance with appropriately localized orbitals.

Remark 3.14. Since (3.31) yields representations of general elements of $\mathcal{F}^{K}$, the constraint to the $N$-particle eigenspace $\mathcal{F}_{N}^{K}$ of $\mathbf{P}$ needs to be handled appropriately. There are certain cases where this constraint does not need to be enforced, since it is automatically satisfied by the sought eigenvectors of $\mathbf{H}$. In general, however, $\mathbf{x} \in \mathcal{F}^{K}$ is in $\mathcal{F}_{N}^{K}$ if and only if it has a representation (3.31) where each of the cores $\mathbf{X}_{i}, i=1, \ldots, K$, has a certain block-sparse structure. In a physics context, this observation goes back to Singh, Pfeifer and Vidal (2011). A mathematical proof, extending this to more general operators than $\mathbf{P}$ with similar structure, is given in Bachmayr et al. (2022). As a consequence, membership in $\mathcal{F}_{N}^{K}$ can be enforced by working directly on the non-zero blocks of such block-sparse representations. The operator $\mathbf{H}$ can be applied to such block-sparse representations in a matrix-free manner that preserves the block-sparse structure (Bachmayr et al. 2022, Sec. 5.2).

### 3.4. Multilevel tensorized representations

We now turn to a different way of approximating solutions of PDEs with low-rank tensors that is aimed at highly compressed approximations of simple discretizations. Here grid vectors are efficiently encoded in low-rank format such that the use of extremely fine uniform grids becomes feasible. For PDEs on domains of moderate dimensionality, this leads to generic approximations for problems with singularities or multiscale structure that otherwise require specialized approaches. In this case tensor modes correspond to scales in the approximation problem.

The basic idea is that hierarchical tensor representations can be applied to arbitrary vectors when these are reinterpreted as high-order tensors. This is achieved by a suitable encoding of indices. For instance, for $L \in \mathbb{N}$, we introduce the isomorphism

$$
\begin{equation*}
\mathcal{J}:\{0,1\}^{L} \rightarrow\left\{1, \ldots, 2^{L}\right\}, \quad\left(b_{1}, \ldots, b_{L}\right) \mapsto \sum_{\ell=1}^{L} b_{\ell} 2^{L-\ell}+1 \tag{3.32}
\end{equation*}
$$

For a given vector $\mathbf{v}$ of length $2^{L}$, we define the tensor $\mathbf{u} \in \bigotimes_{\ell=1}^{L} \mathbb{R}^{2}$ indexed by $\{0,1\}^{L}$ as

$$
\begin{equation*}
\mathbf{u}\left[b_{1}, \ldots, b_{L}\right]=\mathbf{v}\left[\mathcal{J}\left(b_{1}, \ldots, b_{L}\right)\right], \quad b \in\{0,1\}^{L} \tag{3.33}
\end{equation*}
$$

In an analogous manner, a vector in $\mathbb{R}^{N}$ with $N=\prod_{\ell=1}^{L} n_{\ell}$ can be mapped to a tensor of mode sizes $n_{1} \times \cdots \times n_{L}$.

We can then approximate $\mathbf{u}$ in a suitable tensor format to obtain a compressed version of $\mathbf{v}$. Specifically, we can obtain a tensor train representation $V=\left(\mathbf{V}_{1}, \ldots, \mathbf{V}_{L}\right)$ of $\mathbf{u}$ of the form

$$
\begin{equation*}
\mathbf{u}=\tau(\mathrm{V})=\mathbf{V}_{1} \bowtie \mathbf{V}_{2} \bowtie \cdots \bowtie \mathbf{V}_{L} \tag{3.34}
\end{equation*}
$$

Since individual modes are associated to scales, such a representation can be interpreted as a type of multilevel TT representation of the underlying function. In the
literature this is referred to as quantized tensor train (QTT) or as tensorized representation, introduced by Khoromskij (2011) and Grasedyck (2010). Precursors can be found in similar ideas based on canonical tensor representations by Tyrtyshnikov (2003) and in numerical observations on tensor train representations of matrices by Oseledets (2009a). Detailed accounts of these techniques are given in Khoromskij (2018) and Hackbusch (2019, Chap. 14).

Example 3.15. Consider the vector $\mathbf{v} \in \mathbb{R}^{N}$ with $N=2^{L}, L \in \mathbb{N}$, defined by $\mathbf{v}[j]=\exp \left(\lambda x_{j}\right)$ with $\lambda \in \mathbb{C}$ and $x_{j}=(j-1) / N$ for $j=1, \ldots, N$, corresponding to grid values of an exponential function on [0, 1]. Then, as observed by Khoromskij (2011), for the tensor $\mathbf{u}$ defined in (3.33), we have the rank-one representation

which requires only $2 L$ coefficients as opposed to the $2^{L}$ entries of the vector $\mathbf{v}$. As an immediate consequence, grid values of trigonometric functions have representations (3.34) of ranks at most two.

Remark 3.16. A variety of further results on representations and approximations of more general functions in the multilevel form (3.34) have been obtained, including the following.
(i) Values of polynomials of degree $p$ on uniform grids can be represented in analogous form with ranks bounded by $p+1$; see Grasedyck (2010) and Khoromskij (2011).
(ii) In Grasedyck (2010), this representation of polynomials is applied on subintervals in combination with local refinement to obtain approximations of the form (3.34) of functions on intervals with isolated singularities. These approximations converge exponentially, uniformly on an equispaced grid of points, with respect to the number of parameters in the representation, similarly to $h p$-type piecewise polynomial approximations.
(iii) Explicit representations (3.34) of continuous multivariate functions, with functions as components, have been constructed in Oseledets (2013).
(iv) A general formulation of tensorized approximations on a function space level has been developed in Ali and Nouy (2023), where convergence rates for functions with Besov regularity are also obtained.

Various transformations with particular structure have been shown to be efficiently realizable for multilevel representations as in (3.34), including convolutions (Hackbusch 2011), Toeplitz matrices (Kazeev, Khoromskij and Tyrtyshnikov 2013a) and wavelet transforms (Khoromskij and Miao 2014).

Finding representation (3.34) as approximate solutions of differential equations requires suitable representations of discretized differential operators in a compatible
format. In the case of the second derivative on an interval, we have the following basic result.

Example 3.17. As shown by Kazeev and Khoromskij (2012), the matrix $\boldsymbol{\Delta}_{2}$ defined in (3.4) that appears in the finite difference or finite element discretization of the second derivative can be represented as $\mathbf{L}_{1} \bowtie \cdots \bowtie \mathbf{L}_{L}$, where

$$
\mathbf{L}_{1}=\left[\begin{array}{lll}
\mathbf{I} & \mathbf{J}^{\top} & \mathbf{J}
\end{array}\right], \quad \mathbf{L}_{2}=\cdots=\mathbf{L}_{L-1}=\left[\begin{array}{ccc}
\mathbf{I} & \mathbf{J}^{\top} & \mathbf{J} \\
0 & \mathbf{J} & 0 \\
0 & 0 & \mathbf{J}^{\top}
\end{array}\right], \quad \mathbf{L}_{L}=\left[\begin{array}{c}
2 \mathbf{I}-\mathbf{J}-\mathbf{J}^{\top} \\
-\mathbf{J} \\
-\mathbf{J}^{\top}
\end{array}\right]
$$

with component matrices as in (3.26). This representation can be applied to grid vectors represented in the format (3.34).

Concerning the operator representation condition number, as verified in Bachmayr and Kazeev (2020, Prop. 4.10), the above representation $L=\left(\mathbf{L}_{1}, \ldots, \mathbf{L}_{L}\right)$ satisfies

$$
\max _{\ell=1, \ldots, L} \operatorname{oprcond}_{\ell}(\mathrm{L}) \approx 2^{2 L}
$$

This means that applying the representation $L$ to some $V$ as in (3.34) may generate redundancies leading to cancellations (as in Example 2.43) with a loss of relative precision by a factor proportional to $2^{2 L}$. This is indeed confirmed by numerical tests (Bachmayr and Kazeev 2020, Ex. 4.3), which means that the direct application of $L$ is limited to moderately large $L$ (where with double precision, all digits of accuracy are lost when $L \approx 25$ ). This issue, which affects such representations of differential operators in general, can be dealt with by suitable preconditioners that also improve the representation condition number of the representation, rather than just the usual condition number of the represented matrix. A construction of such preconditioners is considered in Section 4.5.

Remark 3.18. The concepts outlined above for univariate functions can also be applied in the multivariate case. For problems in $d$ dimensions, one can devise an isomorphism analogously to (3.32) but depending on parameters $\left(b_{i, 1}, \ldots, b_{i, L}\right)$ for $i=1, \ldots, d$. There are two canonical ways of ordering this larger set of parameters in an isomorphism as in (3.32): either the modes in the TT representation are ordered in the form

$$
\left(b_{1,1}, \ldots, b_{1, L}, b_{2,1}, \ldots, b_{2, L}, \ldots, b_{d, 1}, \ldots, b_{d, L}\right)
$$

which has been used, for instance, in Khoromskij and Oseledets (2011), closely related to a standard tensor representation in $d$ dimensions where multilevel TT representations are applied in each mode (Dolgov and Khoromskij 2013); or the modes are arranged in the so-called transposed ordering

$$
\left(b_{1,1}, b_{2,1}, \ldots, b_{d, 1}, b_{1,2}, b_{2,2}, \ldots, b_{d, 2}, \ldots, b_{1, L}, \ldots, b_{d, L}\right)
$$

used, for instance, by Kazeev and Schwab (2018), which corresponds to a parametrization in terms of a hierarchy of $d$-dimensional patches. This latter form has
proved to be particularly suitable for efficient approximations in dimensions $d=2$ or $d=3$; see also Ali and Nouy (2021).

Corresponding representations of operators in the general $d$-dimensional case are studied in Khoromskij and Oseledets (2011), Dolgov and Khoromskij (2013), Kazeev, Reichmann and Schwab (2013b) and Bachmayr and Kazeev (2020). Applications to elliptic boundary value problems are considered in further detail in Section 5.5.3.

## 4. Approximate inverses and preconditioning of elliptic operators

For bounded elliptic operators $A: V \rightarrow V^{\prime}$ on Hilbert spaces $V$, where we focus in particular on the case $V=H_{0}^{1}\left(\square_{d}\right)$ as in (3.9), in this section we consider two closely related tasks. The first is finding approximate inverses of $A$, that is, approximating $A^{-1}$ by low-rank expressions in suitable operator norms. Such approximations can be used for the direct construction of approximate solutions in low-rank format; however, they are applicable only to operators with particular structure such as the Laplacian-type form (1.6).

The second is the construction of preconditioners, where for a discretization or basis representation $\mathbf{A}$ of $A$, we aim to find $\mathbf{C}$ such that for all $\mathbf{v}$,

$$
c\langle\mathbf{C v}, \mathbf{v}\rangle \leq\left\langle\mathbf{A}^{-1} \mathbf{v}, \mathbf{v}\right\rangle \leq C\langle\mathbf{C} \mathbf{v}, \mathbf{v}\rangle
$$

with uniform constants $c, C>0$, where $\mathbf{C}$ can be applied efficiently. Although this second aim of establishing a spectral equivalence can be accomplished by sufficiently accurate approximate inverses, more direct constructions of preconditioners can generally provide more efficiency, in particular concerning more favourable representation ranks. Low-rank preconditioners are again obtained for operators with Laplacian-type structure, but are equally applicable to more general elliptic operators.

### 4.1. Exponential sums

Let us first consider the case of a tensor product Hilbert space $H=\bigotimes_{i=1}^{d} H_{i}$, with $A: \operatorname{dom}(A) \subseteq H \rightarrow H$ a potentially unbounded densely defined, self-adjoint, positive definite operator with Laplacian-type structure as in (1.6), that is,

$$
A=\sum_{i=1}^{d} \bigotimes_{j=1}^{d} B_{i, j}, \quad B_{i, j}= \begin{cases}A_{i}, & i=j,  \tag{4.1}\\ \mathrm{I}, & i \neq j,\end{cases}
$$

with densely defined, self-adjoint and positive definite $A_{i}: \operatorname{dom}\left(A_{i}\right) \subseteq H_{i} \rightarrow H_{i}$. An instance of such an operator would be the Laplacian regarded as an unbounded operator on $L_{2}\left(\square_{d}\right)$, but we will also apply the following to operators on $\ell_{2}\left(\mathbb{N}^{d}\right)$ arising from basis representations. The observation that is fundamental henceforth
is that for $\alpha \geq 0$,

$$
e^{-\alpha A}=\bigotimes_{i=1}^{d} e^{-\alpha A_{i}}
$$

The basic strategy is now to find $\omega_{k}, \alpha_{k}>0$ for $k=1, \ldots, r$ such that

$$
\begin{equation*}
t^{-1} \approx \sum_{k=1}^{r} \omega_{k} e^{-\alpha_{k} t} \quad \text { for } t \in \sigma(A) \tag{4.2}
\end{equation*}
$$

in an appropriate sense, and by applying this exponential sum approximation to $A$ via its spectral decomposition to conclude

$$
A^{-1} \approx \sum_{k=1}^{r} \omega_{k} e^{-\alpha_{k} A}=\sum_{k=1}^{r} \omega_{k} \bigotimes_{i=1}^{d} e^{-\alpha_{k} A_{i}}
$$

which provides an approximation in terms of a sum of $r$ tensor product operators. By the continuous spectral calculus for self-adjoint operators (or, in the finitedimensional case, the eigenvalue decomposition of $A$ ), we have the following estimates.

Proposition 4.1. Let $A$ be as in (4.1) and let $f, \tilde{f}: \mathbb{R}^{+} \rightarrow \mathbb{R}^{+}$be continuous functions on an interval containing the spectrum $\sigma(A)$.
(i) Under these conditions, we have the estimate

$$
\|f(A)-\tilde{f}(A)\|_{H \rightarrow H} \leq \sup _{t \in \sigma(A)}|f(t)-\tilde{f}(t)| .
$$

(ii) Assume that for some $\delta \in(0,1)$ we have $|f(t)-\tilde{f}(t)| \leq \delta f(t)$ for all $t \in \sigma(A)$. Then

$$
(1-\delta)\langle f(A) v, v\rangle \leq\langle\tilde{f}(A) v, v\rangle \leq(1+\delta)\langle f(A) v, v\rangle \quad \text { for all } v \in H
$$

For related results under slightly different assumptions, see also Hackbusch (2014, Sec. 5.9).

### 4.1.1. Inverse Laplace transform and quadrature approximation

A first construction principle for exponential sum approximations as in (4.2) for a given function $f$, such as $f(t)=t^{-1}$, is applying a quadrature rule to its representation as a Laplace transform of a function $g$,

$$
f(t)=\int_{0}^{\infty} e^{-s t} g(s) \mathrm{d} s
$$

This can be done by sinc quadrature methods: after applying a suitable substitution $\sigma: \mathbb{R} \rightarrow \mathbb{R}^{+}$, we use the trapezoidal rule on $\mathbb{R}$ with step size $\tau>0$ to obtain the approximations

$$
\begin{equation*}
f(t)=\int_{-\infty}^{\infty} e^{-\sigma(x) t} g(\sigma(x)) \sigma^{\prime}(x) \mathrm{d} x \approx \tau \sum_{k \in \mathbb{Z}} e^{-\sigma(k \tau) t} g(\sigma(k \tau)) \sigma^{\prime}(k \tau) \tag{4.3}
\end{equation*}
$$

We then exploit decay properties of the integrand to truncate the sum on the right.
For $f(t)=t^{-\lambda}$ with $\lambda>0$, the following result of this kind, which provides the conditions of Proposition 4.1(ii), is obtained in Beylkin and Monzón (2010); see also Beylkin and Monzón (2005) and Scholz and Yserentant (2017).

Theorem 4.2. Let $\tau, \lambda>0$ and

$$
\delta(\tau, \lambda)=\frac{1}{\Gamma(\lambda)} \sum_{\substack{n \in \mathbb{Z} \\ n \neq 0}}|\Gamma(\lambda+2 \pi \mathrm{i} n / \tau)|
$$

Then, for all $t>0$,

$$
\begin{equation*}
\left|\frac{1}{t^{\lambda}}-\frac{\tau}{\Gamma(\lambda)} \sum_{k \in \mathbb{Z}} e^{\lambda k \tau} \exp \left(-e^{k \tau} t\right)\right| \leq \frac{\delta(\tau, \lambda)}{t^{\lambda}} \tag{4.4}
\end{equation*}
$$

and for each $\hat{\delta} \in(0,1]$, if

$$
\tau \leq \frac{2 \pi}{\ln 3+\lambda|\ln (\cos 1)|+|\ln \hat{\delta}|}
$$

then $\delta(\tau, \lambda) \leq \hat{\delta}$.
We thus obtain approximations on all of $(0, \infty)$ with infinitely many exponential terms. In the restriction to finitely many summands, we arrive at approximations with prescribed relative error on finite intervals as required in Proposition 4.1(ii). The following result is adapted from Bachmayr and Dahmen (2016b).

Corollary 4.3. Let $\lambda \in(0,1], \delta=\delta_{0}+\eta<1$ with $\delta_{0}, \eta \in(0,1)$, and $T>1$. With

$$
\begin{align*}
\tau & =\frac{2 \pi}{\ln 3+\lambda|\ln (\cos 1)|+\left|\ln \left(\delta_{0} / 2\right)\right|}  \tag{4.5}\\
m & =\left\lceil\tau^{-1} \ln \left|\ln \left(\delta_{0} / 2\right)\right|\right\rceil, \quad n=\left\lceil\tau^{-1}\left(\lambda^{-1}|\ln \eta|+\ln T\right)\right\rceil
\end{align*}
$$

define the truncated exponential sums

$$
\begin{aligned}
S_{\lambda, \delta_{0}}(t) & =\frac{\tau}{\Gamma(\lambda)} \sum_{k=-\infty}^{m} e^{\lambda k \tau} \exp \left(-e^{k \tau} t\right), \\
S_{\lambda, \delta_{0}, \eta, T}(t) & =\frac{\tau}{\Gamma(\lambda)} \sum_{k=-n}^{m} e^{\lambda k \tau} \exp \left(-e^{k \tau} t\right) .
\end{aligned}
$$

Then

$$
\begin{align*}
\left|t^{-\lambda}-S_{\lambda, \delta_{0}}(t)\right| \leq \delta_{0} t^{-\lambda}, & \text { for all } t \in[1, \infty),  \tag{4.6a}\\
\left|S_{\lambda, \delta_{0}}(t)-S_{\lambda, \delta_{0}, \eta, T}(t)\right| \leq \eta t^{-\lambda}, & \text { for all } t \in[1, T]  \tag{4.6b}\\
\left|t^{-\lambda}-S_{\lambda, \delta_{0}, \eta, T}(t)\right| \leq \delta t^{-\lambda}, & \text { for all } t \in[1, T] \tag{4.6c}
\end{align*}
$$

Proof. For $m$ as in (4.5),

$$
\tau \sum_{k>m} e^{\lambda k \tau} \exp \left(-e^{k \tau} t\right) \leq \int_{m \tau}^{\infty} \exp \left(-t e^{x}+\lambda x\right) \mathrm{d} x=\frac{1}{t^{\lambda}} \int_{t \exp (m \tau)}^{\infty} y^{\lambda-1} e^{-y} \mathrm{~d} y
$$

where we have used monotonicity and the substitution $x=\ln (y / t)$. Since $\lambda \leq 1$, for $t \geq 1$ we have

$$
\int_{t \exp (m \tau)}^{\infty} y^{\lambda-1} e^{-y} \mathrm{~d} y \leq \exp \left(-e^{m \tau}\right)
$$

and thus (4.6a) follows from Theorem 4.2 and the choice of $\tau$ and $m$. With the bound

$$
\tau \sum_{k<-n} e^{\lambda k \tau} \exp \left(-e^{k \tau} t\right) \leq \int_{n \tau}^{\infty} e^{-\lambda x} \mathrm{~d} x=\frac{1}{\lambda} e^{-n \lambda \tau}
$$

and using $\lambda \Gamma(\lambda)=\Gamma(\lambda+1) \in(0,1]$, we obtain the two further estimates.
Note that with $\tau, m, n$ as in (4.5), for the cases $\lambda=1$ and $\lambda=\frac{1}{2}$ of particular interest we obtain the exponential sums

$$
S_{1, \delta_{0}, \eta, T}=\tau \sum_{k=-n}^{m} e^{k \tau} \exp \left(-e^{k \tau} t\right), \quad S_{1 / 2, \delta_{0}, \eta, T}=\frac{\tau}{\sqrt{\pi}} \sum_{k=-n}^{m} e^{(1 / 2) k \tau} \exp \left(-e^{k \tau} t\right)
$$

with $m+n+1$ terms. The parameters $\tau, m$ and $n$ can be optimized further for particular cases; see Scholz and Yserentant (2017, Sec. 5).

Results analogous to Corollary 4.3 are obtained in Bachmayr and Dahmen $(2016 a, b)$. The construction used there, however, is slightly different from the one in (4.4) and based on general results on sinc quadrature by Stenger (1993). These results can be applied whenever the integrands in (4.3) are holomorphic on $\{z \in \mathbb{C}:|\operatorname{Im} z|<\eta\}$ for some $\eta>0$ and integrable on the boundary of this strip. This technique can also be used for obtaining approximations with bounds on $\sup _{t \in[a, b]}\left|t^{-\lambda}-S_{r}(t)\right|$ for a given interval $[a, b]$; see Hackbusch and Khoromskij (2006a) and Hackbusch (2015a, Sec. D.4). However, stronger results for uniform approximation can be obtained by a different approach, which we consider next.

### 4.1.2. Uniform best approximation

The application of Proposition 4.1(i) requires uniform approximations on intervals. To this end, rather than aiming for explicit constructions, we can consider best approximations by exponential sums. These are in general not analytically computable, but lead to stronger results. The existence of such best approximations for a class of functions including $f(t)=t^{-\lambda}$ for any $\lambda>0$ is shown in Braess (1986). The following result on exponential convergence is proved in Braess and Hackbusch $(2005,2009)$.

Theorem 4.4. For each $R>1, \lambda>0$ and $r \in \mathbb{N}$, there exist $\omega_{k}, \alpha_{k}>0$, $k=1, \ldots, r$, such that

$$
\sup _{t \in[1, R]}\left|\frac{1}{t^{\lambda}}-\sum_{k=1}^{r} \omega_{k} e^{-\alpha_{k} t}\right| \leq 2^{\lambda} 8 \exp \left(-\frac{\pi^{2} r}{\ln (8 R)}\right)
$$

In contrast to Theorem 4.2, we thus obtain a bound on the absolute rather than the relative approximation error. By choosing suitable $R$ and rescaling, from the above we obtain the following result on intervals $[a, \infty)$ with $a>0$; see Braess and Hackbusch (2009) and Dahmen, DeVore, Grasedyck and Süli (2016).

Corollary 4.5. Let $a>0$. Then, for each $r \in \mathbb{N}$, there exist $\omega_{k}, \alpha_{k}>0$, $k=1, \ldots, r$ such that with $\hat{S}_{r}(t)=\sum_{k=1}^{r} \omega_{k} e^{-\alpha_{k} t}$,

$$
\sup _{t \in[a, \infty)}\left|\frac{1}{t}-\hat{S}_{r}(t)\right| \leq \frac{16}{a} \exp (-\pi \sqrt{r})
$$

Such best approximations are computed in Hackbusch (2005) by a Remez-type algorithm based on alternation properties of uniform best approximation errors.

### 4.2. Low-rank approximate inverses of Laplacian-type operators

For the finite difference discretization matrix $\mathbf{A}_{\mathrm{FD}}$ as in (3.3), combining Corollary 4.5 and Proposition 4.1(i), we obtain a bound in the spectral norm and thus

$$
\begin{equation*}
\left\|\mathbf{A}_{\mathrm{FD}}^{-1} \mathbf{f}-\hat{S}_{r}\left(\mathbf{A}_{\mathrm{FD}}\right) \mathbf{f}\right\|_{2} \leq \frac{16}{\min \sigma\left(\mathbf{A}_{\mathrm{FD}}\right)} \exp (-\pi \sqrt{r})\|\mathbf{f}\|_{2} \tag{4.7}
\end{equation*}
$$

for all right-hand sides $\mathbf{f}$, where

$$
\hat{S}_{r}\left(\mathbf{A}_{\mathrm{FD}}\right)=\sum_{k=1}^{r} \omega_{k} \bigotimes_{i=1}^{d} \exp \left(-\alpha_{k} h^{-2} \boldsymbol{\Delta}_{2}\right)
$$

Results of this type, including further matrix functions, are shown in Grasedyck (2004), Gavrilyuk, Hackbusch and Khoromskij (2005), Hackbusch and Khoromskij (2006b) and Khoromskij (2009). As proposed in these works, the required matrix exponentials $\exp \left(-\alpha_{k} h^{-2} \boldsymbol{\Delta}_{2}\right)$ for each $k$ can be evaluated, for instance, by diagonalization or by hierarchical matrix approximations (Hackbusch 2015a).

However, estimates as in (4.7) do not reflect the natural norms of the problem. This can also be seen on the level of function spaces: for $A=-\Delta: H_{0}^{1}\left(\square_{d}\right) \rightarrow$ $H^{-1}\left(\square_{d}\right)$, considering $A$ as an unbounded operator on $\square_{d}$ of the form (4.1), from Proposition 4.1(i) we also obtain

$$
\left\|A^{-1} f-\hat{S}_{r}(A) f\right\|_{L_{2}\left(\square_{d}\right)} \leq \frac{16}{d \pi^{2}} \exp (-\pi \sqrt{r})\|f\|_{L_{2}\left(\square_{d}\right)} \quad \text { for all } f \in L_{2}\left(\square_{d}\right)
$$

Here the norm for $f$ is too strong ( $L_{2}$ rather than $H^{-1}$ ), whereas the approximation error is measured in a norm that is too weak ( $L_{2}$ rather than $H^{1}$ ); in particular, $A^{-1}$ is compact as an operator on $L_{2}\left(\square_{d}\right)$.

The approximation error achieved by $\hat{S}_{r}$ in appropriate norms is investigated in Dahmen et al. (2016). Specializing a result from this work that holds for more general scales of norms (Dahmen et al. 2016, Prop. 1), for the present choice of $A$ and approximation in $H^{1}$-norm, we obtain the following.

Theorem 4.6. For $s \in(0,1)$, for all $f \in H^{-1+s}\left(\square_{d}\right)$,

$$
\left\|A^{-1} f-\hat{S}_{r}(A) f\right\|_{H^{-1+s}\left(\square_{d}\right) \rightarrow H_{0}^{1}\left(\square_{d}\right)} \leq 8 \exp \left(-\frac{s}{2} \pi \sqrt{r}\right)\|f\|_{H^{-1+s}\left(\square_{d}\right)}
$$

The convergence rate is thus strongly dependent on additional regularity of corresponding right-hand sides $f$ higher than $H^{-1}$, and thus on compactness of $A^{-1}$. In Dahmen et al. (2016), the approximations $\hat{S}_{r}(A)$ are combined with quadrature approximations of contour integral representations of required operator exponentials to obtain bounds on the total computational complexity of approximating $A^{-1} f$ in $H_{0}^{1}\left(\square_{d}\right)$; provided that $f$ has suitable low-rank structure, these bounds are in particular polynomial with respect to $d$.

### 4.3. Low-rank preconditioners for finite difference and finite element discretizations

Let us again consider $\mathbf{A}_{\mathrm{FD}} \in \mathbb{R}^{N \times N}$ as in (3.3), rescaled without loss of generality such that $\sigma\left(\mathbf{A}_{\mathrm{FD}}\right) \subset\left[1, \operatorname{cond}\left(\mathbf{A}_{\mathrm{FD}}\right)\right]$; recall that $\operatorname{cond}\left(\mathbf{A}_{\mathrm{FD}}\right)=O\left(h^{-2}\right)$ with $h=$ $1 /(N+1)$. We then obtain the following from Proposition 4.1(ii) and Corollary 4.3.

Proposition 4.7. With $S_{\lambda, \delta_{0}, \eta, T}$ as in Corollary 4.3, with $\delta \in(0,1)$ and $T=$ $\operatorname{cond}\left(\mathbf{A}_{\mathrm{FD}}\right)$, for $\lambda \in(0,1]$ let

$$
\mathbf{C}_{\lambda}=S_{\lambda, \delta / 2, \delta / 2, T}\left(\mathbf{A}_{\mathrm{FD}}\right)
$$

Then, for $\lambda=1$ and $\lambda=1 / 2$, we have

$$
\operatorname{cond}\left(\mathbf{C}_{1} \mathbf{A}_{\mathrm{FD}}\right) \leq \frac{1+\delta}{1-\delta}, \quad \operatorname{cond}\left(\mathbf{C}_{1 / 2} \mathbf{A}_{\mathrm{FD}} \mathbf{C}_{1 / 2}\right) \leq\left(\frac{1+\delta}{1-\delta}\right)^{2}
$$

where $\mathbf{C}_{1}$ and $\mathbf{C}_{1 / 2}$ are each sums of $O(|\log \delta|+|\log h|)$ elementary tensor products.
Exponential sum preconditioners have also been constructed in Khoromskij (2009) based on uniform approximation bounds as in Section 4.1.2; however, this leads to ranks scaling with respect to $h$ as $O\left(\log ^{2} h\right)$ instead of $O(|\log h|)$.

Remark 4.8. Compared to the commonly used left preconditioning $\mathbf{C}_{1} \mathbf{A}_{\mathrm{FD}}$, the two-sided preconditioning $\mathbf{C}_{1 / 2} \mathbf{A}_{\mathrm{FD}} \mathbf{C}_{1 / 2}$ based on the exponential sum for $t \mapsto t^{-1 / 2}$ can be of particular interest in the context of low-rank methods, for the following reason: starting from the original system $\mathbf{A}_{\mathrm{FD}} \mathbf{u}_{h}=\mathbf{f}_{h}$, when using two-sided
preconditioning we solve the equivalent system $\mathbf{C}_{1 / 2} \mathbf{A}_{\mathrm{FD}} \mathbf{C}_{1 / 2} \tilde{\mathbf{u}}_{h}=\mathbf{C}_{1 / 2} \mathbf{f}_{h}$ with $\tilde{\mathbf{u}}_{h}=\mathbf{C}_{1 / 2}^{-1} \mathbf{u}_{h}$. This means that $\left\|\tilde{\mathbf{u}}_{h}\right\|_{2}^{2} \approx\left\langle\mathbf{A}_{\mathrm{FD}} \mathbf{u}_{h}, \mathbf{u}_{h}\right\rangle$; similarly to Riesz basis representations as in Remark 3.11, when working with two-sided preconditioning, the Euclidean norm of solution coefficients (which is controlled in HSVD rank truncations) is thus equivalent to the discrete energy norm induced by $\mathbf{A}_{\mathrm{FD}}$.

Remark 4.9. As an alternative preconditioner for finite difference and finite element discretizations, standard multigrid methods performed in tensor representation are proposed in Ballani and Grasedyck (2013) and Hackbusch (2015b). When such methods are used on a grid hierarchy on product domains such as $\square_{d}$, prolongation and restriction operations are elementary tensor products of corresponding one-dimensional grid hierarchies. Addition of fine- and coarse-grid corrections as well as iterative smoothers performed in tensor representations generally increase the representation ranks of iterates, and these operations thus need to be combined with rank truncations. The convergence analysis of this approach with truncation errors is an open problem.

### 4.4. Low-rank diagonal preconditioners

Recall the wavelet Riesz basis representation $\mathbf{A}$ of the Laplacian defined in (3.21) with the decomposition $\mathbf{A}=\mathbf{D T D}$. Here $\mathbf{T}$ is the representation of the Laplacian with respect to the orthonormal basis

$$
\left\{\bigotimes_{i=1}^{d} \psi_{v_{i}}\right\}_{v \in \mathrm{~V}^{d}}
$$

of $L_{2}\left(\square_{d}\right)$, and

$$
\begin{equation*}
\mathbf{D}\left[v, v^{\prime}\right]=\left(\sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right)^{-1 / 2}, \quad v, v^{\prime} \in \mathrm{V}^{d} \tag{4.8}
\end{equation*}
$$

In this setting, $\mathbf{D}$ thus acts as a diagonal preconditioner on the sequence space $\ell_{2}\left(\mathrm{~V}^{d}\right)$, ensuring that $\mathbf{A}: \ell_{2}\left(\mathrm{~V}^{d}\right) \rightarrow \ell_{2}\left(\mathrm{~V}^{d}\right)$ is an isomorphism. The issue we are facing here is that (4.8) does not have low-rank form. We thus consider how $\mathbf{D}$ can be replaced by low-rank substitutes.

Remark 4.10. For any sequence $\left(d_{v}\right)_{v \in \mathrm{~V} d}$ of positive numbers with

$$
d_{v}=\left(\sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right)^{-1 / 2}, \quad v \in \mathrm{~V}^{d}
$$

$\left\{d_{\nu} \bigotimes_{i=1}^{d} \psi_{\nu_{i}}\right\}_{\nu \in \mathrm{V}^{d}}$ is a Riesz basis of $H_{0}^{1}\left(\square_{d}\right)$.
We now choose such $d_{v}$ to obtain a substitute with efficient low-rank approximations for D. For any fixed $\delta_{0} \in(0,1)$, with $S_{1 / 2, \delta_{0}}$ as in Corollary 4.3 and
defining

$$
M_{0}=\min _{v \in \mathrm{~V}^{d}} \sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}
$$

let

$$
d_{\nu}=M_{0}^{-1 / 2} S_{1 / 2, \delta_{0}}\left(M_{0}^{-1} \sum_{i=1}^{d}\left\|\psi_{\nu_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right)
$$

By Corollary 4.3,

$$
\left|\left(\sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right)^{-1 / 2}-d_{v}\right| \leq \delta_{0}\left(\sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right)^{-1 / 2}, \quad v \in \vee^{d}
$$

Thus $\left\{\tilde{\Psi}_{v}\right\}_{\nu \in \mathrm{V}^{d}}$ with $\tilde{\Psi}_{v}=d_{v} \bigotimes_{i=1}^{d} \psi_{\nu_{i}}$ is indeed a Riesz basis of $H_{0}^{1}\left(\square_{d}\right)$.
Representing $A$ in terms of this basis, we obtain the new infinite matrix representation

$$
\tilde{\mathbf{A}}=\left(\left\langle A \tilde{\Psi}_{\nu^{\prime}}, \tilde{\Psi}_{\nu}\right\rangle\right)_{\nu, v^{\prime} \in \mathrm{V}^{d}}=\tilde{\mathbf{D}} \mathbf{T} \tilde{\mathbf{D}},
$$

where

$$
\tilde{\mathbf{D}}\left[v, v^{\prime}\right]=d_{v} \delta_{v, v^{\prime}} \quad \text { for all } v, v^{\prime} \in \mathrm{V}^{d} .
$$

Note that with $\tau$ and $m$ as in (4.5), with

$$
\omega_{k}=\tau \sqrt{\frac{e^{k \tau}}{M_{0} \pi}}, \quad \alpha_{k}=\frac{e^{k \tau}}{M_{0}}
$$

we have

$$
\begin{equation*}
\tilde{\mathbf{D}}[v, v]=d_{v}=\sum_{k=-\infty}^{m} \omega_{k} \prod_{i=1}^{d} e^{-\alpha_{k}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}}, \tag{4.9}
\end{equation*}
$$

in other words $\tilde{\mathbf{D}}$ has an explicit expansion in terms of tensor product diagonal operators with infinitely many terms.

The strategy pursued in Bachmayr and Dahmen (2016b) is now to formulate the problem in terms of the basis $\left\{\tilde{\Psi}_{\nu}\right\}_{\nu \in \mathrm{V}^{d}}$ with operator representation $\tilde{\mathbf{D} T} \mathbf{D}$, with the representation (3.23) of the operator equation $A u=f \in H^{-1}\left(\square_{d}\right)$ modified to

$$
\begin{equation*}
\tilde{\mathbf{D} T \mathbf{D}} \mathbf{u}=\tilde{\mathbf{D}} \mathbf{g}, \tag{4.10}
\end{equation*}
$$

and introduce additional low-rank approximations of $\tilde{\mathbf{D}}$. This can be done for any finite subset $\Lambda \subset \vee^{d}$, for which we define

$$
T(\Lambda)=M_{0}^{-1} \max _{v \in \Lambda} \sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}
$$

For $\eta>0$ such that $\delta_{0}+\eta<1$, we define the diagonal operator

$$
\begin{align*}
\tilde{\mathbf{D}}_{\eta, \Lambda}[v, v] & =M_{0}^{-1 / 2} S_{1 / 2, \delta_{0}, \eta, T(\Lambda)}\left(M_{0}^{-1} \sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right) \\
& =\sum_{k=-n}^{m} \omega_{k} \prod_{i=1}^{d} e^{-\alpha_{k}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}}, \tag{4.11}
\end{align*}
$$

with $S_{1 / 2, \delta_{0}, \eta, T(\Lambda)}$ and $n$ as in Corollary 4.3, where $n$ from (4.5) satisfies

$$
n=\left\lceil\tau^{-1}(2|\ln \eta|+\ln T(\Lambda))\right\rceil
$$

and the $\omega_{k}, \alpha_{k}$ are as in (4.9). The modified rescaling $\tilde{\mathbf{D}}$ and its low-rank approximations $\tilde{\mathbf{D}}_{\eta, \Lambda}$ have the following properties by Corollary 4.3 and Proposition 4.1.

Proposition 4.11. Let $\delta_{0} \in(0,1)$; then, for $\mathbf{D}$ as in (4.8) and $\tilde{\mathbf{D}}$ defined in (4.9), we have $\left\|\mathbf{D}^{-1} \tilde{\mathbf{D}}\right\|_{\ell_{2} \rightarrow \ell_{2}} \leq 1+\delta_{0},\left\|\tilde{\mathbf{D}}^{-1} \mathbf{D}\right\|_{\ell_{2} \rightarrow \ell_{2}} \leq\left(1-\delta_{0}\right)^{-1}$, and

$$
\langle\tilde{\mathbf{D}} \mathbf{T} \tilde{\mathbf{D}} \mathbf{v}, \mathbf{v}\rangle \approx\|\mathbf{v}\|_{\ell_{2}}^{2} \quad \text { for all } \mathbf{v} \in \ell_{2}\left(\mathbf{V}^{d}\right) .
$$

Moreover, for $\eta>0$ such that $\delta=\delta_{0}+\eta<1$ and finite $\Lambda \subset \vee^{d}$, let $\tilde{\mathbf{D}}_{\eta, \Lambda}$ be defined as in (4.11). Then

$$
\left\langle\tilde{\mathbf{D}}_{\eta, \Lambda} \mathbf{T} \tilde{\mathbf{D}}_{\eta, \Lambda} \mathbf{v}, \mathbf{v}\right\rangle \approx\|\mathbf{v}\|_{\ell_{2}}^{2} \quad \text { for all } \mathbf{v} \in \ell_{2}\left(\vee^{d}\right) \text { with } \operatorname{supp} \mathbf{v} \subseteq \Lambda
$$

with constants independent of $\Lambda$. In addition,

$$
\sup _{\substack{\|\mathbf{v}\|_{\ell_{2}=1} \\ \operatorname{supp} \mathbf{v} \subseteq \Lambda}}\left\|\mathbf{D}^{-1}\left(\tilde{\mathbf{D}}-\tilde{\mathbf{D}}_{\eta, \Lambda}\right) \mathbf{v}\right\|_{\ell_{2}} \leq \eta, \quad \sup _{\substack{\|\mathbf{v}\|_{e_{2}}=1 \\ \operatorname{supp} \mathbf{v} \subseteq \Lambda}}\left\|\mathbf{D}^{-1}\left(\mathbf{D}-\tilde{\mathbf{D}}_{\eta, \Lambda}\right) \mathbf{v}\right\|_{\ell_{2}} \leq \delta .
$$

For the number of summands in $\tilde{\mathbf{D}}_{\eta, \Lambda}$ in terms of $\eta$ and $\Lambda$, we have

$$
\begin{equation*}
m+n+1 \lesssim 1+|\log \eta|+\log \left(\max _{v \in \Lambda} \max _{i=1, \ldots, d}\left|v_{i}\right|\right) \tag{4.12}
\end{equation*}
$$

In summary, we obtain low-rank approximations of the wavelet Riesz basis representation of $A$. In this particular case, the evaluation of exponential sums requires only trivial matrix exponentials of diagonal matrices that can be evaluated entrywise as in (4.11).

Remark 4.12. Note that by (4.12) we have a dependence of the representation rank of $\tilde{\mathbf{D}}_{\eta, \Lambda}$ on the maximum wavelet level $\max _{i}\left|v_{i}\right|$ in the index set $\Lambda \subset \mathrm{V}^{d}$. This dependence becomes more transparent when using, based on the observation (3.24), the alternative rescaling by $d_{v}=2^{-\max _{i}\left|v_{i}\right|}, v \in \mathrm{~V}^{d}$, which satisfies

$$
2^{-\max _{i}\left|v_{i}\right|} \leq\left(\sum_{i=1}^{d}\left\|\psi_{v_{i}}^{\prime}\right\|_{L_{2}(0,1)}^{2}\right)^{-1 / 2} \leq \sqrt{d} 2^{-\max _{i}\left|v_{i}\right|}
$$

Note that the spectral bounds for the resulting operator representations may then
deteriorate by a factor $\sqrt{d}$, but for any finite $\Lambda \subset \vee^{d}$, we obtain a simple explicit low-rank representation as follows. For $S \subset \vee$, let $\mathbf{1}_{S} \in \ell_{2}(V)$ denote the sequence with entries equal to one on $S$, and zero otherwise. Moreover, let $J(\Lambda)=\max _{v \in \Lambda} \max _{i}\left|v_{i}\right|$. Then

$$
\left(2^{-\max _{i}\left|v_{i}\right|}\right)_{v \in \mathrm{~V}^{d}}=\frac{1}{2} \sum_{l=1}^{J(\Lambda)-1} 2^{-l} \bigotimes_{i=1}^{d} \mathbf{1}_{\{v \in \mathrm{~V}:|v| \leq l\}}+2^{-J(\Lambda)} \bigotimes_{i=1}^{d} \mathbf{1}_{\{v \in \mathrm{~V}:|v| \leq J(\Lambda)\}}
$$

with $\max _{v \in \Lambda} \max _{i}\left|v_{i}\right|$ separable terms.
A further approach based on this alternative rescaling sequence is to construct a partition of $\vee^{d}=\bigcup_{k=1}^{\infty} \Lambda_{k}$ such that for each $\Lambda_{k}$, the vector $\left(2^{-\max _{i}\left|\nu_{i}\right|}\right)_{v \in \Lambda_{k}}$ is an elementary tensor product. When using a separate tensor representation for each $\Lambda_{k}$, applying these rescalings then leaves the respective ranks unchanged. Partitions with this property of finite subsets $\left\{v \in \vee^{d}: \max _{i}\left|v_{i}\right| \leq J\right\} \subset \vee^{d}$ for any $J \in \mathbb{N}$ into $O\left(J^{d-2} 2^{J}\right)$ elements are constructed in Bachmayr (2012a).

### 4.5. Preconditioning tensorized discretizations

Let us now return to the tensorized multilevel representations considered in Section 3.4. As noted in Example 3.17, for high discretization levels $L$ corresponding (in $d$-dimensional problems) to grids of size $O\left(2^{d L}\right)$, naive tensorized representations of differential operators have large representation condition numbers according to Definition 2.45.

In certain special cases, this issue can be circumvented by explicit inverses, constructed for instance for the matrix $\boldsymbol{\Delta}_{\mathbf{2}}$ considered in Example 3.17 in Kazeev and Khoromskij (2012), but for more general problems these are not available and suitable preconditioning is required.

The unusual aspect here is that such a preconditioner needs to remedy not only the large standard condition number of the represented matrix but also, as noted in Example 3.17, the equally large representation condition number of the underlying low-rank representation. Such a preconditioner is proposed and analysed in Bachmayr and Kazeev (2020), with finite element discretizations of diffusion problems on $\square_{d}$ (where in this context we aim at $d=2$ or $d=3$ ) with diffusion coefficient $\hat{M}$ as model problems. With a dyadically refined hierarchy of tensor product finite element discretizations with nested spaces $V_{\ell}$ on levels $\ell=0, \ldots, L$, where $\left\{\varphi_{i}\right\}_{i \in \mathcal{N}_{L}}$ with $\# \mathcal{N}_{L}=O\left(2^{d L}\right)$ is the basis of hat functions on level $L$, the considered discretization matrices read

$$
\mathbf{A}=\left(\int_{\square_{d}} \hat{M} \nabla \varphi_{i} \cdot \nabla \varphi_{j} \mathrm{~d} x\right)_{i, j \in \mathcal{N}_{L}}
$$

with a uniformly positive definite $\hat{M} \in \mathbb{R}^{d \times d}$. These discretization matrices are represented in multilevel tensorized format in transposed ordering, according to Remark 3.18, in the form

$$
\mathbf{A}=\tau\left(\mathbf{A}_{0}, \ldots, \mathbf{A}_{L}\right)=\mathbf{A}_{0} \bowtie \cdots \bowtie \mathbf{A}_{L},
$$

where each $\mathbf{A}_{\ell}$ for $\ell=0, \ldots, L$ acts on modes of size $2^{d}$ with indices in $\{0,1\}^{d}$.
We have a decomposition $\mathbf{A}=\mathbf{G}^{\top} \mathbf{M G}$, where $\mathbf{M}=\tau(\mathrm{M})$ is a matrix depending on the diffusion coefficient $\hat{M}$ with well-conditioned representation $M$, and $\mathbf{G}=\tau(\mathrm{G})$ has an ill-conditioned representation similarly to Example 3.17, where $\max _{\ell} \operatorname{oprcond}_{\ell}(\mathrm{G}) \approx 2^{L}$. Here $\mathbf{G}$ can be interpreted as a discrete gradient that maps $\operatorname{span}\left\{\varphi_{i}\right\}_{i \in \mathcal{N}_{L}}$ to a space of discontinuous piecewise polynomials.

As a first step towards a suitable preconditioner, we have the following result (Bachmayr and Kazeev 2020, Thm 2, Thm 3).

Theorem 4.13. Let $\mathbf{C}=\sum_{\ell=0}^{L} 2^{-\ell} \mathbf{P}_{\ell, L} \mathbf{P}_{\ell, L}^{\top}$, where $\mathbf{P}_{\ell, L}$ is the matrix representation of the canonical prolongation operator from $V_{\ell}$ to $V_{L}$. Then $\langle\mathbf{C A C v}, \mathbf{v}\rangle \approx\|\mathbf{v}\|_{2}^{2}$ for all $\mathbf{v}$, with constants that depend only on $M$ and $d$. Moreover, there exists C with $\mathbf{C}=\tau(\mathrm{C})$ with ranks bounded independently of $L$.

Note that the matrix $\mathbf{C}$ provides a two-sided version of the classical BPX preconditioner (Bramble, Pasciak and Xu 1990). Although CAC is well-conditioned as a matrix uniformly in $L$, performing the multiplication by CAC by separately applying $\mathbf{C}$ and $\mathbf{A}$ does not remove the representation ill-conditioning. To arrive at a well-conditioned representation, however, we can combine $G$ and $C$ into a new rank-reduced representation B (again with ranks bounded independently of $L$ ) such that $\tau(\mathrm{B})=\mathbf{G C}$. We arrive at

$$
\mathbf{C A C}=\tau(\mathrm{B})^{\top} \tau(\mathrm{M}) \tau(\mathrm{B})
$$

which provides a well-conditioned representation of CAC for all $L$. For the technical details of the construction, we refer to Bachmayr and Kazeev (2020). Based on this preconditioner, numerically stable computations with $L \approx 50$, corresponding to represented grid sizes near machine precision, are demonstrated to be feasible.

## 5. Solving operator equations in low-rank format

In this section we turn to the central task of computing low-rank approximations of solutions of elliptic operator equations $A u=f$, where $A: V \rightarrow V^{\prime}, f \in V^{\prime}$ with a separable Hilbert space $V$. Here our aim is, for each given error tolerance $\varepsilon>0$, to find an approximation $u_{\varepsilon}$ to the exact solution $u \in V$ such that $\left\|u-u_{\varepsilon}\right\|_{V} \leq \varepsilon$. A guiding model case will again be the Poisson problem as in Example 3.1, where $V=H_{0}^{1}\left(\square_{d}\right)$.

In the context of low-rank approximations, guaranteeing convergence to the exact solution with respect to the $V$-norm involves two central issues: on the one hand, we need to find highly nonlinear representations of approximate solutions, where the required rank parameters are not known a priori. This is addressed by iterative solvers operating on suitable basis representations of the PDE problem. On the other hand, we need to ensure sufficiently small discretization errors, which requires a posteriori error estimates computable in high dimensions as well as adaptive discretization refinement that is compatible with the low-rank format.

Convergence guarantees alone, however, are not very informative: to give an extreme example, one could use a standard solver to obtain a sufficiently accurate approximate solution in full format, and only then compute a compressed lowrank approximation. With such an approach, one can easily ensure convergence, but without any benefits concerning computational complexity of using low-rank approximations. An essential question is therefore how the computational costs of a low-rank solver compare to the minimal required costs, which depend on the intrinsic low-rank approximability of solutions and problem data. In particular, can the curse of dimensionality be avoided in the total computational complexity?

In Section 5.1 we first survey different approaches for computing low-rank approximations of solutions of linear systems of equations. In Section 5.2 we consider iterative methods for finding low-rank approximate solutions with two main features: first, they can always be guaranteed to converge, and second, for certain typical model classes of low-rank approximability, the approximations produced by the iterations have near-optimal ranks in relation to the currently achieved accuracy. Two approaches for achieving this are discussed, based on HSVD truncation (see Section 2.7.1) and on soft thresholding of hierarchical tensors (see Section 2.7.2), respectively. At this point, these methods are formulated as operating on fixed discretizations or on function spaces.

In Section 5.3 we turn to the issue of combining rank adaptation with adaptive refinement of discretizations. Again we obtain methods that can be guaranteed to converge, and that, for solutions in certain model classes of approximability, exhibit near-optimal performance in terms of ranks and discretization sizes. Note that the resulting algorithms are universal: the respective degree of approximability of the exact solution does not need to be known but is automatically detected by the adaptive scheme. Altogether, this leads to bounds on the total computational complexity of approximately solving elliptic PDEs in high dimensions by such methods, which we discuss in Section 5.4. In Section 5.5 we survey what is known on the low-rank approximability of particular classes of elliptic problems.

### 5.1. Iterative solvers in low-rank format

Iterative solvers operating on low-rank representations are often introduced in the literature for discretized problems on finite-dimensional spaces, such as the finite difference discretization $\mathbf{A}_{\mathrm{FD}} \mathbf{u}_{h}=\mathbf{f}_{h}$ discussed in Section 3.1.1. However, the methods considered below can also be formulated to operate on tensors in sequence spaces $\ell_{2}\left(\mathbb{N}^{d}\right)$, which with an appropriate choice of bases is equivalent to performing iterations on function spaces. This is the case with the Riesz basis representations $\mathbf{A u}=\mathbf{f}$, where $\mathbf{A}$ is an isomorphism on $\ell_{2}\left(\mathbb{N}^{d}\right)$, as obtained in Section 3.2. Based on such formulations, we obtain iterative methods that are robust with respect to the choice of discretization. Note that by ellipticity of $A$ we also have ellipticity of $\mathbf{A}$ on $\ell_{2}\left(\mathbb{N}^{d}\right)$, that is,

$$
\langle\mathbf{A v}, \mathbf{v}\rangle \gtrsim\|\mathbf{v}\|_{\ell_{2}}^{2} \quad \text { for all } \mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)
$$

### 5.1.1. Krylov subspace methods with rank truncation

In Krylov subspace methods, approximate solutions are generated by repeated application of the linear mapping $\mathbf{A}_{\text {FD }}$ or $\mathbf{A}$ under consideration. One of the simplest instances of such methods, outlined in Section 1.7, is Richardson iteration. Applied to a fixed discretization $\mathbf{A}_{\mathrm{FD}}$, with a suitable left preconditioner $\mathbf{C}$, in its basic form the method reads

$$
\begin{equation*}
\mathbf{u}_{h}^{n+1}=\mathbf{u}_{h}^{n}-\omega \mathbf{C}\left(\mathbf{A}_{\mathrm{FD}} \mathbf{u}_{h}^{n}-\mathbf{f}_{h}\right) \tag{5.1}
\end{equation*}
$$

As described in Section 1.7, when applying the method in this form directly in lowrank format, the representation ranks of $\mathbf{u}_{h}^{n}$ will in general increase exponentially with respect to $n$. It is thus natural to combine the basic iteration (5.1) with a suitable rank reduction operation $\mathcal{R}_{n}$,

$$
\begin{equation*}
\mathbf{u}_{h}^{n+1}=\mathcal{R}_{n}\left(\mathbf{u}_{h}^{n}-\omega \mathbf{C}\left(\mathbf{A}_{\mathrm{FD}} \mathbf{u}_{h}^{n}-\mathbf{f}_{h}\right)\right) \tag{5.2}
\end{equation*}
$$

One option is to choose $\mathcal{R}_{n}(\mathbf{v})$ as the truncation of the HSVD of the given $\mathbf{v} \in$ $\ell_{2}\left(\mathbb{N}^{d}\right)$ to fixed hierarchical rank. With different tensor formats, such strategies are used in Khoromskij and Schwab (2011) and Billaud-Friess, Nouy and Zahm (2014). However, the resulting truncated iterative scheme can generally be guaranteed to converge only under extremely restrictive conditions on the error reduction achieved in each step of (5.1); see Hackbusch, Khoromskij and Tyrtyshnikov (2008), BillaudFriess et al. (2014) and Bachmayr et al. (2016, Prop. 6.1).

A generally more practical alternative is to choose $\mathcal{R}_{n}$ as the truncation of the HSVD to lower ranks up to a prescribed error tolerance, potentially dependent on $n$, based on the computable error bounds in terms of matricization singular values of Theorem 2.36. In combination with Richardson iteration, such a strategy is used in Matthies and Zander (2012). Error-controlled truncations can also be combined with more advanced Krylov space methods, such as the preconditioned conjugate gradient (PCG) method with rank truncation proposed in Kressner and Tobler (2011a) and Tobler (2012) for positive definite problems, or the generalized minimal residual (GMRES) method with truncation proposed in Ballani and Grasedyck (2013). The variants are typically observed to require fewer iterations than the truncated Richardson iteration, although the orthogonality properties that are crucial for the standard convergence analysis of PCG and GMRES are not preserved by the truncations. However, the analysis by Ali and Urban (2020) of rank-truncated PCG with appropriate truncation tolerances as a perturbed descent method shows that worst-case performance comparable to (5.2) can still be ensured.

In iterations with truncations up to specified tolerances, however, the crucial question is how the ranks of the iterates evolve, that is, whether the truncation tolerances are sufficiently large to ensure the appropriate rank reduction. Ideally, the ranks of iterates should behave similarly to the ranks of best approximations of comparable accuracy. How this can be achieved in a general setting is treated in Section 5.2.

To enable a useful comparison to low-rank best approximations, rank truncations need to be performed with errors with respect to the appropriate norm, in our present setting in $V$-norm. As discussed in Remark 4.8, this can be achieved by two-sided preconditioning and is inherent in Richardson iteration performed on a Riesz basis representation $\mathbf{A u}=\mathbf{f}$ with $\mathbf{A}=\mathbf{D T D}, \mathbf{f}=\mathbf{D g}$. For such a Riesz basis representation, $\left\|\mathbf{u}-\mathbf{u}^{n}\right\|_{\ell_{2}} \approx\left\|\mathbf{f}-\mathbf{A} \mathbf{u}^{n}\right\|_{\ell_{2}}$, and both quantities are proportional to the error of approximating the exact solution $u$ in $V$-norm. In this case (5.2) takes the form

$$
\mathbf{u}^{n+1}=\mathcal{R}_{n}\left(\mathbf{u}^{n}-\omega \mathbf{r}^{n}\right),
$$

where $\mathbf{r}^{n}$ is a sufficiently accurate low-rank approximation of the residual $\mathbf{A u}^{n}$ $\mathbf{f}=\mathbf{D}\left(\mathbf{T D u} \mathbf{u}^{n}-\mathbf{g}\right)$, with the infinite diagonal scaling matrix $\mathbf{D}$ acting as a twosided preconditioner. In this form, Richardson iteration serves as the basis of low-rank methods with adaptively refined discretizations, considered in detail in Section 5.3.2.

### 5.1.2. Alternating optimization

Solution strategies that sequentially optimize individual components of tensor representations are based on reformulations of the original systems of equations as minimization problems. In particular, when $\mathbf{A}$ is symmetric, the solution $\mathbf{u}$ of $\mathbf{A u}=\mathbf{f}$ can be characterized as the unique minimizer of the functional $J$ defined by

$$
\begin{equation*}
J(\mathbf{v})=\frac{1}{2}\langle\mathbf{A v}, \mathbf{v}\rangle-\langle\mathbf{f}, \mathbf{v}\rangle . \tag{5.3}
\end{equation*}
$$

Methods of this kind have been considered for the canonical tensor format (1.9), where the issues discussed in Section 1.4 apply; see for instance Uschmajew (2012). In the context of subspace tensor formats, although alternating optimization approaches are applicable to general hierarchical tensors, they have been investigated primarily for the tensor train format, for which they are easiest to formulate.

A first classical method of this type is known, depending on context, as alternating least squares (ALS), the alternating linear scheme (Holtz, Rohwedder and Schneider 2012a) or as one-site density matrix renormalization group (White 2005). Starting from a tensor train representation

$$
\begin{equation*}
\mathbf{u}^{n}=\tau\left(\mathbf{X}_{1}^{n}, \ldots, \mathbf{X}_{d}^{n}\right)=\mathbf{X}_{1}^{n} \bowtie \cdots \bowtie \mathbf{X}_{d}^{n} \tag{5.4}
\end{equation*}
$$

the next iterate $\mathbf{u}^{n+1}=\mathbf{X}_{1}^{n+1} \bowtie \cdots \bowtie \mathbf{X}_{d}^{n+1}$ is obtained by sequentially solving the substeps

$$
\begin{align*}
\mathbf{X}_{1}^{n+1} & =\underset{\mathbf{Y}}{\arg \min } J\left(\tau\left(\mathbf{Y}, \mathbf{X}_{2}^{n}, \mathbf{X}_{3}^{n}, \ldots, \mathbf{X}_{d}^{n}\right)\right) \\
\mathbf{X}_{2}^{n+1} & =\underset{\mathbf{Y}}{\arg \min } J\left(\tau\left(\mathbf{X}_{1}^{n+1}, \mathbf{Y}, \mathbf{X}_{3}^{n}, \ldots, \mathbf{X}_{d}^{n}\right)\right),  \tag{5.5}\\
& \vdots \\
\mathbf{X}_{d}^{n+1} & =\underset{\mathbf{Y}}{\arg \min } J\left(\tau\left(\mathbf{X}_{1}^{n+1}, \mathbf{X}_{2}^{n+1}, \mathbf{X}_{3}^{n+1}, \ldots, \mathbf{Y}\right)\right) .
\end{align*}
$$

Note that since $\tau$ is multilinear and $J$ is quadratic, the solution of the minimization problem in each substep amounts to a linear system of equations for $\mathbf{Y}$. In this method, only one component is modified at a time and all rank parameters are determined by the remaining components. As a consequence, in this basic form, there is no mechanism for the adaptation of approximation ranks, and the ranks of $\mathbf{u}^{n}$ remain bounded by those of the initial iterate $\mathbf{u}^{0}$ for all $n$. Effectively, we are thus minimizing $J$ under a rank constraint. The iteration may in general fail to converge to the constrained solution. Convergence can be shown under additional conditions that, however, cannot generally be verified; see Rohwedder and Uschmajew (2013), Wang and Chu (2014) and Oseledets, Rakhuba and Uschmajew (2018).

In the (two-site) density matrix renormalization group (DMRG) algorithm introduced by White (1992) in the context of eigenvalue problems, and formulated by Vidal (2003) in terms of matrix product states, the iteration (5.5) is modified by optimizing two components at a time as follows: in the first substep, starting again from $\mathbf{u}^{n}$ as in (5.4), we combine $\mathbf{X}_{1}^{n} \bowtie \mathbf{X}_{2}^{n}$ into a single component carrying two modes of the represented tensor, and minimize over this fourth-order component to obtain

$$
\tilde{\mathbf{X}}_{1,2}^{n+1}=\underset{\tilde{\mathbf{Y}}}{\arg \min } J\left(\tilde{\mathbf{Y}} \bowtie \mathbf{X}_{3}^{n} \bowtie \mathbf{X}_{4}^{n} \bowtie \cdots \bowtie \mathbf{X}_{d}^{n}\right) .
$$

We then perform a low-rank approximation, for instance by truncated SVD, to split the intermediate result again into two components:

$$
\tilde{\mathbf{X}}_{1,2}^{n+1} \approx \mathbf{X}_{1}^{n+1} \bowtie \tilde{\mathbf{X}}_{2}^{n+1}
$$

In the next substep, we combine $\tilde{\mathbf{X}}_{2}^{n+1} \bowtie \mathbf{X}_{3}^{n}$, and so forth. As a solver for linear systems formulated as a minimization problem (5.3), this approach is considered in Holtz et al. (2012a) and Oseledets and Dolgov (2012); see also Oseledets (2011a). Although this approach offers a mechanism for the adaptation of approximation ranks, it may fail to converge in certain cases. An explicit example is given in Pfeffer (2018, Ex. 5.12).

A further variant that can achieve guaranteed convergence is the alternating minimal energy (AMEn) method of Dolgov and Savostyanov (2014). The basic iteration is performed as in (5.5), but in outer iteration $n$, before performing the $i$ th substep for $i=2, \ldots, d$, the previously updated component $\mathbf{X}_{i-1}^{n+1}$ is augmented by a certain number of vectors from a projected residual approximation. The available analysis of this methods guarantees convergence, but with an error reduction that may in the worst case deteriorate exponentially with respect to the tensor order d. With additional practical modifications proposed in Dolgov and Savostyanov (2014), including rank truncation of iterates, the method is demonstrated to perform far better in practice than indicated by the analysis, in particular on test problems with coarse discretizations and low ranks. Although with fine discretizations its performance can degrade for ill-conditioned problems requiring large ranks
(Bachmayr and Dahmen 2016a, Sec. 4.3), it shows very good practical efficiency when combined with optimal preconditioning (Bachmayr and Kazeev 2020, Sec. 7).

### 5.1.3. Riemannian optimization

Solving linear systems by minimization of functionals $J$ as in (5.3) can also be approached differently, using the manifold structure of tensors of fixed hierarchical ranks discussed in Section 2.10. Specifically, for a given dimension tree $\mathbb{T}$, for each $r \in \operatorname{Ranks}(\mathbb{E})$, the set $\left\{\mathbf{v} \in \ell_{2}\left(\mathcal{I}_{1} \times \cdots \times \mathcal{I}_{d}\right): \operatorname{rank}_{\mathbb{E}}(\mathbf{v})=r\right\}$ can be regarded as an embedded submanifold $\mathcal{M}_{r}$ with Riemannian metric inherited from the ambient space $\ell_{2}\left(\mathcal{I}_{1} \times \cdots \times \mathcal{I}_{d}\right)$. Note that while results using this manifold structure are in most cases formulated for finite $\mathcal{I}_{1}, \ldots, \mathcal{I}_{d}$, the approach can be generalized to the case of general separable Hilbert spaces corresponding to countable index sets.

For given $\mathbf{v} \in \mathcal{M}_{r}$ we let $\mathcal{T}_{\mathbf{v}} \mathcal{M}_{r}$ denote the tangent space of $\mathcal{M}_{r}$ at $\mathbf{v}$, and let $\Pi_{\mathrm{v}}$ denote the orthogonal projection onto this tangent space. Basic versions of Riemannian optimization methods can be obtained using, for a given iterate $\mathbf{u}^{n}$, the negative Riemannian gradient

$$
\mathbf{p}^{n}=-\Pi_{\mathbf{u}^{n}} \nabla J\left(\mathbf{u}^{n}\right)
$$

as a search direction. To obtain a new iterate on the manifold, we then use a retraction, that is, a mapping $R_{\mathcal{M}_{r}}: \mathcal{M}_{r} \times \mathcal{T}_{\mathbf{v}} \mathcal{M}_{r} \rightarrow \mathcal{M}_{r}$ with the property

$$
\left\|\mathbf{v}+\mathbf{p}-R_{\mathcal{M}_{\mathrm{r}}}(\mathbf{v}, \mathbf{p})\right\|=o(\|\mathbf{p}\|) \quad \text { for } \mathbf{v} \in \mathcal{M}_{\mathrm{r}}, \mathbf{p} \in \mathcal{T}_{\mathbf{v}} \mathcal{M}_{\mathrm{r}}
$$

With such a mapping, we set $\mathbf{u}^{n+1}=R_{\mathcal{M}_{\mathrm{r}}}\left(\mathbf{u}^{n}, \omega_{n} \mathbf{p}^{n}\right)$ with a suitable damping parameter $\omega_{n}>0$. As shown in Steinlechner (2016), the HSVD truncation to rank $r$ is a retraction. For a comparative overview of several different types of retractions in the case of low-rank matrices, see Absil and Oseledets (2015). In addition to gradient information, the Riemannian Hessian can be used for the construction of Newton-type methods.

For an overview of such methods in the case of low-rank matrices, see Absil et al. (2008). In their basic form operating on fixed ranks, they are applied to tensor completion problems in Kressner, Steinlechner and Vandereycken (2014b) and Da Silva and Herrmann (2015), where tensors are to be recovered from a limited number of entries under rank constraints. Riemannian optimization methods for solving linear systems, in particular finite difference discretizations, are studied in Kressner, Steinlechner and Vandereycken (2016). In this case rank adaptation is performed by applying the solver repeatedly with successively increased ranks.

While the representation ranks that can arise during the iteration are fixed, this comes at the price that convergence to the minimizer of $J$ with ranks constrained to $r$ is not generally guaranteed. For a detailed discussion and numerical comparisons demonstrating accelerated convergence of Riemannian methods in certain cases, see Uschmajew and Vandereycken (2020). In the case of finite-dimensional matrices, systematic approaches to rank adaptation are considered in Uschmajew
and Vandereycken (2015), Schneider and Uschmajew (2015), Zhou et al. (2016) and Gao and Absil (2022).

### 5.1.4. Greedy methods

A further construction principle of optimization-based methods is given by greedy strategies of the basic form

$$
\mathbf{u}^{n+1}=\mathbf{u}^{n}+\mathbf{w}^{n}, \quad \mathbf{w}^{n}=\min _{\mathbf{w} \in \mathcal{S}} J\left(\mathbf{u}^{n}+\mathbf{w}\right),
$$

where $J$ can be quadratic as in (5.3) or a more general functional, and $\mathcal{S}$ is a restricted set of tensors, for instance the set of elementary tensor products. Methods of this type are also known as proper generalized decomposition; see Ammar, Chinesta and Falcó (2010) and Falcó and Nouy (2011, 2012). Closely related approaches are analysed in Cancès, Ehrlacher and Lelièvre (2011, 2013, 2014). Convergence of such methods can be obtained from variational arguments for fairly general classes of functionals on Banach spaces. Unless combined with an additional rank reduction strategy, the ranks of approximations produced by this approach may in general be substantially higher than best approximation ranks for similar accuracies.

### 5.2. Maintaining quasi-optimal ranks

A crucial point when analysing the computational complexity of low-rank solvers is to derive bounds on the ranks of iterates. These have a major impact on the computational costs, since for hierarchical tensor representations, the costs of orthogonalization and HSVD - as in (2.23) and (2.24) - scale like the fourth power of the maximum hierarchical ranks. The natural point of reference for the ranks arising in an iterative method is given by the ranks required for low-rank best approximations of comparable accuracy. The best behaviour of the ranks of iterates one can hope for is that they remain of similar size to these best approximation ranks.

We next consider two constructions of methods that achieve this goal. We formulate these on $\ell_{2}\left(\mathbb{N}^{d}\right)$, corresponding to methods operating on the infinitedimensional function spaces under consideration. The methods then automatically remain robust under a suitable selection of discretizations. Starting from an initial iterate $\mathbf{u}^{0}$ represented in a subspace tensor format, both constructions take the basic form

$$
\mathbf{u}^{n+1}=\left(\mathcal{R}_{n} \circ \mathcal{F}_{n}\right)\left(\mathbf{u}^{n}\right) .
$$

Here, for each $n \in \mathbb{N}_{0}, \mathcal{F}_{n}$ provides an error reduction with respect to the exact solution $\mathbf{u} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ of $\mathbf{A u}=\mathbf{f}$ in the sense that

$$
\begin{equation*}
\left\|\mathbf{u}-\mathcal{F}_{n}\left(\mathbf{u}^{n}\right)\right\|_{\ell_{2}} \leq \rho\left\|\mathbf{u}-\mathbf{u}^{n}\right\|_{\ell_{2}} \tag{5.6}
\end{equation*}
$$

with $\rho \in(0,1)$ independent of $n$. The mapping $\mathcal{R}_{n}$ provides a rank reduction that may increase the error with respect to $\mathbf{u}$.

In principle, the error reduction operation $\mathcal{F}_{n}$ can be realized by any iterative solver that guarantees the property (5.6). However, to eventually allow for complexity estimates, this reduction should be sufficiently robust with respect to the dimensionality and other features of the problem. A simple suitable choice for $\mathcal{F}_{n}$ - as used in Bachmayr (2012a), Bachmayr and Dahmen (2015) and Bachmayr and Schneider (2017) - is one or more steps of Richardson iteration, which may also be performed inexactly such that (5.6) is guaranteed.

### 5.2.1. HSVD truncation

We first consider an approach for obtaining approximations with near-optimal ranks by HSVD truncation with sufficiently large error tolerance. In the case of low-rank tensor approximation, this was analysed in Bachmayr and Dahmen (2015) based on the iterative scheme developed in Bachmayr (2012a). The basic idea for establishing a connection to ranks of best approximations is applicable in a more general context, and we therefore first state this result in abstract terms.

Let $H$ be a Hilbert space and let $\left(\mathcal{B}_{r}\right)_{r \in \mathbb{N}}$ with $\mathcal{B}_{r} \subset H$ be a sequence of subsets, with $\bigcup_{r \in \mathbb{N}} \mathcal{B}_{r}$ dense in $H$, from which approximations are selected, where increasing $r$ corresponds to increasing complexity of approximations. For $v \in H$ and $\varepsilon>0$, we define

$$
b(v, \varepsilon)=\min \left\{r \in \mathbb{N}: \exists w \in \mathcal{B}_{r}:\|v-w\|_{H} \leq \varepsilon\right\}
$$

In other words, $b(v, \varepsilon)$ is the minimal complexity parameter such that $\mathcal{B}_{b(v, \varepsilon)}$ contains a best approximation of accuracy $\varepsilon$. In addition, for each $v \in H$ and $\varepsilon>0$, we assume that we have a linear orthogonal projection $B_{v, \varepsilon}$ onto $\mathcal{B}_{b(v, \varepsilon)}$ such that $\left\|v-B_{v, \varepsilon} v\right\|_{H} \leq \varepsilon$.

Depending on the sets $\mathcal{B}_{r}$, it may be infeasible to compute best approximations, or even to exactly evaluate corresponding approximation errors. We thus assume that for all $r \in \mathbb{N}$ we have mappings $A_{r}: H \rightarrow \mathcal{B}_{r}$ and $E_{r}: H \rightarrow[0, \infty)$ such that for some $\kappa \geq 1$,

$$
\begin{equation*}
\left\|v-A_{r}(v)\right\|_{H} \leq E_{r}(v) \leq \kappa \inf _{w \in \mathcal{B}_{r}}\|v-w\|_{H} \tag{5.7}
\end{equation*}
$$

Here $E_{r}$ plays the role of a computable quasi-optimal error bound for the approximations provided by $A_{r}$, which play the role of computable substitutes of best approximations. In our present setting of hierarchical tensor approximations, the HSVD truncation provides such a quasi-optimal substitute for low-rank best approximations, where a computable error bound is provided by the hierarchical singular values as in Theorem 2.36.

Lemma 5.1. For each $\varepsilon>0$, let $a(v, \varepsilon) \in \mathbb{N}$ be chosen as the minimal integers such that $E_{a(v, \varepsilon)}(v) \leq \varepsilon$. Then, for any $\alpha, \eta>0$ and any $u, v \in H$ with $\|u-v\|_{H} \leq \eta$,

$$
\begin{align*}
\left\|u-A_{a(v, \kappa(1+\alpha) \eta)}(v)\right\|_{H} & \leq(1+\kappa(1+\alpha)) \eta  \tag{5.8}\\
a(v, \kappa(1+\alpha) \eta) & \leq b(u, \alpha \eta)
\end{align*}
$$

Proof. Concerning the first estimate in (5.8), by the triangle inequality

$$
\begin{aligned}
\left\|u-A_{a(v, \kappa(1+\alpha) \eta)}(v)\right\|_{H} & \leq\|u-v\|_{H}+\left\|v-A_{a(v, \kappa(1+\alpha) \eta)}(v)\right\|_{H} \\
& \leq \eta+E_{a(v, \kappa(1+\alpha) \eta}(v)
\end{aligned}
$$

and $E_{a(v, \kappa(1+\alpha) \eta)}(v) \leq \kappa(1+\alpha) \eta$ by definition.
For the proof of the second estimate, note that by linearity and orthogonality of $B_{u, \alpha \eta}$,

$$
\left\|v-B_{u, \alpha \eta} v\right\|_{H} \leq\left\|(v-u)-B_{u, \alpha \eta}(v-u)\right\|_{H}+\left\|u-B_{u, \alpha \eta} u\right\|_{H} \leq(1+\alpha) \eta .
$$

With (5.7), we obtain

$$
E_{b(u, \alpha \eta)}(v) \leq \kappa \inf _{w \in \mathcal{B}_{b(u, \alpha \eta)}}\|v-w\|_{H} \leq \kappa\left\|v-B_{u, \alpha \eta} v\right\|_{H} \leq \kappa(1+\alpha) \eta .
$$

As a consequence,

$$
a(v, \kappa(1+\alpha) \eta)=\min \left\{r \in \mathbb{N}: E_{r}(v) \leq \kappa(1+\alpha) \eta\right\} \leq b(u, \alpha \eta)
$$

which is the second estimate in (5.8).
In other words, Lemma 5.1 states that when an approximation $v$ to $u$ is reapproximated with a sufficiently large error tolerance, the complexity parameter of this reapproximation can be estimated by that of a best approximation of $u$ with proportional error.

In its basic form, Lemma 5.1 goes back to Cohen, Dahmen and DeVore (2001, 2002), where it was used in the context of best n-term approximation; see also Cohen (2003, Thm 4.9.1). In this case we take $H=\ell_{2}(\mathbb{N})$ and for each $n \in \mathbb{N}$, $\mathcal{B}_{n}=\left\{\mathbf{v} \in \ell_{2}(\mathbb{N}): \# \operatorname{supp} \mathbf{v} \leq n\right\}$. The mapping $A_{n}(\mathbf{v})$ is realized by retaining $n$ entries of $\mathbf{v}$ of largest absolute value, so that here

$$
\begin{equation*}
\left\|\mathbf{v}-A_{n}(\mathbf{v})\right\|_{\ell_{2}}=\inf _{\mathbf{v} \in \mathcal{B}_{n}}\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}} \tag{5.9}
\end{equation*}
$$

and thus (5.7) holds with $\kappa=1$. The orthogonal projections realizing best approximations $B_{v, \varepsilon}$ in Lemma 5.1 are then the mappings that restrict to the respective supports and extend by zero.

For later use, we state the definition of approximation spaces for quantifying algebraic rates of best $n$-term approximation; see for example Dahmen (1997), DeVore (1998) and Cohen et al. (2001).

Definition 5.2. For $s>0$ and $\mathbf{v} \in \ell_{2}(\mathbb{N})$, we set

$$
|\mathbf{v}|_{\mathcal{A}^{s}(\mathbb{N})}=\sup _{n \in \mathbb{N}_{0}} n^{s} \min \left\{\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}}: \# \operatorname{supp} \mathbf{w} \leq n\right\}
$$

and $\|\mathbf{v}\|_{\mathcal{A}^{s}(\mathbb{N})}=\|\mathbf{v}\|_{\ell_{2}(\mathbb{N})}+|\mathbf{v}|_{\mathcal{A}^{s}(\mathbb{N})}$, and define

$$
\mathcal{A}^{s}(\mathbb{N})=\left\{\mathbf{v} \in \ell_{2}(\mathbb{N}):|\mathbf{v}|_{\mathcal{A}^{s}(\mathbb{N})}<\infty\right\} .
$$

The same definition applies to other countable index sets, and we simply write $\mathcal{A}^{s}$ when no confusion can arise. In other words, the linear spaces $\mathcal{A}^{s}$ comprise those elements $\mathbf{v} \in \ell_{2}(\mathbb{N})$ that satisfy

$$
\min _{\# \operatorname{supp} \mathbf{w} \leq n}\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}} \leq C n^{-s}
$$

for some $C>0$, where the smallest possible such $C$ is precisely $|\mathbf{v}|_{\mathcal{A}^{s}(\mathbb{N})}$. Conversely, these are the sequences that can be approximated to $\ell_{2}$-error $\varepsilon>0$ using $|\mathbf{v}|_{\mathcal{A}^{s}(\mathbb{N})}^{1 / s} \varepsilon^{-1 / s}$ non-zero coefficients. As a particular conclusion from Lemma 5.1, for elements of $\mathcal{A}^{s}$ we obtain approximations using at most a fixed multiple of the optimal number of coefficients.

In our setting of low-rank approximations by hierarchical tensors, we now apply Lemma 5.1 on the Hilbert space $\ell_{2}\left(\mathbb{N}^{d}\right)$ with

$$
\mathcal{B}_{r}=\left\{\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right):\left\|\operatorname{rank}_{\mathbb{E}}(\mathbf{v})\right\|_{\infty} \leq r\right\}
$$

A quasi-optimal error bound of the form (5.7) is then provided for HSVD truncation by Theorem 2.36. For $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ and $r=\left(r_{e}\right)_{e \in \mathbb{E}}$, it yields

$$
\begin{align*}
\left\|\mathbf{v}-\mathrm{P}_{\mathbf{v}, \mathbf{r}}\right\|_{\ell_{2}} & \leq\left(\sum_{e \in \mathbb{E}} \sum_{k>r_{e}}\left(\sigma_{k}^{[e]}(\mathbf{v})\right)^{2}\right)^{1 / 2} \\
& \leq \kappa_{\mathbb{E}} \min \left\{\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}}: \mathbf{w} \in \ell_{2}\left(\mathbb{N}^{d}\right), \operatorname{rank}_{\mathbb{E}}(\mathbf{w}) \leq \mathrm{r}\right\} \tag{5.10}
\end{align*}
$$

with $\kappa_{\mathbb{E}}=\sqrt{\# \mathbb{E}}$, and where $P_{\mathbf{v}, \mathbf{v}}$ is the HSVD truncation of $\mathbf{v}$ to hierarchical rank $r$. We adapt this to ensure prescribed truncation errors as follows.

Definition 5.3. For each $\eta>0$ and $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$, we choose $r=\left(r_{e}\right)_{e \in \mathbb{E}} \in \operatorname{Ranks}(\mathbb{E})$ with minimal $\|r\|_{\infty}$ such that

$$
\begin{equation*}
\left(\sum_{e \in \mathbb{E}} \sum_{k>r_{e}}\left(\sigma_{k}^{[e]}(\mathbf{v})\right)^{2}\right)^{1 / 2} \leq \eta \tag{5.11}
\end{equation*}
$$

With such $r$, we define $\operatorname{Trunc}_{\eta}(\mathbf{v})$ as the truncation of the HSVD representation of $\mathbf{v}$ to rank r .

Our point of reference is provided by the maximum hierarchical ranks of best approximations with error at most $\eta>0$,

$$
r_{\text {best }}(\mathbf{v}, \eta)=\min \left\{r \in \mathbb{N}_{0}:\left(\exists \mathbf{w} \in \ell_{2}\left(\mathbb{N}^{d}\right):\left\|\operatorname{rank}_{\mathbb{E}}(\mathbf{w})\right\|_{\infty} \leq r \wedge\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}} \leq \eta\right)\right\}
$$

where best approximations can trivially be written as orthogonal projections. Then, for the computable HSVD truncation operation defined in Definition 5.3, combining (5.10) and Lemma 5.1, we obtain the following rank bound in terms of best approximation ranks. Note that here the role of the quasi-optimal computable error bound is played by the expression on the left in (5.11), and the parameter that controls the approximation complexity is the maximum entry of the hierarchical rank as in Definition 5.3.

Lemma 5.4. Let $\kappa_{\mathbb{E}}=\sqrt{\# \mathbb{E}}$ and $\alpha>0$. Then, for any $\mathbf{u}, \mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ with $\|\mathbf{u}-\mathbf{v}\|_{\ell_{2}} \leq \eta$ and

$$
\hat{\mathbf{v}}=\operatorname{Trunc}_{\kappa \mathbb{E}}(1+\alpha) \eta(\mathbf{v}),
$$

we have

$$
\|\mathbf{u}-\hat{\mathbf{v}}\|_{\ell_{2}} \leq\left(1+\kappa_{\mathbb{E}}(1+\alpha)\right) \eta, \quad\left\|\operatorname{rank}_{\mathbb{E}}(\hat{\mathbf{v}})\right\|_{\infty} \leq r_{\text {best }}(\mathbf{u}, \alpha \eta) .
$$

For low-rank approximations of operators corresponding to $d=2$, a similar result is found independently in Dölz, Egger and Schlottbom (2021, Lem. 2.7).

We now turn to the construction of an iterative scheme using Lemma 5.4 based on a mapping $\mathcal{F}_{n}$ with the error reduction property (5.6). Here we assume that $\mathcal{F}_{n}$ satisfies (5.6) with a reduction factor

$$
\begin{equation*}
\rho<\frac{1}{2\left(1+(1+\alpha) \kappa_{\mathbb{E}}\right)}, \tag{5.12}
\end{equation*}
$$

where $\kappa_{\mathbb{E}}=\sqrt{\# \mathbb{E}}=\sqrt{2 d-3}$ in the case of the hierarchical tensor format. This can always be achieved when $\mathcal{F}_{n}$ is realized by sufficiently many steps of (inexact) Richardson iteration.

Then, starting from $\mathbf{u}^{0}=0$ with $\left\|\mathbf{u}-\mathbf{u}^{0}\right\|_{\ell_{2}}=\|\mathbf{u}\|_{\ell_{2}} \leq\left\|\mathbf{A}^{-1}\right\|_{\ell_{2} \rightarrow \ell_{2}}\|\mathbf{f}\|_{\ell_{2}}=\eta_{0}$, for $n=0,1,2, \ldots$, iterate

$$
\begin{equation*}
\mathbf{u}^{n+1}=\operatorname{Trunc}_{\theta \eta_{n}}\left(\mathcal{F}_{n}\left(\mathbf{u}^{n}\right)\right) \quad \text { with } \theta=\frac{(1+\alpha) \kappa_{\mathbb{E}}}{2\left(1+(1+\alpha) \kappa_{\mathbb{E}}\right)}, \tag{5.13}
\end{equation*}
$$

and set $\eta_{n+1}=\frac{1}{2} \eta_{n}$. We then have the following convergence result with rank estimates, which is used implicitly in Bachmayr and Dahmen (2015).

Theorem 5.5. For the sequence $\left(\mathbf{u}^{n}\right)_{n \in \mathbb{N}_{0}}$ defined by (5.13), for each $n \in \mathbb{N}_{0}$ we have

$$
\begin{equation*}
\left\|\mathbf{u}-\mathbf{u}^{n}\right\|_{\ell_{2}} \leq \eta_{n}, \quad\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty} \leq r_{\text {best }}\left(\mathbf{u},\left(1+(1+\alpha) \kappa_{\mathbb{E}}\right)^{-1} \alpha \eta_{n}\right) \tag{5.14}
\end{equation*}
$$

Proof. For $n=0$, the statement is clear. Assume now that it holds for an $n \in \mathbb{N}$. By (5.12), we have

$$
\left\|\mathbf{u}-\mathcal{F}_{n}\left(\mathbf{u}^{n}\right)\right\|_{\ell_{2}} \leq \frac{\eta_{n}}{2\left(1+(1+\alpha) \kappa_{\mathbb{E}}\right)}
$$

Then, by combining (5.13) and Lemma 5.4, we obtain $\left\|\mathbf{u}-\mathbf{u}^{n+1}\right\|_{\ell_{2}} \leq \frac{1}{2} \eta_{n}=\eta_{n+1}$ as well as

$$
\begin{aligned}
\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n+1}\right)\right\|_{\infty} & \leq r_{\text {best }}\left(\mathbf{u}, \frac{1}{2}\left(1+(1+\alpha) \kappa_{\mathbb{E}}\right)^{-1} \alpha \eta_{n}\right) \\
& \leq r_{\text {best }}\left(\mathbf{u},\left(1+(1+\alpha) \kappa_{\mathbb{E}}\right)^{-1} \alpha \eta_{n+1}\right),
\end{aligned}
$$

and the statement thus follows by induction.
Note that Theorem 5.5 yields rank bounds only for each $\mathbf{u}^{n}, n \in \mathbb{N}_{0}$, but not for the intermediate results produced by $\mathcal{F}_{n}$. The possible growth of such ranks of
intermediate results depends strongly on the problem and on the particular method used for $\mathcal{F}_{n}$. We return to this point in Section 5.2.3.

Remark 5.6. The precise meaning of the rank bound in (5.14) depends on the low-rank approximability of $\mathbf{u}$.
(i) Under the assumption of algebraic decay of best low-rank approximation errors with respect to ranks, that is,
with $C, s>0$, from (5.14) we obtain

$$
\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty} \lesssim d^{1 /(2 s)} \eta_{n}^{-1 / s}
$$

Ranks are thus ensured to increase at the optimal rate with respect to the achieved accuracy, where a moderate algebraic dependence on $d$ is possible.
(ii) The most interesting case for low-rank methods is that of best approximation errors with exponential-type decay with respect to ranks, as obtained for certain elliptic problems in Section 4.2. Assuming that

$$
\begin{equation*}
\inf _{\left\|\operatorname{rank}_{\mathbb{E}}(\mathbf{w})\right\|_{\infty} \leq r}\|\mathbf{u}-\mathbf{w}\|_{\ell_{2}} \leq C e^{-c r^{\beta}} \tag{5.15}
\end{equation*}
$$

with some $c, C>0$ and $\beta>0$, recalling that $\kappa_{\mathbb{E}}=\sqrt{2 d-3}$, we obtain

$$
\begin{aligned}
\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty} & \leq\left[\frac{1}{c} \ln \left(C(1+(1+\alpha) \sqrt{2 d-3})\left(\alpha \eta_{n}\right)^{-1}\right)\right]^{1 / \beta} \\
& \lesssim\left(\ln d+\left|\ln \eta_{n}\right|\right)^{1 / \beta}
\end{aligned}
$$

Remark 5.7. To quantify low-rank approximability, we can also introduce approximation classes generalizing Definition 5.2 based on strictly increasing sequences $\gamma=\left(\gamma_{n}\right)_{n \in \mathbb{N}_{0}}$ with $\gamma_{0}=1, \gamma_{n} \rightarrow \infty$, with membership of $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ in the corresponding class $\mathcal{A}_{\mathbb{E}}(\gamma)$ determined by the condition

$$
\|\mathbf{v}\|_{\mathcal{A}_{\mathbb{E}}(\gamma)}=\sup _{r \in \mathbb{N}_{0}} \gamma_{r} \inf \left\{\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}}:\left\|\operatorname{rank}_{\mathbb{E}}(\mathbf{w})\right\|_{\infty} \leq r\right\}<\infty .
$$

This allows for a more fine-grained analysis of low-rank approximability: provided that $\gamma$ grows at most exponentially, which corresponds to a restriction to $\beta \leq 1$ in (5.15), we have in particular the additional estimate

$$
\left\|\mathbf{u}^{n}\right\|_{\mathcal{A}_{\mathbb{B}}(\gamma)} \lesssim\|\mathbf{u}\|_{\mathcal{A}_{\mathbb{E}}(\gamma)}, \quad n \in \mathbb{N}_{0}
$$

which shows that the generated low-rank approximations inherit the low-rank approximability of $\mathbf{u}$. For details, see Bachmayr and Dahmen (2015, Sec. 3.1).

### 5.2.2. Soft thresholding

An alternative approach to controlling hierarchical ranks of iterates, proposed and analysed in Bachmayr and Schneider (2017), is based on the soft thresholding
procedure $\mathrm{ST}_{\delta}$, with $\delta>0$, discussed in Section 2.7.2. In this case the analysis is more strongly tied to the underlying fixed point iteration, and for the basic analysis we thus assume that $\mathcal{F}_{n}=\mathcal{F}$, where $\mathbf{u}=\mathcal{F}(\mathbf{u})$ and

$$
\begin{equation*}
\|\mathcal{F}(\mathbf{v})-\mathcal{F}(\mathbf{w})\|_{\ell_{2}} \leq \rho\|\mathbf{v}-\mathbf{w}\|_{\ell_{2}} \quad \text { for all } \mathbf{v}, \mathbf{w} \in \ell_{2}\left(\mathbb{N}^{d}\right) \tag{5.16}
\end{equation*}
$$

with $\rho \in(0,1)$. This immediately implies

$$
\begin{equation*}
\|\mathbf{u}-\mathcal{F}(\mathbf{v})\|_{\ell_{2}} \leq \rho\|\mathbf{u}-\mathbf{v}\|_{\ell_{2}} \quad \text { for all } \mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right) \tag{5.17}
\end{equation*}
$$

Note that the following applies with any $\rho \in(0,1)$, and we do not need a stronger error reduction assumption as in (5.12). Thus, in the present case, $\mathcal{F}$ satisfying (5.17) can always be realized by a single step of Richardson iteration.

Let us first observe that by the non-expansiveness of $\mathrm{ST}_{\delta}$ according to Proposition 2.40, the mapping $\mathrm{ST}_{\delta} \circ \mathcal{F}$ retains the same contractivity property as $\mathcal{F}$ for any $\delta>0$. This leads to the following result; for a proof, see Bachmayr and Schneider (2017, Lem. 4.1).

Lemma 5.8. Let $\mathcal{F}$ satisfy (5.16) with $\rho \in(0,1)$, and let $\delta>0$. Then $\mathrm{ST}_{\delta} \circ \mathcal{F}$ is also a contraction with Lipschitz constant $\rho$ on $\ell_{2}\left(\mathbb{N}^{d}\right)$, and the unique fixed points $\mathbf{u}$ of $\mathcal{F}$ and $\mathbf{u}^{\delta}$ of $\mathrm{ST}_{\delta} \circ \mathcal{F}$ satisfy

$$
(1+\rho)^{-1}\left\|\mathbf{u}-\operatorname{ST}_{\delta}(\mathbf{u})\right\|_{\ell_{2}} \leq\left\|\mathbf{u}-\mathbf{u}^{\delta}\right\|_{\ell_{2}} \leq(1-\rho)^{-1}\left\|\mathbf{u}-\mathrm{ST}_{\delta}(\mathbf{u})\right\|_{\ell_{2}} .
$$

As a consequence, on the one hand, by contractivity of $\mathrm{ST}_{\delta} \circ \mathcal{F}$ for each fixed $\delta>0$ we obtain

$$
\left\|\mathbf{u}^{\delta}-\left(\mathrm{ST}_{\delta} \circ \mathcal{F}\right)(\mathbf{v})\right\|_{\ell_{2}} \leq \rho\left\|\mathbf{u}^{\delta}-\mathbf{v}\right\|_{\ell_{2}} \quad \text { for all } \mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)
$$

In other words, the thresholded iteration always converges at the same rate as the original one, but to a modified fixed point. On the other hand, combining Lemma 5.8 with Theorem 2.42 yields

$$
\left\|\mathbf{u}-\mathbf{u}^{\delta}\right\|_{\ell_{2}} \lesssim\left\|\mathbf{u}-\mathrm{ST}_{\delta}(\mathbf{u})\right\|_{\ell_{2}} \rightarrow 0 \quad \text { as } \delta \rightarrow 0
$$

Rank bounds can be obtained with the aid of the following lemma, which is obtained in Bachmayr and Schneider (2017, Lem. 4.3) by a generalization of an observation in Cohen et al. (2001, Lem. 5.1), adapted to soft thresholding of sequences as in Dahlke, Fornasier and Raasch (2012).

Lemma 5.9. Let $\mathbf{u}, \mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$, let $\delta>0$ and let $\eta>0$ be such that $\|\mathbf{u}-\mathbf{v}\|_{\ell_{2}} \leq \eta$. Then, for all $e \in \mathbb{E}$,

$$
\begin{equation*}
\operatorname{rank}_{[e]}\left(\operatorname{ST}_{\delta}(\mathbf{v})\right) \leq \frac{4 \eta^{2}}{\delta^{2}}+\#\left\{k \in \mathbb{N}: \sigma_{k}^{[e]}(\mathbf{u})>\delta / 2\right\} \tag{5.18}
\end{equation*}
$$

This means that rank estimates can be obtained by balancing the two terms on the right in (5.18) depending on the currently achieved accuracy $\eta$. The second term is a function of $\delta$ that depends on decay of matricization singular values of the exact solution $\mathbf{u}$.

A scheme for appropriate adjustment of thresholding parameters that does not require explicit knowledge of the singular value decay is given in Bachmayr and Schneider (2017). It requires a functional $\mathcal{E}: \ell_{2}\left(\mathbb{N}^{d}\right) \rightarrow[0, \infty)$ such that with some $\gamma, \Gamma>0$, we have

$$
\gamma\|\mathbf{u}-\mathbf{v}\|_{\ell_{2}} \leq \mathcal{E}(\mathbf{v}) \leq \Gamma\|\mathbf{u}-\mathbf{v}\|_{\ell_{2}} \quad \text { for all } \mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)
$$

In particular, when solving $\mathbf{A u}=\mathbf{f}$, where $\mathbf{A}$ is an isomorphism on $\ell_{2}\left(\mathbb{N}^{d}\right)$, we may take $\mathcal{E}(\mathbf{v})=\|\mathbf{A v}-\mathbf{f}\|_{\ell_{2}}$ with $\gamma=\left\|\mathbf{A}^{-1}\right\|_{\ell_{2} \rightarrow \ell_{2}}^{-1}, \Gamma=\|\mathbf{A}\|_{\ell_{2} \rightarrow \ell_{2}}$.

The scheme proceeds as follows. With $\mathbf{u}^{0}=0$ and $\delta_{0}=(\# \mathbb{E})^{-1}\left\|\mathcal{F}\left(\mathbf{u}^{0}\right)\right\|_{\ell_{2}}$, for each $n=0,1,2, \ldots$ :

$$
\begin{align*}
& \text { Set } \mathbf{u}^{n+1}=\left(\mathrm{ST}_{\delta_{n}} \circ \mathcal{F}\right)\left(\mathbf{u}^{n}\right)  \tag{5.19a}\\
& \text { If }\left\|\mathbf{u}^{n+1}-\mathbf{u}^{n}\right\|_{\ell_{2}} \leq \frac{(1-\rho)}{2 \Gamma \rho} \mathcal{E}\left(\mathbf{u}^{n+1}\right) \text {, then set } \delta_{n+1}=\frac{\delta_{n}}{2}  \tag{5.19b}\\
& \text { otherwise, set } \delta_{n+1}=\delta_{n}
\end{align*}
$$

The basic idea is that by contractivity of $\mathcal{F}$,

$$
\frac{\rho}{1-\rho}\left\|\mathbf{u}^{n+1}-\mathbf{u}^{n}\right\|_{\ell_{2}}
$$

is an upper bound for $\left\|\mathbf{u}^{\delta_{n}}-\mathbf{u}^{n+1}\right\| \ell_{2}$, whereas $\mathcal{E}\left(\mathbf{u}^{n+1}\right)$ is proportional to $\left\|\mathbf{u}-\mathbf{u}^{n+1}\right\| \ell_{2}$. The condition in (5.19b) thus ensures that the thresholding parameter is adjusted when the iterates are closer to $\mathbf{u}^{\delta_{n}}$ than to $\mathbf{u}$. The following convergence result with rank estimates for this iteration is shown in Bachmayr and Schneider (2017, Thm 5.1, Remark 5.6). On $\mathcal{F}$, we assume here only (5.16) with any $\rho \in(0,1)$.

Theorem 5.10. The iterates defined by (5.19) satisfy $\left\|\mathbf{u}-\mathbf{u}^{n}\right\|_{\ell_{2}} \rightarrow 0$. In addition we have the following.
(i) If $\sigma^{[e]}(\mathbf{u}) \in \mathcal{A}^{s}$ for an $s>0$ for all $e \in \mathbb{E}$, then

$$
\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty} \leq C_{1} d^{2+1 / s} \max _{e \in \mathbb{E}}\left\|\sigma^{[e]}(\mathbf{u})\right\|_{\mathcal{A}^{s}}^{1 / s}\left\|\mathbf{u}-\mathbf{u}^{n}\right\|_{\ell_{2}}^{-1 / s}
$$

where $C_{1}$ depends on $\rho, \gamma, \Gamma$ and $s$.
(ii) If there exist $c, C, \beta>0$ such that $\sigma_{k}^{[e]}(\mathbf{u}) \leq C e^{-c k^{\beta}}$ for all $k$ and $e \in \mathbb{E}$, then

$$
\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty} \leq C_{2} d^{2}\left(\ln d+\left|\ln \left\|\mathbf{u}-\mathbf{u}^{n}\right\|_{\ell_{2}}\right|\right)^{1 / \beta}
$$

where $C_{2}$ depends on $\rho, \gamma, \Gamma, c, C$ and $\beta$.
In both cases (i) and (ii), there exists $C_{3}$ depending on the same parameters as $C_{1}$ and $C_{2}$, respectively, such that for each $\varepsilon>0$ we have $\left\|\mathbf{u}-\mathbf{u}^{n}\right\|_{\ell_{2}} \leq \varepsilon$ when $n \geq C_{3}(|\ln \varepsilon|+\ln d) \ln d$.

Note that by Theorem 2.36, the assumptions on singular value decay made in Theorem 5.10 are closely connected to the assumptions on the decay of total approximation errors made in Remark 5.6. Theorem 5.10 thus yields very similar
near-optimal rank estimates to Theorem 5.5, but, in view of Remark 5.6, involving less favourable additional constants. Moreover, (5.19) is linearly convergent, where the number of required iterations may deteriorate logarithmically with respect to $d$.

Remark 5.11. Theorem 5.10 can be extended to the case where the evaluation of $\mathcal{F}$ and $\mathcal{E}$ is done not exactly but with controlled relative errors, for instance when the action of $\mathbf{A}$ is performed only approximately (as required when working with Riesz basis representations on sequence spaces with approximate rescaling as in Section 4.4). For details we refer to Bachmayr and Schneider (2017, Sec. 5.1).

### 5.2.3. Comparison of rank control strategies

The two strategies for obtaining rank bounds discussed in Sections 5.2.1 and 5.2.2 can both be combined with quite general error reduction steps. As a basic example, let us now consider solving a linear system $\mathbf{A u}=\mathbf{f}$ on $\ell_{2}\left(\mathbb{N}^{d}\right)$, where $\mathbf{A}$ is a symmetric elliptic isomorphism on $\ell_{2}\left(\mathbb{N}^{d}\right)$, with error reduction done by Richardson iteration. Note that non-symmetric elliptic operators can be treated completely analogously to what follows, but lead to different convergence estimates for Richardson iteration.

In the case of the iteration (5.19) with soft thresholding, we may simply take

$$
\begin{equation*}
\mathcal{F}(\mathbf{v})=\mathbf{v}-\omega(\mathbf{A} \mathbf{v}-\mathbf{f}), \tag{5.20}
\end{equation*}
$$

which, with appropriately chosen step size parameter $\omega>0$, yields (5.16) and (5.17) with a fixed error reduction factor $\rho \in(0,1)$. Thus, while applying $\mathcal{F}$ as in (5.20) generally increases the ranks of iterates, the subsequent soft thresholding step immediately returns this intermediate result to a new iterate that satisfies the rank bounds of Theorem 5.10.

The situation is more delicate in the case of the iteration (5.13) with HSVD truncation, since here $\rho$ needs to satisfy condition (5.12), which becomes more stringent with increasing $d$. To satisfy this requirement, for each $n$ we can define the mapping $\mathcal{F}_{n}$ in (5.13) by

$$
\begin{equation*}
\mathcal{F}_{n}(\mathbf{v})=\mathbf{v}^{J}, \quad \text { where } \mathbf{v}^{j+1}=\mathbf{v}^{j}-\omega\left(\mathbf{A} \mathbf{v}^{j}-\mathbf{f}\right) \text { for } j \in \mathbb{N}_{0} \text { with } \mathbf{v}^{0}=\mathbf{v}, \tag{5.21}
\end{equation*}
$$

where the stopping index $J$ is chosen to achieve the required total error reduction. With appropriately chosen $\omega$, by the standard convergence analysis of Richardson iteration, (5.12) is satisfied when

$$
J=\zeta \operatorname{cond}(\mathbf{A}) \ln d
$$

with some $\zeta>0$. In this form, (5.13) consists of an outer iteration with rank bounds and the inner iteration (5.21). Unlike those for the iterates of the outer iteration, the rank bounds that can be achieved for the iterates $\mathbf{v}^{1}, \ldots, \mathbf{v}^{J}$ in (5.21) depend strongly on the available knowledge on $\mathbf{A}$ and $\mathbf{f}$.

Remark 5.12. When $\mathbf{A}$ is a basis representation of an operator, in general each column of $\mathbf{A}$ has infinitely many non-zero entries. For this reason, in this case
the iterations (5.20) and (5.21), which are then equivalent to iterating on the entire represented function spaces, cannot be carried out in practice in this form. However, as one possible strategy for the adaptive refinement of discretizations, these iterations can be carried out approximately, with the result of applying $\mathbf{A}$ (as well as $\mathbf{f}$ ) approximated in each step by a sequence with finitely many non-zero entries. Such adaptive methods are discussed in Section 5.3.2.

Example 5.13. To give a simple example, let $\mathbf{A}$ and $\mathbf{f}$ be represented with fixed rank parameters bounded componentwise by $R_{\mathbf{A}} \in \mathbb{N}$ and $r_{\mathbf{f}} \in \mathbb{N}$, respectively. Then the largest representation ranks of an intermediate result arising with soft thresholding via (5.20) are bounded for each $n$ by $\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty}\left(R_{\mathbf{A}}+1\right)+r_{\mathbf{f}}$.

In the case of (5.13), for the largest representation rank of an iterate arising in (5.21), in general we only have the bound

$$
\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty} \frac{R_{\mathbf{A}}^{J+1}-1}{R_{\mathbf{A}}-1}+r_{\mathbf{f}} \frac{R_{\mathbf{A}}^{J}-1}{R_{\mathbf{A}}-1} \lesssim\left(\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty}+r_{\mathbf{f}}\right) R_{\mathbf{A}}^{\zeta \operatorname{cond}(\mathbf{A}) \ln d}
$$

Note that $R_{\mathbf{A}}^{\zeta \operatorname{cond}(\mathbf{A}) \ln d}=d^{\zeta \operatorname{cond}(\mathbf{A}) \ln R_{\mathbf{A}}}$, which means that in this case rank bounds that depend polynomially on $d$ require $d$-independent $\operatorname{cond}(\mathbf{A})$ and $R_{\mathbf{A}}$.

Remark 5.14. The above discussion with uniform representation rank bounds for operators does not apply directly to preconditioned operators on Sobolev spaces. Recall that in terms of Riesz bases of $H_{0}^{1}\left(\square_{d}\right)$ as in Remark 3.10, for $A: H_{0}^{1}\left(\square_{d}\right) \rightarrow$ $H^{-1}\left(\square_{d}\right)$ we obtain operator representations of the form $\mathbf{A}=\mathbf{D T D}$, where the diagonal operator $\mathbf{D}$ does not have an explicit low-rank form. By Proposition 4.11, D can be replaced by an equivalent diagonal scaling operator $\tilde{\mathbf{D}}$ that is still of infinite rank but has efficient low-rank approximations $\tilde{\mathbf{D}}_{\eta, \Lambda}$ of relative error $\eta>0$ on the selected subset of basis indices $\Lambda$. We thus work with the low-rank representations $\tilde{\mathbf{D}}_{\eta, \Lambda} \mathbf{T} \tilde{\mathbf{D}}_{\eta, \Lambda}$ approximating $\tilde{\mathbf{D}} \mathbf{T} \tilde{\mathbf{D}}$. This means, on the one hand, that the error reduction procedures discussed above necessarily need to be applied inexactly. On the other hand, since the representation rank of $\tilde{\mathbf{D}}_{\eta, \Lambda}$ depends on $\Lambda$, it is desirable to refine the discretization determined by $\Lambda$ in parallel to the low-rank approximation of the solution to avoid unnecessarily large ranks.

### 5.3. Adaptive refinement of discretizations

Refining discretizations based on computable a posteriori error estimates is of particular interest in high-dimensional problems: in addition to the necessity of adaptively resolving non-smooth solutions, known a priori estimates are often difficult to apply in the high-dimensional case due to unknown dimension dependences in the arising constants.

An important requirement for the practicality of a posteriori estimates for large $d$, however, is that both the costs of their evaluation and the constants arising in the error bounds depend only moderately on $d$. For error estimates that are $d$-robust in
this sense, the next question is whether one can use them to drive a similarly robust adaptive refinement of discretizations.

Remark 5.15. Whether standard mesh-based approaches for error estimation in finite element discretizations, such as the classical residual error estimator, can be suitably adapted to high-dimensional finite elements in low-rank formats remains largely an open problem. A key difficulty, closely related to Remark 3.2, is establishing $d$-robustness of estimates. A finite element error estimate for tensor train representations using a global auxiliary flux reconstruction problem was proposed in Dolgov and Vejchodský (2021); it has the disadvantage, however, that for large $d$ the auxiliary problem can be substantially more expensive to solve than the original problem.

In what follows, we focus on error estimation and discretization refinement for low-rank methods using basis expansions, which are comparably well studied. Here we start from a Riesz basis representation of the problem on $\ell_{2}\left(\mathbb{N}^{d}\right)$. We then select finite index sets $\Lambda \subset \mathbb{N}^{d}$ such that the span of the corresponding basis elements provides sufficiently accurate approximations.

These concepts have been used in the context of low-rank approximations, for instance in Bachmayr (2012a), Bachmayr and Dahmen (2015, 2016a,b), and Ali and Urban (2020) with wavelet Riesz bases. They appear in a similar form in Eigel, Pfeffer and Schneider (2017), Bachmayr, Cohen and Dahmen (2018) and Eigel, Marschall, Pfeffer and Schneider (2020) with orthonormal polynomial expansions as in Remark 3.3. For an overview, see also Bachmayr and Dahmen (2020).

Remark 5.16. As an additional constraint in our present setting of tensor representations, we restrict ourselves to discretizations defined by subsets of basis indices that have Cartesian product structure. For instance, discretizations in terms of basis representations on $\ell_{2}\left(\mathbb{N}^{d}\right)$ can be defined by $\Lambda \subset \mathbb{N}^{d}$ of the form

$$
\begin{equation*}
\Lambda=\searrow_{i=1}^{d} \Lambda^{(i)} \tag{5.22}
\end{equation*}
$$

with finite $\Lambda^{(i)} \subset \mathbb{N}$ for $i=1, \ldots, d$. In principle it is also conceivable to adapt discretizations individually to each basis vector in the mode frames of a tensor representation. However, since algorithms using HSVD representations require repeated orthogonalizations of mode frames, such a further adaptation is impractical: since the discretization subspaces of all vectors are merged by orthogonalizations, there is no benefit in terms of computational costs.

In what follows, we consider basic techniques for error estimation and discretization refinement in the context of low-rank approximation, as well as their use within iterative solvers based on HSVD truncation. The approaches for generating sparse approximations with tensor structure discussed in Section 5.3.1 and the residual approximations considered in Section 5.3.3 are in fact independent of any
particular adaptive scheme. The computational methods treated in Sections 5.3.2 and 5.3.4 both rely on an extension of the strategy based on HSVD truncation in Section 5.2.1. The combination with adaptive discretization refinement of the alternative soft thresholding-based approach for rank control of Section 5.2.2 is still open.

### 5.3.1. Mode-wise sparse approximations

We first consider a method for selecting index sets of the form (5.22) that capture the entries of largest absolute values of a given tensor, but avoid looking at all individual entries (which would necessarily scale exponentially with respect to the tensor order $d$ ). Here we assume tensors on $\mathbb{N}^{d}$ for ease of presentation, but the following is applicable to general countable Cartesian product index sets.

The following notion was introduced in Bachmayr (2012a) and analysed further in Bachmayr and Dahmen (2015).
Definition 5.17. For $i=1, \ldots, d$ and $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$, we define the sequences $\pi^{(i)}(\mathbf{v})=\left(\pi_{v}^{(i)}(\mathbf{v})\right)_{v \in \mathbb{N}}$, called contractions of $\mathbf{v}$, by

$$
\begin{equation*}
\pi_{v_{i}}^{(i)}(\mathbf{v})=\left(\sum_{v_{1}=1}^{\infty} \ldots \sum_{v_{i-1}=1}^{\infty} \sum_{v_{i+1}=1}^{\infty} \ldots \sum_{v_{d}=1}^{\infty}\left|\mathbf{v}\left[v_{1}, \ldots, v_{i-1}, v_{i}, v_{i+1}, \ldots, v_{d}\right]\right|^{2}\right)^{1 / 2} \tag{5.23}
\end{equation*}
$$

and in addition, we define

$$
\operatorname{supp}_{i} \mathbf{v}=\operatorname{supp} \pi^{(i)}(\mathbf{v}) \subseteq \mathbb{N}
$$

The contractions $\pi^{(i)}(\mathbf{v}), i=1, \ldots, d$, can also be interpreted as the sequences of $\ell_{2}$-norms of the rows of the matricizations mat ${ }_{i}(\mathbf{v})$ of $\mathbf{v}$. They can serve as criteria for assessing where the $\ell_{2}$-norm of $\mathbf{v}$ is concentrated on the high-dimensional index set $\mathbb{N}^{d}$ by considering only sequences on $\mathbb{N}$. An important point in this regard is the following observation, which states that the high-dimensional summation on the right in Definition 5.17 can be avoided by exploiting orthogonality properties of the HSVD form.

Proposition 5.18. For $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$, let $\mathbf{U}^{\alpha}, \alpha \in \mathbb{T}$, be the mode frames of an HSVD of $\mathbf{v}$, and let $\sigma_{k}^{\alpha}(\mathbf{v})$ be the associated sequences of singular values. Then

$$
\pi_{v}^{(i)}(\mathbf{v})=\left(\sum_{k=1}^{\operatorname{rank}_{i}(\mathbf{v})}\left|\mathbf{U}_{k}^{\{i\}}[v] \sigma_{k}^{\{i\}}\right|^{2}\right)^{1 / 2}, \quad v \in \mathbb{N}
$$

The following result from Bachmayr and Dahmen (2015) shows how the jointly sorted sequences $\pi^{(i)}(\mathbf{v})$ can be used to obtain finite product index sets of the form (5.22) for approximating $\mathbf{v}$ in $\ell_{2}$. These index sets are quasi-optimal among all those of product structure with respect to the sum of the numbers of indices in each tensor mode. Here, for $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ and $\Lambda \subset \mathbb{N}^{d}$, we define the restriction $\mathrm{R}_{\Lambda} \mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ as the sequence that is equal to $\mathbf{v}$ on $\Lambda$ and vanishes on $\mathbb{N}^{d} \backslash \Lambda$.

Proposition 5.19. For any $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ and $N \in \mathbb{N}$, select $N$ largest values in $\left\{\pi_{v}^{(i)}(\mathbf{v}): i=1, \ldots, d, v \in \mathbb{N}\right\}$, collect the corresponding indices $v \in \mathbb{N}$ of selected $\pi_{\nu}^{(i)}(\mathbf{v})$ in the respective sets $\Lambda_{N}^{(i)} \subset \mathbb{N}$ for $i=1, \ldots, d$, and define

$$
\Lambda_{N}(\mathbf{v})=\searrow_{i=1}^{d} \Lambda_{N}^{(i)}(\mathbf{v})
$$

Then, for each $N$,

$$
\begin{align*}
\left\|\mathbf{v}-\mathrm{R}_{\Lambda_{N}(\mathbf{v})} \mathbf{v}\right\|_{\ell_{2}} & \leq\left(\sum_{i=1}^{d} \sum_{v \in \mathbb{N} \backslash \Lambda_{N}^{(i)}(\mathbf{v})}\left|\pi_{v}^{(i)}(\mathbf{v})\right|^{2}\right)^{1 / 2} \\
& \leq \sqrt{d} \min \left\{\left\|\mathbf{v}-\mathrm{R}_{\hat{\Lambda}} \mathbf{v}\right\|_{\ell_{2}}: \hat{\Lambda}=\varliminf_{i=1}^{d} \hat{\Lambda}^{(i)}, \sum_{i=1}^{d} \# \hat{\Lambda}^{(i)} \leq N\right\} \tag{5.24}
\end{align*}
$$

Proof. The first estimate follows from the observation that for any $\Lambda=X_{i=1}^{d} \Lambda^{(i)}$,

$$
\begin{aligned}
\left\|\mathbf{v}-\mathrm{R}_{\Lambda} \mathbf{v}\right\|_{\ell_{2}}^{2} & \leq\left\|\mathbf{v}-\mathrm{R}_{\Lambda^{(1)} \times \mathbb{N}^{d-1}} \mathbf{v}\right\|_{\ell_{2}}^{2}+\cdots+\left\|\mathbf{v}-\mathrm{R}_{\mathbb{N}^{d-1} \times \Lambda^{(d)}} \mathbf{v}\right\|_{\ell_{2}}^{2} \\
& =\sum_{i=1}^{d} \sum_{v \in \mathbb{N} \backslash \Lambda^{(i)}}\left|\pi_{v}^{(i)}(\mathbf{v})\right|^{2} .
\end{aligned}
$$

By construction of $\Lambda_{N}(\mathbf{v})$,

$$
\sum_{i=1}^{d} \sum_{v \in \mathbb{N} \backslash \Lambda_{N}^{(i)}(\mathbf{v})}\left|\pi_{v}^{(i)}(\mathbf{v})\right|^{2}=\min _{\sum_{i=1}^{d} \# \hat{\Lambda}^{(i)} \leq N} \sum_{i=1}^{d} \sum_{v \in \mathbb{N} \backslash \hat{\Lambda}^{(i)}}\left|\pi_{v}^{(i)}(\mathbf{v})\right|^{2}
$$

Finally, we note that for $\hat{\Lambda}=X_{i=1}^{d} \hat{\Lambda}^{(i)}$, we have

$$
\sum_{\nu \in \mathbb{N} \backslash \hat{\Lambda}^{(i)}}\left|\pi_{\nu}^{(i)}(\mathbf{v})\right|^{2} \leq\left\|\mathbf{v}-\mathrm{R}_{\hat{\Lambda}^{\mathbf{v}}}\right\|_{\ell_{2}}^{2}
$$

for each $i \in\{1, \ldots, d\}$.
Based on Proposition 5.19, we define a procedure for producing approximations with finitely many non-zero entries of any desired accuracy $\eta>0$ for arbitrary input sequences.

Definition 5.20. For $\mathbf{v} \in \ell_{2}(\mathbf{v})$ and $\eta>0$, with $\Lambda_{N}(\mathbf{v})$ defined as in Proposition 5.19 , and with $N$ chosen as

$$
N=\min \left\{N \in \mathbb{N}:\left(\sum_{i=1}^{d} \sum_{v \in \mathbb{N} \backslash \Lambda_{N}^{(i)}(\mathbf{v})}\left|\pi_{v}^{(i)}(\mathbf{v})\right|^{2}\right)^{1 / 2} \leq \eta\right\},
$$

we define

$$
\operatorname{Coarse}_{\eta}(\mathbf{v})=\mathrm{R}_{\Lambda_{N}(\mathbf{v})} \mathbf{v}
$$

Note that Coarse $_{\eta}(\mathbf{v})$ requires the HSVD form of $\mathbf{v}$ for evaluating the contractions $\pi^{(i)}(\mathbf{v})$ via Proposition 5.18. The costs of computing the HSVD form dominate those for subsequently selecting $\Lambda_{N}(\mathbf{v})$ with the appropriate $N$. With Lemma 5.1, the quasi-optimality property (5.24) immediately yields a result analogous to the one for $\mathrm{Trunc}_{\eta}$, as in Lemma 5.4 for the support sizes produced by Coarse ${ }_{\eta}$. More importantly for the present purposes, however, the two procedures can also be combined, which leads to a construction of adaptive methods, considered next.

### 5.3.2. Adaptive basis refinement by approximate Richardson iteration

We now come to our first method to combine rank adaptation with adaptive discretization refinement, proposed in Bachmayr (2012a) and analysed in Bachmayr and Dahmen (2015). Here we consider general Riesz basis representations $\mathbf{A u}=\mathbf{f}$ on $\ell_{2}\left(\mathbb{N}^{d}\right)$, where $\mathbf{A}: \ell_{2}\left(\mathbb{N}^{d}\right) \rightarrow \ell_{2}\left(\mathbb{N}^{d}\right)$ is an elliptic isomorphism.

As noted in Remark 5.12, in this setting, the rank-controlled iteration (5.13) combined with error reduction by an inner Richardson iteration (5.21) becomes computationally feasible only when the required residuals are approximated by finitely supported sequences. We thus assume for the moment that we have some routine $\operatorname{Res}_{\eta}$ at our disposal that for each $\eta>0$ and $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ with \# supp $\mathbf{v}<\infty$ produces an approximation of the residual $\mathbf{A v}-\mathbf{f}$ (which generally has infinite support) such that

$$
\begin{equation*}
\left\|\operatorname{Res}_{\eta}(\mathbf{v})-(\mathbf{A v}-\mathbf{f})\right\|_{\ell_{2}} \leq \eta, \tag{5.25}
\end{equation*}
$$

and where $\operatorname{Res}_{\eta}(\mathbf{v})$ is in low-rank representation with finite support.
The basic idea is now to perform Richardson iteration on the Riesz basis representation, but with each required residual evaluation replaced by $\operatorname{Res}_{\eta}$ with appropriate $\eta>0$. Similarly to iterations in low-rank format, we face the problem that when the method is carried out in this form, the number of non-zero basis indices that are active in each tensor mode may increase too strongly. This can be remedied by complementing the rank truncation in (5.13) by an additional basis coarsening by the routine Coarse ${ }_{\eta}$ from Definition 5.20.

For this combination we have the following generalization of Lemma 5.4; for the proof, we refer to Bachmayr and Dahmen (2015, Thm 7). Note that for the exact solution $\mathbf{u}$, it is natural to assume that $\pi^{(i)}(\mathbf{u}) \in \mathcal{A}^{s}(\mathbb{N})$ for $i=1, \ldots, d$ with some $s>0$, since wavelet basis expansions as discussed in Section 3.2 typically lead to such algebraic convergence rates. These assumptions can be modified for more rapidly converging spectral approximations.

Lemma 5.21. Let $\mathbf{u}, \mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ with $\pi^{(i)}(\mathbf{u}) \in \mathcal{A}^{s}(\mathbb{N})$ for some $s>0$ for $i=$ $1, \ldots, d$, such that $\|\mathbf{u}-\mathbf{v}\|_{\ell_{2}} \leq \eta$. Let $\kappa_{\mathbb{E}}=\sqrt{\# \mathbb{E}}, \kappa_{d}=\sqrt{d}$. Let $\alpha>0$; then for

$$
\mathbf{w}=\operatorname{Coarse}_{\kappa_{d}\left(\kappa_{\mathrm{B}}+1\right)(1+\alpha) \eta}\left(\operatorname{Trunc}_{\kappa_{\mathbb{B}}}(1+\alpha) \eta(\mathbf{v})\right)
$$

we have

$$
\|\mathbf{u}-\mathbf{w}\|_{\ell_{2}} \leq C_{1}\left(\alpha, \kappa_{\mathbb{E}}, \kappa_{d}\right) \eta
$$

where $C_{1}\left(\alpha, \kappa_{\mathbb{E}}, \kappa_{d}\right)=\left(1+\kappa_{\mathbb{E}}(1+\alpha)+\kappa_{d}\left(\kappa_{\mathbb{E}}+1\right)(1+\alpha)\right)$, and moreover

$$
\left\|\operatorname{rank}_{\mathbb{E}}(\mathbf{w})\right\|_{\infty} \leq r_{\text {best }}(\mathbf{u}, \alpha \eta)
$$

as well as

$$
\begin{aligned}
& \sum_{i=1}^{d} \# \operatorname{supp}_{i} \mathbf{w} \leq 2 d \alpha^{-1 / s}\left(\sum_{i=1}^{d}\left\|\pi^{(i)}(\mathbf{u})\right\|_{\mathcal{A}^{s}}\right)^{1 / s} \eta^{-1 / s}, \\
& \sum_{i=1}^{d}\left\|\pi^{(i)}(\mathbf{w})\right\|_{\mathcal{A}^{s}} \leq C_{2}\left(\alpha, \kappa_{\mathbb{E}}, \kappa_{d}\right) \sum_{i=1}^{d}\left\|\pi^{(i)}(\mathbf{u})\right\|_{\mathcal{A}^{s}},
\end{aligned}
$$

where $C_{2}\left(\alpha, \kappa_{\mathbb{E}}, \kappa_{d}, d, s\right)=2^{s}\left(1+3^{s}\right)+2^{4 s} \alpha^{-1} C_{1}\left(\alpha, \kappa_{\mathbb{E}}, \kappa_{d}\right) d^{\max \{1, s\}}$.
The resulting rank- and basis-adaptive iterative method can be regarded as a generalization of the adaptive method or sparse approximation introduced in Cohen et al. (2002), which follows the same basic strategy. Let $\omega>0$ be chosen such that $\|\mathbf{I}-\omega \mathbf{A}\|_{\ell_{2} \rightarrow \ell_{2}} \leq \rho<1$. With $\kappa_{\mathbb{E}}=\sqrt{\# \mathbb{E}}=\sqrt{2 d-3}, \kappa_{d}=\sqrt{d}$ and $\alpha>0$, let

$$
\begin{aligned}
\hat{\rho} & =\frac{1}{1+(1+\alpha)\left(\kappa_{\mathbb{E}}+\kappa_{d}+\kappa_{\mathbb{E}} \kappa_{d}\right)}, \\
\theta_{1} & =\kappa_{d}\left(\kappa_{\mathbb{E}}+1\right)(1+\alpha) \hat{\rho}, \quad \theta_{2}=(1+\alpha) \kappa_{\mathbb{E}} \hat{\rho},
\end{aligned}
$$

as well as

$$
J=\min \left\{j \in \mathbb{N}: \rho^{j}(1+\omega j) \leq \frac{1}{2} \hat{\rho}\right\}
$$

Starting with $\mathbf{u}^{0}=0$, setting $\eta_{0}=\left\|\mathbf{A}^{-1}\right\|_{\ell_{2} \rightarrow \ell_{2}}\|\mathbf{f}\|_{\ell_{2}}$ and $\eta_{n}=2^{-n} \eta_{0}$ for $n=$ $0,1,2, \ldots$,

$$
\begin{align*}
\mathbf{v}^{n, 0} & =\mathbf{u}^{n}, \quad \mathbf{v}^{n, j+1}=\mathbf{v}^{n, j}-\omega \operatorname{Res}_{\rho^{j} \eta_{n}}\left(\mathbf{v}^{n, j}\right), j=0,1, \ldots, J-1,  \tag{5.26a}\\
\mathbf{u}^{n+1} & =\operatorname{Coarse}_{\theta_{1} \eta_{n+1}}\left(\operatorname{Trunc}_{\theta_{2} \eta_{n+1}}\left(\mathbf{v}^{n, J}\right)\right) . \tag{5.26b}
\end{align*}
$$

Remark 5.22. The inner iterations (5.26a) can be performed with additional rank truncation and basis coarsening steps with sufficiently small tolerances; see Bachmayr and Dahmen (2015, Alg. 5.1). These do not influence the analysis based on Lemma 5.21 but in general improve the quantitative performance.

The proof of the following convergence result with rank and discretization size bounds for (5.26) is given in Bachmayr and Dahmen (2015).

Theorem 5.23. The sequence $\left(\mathbf{u}^{n}\right)_{n \in \mathbb{N}_{0}}$ defined by (5.26) for each $n \in \mathbb{N}_{0}$ satisfies $\left\|\mathbf{u}-\mathbf{u}^{n}\right\|_{\ell_{2}} \leq \eta_{n}=2^{-n} \eta_{0}$ and

$$
\begin{equation*}
\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty} \leq r_{\text {best }}\left(\mathbf{u}, \alpha \hat{\rho} \eta_{n}\right) \tag{5.27}
\end{equation*}
$$

as well as

$$
\begin{align*}
& \sum_{i=1}^{d} \# \operatorname{supp}_{i} \mathbf{u}^{n} \lesssim\left(\sum_{i=1}^{d}\left\|\pi^{(i)}(\mathbf{u})\right\|_{\mathcal{A}^{s}}\right)^{1 / s} \eta_{n}^{-1 / s},  \tag{5.28}\\
& \sum_{i=1}^{d}\left\|\pi^{(i)}\left(\mathbf{u}^{n}\right)\right\|_{\mathcal{A}^{s}} \lesssim \sum_{i=1}^{d}\left\|\pi^{(i)}(\mathbf{u})\right\|_{\mathcal{A}^{s}},
\end{align*}
$$

with constants that depend only on $\alpha, s$ and algebraically on $d$.
Concerning the implications of the rank bound (5.27) when assuming algebraically or exponentially decaying best low-rank approximation errors, Remark 5.6 applies also in this case, with minor modifications to the $d$-dependence of the bounds.

Altogether, (5.26) provides a numerical scheme that operates on finitely many indices of the underlying Riesz basis (that is, on finitely supported coefficient sequences) in each step, and yet is guaranteed to converge to the exact solution of the PDE problem, represented by the coefficient sequence $\mathbf{u}$. After every rank truncation and basis coarsening step (5.26b), the resulting iterates $\mathbf{u}^{n}$ are guaranteed to have near-optimal ranks and support sizes compared to the respective best approximations of $\mathbf{u}$. A crucial point for the total computational costs of the method, however, is how precisely the residual approximation $\operatorname{Res}_{\eta}$ is realized. Similarly to Example 5.13, in particular, the support sizes should not increase too strongly during the inner iteration (5.26a).

### 5.3.3. Residual approximations

The adaptive algorithm (5.26) relies on a routine $\operatorname{Res}_{\eta}(\mathbf{v})$ that for each $\eta>0$ produces a finitely supported approximation of the residual $\mathbf{A v}-\mathbf{f}$ (which is generally a sequence on $\mathbb{N}^{d}$ with infinitely many non-zero entries) with error at most $\eta$. In addition to the direct use of $\operatorname{Res}_{\eta}(\mathbf{v})$ in an iterative scheme, this procedure also yields computable a posteriori error estimates. For example, for given $\mathbf{v}$, starting from a sufficiently large value of $\eta$, by successively halving $\eta$ until $\eta \leq \frac{1}{2}\left\|\operatorname{Res}_{\eta}(\mathbf{v})\right\|_{\ell_{2}}$ holds, we obtain

$$
\begin{equation*}
\left\|\operatorname{Res}_{\eta}(\mathbf{v})\right\|_{\ell_{2}} \approx\|\mathbf{A v}-\mathbf{f}\|_{\ell_{2}} \approx\|\mathbf{v}-\mathbf{u}\|_{\ell_{2}} \tag{5.29}
\end{equation*}
$$

Therefore, due to the Riesz basis property, such a residual approximation also yields a computable quantity that remains proportional to the total error in approximating the exact solution of the underlying PDE.

Whereas suitable approximations of $\mathbf{f}$ are problem-dependent but can typically be derived from properties of the represented functional $f$, the approximate application of operators can be performed by universal strategies, which we focus on here. For the efficient realization of such approximations, we consider a standard construction introduced in Cohen et al. (2001) that is adapted to the adaptivity in the low-rank context in Bachmayr and Dahmen (2015, 2016b). It is based on the following notion of sparse approximability of operators represented by bi-infinite matrices.

Definition 5.24. Let $\mathcal{I}$ be a countable index set, and let $s^{*}>0$. An operator $\mathbf{B}: \ell_{2}(\mathcal{I}) \rightarrow \ell_{2}(\mathcal{I})$ is called $s^{*}$-compressible if for any $s \in\left(0, s^{*}\right)$ there exist summable positive sequences $\left(\alpha_{q}\right)_{q \in \mathbb{N}_{0}},\left(\beta_{q}\right)_{q \in \mathbb{N}_{0}}$ such that for each $q \in \mathbb{N}_{0}$ there exists $\mathbf{B}_{q}$ with at most $\alpha_{q} 2^{q}$ non-zero entries per row and column satisfying $\| \mathbf{B}$ $\mathbf{B}_{q} \| \leq \beta_{q} 2^{-s q}$.

In particular, such approximations can be obtained for representations of differential (or integral) operators with respect to wavelet Riesz bases; see Cohen et al. (2001). For the adaptive compressibility of operators, the following generic strategy for sparse approximations is developed there based on Definition 5.24.

Remark 5.25. To obtain approximations $\mathbf{w}_{\eta}$ of $\mathbf{B v}$ such that $\left\|\mathbf{w}_{\eta}-\mathbf{B v}\right\|_{\ell_{2}} \leq \eta$, for $\mathbf{B}$ as in Definition 5.24 and $\mathbf{v} \in \ell_{2}(\mathcal{I})$ with $N=\# \operatorname{supp} \mathbf{v}<\infty$, for each fixed $s \in\left(0, s^{*}\right)$ we can proceed as follows. First, let $\mathcal{I}_{[q]} \subset \mathcal{I}$ be the supports of the best $2^{q}$-term approximations of $\mathbf{v}$ for $q=0,1, \ldots,\left\lceil\log _{2} N\right\rceil$, so that $\| \mathbf{v}-$ $\mathrm{R}_{\mathcal{I}_{[q]}} \mathbf{v} \|_{\ell_{2}} \leq|\mathbf{v}|_{\mathcal{A}^{s}} 2^{-s q}$. Then, for each $Q \leq\left\lceil\log _{2} N\right\rceil$, with $\mathcal{I}_{[-1]}=\emptyset$, consider the approximations

$$
\hat{\mathbf{B}}_{Q}(\mathbf{v})=\sum_{q=0}^{Q} \mathbf{B}_{Q-q} \mathrm{R}_{\mathcal{I}_{[q]} \backslash \mathcal{I}_{[q-1]}} \mathbf{v},
$$

where $\mathbf{B}_{Q-q}$ are the approximations of $\mathbf{B}$ from Definition 5.24. Note that in this approximation the most accurate (and expensive) approximations of the operator $\mathbf{B}$ are applied only to the largest entries of the input vector $\mathbf{v}$. By the triangle inequality, we have the error bound

$$
\begin{aligned}
& \left\|\mathbf{B} \mathbf{v}-\hat{\mathbf{B}}_{Q}(\mathbf{v})\right\|_{\ell_{2}} \\
& \quad \leq\|\mathbf{B}\|_{\ell_{2} \rightarrow \ell_{2}}\left\|\mathbf{v}-\mathrm{R}_{\mathcal{I}_{[Q]}} \mathbf{v}\right\|_{\ell_{2}}+\sum_{q=0}^{Q}\left\|\mathbf{B}-\mathbf{B}_{Q-q}\right\|_{\ell_{2} \rightarrow \ell_{2}}\left\|\mathrm{R}_{\left.\mathcal{I}_{[q]}\right] \mathcal{I}_{[q-1]}} \mathbf{v}\right\|_{\ell_{2}} .
\end{aligned}
$$

We now insert the estimates $\left\|\mathrm{R}_{\mathcal{I}_{[q]} \backslash \mathcal{I}_{[q-1]}} \mathbf{v}\right\|_{\ell_{2}} \leq\left\|\mathbf{v}-\mathrm{R}_{\mathcal{I}_{[q-1]}} \mathbf{v}\right\|_{\ell_{2}} \leq|\mathbf{v}|_{\mathcal{A}^{s}} 2^{-s(q-1)}$ as well as $\left\|\mathbf{B}-\mathbf{B}_{Q-q}\right\|_{\ell_{2} \rightarrow \ell_{2}} \leq \beta_{Q-q} 2^{-s(Q-q)}$ to obtain $\left\|\mathbf{B v}-\hat{\mathbf{B}}_{Q}(\mathbf{v})\right\|_{\ell_{2}} \lesssim|\mathbf{v}|_{\mathcal{A}^{s}} 2^{-s Q}$. Taking $Q$ as the smallest value such that $\left\|\mathbf{B v}-\hat{\mathbf{B}}_{Q}(\mathbf{v})\right\|_{\ell_{2}} \leq \eta$, we set $\mathbf{w}_{\eta}=\hat{\mathbf{B}}_{Q}(\mathbf{v})$. Then

$$
\begin{aligned}
\# \operatorname{supp} \mathbf{w}_{\eta} & \leq \sum_{q=0}^{Q} 2^{Q-q} \#\left(\mathcal{I}_{[q]} \backslash \mathcal{I}_{[q-1]}\right) \\
& \leq \sum_{q=1}^{Q} 2^{Q-q} 2^{q-1}+2^{Q} \lesssim 2^{Q} \lesssim|\mathbf{v}|_{\mathcal{A}^{s}}^{1 / s} \eta^{-1 / s} .
\end{aligned}
$$

In addition, we verify as in Cohen et al. (2001) that

$$
\left\|\mathbf{w}_{\eta}\right\|_{\mathcal{A}^{s}} \leq C_{\mathcal{A}^{s}}\|\mathbf{v}\|_{\mathcal{A}^{s}}
$$

with a constant $C_{\mathcal{A}^{s}}>0$ independent of $\mathbf{v}$ and $\eta$.

Next we turn to the generalization of this strategy to the setting of low-rank approximations, where we consider $\mathbf{A}: \ell_{2}\left(\mathbb{N}^{d}\right) \rightarrow \ell_{2}\left(\mathbb{N}^{d}\right)$ of the form

$$
\begin{equation*}
\mathbf{A}=\sum_{n_{1}=1}^{R_{1}} \cdots \sum_{n_{d}=1}^{R_{d}} \mathbf{c}\left[n_{1}, \ldots, n_{d}\right] \bigotimes_{i=1}^{d} \mathbf{A}_{n_{i}}^{(i)} \tag{5.30}
\end{equation*}
$$

with the core tensor $\mathbf{c}$ represented in hierarchical format. Our aim is to approximate Av for finitely supported $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ that is given in the same hierarchical tensor format. We assume that all component operators $\mathbf{A}_{n}^{(i)}: \ell_{2}(\mathbb{N}) \rightarrow \ell_{2}(\mathbb{N})$ for $i=1, \ldots, d, n=1, \ldots, R_{i}$, are $s^{*}$-compressible, with the corresponding approximations according to Definition 5.24 denoted by $\mathbf{A}_{n, q}^{(i)}, q \in \mathbb{N}_{0}$. We then consider approximations of the form

$$
\begin{equation*}
\hat{\mathbf{A}}=\sum_{n_{1}=1}^{R_{1}} \ldots \sum_{n_{d}=1}^{R_{d}} \mathbf{c}\left[n_{1}, \ldots, n_{d}\right] \bigotimes_{i=1}^{d} \hat{\mathbf{A}}_{n_{i}}^{(i)} \tag{5.31}
\end{equation*}
$$

For the definition of $\hat{\mathbf{A}}_{n}^{(i)}$ for $i=1, \ldots, d, n=1, \ldots, R_{i}$, we choose suitable $\Lambda_{q}^{(i)} \subset \mathbb{N}$ for $q=0, \ldots, Q_{i}$ such that $\bigcup_{q=0}^{Q_{i}} \Lambda_{q}^{(i)}=\operatorname{supp}_{i} \mathbf{v}$ and $\Lambda_{q}^{(i)} \cap \Lambda_{p}^{(i)}=\emptyset$ whenever $q \neq p$. We now take

$$
\begin{equation*}
\hat{\mathbf{A}}^{(i)}=\sum_{q=0}^{Q_{i}} \mathbf{A}_{n, q}^{(i)} \mathrm{R}_{\Lambda_{q}^{(i)}} \tag{5.32}
\end{equation*}
$$

For approximations of the form (5.31), (5.32), we have the following simple observation.

Proposition 5.26. Let $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ with $\# \operatorname{supp} \mathbf{v}<\infty$, let $\mathbf{A}$ be as in (5.30) and $\hat{\mathbf{A}}$ as in (5.31), (5.32). Then

$$
\|\mathbf{A v}-\hat{\mathbf{A}} \mathbf{v}\|_{\ell_{2}(\mathbb{N} d)} \leq \sum_{i=1}^{d} C_{\mathbf{A}}^{(i)} \sum_{n=1}^{R_{i}} \sum_{q=0}^{Q_{i}}\left\|\mathbf{A}_{n}^{(i)}-\mathbf{A}_{n, q}^{(i)}\right\|_{\ell_{2}(\mathbb{N}) \rightarrow \ell_{2}(\mathbb{N})}\left\|\mathrm{R}_{\Lambda_{q}^{(i)}} \pi^{(i)}(\mathbf{v})\right\|_{\ell_{2}(\mathbb{N})}
$$

with $C_{\mathbf{A}}^{(i)}>0, i=1, \ldots, d$, which can be determined from $\mathbf{A}$.

For a proof and a discussion of the constants $C_{\mathrm{A}}^{(i)}$, we refer to Bachmayr and Dahmen (2015, Lem. 4). We now follow a similar strategy to Remark 5.25, combining the bounds from Definition 5.24 for $\left\|\mathbf{A}_{n}^{(i)}-\mathbf{A}_{n, q}^{(i)}\right\|_{\ell_{2}(\mathbb{N}) \rightarrow \ell_{2}(\mathbb{N})}$ with choosing the sets $\Lambda_{q}^{(i)}$ for $i=1, \ldots, d$ based on $2^{q}$-term approximations of the respective (lower-dimensional) contractions $\pi^{(i)}(\mathbf{v})$. Proceeding in this manner, assuming that $\pi^{(i)}(\mathbf{v}) \in \mathcal{A}^{s}$ for $i=1, \ldots, d$ with an $s<s^{*}$, for each $\eta>0$ we obtain a
finitely supported $\mathbf{w}_{\eta}=\hat{\mathbf{A}} \mathbf{v}$ that maintains the structure and ranks of the low-rank representation of $\mathbf{A v}$, such that $\left\|\mathbf{w}_{\eta}-\mathbf{A v}\right\|_{\ell_{2}\left(\mathbb{N}^{d}\right)} \leq \eta$ and

$$
\begin{align*}
\# \operatorname{supp}_{i} \mathbf{w}_{\eta} & \leq C_{1}\left(\sum_{i=1}^{d}\left\|\pi^{(i)}(\mathbf{v})\right\|_{\mathcal{A}^{s}}\right)^{1 / s} \eta^{-1 / s},  \tag{5.33}\\
\left\|\pi^{(i)}\left(\mathbf{w}_{\eta}\right)\right\|_{\mathcal{A}^{s}} & \leq C_{2}\left\|\pi^{(i)}(\mathbf{v})\right\|_{\mathcal{A}^{s}}
\end{align*}
$$

for $i=1, \ldots, d$, where $C_{1}, C_{2}>0$ depend only on $\mathbf{A}$ and $s$ (Bachmayr and Dahmen 2015, Thm 8).

Altogether, the computation of approximate residuals is thus reduced to operations on single tensor components. Note that in this case the entries of $\pi^{(i)}(\mathbf{v})$ determine to what accuracy the respective individual columns of the component operators of $\mathbf{A}$ in mode $i$ are approximated.

Remark 5.27. Some further modifications are required for wavelet Riesz basis representations on $\ell_{2}\left(\mathrm{~V}^{d}\right)$ of problems on Sobolev spaces, which take the form (3.23). As noted in Remark 5.14, this can be replaced by (4.10) with the operator representation $\mathbf{D} \mathbf{T} \tilde{\mathbf{D}}$ with diagonal operators $\tilde{\mathbf{D}}$ from Proposition 4.11. Here $\mathbf{T}$ depends on the coefficients in the differential operator. In the case of the Laplacian, T has the explicit low-rank representation (3.22) with rank two, but its lower-dimensional component operators $\mathbf{T}_{1}$ generally have infinitely many non-zero entries in each column. For given $\mathbf{v} \in \ell_{2}\left(\mathrm{~V}^{d}\right)$, we thus consider approximations $\mathbf{w}_{\eta}$ such that $\left\|\tilde{\mathbf{D}} \mathbf{T} \tilde{\mathbf{D}} \mathbf{v}-\mathbf{w}_{\eta}\right\| \leq \eta$ of the form

$$
\begin{equation*}
\mathbf{w}_{\eta}=\tilde{\mathbf{D}}_{\eta_{1}, \Lambda_{1}} \tilde{\mathbf{T}} \tilde{\mathbf{D}}_{\eta_{2}, \Lambda_{2}} \mathbf{v} \tag{5.34}
\end{equation*}
$$

Here $\tilde{\mathbf{D}}_{\eta_{j}, \Lambda_{j}}, j=1,2$, is defined as in Proposition 4.11 with suitable $\eta_{j} \approx \eta$, where $\Lambda_{2} \supseteq \operatorname{supp} \mathbf{v}$ and $\Lambda_{1} \supseteq \operatorname{supp}\left(\tilde{\mathbf{T}} \tilde{\mathbf{D}}_{\eta_{2}, \Lambda_{2}} \mathbf{v}\right)$. The operator $\tilde{\mathbf{T}}$ is obtained by componentwise sparse approximations analogously to (5.31), (5.32), additionally taking into account the effect of the basis rescaling by $\tilde{\mathbf{D}}$. We thus combine sparse approximations of $\mathbf{T}$ with low-rank approximations of $\tilde{\mathbf{D}}$. For details and the corresponding estimates replacing (5.33), we refer to Bachmayr and Dahmen (2016b).

Remark 5.28. The estimates (5.33) can now be combined with the estimates from Theorem 5.23 for the iterates of the adaptive scheme (5.26). As noted above, for the evolution of the ranks of the iterates $\mathbf{v}^{n, j}$ of the inner iterations (5.26a), the conclusions of Remark 5.6 apply. As shown in Bachmayr and Dahmen (2016b, Sec. 6.5), if also $\left\|\pi^{(i)}(\mathbf{f})\right\|_{\mathcal{A}^{s}}<\infty$ for $i=1, \ldots, d$ and certain technical assumptions on the low-rank representation of $\mathbf{A}$ hold true (which are satisfied in the case of

Remark 5.27), combining (5.28) with (5.33) we obtain estimates

$$
\begin{aligned}
& \quad \sum_{i=1}^{d} \# \operatorname{supp}_{i} \mathbf{v}^{n, j} \leq C_{1, j}\left(\sum_{i=1}^{d} \max \left\{\left\|\pi^{(i)}(\mathbf{u})\right\|_{\mathcal{A}^{s}},\left\|\pi^{(i)}(\mathbf{f})\right\|_{\mathcal{A}^{s}}\right\}\right)^{1 / s} \eta_{n}^{-1 / s}, \\
& \sum_{i=1}^{d}\left\|\pi^{(i)}\left(\mathbf{v}^{n, j}\right)\right\|_{\mathcal{A}^{s}} \leq C_{2, j} \sum_{i=1}^{d} \max \left\{\left\|\pi^{(i)}(\mathbf{u})\right\|_{\mathcal{A}^{s}},\left\|\pi^{(i)}(\mathbf{f})\right\|_{\mathcal{A}^{s}}\right\} .
\end{aligned}
$$

Here $C_{1, j}, C_{2, j}>0$ depend polynomially on $d$, but may in general increase exponentially with respect to the inner iteration number $j$. After a fixed number of inner iteration steps, however, by the complexity reduction step (5.26b) we return to the estimates (5.28). In summary, we thus retain quasi-optimal sums of mode-wise support sizes throughout the entire iteration (5.26).

### 5.3.4. Iteratively refined Galerkin discretizations

An alternative construction of adaptive methods that is conceptually closer to adaptive finite element methods is based on solving successively refined Galerkin discretizations. Methods of this type with quasi-optimal computational costs for sparse approximation by wavelets were developed and analysed in Cohen et al. (2001).

We first sketch the basic approach. Here we need to assume, in addition to boundedness and ellipticity, that $\mathbf{A}$ is a symmetric operator, so $\|\mathbf{v}\|_{\mathbf{A}}=\langle\mathbf{A v}, \mathbf{v}\rangle^{1 / 2}$ for $\mathbf{v} \in \ell_{2}\left(\mathbb{N}^{d}\right)$ defines a norm on $\ell_{2}\left(\mathbb{N}^{d}\right)$; note that symmetry is not required for the scheme (5.26). The method proceeds by stepwise refinement of finite index sets $\Lambda^{n} \subset \mathbb{N}^{d}, n=0,1,2, \ldots$, starting from a given $\Lambda^{0}$. For each $n$, we first determine the Galerkin solution $\mathbf{u}_{\Lambda^{n}}$, which is uniquely defined by the conditions $\operatorname{supp} \mathbf{u}_{\Lambda^{n}} \subseteq \Lambda^{n}$ and

$$
\begin{equation*}
\mathrm{R}_{\Lambda^{n}} \mathbf{A} \mathbf{u}_{\Lambda^{n}}=\mathrm{R}_{\Lambda^{n}} \mathbf{f} \tag{5.35a}
\end{equation*}
$$

With $\mathbf{r}_{\Lambda^{n}}=\mathbf{A} \mathbf{u}_{\Lambda^{n}}-\mathbf{f}$ and a fixed $\alpha \in(0,1]$, we select the smallest $\hat{\Lambda} \supseteq \Lambda^{n}$ such that

$$
\begin{equation*}
\left\|\mathrm{R}_{\hat{\Lambda}} \mathbf{r}_{\Lambda^{n}}\right\|_{\ell_{2}} \geq \alpha\left\|\mathbf{r}_{\Lambda^{n}}\right\|_{\ell_{2}} \tag{5.35b}
\end{equation*}
$$

and set $\Lambda^{n+1}=\hat{\Lambda}$. Here, for computational purposes, $\mathbf{r}_{\Lambda^{n}}$ is again replaced by a finitely supported approximation. As shown in Cohen et al. (2001, Lem. 4.1), (5.35b) ensures

$$
\left\|\mathbf{u}-\mathbf{u}_{\Lambda^{n+1}}\right\|_{\mathbf{A}} \leq \rho\left\|\mathbf{u}-\mathbf{u}_{\Lambda^{n}}\right\|_{\mathbf{A}}
$$

with $\rho \in(0,1)$ that depends on $\alpha$ and cond(A). Cohen et al. (2001) combined this error reduction with a basis coarsening by application of Lemma 5.1 to best $n$-term approximation as in (5.9). Assuming $\mathbf{u} \in \mathcal{A}^{s}$ for some $s>0$, this again leads to a method using a quasi-optimal number of non-zero coefficients for each iterate.

Gantumur, Harbrecht and Stevenson (2007) adapted the method to avoid the coarsening steps: provided that $\alpha$ in (5.35b) is chosen in $\left(0, \operatorname{cond}^{-1 / 2}(\mathbf{A})\right)$, the
index sets $\Lambda^{n}$ obtained from (5.35b) have quasi-optimal cardinality without further coarsening. This variation of the method, termed the adaptive wavelet-Galerkin method, yields improved quantitative performance. For an overview, see also Stevenson (2009).

The above approach, based on successively refined Galerkin discretizations, has been adapted to the setting of low-rank tensor approximations in Ali and Urban (2020). The Galerkin solutions $\mathbf{u}_{\Lambda^{n}}$ are approximated by a rank-truncated PCG method, and finitely supported low-rank approximations of $\mathbf{r}_{\Lambda^{n}}$ are computed by the techniques discussed in Section 5.3.3. Index sets $\hat{\Lambda}$ with product structure in (5.35b) are obtained by applying Coarse $\eta_{\eta}$ to $\mathbf{r}_{\Lambda^{n}}$ with a suitable tolerance $\eta>0$, as in Definition 5.3.

To obtain rank and discretization size estimates as in Theorem 5.23, this variant also relies on the combination of Coarse $\eta_{\eta}$ and $\mathrm{Trunc}_{\eta}$ as in (5.26b). In principle, the scheme is thus of a similar structure to the adaptive low-rank Richardson iteration (5.26), but with the inner iteration (5.26a) replaced by sufficiently many steps of solving a Galerkin discretization and selecting new degrees of freedom via (5.35b). The conclusions concerning rank and discretization size bounds for all iterates obtained in Ali and Urban (2020) are analogous to those for (5.26), but with quantitative advantages concerning the practical performance.

In contrast to the setting of best $n$-term approximations, the question remains open whether it is possible to obtain quasi-optimal discretization sizes without applying a basis coarsening step as with Coarse ${ }_{\eta}$ in (5.26b). The reason is that the effective value of $\alpha$ obtained in (5.35b) when selecting $\hat{\Lambda}$ using Coarse ${ }_{\eta}$ cannot in general be chosen sufficiently small for larger $d$, since the selected index sets are only quasi-optimal up to a factor $\sqrt{d}$ as in (5.24). For a detailed discussion, see Ali and Urban (2020, Sec. 4.7).

### 5.4. Computational complexity

In Section 5.2 we consider two strategies, based on HSVD truncation and on soft thresholding, for maintaining quasi-optimal ranks in iterative low-rank solvers. While these approaches can be formulated on infinite-dimensional sequence spaces, in computationally realizable methods they need to be combined with suitable discretizations. A basic approach for guaranteeing prescribed discretization errors in high dimensions is given in Section 5.3, based on combined low-rank and sparse approximations of residuals. On this basis, iterative solvers that complement HSVD truncation with basis coarsening can be constructed, with error reduction realized by approximate Richardson iteration as in (5.26) or by solving successively refined Galerkin discretizations as in (5.35).

Remark 5.29. An important feature of these adaptive solvers is their universality, in the sense that they can be carried out without any knowledge of the approximability of the exact solution $\mathbf{u}$ : we only need to know suitable approximability of
the problem data $\mathbf{A}$ and $\mathbf{f}$. In particular, plain convergence of the methods to $\mathbf{u}$ is obtained without any assumptions on low-rank or sparse approximability of $\mathbf{u}$.

Further statements on the performance of the methods then require additional assumptions on the approximability of $\mathbf{u}$. As a consequence of Theorem 5.23 and Remark 5.28, if $\mathbf{u}$ has a certain low-rank and sparse approximability, this will be reflected in the ranks and support sizes of iterates. The two crucial approximability assumptions concern intrinsic properties of $\mathbf{u}$ : on the one hand, convergence of best approximations of prescribed maximum hierarchical rank, which is determined by the decay of singular value sequences $\sigma^{[e]}(\mathbf{u})$ for $e \in \mathbb{E}$, and on the other hand, near-sparsity of the contractions $\pi^{(i)}(\mathbf{u}), i=1, \ldots, d$, quantified by finiteness of $\mathcal{A}^{s}$-norms for some $s>0$.

Based on such rank and support size estimates for iterates, we arrive at estimates for the total computational costs of these adaptive solvers; in other words, we obtain estimates for the computational complexity of deterministically solving high-dimensional PDEs up to a guaranteed total approximation error.

Remark 5.30. The computational costs of the iteration (5.26) are generally dominated by those of the computation of HSVD forms of the iterates $\mathbf{v}^{n, j}$. As noted in (2.24), the number of arithmetic operations this requires for each $n$ and $j$ is bounded by

$$
\begin{equation*}
C\left(d\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{v}^{n, j}\right)\right\|_{\infty}^{4}+\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{v}^{n, j}\right)\right\|_{\infty}^{2} \sum_{i=1}^{d} \# \operatorname{supp}_{i} \mathbf{v}^{n, j}\right) \tag{5.36}
\end{equation*}
$$

with a fixed $d$-independent constant $C>0$. Assuming $\pi^{(i)}(\mathbf{u}) \in \mathcal{A}^{s}, i=1, \ldots, d$, for some $s>0$, we compare the implications of this bound for the two different types of low-rank approximability considered in Remark 5.6. For the iterates $\mathbf{u}^{n}$ with $\left\|\mathbf{u}-\mathbf{u}^{n}\right\|_{\ell_{2}} \leq \eta_{n}$ of the outer iteration (5.26b), as a consequence of Theorem 5.23, the complexity bound (5.36) takes the following forms.
(i) Assuming $\inf _{\left\|\operatorname{rank}_{\mathbb{E}}(\mathbf{w})\right\|_{\infty} \leq r}\|\mathbf{u}-\mathbf{w}\|_{\ell_{2}} \leq C r^{-p}$ with $C, p>0$, we obtain $\left\|\operatorname{rank}_{\mathbb{E}}\left(\mathbf{u}^{n}\right)\right\|_{\infty} \lesssim d^{1 /(2 p)} \eta_{n}^{-1 / p}$, and thus the HSVD of $\mathbf{u}^{n}$ costs at most

$$
C_{1}\left(\eta_{n}^{-4 / p}+\eta_{n}^{-2 / p-1 / s}\right)
$$

arithmetic operations, where $C_{1}$ depends on $\mathbf{u}$ and polynomially on $d$. Depending on $p$, the costs may thus scale much less favourably than the sizes $O\left(\eta_{n}^{-1 / s}\right)$ of the discretizations in each tensor mode.
(ii) Under the stronger assumption

$$
\begin{equation*}
\inf _{\left\|\operatorname{rank}_{\mathbb{E}}(\mathbf{w})\right\|_{\infty} \leq r}\|\mathbf{u}-\mathbf{w}\|_{\ell_{2}} \leq C e^{-c r^{\beta}} \tag{5.37}
\end{equation*}
$$

with $c, C, \beta>0$, the costs of the HSVD of $\mathbf{u}^{n}$ can be estimated by

$$
C_{2}\left(\left|\log \eta_{n}\right|^{4}+\left|\log \eta_{n}\right|^{2} \eta_{n}^{-1 / s}\right)
$$

where $C_{2}$ depends again on $\mathbf{u}$ and polynomially on $d$. In this more favourable case, the second summand in the above bound dominates and the costs differ from those of the lower-dimensional discretizations in each tensor mode only by a logarithmic factor.

The ranks of the iterates $\mathbf{v}^{n, j}$ of the inner iteration additionally depend on the representation ranks of residual approximations and are thus more problem-dependent.

Example 5.31. To consider the total computational costs, we return to the example of a second-order elliptic operator $A: H_{0}^{1}\left(\square_{d}\right) \rightarrow H^{-1}\left(\square_{d}\right)$ with wavelet Riesz basis representation $\tilde{\mathbf{A}}=\tilde{\mathbf{D}} \mathbf{T} \tilde{\mathbf{D}}$ on $\ell_{2}\left(\mathrm{~V}^{d}\right)$ and residual approximation as outlined in Remark 5.27. We assume $\mathbf{T}$ to be of fixed rank such that cond( $\tilde{\mathbf{A}})$ is bounded independently of $d$. Moreover, we assume $\mathbf{u}$, motivated by the results for the Poisson problem from Section 4.2, to have low-rank best approximations with exponential-type convergence (5.37). In addition, we assume $\mathbf{f}$ to have the same type of low-rank approximability, and that for some $s>0,\left\|\pi^{(i)}(\mathbf{u})\right\|_{\mathcal{A}^{s}},\left\|\pi^{(i)}(\mathbf{f})\right\|_{\mathcal{A}^{s}}<\infty$ for $i=1, \ldots, d$. With some further technical assumptions, a bound for the total computational costs in this setting is obtained in Bachmayr and Dahmen (2016b). Here, a major issue is taking into account the connection between discretization and preconditioning ranks in the residual approximations (5.34). Altogether, (5.26) under these conditions is guaranteed to produce for each given $\varepsilon>0$ an approximation $\mathbf{u}_{\varepsilon}$ such that $\left\|\mathbf{u}-\mathbf{u}_{\varepsilon}\right\|_{\ell_{2}} \leq \varepsilon$ using a total number of floating-point operations bounded by

$$
\begin{equation*}
C_{0} d^{c_{1} \ln d}(1+|\ln \varepsilon|)^{c_{2} \ln d+2 / \beta} \varepsilon^{-1 / s} \tag{5.38}
\end{equation*}
$$

where $C_{0}, c_{1}, c_{2}>0$ are independent of $d$ and $\varepsilon$. Recall that for the functions $u, u_{\varepsilon} \in H^{1}\left(\square_{d}\right)$ represented by $\mathbf{u}$ and $\mathbf{u}_{\varepsilon}$, respectively, we have $\left\|u-u_{\varepsilon}\right\|_{H^{1}} \bar{\sim}$ $\left\|\mathbf{u}-\mathbf{u}_{\varepsilon}\right\|_{\ell_{2}}$. Up to logarithmic factors, we thus obtain the costs $O\left(\varepsilon^{-1 / s}\right)$ that are optimal for the approximation of each one-dimensional component. Moreover, the curse of dimensionality is avoided: the leading $d$-dependence $d^{c_{1} \ln d}=e^{c_{1} \ln ^{2} d}$ grows faster than any polynomial, but substantially slower than exponentially in $d$.

This shows that for this typical class of problems, the curse of dimensionality can be avoided while achieving (deterministic) computable a posteriori error bounds for the approximate solutions. However, the above complexity bounds are likely not sharp: the numerical tests in Bachmayr and Dahmen (2016b) indicate that for the Poisson problem and certain variants, the total numerical costs exhibit low-order polynomial scaling with respect to $d$.

For the approach using Galerkin discretizations, analogous conclusions for Example 5.31 are obtained in Ali and Urban (2020). For the alternative rank control by soft thresholding discussed in Section 5.2.2, the analysis of a combination with adaptively refined discretizations is still open, but one may conjecture that this can
lead to an improved $d$-dependence in the total cost estimates compared to (5.38). For discussions of the total computational complexity of (5.26) under different assumptions, see also Bachmayr and Dahmen (2015, 2020).

### 5.5. Approximability and applications to high-dimensional elliptic problems

We now turn to applications to elliptic PDEs beyond the guiding model problems considered so far. As we have seen, the efficiency of low-rank solvers is determined in particular by the low-rank approximability of the corresponding exact solutions. On the one hand, the approximability of solutions can be studied experimentally by solvers with rank guarantees, as in Theorems 5.5 or 5.10 . On the other hand, it is also of interest to obtain analytical results that yield efficient low-rank approximations. Such investigations can also shed some light on the mechanisms and structural features of solutions that lead to rapidly convergent best low-rank approximations.

General low-rank approximability results for functions of classical or mixed Sobolev regularity are obtained in Schneider and Uschmajew (2014), Griebel and Harbrecht $(2014,2019,2023)$ and Bachmayr, Nouy and Schneider (2021b), where results for functions with compositional structure are also obtained, and in Griebel, Harbrecht and Schneider (2022). These results based on regularity properties only yield algebraic convergence of low-rank best approximations. In many applications, however, we observe exponential-type decay of best approximation errors, which in view of Remark 5.30 leads to substantially more favourable complexity of solvers. Such stronger approximability is often tied to particular structural features of solutions.

### 5.5.1. High-dimensional diffusion problems

The results discussed in Section 4.2, especially Theorem 4.6 from Dahmen et al. (2016), show low-rank approximability of solutions of Poisson problems as in Example 3.1 on $\square_{d}$ with potentially large $d$. In this case we obtain exponentialtype decay of best approximation errors with respect to ranks, or conversely, best approximation ranks that grow like a power of a logarithm of the achieved error.

More generally, one may be interested in problems of finding $u \in H_{0}^{1}\left(\square_{d}\right)$ such that

$$
\int_{\square_{d}} \hat{M} \nabla u \cdot \nabla v \mathrm{~d} x=\int_{\square_{d}} f v \mathrm{~d} x \quad \text { for all } v \in H_{0}^{1}\left(\square_{d}\right)
$$

with a diffusion coefficient $\hat{M} \in \mathbb{R}^{d \times d}$. The above-mentioned results apply only to diagonal $\hat{M}$. The results for discretized problems in Kressner and Uschmajew (2016) are applicable to more general $\hat{M}$, but yield only algebraic decay of best approximation errors with bounds that deteriorate under discretization refinement. Discretization-independent bounds for the approximability of solutions with general $\hat{M}$ are still an open problem. Experimentally, for problems with certain tridiagonal $\hat{M}$, Bachmayr and Dahmen (2016b) and Kressner et al. (2016) have
observed low-rank approximability similar to that for diagonal $\hat{M}$. Slightly more efficient approximations are obtained when errors are controlled only in $L_{2}\left(\square_{d}\right)$ rather than in $H_{0}^{1}\left(\square_{d}\right)$; see Bachmayr and Dahmen (2016a).

### 5.5.2. Parameter-dependent elliptic problems

In parameter-dependent PDEs as outlined in (1.5), the problems considered are typically posed on domains $\Omega \subset \mathbb{R}^{d}$ with $d \in\{1,2,3\}$, but with solutions that are still high-dimensional functions due to an additional dependence on a parameter from some set $Y$. A frequently treated model case is that of second-order elliptic problems with diffusion coefficients $a(\cdot, y) \in L_{\infty}(\Omega)$ with a parameter $y \in Y$, where for each $y$, the solution $u(\cdot, y) \in H_{0}^{1}(\Omega)$ for fixed right-hand side $f \in L_{2}(\Omega)$ solves

$$
\begin{equation*}
\int_{\Omega} a(x, y) \nabla u(x, y) \cdot \nabla v(x) \mathrm{d} x=\int_{\Omega} f(x) v(x) \mathrm{d} x \quad \text { for all } v \in H_{0}^{1}(\Omega) \tag{5.39}
\end{equation*}
$$

Different types of parameter dependences of $a$ are considered in the literature.
Example 5.32. A frequently used class of model problems is that of coefficients with affine parametrizations of the form

$$
\begin{equation*}
a(x, y)=\bar{a}(x)+\sum_{j \in \mathcal{N}} y_{j} \theta_{j}(x) \tag{5.40}
\end{equation*}
$$

with $y=\left(y_{1}, y_{2}, \ldots\right) \in Y=[-1,1]^{\mathcal{N}}$, where $\mathcal{N}$ may be finite or countably infinite and $\bar{a}$ as well as $\theta_{j}, j \in \mathcal{N}$, are fixed functions in $L_{\infty}(\Omega)$. Two examples in which ellipticity holds uniformly in $y \in Y$ are as follows:
(A) $\# \mathcal{N}<\infty, \bar{a}=1$ and $\theta_{j}=c_{j} \chi_{\Omega_{j}}$, where $\Omega_{j} \subset \Omega$ are disjoint subdomains and $c_{j}$ are constants with $\left|c_{j}\right|<1$;
(B) $\mathcal{N} \simeq \mathbb{N}$ and $\left\|\theta_{j}\right\|_{L_{\infty}} \lesssim j^{-r}$ with some $r>0$ such that $\bar{a}-\sum_{j \in \mathcal{N}}\left|\theta_{j}\right| \geq a_{0}>0$ in $\Omega$ for some $a_{0} \in \mathbb{R}$.

Particularly in a stochastic context, where the entries of $y$ correspond to scalar random variables, it is natural to consider weak formulations in the parametric domain as well. For simplicity, we consider the uniform distribution on $Y$ : let $\mu_{1}$ be the uniform measure on $[-1,1]$, so that $\mu=\bigotimes_{i \in \mathcal{N}} \mu_{1}$ is the uniform measure on $Y=[-1,1]^{\mathcal{N}}$. We then seek $u \in \mathcal{V}=H_{0}^{1}(\Omega) \otimes L_{2}(Y)$ such that for all $v \in \mathcal{V}$,

$$
\begin{equation*}
\int_{Y} \int_{\Omega} a(x, y) \nabla u(x, y) \cdot \nabla v(x, y) \mathrm{d} x \mathrm{~d} \mu(y)=\int_{Y} \int_{\Omega} f(x) v(x, y) \mathrm{d} x \mathrm{~d} \mu(y) . \tag{5.41}
\end{equation*}
$$

Here the space $\mathcal{V}$ is a tensor product; in the case $\# \mathcal{N}<\infty$, we have in particular

$$
\mathcal{V}=H_{0}^{1}(\Omega) \otimes L_{2}(Y)=H_{0}^{1}(\Omega) \otimes \bigotimes_{i \in \mathcal{N}} L_{2}(-1,1)
$$

A natural choice of basis for $L_{2}(Y)$ is that of tensor product orthonormal polynomials as in Example 3.3, which can be combined with any discretization or Riesz
basis for $H_{0}^{1}(\Omega)$. Since the spatial coordinate is usually treated as a single tensor mode here, in solving (5.41) one can work with tensor product Riesz bases of $\mathcal{V}$. This means that the issues of low-rank preconditioning discussed in Section 4.4 do not arise here.

The low-rank approximability of solutions depends strongly on the particular problem, especially on the structure of the parametrization of $a$. Bachmayr and Cohen (2017) have obtained bounds on best low-rank approximation errors for case (A) of Example 5.32. They show favourable exponential decay of errors and, for problems with certain structural features, a substantial improvement over direct product polynomial approximations in the parametric variable of the solution. This is confirmed by numerical tests in Bachmayr et al. (2018).

The situation is different, however, in problems with infinitely many parameters, as in case (B) of Example 5.32. The anisotropy in the parametric variables, due to the algebraic decay of $\left\|\theta_{j}\right\|_{L_{\infty}}$, can be exploited by sparse product polynomial approximations that converge independently of any parametric dimensionality parameter; see Cohen, DeVore and Schwab $(2010,2011)$ and Bachmayr, Cohen, Dũng and Schwab (2017). Low-rank approximability of such problems is considered in Bachmayr et al. (2018). As these results show, the singular values of the particular matricization with spatial degrees of freedom in its rows and parametric ones in its columns can generally be expected to decay only algebraically. In such cases, concerning the asymptotic total computational complexity of solvers, we can generally not expect low-rank approximations to offer an advantage over direct sparse approximations.

Low-rank approximation methods based on stochastic Galerkin formulations (5.41) with fixed discretizations are considered, for instance, in Khoromskij and Schwab (2011), Kressner and Tobler (2011a), Tobler (2012), Matthies and Zander (2012) and Lee and Elman (2017). Methods with adaptive refinement of spatial and parametric discretizations are considered in Eigel et al. (2017) based on adaptive finite elements for the spatial degrees of freedom, and in Bachmayr et al. (2018) using wavelets in space. The more involved case of lognormal coefficient parametrizations is considered in Eigel et al. (2020).

Remark 5.33. As an alternative to variational formulations (5.41), approximations of parameter-dependent solutions can also be constructed from point evaluations with respect to $y$, that is, from solutions of (5.39) for certain well-chosen parameter values. In the low-rank context, this can be done by the TT-cross method for tensor trains (Oseledets and Tyrtyshnikov 2010) or its counterpart for hierarchical tensors (Ballani, Grasedyck and Kluge 2013). Such approaches are used for parametric problems, for instance in Ballani and Grasedyck (2015) and Dolgov and Scheichl (2019). A different approach based on interpolation of parametric point evaluations is taken in Khoromskij and Oseledets (2010). Compared to (5.41), such methods offer additional flexibility, especially in problems with $a$ depending
nonlinearly on $y$, but we do not have the mechanisms for error control as offered by (5.41).

Remark 5.34. The application of low-rank methods to parameter-dependent problems is strongly connected to certain methods for model order reduction, in particular to proper orthogonal decomposition (Kahlbacher and Volkwein 2007) and to reduced basis methods (Prud'homme and Patera 2004, Rozza, Huynh and Patera 2008, Hesthaven, Rozza and Stamm 2016); for a recent overview, see Benner, Cohen, Ohlberger and Willcox (2017). In reduced basis methods, case (A) in Example 5.32 is a typical model problem. With $N$ well-chosen parameter values $y^{1}, \ldots, y^{N} \in Y$, we compute the so-called solution snapshots $v_{k}=u\left(\cdot, y^{k}\right)$, $k=1, \ldots, N$, and obtain efficient approximate solutions for arbitrary parameters $y \in Y$ by solving the Galerkin projections of (5.39) onto $\operatorname{span}\left\{v_{k}\right\}_{k=1, \ldots, N}$. This leads to approximations of the form

$$
u(x, y) \approx \sum_{k=1}^{N} v_{k}(x) \phi_{k}(y)
$$

where the functions $\phi_{k}, k=1, \ldots, N$, are defined implicitly by Galerkin projection.

### 5.5.3. Tensorized approximations for boundary value problems

We briefly comment on second-order elliptic boundary value problems on domains $\Omega \subset \mathbb{R}^{d}$ with $d \in\{1,2,3\}$. The efficient approximation of solutions of such problems can require specialized methods, such as $h p$-finite elements for problems with gradient singularities caused by non-smooth problem data or domains, or homogenization approaches for problems with high-frequency oscillations. The technique of multilevel tensorized approximations, as discussed in Sections 3.4 and 4.5, can be used to obtain generic approximations for such problems by providing efficient parametrizations of simple low-order methods (such as piecewise multilinear finite elements) on extremely fine grids. For diffusion problems parametrized in tensorized low-rank form as outlined in Section 4.5, this is considered by Kazeev (2015), Kazeev and Schwab (2018) and Marcati, Rakhuba and Schwab (2022a). Similarly to convergence results for $h p$-FEM, they establish exponential-type convergence of approximate solutions with respect to the total number of representation parameters. Comparably efficient approximations are obtained for multiscale problems in Kazeev, Oseledets, Rakhuba and Schwab $(2017,2022)$ and for convection-diffusion problems on intervals in Marcati, Rakhuba and Ulander (2022b). The technique is extended to isogeometric analysis in Markeeva, Tsybulin and Oseledets (2021).

Combined with preconditioning as described in Section 4.5, low-rank parametrizations of approximations on very fine grids (corresponding to grid spacings near machine precision) can be used, so that adaptive refinement of discretizations typically becomes unnecessary in this context. Approximate solutions can be obtained by the methods discussed in Section 5.1 with preconditioning or by the more problem-specific alternating direction implicit method of Rakhuba (2021).

## 6. Eigenvalue problems

Solvers based on low-rank approximations for eigenvalue problems are especially important in quantum physics applications, where they are applied both in the classical formulation in terms of particle coordinates, as in (3.25), and in the occupation number formulation discussed in Section 3.3. In the latter case, lowrank structures in wavefunction approximations are closely connected to measures of entanglement in quantum systems (Orús 2014). A physical motivation for the application of low-rank methods in this context is that in view of so-called area laws (Eisert, Cramer and Plenio 2010), the most relevant states can be expected to be those of relatively low ranks in tensor representations.

Many of the considerations of the preceding sections for elliptic partial differential equations apply equally to associated eigenvalue problems. The corresponding convergence theory, however, is not as well developed in the latter case, especially concerning error estimation and adaptive refinement of discretizations.

### 6.1. Applications in quantum physics

A high-dimensional eigenvalue problem that is of central importance in quantum chemistry is the electronic Schrödinger equation. The corresponding Hamiltonian operator is of the general form (3.25). Writing $x_{1}, \ldots, x_{N} \in \mathbb{R}^{3}$ for the spatial coordinates of the $N$ electrons in the system under consideration, for a molecule composed of $K$ atomic nuclei with positions $R_{k} \in \mathbb{R}^{3}$ and charges $Z_{k}>0$ for $k=1, \ldots, K$, this Hamiltonian reads

$$
H=-\frac{1}{2} \sum_{i=1}^{N} \Delta_{x_{i}}-\sum_{i=1}^{N} \sum_{k=1}^{K} \frac{Z_{k}}{\left\|x_{i}-R_{k}\right\|_{2}}+\frac{1}{2} \sum_{\substack{i, j=1 \\ i \neq j}}^{N} \frac{1}{\left\|x_{i}-x_{j}\right\|_{2}}
$$

For an in-depth treatment of the analysis of electronic Schrödinger eigenvalue problems, we refer to Yserentant (2010). Here, the eigenfunctions (called wavefunctions) of interest are typically those corresponding to the lowest eigenvalues, that is, to the lowest energies. By standard techniques, approximation errors in eigenvalues can be estimated in terms of $H^{1}$-errors in the corresponding eigenfunction approximations.

As noted in Section 3.3, fermionic wavefunctions need to be antisymmetric under exchange of electron coordinates. Although the direct application of the hierarchical tensor format is problematic due to this requirement, direct approximations of wavefunctions by antisymmetrized canonical tensor representations are considered in Beylkin, Mohlenkamp and Pérez (2008) and Mohlenkamp (2010). However, with this approach we face the difficulties outlined in Section 1.4.

There exist a variety of approximations of the electronic Schrödinger equation by simplified models. The Hartree-Fock (HF) equations are derived by forming the Rayleigh quotient of $H$ over a single $N$-electron Slater determinant and minimizing over the corresponding $N$ orbitals $\phi_{1}, \ldots, \phi_{N}$, which are functions of a single
electron coordinate. Constraining the orbitals $\phi_{1}, \ldots, \phi_{N}$ to be orthonormal, the necessary optimality conditions lead to the HF equations

$$
\begin{equation*}
F\left(\phi_{1}, \ldots, \phi_{N}\right) \phi_{i}=\mu_{i} \phi_{i}, \quad i=1, \ldots, N \tag{6.1}
\end{equation*}
$$

with the Fock operator $F\left(\phi_{1}, \ldots, \phi_{N}\right)$, which is an integrodifferential operator acting on functions on $\mathbb{R}^{3}$.

While standard methods for (6.1) are based on approximations by a small number of atom-centred Gaussian functions (Helgaker et al. 2000, Chap. 8), numerical methods that can reach higher accuracies are frequently of interest. In this case grid-based methods for the HF equations (and related problems in density functional theory) that approximate the orbitals as low-rank third-order tensors are proposed in Khoromskij, Khoromskaia, Chinnamsetty and Flad (2009), Khoromskij, Khoromskaia and Flad (2011) and Rakhuba and Oseledets (2016). In Bischoff and Valeev (2011), approximations in low-rank tensor format of an integral equation formulation of (6.1) are combined with locally refined discontinuous wavelets expansions for one-dimensional components; these wavelets, however, are not sufficiently regular for a rigorous error control of computed energies. In Bachmayr (2012a), adaptively refined higher-order Daubechies wavelet expansions of oneand two-electron wavefunctions are approximated in Tucker format.

These approaches rely on explicit low-rank approximations of potential terms, which can be obtained from exponential sum approximations as considered in Section 4.1. In particular, applying approximations of $t \mapsto t^{-1 / 2}$ for $t>0$ provided by Theorems 4.2 or Corollary 4.5 leads to approximations of Coulomb potentials by Gaussians,

$$
\frac{1}{\sqrt{x_{1}^{2}+x_{2}^{2}+x_{3}^{2}}} \approx \sum_{k=1}^{r} \omega_{k} \prod_{i=1}^{3} e^{-\alpha_{k} x_{i}^{2}} .
$$

Such approximations are considered, for instance, in Chinnamsetty et al. (2007), Zeiser (2010), Bachmayr (2012a,b) and Scholz and Yserentant (2017). Related issues include the efficient evaluation of two-electron potentials (Khoromskaia, Khoromskij and Schneider 2013) and tensor-based summation of potentials on lattices (Khoromskaia and Khoromskij 2016). For an overview, see also Khoromskaia and Khoromskij $(2015,2018)$.

A variety of refined models for the electronic Schrödinger equations take HF approximations as their starting points. Instances of such so-called post-HartreeFock methods are the Møller-Plesset perturbation theory and the coupled cluster method. The use of low-rank tensor approximations in the context of such methods is investigated in Benedikt, Auer, Espig and Hackbusch (2011), Benedikt, Böhm and Auer (2013b) and Benedikt et al. (2013a).

The concepts of matrix product states (tensor trains) and tree tensor networks (hierarchical tensors), as well as the DMRG algorithm outlined in Section 5.1, have their origins in the study of spin systems. These are quantum systems of particles
on lattices with locally interacting spins, where $N$-particle systems typically lead to wavefunctions on $\{0,1\}^{N}$; concerning such applications, see Schollwöck (2011) and Orús (2014).

Starting with White and Martin (1999), such methods have also been successfully applied to the occupation number representation of the electronic Schrödinger eigenvalue problem as discussed in Section 3.3. For an overview, we refer to Wouters and Van Neck (2014). Such approaches based on tree tensor networks, which can be regarded as multi-configurational self-consistent field methods, are of particular interest in cases where the Hartree-Fock equations yield poor approximations. They can also serve as the basis of modified post-HF methods such as the tailored coupled cluster method (Faulstich et al. 2019).

Concerning the low-rank approximability of operators in this context, see Remark 3.13. The low-rank approximability of electronic wavefunctions in occupation number representation is investigated in Graswald and Friesecke (2021), Friesecke and Graswald (2022) and Friesecke, Graswald and Legeza (2022). Here, the optimization of the underlying orbitals plays an important role in improving the efficiency of low-rank approximations; see for instance Krumnow et al. (2016).

### 6.2. Solving eigenvalue problems in low-rank formats

One strategy for constructing eigensolvers operating on low-rank representations is the adaptation of standard methods by additional rank truncations. Power methods and preconditioned inverse iteration (PINVIT) for discretizations of differential operators are considered by Hackbusch, Khoromskij, Sauter and Tyrtyshnikov (2012). Kressner and Tobler (2011b) proposed an LOBPCG eigensolver modified by additional rank truncations and tested it numerically. Cancès et al. (2014) considered the application of the greedy approach outlined in Section 5.1.4. Rakhuba and Oseledets (2018) proposed a Jacobi-Davidson method adapted to low-rank matrix manifolds. Rakhuba, Novikov and Oseledets (2019) developed a Riemannian version of LOBPCG for higher-order tensors. Concerning the combination of low-rank approximations with adaptivity, Bachmayr (2012a) have shown that a low-rank eigensolver based on PINVIT with adaptive wavelet discretization has guaranteed local convergence in $H^{1}$-norm.

A second construction principle for eigensolvers is given by methods based on alternating optimization. A prominent example is the DMRG algorithm introduced in White $(1992,2005)$ and formulated in terms of matrix product states in Vidal (2003). These methods are investigated from a numerical perspective in Holtz et al. (2012a). The adaptation of such methods to the simultaneous approximation of several eigenspaces by attaching eigenfunction indices as additional tensor modes is considered in Kressner, Steinlechner and Uschmajew (2014a) and Dolgov, Khoromskij, Oseledets and Savostyanov (2014). Krumnow, Pfeffer and Uschmajew (2021) propose direct energy minimization instead of SVD for the rank truncation in DMRG substeps.

## 7. Time-dependent problems

In Section 1.2 we discussed two examples of time-dependent problems of potentially high spatial dimensionality: the parabolic problem (1.1) and the instationary Schrödinger equation (1.2). To obtain approximate solutions in low-rank tensor formats, one strategy is to apply standard time-stepping procedures. For implicit methods, as required for parabolic problems, the problems on the spatial domain that need to be solved in each time step can be treated by low-rank solvers for elliptic problems, as in Section 5. Such approaches are pursued, for instance, in Dolgov, Khoromskij and Oseledets (2012) and Cho, Venturi and Karniadakis (2016).

However, it is a delicate matter to adjust low-rank approximation tolerances in each time step to achieve a desired total error over the given time interval, and no informative bounds on the arising ranks are available. Moreover, in certain problems we are also interested in preserving conservation properties of the continuous problem, which is generally difficult to reconcile with rank truncations in every time step. These limitations are addressed by alternative methods for low-rank approximation of time-dependent problems, on which we focus in this section.

### 7.1. Dynamical low-rank approximation

In dynamical low-rank approximations, evolution problems are approximated on manifolds of fixed-rank matrices or tensors, which leads to a reduction to lowerdimensional evolution problems for the individual components of low-rank representations.

For matrices of fixed rank, this concept is introduced in Koch and Lubich (2007a). Assume that we are given a differential equation for a time-dependent matrix $\mathbf{Y}(t) \in \mathbb{R}^{m \times n}$,

$$
\dot{\mathbf{Y}}(t)=F(\mathbf{Y}(t)), \quad \mathbf{Y}(0)=\mathbf{Y}_{0}
$$

where $F: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ is a Lipschitz continuous function. Discretizations of instationary PDEs on two-dimensional spatial domains lead naturally to matrix differential equations of this structure.

Let $\mathcal{M}$ be the submanifold of matrices of fixed rank $r \in \mathbb{N}$ (see also Sections 2.10 and 5.1.3). We now approximate $\mathbf{Y}(t)$ by $\mathbf{X}(t) \in \mathcal{M}$ defined by the Dirac-Frenkel variational principle (Dirac 1930, Frenkel 1934):

$$
\langle\dot{\mathbf{X}}(t)-F(\mathbf{X}(t)), \boldsymbol{\Xi}\rangle=0 \quad \text { for all } \boldsymbol{\Xi} \in \mathcal{T}_{\mathbf{X}(t)} \mathcal{M}
$$

Since $\dot{\mathbf{X}}(t) \in \mathcal{T}_{\mathbf{X}(t)} \mathcal{M}$, we thus solve the modified problem

$$
\begin{equation*}
\dot{\mathbf{X}}(t)=\Pi_{\mathbf{X}(t)} F(\mathbf{X}(t)), \quad \mathbf{X}(0)=\mathbf{X}_{0}, \tag{7.1}
\end{equation*}
$$

where $\mathbf{X}_{0}$ is the best rank- $r$ approximation of $\mathbf{Y}_{0}$ and where again $\Pi_{\mathbf{X}(t)}$ denotes the orthogonal projection onto $\mathcal{T}_{\mathbf{X}(t)} \mathcal{M}$.

With the representation (2.9) in the form $\mathbf{X}(t)=\mathbf{U}(t) \mathbf{S}(t) \mathbf{V}(t)^{\top}$, where $\mathbf{U}, \mathbf{V}$ are required to have orthonormal columns for all times and $\mathbf{S} \in \mathbb{R}^{r \times r}$ is invertible, as
shown in Koch and Lubich (2007a), (7.1) can be rewritten componentwise as

$$
\begin{align*}
\dot{\mathbf{S}} & =\mathbf{U}^{\top} F(\mathbf{X}) \mathbf{V},  \tag{7.2a}\\
\dot{\mathbf{U}} & =\left(\mathbf{I}-\mathbf{U} \mathbf{U}^{\top}\right) F(\mathbf{X}) \mathbf{V} \mathbf{S}^{-1},  \tag{7.2b}\\
\dot{\mathbf{V}} & =\left(\mathbf{I}-\mathbf{V} \mathbf{V}^{\top}\right) F(\mathbf{X}) \mathbf{U} \mathbf{S}^{-\top} . \tag{7.2c}
\end{align*}
$$

Provided that $F$ has suitable structure, the original problem is thus reduced to differential equations for lower-dimensional components. An analogous scheme for Tucker tensors is considered in Koch and Lubich (2010). Similarly to (7.2), we obtain separate but coupled evolution equations for the core tensor and for the orthonormal mode frames. The method can also be formulated on function spaces, with a natural generalization to Hilbert spaces. A formulation on Banach spaces is given in Falcó et al. (2019).

In the context of quantum dynamics described by instationary Schrödinger equations (1.2), the method for Tucker tensors was introduced as the multi-configuration time-dependent Hartree (MCTDH) method by Meyer, Manthe and Cederbaum (1990); for an overview of such models and related methods, we refer to Lubich (2008) and Meyer, Gatti and Worth (2009). Applied to (1.2), approximations are obtained by solving

$$
\mathrm{i} \partial_{t} u=-\Delta u+\Pi_{u} V u
$$

with appropriate low-rank approximations of initial data, where $\Pi_{u}$ denotes the projection onto the tangent space at $u$ of the manifold of fixed-rank tensors. Note that the Laplacian $\Delta$ maps elements of the manifold to the tangent space, and thus no projection is required on this term. ${ }^{1}$

MCTDH is applied to models for the evolution for distinguishable particles such as atomic nuclei, where we do not have the symmetry or antisymmetry requirements mentioned in Section 3.3. However, similar approaches can be considered for electron dynamics, where the antisymmetry requirement applies, which leads to the multi-configuration time-dependent Hartree-Fock method; see Bardos, Catto, Mauser and Trabelsi $(2009,2010)$ and Koch and Lubich (2011).

Applying the above approach to hierarchical tensor approximations leads to the multilayer MCTDH method (Wang and Thoss 2003). In a general context, dynamical low-rank approximations for tensor trains and hierarchical tensors are analysed in Lubich, Rohwedder, Schneider and Vandereycken (2013), Uschmajew and Vandereycken (2013) and Arnold and Jahnke (2014). The application to the time evolution of matrix product states in spin systems is investigated in Haegeman, Osborne and Verstraete (2013). Projection methods for time integration of such formulations for matrices and tensor trains are analysed in Kieri and Vandereycken (2019).

[^0]For solutions $u$ of time-dependent PDEs with an additional dependence on a parameter $y$ (which may be a random variable), approximations of the low-rank form

$$
\begin{equation*}
u(t, x, y) \approx \bar{u}(t, x)+\sum_{k=1}^{r} v_{k}(t, x) w_{k}(t, y) \tag{7.3}
\end{equation*}
$$

where $\left\{v_{k}(t, \cdot)\right\}_{k=1, \ldots, r}$ are $L_{2}$-orthonormal on the spatial domain for each $t$, can be obtained by a variation of (7.2) termed dynamically orthogonal approximation that is introduced in Sapsis and Lermusiaux (2009). This approach is analysed in Musharbash, Nobile and Zhou (2015) with a particular focus on the treatment of parabolic PDEs with random parameters, and also concerning its close relation to dynamical low-rank approximation (7.2). A general analysis is given in Feppon and Lermusiaux (2018, 2019), including a characterization of dynamical lowrank approximation as a limit of alternating time steps and low-rank truncations. Further applications to Navier-Stokes problems and wave equations are considered in Musharbash and Nobile (2018) and Musharbash, Nobile and Vidličková (2020).

Remark 7.1. In the case of finite-dimensional matrices as in (7.2), local-in-time existence of dynamical low-rank approximations follows by the Picard-Lindelöf theorem. In the case of infinite-dimensional function spaces, the situation is more delicate. Existence, uniqueness and regularity of the MCTDH equations for sufficiently regular initial data are shown in Koch and Lubich (2007b). For approximations of parabolic problems depending on a random parameter $y$ as in (7.3), well-posedness in the mild and strong sense is shown in Kazashi and Nobile (2021). In the case of parabolic problems in two dimensions, existence and uniqueness for a weak formulation of dynamical low-rank approximation are demonstrated in Bachmayr, Eisenmann, Kieri and Uschmajew (2021a), and robustness results for corresponding spatially discretized problems are obtained in Conte (2020).

Remark 7.2. Let $\mathcal{M}$ be a manifold of rank-one tensors in a Hilbert space $H$, and consider

$$
\dot{u}(t)=f, \quad u(0)=u_{0} \in \mathcal{M}
$$

with fixed $f \in \mathcal{M}$. Then, whenever $f$ is orthogonal to $\mathcal{T}_{u_{0}} \mathcal{M}$, the dynamical low-rank approximation equals $u_{0}$ for all $t$. The best rank-one approximation of the exact solution $u(t)=u_{0}+t f$, however, will be $t f$ when $t$ is sufficiently large.

Local quasi-optimality properties of MCTDH approximations, compared to best approximations of the same rank, are established in Lubich (2005), and an error analysis of MCTDH is conducted in Conte and Lubich (2010). The assumptions required in the analysis in particular ensure that the above issue is avoided, but these conditions are typically not verifiable in practice.

### 7.2. Splitting and Basis Update \& Galerkin integrators

The system of differential equations (7.2) can be difficult to solve numerically, especially in the presence of small singular values of $\mathbf{S}$. An integrator for dynamical low-rank approximation with particularly favourable properties can be obtained by a splitting approach. Writing the projector onto the tangent space at $\mathbf{X}=\mathbf{U S V}^{\top}$ as

$$
\Pi_{\mathrm{X}} \mathbf{Z}=\mathbf{Z} \mathbf{V} \mathbf{V}^{\top}-\mathbf{U} \mathbf{U}^{\top} \mathbf{Z} \mathbf{V} \mathbf{V}^{\top}+\mathbf{U} \mathbf{U}^{\top} \mathbf{Z}
$$

and decomposing the right-hand side in (7.1) accordingly, one can apply standard Lie-Trotter or Strang splitting to arrive at the projector-splitting integrator that is proposed and analysed in Lubich and Oseledets (2014). This integrator is shown to enjoy particularly favourable exactness properties with improved robustness concerning small singular values of $\mathbf{S}$.

Integrators based on the same principles are introduced for MCTDH in Lubich (2015) and for general dynamical low-rank approximations of tensor trains in Lubich, Oseledets and Vandereycken (2015). Haegeman et al. (2016) established a connection to the DMRG scheme for matrix product states. The robustness properties of the splitting integrators with respect to small singular values are investigated further in Kieri, Lubich and Walach (2016) and Ostermann, Piazzola and Walach $(2019)$. Einkemmer and Lubich $(2018,2019)$ constructed projectorsplitting integrators for Vlasov-Poisson equations. For random parabolic equations, a scheme of this type is analysed in Kazashi, Nobile and Vidličková (2021). Variants including strategies for the adaptation of approximation ranks are proposed in Yang and White (2020), Dektor, Rodgers and Venturi (2021) and Dunnett and Chin (2021).

Integrators for fixed-rank matrices and Tucker tensors with similar structure and robustness properties, but not derived from a splitting of the projector, are introduced and analysed in Ceruti and Lubich (2020, 2022). In this approach, termed Basis Update \& Galerkin, in the matrix case considered above a time step is applied to the bases given by $\mathbf{U}$ and $\mathbf{V}$, and $\mathbf{S}$ is evolved by a Galerkin method using these bases. An integrator that includes rank adaptivity is developed in Ceruti, Kusch and Lubich (2022) and extended to tree tensor networks in Ceruti et al. (2023); it joins the basis vectors from current and previous time steps in each step for $\mathbf{S}$ followed by SVD-based rank truncation.

### 7.3. Space-time methods

Solving time-dependent problems in formulations on the full space-time domain leads to approximation problems with an additional dimension, compared to the spatial approximations required in time-stepping schemes. Such methods are considered in the literature, in particular for parabolic problems. A classical model problem is the heat equation on the $d$-dimensional unit cube with homogeneous Dirichlet boundary conditions,

$$
\begin{equation*}
\partial_{t} u-\Delta u=f \quad \text { in }(0, T) \times \square_{d},\left.\quad u\right|_{t=0}=u_{0},\left.u\right|_{(0, T) \times \partial \square_{d}}=0, \tag{7.4}
\end{equation*}
$$

as a special case of (1.1).
An algebraic construction of space-time solvers is obtained in Dolgov et al. (2012), where the entire coupled system resulting from a standard time-stepping discretization is solved jointly for all time steps in a multilevel tensorized representation. Although approximability results of solutions are obtained for certain cases, robustness with respect to discretization refinement and spatial dimensionality remain unclear in this approach. A related method based on certain series expansions is considered in Gavrilyuk and Khoromskij (2019). Further heuristic space-time schemes based on Chebyshev interpolation are considered in Sun and Kumar (2014) and Chertkov and Oseledets (2021).

To control discretization errors, methods derived from well-posed space-time weak formulations are of particular interest. Based on corresponding residuals measured in appropriate norms, discretization errors can be estimated jointly in space and time for such formulations. For second-order parabolic problems such as (7.4), a classical weak formulation using the space $\mathcal{W}=L_{2}(0, T ; V) \cap H^{1}\left(0, T ; V^{\prime}\right)$ with $V=H_{0}^{1}\left(\square_{d}\right)$ reads as follows: find $u \in \mathcal{W}$ such that

$$
\begin{align*}
& \int_{0}^{\top}\left\langle\partial_{t} u, v\right\rangle_{V^{\prime} \times V}+\int_{\square_{d}} \nabla u \cdot \nabla v \mathrm{~d} x \mathrm{~d} t+\left.\int_{\square_{d}} u\right|_{t=0} w \mathrm{~d} x \\
& \quad=\int_{0}^{\top}\langle f, v\rangle_{V^{\prime} \times V} \mathrm{~d} t+\int_{\square_{d}} u_{0} w \mathrm{~d} x \tag{7.5}
\end{align*}
$$

for all $v \in L_{2}(0, T ; V)$ and $w \in L_{2}\left(\square_{d}\right)$. The left-hand side of (7.5) can be shown to define an isomorphism from $\mathcal{W}$ to $L_{2}\left(0, T ; V^{\prime}\right) \times L_{2}\left(\square_{d}\right)$; see Schwab and Stevenson (2009).

Remark 7.3. Schwab and Stevenson (2009) have constructed an adaptive wavelet method based on a Riesz basis representation of (7.5). Applying such a Riesz basis representation in the context of low-rank approximations leads to additional difficulties. To see these, assume that $\psi_{\nu}, v \in \vee$, are as in Remark 3.10, and that in addition, $\omega_{\mu} \in H^{1}(0, T)$ are such that $\left\{\omega_{\mu}\right\}_{\mu \in \hat{v}}$ and $\left\{\left\|\omega_{\mu}\right\|_{H^{1}}^{-1} \omega_{\mu}\right\}_{\mu \in \hat{v}}$ are Riesz bases of $L_{2}(0, T)$ and $H^{1}(0,1)$, respectively. Then a Riesz basis for the test space $L_{2}(0, T ; V) \simeq L_{2}(0, T) \otimes V$ is given by $\left\{\omega_{\mu} \otimes \Psi_{\nu}\right\}_{(\mu, v) \in \hat{\mathrm{V}} \times \mathrm{V}^{d}}$ with $\Psi_{\nu}, v \in \mathrm{~V}^{d}$, as in Remark 3.10. This can be handled as described in Section 4.4. For the space $\mathcal{W}$, however, Riesz bases are given by $\left\{d_{\mu, v} \omega_{\mu} \otimes \psi_{v_{1}} \otimes \cdots \otimes \psi_{v_{d}}\right\}_{(\mu, v) \in \hat{\mathrm{V}} \times \mathrm{V}^{d}}$, with the requirement

$$
\begin{equation*}
d_{\mu, v_{1}, \ldots, v_{d}}=\left(\sum_{i=1}^{d} 2^{2\left|v_{i}\right|}+2^{2|\mu|}\left(\sum_{i=1}^{d} 2^{2\left|v_{i}\right|}\right)^{-1}\right)^{-1 / 2}, \quad \mu \in \hat{\mathrm{~V}}, v \in \mathrm{~V}^{d} \tag{7.6}
\end{equation*}
$$

(compare (3.24)). Low-rank approximations of the expressions on the right in (7.6) that are suitable for diagonal preconditioning in low-rank formats are substantially more difficult to construct than in the case of $H^{1}$-norms considered in Section 4.4.

Andreev and Tobler (2015) have solved a fixed discretization of (7.5) by multilinear finite elements in space and time in least-squares form with low-rank approximations in hierarchical format, where time and each spatial variable are treated as separate tensor modes. The formulation is based on the stable Petrov-Galerkin discretization by Andreev (2013), which relies on a careful adjustment of trial and test spaces. The appropriate norms are realized by a BPX preconditioner, which as in the case of Riesz bases - reduces preconditioning to diagonal scaling. However, this also requires scaling by the values in (7.6), which is realized heuristically in Andreev and Tobler (2015).

A space-time method based on a slightly different least-squares formulation is considered in Boiveau, Ehrlacher, Ern and Nouy (2019). In this case the analysis is based on a semidiscretization in the spatial variables. The low-rank approximation uses tensors of order two, separating time from the spatial variables; in this form, this method is thus not aimed at high-dimensional problems. The approximations are constructed by a greedy method as outlined in Section 5.1.4.

A different approach combining sparse wavelet approximations in time with adaptive low-rank approximations in space is used in Bachmayr and Faldum (2023). To this end, (7.5) is represented in a Riesz basis representation as in Remark 7.6 in terms of spline wavelets, using $\left\{\tilde{\Psi}_{\mu, \nu}\right\}_{(\mu, \nu) \in \hat{\mathrm{V}} \times \mathrm{V}^{d}}$ with $\tilde{\Psi}_{\mu, \nu}=d_{\mu, \nu} \omega_{\mu} \otimes \psi_{\nu_{1}} \otimes$ $\cdots \otimes \psi_{v_{d}}$, and $d_{\mu, v}$ satisfying (7.6), as a Riesz basis of $\mathcal{W}$. The solution $u$ of (7.5) then has the expansion

$$
u=\sum_{\mu \in \hat{V}} \sum_{\nu \in \mathrm{V}^{d}} \mathbf{u}\left[\mu, v_{1}, \ldots, v_{d}\right] \tilde{\Psi}_{\mu, v}
$$

With a fixed dimension tree $\mathbb{T}$ and a suitable index set $\hat{\Lambda} \subset \hat{V}$ defining a temporal discretization, we now aim to find separately for each $\mu \in \hat{\Lambda}$ hierarchical tensor approximations $\tilde{\mathbf{u}}_{\mu}$ of the $d$ th-order tensors

$$
\mathbf{u}_{\mu}=\left(\mathbf{u}\left[\mu, v_{1}, \ldots, v_{d}\right]\right)_{v_{1}, \ldots, v_{d} \in \mathrm{~V}}
$$

These approximations are chosen with supp $\tilde{\mathbf{u}}_{\mu} \subseteq \Lambda_{\mu}=X_{i=1}^{d} \Lambda_{\mu}^{(i)} \subset \vee^{d}$, where the finite index sets $\Lambda_{\mu}$ also need to be determined.

As shown in Bachmayr and Faldum (2023), the basic techniques described in Section 5.3 can be modified for computing such families $\left(\tilde{\mathbf{u}}_{\mu}\right)_{\mu \in \hat{\Lambda}}$ of hierarchical tensor approximations with adaptively selected discretizations. This approach requires efficient low-rank approximations of (7.6) only separately for each $\mu \in$ $\hat{v}$, which can be achieved by exponential sum approximations similar to those considered in Section 4.1.1. The resulting method yields guaranteed convergence to $u$ with computable error bounds in the $\mathcal{W}$-norm, as well as bounds on the arising ranks, discretization sizes and total computational costs analogous to those for the elliptic case discussed in Section 5.4. The method is demonstrated to be applicable to large spatial dimensions $d$ in practice.

## Acknowledgements

The author thanks Wolfgang Dahmen, Henrik Eisenmann, Manfred Faldum, Vladimir Kazeev, Sebastian Krämer, Christian Lubich, Reinhold Schneider and André Uschmajew for reading earlier versions of this article and for important hints and suggestions for improvements.

The author acknowledges funding by Deutsche Forschungsgemeinschaft (German Research Foundation), project number 442047500, through the Collaborative Research Center ‘Sparsity and Singular Structures’ (SFB 1481).

## References

P.-A. Absil and I. V. Oseledets (2015), Low-rank retractions: A survey and new results, Comput. Optim. Appl. 62, 5-29.
P.-A. Absil, R. Mahony and R. Sepulchre (2008), Optimization Algorithms on Matrix Manifolds, Princeton University Press.
M. Ali and A. Nouy (2020a), Singular value decomposition in Sobolev spaces, Part I, Z. Anal. Anwend. 39, 349-369.
M. Ali and A. Nouy (2020b), Singular value decomposition in Sobolev spaces, Part II, Z. Anal. Anwend. 39, 371-394.
M. Ali and A. Nouy (2021), Approximation theory of tree tensor networks: Tensorized multivariate functions. Available at arXiv:2101.11932.
M. Ali and A. Nouy (2023), Approximation theory of tree tensor networks: Tensorized univariate functions, Constr. Approx. Available at doi:10.1007/s00365-023-09620-w.
M. Ali and K. Urban (2020), HT-AWGM: A hierarchical Tucker-adaptive wavelet Galerkin method for high-dimensional elliptic problems, Adv. Comput. Math. 46, 59.
A. Ammar, F. Chinesta and A. Falcó (2010), On the convergence of a greedy rank-one update algorithm for a class of linear systems, Arch. Comput. Methods Engrg 17, 473486.
R. Andreev (2013), Stability of sparse space-time finite element discretizations of linear parabolic evolution equations, IMA J. Numer. Anal. 33, 242-260.
R. Andreev and C. Tobler (2015), Multilevel preconditioning and low-rank tensor iteration for space-time simultaneous discretizations of parabolic PDEs, Numer. Linear Algebra Appl. 22, 317-337.
A. Arnold and T. Jahnke (2014), On the approximation of high-dimensional differential equations in the hierarchical Tucker format, BIT Numer. Math. 54, 305-341.
M. Bachmayr (2012a), Adaptive low-rank wavelet methods and applications to two-electron Schrödinger equations. PhD thesis, RWTH Aachen.
M. Bachmayr (2012b), Hyperbolic wavelet discretization of the two-electron Schrödinger equation in an explicitly correlated formulation, ESAIM Math. Model. Numer. Anal. 46, 1337-1362.
M. Bachmayr and A. Cohen (2017), Kolmogorov widths and low-rank approximations of parametric elliptic PDEs, Math. Comp. 86, 701-724.
M. Bachmayr and W. Dahmen (2015), Adaptive near-optimal rank tensor approximation for high-dimensional operator equations, Found. Comput. Math. 15, 839-898.
M. Bachmayr and W. Dahmen (2016a), Adaptive low-rank methods for problems on Sobolev spaces with error control in L2, ESAIM Math. Model. Numer. Anal. 50, 11071136.
M. Bachmayr and W. Dahmen (2016b), Adaptive low-rank methods: Problems on Sobolev spaces, SIAM J. Numer. Anal. 54, 744-796.
M. Bachmayr and W. Dahmen (2020), Adaptive low-rank approximations for operator equations: Accuracy control and computational complexity, in 75 Years of Mathematics of Computation (S. C. Brenner et al., eds), Vol. 754 of Contemporary Mathematics, American Mathematical Society, pp. 1-44.
M. Bachmayr and M. Faldum (2023), A space-time adaptive low-rank method for highdimensional parabolic partial differential equations. Available at arXiv:2302.01658.
M. Bachmayr and V. Kazeev (2020), Stability of low-rank tensor representations and structured multilevel preconditioning for elliptic PDEs, Found. Comput. Math. 20, 11751236.
M. Bachmayr and R. Schneider (2017), Iterative methods based on soft thresholding of hierarchical tensors, Found. Comput. Math. 17, 1037-1083.
M. Bachmayr, A. Cohen and W. Dahmen (2018), Parametric PDEs: Sparse or low-rank approximations?, IMA J. Numer. Anal. 38, 1661-1708.
M. Bachmayr, A. Cohen, D. Dũng and C. Schwab (2017), Fully discrete approximation of parametric and stochastic elliptic PDEs, SIAM J. Numer. Anal. 55, 2151-2186.
M. Bachmayr, H. Eisenmann, E. Kieri and A. Uschmajew (2021a), Existence of dynamical low-rank approximations to parabolic problems, Math. Comp. 90, 1799-1830.
M. Bachmayr, M. Götte and M. Pfeffer (2022), Particle number conservation and block structures in matrix product states, Calcolo 59, 24.
M. Bachmayr, A. Nouy and R. Schneider (2021b), Approximation by tree tensor networks in high dimensions: Sobolev and compositional functions. Available at arXiv:2112.01474. To appear in Pure Appl. Funct. Anal.
M. Bachmayr, R. Schneider and A. Uschmajew (2016), Tensor networks and hierarchical tensors for the solution of high-dimensional partial differential equations, Found. Comput. Math. 16, 1423-1472.
J. Ballani and L. Grasedyck (2013), A projection method to solve linear systems in tensor format, Numer. Linear Algebra Appl. 20, 27-43.
J. Ballani and L. Grasedyck (2015), Hierarchical tensor approximation of output quantities of parameter-dependent PDEs, SIAM/ASA J. Uncertain. Quantif. 3, 852-872.
J. Ballani, L. Grasedyck and M. Kluge (2013), Black box approximation of tensors in hierarchical Tucker format, Linear Algebra Appl. 438, 639-657.
C. Bardos, I. Catto, N. Mauser and S. Trabelsi (2010), Setting and analysis of the multiconfiguration time-dependent Hartree-Fock equations, Arch. Ration. Mech. Anal. 198, 273-330.
C. Bardos, I. Catto, N. J. Mauser and S. Trabelsi (2009), Global-in-time existence of solutions to the multiconfiguration time-dependent Hartree-Fock equations: A sufficient condition, Appl. Math. Lett. 22, 147-152.
U. Benedikt, A. A. Auer, M. Espig and W. Hackbusch (2011), Tensor decomposition in post-Hartree-Fock methods, I: Two-electron integrals and MP2, J. Chem. Phys. 134, 054118.
U. Benedikt, H. Auer, M. Espig, W. Hackbusch and A. A. Auer (2013a), Tensor representation techniques in post-Hartree-Fock methods: Matrix product state tensor format, Molecular Phys. 111, 2398-2413.
U. Benedikt, K.-H. Böhm and A. A. Auer (2013b), Tensor decomposition in post-HartreeFock methods, II: CCD implementation, J. Chem. Phys. 139, 224101.
P. Benner, A. Cohen, M. Ohlberger and K. Willcox, eds (2017), Model Reduction and Approximation: Theory and Algorithms, Vol. 15 of Computational Science \& Engineering, SIAM.
G. Beylkin and M. J. Mohlenkamp (2002), Numerical operator calculus in higher dimensions, Proc. Natl. Acad. Sci. USA 99, 10246-10251.
G. Beylkin and M. J. Mohlenkamp (2005), Algorithms for numerical analysis in high dimensions, SIAM J. Sci. Comput. 26, 2133-2159.
G. Beylkin and L. Monzón (2005), On approximation of functions by exponential sums, Appl. Comput. Harmon. Anal. 19, 17-48.
G. Beylkin and L. Monzón (2010), Approximation by exponential sums revisited, Appl. Comput. Harmon. Anal. 28, 131-149.
G. Beylkin, M. J. Mohlenkamp and F. Pérez (2008), Approximating a wavefunction as an unconstrained sum of Slater determinants, J. Math. Phys. 49, 032107.
D. Bigoni, A. P. Engsig-Karup and Y. M. Marzouk (2016), Spectral tensor-train decomposition, SIAM J. Sci. Comput. 38, A2405-A2439.
M. Billaud-Friess, A. Nouy and O. Zahm (2014), A tensor approximation method based on ideal minimal residual formulations for the solution of high-dimensional problems, ESAIM Math. Model. Numer. Anal. 48, 1777-1806.
F. A. Bischoff and E. F. Valeev (2011), Low-order tensor approximations for electronic wave functions: Hartree-Fock method with guaranteed precision, J. Chem. Phys. 134, 104104.
T. Boiveau, V. Ehrlacher, A. Ern and A. Nouy (2019), Low-rank approximation of linear parabolic equations by space-time tensor Galerkin methods, ESAIM Math. Model. Numer. Anal. 53, 635-658.
D. Braess (1986), Nonlinear Approximation Theory, Vol. 7 of Springer Series in Computational Mathematics, Springer.
D. Braess and W. Hackbusch (2005), Approximation of $1 / x$ by exponential sums in $[1, \infty$ ), IMA J. Numer. Anal. 25, 685-697.
D. Braess and W. Hackbusch (2009), On the efficient computation of high-dimensional integrals and the approximation by exponential sums, in Multiscale, Nonlinear and Adaptive Approximation (R. DeVore and A. Kunoth, eds), Springer, pp. 39-74.
J. H. Bramble, J. E. Pasciak and J. Xu (1990), Parallel multilevel preconditioners, Math. Comp. 55, 1-22.
J.-F. Cai, E. J. Candès and Z. Shen (2010), A singular value thresholding algorithm for matrix completion, SIAM J. Optim. 20, 1956-1982.
E. Cancès, V. Ehrlacher and T. Lelièvre (2011), Convergence of a greedy algorithm for high-dimensional convex nonlinear problems, Math. Models Methods Appl. Sci. 21, 2433-2467.
E. Cancès, V. Ehrlacher and T. Lelièvre (2013), Greedy algorithms for high-dimensional non-symmetric linear problems, ESAIM Proc. 41, 95-131.
E. Cancès, V. Ehrlacher and T. Lelièvre (2014), Greedy algorithms for high-dimensional eigenvalue problems, Constr. Approx. 40, 387-423.
G. Ceruti and C. Lubich (2020), Time integration of symmetric and anti-symmetric lowrank matrices and Tucker tensors, BIT Numer. Math. 60, 591-614.
G. Ceruti and C. Lubich (2022), An unconventional robust integrator for dynamical lowrank approximation, BIT Numer. Math. 62, 23-44.
G. Ceruti, J. Kusch and C. Lubich (2022), A rank-adaptive robust integrator for dynamical low-rank approximation, BIT Numer. Math. 62, 1149-1174.
G. Ceruti, C. Lubich and D. Sulz (2023), Rank-adaptive time integration of tree tensor networks, SIAM J. Numer. Anal. 61, 194-222.
Z. Chen, K. Batselier, J. A. K. Suykens and N. Wong (2017), Parallelized tensor train learning of polynomial classifiers, IEEE Trans. Neural Netw. Learn. Syst. 29, 46214632.
A. Chertkov and I. V. Oseledets (2021), Solution of the Fokker-Planck equation by cross approximation method in the tensor train format, Front. Artif. Intell. 4, 668215.
S. R. Chinnamsetty, M. Espig, B. N. Khoromskij, W. Hackbusch and H.-J. Flad (2007), Tensor product approximation with optimal rank in quantum chemistry, J. Chem. Phys. 127, 084110.
H. Cho, D. Venturi and G. E. Karniadakis (2016), Numerical methods for high-dimensional probability density function equations, J. Comput. Phys. 305, 817-837.
A. Cichocki, D. Mandic, L. De Lathauwer, G. Zhou, Q. Zhao, C. Caiafa and H. A. Phan (2015), Tensor decompositions for signal processing applications: From two-way to multiway component analysis, IEEE Signal Process. Mag. 32, 145-163.
A. Cohen (2003), Numerical Analysis of Wavelet Methods, Vol. 32 of Studies in Mathematics and its Applications, North-Holland.
A. Cohen and R. DeVore (2015), Approximation of high-dimensional parametric PDEs, Acta Numer. 24, 1-159.
A. Cohen, W. Dahmen and R. DeVore (2001), Adaptive wavelet methods for elliptic operator equations: Convergence rates, Math. Comp. 70, 27-75.
A. Cohen, W. Dahmen and R. DeVore (2002), Adaptive wavelet methods, II: Beyond the elliptic case, Found. Comput. Math. 2, 203-245.
A. Cohen, I. Daubechies and J.-C. Feauveau (1992), Biorthogonal bases of compactly supported wavelets, Commun. Pure Appl. Math. 45, 485-560.
A. Cohen, R. DeVore and C. Schwab (2010), Convergence rates of best $N$-term Galerkin approximations for a class of elliptic sPDEs, Found. Comput. Math. 10, 615-646.
A. Cohen, R. DeVore and C. Schwab (2011), Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDE's, Anal. Appl. (Singap.) 9, 11-47.
D. Conte (2020), Dynamical low-rank approximation to the solution of parabolic differential equations, Appl. Numer. Math. 156, 377-384.
D. Conte and C. Lubich (2010), An error analysis of the multi-configuration time-dependent Hartree method of quantum dynamics, M2AN Math. Model. Numer. Anal. 44, 759-780.
G. M. Crosswhite and D. Bacon (2008), Finite automata for caching in matrix product algorithms, Phys. Rev. A 78, 012356.
C. Da Silva and F. J. Herrmann (2015), Optimization on the hierarchical Tucker manifold: Applications to tensor completion, Linear Algebra Appl. 481, 131-173.
S. Dahlke, M. Fornasier and T. Raasch (2012), Multilevel preconditioning and adaptive sparse solution of inverse problems, Math. Comp. 81, 419-446.
W. Dahmen (1997), Wavelet and multiscale methods for operator equations, Acta Numer. 6, 55-228.
W. Dahmen, R. DeVore, L. Grasedyck and E. Süli (2016), Tensor-sparsity of solutions to high-dimensional elliptic partial differential equations, Found. Comput. Math. 16, 813-874.
W. Dahmen, A. Kunoth and K. Urban (1999), Biorthogonal spline wavelets on the interval: Stability and moment conditions, Appl. Comput. Harmon. Anal. 6, 132-196.
I. Daubechies (1988), Orthonormal bases of compactly supported wavelets, Commun. Pure Appl. Math. 41, 909-996.
I. Daubechies (1992), Ten Lectures on Wavelets, Vol. 61 of CBMS-NSF Regional Conference Series in Applied Mathematics, SIAM.
L. De Lathauwer, B. De Moor and J. Vandewalle (2000), A multilinear singular value decomposition, SIAM J. Matrix Anal. Appl. 21, 1253-1278.
W. De Launey and J. Seberry (1994), The strong Kronecker product, J. Combin. Theory Ser. A 66, 192-213.
V. de Silva and L.-H. Lim (2008), Tensor rank and the ill-posedness of the best low-rank approximation problem, SIAM J. Matrix Anal. Appl. 30, 1084-1127.
A. Defant and K. Floret (1993), Tensor Norms and Operator Ideals, Vol. 176 of NorthHolland Mathematics Studies, North-Holland.
A. Dektor, A. Rodgers and D. Venturi (2021), Rank-adaptive tensor methods for highdimensional nonlinear pdes, J. Sci. Comput. 88, 36.
R. A. DeVore (1998), Nonlinear approximation, Acta Numer. 7, 51-150.
T. J. Dijkema, C. Schwab and R. Stevenson (2009), An adaptive wavelet method for solving high-dimensional elliptic PDEs, Constr. Approx. 30, 423-455.
P. A. M. Dirac (1930), Note on exchange phenomena in the Thomas atom, Math. Proc. Camb. Phil. Soc. 26, 376-385.
S. V. Dolgov and B. N. Khoromskij (2013), Two-level QTT-Tucker format for optimized tensor calculus, SIAM J. Matrix Anal. Appl. 34, 593-623.
S. V. Dolgov and D. V. Savostyanov (2014), Alternating minimal energy methods for linear systems in higher dimensions, SIAM J. Sci. Comput. 36, A2248-A2271.
S. V. Dolgov and R. Scheichl (2019), A hybrid alternating least squares-TT-cross algorithm for parametric PDEs, SIAM/ASA J. Uncertain. Quantif. 7, 260-291.
S. V. Dolgov and T. Vejchodský (2021), Guaranteed a posteriori error bounds for low-rank tensor approximate solutions, IMA J. Numer. Anal. 41, 1240-1266.
S. V. Dolgov, B. N. Khoromskij and I. V. Oseledets (2012), Fast solution of parabolic problems in the tensor train/quantized tensor train format with initial application to the Fokker-Planck equation, SIAM J. Sci. Comput. 34, A3016-A3038.
S. V. Dolgov, B. N. Khoromskij, I. V. Oseledets and D. V. Savostyanov (2014), Computation of extreme eigenvalues in higher dimensions using block tensor train format, Comput. Phys. Commun. 185, 1207-1216.
S. V. Dolgov, D. Kressner and C. Strössner (2021), Functional Tucker approximation using Chebyshev interpolation, SIAM J. Sci. Comput. 43, A2190-A2210.
J. Dölz, H. Egger and M. Schlottbom (2021), A model reduction approach for inverse problems with operator valued data, Numer. Math. 148, 889-917.
G. C. Donovan, J. S. Geronimo and D. P. Hardin (1996), Intertwining multiresolution analyses and the construction of piecewise-polynomial wavelets, SIAM J. Math. Anal. 27, 1791-1815.
G. C. Donovan, J. S. Geronimo and D. P. Hardin (1999), Orthogonal polynomials and the construction of piecewise polynomial smooth wavelets, SIAM J. Math. Anal. 30, 1029-1056.
A. J. Dunnett and A. W. Chin (2021), Efficient bond-adaptive approach for finitetemperature open quantum dynamics using the one-site time-dependent variational principle for matrix product states, Phys. Rev. B 104, 214302.
C. Eckart and G. Young (1936), The approximation of one matrix by another of lower rank, Psychometrika 1, 211-218.
M. Eigel, M. Marschall, M. Pfeffer and R. Schneider (2020), Adaptive stochastic Galerkin FEM for lognormal coefficients in hierarchical tensor representations, Numer. Math. 145, 655-692.
M. Eigel, M. Pfeffer and R. Schneider (2017), Adaptive stochastic Galerkin FEM with hierarchical tensor representations, Numer. Math. 136, 765-803.
L. Einkemmer and C. Lubich (2018), A low-rank projector-splitting integrator for the Vlasov-Poisson equation, SIAM J. Sci. Comput. 40, B1330-B1360.
L. Einkemmer and C. Lubich (2019), A quasi-conservative dynamical low-rank algorithm for the Vlasov equation, SIAM J. Sci. Comput. 41, B1061-B1081.
J. Eisert, M. Cramer and M. B. Plenio (2010), Colloquium: Area laws for the entanglement entropy, Rev. Mod. Phys. 82, 277-306.
K. Fackeldey, M. Oster, L. Sallandt and R. Schneider (2022), Approximative policy iteration for exit time feedback control problems driven by stochastic differential equations using tensor train format, Multiscale Model. Simul. 20, 379-403.
A. Falcó and W. Hackbusch (2012), On minimal subspaces in tensor representations, Found. Comput. Math. 12, 765-803.
A. Falcó and A. Nouy (2011), A proper generalized decomposition for the solution of elliptic problems in abstract form by using a functional Eckart-Young approach, J. Math. Anal. Appl. 376, 469-480.
A. Falcó and A. Nouy (2012), Proper generalized decomposition for nonlinear convex problems in tensor Banach spaces, Numer. Math. 121, 503-530.
A. Falcó, W. Hackbusch and A. Nouy (2019), On the Dirac-Frenkel variational principle on tensor Banach spaces, Found. Comput. Math. 19, 159-204.
A. Falcó, W. Hackbusch and A. Nouy (2021), Tree-based tensor formats, SeMA J. 78, 159-173.
F. M. Faulstich, A. Laestadius, O. Legeza, R. Schneider and S. Kvaal (2019), Analysis of the tailored coupled-cluster method in quantum chemistry, SIAM J. Numer. Anal. 57, 2579-2607.
F. Feppon and P. F. J. Lermusiaux (2018), A geometric approach to dynamical model order reduction, SIAM J. Matrix Anal. Appl. 39, 510-538.
F. Feppon and P. F. J. Lermusiaux (2019), The extrinsic geometry of dynamical systems tracking nonlinear matrix projections, SIAM J. Matrix Anal. Appl. 40, 814-844.
J. Frenkel (1934), Wave Mechanics: Advanced General Theory, Clarendon Press.
G. Friesecke and B. R. Graswald (2022), Two-electron wavefunctions are matrix product states with bond dimension three, J. Math. Phys. 63, 091901.
G. Friesecke, B. R. Graswald and Ö. Legeza (2022), Exact matrix product state representation and convergence of a fully correlated electronic wavefunction in the infinite-basis limit, Phys. Rev. B 105, 165144.
T. Gantumur, H. Harbrecht and R. Stevenson (2007), An optimal adaptive wavelet method without coarsening of the iterands, Math. Comp. 76, 615-629.
B. Gao and P.-A. Absil (2022), A Riemannian rank-adaptive method for low-rank matrix completion, Comput. Optim. Appl. 81, 67-90.
I. Gavrilyuk and B. N. Khoromskij (2019), Quasi-optimal rank-structured approximation to multidimensional parabolic problems by Cayley transform and Chebyshev interpolation, Comput. Methods Appl. Math. 19, 55-71.
I. P. Gavrilyuk, W. Hackbusch and B. N. Khoromskij (2005), Hierarchical tensor-product approximation to the inverse and related operators for high-dimensional elliptic problems, Computing 74, 131-157.
G. H. Golub and W. Kahan (1965), Calculating the singular values and pseudo-inverse of a matrix, J. Soc. Indust. Appl. Math. Ser. B Numer. Anal. 2, 205-224.
G. H. Golub and C. Reinsch (1970), Handbook Series Linear Algebra: Singular value decomposition and least squares solutions, Numer. Math. 14, 403-420.
A. Gorodetsky, S. Karaman and Y. Marzouk (2019), A continuous analogue of the tensortrain decomposition, Comput. Methods Appl. Mech. Engrg 347, 59-84.
L. Grasedyck (2004), Existence and computation of low Kronecker-rank approximations for large linear systems of tensor product structure, Computing 72, 247-265.
L. Grasedyck (2009/10), Hierarchical singular value decomposition of tensors, SIAM J. Matrix Anal. Appl. 31, 2029-2054.
L. Grasedyck (2010), Polynomial approximation in hierarchical Tucker format by vectortensorization. DFG SPP 1324 Preprint 43.
L. Grasedyck and W. Hackbusch (2011), An introduction to hierarchical ( $\mathcal{H}-$ ) rank and TT-rank of tensors with examples, Comput. Methods Appl. Math. 11, 291-304.
L. Grasedyck, D. Kressner and C. Tobler (2013), A literature survey of low-rank tensor approximation techniques, GAMM-Mitt. 36, 53-78.
B. R. Graswald and G. Friesecke (2021), Electronic wavefunction with maximally entangled mps representation, Europ. Phys. J. D 75, 1-4.
M. Griebel and H. Harbrecht (2014), Approximation of bi-variate functions: Singular value decomposition versus sparse grids, IMA J. Numer. Anal. 34, 28-54.
M. Griebel and H. Harbrecht (2019), Singular value decomposition versus sparse grids: Refined complexity estimates, IMA J. Numer. Anal. 39, 1652-1671.
M. Griebel and H. Harbrecht (2023), Analysis of tensor approximation schemes for continuous functions, Found. Comput. Math. 23, 219-240.
M. Griebel, H. Harbrecht and R. Schneider (2022), Low-rank approximation of continuous functions in Sobolev spaces with dominating mixed smoothness. Available at arXiv:2203.04100.
M. Gu and S. C. Eisenstat (1995), A divide-and-conquer algorithm for the bidiagonal SVD, SIAM J. Matrix Anal. Appl. 16, 79-92.
W. Hackbusch (2005), Entwicklungen nach Exponentialsummen. Technical Report 4, MPI MIS Leipzig.
W. Hackbusch (2011), Tensorisation of vectors and their efficient convolution, Numer. Math. 119, 465-488.
W. Hackbusch (2014), Numerical tensor calculus, Acta Numer. 23, 651-742.
W. Hackbusch (2015a), Hierarchical Matrices: Algorithms and Analysis, Vol. 49 of Springer Series in Computational Mathematics, Springer.
W. Hackbusch (2015b), Solution of linear systems in high spatial dimensions, Comput. Vis. Sci. 17, 111-118.
W. Hackbusch (2018), On the representation of symmetric and antisymmetric tensors, in Contemporary Computational Mathematics: A Celebration of the 80th Birthday of Ian Sloan (J. Dick et al., eds), Springer, pp. 483-515.
W. Hackbusch (2019), Tensor Spaces and Numerical Tensor Calculus, Vol. 56 of Springer Series in Computational Mathematics, second edition, Springer.
W. Hackbusch and B. N. Khoromskij (2006a), Low-rank Kronecker-product approximation to multi-dimensional nonlocal operators, I: Separable approximation of multi-variate functions, Computing 76, 177-202.
W. Hackbusch and B. N. Khoromskij (2006b), Low-rank Kronecker-product approximation to multi-dimensional nonlocal operators, II: HKT representation of certain operators, Computing 76, 203-225.
W. Hackbusch and S. Kühn (2009), A new scheme for the tensor representation, J. Fourier Anal. Appl. 15, 706-722.
W. Hackbusch, B. N. Khoromskij and E. E. Tyrtyshnikov (2008), Approximate iterations for structured matrices, Numer. Math. 109, 365-383.
W. Hackbusch, B. N. Khoromskij, S. Sauter and E. E. Tyrtyshnikov (2012), Use of tensor formats in elliptic eigenvalue problems, Numer. Linear Algebra Appl. 19, 133-151.
J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken and F. Verstraete (2016), Unifying time evolution and optimization with matrix product states, Phys. Rev. B 94, 165116.
J. Haegeman, T. J. Osborne and F. Verstraete (2013), Post-matrix product state methods: To tangent space and beyond, Phys. Rev. B 88, 075133.
T. Helgaker, P. Jørgensen and J. Olsen (2000), Molecular Electronic-Structure Theory, Wiley.
U. Helmke and M. A. Shayman (1995), Critical points of matrix least squares distance functions, Linear Algebra Appl. 215, 1-19.
J. S. Hesthaven, G. Rozza and B. Stamm (2016), Certified Reduced Basis Methods for Parametrized Partial Differential Equations, SpringerBriefs in Mathematics, Springer; BCAM Basque Center for Applied Mathematics, Bilbao.
C. J. Hillar and L.-H. Lim (2013), Most tensor problems are NP-hard, J. Assoc. Comput. Mach. 60, Art. 45, 39.
F. L. Hitchcock (1927), The expression of a tensor or a polyadic as a sum of products, J. Math. Phys. 6, 164-189.
F. L. Hitchcock (1928), Multiple invariants and generalized rank of a $p$-way matrix or tensor, J. Math. Phys. 7, 39-79.
S. Holtz, T. Rohwedder and R. Schneider (2012a), The alternating linear scheme for tensor optimization in the tensor train format, SIAM J. Sci. Comput. 34, A683-A713.
S. Holtz, T. Rohwedder and R. Schneider (2012b), On manifolds of tensors of fixed TT-rank, Numer. Math. 120, 701-731.
M. Janzamin, R. Ge, J. Kossaifi, A. Anandkumar et al. (2019), Spectral learning on matrices and tensors, Found. Trends Mach. Learn. 12, 393-536.
M. Kahlbacher and S. Volkwein (2007), Galerkin proper orthogonal decomposition methods for parameter dependent elliptic systems, Discuss. Math. Differ. Incl. Control Optim. 27, 95-117.
Y. Kazashi and F. Nobile (2021), Existence of dynamical low rank approximations for random semi-linear evolutionary equations on the maximal interval, Stoch. Partial Differ. Equ. Anal. Comput. 9, 603-629.
Y. Kazashi, F. Nobile and E. Vidličková (2021), Stability properties of a projector-splitting scheme for dynamical low rank approximation of random parabolic equations, Numer. Math. 149, 973-1024.
V. A. Kazeev (2015), Quantized tensor structured finite elements for second-order elliptic PDEs in two dimensions. PhD thesis, ETH Zürich.
V. A. Kazeev and B. N. Khoromskij (2012), Low-rank explicit QTT representation of the Laplace operator and its inverse, SIAM J. Matrix Anal. Appl. 33, 742-758.
V. A. Kazeev and C. Schwab (2018), Quantized tensor-structured finite elements for secondorder elliptic PDEs in two dimensions, Numer. Math. 138, 133-190.
V. A. Kazeev, B. N. Khoromskij and E. E. Tyrtyshnikov (2013a), Multilevel Toeplitz matrices generated by tensor-structured vectors and convolution with logarithmic complexity, SIAM J. Sci. Comput. 35, A1511-A1536.
V. A. Kazeev, I. V. Oseledets, M. Rakhuba and C. Schwab (2017), QTT-finite-element approximation for multiscale problems, I: Model problems in one dimension, $A d v$. Comput. Math. 43, 411-442.
V. A. Kazeev, I. V. Oseledets, M. V. Rakhuba and C. Schwab (2022), Quantized tensor FEM for multiscale problems: Diffusion problems in two and three dimensions, Multiscale Model. Simul. 20, 893-935.
V. A. Kazeev, O. Reichmann and C. Schwab (2013b), Low-rank tensor structure of linear diffusion operators in the TT and QTT formats, Linear Algebra Appl. 438, 4204-4221.
S. Keller, M. Dolfi, M. Troyer and M. Reiher (2015), An efficient matrix product operator representation of the quantum chemical Hamiltonian, J. Chem. Phys. 143, 244118.
V. Khoromskaia and B. N. Khoromskij (2015), Tensor numerical methods in quantum chemistry: From Hartree-Fock to excitation energies, Phys. Chem. Chem. Phys. 17, 31491-31509.
V. Khoromskaia and B. N. Khoromskij (2016), Fast tensor method for summation of longrange potentials on 3D lattices with defects, Numer. Linear Algebra Appl. 23, 249-271.
V. Khoromskaia and B. N. Khoromskij (2018), Tensor Numerical Methods in Quantum Chemistry, De Gruyter.
V. Khoromskaia, B. N. Khoromskij and R. Schneider (2013), Tensor-structured factorized calculation of two-electron integrals in a general basis, SIAM J. Sci. Comput. 35, A987A1010.
B. N. Khoromskij (2009), Tensor-structured preconditioners and approximate inverse of elliptic operators in $\mathbb{R}^{d}$, Constr. Approx. 30, 599-620.
B. N. Khoromskij (2011), $O(d \log N)$-quantics approximation of $N-d$ tensors in highdimensional numerical modeling, Constr. Approx. 34, 257-280.
B. N. Khoromskij (2015), Tensor numerical methods for multidimensional PDEs: Theoretical analysis and initial applications, ESAIM Proc. Surveys 48, 1-28.
B. N. Khoromskij (2018), Tensor Numerical Methods in Scientific Computing, Vol. 19 of Radon Series on Computational and Applied Mathematics, De Gruyter.
B. N. Khoromskij and S. Miao (2014), Superfast wavelet transform using quantics-TT approximation, I: Application to Haar wavelets, Comput. Methods Appl. Math. 14, 537-553.
B. N. Khoromskij and I. V. Oseledets (2010), Quantics-TT collocation approximation of parameter-dependent and stochastic elliptic PDEs, Comput. Methods Appl. Math. 10, 376-394.
B. N. Khoromskij and I. V. Oseledets (2011), QTT approximation of elliptic solution operators in higher dimensions, Russian J. Numer. Anal. Math. Modelling 26, 303-322.
B. N. Khoromskij and C. Schwab (2011), Tensor-structured Galerkin approximation of parametric and stochastic elliptic PDEs, SIAM J. Sci. Comput. 33, 364-385.
B. N. Khoromskij, V. Khoromskaia and H.-J. Flad (2011), Numerical solution of the Hartree-Fock equation in multilevel tensor-structured format, SIAM J. Sci. Comput. 33, 45-65.
B. N. Khoromskij, V. Khoromskaia, S. R. Chinnamsetty and H.-J. Flad (2009), Tensor decomposition in electronic structure calculations on 3D Cartesian grids, J. Comput. Phys. 228, 5749-5762.
E. Kieri and B. Vandereycken (2019), Projection methods for dynamical low-rank approximation of high-dimensional problems, Comput. Methods Appl. Math. 19, 73-92.
E. Kieri, C. Lubich and H. Walach (2016), Discretized dynamical low-rank approximation in the presence of small singular values, SIAM J. Numer. Anal. 54, 1020-1038.
O. Koch and C. Lubich (2007a), Dynamical low-rank approximation, SIAM J. Matrix Anal. Appl. 29, 434-454.
O. Koch and C. Lubich (2007b), Regularity of the multi-configuration time-dependent Hartree approximation in quantum molecular dynamics, M2AN Math. Model. Numer. Anal. 41, 315-331.
O. Koch and C. Lubich (2010), Dynamical tensor approximation, SIAM J. Matrix Anal. Appl. 31, 2360-2375.
O. Koch and C. Lubich (2011), Variational-splitting time integration of the multiconfiguration time-dependent Hartree-Fock equations in electron dynamics, IMA J. Numer. Anal. 31, 379-395.
T. G. Kolda and B. W. Bader (2009), Tensor decompositions and applications, SIAM Rev. 51, 455-500.
D. Kressner and C. Tobler (2011a), Low-rank tensor Krylov subspace methods for parametrized linear systems, SIAM J. Matrix Anal. Appl. 32, 1288-1316.
D. Kressner and C. Tobler (2011b), Preconditioned low-rank methods for high-dimensional elliptic PDE eigenvalue problems, Comput. Methods Appl. Math. 11, 363-381.
D. Kressner and A. Uschmajew (2016), On low-rank approximability of solutions to highdimensional operator equations and eigenvalue problems, Linear Algebra Appl. 493, 556-572.
D. Kressner, M. Steinlechner and A. Uschmajew (2014a), Low-rank tensor methods with subspace correction for symmetric eigenvalue problems, SIAM J. Sci. Comput. 36, A2346-A2368.
D. Kressner, M. Steinlechner and B. Vandereycken (2014b), Low-rank tensor completion by Riemannian optimization, BIT Numer. Math. 54, 447-468.
D. Kressner, M. Steinlechner and B. Vandereycken (2016), Preconditioned low-rank Riemannian optimization for linear systems with tensor product structure, SIAM J. Sci. Comput. 38, A2018-A2044.
C. Krumnow, M. Pfeffer and A. Uschmajew (2021), Computing eigenspaces with low rank constraints, SIAM J. Sci. Comput. 43, A586-A608.
C. Krumnow, L. Veis, O. Legeza and J. Eisert (2016), Fermionic orbital optimization in tensor network states, Phys. Rev. Lett. 117, 210402.
S. Kühn (2012), Hierarchische Tensordarstellung. PhD thesis, Universität Leipzig.
J. M. Landsberg (2012), Tensors: Geometry and Applications, American Mathematical Society.
J. M. Landsberg, Y. Qi and K. Ye (2012), On the geometry of tensor network states, Quantum Inf. Comput. 12, 346-354.
K. Lee and H. C. Elman (2017), A preconditioned low-rank projection method with a rank-reduction scheme for stochastic partial differential equations, SIAM J. Sci. Comput. 39, S828-S850.
W. A. Light and E. W. Cheney (1985), Approximation Theory in Tensor Product Spaces, Vol. 1169 of Lecture Notes in Mathematics, Springer.
L.-H. Lim (2021), Tensors in computations, Acta Numer. 30, 555-764.
C. Lubich (2005), On variational approximations in quantum molecular dynamics, Math. Comp. 74, 765-779.
C. Lubich (2008), From Quantum to Classical Molecular Dynamics: Reduced Models and Numerical Analysis, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS).
C. Lubich (2015), Time integration in the multiconfiguration time-dependent Hartree method of molecular quantum dynamics, Appl. Math. Res. Express. AMRX 2015, 311328.
C. Lubich and I. V. Oseledets (2014), A projector-splitting integrator for dynamical lowrank approximation, BIT Numer. Math. 54, 171-188.
C. Lubich, I. V. Oseledets and B. Vandereycken (2015), Time integration of tensor trains, SIAM J. Numer. Anal. 53, 917-941.
C. Lubich, T. Rohwedder, R. Schneider and B. Vandereycken (2013), Dynamical approximation by hierarchical Tucker and tensor-train tensors, SIAM J. Matrix Anal. Appl. 34, 470-494.
C. Marcati, M. Rakhuba and C. Schwab (2022a), Tensor rank bounds for point singularities in $\mathbb{R}^{3}, A d v$. Comput. Math. 48, 18.
C. Marcati, M. Rakhuba and J. E. M. Ulander (2022b), Low-rank tensor approximation of singularly perturbed boundary value problems in one dimension, Calcolo 59, 2.
L. Markeeva, I. Tsybulin and I. V. Oseledets (2021), QTT-isogeometric solver in two dimensions, J. Comput. Phys. 424, 109835.
H. G. Matthies and E. Zander (2012), Solving stochastic systems with low-rank tensor compression, Linear Algebra Appl. 436, 3819-3838.
H.-D. Meyer, F. Gatti and G. A. Worth (2009), Multidimensional Quantum Dynamics: MCTDH Theory and Applications, Wiley.
H.-D. Meyer, U. Manthe and L. S. Cederbaum (1990), The multi-configurational timedependent Hartree approach, Chem. Phys. Lett. 165, 73-78.
B. Michel and A. Nouy (2022), Learning with tree tensor networks: Complexity estimates and model selection, Bernoulli 28, 910-936.
M. J. Mohlenkamp (2010), A center-of-mass principle for the multiparticle Schrödinger equation, J. Math. Phys. 51, 022112.
J.-J. Moreau (1965), Proximité et dualité dans un espace hilbertien, Bull. Soc. Math. France 93, 273-299.
V. Murg, F. Verstraete, O. Legeza and R. M. Noack (2010), Simulating strongly correlated quantum systems with tree tensor networks, Phys. Rev. B 82, 205105.
E. Musharbash and F. Nobile (2018), Dual dynamically orthogonal approximation of incompressible Navier Stokes equations with random boundary conditions, J. Comput. Phys. 354, 135-162.
E. Musharbash, F. Nobile and E. Vidličková (2020), Symplectic dynamical low rank approximation of wave equations with random parameters, BIT Numer. Math. 60, 11531201.
E. Musharbash, F. Nobile and T. Zhou (2015), Error analysis of the dynamically orthogonal approximation of time dependent random PDEs, SIAM J. Sci. Comput. 37, A776-A810.
A. Novikov, M. Trofimov and I. Oseledets (2018), Exponential machines, Bull. Pol. Acad. Sci. Math. 66, 789-797.
R. Orús (2014), A practical introduction to tensor networks: Matrix product states and projected entangled pair states, Ann. Phys. 349, 117-158.
I. V. Oseledets (2009a), Approximation of matrices with logarithmic number of parameters, Dokl. Math. 80, 653-654.
I. V. Oseledets (2009b), On a new tensor decomposition, Dokl. Akad. Nauk 427, 168-169. In Russian; English translation in Dokl. Math. 80, 495-496 (2009).
I. V. Oseledets (2011a), DMRG approach to fast linear algebra in the TT-format, Comput. Methods Appl. Math. 11, 382-393.
I. V. Oseledets (2011b), Tensor-train decomposition, SIAM J. Sci. Comput. 33, 2295-2317.
I. V. Oseledets (2013), Constructive representation of functions in low-rank tensor formats, Constr. Approx. 37, 1-18.
I. V. Oseledets and S. V. Dolgov (2012), Solution of linear systems and matrix inversion in the TT-format, SIAM J. Sci. Comput. 34, A2718-A2739.
I. V. Oseledets and E. E. Tyrtyshnikov (2009a), Breaking the curse of dimensionality, or how to use SVD in many dimensions, SIAM J. Sci. Comput. 31, 3744-3759.
I. V. Oseledets and E. E. Tyrtyshnikov (2009b), Recursive decomposition of multidimensional tensors, Dokl. Akad. Nauk 427, 14-16. In Russian; English translation in Dokl. Math. 80, 460-462 (2009).
I. V. Oseledets and E. E. Tyrtyshnikov (2010), TT-cross approximation for multidimensional arrays, Linear Algebra Appl. 432, 70-88.
I. V. Oseledets, M. V. Rakhuba and A. Uschmajew (2018), Alternating least squares as moving subspace correction, SIAM J. Numer. Anal. 56, 3459-3479.
M. Oster, L. Sallandt and R. Schneider (2022), Approximating optimal feedback controllers of finite horizon control problems using hierarchical tensor formats, SIAM J. Sci. Comput. 44, B746-B770.
A. Ostermann, C. Piazzola and H. Walach (2019), Convergence of a low-rank Lie-Trotter splitting for stiff matrix differential equations, SIAM J. Numer. Anal. 57, 1947-1966.
B. N. Parlett (1998), The Symmetric Eigenvalue Problem, Vol. 20 of Classics in Applied Mathematics, SIAM. Corrected reprint of the 1980 original.
M. Pfeffer (2018), Tensor methods for the numerical solution of high-dimensional parametric partial differential equations. PhD thesis, Technische Universität Berlin.
M. Primbs (2010), New stable biorthogonal spline-wavelets on the interval, Results Math. 57, 121-162.
C. Prud'homme and A. T. Patera (2004), Reduced-basis output bounds for approximately parameterized elliptic coercive partial differential equations, Comput. Vis. Sci. 6, 147162.
M. Rakhuba (2021), Robust alternating direction implicit solver in quantized tensor formats for a three-dimensional elliptic PDE, SIAM J. Sci. Comput. 43, A800-A827.
M. Rakhuba, A. Novikov and I. V. Oseledets (2019), Low-rank Riemannian eigensolver for high-dimensional Hamiltonians, J. Comput. Phys. 396, 718-737.
M. V. Rakhuba and I. V. Oseledets (2016), Grid-based electronic structure calculations: The tensor decomposition approach, J. Comput. Phys. 312, 19-30.
M. V. Rakhuba and I. V. Oseledets (2018), Jacobi-Davidson method on low-rank matrix manifolds, SIAM J. Sci. Comput. 40, A1149-A1170.
L. Richter, L. Sallandt and N. Nüsken (2021), Solving high-dimensional parabolic PDEs using the tensor train format, in 38th International Conference on Machine Learning, Vol. 139 of Proceedings of Machine Learning Research, PMLR, pp. 8998-9009.
T. Rohwedder and A. Uschmajew (2013), On local convergence of alternating schemes for optimization of convex problems in the tensor train format, SIAM J. Numer. Anal. 51, 1134-1162.
G. Rozza, D. B. P. Huynh and A. T. Patera (2008), Reduced basis approximation and $a$ posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: Application to transport and continuum mechanics, Arch. Comput. Methods Engrg 15, 229-275.
R. A. Ryan (2002), Introduction to Tensor Products of Banach Spaces, Springer Monographs in Mathematics, Springer.
T. P. Sapsis and P. F. Lermusiaux (2009), Dynamically orthogonal field equations for continuous stochastic dynamical systems, Phys. D 238, 2347-2360.
E. Schmidt (1907), Zur Theorie der linearen und nichtlinearen Integralgleichungen, Math. Ann. 63, 433-476.
R. Schneider and A. Uschmajew (2014), Approximation rates for the hierarchical tensor format in periodic Sobolev spaces, J. Complexity 30, 56-71.
R. Schneider and A. Uschmajew (2015), Convergence results for projected line-search methods on varieties of low-rank matrices via Łojasiewicz inequality, SIAM J. Optim. 25, 622-646.
U. Schollwöck (2011), The density-matrix renormalization group in the age of matrix product states, Ann. Phys. 326, 96-192.
S. Scholz and H. Yserentant (2017), On the approximation of electronic wavefunctions by anisotropic Gauss and Gauss-Hermite functions, Numer. Math. 136, 841-874.
C. Schwab and C. J. Gittelson (2011), Sparse tensor discretizations of high-dimensional parametric and stochastic PDEs, Acta Numer. 20, 291-467.
C. Schwab and R. Stevenson (2009), Space-time adaptive wavelet methods for parabolic evolution problems, Math. Comp. 78, 1293-1318.
Y.-Y. Shi, L.-M. Duan and G. Vidal (2006), Classical simulation of quantum many-body systems with a tree tensor network, Phys. Rev. A 74, 022320.
N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis and C. Faloutsos (2017), Tensor decomposition for signal processing and machine learning, IEEE Trans. Signal Process. 65, 3551-3582.
M. Signoretto, Q. Tran Dinh, L. De Lathauwer and J. A. K. Suykens (2014), Learning with tensors: A framework based on convex optimization and spectral regularization, Mach. Learn. 94, 303-351.
S. Singh, R. N. C. Pfeifer and G. Vidal (2011), Tensor network states and algorithms in the presence of a global $U(1)$ symmetry, Phys. Rev. B 83, 115125.
M. Steinlechner (2016), Riemannian optimization for high-dimensional tensor completion, SIAM J. Sci. Comput. 38, S461-S484.
F. Stenger (1993), Numerical Methods Based on Sinc and Analytic Functions, Vol. 20 of Springer Series in Computational Mathematics, Springer.
R. Stevenson (2009), Adaptive wavelet methods for solving operator equations: An overview, in Multiscale, Nonlinear and Adaptive Approximation (R. DeVore and A. Kunoth, eds), Springer, pp. 543-597.
E. Stoudenmire and D. J. Schwab (2016), Supervised learning with tensor networks, in Advances in Neural Information Processing Systems 29 (D. Lee et al., eds), Curran Associates.
Y. Sun and M. Kumar (2014), Numerical solution of high dimensional stationary FokkerPlanck equations via tensor decomposition and Chebyshev spectral differentiation, Comput. Math. Appl. 67, 1960-1977.
S. Szalay, M. Pfeffer, V. Murg, G. Barcza, F. Verstraete, R. Schneider and Ö. Legeza (2015), Tensor product methods and entanglement optimization for $a b$ initio quantum chemistry, Internat. J. Quantum Chem. 115, 1342-1391.
C. Tobler (2012), Low-rank tensor methods for linear systems and eigenvalue problems. PhD thesis, ETH Zürich.
L. R. Tucker (1964), The extension of factor analysis to three-dimensional matrices, in Contributions to Mathematical Psychology (H. Gulliksen and N. Frederiksen, eds), Holt, Rinehart \& Winston, pp. 109-127.
L. R. Tucker (1966), Some mathematical notes on three-mode factor analysis, Psychometrika 31, 279-311.
E. E. Tyrtyshnikov (2003), Tensor approximations of matrices generated by asymptotically smooth functions, Mat. Sb. 194, 147-160.
A. Uschmajew (2010), Well-posedness of convex maximization problems on Stiefel manifolds and orthogonal tensor product approximations, Numer. Math. 115, 309-331.
A. Uschmajew (2012), Local convergence of the alternating least squares algorithm for canonical tensor approximation, SIAM J. Matrix Anal. Appl. 33, 639-652.
A. Uschmajew (2013), Zur Theorie der Niedrigrangapproximation in Tensorprodukten von Hilberträumen. PhD thesis, Technische Universität Berlin.
A. Uschmajew and B. Vandereycken (2013), The geometry of algorithms using hierarchical tensors, Linear Algebra Appl. 439, 133-166.
A. Uschmajew and B. Vandereycken (2015), Greedy rank updates combined with Riemannian descent methods for low-rank optimization, in 2015 International Conference on Sampling Theory and Applications (SampTA), pp. 420-424.
A. Uschmajew and B. Vandereycken (2020), Geometric methods on low-rank matrix and tensor manifolds, in Handbook of Variational Methods for Nonlinear Geometric Data (P. Grohs et al., eds), Springer, pp. 261-313.
G. Vidal (2003), Efficient classical simulation of slightly entangled quantum computations, Phys. Rev. Lett. 91, 147902.
H. Wang and M. Thoss (2003), Multilayer formulation of the multiconfiguration timedependent Hartree theory, J. Chem Phys. 119, 1289-1299.
L. Wang and M. T. Chu (2014), On the global convergence of the alternating least squares method for rank-one approximation to generic tensors, SIAM J. Matrix Anal. Appl. 35, 1058-1072.
J. Weidmann (1980), Linear Operators in Hilbert Spaces, Vol. 68 of Graduate Texts in Mathematics, Springer.
S. R. White (1992), Density matrix formulation for quantum renormalization groups, Phys. Rev. Lett. 69, 2863-2866.
S. R. White (2005), Density matrix renormalization group algorithms with a single center site, Phys. Rev. B 72, 180403.
S. R. White and R. L. Martin (1999), Ab initio quantum chemistry using the density matrix renormalization group, J. Chem. Phys. 110, 4127-4130.
J. H. Wilkinson (1968), Global convergence of tridiagonal QR algorithm with origin shifts, Linear Algebra Appl. 1, 409-420.
S. Wouters and D. Van Neck (2014), The density matrix renormalization group for ab initio quantum chemistry, Europ. Phys. J. D 68, 1-20.
M. Yang and S. R. White (2020), Time-dependent variational principle with ancillary Krylov subspace, Phys. Rev. B 102, 094315.
H. Yserentant (2010), Regularity and Approximability of Electronic Wave Functions, Springer.
A. Zeiser (2010), Direkte Diskretisierung der Schrödinger-Gleichung auf dünnen Gittern. PhD thesis, TU Berlin.
G. Zhou, W. Huang, K. A. Gallivan, P. Van Dooren and P.-A. Absil (2016), A Riemannian rank-adaptive method for low-rank optimization, Neurocomput. 192, 72-80.


[^0]:    ${ }^{1}$ This corresponds to the observation that for operators of the form (1.6), the solution (1.7) is exactly reproduced by dynamical low-rank approximation.

