

Iterative algorithms for the post-processing of high-dimensional data

Mike Espig

Westfälische Hochschule Zwickau, Dr.-Friedrichs-Ring 2A, D-08012 Zwickau, Germany

Wolfgang Hackbusch

Max-Planck-Institute for Mathematics in the Sciences (MIS), Inselstr. 22, D-04103 Leipzig, Germany

Alexander Litvinenko*

Rheinisch-Westfälische Technische Hochschule (RWTH), Kackertstr. 9, D-52072 Aachen, Germany

Hermann G. Matthies and Elmar Zander*

Technische Universität Braunschweig, Mühlentorstrasse 23, D-38106 Braunschweig, Germany

Abstract

Scientific computations or measurements may result in huge volumes of high-dimensional data, for instance 10^{20} or 100^{300} elements. Often these can be thought of representing a real-valued function on a high-dimensional domain. In this and also in most other cases the data can be conceptually arranged in the format of a tensor of high degree, and stored in some truncated or lossy compressed format. We look at some common post-processing tasks which are too time and storage consuming in the uncompressed data format and not obvious in the compressed format, as such huge data sets can not be stored in their entirety, and the value of an element is not readily accessible through simple look-up. The tasks we consider are finding the location of maximum or minimum, or finding all elements in some interval — i.e. level sets, or the number of elements with a value in such a level set, the probability of an element being in a particular level set, and the mean and variance of the total collection. The algorithms to be described are fixed point iterations of particular point-wise functions of the data, which will then exhibit the desired result. To formulate these algorithms, the data is considered as an element of a commutative algebra, and in an abstract sense, the algorithms, which only use these algebraic operations, are independent of the representation. We allow the actual computational representation to be a lossy compression, and we allow the algebra operations to be performed in an approximate fashion, so as to maintain a high compression level. One such example format which is addressed explicitly and described in some detail is the representation of the data as a tensor with compression in the form of a low-rank representation.

Keywords: post-processing, fixed point iteration, high-dimensional data, compression, eigenvalue computation, low-rank tensor representation

1. Introduction

Many scientific and engineering computations or measurements, as well as economic or financial applications, produce large volumes of data. Often these data are numerical values indexed by different parameters (here called dimensions d) and present for many different parameter values (here called sample points n). Some such high-dimensional problems from chemistry and physics are the Hartree-Fock-, Schrödinger-, or Master-equations, and low-rank tensor methods of their solution are for example considered in [43, 52, 53, 16]. In many chemical applications one has $n = 100$, and d is a few hundreds or even thousands [10], making in total $N = n^d$ entries. Such data can conveniently be thought of a multi-indexed array, or in other words a tensor. Also discretisations of stochastic partial

*Corresponding author

Email address: `litvinenko@uq.rwth-aachen.de` (Alexander Litvinenko)

differential equations, or uncertainty quantification problems in general, may result in very high-dimensional problems [65, 64, 19, 56, 83, 87]. The reason is that some random or uncertain coefficient fields like the porosity or permeability may require many hundreds of independent random variables to represent it accurately [23, 55, 6, 94, 66, 15, 20].

For simplicity, assume that we have computed or observed large quantities of one real-valued variate. Assume further that the amount of data is so large that it can not be held or stored in its entirety and has to be compressed in some way.

Frequently, the task arises for example to find the location with the maximum value of the data set, or to find the locations of all values which lie in a given interval — we call this a level set. Other, similar tasks we consider are finding the number of values in a given interval, the probability of being in a given interval, or finding the mean or the variance of the data set. All these tasks can be used, for example, to construct high-dimensional histograms and probability density functions. These tasks above are trivial for small data sets, typically performed by inspecting each datum. But if we consider truly huge amounts of data which can not be stored in full because this would exceed the capacity of any storage device, but only in a compressed manner, and where additionally it may not be possible due to time constraints to inspect each element, these tasks are not trivial any more. This is due to the fact that in the compressed representation the data values are normally not directly accessible and require additional processing. Such compression will in general be “lossy”, so that not each value can be restored exactly, but only up to a certain reconstruction error.

Assuming the data as an element of some set \mathcal{T} , the algorithms are independent of the representation of the data, as well as from the compression and reconstruction technique used, subject only to the possibility of approximately performing the operations of an Euclidean associative commutative algebra. This means that we assume that \mathcal{T} is a vector space with an inner product, and additionally an associative commutative bilinear multiplication, making it into an algebra. The simplest example of such a structure is to envisage the data points as a vector $\mathbf{w} \in \mathbb{R}^N$ with an appropriate $N \in \mathbb{N}$, reflecting the amount of data. The vector space operations are clear, as is the canonical Euclidean product, and for the commutative multiplication consider the point-wise or Hadamard product, i.e. the component-wise multiplication of two vectors. We shall allow that all these operations are performed only approximately, in order to maintain a high compression level.

Large volumes of data, especially when they can be thought of as samples or discrete values of some real-valued function on a high-dimensional space, can often be arranged in form of a tensor [57, 58, 13, 47]. This offers the possibility to use approximate or compressed tensor representations. Here we will especially show one possibility, namely low-rank representations, which generalise the truncated singular value decomposition for matrices. For the sake of completeness, we shall show the possible implementation of the above mentioned algebraic operations for some of the more common low-rank formats, and the approximated, compressed, or truncated representation of their results; see also [75, 29].

The proposed algorithms are iterative in nature, and the convergence tolerance can be adapted to the reconstruction error. The basic idea for the algorithms, which operate only on the algebraic structure, is an iterative scheme which converges to a result which solves the desired problem in some way. Such iterations typically destroy the compressed representation, so they have to be combined with re-compression or truncation.

1.1. An example

Let us give an example which motivates much of the following formulation and development. In uncertainty quantification (UQ) [65, 64, 19, 56, 83], one may have to solve a partial differential equation which may contain uncertain or random coefficient fields, see Subsection 4.2 for a small numerical example of this kind. Such problems lead naturally to the concept that the solution is to be sought in a tensor product space. An example of a solution could be a function $v(\mathbf{x}, t, \boldsymbol{\mu})$, where \mathbf{x} is a spatial variable, t is time, and $\boldsymbol{\mu}$ is a vector of parameters describing the random field, i.e. this is the solution of a parametric partial differential equation. Numerical approximations can be computed through various kinds of discretisation. Quite often a goal is to compute a so-called surrogate model, an approximation which allows an evaluation of $v(\mathbf{x}, t, \boldsymbol{\mu})$ for different values of $\boldsymbol{\mu}$ without the cost of each time

numerically solving a partial differential equation. In the simplest discrete setting, one may want to compute such a solution at distinct values or samples of these variables. If we think of samples of $v(\mathbf{x}, t, \boldsymbol{\mu})$ at space points $(\mathbf{x}_k)_{k=1}^K$, time instances $(t_j)_{j=1}^J$ and parameter values $(\boldsymbol{\mu}_i)_i = (\mu_{i_1}, \dots, \mu_{i_L})$, with a multi-index $\mathbf{i} = (i_1, \dots, i_L) \in \mathbb{N}^L$ with $1 \leq i_\ell \leq I_\ell$, then the samples of the solution form a multi-indexed array

$$\mathbf{v} = (v(x_k, t_j, \mu_{i_1}, \dots, \mu_{i_L})) = (v_{kj i_1 \dots i_L}) \in \mathbb{R}^K \otimes \mathbb{R}^J \otimes \mathbb{R}^{I_1} \otimes \dots \otimes \mathbb{R}^{I_L}, \quad (1)$$

or in other words a high order or degree *tensor* [47]. Similar high dimensional functions and approximations appear in Bayesian inverse problems, see [87] for a review, or in probabilistic estimation with random fields like ‘‘Kriging’’, see [20] for a recent example in the context of low-rank tensor representations.

For the sake of simplicity, we will treat all arguments of the function v in Eq. (1) equally, denoting them in that case by $\mathbf{p} := (\mathbf{x}, t, \boldsymbol{\mu}) =: (p_1, \dots, p_d)$. We consider a real-valued function $\mathbf{p} \mapsto w(\mathbf{p}) = w(p_1, \dots, p_d)$ and its samples

$$\mathbf{w} = (w(\mathbf{p}_m)) = (w(p_{m_1}, \dots, p_{m_d})) = (w_m) = (w_{m_1 \dots m_d}) \in \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}, \quad (2)$$

which form a tensor of order or degree d , representing a total of $N := \prod_{\ell=1}^d M_\ell$ values, and where for the sake of convenience and compactness of notation multi-indices $\mathbf{m} \in \mathcal{M} := \times_{\ell=1}^d \{1, \dots, M_\ell\} \subset \mathbb{N}^d$ have been introduced.

As one may see, a tensor can be simply defined as a high-order matrix or multi-index array, where multi-indices are used instead of indices. As an example, assume that we have $N = 10^{16}$ data points. Now a vector $\mathbf{w} \in \mathbb{R}^N$ is just a tensor of first order, but one may *reshape* the data, say into a matrix $\mathbf{W} \in \mathbb{R}^{10^8 \times 10^8} \cong \mathbb{R}^{10^8} \otimes \mathbb{R}^{10^8}$ — a tensor of 2nd order. Reshaping further, one may take $\mathbb{R}^{10^4 \times 10^4 \times 10^4 \times 10^4} \cong \bigotimes_{k=1}^4 \mathbb{R}^{10^4}$ — a tensor of 4th order — as well as other combinations, such as $\mathbf{w} \in \bigotimes_{k=1}^{16} \mathbb{R}^{10}$ — an array of size $10 \times \dots \times 10$ (16 times) — which is a tensor of 16th order; and finally all the way to $\mathbf{w} \in \bigotimes_{k=1}^{16} (\mathbb{R}^2 \otimes \mathbb{R}^5)$ — a tensor of order 32.

The tensors obtained in this way contain not only rows and columns, but also *slices* and *fibres* [57, 58, 13, 47]. These slices and fibres can be analysed for linear dependencies, super symmetry, or sparsity, and may result in a strong data compression. To have a first glimpse of possible compression techniques, assume that in the above example the data has been stored in the matrix $\mathbf{W} \in \mathbb{R}^{M \times M}$, where $M = 10^8 = \sqrt{N}$, and consider its singular value decomposition (SVD) $\mathbf{W} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^\top = \sum_{m=1}^M \varsigma_m \mathbf{u}_m \mathbf{v}_m^\top$, where $\boldsymbol{\Sigma} = \text{diag}(\varsigma_1, \dots, \varsigma_M)$ is the diagonal matrix of singular values $\varsigma_m \geq 0$, assumed arranged by decreasing value, and $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_M]$, $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_M]$ collect the right and left singular vectors. Often there is a number $r \ll M$, such that for some small $\varepsilon > 0$ one has $\varsigma_m \leq \varepsilon$ for all $m > r$. Then one can formulate a compressed or truncated version

$$\mathbf{W} \approx \mathbf{W}_r = \mathbf{U}_r \boldsymbol{\Sigma}_r \mathbf{V}_r^\top = \sum_{m=1}^r \varsigma_m \mathbf{u}_m \mathbf{v}_m^\top,$$

where $\boldsymbol{\Sigma}_r = \text{diag}(\varsigma_1, \dots, \varsigma_r)$, $\mathbf{U}_r = [\mathbf{u}_1, \dots, \mathbf{u}_r]$, $\mathbf{V}_r = [\mathbf{v}_1, \dots, \mathbf{v}_r]$, and $\mathbf{w}_m^{(1)} = \sqrt{\varsigma_m} \mathbf{u}_m$, $\mathbf{w}_m^{(2)} = \sqrt{\varsigma_m} \mathbf{v}_m$. This reduced SVD is a special case of Eq. (3). For the sake of a concrete example, assume that $r = 100$. One may observe that the required storage has been reduced from $N = M^2 = 10^{16}$ to $2 \times r \times M = 2 \times 10^{10}$ for the $2r$ vectors $\{\mathbf{w}_m^{(1)}, \mathbf{w}_m^{(2)}\}_{m=1}^r$.

The set of possible objects we want to work with will be denoted by \mathcal{T} . We assume that this set is equipped with the operations of an associative and commutative real algebra, and carries an inner product. It is a well known result [84] that such algebras — modulo some technical details — are isomorphic via the Gel’fand representation to a function algebra — to be more precise continuous real valued functions on a compact set. Under point-wise addition and multiplication by scalars, such functions clearly form a vector space, and if one includes point-wise multiplication of functions, they form an associative and commutative real algebra. In our case this is simply the function algebra $(\mathcal{M} = \times_{\ell=1}^d \{1, \dots, M_\ell\} \rightarrow \mathbb{R})$, which is obviously the same as $\bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$. The advantage of this abstract algebraic formulation is not only that it applies to any kind of data representation on which one may define these algebraic operations, but that one may use algorithms [48] which have been developed for the algebra of real $N \times N$ matrices

$\mathfrak{gl}(\mathbb{R}, N) = \mathbb{R}^{N \times N} \cong \mathbb{R}^N \otimes \mathbb{R}^N$. It may be noted here that this abstract formulation applies also to continuous and adaptive tensor decompositions of functions, which employ similar ideas and principles [73, 8, 36] to those used in the discrete case. This is thus a computational approach which gives an operational functional calculus, and spans the range from algebras of self-adjoint operators and their spectral calculus via function algebras [84] all the way to the discrete setting of high-dimensional tensors and their low-rank approximations as considered here in the numerical examples in Section 4, and described in Section 3 and in Appendix A.

In our concrete examples we work directly with the data represented as tensors $\mathcal{T} := \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$. Even without the identification with $(\mathcal{M} \rightarrow \mathbb{R})$, it is obvious that this is naturally a vector space. As also $\mathcal{T} \cong \mathbb{R}^N$, it is clear that \mathcal{T} can be equipped with the canonical Euclidean inner product from \mathbb{R}^N . The associative and commutative algebra product comes from the identification with the function algebra $(\mathcal{M} \rightarrow \mathbb{R})$, i.e. the point-wise product, which is also known as the *Hadamard* product [47]. Hence it is clear that for data $\mathbf{w} \in \mathbb{R}^N$ it is not difficult at all to define the algebraic operations, but rather how to perform them when the data is in a compressed format.

To simplify later notation on the number of operations and amount of storage needed, we will often make the assumption that $M_1 = \dots = M_d = n$, so that the tensor in Eq. (2) represents $N = n^d$ values. As already mentioned, our example of data compression is based on low-rank approximations to elements in \mathcal{T} . We would like to stress that the algorithms presented here are formulated purely in terms of the abstract algebraic structure, and are thus independent of the particular representation, although the only compressed formats considered here explicitly are low-rank tensor approximations.

Whereas a general element $\mathbf{w} \in \mathcal{T}$ hence has $N = n^d$ terms, a compressed representation — some low-rank versions of which will be discussed in more detail in Section 3 — will have significantly fewer terms. For example, the *CP*-decomposition (*canonical polyadic*, [49, 58]) representation, truncated to r terms

$$\mathbf{w} \approx \mathbf{w}_r = \sum_{i=1}^r \bigotimes_{k=1}^d \mathbf{w}_i^{(k)}; \quad \mathbf{w}_i^{(k)} \in \mathbb{R}^n; \quad (3)$$

has $r \times d \times n$ entries in \mathbf{w}_r , since $\mathbf{w}_i^{(k)}$ is a n -vector. If the *rank* r is reasonably small compared to $N = n^d$ independent of n and d , then we have an approximation $\mathbf{w}_r \in \mathcal{T}$ with much less storage, which also only depends on the *product* of n and d and is not exponential in the dimension d . But now the maximum can not be easily looked up; to get a particular element, the expression Eq. (3) has to be evaluated for that index combination. In the following we shall assume that we work with approximations such as in Eq. (3) which need much less storage. One has to make sure then that the *algebraic* operations of the Hadamard algebra — the vector space operations on \mathcal{T} together with the Hadamard element- or point-wise product — can be performed efficiently in an approximative manner, so that they have much lower complexity than the obvious $\mathcal{O}(n^d)$, although the single elements of \mathbf{w} which are usually needed for point-wise operations — and post-processing, see Subsection 1.2 — are not directly available. Thus the motivating factors for applying compression, and in particular low-rank tensor techniques, include the following:

- The storage cost is reduced, depending on the tensor format, from $\mathcal{O}(n^d)$ to $\mathcal{O}(drn)$, where $d > 1$.
- The low-rank tensor approximation is relatively new, but already a well-studied technique with free software libraries available. Other data compression techniques are available.
- The approximation accuracy is fully controlled by the tensor rank. The full rank gives an exact representation.
- Even more complicated operations like the Fourier transform can be performed efficiently in low-rank format. The basic fast Fourier transform on $\mathbf{w} \in \mathcal{T}$ would have complexity $\mathcal{O}(n^d \log(n^d))$. These low-rank techniques can either be combined with the fast Fourier transform giving a complexity of $\mathcal{O}(drn \log n)$, or can even be further accelerated at the price of an additional approximation error yielding a *superfast* Fourier transform [68, 22].

On the other hand, general limitations of a low-rank tensor compression technique are that

- it could be time consuming to compute a compression, or in particular a low-rank tensor decomposition;

- if a function approximation is sought, it requires an axes-parallel mesh;
- although many functions have a low-rank representation, in practice only some theoretical estimates exist.

But the fact still remains that there are situations where storage of all items is not feasible, and some kind of compression has to be employed.

1.2. Post-processing tasks

This work is about exploiting the compression together with the structure of a commutative algebra and the ability to perform the algebraic operations in the compressed format, at least approximately. The tensor product structure, with the help of expressions such as Eq. (3) or other compression techniques like the tensor train in AppendixA, is one example of a compression technique which allows one to carry out these algebra operations in an efficient manner [47], and these are the only compression techniques explicitly considered here.

Such a tensor product structure is typical for parametric problems [67], which were alluded to in the description in Subsection 1.1 preceding Eq. (1), where the function $v(\mathbf{x}, t, \boldsymbol{\mu})$ is the state of some system, controlled typically by a parametric partial differential equation, which is numerically costly to solve. What is often desired is a representation which allows for the approximate evaluation of the state Eq. (1) as function of $\boldsymbol{\mu} \in \mathcal{M}$, but avoids the numerically costly solution of the partial differential equation. Furthermore, one would like this representation to be inexpensive to evaluate, and for it to be convenient for certain post-processing tasks, for example like finding the minimum or maximum value over some or all parameter values.

The tasks we consider here are

1. the indices or values of *maxima*, *minima* of \mathbf{w} ,
2. the indices or values of *maxima*, *minima* of some function f of \mathbf{w} ,
3. the index or value of the datum of \mathbf{w} *closest* to a given number,
4. level sets, i.e. (the indices of) all $w_{m_1 \dots m_d} \in [\omega_0, \omega_1]$,
5. the *number of indices* such that $w_{m_1 \dots m_d} \in [\omega_0, \omega_1]$,
6. the mean (sum of all items) and variance (sum of all items squared),
7. some auxiliary functions, such as the algebraic inverse $\mathbf{w}^{\odot -1}$ of \mathbf{w} — in our case the Hadamard inverse — and others like the $\text{sign}(\mathbf{w})$ function which will be seen to be needed for the above tasks.

As an example, consider finding the maximum value. A naïve approach to compute the maximum would be to visit and inspect each element, but then the number of visits is $\mathcal{O}(n^d)$, exponential in d , which may be unacceptable for high d . Such phenomena when the total complexity/storage cost depends exponentially on the problem dimension have been called the curse of dimensionality [47].

The idea for such iterative post-processing algorithms, in the form of finding the maximum, was first presented in [24] for low-rank tensor approximations. Some further post-processing tasks for such low-rank tensor representations were investigated in [28]. To give an example from [24] of the possible savings in space and time, assume that one looks at the solution of a d -dimensional Poisson equation:

$$-\nabla^2 u = f \quad \text{on } \mathcal{G} = [0, 1]^d \quad \text{with } u|_{\partial \mathcal{G}} = 1,$$

and right-hand-side

$$f(x_1, \dots, x_d) \propto \sum_{k=1}^d \prod_{\ell=1, \ell \neq k}^d x_\ell (1 - x_\ell).$$

Assume further that this is solved numerically by a standard finite-difference method with $n = 100$ grid-points in each direction, so that the solution vector \mathbf{u} has $N = n^d$ data points in \mathbb{R}^N , but can also naturally be viewed as a tensor \mathbf{u} of degree d in $\bigotimes_{k=1}^d \mathbb{R}^{100}$. So for full storage, one needs $N = n^d$ storage locations, and if one is looking for the maximum by inspecting each element, one would have to inspect N elements.

Table 1: Computing times (4th column) on 2 GHz dual-core CPU to find the maximum.

d	# loc's.: $N = n^d$	\approx years [a] inspect. N	actual [24] time [s] of Algorithm 2
25	10^{50}	1.6×10^{33}	0.16
50	10^{100}	1.6×10^{83}	0.42
75	10^{150}	1.6×10^{133}	1.16
100	10^{200}	1.6×10^{183}	2.58
125	10^{250}	1.6×10^{233}	4.97
150	10^{300}	1.6×10^{283}	8.56

We mention in passing that of course actually solving the discrete equation is a challenge for $d > 3$, but that is a different story, and for this we refer to [24] and the references therein.

Such problem sizes like in Table 1 may appear in the computation of the electron density and the Hartree potential of small molecules, see [10, 24]. Now assume for the sake of simplicity that one can inspect 2×10^9 elements per second — on an ideal 2 GHz CPU with one inspection per cycle. Then for $\mathbf{u} \in (\mathbb{R}^n)^{\otimes d} \cong \mathbb{R}^{n^d} = \mathbb{R}^N$ the times needed to find the maximum for the full data-set are shown in Table 1 in the third column—assuming that it were somehow possible to store all the values indicated in the second column—whereas the actual computation with a compressed format is shown in the last and fourth column.

It is obvious that for growing d for the full representation the computational complexity and the storage requirements quickly become not only unacceptable, but totally impossible to satisfy. The second and third column behave like $\mathcal{O}(n^d)$ and grow exponentially with d , whereas for a low-rank representation \mathbf{u}_r of \mathbf{u} not only can the data be stored on a modest laptop, but for the simple Algorithm 2 — to be explained later in Subsection 2.4 — the computing times in the fourth column are in terms of seconds and behave like $\mathcal{O}(n d^3)$.

1.3. State of the art

As already stated, the algorithms considered here use only the structure of the Hadamard algebra and are hence applicable to any compression format which approximately supports those algebraic operations. The authors are presently only aware of low-rank tensor representations which satisfy this requirement. Therefore in this short review the focus is on algorithms which have been developed for various low-rank tensor formats. Naturally, for specific compression formats, algorithms which use more information about the format may be formulated, but these may in turn only be applicable to the specific format.

This is not a discussion of the state of the art regarding general tensor formats and their low-rank approximations, in quantum physics also known as *tensor networks*. For the physical motivations and numerical developments from there see [91, 81, 32, 69, 9, 7]. For the mathematical and numerical view we refer to the review [58], the monographs [47, 52, 53], and to the literature survey on low-rank approximations [40]. In the following, we concentrate on the question of post-processing such data.

The idea of finding the largest element and the corresponding location of a tensor by solving an eigenvalue problem, where the matrix and the vectors are tensors given in the CP tensor format, was introduced in [24]. Additionally, the first numerical schemes to compute the point-wise inverse and the sign function were introduced, as well as the rank-truncation procedure for tensors given in the CP format. Later, in [28], these ideas were extended and applied to tensors which were obtained after discretisation and solution of elliptic PDEs with uncertain coefficients. Another group of authors, in [17], by combining the advances of the density matrix renormalisation group and the variational numerical renormalisation group methods, approximated several low-lying eigenpairs of large Hermitian matrices simultaneously in the block version of the tensor train (TT) format [72] via the alternating minimisation of the block Rayleigh quotient sequentially for all TT cores. A new tensor conjugate-gradient-type method for the Rayleigh quotient minimisation in the block QTT-format was developed in [61]. In [19, 18], the authors suggested methods to compute the mean, the variance, and sensitivity indices in the TT-format with applications to stochastic PDEs, whereas the diagonalisation of large Toeplitz or circulant matrices via combination of the fast Fourier and the

CP tensor format was shown in [68]. An investigation of approximations to eigenfunctions of a certain class of elliptic operators in \mathbb{R}^d by finite sums of products of functions with separated variables is the topic of [44]. An alternative to eigenvalue computations may be the *maxvol* algorithm, which is claimed [35] to be a robust method to find extrema for practical tensor formats; it can be applied to matrices and to tensors [88].

Next to solving eigenvalue problems, the other tasks considered here require a point-wise function of the data to be computed. General methods for such computations in a commutative algebra are — next to the algebraic iterative algorithms used here — series, in particular power series, or the use of Cauchy’s integral formula, see e.g. [47]. Such general algorithms are closely connected with spectral calculus of self-adjoint operators in a Hilbert space, see e.g. [84]. But here the only such point-wise function essentially needed in the sequel is the sign-function, which is discontinuous and thus power series or Cauchy’s integral formula can not be used directly. Luckily, for the sign-function there is a Newton-, or rather Newton-Schultz iteration, which is well known for matrices [48], and which is used here. Having the sign function computed, one can compute characteristic- and level-set-functions, and with it (high-dimensional) histograms, which are needed for approximating high-dimensional probability density functions, cumulative distribution functions, quantiles, and exceedance probabilities. As before in the case of finding the maximum element, also here there are particular algorithms which are tailored and only suited for certain tensor representations. Some positive experience [71, 25, 3] was gained in applying cross approximation algorithms for computing functions of low-rank tensors, especially those that lack a convenient Newton-Schultz iteration. The authors state that it is robust to use. Specifically, some of these methods are: the rank-adaptive DMRG (renormalization) cross method [82], the multivariate adaptive cross approximation (ACA) [5], the hierarchical Tucker (HT) cross methods [2, 4, 3], and the parallel cross algorithms in HT- [41] and TT-formats [21]. The fast multidimensional convolution in low-rank tensor formats via cross approximation was developed in [79]. In addition to cross methods, an important part of high-dimensional data science is tensor completion, for example, [86, 12, 39]. Specifically the Hadamard product of low-rank tensors can be accelerated by ALS/DMRG algorithms [70] and cross methods [33, 79].

1.4. Outline of the paper

In the following Section 2 the necessary material for abstract algebras is quickly reviewed. Then the algorithms and functions on the algebra \mathcal{T} used to compute the post-processing tasks outlined in Subsection 1.2 are formulated in an abstract fashion, independent of the representation chosen for the data. Even as the formulation uses only the abstract algebra operations, we point out and motivate what this means in terms of the Hadamard algebra. But also for the Hadamard algebra, the presentation is independent of the tensor format to be chosen. Furthermore, for the iterative algorithm — a fixed point iteration — it is discussed how the possible truncation operation influences the convergence.

In Section 3 we present some concrete examples of compression of high-dimensional data $\mathbf{w} \in \mathbb{R}^N$, once it is identified with a tensor $\mathbf{w} \in \mathcal{T} := \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$ with $N = \prod_{\ell=1}^d M_\ell$. For such tensors of degree d we discuss the *canonical polyadic* (CP) representation, the *tensor-train* (TT) representation (in appendix), and the Tucker tensor formats [29]. In particular, we show how the algebra operations can be carried out in the different tensor formats, the numerical effort involved, and the effect these algebraic operations have on the compression or low-rank representation. As the compression level may deteriorate, i.e. the rank of the approximation may grow, it is important that one is able to re-compress. Pointers to such methods in the literature are included as well.

The Section 4 contains a few numerical examples and a discussion of their results used to illustrate the algorithms of Section 2 and the representations from Section 3. These examples are solutions of elliptic high-dimensional or parametric resp. stochastic partial differential equations, a domain where the data is naturally in a tensor format. But as already pointed out, for any data it is just a — often only conceptual — *reshape* to consider it as an element of a tensor space. The conclusion is contained in Section 5.

2. Algorithms for post-processing

Given a data-set \mathbf{w} in some compressed format, where looking up each element is not trivial, and where additionally it may not be computationally feasible to look at each element, we still want to be able to perform tasks which essentially require looking at each datum. The *tasks* we consider are outlined in Subsection 1.2. These tasks and auxiliary functions will be computed by finding an *iteration* (mapping) Φ_P for each post-processing task P , such that the *fixed-point* \mathbf{v}_* of Φ_P — i.e. $\Phi_P(\mathbf{v}_*) = \mathbf{v}_*$ — is either the sought solution for the task, or computes one of the auxiliary functions.

2.1. Iteration with Truncation

When performing computations using the operations of the algebra — like the Hadamard algebraic operations on tensors in some compressed format — after the operation the compression will typically be sub-optimal, which means that the result has to be compressed again. Thus, when performing the algebraic operations for a fixed-point iteration [59], the compression would either get worse and worse in each iteration, or one has to re-compress or *truncate* the result again to a good compression level. Therefore a re-compression after each or after a number of algebraic operations may be necessary, and we thus allow that the algebraic operations are possibly only executed approximately [61].

In our example case the low-rank representation of tensors explained in Section 3 acts as compression, and the rank may increase through the algebraic operations, and hence we will use *truncation* T_ϵ to *low rank* r with error ϵ [47, 6, 61] to avoid the problem of a deteriorating compression level.

In other words, with a general compressed representation \mathbf{w}_r of \mathbf{w} the computation will be a *truncated* or *perturbed iteration* [45, 61, 66]. If we denote the general compression mapping by T_ϵ — meaning compression with an accuracy ϵ — the iteration map is changed from Φ_P to $T_\epsilon \circ \Phi_P$.

The general structure of the iterative algorithms for a post-processing task P or for an auxiliary function is shown in Algorithm 1. Here we want to collect some results for this kind of iteration. For such a truncated or perturbed

Algorithm 1 Iteration with truncation

```

1: Start with some initial compressed guess  $\mathbf{v}_0$  depending on task  $P$ .
2:  $i \leftarrow 0$ 
3: while no convergence do
4:    $\mathbf{z}_i \leftarrow \Phi_P(\mathbf{v}_i)$ ; ▷ See Remark below
▷ the iterator  $\Phi_P$  may deteriorate the compression level
5:   if compression level of  $\mathbf{z}_i$  is too bad then
6:      $\mathbf{v}_{i+1} \leftarrow T_\epsilon(\mathbf{z}_i)$ ; ▷ use truncation  $T_\epsilon$  to compress  $\mathbf{z}_i$  with error  $\epsilon$ 
7:   else
8:      $\mathbf{v}_{i+1} \leftarrow \mathbf{z}_i$ ;
9:   end if
10:   $i \leftarrow i + 1$ 
11: end while

```

iteration as in Algorithm 1, it is known that

1. if the iteration by Φ_P is *super-linearly* convergent, the *truncated iteration* $T_\epsilon \circ \Phi_P$ will still *converge* super-linearly, but finally *stagnate* in an ϵ -neighbourhood of the fixed point \mathbf{v}_* [45, 59]. One could loosely say that the super-linear convergence of Φ_P is stronger than the truncation by T_ϵ .
2. if the iteration by Φ_P is *linearly* convergent with *contraction* factor q , the *truncated iteration* $T_\epsilon \circ \Phi_P$ will still *converge* linearly, but finally *stagnate* in an $\epsilon/(1-q)$ -neighbourhood of \mathbf{v}_* [66]. Again, one could loosely say that iteration by Φ_P and truncation by T_ϵ balance each other, thus resulting in a larger neighbourhood of stagnation.

We shall assume that the truncation level has thus been chosen according to the desired re-construction accuracy and taking into account the possible influence due to the convergence behaviour of the iterator Φ_P .

Remark 2.1. In Algorithm 1 we assume that the tensor rank of \mathbf{z}_i is not increasing too much after applying the mapping Φ_P to the tensor \mathbf{v}_i . This could be a strong assumption. Additionally, we assume that there is a mapping T_ϵ , which is able to truncate the rank of \mathbf{z}_i . In some cases the ranks of intermediate iterations can be estimated [1].

2.2. Preliminaries and basic algebraic operations

As already pointed out, we assume in general that the set \mathcal{T} is a vector space. Our example of the space $\mathcal{T} = \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$ of tensors of interest was already introduced in Subsection 1.1. This is clearly a vector space, so we are able to add two such tensors, and multiply each by some real number, in other words we may form linear combinations. The additive neutral element — the zero tensor — will be denoted as \mathbf{o} . It is important that the compressed storage format chosen allows to perform the vector space operations without going into the *full* representation. This will be shown for some of the familiar tensor formats in Section 3. Furthermore, as $\mathcal{T} := \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell} \cong \mathbb{R}^{M_1 \times \dots \times M_d} \cong \mathbb{R}^N$ as vector spaces, we can carry the canonical Euclidean inner product on \mathbb{R}^N to \mathcal{T} , which for $\mathbf{u}, \mathbf{v} \in \mathcal{T}$ is denoted by

$$\langle \mathbf{u} | \mathbf{v} \rangle_{\mathcal{T}} := \sum_{m_1=1, \dots, m_d=1}^{M_1, \dots, M_d} u_{m_1, \dots, m_d} \cdot v_{m_1, \dots, m_d} = \sum_{\mathbf{m} \in \mathcal{M}} u_{\mathbf{m}} \cdot v_{\mathbf{m}}.$$

This makes \mathcal{T} into a Euclidean or Hilbert space. We assume that the computation of this inner product is feasible when both \mathbf{u} and \mathbf{v} are in a compressed representation; again this will be demonstrated in Section 3 for the low-rank representations we shall consider as examples.

On this space \mathcal{T} an additional structure is needed, namely a *multiplication* which will make it into a unital associative and commutative algebra. In our example this will be the *Hadamard* multiplication, which apparently goes back to *Schur*, which is the point- or component-wise multiplication on $\bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell} \cong \mathbb{R}^{\times_{\ell=1}^d M_\ell}$. It is the usual point-wise multiplication of functions, in this case the functions $\times_{\ell=1}^d M_\ell \rightarrow \mathbb{R}$. For $\mathbf{u}, \mathbf{v} \in \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$, define the Hadamard product as

$$\odot : \mathcal{T} \times \mathcal{T} \rightarrow \mathcal{T}, \quad \mathbf{u} \odot \mathbf{v} \mapsto \mathbf{w} = (w_{m_1 \dots m_d}) := (u_{m_1 \dots m_d} \cdot v_{m_1 \dots m_d}). \quad (4)$$

It is a bi-linear operation (linear in each entry), and it is commutative. Thus this product makes \mathcal{T} into an *associative* and *commutative algebra*. The symbol \odot will be used both for the abstract algebra product on \mathcal{T} , as well as for the Hadamard product on $\bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$. As is any such algebra, for any $\mathbf{w} \in \mathcal{T}$ it holds that $\mathbf{w} \odot \mathbf{o} = \mathbf{o}$.

Especially, as a function of \mathbf{v} , the product $\mathbf{u} \odot \mathbf{v}$ is a linear map $\mathbf{L}_u \in \mathcal{L}(\mathcal{T})$ on \mathcal{T} . This is the familiar canonical representation $\mathcal{T} \ni \mathbf{v} \mapsto \mathbf{L}_v \in \mathcal{L}(\mathcal{T})$ of an associative algebra as linear maps on itself, i.e. in this case in a commutative sub-algebra of the algebra of all linear maps $\mathcal{L}(\mathcal{T})$ with concatenation as product.

It easy to see that this Hadamard algebra has a unit element $\mathbf{1} \in \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$ for the product,

$$\mathbf{1} := (1_{m_1 \dots m_d}), \quad \forall \mathbf{w} \in \mathcal{T} : \mathbf{w} = \mathbf{1} \odot \mathbf{w} = \mathbf{w} \odot \mathbf{1}, \quad (5)$$

a tensor with all entries equal to unity — this makes $\bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$ into a *unital* algebra. Observe that by defining for all $1 \leq k \leq d$ the *all-ones* vectors $\mathbf{1}_{M_k} = [1, \dots, 1]^T \in \mathbb{R}^{M_k}$, the Hadamard unit has rank *one* according to Eq. (3): $\mathbf{1} = \bigotimes_{k=1}^d \mathbf{1}_{M_k}$. Obviously one has $\mathbf{L}_1 = I_{\mathcal{T}}$, the identity in $\mathcal{L}(\mathcal{T})$.

Having defined a unit or neutral element for multiplication, we say that $\mathbf{w} \in \mathcal{T}$ has a *multiplicative inverse* iff there is an element, denoted by $\mathbf{w}^{\odot-1} \in \mathcal{T}$, such that

$$\mathbf{1} = \mathbf{w}^{\odot-1} \odot \mathbf{w} = \mathbf{w} \odot \mathbf{w}^{\odot-1}. \quad (6)$$

Obviously not all elements have an inverse, e.g. the zero element $\mathbf{o} \in \mathcal{T}$ is never invertible. In any case, not every $\mathbf{w} \in \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$ has a Hadamard inverse, for this it is necessary that all entries $w_{m_1 \dots m_d} \neq 0$ are non-zero, and then $\mathbf{w}^{\odot-1} = (w_{m_1 \dots m_d}^{-1}) = ((1/w)_{m_1 \dots m_d})$. Note that $(\mathbf{w}^{\odot-1})^{\odot-1} = \mathbf{w}$ and $\mathbf{1}^{\odot-1} = \mathbf{1}$ as it should be, and it is easily seen that $\mathbf{L}_{\mathbf{w}^{\odot-1}} = \mathbf{L}_{\mathbf{w}}^{-1}$, the inverse of $\mathbf{L}_{\mathbf{w}}$ in $\mathcal{L}(\mathcal{T})$.

Elements of the algebra which can be written as a square $\mathbf{w} = \mathbf{u} \odot \mathbf{u} = \mathbf{u}^{\odot 2}$ are called positive, they form a convex cone $\mathcal{T}_+ \subset \mathcal{T}$; note that obviously the zero tensor \mathbf{o} and the multiplicative unit $\mathbf{1}$ are positive, and that if an invertible \mathbf{w} is positive, so is its inverse $\mathbf{w}^{\odot -1}$. As usual, this defines an order relation $\mathbf{u} \leq \mathbf{v} \Leftrightarrow \mathbf{v} - \mathbf{u} \in \mathcal{T}_+$, which implies that for any positive $\mathbf{w} \in \mathcal{T}_+$ one has $\mathbf{o} \leq \mathbf{w}$. Further one may observe that $\mathbf{u}, \mathbf{v} \in \mathcal{T}_+ \Rightarrow \mathbf{u} \odot \mathbf{v} \in \mathcal{T}_+$.

A certain compatibility of the inner product and the algebra product is needed, as we require that

$$\langle \mathbf{w} \odot \mathbf{u} | \mathbf{v} \rangle_{\mathcal{T}} = \langle \mathbf{L}_{\mathbf{w}} \mathbf{u} | \mathbf{v} \rangle_{\mathcal{T}} = \langle \mathbf{u} | \mathbf{L}_{\mathbf{w}} \mathbf{v} \rangle_{\mathcal{T}} = \langle \mathbf{u} | \mathbf{w} \odot \mathbf{v} \rangle_{\mathcal{T}}, \quad (7)$$

i.e. that the action of the algebra product is *self-adjoint*, or, in other words, that the maps $\mathbf{L}_{\mathbf{w}}$ are self adjoint. This condition is satisfied for the Hadamard algebra.

The inner product with the unit element $\phi(\mathbf{w}) := \langle \mathbf{w} | \mathbf{1} \rangle_{\mathcal{T}}$ defines, as function of \mathbf{w} , a positive linear functional ϕ , as for $\mathbf{w} \in \mathcal{T}_+$ one has

$$\phi(\mathbf{w}) = \phi(\mathbf{u} \odot \mathbf{u}) = \langle \mathbf{u} \odot \mathbf{u} | \mathbf{1} \rangle_{\mathcal{T}} = \langle \mathbf{u} | \mathbf{u} \rangle_{\mathcal{T}} \geq 0.$$

It is a kind of “trace” or un-normalised state functional on the algebra, as $\phi(\mathbf{1}) = N = \dim \mathcal{T}$ in the Hadamard algebra. In particular

$$\phi(\mathbf{u} \odot \mathbf{v}) = \langle \mathbf{u} \odot \mathbf{v} | \mathbf{1} \rangle_{\mathcal{T}} = \langle \mathbf{u} | \mathbf{v} \rangle_{\mathcal{T}}. \quad (8)$$

Conversely, if the algebra comes equipped with such an un-normalised positive state functional, the Eq. (8) may be taken as the definition of an inner product, which automatically satisfies Eq. (7).

We shall assume that we can compute the algebraic operations — at least approximately — in compressed representation (this will be shown for the Hadamard algebra for the low-rank formats in Section 3), and that we can compute the multiplicative inverse also in compressed representation — again at least approximately.

As in any unital algebra, if for $\lambda \in \mathbb{C}$ the element $\mathbf{w} - \lambda \mathbf{1}$ fails to be invertible — this means that $\mathbf{L}_{\mathbf{w}} - \lambda \mathbf{I}_{\mathcal{T}}$ is not invertible in $\mathcal{L}(\mathcal{T})$ — we shall say that λ is in the *spectrum* of \mathbf{w} ; here — as we are in finite dimensional spaces — it is an *eigenvalue*, i.e. there is an *eigenvector* $\mathbf{v}_{\lambda} \in \mathcal{T}$ such that $\mathbf{L}_{\mathbf{w}} \mathbf{v}_{\lambda} = \mathbf{w} \odot \mathbf{v}_{\lambda} = \lambda \mathbf{v}_{\lambda} = \lambda \mathbf{I}_{\mathcal{T}} \mathbf{v}_{\lambda}$. From this discussion it is immediate that in the Hadamard algebra case, if λ is an eigenvalue of $\mathbf{w} = (w_{\mathbf{m}}) \in \mathcal{T}$ or $\mathbf{L}_{\mathbf{w}} \in \mathcal{L}(\mathcal{T})$, there must be a multi-index $\mathbf{m}_{\lambda} \in \mathcal{M}$ such that $w_{\mathbf{m}_{\lambda}} = \lambda$, and hence in the real Hadamard algebra \mathcal{T} one has that always $\lambda \in \mathbb{R}$. Thus the spectrum of \mathbf{w} resp. $\mathbf{L}_{\mathbf{w}}$ is $\sigma(\mathbf{w}) = \sigma(\mathbf{L}_{\mathbf{w}}) = \{w_{\mathbf{m}} \mid \mathbf{m} \in \mathcal{M}\}$. Observe that according to Eq. (7), each representing map $\mathbf{L}_{\mathbf{w}}$ from above is self-adjoint, which means that also in general all the spectra are real — $\sigma(\mathbf{w}) \subset \mathbb{R}$.

Specifically, the Euclidean Hadamard algebra $\mathcal{T} = \bigotimes_{\ell=1}^d \mathbb{R}^{M_{\ell}}$ which we have constructed is obviously isomorphic — as a Euclidean algebra — to \mathbb{R}^N with the canonical Euclidean inner product when equipped with the point- or element-wise Hadamard product \odot_N . Let us denote this isomorphism by $V : \mathcal{T} \rightarrow \mathbb{R}^N$; it implies some ordering of the terms of each $\mathbf{w} \in \mathcal{T}$.

More enlightening and less obvious may be the unital algebra isomorphism with the commutative sub-algebra of *diagonal* matrices $\mathfrak{diag}(\mathbb{R}, N)$ in the full matrix algebra $\mathfrak{gl}(\mathbb{R}, N) = \mathbb{R}^{N \times N} \cong \mathbb{R}^N \otimes \mathbb{R}^N \cong \mathcal{L}(\mathbb{R}^N)$ with the usual matrix multiplication. Let us denote this isomorphism by $M : \mathcal{T} \rightarrow \mathfrak{diag}(\mathbb{R}, N)$. If $\mathbf{w} = V(\mathbf{w}) \in \mathbb{R}^N$ is the vector containing all the elements of the tensor \mathbf{w} , then $M(\mathbf{w}) := \mathbf{W} = \text{diag}(\mathbf{w}) = \text{diag}(V(\mathbf{w})) \in \mathfrak{diag}(\mathbb{R}, N) \subset \mathfrak{gl}(\mathbb{R}, N)$ is the corresponding diagonal matrix, i.e. $M = \text{diag} \circ V$. Hence, by choosing the canonical Euclidean basis in \mathbb{R}^N and its image by V^{-1} as a basis in \mathcal{T} , the canonical representation $\mathbf{L}_{\mathbf{w}} \in \mathcal{L}(\mathcal{T})$ is itself represented by the matrix $\mathbf{W} = \text{diag}(\mathbf{w}) = M(\mathbf{w}) \in \mathfrak{diag}(\mathbb{R}, N) \subset \mathfrak{gl}(\mathbb{R}, N) \cong \mathcal{L}(\mathbb{R}^N)$.

Let $\mathbf{v} \in \mathcal{T}$ be another tensor, with associated $V(\mathbf{v}) = \mathbf{v} \in \mathbb{R}^N$ and $M(\mathbf{v}) = \mathbf{V} = \text{diag}(V(\mathbf{v})) = \text{diag}(\mathbf{v}) \in \mathfrak{diag}(\mathbb{R}, N)$, then from the definition of the isomorphism M :

$$\begin{aligned} M(\mathbf{w} \odot \mathbf{v}) &= M(\mathbf{w})M(\mathbf{v}) = \mathbf{W}\mathbf{V} = \text{diag}(V(\mathbf{w})) \text{diag}(V(\mathbf{v})) \\ &= \text{diag}(\mathbf{w}) \text{diag}(\mathbf{v}) = \text{diag}(M(\mathbf{w})V(\mathbf{v})) = \text{diag}(\mathbf{W}\mathbf{v}) = \text{diag}(\text{diag}(\mathbf{w})\mathbf{v}), \end{aligned} \quad (9)$$

which contains the important equality

$$\mathbf{w} \odot \mathbf{v} = M^{-1}(\mathbf{W}\mathbf{V}) = M^{-1}(\text{diag}(\mathbf{W}\mathbf{v})) = M^{-1}(\text{diag}(\mathbf{w}) \text{diag}(\mathbf{v})) = M^{-1}(\text{diag}(M(\mathbf{w})V(\mathbf{v}))) = V^{-1}(\mathbf{W}\mathbf{v}) = \mathbf{L}_w \mathbf{v}, \quad (10)$$

from where one sees that

$$V(\mathbf{w} \odot \mathbf{v}) = \mathbf{W}\mathbf{v} = \text{diag}(\mathbf{w})\mathbf{v} = M(\mathbf{w})V(\mathbf{v}), \quad (11)$$

which helps in understanding the following algorithms. Note that this shows that the algebra $\bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$ with the Hadamard product is already jointly diagonalised and is essentially in its *Gel'fand* representation [84], which is an abstract way of saying what was already stated above, namely that the spectrum $\sigma(\mathbf{w})$ of an element $\mathbf{w} \in \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$ are exactly the individual terms or data stored in the tensor. If in some abstract unital associative and commutative algebra the multiplication is *not* the point-wise multiplication — e.g. think of convolution — then it is known that modulo some technicalities [84] it is isomorphic via the Gel'fand “diagonalisation” morphism to a function algebra, and all the algorithms to follow would deal with the spectrum $\sigma(\mathbf{w})$ of \mathbf{w} instead of with the “values of \mathbf{w} ”. These two notions only coincide for a function algebra.

2.3. Post-processing Tasks

Here the different post-processing tasks are explained together with how they will be computed. This may involve a number of auxiliary functions. The computation of these — again through truncated iteration Algorithm 1 — and the way in which some of the computations can be enhanced or accelerated is shown in Subsection 2.4.

Finding the maximum or minimum of $\mathbf{w} \in \mathcal{T} = \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$ is the first task we consider. Observe that the element of maximum modulus of the tensor \mathbf{w} is also equal to the ∞ -norm $\|\mathbf{w}\|_\infty$.

Finding the maximum means finding the index $\hat{\mathbf{m}} = (\hat{m}_1, \dots, \hat{m}_d) \in \mathcal{M}$ where the maximum \hat{w} of the elements in \mathbf{w} occurs:

$$\hat{w} := w_{\hat{\mathbf{m}}} := w_{\hat{m}_1, \dots, \hat{m}_d} := \max \{ w_{\mathbf{m}} : \mathbf{m} = (m_1, \dots, m_d) \in \mathcal{M} \}. \quad (12)$$

It was already established that each value of \mathbf{w} is an eigenvalue. Defining for each $1 \leq m \leq M_\ell$ the canonical unit basis vectors $\mathbf{e}_{M_\ell}^{(m)}$ in each \mathbb{R}^{M_ℓ} in the tensor product $\mathcal{T} = \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$ as $\mathbf{e}_{M_\ell}^{(m)} := (\delta_n^{(m)})_{n=1, \dots, M_\ell} \in \mathbb{R}^{M_\ell}$ via the Kronecker- δ -symbol, and similarly for each $\mathbf{m} \in \mathcal{M}$ the canonical unit basis vectors $\mathbf{e}^{(\mathbf{m})} = (\delta_{\mathbf{n}}^{(\mathbf{m})})_{\mathbf{n} \in \mathcal{M}} \in \mathcal{T} = \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$, then for *any* $\mathbf{m} = (m_1, \dots, m_d) \in \mathcal{M}$ the element $w_{\mathbf{m}}$ satisfies an *eigenvalue* equation:

$$\mathbf{L}_w \mathbf{e}^{(\mathbf{m})} := V^{-1} \left(M(\mathbf{w})V(\mathbf{e}^{(\mathbf{m})}) \right) = \mathbf{w} \odot \mathbf{e}^{(\mathbf{m})} = w_{\mathbf{m}} \mathbf{e}^{(\mathbf{m})}, \quad (13)$$

and the eigenvectors $\mathbf{e}^{(\mathbf{m})} = \bigotimes_{\ell=1}^d \mathbf{e}_{M_\ell}^{(m_\ell)}$ are evidently of rank one. This shows that each datum $w_{\mathbf{m}}$ of \mathbf{w} may be found through *eigenvalue* computation. The eigenvalue is the data entry $w_{\mathbf{m}}$ with the corresponding eigenvector $\mathbf{e}^{(\mathbf{m})}$ indicating the location resp. the index $\mathbf{m} \in \mathcal{M}$.

Although quite obvious, it may be worthwhile pointing out that the linear span of any nonempty collection of eigen- resp. unit vectors $\mathcal{T}_{\mathcal{N}} := \text{span}\{\mathbf{e}^{(\mathbf{n})} \mid \mathbf{n} \in \mathcal{N}\}$ ($\mathcal{N} \subseteq \mathcal{M}$) is an invariant subspace, and the orthogonal projector on $\mathcal{T}_{\mathcal{N}}$ is $\mathbf{L}_{\mathbf{p}^{(\mathcal{N})}} := \sum_{\mathbf{n} \in \mathcal{N}} \mathbf{L}_{\mathbf{e}^{(\mathbf{n})}}$ with $\mathbf{p}^{(\mathcal{N})} = \sum_{\mathbf{n} \in \mathcal{N}} \mathbf{e}^{(\mathbf{n})}$. With this, the spectral resolution of the identity for any and all $\mathbf{L}_w \in \mathcal{L}(\mathcal{T})$ for $\mathbf{w} \in \mathcal{T}$ is $\mathbf{I}_{\mathcal{T}} = \sum_{\mathbf{m} \in \mathcal{M}} \mathbf{L}_{\mathbf{e}^{(\mathbf{m})}}$ corresponding to $\mathbf{1} = \sum_{\mathbf{m} \in \mathcal{M}} \mathbf{e}^{(\mathbf{m})}$, giving the spectral resolution

$$\mathbf{L}_w = \sum_{\mathbf{m} \in \mathcal{M}} w_{\mathbf{m}} \mathbf{L}_{\mathbf{e}^{(\mathbf{m})}},$$

$$\text{corresponding to } \mathbf{w} = \sum_{\mathbf{m} \in \mathcal{M}} w_{\mathbf{m}} \mathbf{e}^{(\mathbf{m})} = \sum_{\mathbf{m} \in \mathcal{M}} w_{\mathbf{m}} \bigotimes_{\ell=1}^d \mathbf{e}_{M_\ell}^{(m_\ell)}, \quad (14)$$

which may be seen as the trivial basis representation of \mathbf{w} .

As the basic algorithms for eigenvalue computation like power iteration for self-adjoint linear maps converge to the eigenvalue of maximum absolute value or modulus, it is now clear that if we assume that the maximum is also the element of maximum absolute value, we only have to perform power iteration. This gives us the value $\hat{w} = w_{\hat{m}}$, and the location through its corresponding eigenvector $\mathbf{e}^{(\hat{m})}$. How to find the minimum in this case, or what to do when the element of maximum modulus is the minimum will be explained later in Subsection 2.4; it will all be accomplished with spectral shifts familiar from eigenvalue computations.

Finding the maximum or minimum of some function f of \mathbf{w} , where it is assumed that $f : \sigma(\mathbf{w}) \rightarrow \mathbb{R}$, is not very difficult if one can compute $f(\mathbf{w})$, which will be simply understood as $f(\mathbf{w}) := (f(w_{\mathbf{m}}))_{\mathbf{m} \in \mathcal{M}}$. This definition of $f(\mathbf{w})$ is also the one which comes via the isomorphy with the algebra of diagonal matrices $\mathfrak{diag}(\mathbb{R}, N) \subset \mathfrak{gl}(\mathbb{R}, N)$, where $f(\mathbf{w})$ is defined [48] through the corresponding matrix function $f(M(\mathbf{w})) = f(\mathbf{W})$, and is also the same as the one which comes from the abstract functional calculus on the Banach algebra \mathcal{T} [84].

Observe that one may compute $f(\mathbf{w})$ for any real or complex function f defined on the spectrum $\sigma(\mathbf{w})$ — which is a finite set — at least in principle, through an interpolating polynomial $p(w) = \sum_k \alpha_k w^k$. As usual, the number $w \in \sigma(\mathbf{w})$ is replaced by \mathbf{w} (or $\mathbf{L}_{\mathbf{w}}$) in the polynomial, resulting in $f(\mathbf{w}) = p(\mathbf{w}) = \sum_k \alpha_k \mathbf{w}^{\odot k}$. This may be highly inefficient, as in general the degree of the polynomial would be equal to the number of data, in this case the huge number $N - 1$. For the functions which will be needed for the post-processing tasks, the computation of the function will be done differently, namely iteratively via Algorithm 1.

Assume now that $f(\mathbf{w})$ has been computed, at least approximately. To find maxima or minima of $f(\mathbf{w})$, one now has to simply apply the considerations of the preceding paragraph to the tensor $f(\mathbf{w})$. Obviously, $f(\mathbf{w})$ has the same spectral resolution as \mathbf{w} in Eq. (14):

$$f(\mathbf{L}_{\mathbf{w}}) = \sum_{\mathbf{m} \in \mathcal{M}} f(w_{\mathbf{m}}) \mathbf{L}_{\mathbf{e}^{(\mathbf{m})}}, \quad \text{corresponding to}$$

$$f(\mathbf{w}) = \sum_{\mathbf{m} \in \mathcal{M}} f(w_{\mathbf{m}}) \mathbf{e}^{(\mathbf{m})} = \sum_{\mathbf{m} \in \mathcal{M}} f(w_{\mathbf{m}}) \bigotimes_{\ell=1}^d \mathbf{e}_{M_{\ell}}^{(m_{\ell})}. \quad (15)$$

Finding the index and value closest to a given number $\rho \in \mathbb{R}$ is now simply finding the eigenvector and eigenvalue of

$$(\mathbf{L}_{\mathbf{w}} - \rho \mathbf{I}_{\mathcal{T}})^{-1} = \mathbf{L}_{\mathbf{y}} \quad \text{with} \quad \mathbf{y} = (\mathbf{w} - \rho \mathbf{1})^{\odot -1}. \quad (16)$$

This is a special case of the preceding paragraph for the function $f : t \mapsto (t - \rho)^{-1}$. Therefore, if $\lambda_{\mathbf{w}}$ is any eigenvalue of $\mathbf{L}_{\mathbf{w}}$, then the corresponding eigenvalue of $\mathbf{L}_{\mathbf{y}}$ is the transformed one $\lambda_{\mathbf{y}} = (\lambda_{\mathbf{w}} - \rho)^{-1}$. Hence the value of \mathbf{w} closest to ρ is the eigenvalue of maximum modulus of $\mathbf{L}_{\mathbf{y}}$, and the element of largest magnitude of $(\mathbf{w} - \rho \mathbf{1})^{\odot -1}$. To find the element of smallest magnitude, or even the vanishing elements of \mathbf{w} , one would use $\rho = 0$. Thus one may use the same algorithms as in the maximum search above. Observe that this operation here requires the Hadamard inverse, which will be one of the auxiliary functions.

Finding the elements in a subset is achieved by computing the so-called *characteristic* function for a subset $S \subset \mathbb{R}$ — we shall only look at intervals — $\chi_S : \mathcal{T} \rightarrow \mathcal{T}$, defined component-wise:

$$\mathcal{T} \ni (\chi_S(\mathbf{w})_{m_1, \dots, m_d}) := \begin{cases} 1, & \text{if } w_{m_1, \dots, m_d} \in S; \\ 0, & \text{if } w_{m_1, \dots, m_d} \notin S. \end{cases} \quad (17)$$

The characteristic function for intervals will be computed in terms of the *sign* function, where $\text{sign} : \mathcal{T} \rightarrow \mathcal{T}$ is also defined component-wise as

$$\mathcal{T} \ni (\text{sign}(\mathbf{w})_{m_1, \dots, m_d}) := \begin{cases} 1, & \text{if } w_{m_1, \dots, m_d} > 0; \\ 0, & \text{if } w_{m_1, \dots, m_d} = 0; \\ -1, & \text{if } w_{m_1, \dots, m_d} < 0. \end{cases} \quad (18)$$

This is again a special case of the previous one with a general function f .

With these two auxiliary functions, it is possible to define the characteristic function of a level set, i.e. all values between $\omega_1, \omega_2 \in \mathbb{R}$. We then have with $-\infty < \omega_1 < \omega_2 < \infty$:

$$(\chi_S(\mathbf{w}))_{m_1, \dots, m_d} := \begin{cases} \frac{1}{2}(\mathbf{1} + \text{sign}(\omega_2 \mathbf{1} - \mathbf{w})), & \text{if } S =]-\infty, \omega_2[; \\ \frac{1}{2}(\mathbf{1} - \text{sign}(\omega_1 \mathbf{1} - \mathbf{w})), & \text{if } S =]\omega_1, +\infty[; \\ \frac{1}{2}(\text{sign}(\omega_2 \mathbf{1} - \mathbf{w}) - \text{sign}(\omega_1 \mathbf{1} - \mathbf{w})), & \text{if } S =]\omega_1, \omega_2[; \end{cases} \quad (19)$$

Each case is easily computed with the sign function from Eq. (18). Hence, the indices of all components $w_{m_1 \dots m_d} \in]\omega_1, \omega_2[$ are provided by $\chi_{] \omega_1, \omega_2 [}(\mathbf{w})$. One may additionally define the *level set* function $\mathcal{L}_S(\cdot)$ of a subset $S \subset \mathbb{R}$ as $\mathcal{L}_S(\mathbf{w}) := \chi_S(\mathbf{w}) \odot \mathbf{w}$, which together with the indices provides the values in the subset S .

Finding the number of indices in a level set is accomplished through consideration of the *support*

$$\text{supp } \chi_S := \{\mathbf{m} \in \mathcal{M} : \chi_S(\mathbf{w})_{\mathbf{m}} \neq 0\} \subset \mathcal{M} \quad (20)$$

of a characteristic function χ_S as the subset of those indices where it is *non-zero*. Its *cardinality* is

$$\#(\text{supp } \chi_S(\mathbf{w})) = |\text{supp } \chi_S(\mathbf{w})| = \langle \chi_S(\mathbf{w}) | \mathbf{1} \rangle_{\mathcal{T}}, \quad (21)$$

the number of non-zero positions in the characteristic function; requiring only the computation of one inner product with the Hadamard multiplicative unit.

Computing the probability of being in a level, as well as the mean and the variance is carried out under the assumption that each index in \mathbf{w} carries the same probability, as then the *probability* of a value of \mathbf{w} being in a subset S is simply

$$\mathbb{P}_{\mathbf{w}}(S) := \frac{\#(\text{supp } \chi_S(\mathbf{w}))}{N} = \frac{\langle \chi_S(\mathbf{w}) | \mathbf{1} \rangle_{\mathcal{T}}}{N}. \quad (22)$$

The *mean* or *average* and *variance* of \mathbf{w} is then

$$\mathbb{E}(\mathbf{w}) := \bar{w} := \frac{1}{N} \langle \mathbf{w} | \mathbf{1} \rangle_{\mathcal{T}}; \quad \text{var}(\mathbf{w}) = \frac{1}{N} \langle \tilde{\mathbf{w}} | \tilde{\mathbf{w}} \rangle_{\mathcal{T}}, \quad \text{where } \tilde{\mathbf{w}} := \mathbf{w} - \bar{w} \mathbf{1}. \quad (23)$$

Given the level set function from above, one may also compute the conditional mean, conditioned on being in the set S :

$$\mathbb{E}(\mathbf{w} | S) := \bar{w}_{|S} := \frac{\langle \mathcal{L}_S(\mathbf{w}) | \mathbf{1} \rangle_{\mathcal{T}}}{\langle \chi_S(\mathbf{w}) | \mathbf{1} \rangle_{\mathcal{T}}}. \quad (24)$$

2.4. Auxiliary functions and algorithmic details

Here we show how to compute the auxiliary functions used in the previous Subsection 2.3. These are functions defined on the Euclidean Hadamard algebra $\mathcal{T} = \bigotimes_{\ell=1}^d \mathbb{R}^{M_{\ell}}$, and it is worth while pointing out that the iterative algorithms will be operating on the algebra, which, as we saw in Subsection 2.2, is isomorphic to an algebra of diagonal matrices, $\text{diag}(\mathbb{R}, N) \subset \mathfrak{gl}(\mathbb{R}, N)$, and these matrices are real symmetric. This means in particular that all the known algorithms for computing matrix functions [48] can be used on the Hadamard algebra as well. Of those needed here, the most basic one and the one with the simplest connection to the algebra turns out to be the inverse $\mathbf{w}^{\odot -1}$. Let us remark once again that the algorithms are valid in any abstract algebra, but we are mainly concerned with the example of the “function algebra” $\bigotimes_{\ell=1}^d \mathbb{R}^{M_{\ell}}$.

The *Hadamard inverse* $\mathbf{w}^{\odot -1}$ of a Hadamard invertible $\mathbf{w} \in \mathcal{T}$ is needed for *inverse* iteration and other post-processing tasks. Although the Hadamard inverse is the application of the function $f : t \mapsto t^{-1}$ to \mathbf{w} and could in principle be computed through a polynomial, it is often advantageous and much more effective to use other algorithms. The

algorithm for the inverse can be given simply by referring to the quadratically convergent Newton algorithm for matrices [48] for computing inverses, i.e. we use *Newton's method* for $(P \leftarrow (\odot - 1))$ and apply it to the equation $F(\mathbf{v}) := \mathbf{v}^{\odot - 1} - \mathbf{w} = \mathbf{0}$ to solve for \mathbf{v} , the only solution of which is $\mathbf{v} = \mathbf{w}^{\odot - 1}$. Hence the iteration function to be used in Algorithm 1 with starting vector $\mathbf{v}_0 := \mathbf{w}$ is

$$\Phi_{(\odot - 1)}(\mathbf{v}) := \mathbf{v} \odot (2 \cdot \mathbf{1} - \mathbf{w} \odot \mathbf{v}). \quad (25)$$

It is well known that the iteration converges locally *quadratically* [48], and according to what was explained in Subsection 2.1 this is true even with *truncation*.

The *sign function* is defined in Eq. (18). Many of the tasks explained in the previous Subsection 2.3 involve the sign function. Given the sign function, the characteristic, support, and level set functions are easily computed by the basic operations of the algebra.

From the fact that each value in a tensor \mathbf{w} is an eigenvalue in the Hadamard algebra, and from the isomorphy with the algebra of diagonal matrices $\text{diag}(\mathbb{R}, N) \subset \mathfrak{gl}(\mathbb{R}, N)$ where each value on the diagonal is also obviously an eigenvalue, one sees that this is actually the same definition as the one for matrices [48], and is also the one which follows from general functional calculus in abstract algebras [84]. This means that $\text{sign}(\mathbf{w})$ is the application of the function $f : t \mapsto \text{sign}(t)$ to $\mathbf{w} \in \mathcal{T}$, and again could in principle be computed with a polynomial. To compute it via the iterative Algorithm 1, i.e. $(P \leftarrow \text{sign})$, one uses the same algorithm [48] — the Roberts-Newton algorithm — as for matrices. It can be derived by applying *Newton's method* to the equation $F(\mathbf{v}) := \mathbf{v} \odot \mathbf{v} - \mathbf{1} = \mathbf{0}$ with starting value $\mathbf{v}_0 := \mathbf{w}$. This yields the iteration function

$$\Phi_{\text{sign}}(\mathbf{v}) := \frac{1}{2} (\mathbf{v} + \mathbf{v}^{\odot - 1}). \quad (26)$$

Observe that this iteration function in a slightly more general form $\Phi_{(\sqrt{\mathbf{w}})}(\mathbf{v}) := \frac{1}{2} (\mathbf{v} + \mathbf{v}^{\odot - 1} \odot \mathbf{w})$ is the ancient *Babylonian* method [48] — in modern parlance Newton's method — to find the square root $\mathbf{w}^{1/2}$ of \mathbf{w} , as by inserting $\mathbf{w} = \mathbf{1}$ in the Babylonian iteration $\Phi_{(\sqrt{\mathbf{w}})}$, one obtains Eq. (26). This means that one is iterating to compute the square root of the unit element $\mathbf{1}$ with the specific starting value $\mathbf{v}_0 = \mathbf{w}$. We will come back to this point of view shortly.

It is known [48] that the iteration converges *quadratically*, and according to what was explained in Subsection 2.1, it does so even with *truncation*; observe that it needs the Hadamard inverse from the previous paragraph. Therefore it is not really practical as this would mean that it can only be applied to invertible $\mathbf{w} \in \mathcal{T} = \bigotimes_{\ell=1}^d \mathbb{R}^{M_\ell}$, but it also actually needs a new inverse $\mathbf{v}_i^{\odot - 1}$ in each iteration. This last fact would computationally result in a nested iteration — using the algorithm from the previous paragraph in an inner iteration to compute the inverse $\mathbf{v}_i^{\odot - 1}$ in each sweep of the outer iteration for the sign function — and such procedures are seldom computationally efficient.

It is thus simpler to use a device employed also for matrix sign computations [48], namely to replace the explicit inverse in Eq. (26) by one step of the iteration for the inverse in Eq. (25). The resulting Newton-Schulz algorithm [48] has the same starting point $\mathbf{v}_0 := \mathbf{w}$ as before, but uses the iteration function

$$\Phi_{\text{N-S}}(\mathbf{v}) := \frac{1}{2} \cdot \mathbf{v} \odot (3 \cdot \mathbf{1} - \mathbf{v}^{\odot 2}) \quad (27)$$

in the Algorithm 1. It is now a simple iteration which still converges *quadratically* [48], even with *truncation* according to the explanations in Subsection 2.1. We refer to the discussion in [48] on how to get even faster algorithms using Padé approximations and other ways to accelerate the convergence through scaling, which is especially important in the initial stages.

Remark 2.2. *As the convergence of Newton's method is quadratic only locally in the vicinity of the solution, that in the first iterations the sign of small elements may be flipped due to truncation errors, and the method would converge to a qualitatively wrong answer for these elements. A possible remedy would be to define the truncation tolerance*

below the minimum absolute tensor element, so that no sign flip occurs. Depending on the data \mathbf{w} , it may be that this possibility can be ignored as one may assume that the total number of such small elements is negligible.

Coming back to the somewhat curious idea of iterating for the square root of the unit $\mathbf{1}$, one may observe that although the defining equation $F(\mathbf{v}) := \mathbf{v} \odot \mathbf{v} - \mathbf{1} = \mathbf{0}$ does not cover the case that a datum with the exact value of zero occurs — a non-invertible element would not satisfy the defining equation — the Newton-Schulz iteration with the function in Eq. (27) with starting value $\mathbf{v}_0 := \mathbf{w}$ takes care of this. A value which vanishes in \mathbf{v}_i also vanishes in $\mathbf{v}_{i+1} = \Phi_{\text{N-S}}(\mathbf{v}_i)$, due to the product with \mathbf{v}_i in Eq. (27).

Stopping criteria in case of quadratic convergence in an iteration like in Algorithm 1 are well known. First, if one wants to solve $F(\mathbf{v}) = \mathbf{0}$ or a fixed point equation $F(\mathbf{v}) := \Phi(\mathbf{v}) - \mathbf{v} = \mathbf{0}$, a natural criterion is the size of the residuum at step i :

$$\|F(\mathbf{v}_i)\|_{\mathcal{T}} < \eta_F. \quad (28)$$

But this checks only how well the equation is satisfied, and not directly how accurate the iterate \mathbf{v}_i is. Regarding this latter issue, specifically for quadratic convergence, a natural criterion [48] to check for the accuracy of \mathbf{v}_i at step i is

$$\delta_i := \frac{\|\mathbf{v}_i - \mathbf{v}_{i-1}\|_{\mathcal{T}}}{\|\mathbf{v}_i\|_{\mathcal{T}}} < \eta_v. \quad (29)$$

Further, referring to the discussion specifically for the sign function in [48], there are arguments to check

$$\delta_i < \|\mathbf{v}_i\|_{\mathcal{T}}^p \eta_v \quad (30)$$

for the exponents $p = 0, 1, 2$. For $p = 0$, this is the original general criterion Eq. (29), whereas the other values of p take specific consideration of the sign function.

Eigenvalue computations are involved in the first three tasks described in Subsection 2.3. One may use any algorithm developed for large scale matrices [34, 78, 80, 92] which only uses the action of the matrix on a vector in the computation — this is the Hadamard product in our case — and we shall here only explain the main idea and a few variations.

The simplest algorithm is power iteration ($P \leftarrow \text{pow-it}$), and the iteration map $\Phi_{\text{pow-it}}$ is given in Algorithm 2, to be used in Algorithm 1. Assume that the datum $w_{\hat{\mathbf{m}}} > 0$ of \mathbf{w} with maximum absolute value or maximum modulus is indeed a (positive) maximum. Assume also that this occurs at one unique index $\hat{\mathbf{m}}$, i.e. the datum $\lambda_1 := w_{\hat{\mathbf{m}}}$ is a simple eigenvalue, and denote the next by absolute value smaller element / eigenvalue by λ_2 . Assume the further eigenvalues ordered by decreasing absolute value.

Algorithm 2 One step power iteration $\Phi_{\text{pow-it}}$

- | | |
|---|--|
| 1: Input iterate \mathbf{v}_i ; | ▷ assume \mathbf{v}_i of unit length. |
| 2: $\mathbf{u} \leftarrow \mathbf{w} \odot \mathbf{v}_i$; | ▷ $\mathbf{u} = \mathbf{L}_w \mathbf{v}_i$. |
| 3: $\gamma \leftarrow \langle \mathbf{u} \mathbf{u} \rangle_{\mathcal{T}}^{-1/2}$; | ▷ inverse length of \mathbf{u} . $\lambda_1 \approx \gamma^{-1}$ |
| 4: $\mathbf{z} \leftarrow \gamma \cdot \mathbf{u}$; | ▷ normalise output \mathbf{z} to unit length. |
| 5: Output \mathbf{z} ; | |
-

As is well known [34, 78], the generated sequence \mathbf{v}_i converges linearly to $\pm \mathbf{e}^{(\hat{\mathbf{m}})}$ with a contraction factor of $q = |\lambda_2/\lambda_1|$, and with what was said in the explanations in Subsection 2.1, it still converges with truncation and will stagnate in the vicinity of $\pm \mathbf{e}^{(\hat{\mathbf{m}})}$. Also, as in the algorithm the input \mathbf{v}_i has unit length, the length of \mathbf{u} — given by the inverse γ^{-1} of the scaling factor in line 3 — converges linearly to the desired $\lambda_1 = w_{\hat{\mathbf{m}}}$ with the same rate. Algorithmically, the eigenvalue approximation $\lambda_1 \approx \gamma^{-1}$ is a *side-effect* of Algorithm 2.

This can be immediately enhanced through the additional computation of the *Rayleigh quotient* (RQ)

$$\varrho_w(\mathbf{v}_i) := \frac{\langle \mathbf{L}_w \mathbf{v}_i | \mathbf{v}_i \rangle_{\mathcal{T}}}{\langle \mathbf{v}_i | \mathbf{v}_i \rangle_{\mathcal{T}}} = \frac{\langle \mathbf{w} \odot \mathbf{v}_i | \mathbf{v}_i \rangle_{\mathcal{T}}}{\langle \mathbf{v}_i | \mathbf{v}_i \rangle_{\mathcal{T}}} \quad (31)$$

between line 2 and line 3 of Algorithm 2: $\varrho_{\mathbf{w}}(\mathbf{v}_i) = \langle \mathbf{u} | \mathbf{v}_i \rangle_{\mathcal{T}}$ — no need to divide by the length of the unit vector \mathbf{v}_i . The Rayleigh quotient is stationary at an eigenvalue and has thus a quadratic convergence, it usually represents a much better approximation to $\lambda_1 = w_{\hat{\mathbf{m}}}$ than γ^{-1} ; see Algorithm 3.

Stopping criteria for eigenvalues could be done similarly to normal convergent processes, but it is possible to have a posteriori error estimates which are actual bounds specifically for eigenvalues and -vectors, e.g. see [62, 34, 78]. The simplest seems to be the so-called Krylov-Bogolyubov bound. With an approximate eigenvalue μ and approximate eigenvector \mathbf{x} of $\mathbf{L}_{\mathbf{w}}$ it is given by [62]

$$\min_{\lambda_j \in \sigma(\mathbf{w})} \frac{|\lambda_j - \mu|}{|\lambda_j|} \leq \frac{\|\mathbf{L}_{\mathbf{w}}\mathbf{x} - \mu\mathbf{x}\|_{\mathcal{T}}}{\|\mathbf{x}\|_{\mathcal{T}}}. \quad (32)$$

The right-hand side of Eq. (32) is minimised by $\mu = \varrho_{\mathbf{w}}(\mathbf{x})$. A short computation [62] shows that for the substitutions $\mathbf{x} \leftarrow \mathbf{v}_i$ and $\mu \leftarrow \varrho_{\mathbf{w}}(\mathbf{v}_i)$ the right-hand side of Eq. (32) becomes $(\varrho_{\mathbf{w}^{\odot 2}}(\mathbf{v}_i) - \varrho_{\mathbf{w}}(\mathbf{v}_i)^2)^{1/2}$. As

$$\varrho_{\mathbf{w}^{\odot 2}}(\mathbf{v}_i) = \langle \mathbf{w}^{\odot 2} \odot \mathbf{v}_i | \mathbf{v}_i \rangle_{\mathcal{T}} = \langle \mathbf{w} \odot \mathbf{v}_i | \mathbf{w} \odot \mathbf{v}_i \rangle_{\mathcal{T}} = \langle \mathbf{u} | \mathbf{u} \rangle_{\mathcal{T}},$$

the Krylov-Bogolyubov error bound Eq. (32) can be computed as

$$\min_{\lambda_j \in \sigma(\mathbf{w})} \frac{|\lambda_j - \langle \mathbf{u} | \mathbf{v}_i \rangle_{\mathcal{T}}|}{|\lambda_j|} \leq \varepsilon_{\lambda} := (\langle \mathbf{u} | \mathbf{u} \rangle_{\mathcal{T}} - \langle \mathbf{u} | \mathbf{v}_i \rangle_{\mathcal{T}}^2)^{1/2}. \quad (33)$$

Observe that this is not an a posteriori error estimate, but an actual *bound*. It is the best possible with the available information [62], and it can be very easily computed inside the iteration to control the possible termination. Similar bounds exist for the approximate eigenvector \mathbf{v}_i . Inserting these considerations into Algorithm 2 gives Algorithm 3 for the task power iteration with RQ computation ($P \leftarrow \text{pow-RQ}$) for the iteration map $\Phi_{\text{pow-RQ}}$ to be used in Algorithm 1, which has the more accurate RQ eigenvalue approximation $\lambda_1 \approx \varrho_{\mathbf{w}}$ and error bound $\varepsilon_{\lambda} = (\langle \mathbf{u} | \mathbf{u} \rangle_{\mathcal{T}} - \langle \mathbf{u} | \mathbf{v}_i \rangle_{\mathcal{T}}^2)^{1/2}$ included as side-effects.

Algorithm 3 One step power iteration with Rayleigh quotient (RQ): $\Phi_{\text{pow-RQ}}$

- | | |
|---|--|
| 1: Input iterate \mathbf{v}_i ; | ▷ assume \mathbf{v}_i of unit length. |
| 2: $\mathbf{u} \leftarrow \mathbf{w} \odot \mathbf{v}_i$; | ▷ $\mathbf{u} = \mathbf{L}_{\mathbf{w}}\mathbf{v}_i$. |
| 3: $\varrho_1 \leftarrow \langle \mathbf{u} \mathbf{v}_i \rangle_{\mathcal{T}}$; | ▷ Rayleigh quotient (RQ) $\lambda_1 \approx \varrho_1$ |
| 4: $\varrho_2 \leftarrow \langle \mathbf{u} \mathbf{u} \rangle_{\mathcal{T}}$; | ▷ RQ of $\mathbf{w}^{\odot 2}$. |
| 5: $\varepsilon_{\lambda} \leftarrow (\varrho_2 - \varrho_1^2)^{1/2}$; | ▷ error bound Eq. (33) of λ_1 . |
| 6: $\gamma \leftarrow \varrho_2^{-1/2}$; | ▷ inverse length of \mathbf{u} . |
| 7: $\mathbf{z} \leftarrow \gamma \cdot \mathbf{u}$; | ▷ normalise output \mathbf{z} to unit length. |
| 8: Output \mathbf{z} ; | |
-

Once the determination of the eigenvalue is accurate enough such that one has an interval which contains only *one* eigenvalue, one may use the even tighter Temple-Kato bounds [62].

Starting vectors and other enhancements like deflation and Krylov subspaces are discussed here. Recalling the discussion of invariant subspaces in the paragraph on maxima and minima in Subsection 2.3, the starting vector \mathbf{v}_0 in Algorithm 1 should not have any zero in it, as this would exclude an invariant subspace from the investigation. One possibility is $\mathbf{v}_0 = \mathbf{1}/\sqrt{N}$, being equal in all positions. With this starting vector one has $\mathbf{v}_1 \propto \mathbf{w} \odot \mathbf{1} = \mathbf{w}^{\odot 1}$, and it is easy to see that what is computed are scaled versions of $\mathbf{w}^{\odot i}$ as eigenvector approximation \mathbf{v}_i . This gives another possibility for acceleration [42], which will be discussed later.

Another technique in eigenvalue computations is *deflation* [34, 78]: when one eigenvalue $\lambda_1 = \lambda_{\hat{\mathbf{m}}}$ of largest modulus and corresponding eigenvector $\mathbf{e}^{(\hat{\mathbf{m}})}$ have been located, one may want to compute a further location $\hat{\mathbf{m}}$ where there is an element $w_{\hat{\mathbf{m}}} = \lambda_{\hat{\mathbf{m}}}$ of equal magnitude $|w_{\hat{\mathbf{m}}}| = |\lambda_1| = |\lambda_{\hat{\mathbf{m}}}|$. For this one can use deflation, and one could change line 2 in Algorithm 2 or Algorithm 3 to $\mathbf{u} \leftarrow \mathbf{w} \odot (\mathbf{1} - \mathbf{e}^{(\hat{\mathbf{m}})}) \odot \mathbf{v}_i$ to obtain a new iteration map for deflation. The

factor $(\mathbf{1} - \mathbf{e}^{(\hat{m})})$ projects into the orthogonal complement of the invariant subspace $\text{span}\{\mathbf{e}^{(\hat{m})}\}$, and thus all the eigenvalues are left unchanged, except for $\lambda_{\hat{m}}$, which is mapped to zero. But in this special case — where we know the form of all invariant subspaces — it is even simpler to keep the same iteration map and to choose as starting vector $\mathbf{v}_0 := (\mathbf{1} - \mathbf{e}^{(\hat{m})})/\sqrt{N-1}$. It has a zero at position \hat{m} and is thus in the invariant subspace $(\text{span}\{\mathbf{e}^{(\hat{m})}\})^\perp$, and all iterates will stay in that subspace. Iterating in Algorithm 1 with that starting vector either with the iteration map in Algorithm 2 or in Algorithm 3 would give us either another eigenvalue $\lambda_{\hat{m}}$ of equal magnitude, or the eigenvalue with second largest magnitude λ_2 , and the corresponding eigenvector. In this manner all desired eigenvalues can be computed via deflation.

Some other acceleration techniques should be mentioned briefly: The convergence speed in power iteration is controlled by the ratio $|\lambda_2/\lambda_1|$, where λ_2 is the next smallest eigenvalue in absolute size. Sometimes this ratio can be very close to unity. If instead with a single vector \mathbf{v}_i one iterates with a whole block $\mathbf{v}_i^{(1)}, \dots, \mathbf{v}_i^{(j)}$ of j mutually orthogonal vectors — effectively a subspace — the convergence speed changes to $|\lambda_{1+j}/\lambda_1|$. One has to restore orthogonality though after each iteration sweep [78] of this block- or subspace iteration. This kind of technique also makes it possible to compute multiple eigenvalues, i.e. when the maximum occurs at several places.

Other well-known methods for eigenvalues with even faster convergence build on *Krylov* subspaces [78], here we have symmetric matrices and thus one would use the Lanczos method. These procedures rely on orthogonalisation though, and this may be problematic when combined with truncation. And certainly the Krylov subspace can be combined with the block iteration idea to give block- or subspace-Lanczos methods [63].

“*Exponentiating*” the power iteration can be easily achieved via the clever idea of [42]. Recall that with the starting vector $\mathbf{v}_0 := \mathbf{1}/\sqrt{N}$, power iteration would compute $\mathbf{v}_1 \propto \mathbf{w}^{\odot 1} = \mathbf{w}^{\odot 1} \odot \mathbf{1}$, a scaled version of \mathbf{w} , corresponding to the action of $\mathbf{w}^{\odot 1}$ on $\mathbf{1}$. Starting actually with $\mathbf{v}_0 \propto \mathbf{w}$ and changing line 2 in Algorithm 2 or Algorithm 3 to $\mathbf{u} := \mathbf{v}_i^{\odot 2}$, one has $\mathbf{v}_1 \propto \mathbf{w}^{\odot 2} = \mathbf{w}^{\odot 2} \odot \mathbf{1}$, i.e. the action of $\mathbf{w}^{\odot 2}$ on $\mathbf{1}$. In the next iteration one has $\mathbf{v}_2 \propto \mathbf{w}^{\odot 4}$, and in iteration i one has $\mathbf{v}_i \propto \mathbf{w}^{\odot 2^i}$. Reformulating Algorithm 3 for this new task of “exponentiated” power iteration ($P \leftarrow \text{exp-pow}$) for the new iteration function $\Phi_{\text{exp-pow}}$ results in Algorithm 4, to be used in Algorithm 1 with starting vector $\mathbf{v}_0 := \mathbf{w}/\|\mathbf{w}\|_{\mathcal{T}}$. We also need to keep the auxiliary vector $\mathbf{y} = \mathbf{w}^{\odot 2}$.

Algorithm 4 One step “exponentiated” power iteration with RQ: $\Phi_{\text{exp-pow}}$

- | | |
|---|---|
| 1: Input iterate \mathbf{v}_i ;
2: $\mathbf{u} \leftarrow \mathbf{v}_i \odot \mathbf{v}_i$;
3: $\varrho_1 \leftarrow \langle \mathbf{w} \mathbf{u} \rangle_{\mathcal{T}}$;
4: $\varrho_2 \leftarrow \langle \mathbf{y} \mathbf{u} \rangle_{\mathcal{T}}$;
5: $\varepsilon_\lambda \leftarrow (\varrho_2 - \varrho_1^2)^{1/2}$;
6: $\gamma \leftarrow \langle \mathbf{u} \mathbf{u} \rangle_{\mathcal{T}}^{-1/2}$;
7: $\mathbf{z} \leftarrow \gamma \cdot \mathbf{u}$;
8: Output \mathbf{z} ;
 | \triangleright assume $\mathbf{v}_i \propto \mathbf{w}^{\odot 2^i}$ of unit length.
$\quad \triangleright \mathbf{u} \propto \mathbf{L}_{\mathbf{w}}^{\odot 2^{i+1}} \mathbf{1} = \mathbf{w}^{\odot 2^{i+1}}$.
\triangleright Rayleigh quotient (RQ) $\lambda_1 \approx \varrho_1$
$\quad \triangleright$ RQ of $\mathbf{w}^{\odot 2}$.
\triangleright error bound Eq. (33) of λ_1 .
$\quad \triangleright$ inverse length of \mathbf{u} .
\triangleright normalise output \mathbf{z} to unit length. |
|---|---|
-

Thus the eigenvalue approximation is through the Rayleigh quotient $\varrho_1 = \langle \mathbf{w} | \mathbf{u} \rangle_{\mathcal{T}} = \langle \mathbf{w} | \mathbf{v}_i^{\odot 2} \rangle_{\mathcal{T}} = \langle \mathbf{w} \odot \mathbf{v}_i | \mathbf{v}_i \rangle_{\mathcal{T}}$ in line 3 of Algorithm 4. As $\mathbf{y} = \mathbf{w}^{\odot 2}$, on line 4 the quantity $\varrho_2 = \langle \mathbf{y} | \mathbf{u} \rangle_{\mathcal{T}} = \langle \mathbf{w}^{\odot 2} | \mathbf{v}_i^{\odot 2} \rangle_{\mathcal{T}} = \langle \mathbf{w}^{\odot 2} \odot \mathbf{v}_i | \mathbf{v}_i \rangle_{\mathcal{T}}$ computes the RQ of $\mathbf{w}^{\odot 2}$. The “eigenvector” $\mathbf{v}_i \propto \mathbf{w}^{\odot 2^i}$ converges with the rate $|\lambda_2/\lambda_1|^{2^i}$, and hence one achieves *exponential* convergence. This acceleration technique can be used in all eigenvalue computations described here.

Transformed eigenvalue computations are a well-known device to find e.g. the maximum of the function $\mathbf{y} = f(\mathbf{w})$. After computing $f(\mathbf{w})$, one proceeds as before, but with \mathbf{y} instead of \mathbf{w} . Some functions $f(t)$ are simple enough so that the computation of $f(\mathbf{w})$ is combined with the iteration. Recall that this transforms all eigenvalues $\lambda_{\mathbf{w}}$ to $\lambda_{\mathbf{y}} = f(\lambda_{\mathbf{w}})$ according to spectral calculus, leaving the eigenvectors unchanged.

Shifting and inverse shifting are functions needed to access other than the point of largest modulus in the spectrum of $\mathbf{L}_{\mathbf{w}}$ resp. \mathbf{w} . These functions are the shift $f : t \mapsto t + \beta$ and the inverse shift $f : t \mapsto (t - \rho)^{-1}$.

The simple function of shifting $f(t) = t + \beta$ is needed if the actual element of maximum modulus is a minimum of \mathbf{w} ; then this is a maximum of $-\mathbf{w}$. In the iteration this will be picked up by the Rayleigh-quotient being negative: $\varrho_{\mathbf{w}} < 0$. Also note that for this case — where the maximum has smaller absolute value than the minimum — one may *shift* the tensor by a value of $\beta = -\varrho_{\mathbf{w}} \approx -\lambda_1$ to $\hat{\mathbf{w}} := \mathbf{w} + \beta \cdot \mathbf{1}$ so that the maximum of $\hat{\mathbf{w}}$ has larger absolute value than the minimum of $\hat{\mathbf{w}}$. One then iterates with $\hat{\mathbf{w}}$. After determining the maximum of $\hat{\mathbf{w}}$, we may subtract the number $\beta > 0$ again to obtain the maximum of \mathbf{w} . The index of the maximum of $\hat{\mathbf{w}}$ is of course the same as the one for the maximum of \mathbf{w} . Completely analogous manipulations can be performed to find a minimum which is not of maximum modulus. These and similar techniques are well known from eigenvalue calculations of large / sparse symmetric matrices [34, 78, 80, 92].

The inverse shift $f(t) = (t - \rho)^{-1}$ is needed to access intermediate points in the spectrum $\sigma(\mathbf{w})$. Here one finds the data in \mathbf{w} closest to ρ , as they have the maximum modulus under the transformation. Thus one computes first $\mathbf{y} := (\mathbf{w} - \rho \cdot \mathbf{1})^{\odot -1}$ approximately through Algorithm 1 with the iteration function Eq. (25) — with $\check{\mathbf{w}} := \mathbf{w} - \rho \cdot \mathbf{1}$ instead of \mathbf{w} — followed by the eigenvector computation with iteration function as in Algorithm 4, but with starting vector $\mathbf{v}_0 := \mathbf{y} / \|\mathbf{y}\|_{\mathcal{T}}$.

3. Tensor formats

3.1. Overview of tensor formats

In this section we review definitions and properties of frequently used tensor formats. Many such formats are used in quantum physics under the name *tensor networks*, see [91, 81, 32, 69, 9, 7]. We only look at the *canonical polyadic* (CP) and the *Tensor Train* (TT) format. For the Tucker [89] format see [29]. The CP [49] format has been well known for a long time and is therefore very popular. The TT format was originally developed in quantum physics and chemistry as “matrix product states” (MPS), see [91] and references therein, and rediscovered in [71, 72] as tensor train. A new class of tensor formats, which we do not consider, is the hierarchical tensor (HT) format. It was introduced in [46], and further considered in [37].

Note that the sets of low-rank tensors of fixed rank in the CP format are not closed for $d > 2$, whereas in the Tucker, TT, or HT tensor formats these sets are closed. The computations in the TT and HT formats are based on the *singular value decomposition* (SVD) [76, 37], and in the Tucker format on the *higher order SVD* (HOSVD) [90]. In all tensor formats the tensor rank doubles for addition and squares for the Hadamard product. To avoid unnecessary and harmful rank growth, the rank is usually truncated, e.g. by ALS-like or other optimisation algorithms, for the CP format see e.g. [24, 30], for the TT format [50], and for the Tucker format [90].

A tensor format is described by a parameter vector space $\mathcal{P} = \times_{\nu=1}^d \mathcal{P}_{\nu}$, where $\mathcal{P}_{\nu} = \mathbb{R}^{M_{\nu}}$, and a multilinear map $U : \mathcal{P} \rightarrow \mathcal{T}$ into the tensor space $\mathcal{T} := \otimes_{\ell=1}^d \mathbb{R}^{M_{\ell}}$. For practical implementations of high dimensional problems we need to distinguish between a tensor $\mathbf{w} \in \mathcal{T}$ and its tensor format representation $\mathbf{P} \in \mathcal{P}$, where $\mathbf{w} = U(\mathbf{P})$. There are many possibilities to define tensor formats. Here the canonical (CP) and the tensor train (TT) formats are considered. In the following we briefly repeat definitions and properties of these formats [47]. Note also that the reader can invent his own tensor format, one which will fit especially well to his needs.

The CP format is cheap, it is simpler than the Tucker or TT format, but, compared to others, there are no reliable algorithms to compute CP decompositions for $d > 2$ [14, 47, 53]. The Tucker format has stable algorithms [54], but the storage and complexity costs are $\mathcal{O}(drn + r^d)$, i.e. they grow exponentially with d . The TT format is a bit more complicated, but does not have this disadvantage. CP and Tucker rank-structured tensor formats have been applied in chemometrics and in signal processing [85, 11].

3.2. The canonical polyadic tensor format

The canonical representation of multivariate functions [49] was introduced in 1927. A schema of the CP tensor format for $d = 3$ is shown in Figure 1. The full tensor $\mathbf{w} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ is shown on the left side of the equality sign, and its decomposition on the right. The lines denote the vectors w_{i1} , w_{i2} , and w_{i3} for $i = 1, \dots, r$, respectively.

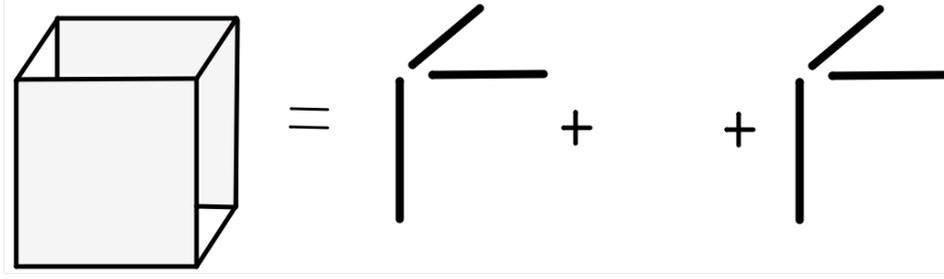


Figure 1: Schema of the CP tensor decomposition of a 3D tensor.

Definition 3.1 (Canonical Polyadic (CP) Tensor Format). *The canonical polyadic tensor format in \mathcal{T} for variable r is defined by the multilinear mapping*

$$U_{\text{CP},r} : \mathcal{P}_{\text{CP},r} := \prod_{\nu=1}^d \mathcal{P}_{\nu}^r \rightarrow \mathcal{T}, \quad \mathcal{P}_{\nu} = \mathbb{R}^{M_{\nu}}, \quad (34)$$

$$\mathcal{P}_{\text{CP},r} \ni \mathbf{P} := (\mathbf{w}_i^{(\nu)} : 1 \leq i \leq r, 1 \leq \nu \leq d) \mapsto$$

$$U_{\text{CP},r}(\mathbf{P}) := \mathbf{w} = \sum_{i=1}^r \bigotimes_{\nu=1}^d \mathbf{w}_i^{(\nu)} \in \mathcal{T}.$$

We call the sum of elementary tensors $\mathbf{w} := U_{\text{CP},r}(\mathbf{P})$ a tensor represented in the canonical tensor format with r terms. The system of vectors $\mathbf{P} = (\mathbf{w}_i^{(\nu)} \in \mathbb{R}^{M_{\nu}} : 1 \leq i \leq r, 1 \leq \nu \leq d)$ is a representation system of \mathbf{w} with representation rank r . One may think of \mathbf{P} as of a vector valued $r \times d$ matrix with the vector $\mathbf{w}_i^{(\nu)} \in \mathbb{R}^{M_{\nu}}$ at index position (i, ν) .

The storage requirement for $\mathbf{w} = U_{\text{CP},r}(\mathbf{P})$ is $r \times \sum_{\nu=1}^d M_{\nu}$, and in the simple case $M_1 = \dots = M_d = n$ it is $\mathcal{O}(r d n)$.

3.2.1. Basic operations with the canonical format

We denote the set of all tensors $\mathbf{w} \in \mathcal{T}$ of rank r by \mathcal{T}^r . The set \mathcal{T}^r is a cone, i.e. $\mathbf{w} \in \mathcal{T}^r$ implies $\alpha \cdot \mathbf{w} \in \mathcal{T}^r$ for $\alpha \in \mathbb{R}$, and at the same time \mathcal{T}^r is not a vector space as for $\mathbf{w}_1, \mathbf{w}_2 \in \mathcal{T}^r$ one has in general $\mathbf{w}_1 + \mathbf{w}_2 \notin \mathcal{T}^r$, but $\mathbf{w}_1 + \mathbf{w}_2 \in \mathcal{T}^{2r}$ [30].

A complete description of fundamental operations in the canonical tensor format and a their numerical cost can be found in [47]. For recent algorithms in the canonical tensor format we refer to [24, 25, 26, 30].

Multiplication by a scalar $\alpha \in \mathbb{R}$ could be done for a $\mathbf{w} = U_{\text{CP},r}(\mathbf{P})$ as in Eq. (34) by multiplying all of the vectors $\{\mathbf{w}_j^{(\nu)}, i = 1, \dots, r\}$ for any ν by α . But to spread the effect equally and keep the vectors balanced in size, we recommend to define $\alpha_{\nu} := \sqrt[d]{|\alpha|}$ for all $\nu > 1$, and $\alpha_1 := \text{sign}(\alpha) \sqrt[d]{|\alpha|}$. Then, with a computational cost of $\mathcal{O}(r n d)$, and without changing the rank,

$$\alpha \cdot \mathbf{w} = \sum_{j=1}^r \alpha \bigotimes_{\nu=1}^d \mathbf{w}_j^{(\nu)} = \sum_{j=1}^r \bigotimes_{\nu=1}^d (\alpha_{\nu} \mathbf{w}_j^{(\nu)}) = \sum_{j=1}^r \bigotimes_{\nu=1}^d \tilde{\mathbf{w}}_j^{(\nu)}, \quad \text{with } \tilde{\mathbf{w}}_j^{(\nu)} = \alpha_{\nu} \mathbf{w}_j^{(\nu)}.$$

The sum of two tensors in the CP format $\mathbf{w} = \mathbf{u} + \mathbf{v}$ can be written as follows

$$\mathbf{w} = \mathbf{u} + \mathbf{v} = \left(\sum_{j=1}^{r_u} \bigotimes_{\nu=1}^d \mathbf{u}_j^{(\nu)} \right) + \left(\sum_{k=1}^{r_v} \bigotimes_{\mu=1}^d \mathbf{v}_k^{(\mu)} \right) = \sum_{j=1}^{r_u+r_v} \bigotimes_{\nu=1}^d \mathbf{w}_j^{(\nu)},$$

where $\mathbf{w}_j^{(\nu)} := \mathbf{u}_j^{(\nu)}$ for $j \leq r_u$ and $\mathbf{w}_j^{(\nu)} := \mathbf{v}_j^{(\nu)}$ for $r_u < j \leq r_u + r_v$. The result generally has rank $r_u + r_v$, and as this operation requires only concatenation of memory it has only a computing cost of $\mathcal{O}(1)$.

The Hadamard product $\mathbf{w} = \mathbf{u} \odot \mathbf{v}$ can be written as follows

$$\mathbf{w} = \mathbf{u} \odot \mathbf{v} = \left(\sum_{j=1}^{r_u} \bigotimes_{\nu=1}^d \mathbf{u}_j^{(\nu)} \right) \odot \left(\sum_{k=1}^{r_v} \bigotimes_{\nu=1}^d \mathbf{v}_k^{(\nu)} \right) = \sum_{j=1}^{r_u} \sum_{k=1}^{r_v} \bigotimes_{\nu=1}^d \left(\mathbf{u}_j^{(\nu)} \odot \mathbf{v}_k^{(\nu)} \right).$$

The new rank is generally $r_u \times r_v$, and the computational cost is $\mathcal{O}(r_u r_v n d)$ arithmetic operations.

The Euclidean inner product is computed as follows:

$$\langle \mathbf{u} | \mathbf{v} \rangle_{\mathcal{T}} = \left\langle \sum_{j=1}^{r_u} \bigotimes_{\nu=1}^d \mathbf{u}_j^{(\nu)} \middle| \sum_{k=1}^{r_v} \bigotimes_{\nu=1}^d \mathbf{v}_k^{(\nu)} \right\rangle_{\mathcal{T}} = \sum_{j=1}^{r_u} \sum_{k=1}^{r_v} \prod_{\nu=1}^d \langle \mathbf{u}_j^{(\nu)} | \mathbf{v}_k^{(\nu)} \rangle_{\mathcal{P}_{\nu}}.$$

The computational cost of the inner product is $\mathcal{O}(r_u r_v n d)$.

Note that in all operations the numerical cost grows only linearly with d , but the representation rank of the resulting tensors may increase. Therefore, a rank truncation procedure is needed, which approximates a given tensor represented in the canonical format with lower rank tensors up to a given accuracy.

3.2.2. Rank truncation in the CP format

Let \mathbf{w} be a tensor of rank R . Truncating \mathbf{w} to a new rank $r < R$ is a fundamental problem [24]. This problem can be formulated as follows:

$$\text{find a } \mathbf{w}^* \text{ with rank } r \text{ such that } \forall \mathbf{u} \text{ with rank } r : \|\mathbf{w} - \mathbf{w}^*\| \leq \|\mathbf{w} - \mathbf{u}\|.$$

Typical methods to solve this problem are the alternating least squares (ALS) method and the Gauss-Newton-method. The ALS method may show slow convergence, see [24, 30] and references therein. The Gauss-Newton method for $d \geq 3$ requires some additional assumptions, and also may not show any convergence at all [24].

As was pointed in Subsection 3.1, see for instance pp. 91–92 in [53] or [14], the set of rank- r CP tensors \mathcal{T}^r is not a closed set in the corresponding tensor product space \mathcal{T} for $d > 2$. Therefore there is no minimising \mathbf{w}^* as required above, and one may look for an ε -solution, i.e. minimising within a deviation of ε from the infimum:

Definition 3.2 (Approximation Problem). *For a given tensor \mathbf{w} in CP format with rank R and $\varepsilon > 0$ we are looking for minimal $r_{\varepsilon} < R$ and a tensor \mathbf{w}^* of rank r_{ε} in CP format, such that:*

$$\|\mathbf{w} - \mathbf{w}^*\| \leq \varepsilon \|\mathbf{w}\|.$$

This problem is discussed in [24, 30], and a collection of CP optimisation methods was compared in [74].

Properties of the TT format are described in Appendix AppendixA, and of the Tucker format in [29].

4. Numerical examples

The algorithms for post-processing high-dimensional data, or large volumes of data are shown on a few illustrative examples. Obviously, in any such application the actual computing times and even the possibility of executing such algorithms depend not only on that the huge amounts of data can be compressed to reasonable size, but also on the possibility of executing the algebraic operations with a reasonable speed on the compressed data. Additionally, as the compression may deteriorate in the course of the computation, one must assume that the intermediate results can be compressed to reasonable size, and that this re-compression can be done efficiently. All this depends very much on the data, and currently there are only very few general results to predict when this might or might not occur. Some results in this direction are already contained in [28, 18, 19] for the low-rank TT-format which was also used in our computations.

4.1. Finding the value and location of the maximum element

Already in Subsection 1.2 one could see an example from [24], finding the maximum of the numerical solution $\mathbf{w} = (w_{m_1, \dots, m_d})$ of a d -dimensional Poisson equation

$$-\nabla^2 u(x_1, \dots, x_d) = f(x_1, \dots, x_d) \propto \sum_{k=1}^d \prod_{\ell=1, \ell \neq k}^d x_\ell (1 - x_\ell) \quad \text{on } \mathcal{G} = [0, 1]^d \text{ with } u|_{\partial\mathcal{G}} = 1, \quad (35)$$

which was discretised with a standard finite-difference method with $n = 100$ grid-points $(x_1^{(m_1)}, \dots, x_d^{(m_d)})$ with $(1 \leq m_\ell \leq n, 1 \leq \ell \leq d)$ in each dimension. So with $N = n^d$ data points, the numerical solution is $\mathcal{T} = \mathbb{R}^N = \mathbb{R}^{n^d} \ni \mathbf{w} = (w_{m_1, \dots, m_d}) = (u(x_1^{(m_1)}, \dots, x_d^{(m_d)}))$. The results are shown in Table 1 in Subsection 1.2 and illustrate the kinds of savings which are possible when compared to the method of inspecting every element. We used low-rank TT-format storage, and Algorithm 2. In no case did it need more than 20 iterations. The actual computing times are shown in the fourth column of Table 1.

4.2. Finding level sets

We show two examples, the first is the same as in the previous Subsection 4.1, whereas the second one is a solution to a stochastic partial differential equation (SPDE).

Level set of the solution to a high-dimensional Poisson equation [24]. Let us consider the same d -dimensional Poisson problem Eq. (35) as in Subsection 4.1 with the same discretisation and the same data tensor $\mathbf{w} \in \mathcal{T} = \mathbb{R}^N = \mathbb{R}^{n^d}$.

Now we want to find the indices of the elements which are in the set $S =]-\infty, 0.25[$, and according to Subsection 2.3 one has to compute the characteristic function χ_S as given in Eq. (19).

For this one needs the sign function given in Eq. (18), computed with the Newton-Schulz iteration function in Eq. (27). The results are shown in Table 2. The first three columns are similar to those in Table 1, and the fourth column are the actual computing times in seconds on a 3 GHz CPU, which behave like $\mathcal{O}(n^d)$.

Table 2: Computing times (4th column) on 3 GHz CPU to compute χ_S .

d	# elements.: $N = n^d$	\approx years [a] inspect. N	actual 3 GHz CPU iteration time [s] for Newton-Schulz
25	10^{50}	10^{33}	0.63
50	10^{100}	10^{83}	1.30
100	10^{200}	10^{133}	2.66
200	10^{400}	10^{233}	6.03
400	10^{800}	10^{433}	12.97
800	10^{1600}	10^{1233}	25.32

Level sets of the solution to a stochastic partial differential equation (SPDE) is the next example. This example is described in a bit more detail in [28].

Two scientific software libraries *sglib* and *TensorCalculus* [93, 31] were used. With the *sglib* procedures we discretise the SPDE Eq. (36), and with the *TensorCalculus* we solve the obtained tensor equation and compute the quantities of interests.

Consider the following diffusion equation with uncertain diffusion coefficient $\kappa(\omega, x)$:

$$\left. \begin{aligned} -\nabla(\kappa(\omega, x)\nabla u(\omega, x)) &= f(\omega, x) && \text{in } \mathcal{G}, \\ u(\omega, x) &= 0 && \text{on } \partial\mathcal{G}, \end{aligned} \right\} \text{a.s. in } \omega \in \Omega, \quad (36)$$

Here \mathcal{G} is the two-dimensional L-shaped domain $[-1, 1]^2 \setminus [0, 1]^2$, and a triangular mesh with 557 mesh points was used for the spatial piece-wise linear finite element discretisation, and, including the boundary values, the discrete spatial solution needs $M_d = 557$ storage locations.

After a stochastic discretisation with the Karhunen-Loève and Polynomial Chaos Expansion, and applying the stochastic Galerkin method as in [65, 64, 18, 19, 28, 27], we can compute the solution $u(\omega, x)$ in a chosen tensor representation.

The random field $\kappa(\omega, x)$ was taken to have a shifted log-normal distribution for, i.e. $\log(\kappa(\omega, x) - 1.1)$ has a normal distribution with parameters $\{\mu = 0.5, \sigma^2 = 1.0\}$. The isotropic and homogeneous covariance function is of Gaussian type with covariance lengths $\ell_x = \ell_y = 0.3$.

Ten Karhunen-Loève terms were taken to represent the mean zero random part of the field $\kappa(\omega, x)$, and for the polynomial chaos expansion (PCE) multivariate Hermite polynomials of maximum second degree were taken.

For the right-hand side $f(\omega, x)$ a β -distribution $\{4, 2\}$ was chosen for the random field. The covariance function is also of Gaussian type with covariance lengths $\ell_x = \ell_y = 0.6$. Again ten terms in the Karhunen-Loève expansion for the zero mean random part of $f(\omega, x)$ with a maximum second degree polynomial chaos were taken.

The total stochastic dimension of the solution $u(\omega, x)$ is 20 — ten random variables from the ten terms of the Karhunen-Loève expansion of the coefficient $\kappa(\omega, x)$, plus ten terms of the Karhunen-Loève expansion of the right hand side $f(\omega, x)$ — i.e. the multi-index α will consist of 20 indices ($\alpha = (\alpha_1, \dots, \alpha_{20})$), and the solution will be represented as a function of 20 independent random variables. Together with the extra dimension from the spatial discretisation — which is lumped into one — one has a discrete version of a function $u(\omega, x)$ on $d = 21$ dimensions. With second degree polynomials, there are $M_s = 3$ polynomials for each random variable, thus the stochastic number degrees of freedom is $N_s = M_s^{20} = 3^{20} = 3,486,784,401 \approx 3.5 \times 10^9$. Thus the total number of entries for the full tensor is $N_d \times N_s = N_d \times M_s^{20} = 557 \times 3,486,784,401 = 1,942,138,911,357 \approx 2 \times 10^{12}$, and the discrete solution $\mathbf{u} \in \mathbb{R}^{557} \otimes \bigotimes_{\mu=1}^{20} \mathbb{R}^3 \cong \mathbb{R}^{1,942,138,911,357}$. Thus, even for this relatively coarse discretisation, full storage of just the solution vector would require approximately 16 TB of memory.

For the numerical solution, a low-rank tensor solver was used [66, 18, 19, 28, 27], and the solution was represented in the CP-format Eq. (34) with a rank of $r = 231$. Thus

$$\mathbf{u} = \sum_{j=1}^{231} \mathbf{u}_{j0} \otimes \bigotimes_{\mu=1}^{20} \mathbf{u}_{j\mu} \in \mathbb{R}^{557} \otimes \bigotimes_{\mu=1}^{20} \mathbb{R}^3,$$

and the storage of the solution vector requires only $231 \times (557 + 20 \times 3) = 142,527$ storage locations, i.e. less than 1.15 MB.

Table 3: Computing $\chi_S(\mathbf{u})$.

$b/\ \mathbf{u}\ _\infty$	rank χ_S	max it rank	k_{\max} its	error
0.2	12	24	12	2.9×10^{-8}
0.4	12	20	20	1.9×10^{-7}
0.6	8	16	12	1.6×10^{-7}
0.8	8	15	8	1.2×10^{-7}

Finally, we computed $\chi_S(\mathbf{u})$ for $S =] - \infty, b[$ by computing $\text{sign}(b\|\mathbf{u}\|_\infty \mathbf{1} - \mathbf{u})$ for $b \in \{0.2, 0.4, 0.6, 0.8\}$, which is shown in the first column of Table 3. The final representation ranks of $\chi_S(\mathbf{u})$ resp. $\text{sign}(b\|\mathbf{u}\|_\infty \mathbf{1} - \mathbf{u})$ are given in the second column. In this numerical example, the ranks are in all cases smaller than 13. The sign-function was computed with the iteration from Algorithm 1. The iteration function is given in Eq. (27), and the iteration terminated after k_{\max} steps, shown in column four. The maximal representation rank of the iterates is documented in the third column. The error $\|\mathbf{1} - \mathbf{u}_{k_{\max}} \odot \mathbf{u}_{k_{\max}}\| / \|(b\|\mathbf{u}\|_\infty \mathbf{1} - \mathbf{u})\|$ is given in the last column. Each computation — each row of the table — took less than 10 min on a 3 GHz CPU.

5. Conclusion

Large volumes of data are used resp. generated in an increasing number of instances in science, technology, and business; some of which were mentioned in Subsection 1.1. Often this data is in the form of a sample of a high dimensional function, or it can be reshaped in this way. It may well happen that it is not possible to store such data in its entirety, so that it has to be compressed in some way. And it is usually not enough to compress these large data sets, efficient numerical algorithms are required to post-process such data in a compressed data format. All these post-processing tasks may be trivial in low dimensions and for small data sets, but they become non-trivial for high-dimensional data, say $d > 5$, which is in some compressed format. We have introduced some numerical methods which may partially close this gap.

In this work we assumed that 1) the data are given or can be approximated in a compressed data format (e.g., in a low-rank tensor format); 2) a recompression (tensor rank truncation in our case) procedure is available and the ranks after linear algebra operations are not increases too much.

Here we have formulated algorithms for such post-processing, which only use the structure of an abstract commutative algebra. The Hadamard algebra is an example of such an algebra, and it can be defined for all real valued data sets. The Hadamard product of two tensors turned out to be a key tool in approximating the variance, point-wise inverse, level sets, frequency, and the maximal element of huge multi-dimensional data sets.

In our example application, the compression is low-rank tensor approximation, and we showed how to perform the algebra operations in some example low-rank tensor formats. We assumed that the initial data set has low-rank tensor representation, and showed that this low-rank property can be preserved during the whole computing process.

The choice of the tensor format is not crucial; it could be, for example, the canonical polyadic format, the Tucker format, the tensor train format, and many others. The only requirement is that it should be possible to compute the operations of the Hadamard algebra in a reasonable (linear in n) time, and the existence of stable rank truncation algorithms to truncate possibly large intermediate ranks. The algorithms needed were among others the multiplicative inverse and the sign function, and this could be fashioned after the algorithms for matrix algebras. Finally, the algorithms are demonstrated on some high dimensional data which comes from the solution of high dimensional or parametric resp. stochastic elliptic partial differential equations.

Appendix A. The Tensor Train format

The tensor train (TT) format is described in [76, 72, 47, 53]. As already noted, it was originally developed in quantum chemistry as “matrix product states” (MPS), see [91] and references therein, and rediscovered later [71, 72]. For a motivation, let us start with a well-known example [51].

Example Appendix A.1. *Consider the d -dimensional Laplacian operator discretised with standard finite differences over a uniform tensor grid with n degrees of freedom in each direction. It has the Kronecker (canonical) rank- d representation:*

$$\mathbf{A} = \mathbf{A} \otimes \overbrace{\mathbf{I} \otimes \cdots \otimes \mathbf{I}}^{d-1 \text{ times}} + \mathbf{I} \otimes \mathbf{A} \otimes \mathbf{I} \cdots \otimes \mathbf{I} + \cdots + \mathbf{I} \otimes \cdots \otimes \mathbf{I} \otimes \mathbf{A} \in \mathbb{R}^{n^d \times n^d}$$

with $\mathbf{A} = \text{tridiag}\{-1, 2, -1\} \in \mathbb{R}^{n \times n}$, and \mathbf{I} being the $n \times n$ identity. However, the same operator in the TT format is explicitly representable with all TT ranks equal to 2 for any dimension [51]. In the element-wise matrix product notation, the Laplace operator reads

$$\mathbf{A}[\mathbf{i}, \mathbf{j}] = \begin{pmatrix} \mathbf{A}[i_1, j_1] & \mathbf{I}[i_1, j_1] \end{pmatrix} \begin{pmatrix} \mathbf{I}[i_2, j_2] & 0 \\ \mathbf{A}[i_2, j_2] & \mathbf{I}[i_2, j_2] \end{pmatrix} \cdots \begin{pmatrix} \mathbf{I}[i_d, j_d] \\ \mathbf{A}[i_d, j_d] \end{pmatrix}.$$

Definition Appendix A.2 (TT-Format, TT-Representation, TT-Ranks). *The TT-tensor format is for variable TT-representation ranks $\mathbf{r} = (r_0, \dots, r_d) \in \mathbb{N}^{d+1}$ — with $r_0 = r_d = 1$ and under the assumption that $d > 2$ — defined by*

the following multilinear mapping

$$\begin{aligned}
U_{\text{TT}} : \mathcal{P}_{\text{TT}, \mathbf{r}} &:= \prod_{\nu=1}^d \mathcal{P}_{\nu}^{r_{\nu-1} \times r_{\nu}} \rightarrow \mathcal{T}, \quad \mathcal{P}_{\nu} = \mathbb{R}^{M_{\nu}} \quad (\nu = 1, \dots, d), \\
\mathcal{P}_{\text{TT}, \mathbf{r}} \ni \mathbf{P} &= (\mathbf{W}^{(\nu)} = (\mathbf{w}_{j_{\nu-1} j_{\nu}}^{(\nu)}) \in \mathcal{P}_{\nu}^{r_{\nu-1} \times r_{\nu}} : 1 \leq j_{\nu} \leq r_{\nu}, 1 \leq \nu \leq d) \\
&\mapsto U_{\text{TT}}(\mathbf{P}) := \mathbf{w} = \sum_{j_0=1}^{r_0} \cdots \sum_{j_d=1}^{r_d} \bigotimes_{\nu=1}^d \mathbf{w}_{j_{\nu-1} j_{\nu}}^{(\nu)} \in \mathcal{T}.
\end{aligned} \tag{A.1}$$

We call $\mathbf{w} := (w_{i_1 \dots i_d}) = U_{\text{TT}, \mathbf{r}}(\mathbf{P})$ a tensor represented in the train tensor format. Note that the TT-cores $\mathbf{W}^{(\nu)}$ may be viewed as a vector valued $r_{\nu-1} \times r_{\nu}$ matrix with the vector $\mathbf{w}_{j_{\nu-1} j_{\nu}}^{(\nu)} \in \mathbb{R}^{M_{\nu}}$ with the components $w_{j_{\nu-1} j_{\nu}}^{(\nu)}[i_{\nu}] : 1 \leq i_{\nu} \leq M_{\nu}$ at index position $(j_{\nu-1}, j_{\nu})$. The representation in components is then

$$(w_{i_1 \dots i_d}) = \sum_{j_0=1}^{r_0} \cdots \sum_{j_d=1}^{r_d} w_{j_0 j_1}^{(1)}[i_1] \cdots w_{j_{\nu-1} j_{\nu}}^{(\nu)}[i_{\nu}] \cdots w_{j_{d-1} j_d}^{(d)}[i_d] \tag{A.2}$$

Alternatively, each TT-core $\mathbf{W}^{(\nu)}$ may be seen as a vector of $r_{\nu-1} \times r_{\nu}$ matrices $\mathbf{W}_{i_{\nu}}^{(\nu)}$ of length M_{ν} , i.e. $\mathbf{W}^{(\nu)} = (\mathbf{W}_{i_{\nu}}^{(\nu)}) : 1 \leq i_{\nu} \leq M_{\nu}$. Then the representation Eq. (A.2) reads

$$(w_{i_1 \dots i_d}) = \prod_{\nu=1}^d \mathbf{W}_{i_{\nu}}^{(\nu)}, \tag{A.3}$$

which explains the name matrix product state. Observe that the first matrix is always a row vector as $r_0 = 1$, and the last matrix is always a column vector as $r_d = 1$. The matrix components $\mathbf{W}_{i_{\nu}}^{(\nu)}$ of the TT-cores are also called “carriages” or “waggon” with “wheels” i_{ν} at the bottom, coupled to the next “carriage” or “waggon” via the matrix product. This explains the tensor train name. If one notes more carefully $\mathbf{W}^{(\nu)} \in \mathcal{P}_{\nu}^{r_{\nu-1} \times r_{\nu}} = \mathbb{R}^{r_{\nu-1}} \otimes \mathbb{R}^{M_{\nu}} \otimes \mathbb{R}^{r_{\nu}}$, then Eq. (A.3) can be written more concisely as

$$\mathbf{w} = U_{\text{TT}}(\mathbf{P}) = \mathbf{W}^{(1)} \times_3^1 \mathbf{W}^{(2)} \times_3^1 \cdots \times_3^1 \mathbf{W}^{(d)}, \tag{A.4}$$

where $\mathbf{U} \times_k^{\ell} \mathbf{V}$ is a contraction of the k -th index of \mathbf{U} with the ℓ -th index of \mathbf{V} , where one often writes just \times_k for \times_k^1 . Thus in Eq. (A.4) the contractions leave the indices from the $\mathbb{R}^{M_{\nu}}$ untouched, so that the tensor \mathbf{w} is formed.

Each TT-core (or block) $\mathbf{W}^{(\nu)}$ is defined by $r_{\nu-1} \times r_{\nu} \times M_{\nu}$ numbers. Assuming $n = M_{\nu}$ for all $\nu = 1, \dots, d$, the total number of entries scales as $\mathcal{O}(dn r^2)$, which is tractable as long as $r = \max\{r_k\}$ is moderate.

A pictorial representation of the schema for the TT tensor format is shown in Figure A.2. It shows d connected waggon with one wheel. The waggon denote the TT-cores, and each wheel denotes the index i_{ν} . The waggon for $\nu = 2, \dots, (d-1)$ are connected with their neighbours by two indices $j_{\nu-1}$ and j_{ν} . The first and the last waggon are connected by only one index, namely j_1 and j_{d-1} respectively. Since by the convention in Definition AppendixA.2 above, $r_0 = r_d = 1$, the indices j_0 and j_d run from 1 to 1, i.e. are purely formal.

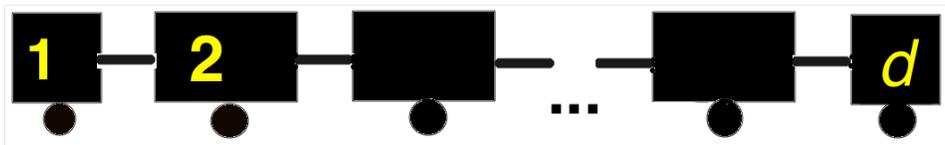


Figure A.2: Schema of the TT tensor decomposition. The waggon denote the TT cores and each wheel denotes the index i_{ν} . Each waggon is connected with neighbours by indices $j_{\nu-1}$ and j_{ν} .

The waggon or carriage 2 — $\mathbf{W}^{(2)}$ — is a tensor of degree 3, described by three indices j_1 (the left hitch), i_2 (the wheel), and j_2 (the right hitch). Multiplication of the third TT-core with the second and forth cores means the tensor contraction by the indices j_1 and j_2 . If we perform tensor contraction of all TT-cores over the indices j_1, \dots, j_{d-1} ,

and disregard the purely formal constant indices j_0 and j_d , then the indices j_0, \dots, j_d — the hitches — will disappear, and only the indices i_1, \dots, i_d — the wheels — will be left.

Appendix A.1. Basic operations with the TT format

We follow to the work of Oseledets [72] and list the major properties of the TT-tensor format.

The multiplication with scalar. α could be simply done by multiplying one of the TT-cores $\mathbf{W}^{(\nu)}$ in the representation Eq. (A.4) for any ν in $\mathbf{w} = \mathbf{W}^{(1)} \times_{\frac{1}{3}} \mathbf{W}^{(2)} \times_{\frac{1}{3}} \dots \times_{\frac{1}{3}} \mathbf{W}^{(d)}$. But to balance the effect better, define $\alpha_\nu := \sqrt[d]{|\alpha|}$ for all $\nu > 1$, and $\alpha_1 := \text{sign}(\alpha) \sqrt[d]{|\alpha|}$. Then $\tilde{\mathbf{w}} = \alpha \cdot \mathbf{w}$ is given by

$$\tilde{\mathbf{w}} = (\alpha_1 \cdot \mathbf{W}^{(1)}) \times_{\frac{1}{3}} (\alpha_2 \cdot \mathbf{W}^{(2)}) \times_{\frac{1}{3}} \dots \times_{\frac{1}{3}} (\alpha_d \cdot \mathbf{W}^{(d)}) = \tilde{\mathbf{W}}^{(1)} \times_{\frac{1}{3}} \dots \times_{\frac{1}{3}} \tilde{\mathbf{W}}^{(d)}.$$

The new cores are given by $\tilde{\mathbf{W}}^{(\nu)} = (\tilde{\mathbf{W}}_{i_\nu}^{(\nu)}) = (\alpha_\nu \mathbf{W}_{i_\nu}^{(\nu)})$, a sequence of new “carriage” matrices. The computational complexity is $\mathcal{O}(dnr^2)$.

Addition of two TT-tensors. Assume two tensors \mathbf{u} and \mathbf{v} are given in the TT-tensor format as in Eq. (A.3), i.e. $(u_{i_1 \dots i_d}) = \prod_{\nu=1}^d \mathbf{U}_{i_\nu}^{(\nu)}$ and $(v_{i_1 \dots i_d}) = \prod_{\nu=1}^d \mathbf{V}_{i_\nu}^{(\nu)}$. The sum $\mathbf{w} = \mathbf{u} + \mathbf{v}$ is given by the new cores $\mathbf{W}_{i_\nu}^{(\nu)}$ such that $(w_{i_1 \dots i_d}) = \prod_{\nu=1}^d \mathbf{W}_{i_\nu}^{(\nu)}$, where

$$\mathbf{W}_{i_\nu}^{(\nu)} = \begin{pmatrix} \mathbf{U}_{i_\nu}^{(\nu)} & \mathbf{0} \\ \mathbf{0} & \mathbf{V}_{i_\nu}^{(\nu)} \end{pmatrix}, \quad 1 \leq i_\nu \leq r_\nu, 2 \leq \nu \leq d-1.$$

and the first and the last cores will be

$$\mathbf{W}_{i_1}^{(1)} = \begin{pmatrix} \mathbf{U}_{i_1}^{(1)} & \mathbf{V}_{i_1}^{(1)} \end{pmatrix} \quad \text{and} \quad \mathbf{W}_{i_d}^{(d)} = \begin{pmatrix} \mathbf{U}_{i_d}^{(d)} \\ \mathbf{V}_{i_d}^{(d)} \end{pmatrix}.$$

As only storage may have to be concatenated, the computational cost is $\mathcal{O}(1)$, but as the carriages resp. TT-cores grow, the final rank will generally be the sum of the ranks.

The Hadamard product. $\mathbf{w} = \mathbf{u} \odot \mathbf{v}$ in the TT format is computed as follows. Assume two tensors \mathbf{u} and \mathbf{v} are given in the TT tensor format as in Eq. (A.3), i.e. $(u_{i_1 \dots i_d}) = \prod_{\nu=1}^d \mathbf{U}_{i_\nu}^{(\nu)}$ and $(v_{i_1 \dots i_d}) = \prod_{\nu=1}^d \mathbf{V}_{i_\nu}^{(\nu)}$. The Hadamard product is

$$(w_{i_1 \dots i_d}) = (u_{i_1 \dots i_d} \cdot v_{i_1 \dots i_d}).$$

The tensor \mathbf{w} has also the TT-tensor format, namely with the new cores

$$\mathbf{W}_{i_\nu}^{(\nu)} = \mathbf{U}_{i_\nu}^{(\nu)} \otimes_{\mathbb{K}} \mathbf{V}_{i_\nu}^{(\nu)}, \quad 1 \leq i_\nu \leq r_\nu, 1 \leq \nu \leq d,$$

where $\otimes_{\mathbb{K}}$ is the Kronecker product of two matrices [47]. The rank of $\mathbf{W}^{(\nu)} = (\mathbf{W}_{i_\nu}^{(\nu)})$ is the product of the ranks of the TT-cores $\mathbf{U}^{(\nu)}$ and $\mathbf{V}^{(\nu)}$.

The Euclidean inner product. of two tensors in the TT-format as in Eq. (A.1)

$$\mathbf{u} = \sum_{j_0=1}^{r_0^u} \dots \sum_{j_d=1}^{r_d^u} \bigotimes_{\nu=1}^d \mathbf{u}_{j_{\nu-1} j_\nu}^{(\nu)}, \quad \mathbf{v} = \sum_{j_0=1}^{r_0^v} \dots \sum_{j_d=1}^{r_d^v} \bigotimes_{\nu=1}^d \mathbf{v}_{j_{\nu-1} j_\nu}^{(\nu)},$$

with ranks r^u and r^v can be computed as follows:

$$\langle \mathbf{u} | \mathbf{v} \rangle_{\mathcal{T}} = \sum_{j_0=1}^{r_0^u} \dots \sum_{j_d=1}^{r_d^u} \sum_{i_0=1}^{r_0^v} \dots \sum_{i_d=1}^{r_d^v} \prod_{\nu=1}^d \langle \mathbf{u}_{j_{\nu-1} j_\nu}^{(\nu)} | \mathbf{v}_{i_{\nu-1} i_\nu}^{(\nu)} \rangle_{\mathcal{P}_\nu}.$$

The computational complexity is $\mathcal{O}(dnr^4)$, and can be reduced further to $\mathcal{O}(dnr^3)$ [77].

Appendix A.2. Rank truncation in the TT format

The rank truncation operation is based on the SVD algorithm and requires $\mathcal{O}(dnr^3)$ operations [38]. The TT-rounding algorithm (p. 2305 in [72]) is based on QR decomposition and costs $\mathcal{O}(dnr^3)$.

Corollary 2.4 in [72] states that for a given tensor \mathbf{w} and rank bounds r_k , the best approximation to \mathbf{w} in the Frobenius norm with TT-ranks bounded by r_k always exist (denote it by \mathbf{w}^*), and the TT-approximation \mathbf{u} computed by the TT-SVD algorithm (p. 2301 in [72]) is quasi-optimal:

$$\|\mathbf{w} - \mathbf{u}\|_F \leq \sqrt{d-1} \|\mathbf{w} - \mathbf{w}^*\|_F. \quad (\text{A.5})$$

In [60] the authors suggested a new re-compression randomised algorithm for Tucker and TT tensor formats.

Acknowledgments

Alexander Litvinenko was supported by the Alexander von Humboldt Foundation. Hermann Matthies and Elmar Zander were partly supported by the “Deutsche Forschungsgemeinschaft” (DFG). We would also like to thank the two anonymous reviewers for their suggestions and comments.

References

- [1] M. Bachmayr and W. Dahmen, *Adaptive near-optimal rank tensor approximation for high-dimensional operator equations*, *Found Comput Math* **15** (2015), 839–898, doi:10.1007/s10208-013-9187-3.
- [2] J. Ballani and L. Grasedyck, *Tree adaptive approximation in the hierarchical tensor format*, *SIAM Journal on Scientific Computing* **36** (2014), A1415–A1431, doi:10.1137/130926328.
- [3] ———, *Hierarchical tensor approximation of output quantities of parameter-dependent pdes*, *SIAM/ASA Journal on Uncertainty Quantification* **3** (2015), 852–872, doi:10.1137/140960980.
- [4] J. Ballani, L. Grasedyck, and M. Kluge, *Black box approximation of tensors in hierarchical Tucker format*, *Linear Algebra and its Applications* **438** (2013), no. 2, 639–657, doi:10.1016/j.laa.2011.08.010.
- [5] M. Bebendorf and C. Kuske, *Separation of variables for function generated high-order tensors*, *J. Sci. Comput.* **61** (2014), no. 1, 145–165, doi:10.1007/s10915-014-9822-4.
- [6] U. Benedikt, H. Auer, M. Espig, W. Hackbusch, and A. Auer, *Tensor representation techniques in post-Hartree-Fock methods: matrix product state tensor format*, *Molecular Physics* **111** (2013), no. 16–17, 2398–2413, doi:10.1080/00268976.2013.798433.
- [7] J. Biamonte and V. Bergholm, *Quantum tensor networks in a nutshell* [online], arXiv: 1708.00006 [quant-ph], 2017, Available from: <https://arxiv.org/abs/1708.00006>, arXiv:1708.00006.
- [8] D. Bigoni, A. P. Engsig-Karup, and Y. M. Marzouk, *Spectral tensor train decomposition*, *SIAM Journal of Scientific Computing* **38** (2016), no. 4, A2405–A2439, doi:10.1137/15M1036919.
- [9] J. C. Bridgeman and C. T. Chubb, *Hand-waving and interpretive dance: An introductory course on tensor networks*, *J. Phys. A: Math. Theor.* **50** (2017), 223001, doi:10.1088/1751-8121/aa6dc3.
- [10] S. R. Chinnamsetty, M. Espig, B. N. Khoromskij, W. Hackbusch, and H. J. Flad, *Tensor product approximation with optimal rank in quantum chemistry*, *The Journal of Chemical Physics* **127** (2007), no. 8, 084–110, doi:10.1063/1.2761871.
- [11] A. Cichocki, , and S. Sh. Amari, *Adaptive blind signal and image processing: Learning algorithms and applications*, Wiley, 2002.

- [12] C. Da Silva and F. J. Herrmann, *Optimization on the hierarchical Tucker manifold – applications to tensor completion*, Linear Algebra and its Applications **481** (2015), 131–173, doi:10.1016/j.laa.2015.04.015.
- [13] L. De Lathauwer, B. D. Moor, and J. Vandewalle, *A multilinear singular value decomposition*, SIAM Journal on Matrix Analysis and Applications **21** (2000), no. 4, 1253–1278, doi:10.1137/S0895479896305696.
- [14] V. de Silva and L.-H. Lim, *Tensor rank and the ill-posedness of the best low-rank approximation problem*, SIAM Journal on Matrix Analysis and Applications **30** (2008), no. 3, 1084–1127, doi:10.1137/06066518X.
- [15] S. Dolgov, K. Anaya-Izquierdo, C. Fox, and R. Scheichl, *Approximation and sampling of multivariate probability distributions in the tensor train decomposition*, Statistics and Computing (2020), 603–625, doi:10.1007/s11222-019-09910-z.
- [16] S. Dolgov and B. Khoromskij, *Simultaneous state-time approximation of the chemical master equation using tensor product formats*, Numerical Linear Algebra with Applications **22** (2015), no. 2, 197–219, doi:10.1002/nla.1942.
- [17] S. Dolgov, B. Khoromskij, D. Savostyanov, and I. Oseledets, *Computation of extreme eigenvalues in higher dimensions using block tensor train formats*, Comp. Phys. Communications **185** (2014), no. 4, 1207–1216.
- [18] S. Dolgov, B. N. Khoromskij, A. Litvinenko, and H. G. Matthies, *Computation of the response surface in the tensor train data format* [online], arXiv: 1406.2816 [math.NA], 2014, Available from: <https://arxiv.org/abs/1406.2816>, arXiv:1406.2816.
- [19] S. Dolgov, B. N. Khoromskij, A. Litvinenko, and H. G. Matthies, *Polynomial chaos expansion of random coefficients and the solution of stochastic partial differential equations in the tensor train format*, SIAM/ASA Journal on Uncertainty Quantification **3** (2015), no. 1, 1109–1135, doi:10.1137/140972536.
- [20] S. Dolgov, A. Litvinenko, and D. Liu, *Kriging in tensor train data format*, Proceedings, 3rd International Conference on Uncertainty Quantification in Computational Sciences and Engineering, 2019, pp. 309–329, Available from: https://files.eccomasproceedia.org/papers/e-books/unccecomp_2019.pdf, doi:10.7712/120219.6343.18651.
- [21] S. Dolgov and D. Savostyanov, *Parallel cross interpolation for high-precision calculation of high-dimensional integrals*, Computer Physics Communications **246** (2020), 106869, doi:10.1016/j.cpc.2019.106869.
- [22] S. Dolgov, B. Khoromskij, and D. V. Savostyanov, *Superfast Fourier transform using QTT approximation*, Journal of Fourier Analysis and Applications **18** (2012), no. 5, 915–953, doi:10.1007/s00041-012-9227-4.
- [23] J. Dölz, H. Harbrecht, and C. Schwab, *Covariance regularity and H-matrix approximation for rough random fields*, Numerische Mathematik **135** (2017), 1045–1071, doi:10.1007/s00211-016-0825-y.
- [24] M. Espig, *Effiziente Bestapproximation mittels Summen von Elementartensoren in hohen Dimensionen*, Ph.D. thesis, Universität Leipzig, Germany, 2008.
- [25] M. Espig, L. Grasedyck, and W. Hackbusch, *Black box low tensor-rank approximation using fiber-crosses*, Constructive approximation **30** (2009), 557–597, doi:10.1007/s00365-009-9076-9.
- [26] M. Espig and W. Hackbusch, *A regularized Newton method for the efficient approximation of tensors represented in the canonical tensor format*, Numerische Mathematik **122** (2012), no. 3, 489–525, doi:10.1007/s00211-012-0465-9.
- [27] M. Espig, W. Hackbusch, A. Litvinenko, H. G. Matthies, and P. Wähnert, *Efficient low-rank approximation of the stochastic Galerkin matrix in tensor formats.*, Computers & Mathematics with Applications **67** (2014), 818–829, doi:10.1016/j.camwa.2012.10.08.

- [28] M. Espig, W. Hackbusch, A. Litvinenko, H. G. Matthies, and E. Zander, *Efficient analysis of high dimensional data in tensor formats*, Sparse Grids and its Applications (J. Garcke and M. Griebel, eds.), LNCSE, vol. 88, Springer-Verlag, 2013, pp. 31–56, doi:10.1007/978-3-642-31703-3_2.
- [29] M. Espig, W. Hackbusch, A. Litvinenko, H. G. Matthies, and E. Zander, *Post-Processing of High-Dimensional Data* [online], arXiv: 1906.05669 [math.NA], 2019, Available from: <https://arxiv.org/abs/1906.05669>, arXiv: 1906.05669.
- [30] M. Espig, W. Hackbusch, T. Rohwedder, and R. Schneider, *Variational calculus with sums of elementary tensors of fixed rank*, Numerische Mathematik **122** (2012), no. 3, 469–488, doi:10.1007/s00211-012-0464-x.
- [31] M. Espig, M. Schuster, A. Killaitis, N. Waldren, P. Waehnert, S. Handschuh, and H. Auer, *Tensorcalculus, C++ library*, 2012, Available from: <http://gitorious.org/tensorcalculus>.
- [32] G. Evenbly and G. Vidal, *Tensor network states and geometry*, J Stat Phys **145** (2011), 891–918, doi:10.1007/s10955-011-0237-4.
- [33] H.-J. Flad, B. Khoromskij, D. Savostyanov, and et al., *Verification of the cross 3d algorithm on quantum chemistry data*, Russian Journal of Numerical Analysis and Mathematical Modelling **23** (2008), no. 4, 329–344, doi:10.1515/RJNAMM.2008.020.
- [34] G. Golub and C. F. van Loan, *Matrix computations*, Johns Hopkins University Press, Baltimore, MD, 1996.
- [35] S. A. Goreinov, I. V. Oseledets, D. V. Savostyanov, E. Tyrtshnikov, and N. L. Zamarashkin, *How to find a good submatrix*, Matrix Methods: Theory, Algorithms and Applications, 2010, pp. 247–256, doi:10.1142/9789812836021_0015.
- [36] A. Gorodetsky, S. Karaman, and Y. M. Marzouk, *A continuous analogue of the tensor-train decomposition*, Computer Methods in Applied Mechanics and Engineering **347** (2019), 59–84, doi:10.1016/j.cma.2018.12.015.
- [37] L. Grasedyck, *Hierarchical singular value decomposition of tensors*, SIAM Journal on Matrix Analysis and Applications **31** (2010), 2029–2054, doi:10.1137/090764189.
- [38] L. Grasedyck and W. Hackbusch, *An introduction to hierarchical (H-) rank and TT-rank of tensors with examples*, Comput. Methods Appl. Math. **11** (2011), no. 3, 291–304, doi:10.2478/cmam-2011-0016.
- [39] L. Grasedyck and S. Krämer, *Stable ALS approximation in the TT-format for rank-adaptive tensor completion*, Numerische Mathematik **143** (2019), no. 4, 855–904, doi:10.1007/s00211-019-01072-4.
- [40] L. Grasedyck, D. Kressner, and C. Tobler, *A literature survey of low-rank tensor approximation techniques*, GAMM-Mitteilungen **36** (2013), no. 1, 53–78.
- [41] L. Grasedyck, R. Kriemann, C. Löbbert, A. Nägel, G. Wittum, and K. Xylouris, *Parallel tensor sampling in the hierarchical Tucker format*, Computing and Visualization in Science **17** (2015), 67–78, doi:10.1007/s00791-015-0247-x.
- [42] L. Grasedyck, L. Juschka, and C. Löbbert, *Finding entries of maximum absolute value in low-rank tensors* [online], arXiv: 1912.02072 [math.NA], 2019, Available from: <https://arxiv.org/abs/1912.02072>, arXiv:1912.02072.
- [43] W. Hackbusch and B. N. Khoromskij, *Tensor-product approximation to operators and functions in high dimensions*, J. Complexity **23** (2007), no. 4-6, 697–714. MR MR2372023 (2008k:65042)
- [44] W. Hackbusch, B. N. Khoromskij, S. Sauter, and E. E. Tyrtshnikov, *Use of tensor formats in elliptic eigenvalue problems*, Numerical Linear Algebra with Applications **19** (2012), no. 1, 133–151, doi:10.1002/nla.793.

- [45] W. Hackbusch, B. N. Khoromskij, and E. E. Tyrtyshnikov, *Approximate iterations for structured matrices*, Numerische Mathematik **109** (2008), no. 3, 365–383, doi:10.1007/s00211-008-0143-0.
- [46] W. Hackbusch and S. Kühn, *A new scheme for the tensor representation*, Journal of Fourier Analysis and Applications **15** (2009), no. 5, 706–722, doi:10.1007/s00041-009-9094-9.
- [47] W. Hackbusch, *Tensor Spaces and Numerical Tensor Calculus*, Springer, Berlin, 2012, doi:10.1007/978-3-642-28027-6.
- [48] N. Higham, *Functions of matrices — theory and computation*, SIAM, Philadelphia, PA, 2008.
- [49] F. L. Hitchcock, *The expression of a tensor or a polyadic as a sum of products*, J. Math. Physics **6** (1927), 164–189.
- [50] S. Holtz, T. Rohwedder, and R. Schneider, *The alternating linear scheme for tensor optimization in the tensor train format*, SIAM Journal on Scientific Computing **34** (2012), A683–A713, doi:10.1137/100818893.
- [51] V. Kazeev and B. Khoromskij, *Low-rank explicit QTT representation of the Laplace operator and its inverse*, SIAM Journal on Matrix Analysis and Applications **33** (2012), 742—758, doi:10.1137/100820479.
- [52] V. Khoromskaia and B. N. Khoromskij, *Tensor numerical methods in quantum chemistry*, Walter de Gruyter GmbH & Co KG, 2018.
- [53] B. N. Khoromskij, *Tensor numerical methods in scientific computing*, Walter de Gruyter GmbH & Co KG, 2018.
- [54] B. N. Khoromskij and V. Khoromskaia, *Low rank Tucker-type tensor approximation to classical potentials*, Cent. Eur. J. Math. **5** (2007), no. 3, 523–550. MR MR2322828 (2008c:65118)
- [55] B. N. Khoromskij, A. Litvinenko, and H. G. Matthies, *Application of hierarchical matrices for computing the Karhunen–Loève expansion*, Computing **84** (2009), no. 1–2, 49–67, doi:10.1007/s00607-008-0018-3.
- [56] B. N. Khoromskij and C. Schwab, *Tensor-structured Galerkin approximation of parametric and stochastic elliptic PDEs*, SIAM Journal on Scientific Computing **33** (2011), 364–385, doi:10.1137/100785715.
- [57] T. Kolda, *Orthogonal tensor decompositions*, SIAM J. Matrix Anal. Appl. **23** (2001), 243–255.
- [58] T. Kolda and B. W. Bader, *Tensor decompositions and applications*, SIAM Review **51** (2009), no. 3, 455–500.
- [59] D. A. Kolesnikov and I. V. Oseledets, *Convergence analysis of projected fixed-point iteration on a low-rank matrix manifold*, Numerical Linear Algebra with Applications **25** (2018), no. 5, e2140, doi:10.1002/nla.2140.
- [60] D. Kressner and L. Perisa, *Recompression of Hadamard products of tensors in Tucker format*, SIAM Journal on Scientific Computing **39** (2017), A1879–A1902, doi:10.1137/16M1093896.
- [61] O. Lebedeva, *Tensor conjugate-gradient-type method for rayleigh quotient minimization in block QTT-format*, Russian Journal of Numerical Analysis and Mathematical Modelling **26** (2011), no. 5, 465–489, doi:10.1515/rjnamm.2011.026.
- [62] H. G. Matthies, *Computable error bounds for the generalized symmetric eigenproblem*, Communications in Applied Numerical Methods **1** (1985), 33–38, doi:10.1002/cnm.1630010107.
- [63] ———, *A subspace Lanczos method for the generalized symmetric eigenproblem*, Computers and Structures **21** (1985), 319–325, doi:10.1016/0045-7949(85)90252-4.
- [64] ———, *Stochastic finite elements: Computational approaches to stochastic partial differential equations*, Zeitschrift für angewandte Mathematik und Mechanik **88** (2008), no. 11, 849–873, doi:10.1002/zamm.200800095.

- [65] H. G. Matthies and A. Keese, *Galerkin methods for linear and nonlinear elliptic stochastic partial differential equations*, *Computer Methods in Applied Mechanics and Engineering* **194** (2005), no. 12–16, 1295–1331, doi:10.1016/j.cma.2004.05.027.
- [66] H. G. Matthies and E. Zander, *Solving stochastic systems with low-rank tensor compression*, *Linear Algebra and its Applications* **436** (2012), 3819–3838, doi:10.1016/j.laa.2011.04.017.
- [67] H. G. Matthies and R. Ohayon, *Analysis of parametric models — Linear methods and approximations*, *Advances in Computational Mathematics* **45** (2019), 2555–2586, doi:10.1007/s10444-019-09735-4.
- [68] W. Nowak and A. Litvinenko, *Kriging and spatial design accelerated by orders of magnitude: combining low-rank covariance approximations with FFT-techniques*, *Mathematical Geosciences* **45** (2013), no. 4, 411–435, doi:10.1007/s11004-013-9453-6.
- [69] R. Orús, *A practical introduction to tensor networks: Matrix product states and projected entangled pair states*, *Annals of Physics* **349** (2014), 117–158, doi:10.1016/j.aop.2014.06.013.
- [70] I. Oseledets, *DMRG approach to fast linear algebra in the TT-format*, *Comput. Methods Appl. Math.* **11** (2011), no. 3, 382–393, doi:10.2478/cmam-2011-0021.
- [71] I. Oseledets and E. Tyrtshnikov, *TT-cross approximation for multidimensional arrays*, *Linear Algebra and its Applications* **432** (2010), 70–88, doi:10.1016/j.laa.2009.07.024.
- [72] I. V. Oseledets, *Tensor-train decomposition*, *SIAM J. Sci. Comput.* **33** (2011), no. 5, 2295–2317, doi:10.1137/090752286.
- [73] I. V. Oseledets, *Constructive representation of functions in low-rank tensor formats*, *Constr. Appr.* **37** (2013), no. 1, 1–18, doi:10.1007/s00365-012-9175-x.
- [74] I. V. Oseledets and D. V. Savostyanov, *Minimization methods for approximating tensors and their comparison*, *Comput. Math. and Math. Phys.* **46** (2006), 1641–1650, doi:10.1134/S0965542506100022.
- [75] I. V. Oseledets, D. V. Savostyanov, and E. E. Tyrtshnikov, *Linear algebra for tensor problems*, *Computing* **85** (2009), 169–188, doi:10.1007/s00607-009-0047-6.
- [76] I. V. Oseledets and E. Tyrtshnikov, *Breaking the curse of dimensionality, or how to use SVD in many dimensions*, *SIAM J. Scientific Computing* **31** (2009), no. 5, 3744–3759.
- [77] I. Oseledets, *Matlab TT-toolbox, version 2.2*, 2011, Available from: <https://github.com/oseledets/TT-Toolbox>.
- [78] B. N. Parlett, *The symmetric eigenvalue problem*, SIAM, Philadelphia, PA, 1998.
- [79] M. V. Rakhuba and I. V. Oseledets, *Fast multidimensional convolution in low-rank tensor formats via cross approximation*, *SIAM Journal on Scientific Computing* **37** (2015), A565–A582, doi:10.1137/140958529.
- [80] Y. Saad, *Numerical methods for large eigenvalue problems: Theory and algorithms*, Manchester University Press, Manchester, 1992.
- [81] S. Sachdev, *Tensor networks—a new tool for old problems*, *Physics* **2** (2009), 90, doi:10.1103/Physics.2.90.
- [82] D. Savostyanov and I. Oseledets, *Fast adaptive interpolation of multi-dimensional arrays in tensor train format*, *The 2011 International Workshop on Multidimensional (nD) Systems*, Sep. 2011, pp. 1–8, doi:10.1109/nDS.2011.6076873.
- [83] C. Schwab and R. A. Todor, *Karhunen-Loève approximation of random fields by generalized fast multipole methods*, *J. Comput. Phys.* **217** (2006), no. 1, 100–122. MR MR2250527 (2007d:65011)

- [84] I. E. Segal and R. A. Kunze, *Integrals and operators*, 2nd ed., Springer, Berlin, 1978.
- [85] A. Smilde, R. Bro, and P. Geladi, *Multi-way analysis with applications in the chemical sciences*, Wiley, 2004.
- [86] M. Steinlechner, *Riemannian optimization for high-dimensional tensor completion*, SIAM Journal on Scientific Computing **38** (2016), S461–S484, doi:10.1137/15M1010506.
- [87] A. M. Stuart, *Inverse problems: A bayesian perspective*, Acta Numerica **19** (2010), 451–559, doi:10.1017/S0962492910000061.
- [88] A. V. Sulimov, D. Zheltkov, I. V. Oferkin, D. C. Kutov, E. V. Katkova, E. E. Tyrtysnikov, and V. B. Sulimov, *Tensor train global optimization: Application to docking in the configuration space with a large number of dimensions*, Communications in Computer and Information Science. Conference: Russian Supercomputing Days, vol. 793, 2017, pp. 151–167, doi:10.1007/978-3-319-71255-0_12.
- [89] L. R. Tucker, *Some mathematical notes on three-mode factor analysis*, Psychometrika **31** (1966), 279–311.
- [90] N. Vannieuwenhoven, R. Vandebril, and K. Meerbergen, *A new truncation strategy for the higher-order singular value decomposition*, SIAM Journal on Scientific Computing **34** (2012), A1027–A1052, doi:10.1137/110836067.
- [91] G. Vidal, *Efficient classical simulation of slightly entangled quantum computations*, Phys. Rev. Lett. **91** (2003), no. 14, 147902–0, doi:10.1103/PhysRevLett.91.147902.
- [92] D. Watkins, *The matrix eigenvalue problem: GR and Krylov subspace methods*, SIAM, Philadelphia, PA, 2007.
- [93] E. Zander, *SGLib - A Matlab Toolbox for Stochastic Galerkin Methods*, February 2010, Available from: <http://github.com/ezander/splib/>.
- [94] E. Zander, *Tensor methods for stochastic problems*, Ph.D. thesis, Technische Universität Braunschweig, Germany, 2012.