# Low-Rank Tensor Methods for High-Dimensional Problems 

Présentée le 31 mars 2023
Faculté des sciences de base
Algorithmes numériques et calcul haute performance - Chaire CADMOS Programme doctoral en mathématiques
pour l'obtention du grade de Docteur ès Sciences
par

## Christoph Max STRÖSSNER

Acceptée sur proposition du jury
Prof. M. Colombo, présidente du jury
Prof. D. Kressner, directeur de thèse
Prof. H. Harbrecht, rapporteur
Prof. A. Uschmajew , rapporteur
Prof. F. Nobile, rapporteur

## Acknowledgements

First and foremost, I am genuinely grateful to my supervisor Daniel Kressner for his advice, encouragement, and patience during the last years. Daniel provided me this great opportunity to explore the fascinating world of tensors and their decompositions. Giving me the freedom and flexibility to work on different problems and topics, encouraging me to travel to various conferences and summer schools, trusting me to handle things independently, and guiding me professionally through the unusual circumstances during a global pandemic characterizes Daniel's outstanding support and is far from standard. I thank Helmut Harbrecht, André Uschmajew and Fabio Nobile for being part of my jury and for taking the time to read my thesis. Special thanks go to Maria Colombo for agreeing to preside over my private defense.
Over the last years, I had the pleasure to collaborate with Sergey Dolgov, Virginie Ehrlacher, Jad Dabaghi and Bonan Sun who provided valuable insights and assistance. I would like to particularly express my gratitude to Virginie Ehrlacher for introducing me to self-diffusion matrices during my six-week stay at CIRM in Luminy as part of CEMRACS 2021 and for inviting me to visit École des ponts ParisTech. I would like to thank Nick Trefethen and the Chebfun team giving me the opportunity to contribute to the official Chebfun package. I am grateful for the help of Behnam Hashemi, Nick Hale, Yuji Nakatsukasa and Alex Townsend.
Moreover, I would like to thank my colleagues and those with whom I had many insightful discussions at EPFL; namely Alice, Amal, Axel, Bonan, David, Francesco, Gianluca, Haoze, Hysan, Ivana, Junli, Kathryn, Lana, Margherita, Minhong, Nico, Stefano, Tingting and Yuxin. Special thanks go to Gianluca who helped proof-reading this thesis.
Last but not least, I would like to thank my friends, my girlfriend, and my family for their immense support over the years.

## Abstract

In this thesis, we propose and analyze novel numerical algorithms for solving three different high-dimensional problems involving tensors. The commonality of these problems is that the tensors can potentially be well approximated in low-rank formats. Identifying and exploiting this low-rank structure allows us to mitigate the curse of dimensionality.
The first problem considered in this thesis is the computation of functional low-rank approximations of multivariate functions defined on tensor product domains. We develop two novel algorithms to compute such approximations by combining tensorized Chebyshev interpolation with low-rank approximations of the coefficient tensor. For most functions, our numerical experiments demonstrate that our algorithms require a lower number of function evaluations and achieve the same approximation error compared to existing methods. In addition, we solve partial differential equations using a novel global spectral method that can potentially be combined with functional low-rank approximations.
The second problem of interest is the solution of multi-marginal optimal transport problems. After adding entropic regularization, these are equivalent to tensor scaling problems that can be solved using the Sinkhorn algorithm. In literature, it has been suggested to accelerate the Sinkhorn algorithm by exploiting either a graphical model structure or a low-rank tensor approximation. We propose to combine these two approaches to accelerate the solution of the tensor scaling problem even further.
The third problem of interest is the computation of the self-diffusion matrix of a tagged particle process defined on a grid. We propose a novel approach based on computing the matrix via solving a high-dimensional tensor-valued optimization problem. We observe numerically that our approach is much less subject to statistical noise compared to classical approaches based on estimating long-time averages of empirical means of deviations of some stochastic processes.

Key words: low-rank approximation, high-dimensional problems, tensor methods, curse of dimensionality, Chebyshev interpolation, functional low-rank approximation, multimarginal optimal transport, self-diffusion matrix, spectral method

## Zusammenfassung

Diese Arbeit befasst sich mit der Entwicklung und Analyse numerischer Algorithmen zur Lösung dreier hochdimensionaler Probleme. Diese haben gemein, dass sie Tensoren beinhalten, die möglicherweise gut in Niedrigrangformaten approximiert werden können. Durch Ausnutzen dieser Niedrigrangstruktur kann der Fluch der Dimensionalität gelindert werden.
Das erste Problem stellt die Approximation multivariater Funktionen, die auf Tensorproduktmengen definiert sind, mithilfe funktionaler Niedrigrangapproximationen dar. Dafür werden zwei neuartige Algorithmen entwickelt, die die gewünschte Approximation durch Kombination von tensorisierter Tschebyscheff-Interpolation und Niedrigrangapproximation des Koeffiziententensors berechnen. In numerischen Experimenten wird gezeigt, dass die neuen Algorithmen für die meisten Funktionen weniger Funktionsauswertungen als existierende Methoden benötigen und dennoch die gleiche Genauigkeit erzielen. Zudem wird eine Spektralmethode zur Lösung partieller Differenzialgleichungen entwickelt, die mit funktionalen Niedrigrangapproximationen kombiniert werden könnte.
Anschliessend befasst sich diese Arbeit mit dem optimalen Transport mehrerer Randverteilungen. Das zugehörige mathematische Problem kann durch entropische Regularisierung in ein Tensorskalierungsproblem umgewandelt werden, das wiederum mit dem SinkhornAlgorithmus gelöst werden kann. In der Literatur existieren Ansätze, diesen Algorithms entweder durch das Ausnutzen eines probabilistischen grafischen Modells oder durch eine Niedrigrangapproximation zu beschleunigen. Diese beiden Ansätze werden in dieser Arbeit kombiniert, um den Sinkhorn-Algorithmus noch stärker zu beschleunigen.
Zuletzt wird das Problem der Berechnung von Selbstdiffusionsmatrizen stochastischer Prozesse mit markiertem Teilchen auf Gittern untersucht. In dieser Arbeit wird zur Lösung dieses Problems ein neuartiger Ansatz basierend auf der Lösung eines hochdimensionalen Minimierungsproblems präsentiert. Im Vergleich zu klassischen Algorithmen, die die langfristige gemittelte Abweichung eines stochastischen Prozesses mit einer Monte-CarloSimulation bestimmen, führt der neuartige Algorithmus in numerischen Experimenten zu deutlich weniger statistischem Rauschen.

Stichwörter: Niedrigrangapproximation, hochdimensionale Probleme, Tensormethoden, Fluch der Dimensionalität, Tschebyscheff-Interpolation, funktionale Niedrigrangapproximation, optimaler Transport mehrerer Randverteilungen, Selbstdiffusionsmatrix, Spektralmethode

## Contents

Acknowledgements ..... i
Abstract ..... iii
Zusammenfassung ..... v
1 Introduction ..... 1
2 Preliminaries ..... 7
2.1 Low-rank approximation ..... 7
2.1.1 Low-rank approximation of matrices ..... 7
2.1.2 Low-rank approximation of tensors ..... 11
2.1.3 Low-rank tensor formats with additional structure ..... 12
2.1.4 Approximation algorithms ..... 13
2.2 Functional low-rank approximation ..... 15
2.3 Chebyshev interpolation ..... 17
2.3.1 Chebyshev interpolation for univariate functions. ..... 17
2.3.2 Tensorized Chebyshev interpolation for multivariate functions ..... 18
2.3.3 Coefficient approximation ..... 20
2.3.4 Connection to functional low-rank approximations ..... 22
2.3.5 Proof of concept ..... 22
3 Functional Tucker approximation ..... 25
3.1 Existing algorithm: Chebfun3 ..... 26
3.1.1 Phase 1: Block term decomposition ..... 26
3.1.2 Phase 2: Refinement ..... 27
3.1.3 Phase 3: Compression ..... 27
3.1.4 Disadvantages ..... 28
3.2 Novel algorithm: Chebfun3F ..... 29
3.2.1 Phase 1: Fiber indices and factor matrices ..... 29
3.2.2 Phase 2: Refinement of the factors ..... 31
3.2.3 Phase 3: Reconstruction of the core tensor ..... 31
3.2.4 Chebfun3F algorithm ..... 32
3.2.5 Existence of a quasi-optimal Chebfun3F approximation ..... 34
3.2.6 Comparison of the theoretical cost ..... 35
3.3 Numerical results ..... 36
3.3.1 Chebfun3 vs. Chebfun3F ..... 36
3.3.2 Comparison to sparse grids ..... 39
3.4 Coefficient decay analysis ..... 41
3.4.1 Observations in a two-dimensional setting ..... 42
3.4.2 Potential solutions in the two-dimensional setting ..... 43
3.4.3 Incorporation of Chebfun3F in Chebfun ..... 44
3.5 Additional insights ..... 45
3.5.1 Approximations in Lebesgue spaces ..... 45
3.5.2 Further compression ..... 46
4 Extended functional tensor train approximation ..... 51
4.1 Extended functional tensor train format ..... 52
4.2 Approximation algorithm ..... 54
4.3 Numerical experiments ..... 56
4.3.1 Comparison to a direct TT approximation ..... 57
4.3.2 Comparison to the FTT approximation algorithm ..... 59
4.A Appendix ..... 63
4.A. 1 Benchmark functions ..... 63
4.A. 2 Genz functions ..... 65
4.A. 3 Parametric PDE problem ..... 65
5 Multi-marginal optimal transport ..... 67
5.1 Multi-marginal optimal transport and the Sinkhorn algorithm ..... 68
5.1.1 Mathematical setting ..... 69
5.1.2 Multi-marginal Sinkhorn algorithm ..... 69
5.2 Impact of approximating the Gibbs kernel ..... 72
5.3 Tensor networks and graphical models ..... 76
5.4 Low-rank approximations in tensor networks ..... 79
5.5 Numerical experiments ..... 80
5.5.1 Proof of concept ..... 80
5.5.2 Application: Color transfer from color barycenters ..... 83
5.5.3 A tensor network with circles ..... 86
6 Self-diffusion matrix ..... 91
6.1 Infinite-dimensional definition ..... 92
6.1.1 Definition as optimization problem ..... 92
6.1.2 Definition as long time mean square deviation ..... 93
6.2 Finite-dimensional approximation ..... 94
6.2.1 Discretized minimization problems ..... 94
6.2.2 Combined minimization problem ..... 95
6.2.3 Estimation of long-time mean square deviation ..... 96
6.3 Low-rank solutions for the optimization problem ..... 97
6.3.1 Fast and stable evaluation ..... 98
6.3.2 Successive minimization ..... 99
6.3.3 Alternating least squares ..... 100
6.3.4 Monte Carlo methods ..... 102
6.3.5 Limitations of the approach ..... 102
6.4 Numerical Experiments ..... 104
6.4.1 Solving the optimization problem ..... 104
6.4.2 Estimation of long-time mean square deviation ..... 105
6.4.3 Comparison of algorithms ..... 107
6.5 Application: Cross-diffusion system ..... 109
6.5.1 Hydrodynamic limit of a lattice-based stochastic hopping model ..... 109
6.5.2 Deterministic resolution of a simplified cross-diffusion system ..... 111
6.5.3 Numerical Experiment ..... 114
7 Fast global spectral method ..... 117
7.1 Problem setting ..... 119
7.1.1 Structure of a linear differential operator ..... 119
7.1.2 Approximation format ..... 119
7.2 Operator discretization ..... 119
7.2.1 One-dimensional differential operators ..... 120
7.2.2 Three-dimensional differential operators ..... 122
7.2.3 Generalization to non-constant coefficients ..... 122
7.3 A spectral method for three-dimensional linear PDEs ..... 125
7.3.1 Boundary condition discretization ..... 126
7.3.2 Incorporating the boundary conditions ..... 127
7.3.3 Solving tensor-valued linear systems ..... 129
7.4 Numerical results ..... 130
7.4.1 Runtime comparison ..... 130
7.4.2 Stationary problems ..... 131
7.4.3 Time-dependent problems ..... 137
7.4.4 Eigenvalue problems ..... 138
8 Conclusions and outlook ..... 141
Bibliography ..... 145
Curriculum Vitae ..... 167

## 1 Introduction

This thesis is concerned with the solution of three different high-dimensional problems: the approximation of multivariate functions using tensorized Chebyshev interpolation, the solution of entropically regularized multi-marginal optimal transport problems, and the computation of self-diffusion matrices via solving a high-dimensional optimization problem.

At a first glace, the problems studied in this thesis might seem vastly different, but at heart all of them involve a tensor $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{d}}$ of order $d$, where $d$ denotes the dimension of the problem. Since the size of $\mathcal{T}$ increases exponentially with respect to $d$, we can not store the tensor $\mathcal{T}$ any more for large values of $d$. To potentially mitigate this so-called curse of dimensionality [33], we need to identify and exploit underlying structures of $\mathcal{T}$. For instance, when $\mathcal{T}$ is sparse, we only need to store the non-zero entries. When $\mathcal{T}$ is a rank- 1 tensor, i.e.

$$
\begin{equation*}
\mathcal{T}_{i_{1}, i_{2}, \ldots, i_{d}}=u_{i_{1}}^{(1)} u_{i_{2}}^{(2)} \cdots u_{i_{d}}^{(d)}, \quad \text { for } i_{\ell}=1, \ldots, n_{\ell}, \quad \ell=1, \ldots, d \tag{1.1}
\end{equation*}
$$

for some vectors $u^{(\ell)} \in \mathbb{R}^{n_{\ell}}, \ell=1, \ldots, d$, we only need to store these vectors. In applications, the tensor $\mathcal{T}$ rarely has rank- 1 structure, but often $\mathcal{T}$ can be well approximated by the sum of $r$ rank-1 tensors, i.e. by a tensor of rank at most $r$ [190]. The resulting so-called low-rank tensor approximations [145, 198] have been applied to solve a plethora of high-dimensional problems in various fields including but not limited to model order reduction [243], uncertainty quantification [203], sensitivity analysis [201], machine learning [179, 291], signal processing [72], the solution of partial differential [22, 188] and time-dependent differential equations [63, 196], data science [302], optimization [73, 274], quantum physics [256, 331], plasma physics [109] and quantum chemistry [312]. In certain applications, it is beneficial to impose additional structure to the sum of separable functions, leading to the so-called Tucker format [327], the tensor train format [249] or general tensor networks [247]. We would like to point out that it is even possible to apply low-rank tensor approximations to low-dimensional problems by reshaping matrices and
vectors into tensors [187, 248].

## Functional low-rank approximation

The basic ideas of low-rank approximations can be generalized from tensors in finite dimensional tensor product spaces to multivariate functions defined on tensor product domains [154]. Instead of sums of rank-1 tensors, we are now interested in sums of separable functions. A function $f:[-1,1]^{d} \rightarrow \mathbb{R}$ is called separable if it can be written in terms of univariate functions $g_{\ell}:[-1,1] \rightarrow \mathbb{R}$ for $\ell=1, \ldots, d$ as

$$
\begin{equation*}
f\left(x_{1}, x_{2}, \ldots, x_{d}\right)=g_{1}\left(x_{1}\right) g_{2}\left(x_{2}\right) \cdots g_{d}\left(x_{d}\right), \quad \text { for } x_{\ell} \in[-1,1], \ell=1, \ldots, d \tag{1.2}
\end{equation*}
$$

For the sake of simplicity, we only consider the domain $[-1,1]^{d}$ throughout this thesis. Generalizations to arbitrary tensor product domains are possible via appropriate domain transformations. Further extensions to spheres are also viable [43, 321]. It is again possible to prescribe additional structure to the sum of separable functions to obtain approximations in the so-called functional Tucker $[100,163]$ or functional tensor train format [41, 141]. The structure of the functional low-rank approximation can then be exploited in quite a few applications, including the solution of time-dependent partial differential equations (PDEs) [91], uncertainty quantification [202], sensitivity analysis [25], optimal control [142] and quantum dynamic simulations [298].

It is important to emphasize that not every function admits a good approximation in lowrank formats with a moderate amount of terms, especially for larger $d$. Generally speaking, smoothness and a restricted (e.g., nearest neighbor) interaction between variables are helpful. Upper bounds for the number of terms needed to attain a certain accuracy are derived in $[41,146]$ for functions in Sobolev spaces. The required number of terms can change significantly when variables are transformed [324] or the order of variables is permuted [92]. For functions in periodic and mixed Sobolev spaces approximation rates can be found in [286] and [148], respectively. Much faster rates can be obtained for functions with special structures such as compositional functions [21] and quantities associated with certain parametric PDEs [20]. A more abstract analysis of functional low-rank approximation can be found in $[4,5,6]$.

In practice, we obtain a fully discrete approximation of $f$ via discretizing each univariate function in the approximation by, e.g. a truncated series expansion, which incurs an additional truncation error [141]. The resulting approximations are intricately related to approximations of $f$ using a truncated expansion in terms of tensorized basis functions of the form

$$
\begin{equation*}
f\left(x_{1}, x_{2}, \ldots, x_{d}\right) \approx \sum_{i_{1}=1}^{n_{1}} \sum_{i_{2}=1}^{n_{2}} \cdots \sum_{i_{d}=1}^{n_{d}} \mathcal{A}_{i_{1}, i_{2}, \ldots, i_{d}} b_{i_{1}}^{(1)}\left(x_{1}\right) b_{i_{2}}^{(2)}\left(x_{2}\right) \cdots b_{i_{d}}^{(d)}\left(x_{d}\right) \tag{1.3}
\end{equation*}
$$

with coefficient tensor $\mathcal{A} \in \mathbb{R}^{n_{1} \times n_{2} \times \ldots n_{d}}$ of order $d$ and univariate basis functions $b_{i}^{(\ell)}:[-1,1] \rightarrow \mathbb{R}$ for $i=1, \ldots, n_{\ell}, \ell=1, \ldots, d$. Any low-rank approximation of $\mathcal{A}$ can be seen as functional low-rank approximation of $f$. For instance, the Chebfun package [104] for computing numerically on the level of functions [260] approximates bi- and trivariate functions $[163,320]$ by computing low-rank approximations of the coefficient tensor obtained via multivariate interpolation [209] with tensorized Chebyshev polynomial basis functions [223]. The univariate functions in the resulting functional low-rank approximation are internally represented in terms of univariate Chebyshev interpolants [38, 165]. Once the functional low-rank approximation is obtained, operations such as computing derivatives are performed by manipulating the Chebyshev coefficients of the univariate interpolants [27]. Note that the interpolation approach in Chebfun requires that we have access to point evaluation on a grid. When only point evaluations at unstructured points are available, we can use supervised learning to obtain functional low-rank approximations [140].

Instead of computing low-rank approximations of $\mathcal{A}$, we can also try to choose the basis functions $b_{i}^{\ell}$ in (1.3) such that most of the information about the function is stored in a low number of entries of $\mathcal{A}$, i.e. we can approximate $\mathcal{A}$ by a sparse tensor. This is the idea of sparse girds [295, 345]. Especially for functions with bounded mixed derivatives, the sparse grid approach offers the potential to mitigate the curse of dimensionality [50].

We briefly want to point out that there exist many other techniques that offer the potential to mitigate or even overcome the curse of dimensionality when approximating a high-dimensional function $f:[-1,1]^{d} \rightarrow \mathbb{R}$. For instance, we could decompose a multivariate function using the ANOVA-decomposition [297], which represents $f$ as truncated sum of functions depending on small subsets of variables [296], or using arithmetic circuit tensor networks [254], which describe the multivariate function as the interaction of univariate functions in terms of a tensor network. Approximations of $f$ obtained using neural networks can implicitly exploit certain structures of $f$ to overcome the curse of dimensionality [19, 262]. Dimensionality reductions techniques such as active subspaces [76], inverse regression [212] and autoencoders [167] can be used to identify whether $f$ actually depends on all variables or whether it is mostly determined by a lowdimensional subspace of $[-1,1]^{d}$. In the latter case, we can replace the high-dimensional $f$ by a low-dimensional surrogate defined directly on the subspace. However, none of these alternative approaches yields an approximations in terms of sums of univariate functions, which would be highly beneficial for further numerical computations [39, 40].

## Multi-marginal optimal transport

Classical optimal transport minimizes the transport cost between $d=2$ probability measures [34, 333]. For discrete measures, this problem can be expressed as the minimization of $\langle C, P\rangle$ for a nonnegative cost matrix $C$. The so-called transport plan $P$ is a nonnegative
matrix that has to satisfy marginal constraints, i.e. the column and row sums of $P$ are prescribed by discrete measures. The pioneering work of Cuturi [82] established a relation between entropy regularized optimal transport and matrix scaling of $\exp (-C / \eta)$, where $\exp$ denotes the elementwise exponential and $\eta>0$ denotes the regularization parameter. Scaling the rows and columns of $\exp (-C / \eta)$ such that marginal constraints are satisfied can be achieved numerically using the Sinkhorn algorithm [294], whose convergence speed can be accelerated using greedy coordinate descent [9, 216], overrelaxation [314] or accelerated gradient descent [106]. This matrix scaling approach allows one to solve much larger optimal transport problems compared to previous attempts based on solving the original linear program, which in turn has impacted various fields including image processing [121, 267], data science [258], engineering [235] and machine learning [134, 199].

The classical optimal transport problem can be generalized to a multi-marginal setting, in which the transport cost between $d \geq 3$ measures is minimized [253]. Such multi-marginal problems arise in the areas of density functional theory [97], generalized incompressible flow [36], neural networks [57], signal processing [110] and Wasserstein barycenters [59]. For discrete measures, the problem can be expressed as finding the nonnegative transport plan tensor $\mathcal{P}$ of order $d$, which minimizes $\langle\mathcal{C}, \mathcal{P}\rangle$ subject to marginal constraints, where $\mathcal{C}$ denotes a given nonnegative cost tensor of order $d$. In analogy to the matrix case, after adding entropic regularization, the problem can equivalently be transformed into a tensor scaling problem for the tensor $\exp (-\mathcal{C} / \eta)$ [37]. The Sinkhorn algorithm can be generalized to solve this multi-marginal problem [58]; acceleration techniques via greedy coordinate descent are described in [127, 217].

The multi-marginal Sinkhorn algorithm crucially relies on the repeated evaluation of marginals of the rescaled tensor $\exp (-\mathcal{C} / \eta)$. The cost of computing such a marginal increases exponentially with respect to $d$. There exist two different approaches in literature to mitigate this curse of dimensionality by exploiting additional structure of $\exp (-\mathcal{C} / \eta)$. The first line of research is to exploit that, in many applications, the structure of $\exp (-\mathcal{C} / \eta)$ allows to specify the transport plan in terms of a graphical model. When this graphical model does not contain circles, marginals can be computed efficiently using the belief propagation algorithm [118, 150] and Fourier-based fast summation [17]. In particular, this includes tree structured cost tensors [32, 151]. When the model contains circles, the junction tree algorithm [173] can be used to evaluate the marginals, but it might still incur a large computational cost. A second approach to possibly attain a complexity reduction is to replace $\exp (-\mathcal{C} / \eta)$ by a low-rank approximation, whose marginals can be evaluated efficiently [8]. For the classical case $d=2$, Altschuler et al. [10] analyze the impact of the approximation error on the solution returned by the Sinkhorn algorithm. For the case $d \geq 3$, asymptotic complexity bounds are derived in [7] for the specific case that $\mathcal{C}$ has low tensor rank [190] and is given explicitly in factored form. Let us point out that low-rank approximations of $\exp (-\mathcal{C} / \eta)$ should not be confused with low-rank approximations of the desired transport plan, as proposed in [283, 284]. Further, we
would like to point out that it is also possible to exploit the structure of $\mathcal{C}$ [348] and sparsity of the solution $[129,130]$ to efficiently compute solutions of the unregularized multi-marginal optimal transport problem.

## Self-diffusion matrix

The self-diffusion matrix of a tagged particle process on a grid [200] identifies the hydrodynamic limit of multi-species symmetric exclusion processes [266]. It is traditionally computed by approximating long-time averages of the mean-square displacement of the tagged particle [42, 89, 193, 266]. In a first step, the computational domain is truncated to a finite-size supercell with periodic boundary conditions. Then, a large number of realisations of the trajectories of the tagged particle are sampled in order to approximate the mean-square displacement by an empirical average computed with a standard MonteCarlo approach [158, 227]. Moreover, the value of a finite final time has to be chosen beforehand to approximate the long-time limit. In [210], it is shown that the error linked to the truncation of the computational domain decays exponentially with the size of the supercell. In contrast, the statistical error of the approximation linked to the use of a finite number of random samples of the trajectories of the tagged particle decays as the inverse of the square root of the number of samples. As a consequence, the main source of error in the practical computation of approximations of the self-diffusion matrix is due to statistical noise; see e.g. [95, 120, 159, 279, 315, 317].

The self-diffusion matrix can equivalently be defined via the solution of a deterministic high-dimensional optimization problem [42, 210, 266] of the form

$$
\min _{\mathcal{T} \in \mathbb{R}^{2 \times 2 \times \ldots \times 2}} f(\mathcal{T}),
$$

where $f: \mathbb{R}^{2 \times 2 \times \cdots \times 2} \rightarrow \mathbb{R}$ denotes the objective function and the order $d$ of $\mathcal{T}$ is equal to the number of grid points in the finite-size supercell. Solving this minimization problem potentially yields the self-diffusion matrix without incurring any statistical noise. In order to solve this minimization problem numerically, we need to reduce the number of degrees of freedom. This can be achieved by assuming that $\mathcal{T}$ can be well approximated by a sum of rank- 1 tensors. We can then iteratively optimize each of the terms in the low-rank representation of $\mathcal{T}$ using a classical alternating linear scheme [171, 274]. Alternatively, we could use Riemannian optimization [44, 205, 264] to solve the optimization problem directly on the manifold of low-rank tensors directly.

## Organisation of the thesis

Chapter 2. We recall the basic concepts of low-rank tensor approximations, low-rank functions and tensorized Chebyshev interpolation and establish the notation for the remainder of this thesis.

Chapters 3 and 4 (based on [100] and [306]). We develop two novel algorithms for computing functional low-rank approximations by combining tensorized Chebyshev interpolation with low-rank approximations of the coefficient tensor in the Tucker and extended tensor train format. Compared to existing algorithms our novel approaches require fewer function evaluations to achieve the same approximation accuracy.

Chapter 5 (based on [305]). We propose to combine graphical models and low-rank approximations to further accelerate the computation of marginals for the Sinkhorn algorithm. Our numerical experiments demonstrate that this combination offers the potential to solve of entropically regularized optimal transport problems even faster.

Chapter 6 (based on [85] and [84]). We develop a novel approach to compute selfdiffusion matrices numerically by solving the high-dimensional minimization problem using an alternating linear scheme instead of sampling long-time averages of the mean-square displacement.

Chapter 7 (based on [304]). We develop a global spectral method to solve threedimensional PDEs using discretizations of the form (1.3).

Chapter 8. We conclude the thesis and provide an outlook on further research directions. In particular, we discuss the potential to combine the spectral method presented in Chapter 7 with the low-rank approximations of the coefficient tensor in Chapters 3 and 4.

## 2 Preliminaries

In this chapter, we recall the notation and basic results for low-rank approximations of matrices, tensors and functions as well as the fundamental concepts for tensorized Chebyshev interpolation.

In Section 2.1, we introduce low-rank matrices and tensors. We establish a notation for low-rank approximations of matrices and tensors in various formats following the book by Golub and Van Loan [137] and the article by Kolda and Bader [198].

In Section 2.2, we generalize the concept of low-rank approximations from tensors to multivariate functions and introduce the corresponding functional low-rank approximation formats.

In Section 2.3, we recall the fundamentals of Chebyshev interpolation for univariate functions from the book by Mason and Handscomb [224] and the book by Trefethen [322]. We extend the concept to multivariate functions using tensorization. The resulting interpolants are computed based on a tensor containing the evaluation of the function on a tensor product grid. We study how approximating this tensor affects the interpolation error and how such approximation are related to functional low-rank approximations.

### 2.1 Low-rank approximation

### 2.1.1 Low-rank approximation of matrices

A matrix $A \in \mathbb{R}^{m \times n}$ is said to be rank-1, when it can be written as $M=u \circ v$ for vectors $u \in \mathbb{R}^{m}, v \in \mathbb{R}^{n}$, where the outer product $\circ$ is defined as

$$
(u \circ v)_{i, j}=u_{i} v_{j}, \quad \text { for } i=1, \ldots, m, j=1, \ldots, n .
$$

In general, the rank $r$ of a matrix $A \in \mathbb{R}^{m \times n}$ is defined as the minimum $r \in \mathbb{N}$ such that $A$ can be written as a sum of $r$ rank-1 matrices. In the following, we focus on the approximation of given matrix $A \in \mathbb{R}^{n \times m}$ by a matrix $B \in \mathbb{R}^{n \times m}$ subject to the constraint that the rank of $B$ is at most $r$ for a fixed $r \in \mathbb{N}$. The matrix $B$ is called low-rank approximation of $A$.

Since $\operatorname{rank}(B) \leq r$, there exist matrices $U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{n \times r}$ such that $B=U V^{T}$. Assume that $r \ll \min (n, m)$. We observe that storing the matrices $U, V$ of the lowrank approximation $B$ requires only $\mathcal{O}(r(m+n))$ storage, whereas the full matrix $A$ requires $\mathcal{O}(m n)$ storage. By replacing the matrix $A$ by the low-rank approximation $B$ in subsequent computations, we can potentially accelerate these computations by exploiting the structure of $B$. For instance, matrix vector products involving $B$ can be performed in $\mathcal{O}(r(m+n))$ operations when $B$ is given in terms of the matrices $U, V$, whereas matrix vector productions involving $A$ require $\mathcal{O}(m n)$ operations. Note that this introduces an error depending on how well $A$ is approximated by $B$.

Let $\|\cdot\|_{2}$ denote the spectral norm and let $\|\cdot\|_{F}$ denote the Frobenius norm. The Eckhart-Young-Mirsky theorem [231] states that an optimal rank- $r$ matrix $B$ in the sense that $\|A-B\|_{2}$ or $\|A-B\|_{F}$ is minimized can be computed by truncating the singular value decomposition of $A$. Finding this so-called best approximation with respect to the maximum norm $\|\cdot\|_{\infty}$ or the entrywise $\ell_{1}$-norm $\|\cdot\|_{1}$ is still subject to research [299, 343]. Note that we do not necessarily need to compute the best rank- $r$ approximation to accelerate subsequent computations. It suffices to find any matrices $U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{n \times r}$ such that $A \approx U V^{T}$. In the following, we present how to compute such low-rank approximations of $A$ efficiently. The presented approximations can be computed faster than the truncated singular value decomposition and can potentially be obtained without evaluating every single entry of $A$. For a detailed overview of low-rank approximation algorithms for matrices, we refer to the review [194]. Further, we would like to point out that it is even possible to compute low-rank approximations when only selected entries of $A$ can be accessed [55, 350].

We now introduce so-called cross approximations [138]. These are rank-r approximations of the form

$$
\begin{equation*}
A \approx A(:, J) A(I, J)^{-1} A(I,:), \tag{2.1}
\end{equation*}
$$

where $I \subset\{1, \ldots, m\}, J \subset\{1, \ldots, n\}$ denote index sets of cardinality $r$ chosen such that $A(I, J)$ is invertible. This approximation format is visualized in Figure 2.1. Throughout this thesis, we use adaptive cross approximation (ACA) [30] to determine suitable index sets. Initially, ACA determines the index tuple $(i, j)$ of the entry with maximal absolute value of $A$ and adds $i$ to the initially empty index set $I$ and $j$ to $J$ respectively. This greedy approach is then repeated iteratively by determining the largest absolute value of the residual of the current cross approximation. The algorithm is stopped once the


Figure 2.1 - Visualization of a cross approximation (2.1). The columns $A(:, J)$ are marked in blue. The rows $A(I,:)$ are marked in red.
largest absolute value of the residual falls below a prescribed tolerance. This adaptively determines the suitable rank-r. We formalize ACA in Algorithm 1. Note that ACA corresponds to greedily maximizing the volume of $A(I, J)[138]$ and comes with theoretical guarantees [80]. An overview of alternative algorithms to determine the index sets is given in [78].

```
Algorithm 1 Adaptive cross approximation
    Input: matrix \(A \in \mathbb{R}^{m \times n}\), tolerance \(\varepsilon\)
    Output: index sets \(I\) and \(J\) s.t. \(A \approx A(:, J) A(I, J)^{-1} A(I,:)\)
    \(I=[], J=[]\)
    while \(\max |A| \geq \varepsilon\)
        \((i, j)=\arg \max |A(i, j)|\)
            \((i, j)\)
        \(I=[I, i], J=[J, j]\)
        \(A=A-A(:, j) A(i,:) / A(i, j)\)
```

Remark 2.1. In Algorithm 1, we present ACA with full pivoting, i.e. we determine the next index tuple based on the largest absolute value in line 4. This requires the evaluation of every single entry of $A$. In practice, we could use any entry with a large absolute value to determine the next index tuple. A commonly used heuristic to determine entries with large absolute value is partial pivoting [29]. An initially random index tuple $(i, j)$ is updated repeatedly by alternatingly setting $i=\arg \max _{i}|A(i, j)|$ and $j=\arg \max _{j}|A(i, j)|$. This process only requires the evaluation of several rows and columns instead of the whole matrix.

Cross approximations can be generalized to so-called CUR-decompositions [45] by allowing for an arbitrary matrix $\tilde{U} \in \mathbb{R}^{r \times r}$ instead of the inverse of $A(I, J)$. This leads to an approximation of the form

$$
\begin{equation*}
A \approx A(:, J) \tilde{U} A(I,:) \tag{2.2}
\end{equation*}
$$

Given the matrix $C=A(:, J)$ containing columns of $A$ and the matrix $R=A(I,:)$ containing rows of $A$, we can project the columns of $A$ onto the span of $C$ and the rows of $A$ onto the span of $R$ using orthogonal projections. This yields $A \approx C C^{T} A R^{T} R$, which corresponds to the CUR decomposition (2.2) with $\tilde{U}=C^{T} A R^{T}$. Note that computing $C^{T} A R^{T}$ requires the evaluation of the full matrix $A$.

In the following, we describe how we can replace the orthogonal projections by oblique projections to obtain a low-rank approximation of $A$ given $C, R$ without fully evaluating $A$. Let $Q_{C}, Q_{R}$ denote the orthogonal matrices in the economic QR decomposition of $C, R$ respectively. If we use oblique projections [300] based on index sets $\tilde{I} \subset\{1, \ldots, m\}, \tilde{J} \subset$ $\{1, \ldots, n\}$ of cardinality $r$ chosen such that $Q_{C}(\tilde{I},:)$ and $Q_{R}(\tilde{J},:)$ are invertible, we obtain

$$
\begin{align*}
A & \approx Q_{C}\left(\Phi_{\tilde{I}}^{T} Q_{C}\right)^{-1} \Phi_{\tilde{\tilde{I}}}^{T} A\left(Q_{R}\left(\Phi_{\tilde{J}}^{T} Q_{R}\right)^{-1} \Phi_{\tilde{\tilde{J}}}^{T}\right)^{T}  \tag{2.3}\\
& =Q_{C}\left(Q _ { C } ( \tilde { I } , : ) ^ { - 1 } A ( \tilde { I } , \tilde { J } ) \left(Q_{R}(\tilde{J},:)^{-T} Q_{R},\right.\right.
\end{align*}
$$

where $\Phi_{\tilde{I}}^{T} A=A(\tilde{I},:)$. Note that (2.3) is a rank- $r$ approximation of $A$, which given $C, R$ can be computed by only evaluating $r^{2}$ entries of $A$. The following lemma summarizes the statement of [66, Lemma 7.3], which plays a critical role in guiding the choice of the indices $\tilde{I}, \tilde{J}$.

Lemma 2.1. Let $A \in \mathbb{R}^{m \times n}, m \geq n$, have orthonormal columns. Consider an index set $I \subset\{1, \ldots, m\}$ of cardinality $n$ such that $\Phi_{I}^{T} A$ is invertible. Then the oblique projection $A\left(\Phi_{I}^{T} A\right)^{-1} \Phi_{I}^{T}$ satisfies

$$
\left\|x-A\left(\Phi_{I}^{T} A\right)^{-1} \Phi_{I}^{T} x\right\|_{2} \leq\left\|\left(\Phi_{I}^{T} A\right)^{-1}\right\|_{2} \cdot\left\|\left(I-A A^{T}\right) x\right\|_{2}, \quad \forall x \in \mathbb{R}^{n}
$$

Lemma 2.1 exhibits the critical role played by the quantity $\left\|\left(\Phi_{I}^{T} Q_{C}\right)^{-1}\right\|_{2} \geq 1$ for oblique projections. Suitable index sets $\tilde{I}, \tilde{J}$ given $Q_{C}, Q_{R}$ can be computed using the discrete empirical interpolation method (DEIM) [65, 66], presented in Algorithm 2. In practice, the computed index sets usually yield good approximations as $\left\|\left(\Phi_{I}^{T} Q_{C}\right)^{-1}\right\|_{2}$ tends to be small. Note that we could in principle apply oblique projections using $C, R$ directly instead of $Q_{C}, Q_{R}$, but ( $\left.\Phi_{I}^{T} C\right)^{-1}$ is potentially ill-conditioned.

```
Algorithm 2 Discrete empirical interpolation method
    Input: orthonormal matrix \(M \in \mathbb{R}^{m \times n}\)
    Output: index set \(I\)
    \(I=[\operatorname{argmax}|A(:, 1)|]\)
    for \(k=2, \ldots, n\)
        \(c=A(I, 1: k-1) \backslash A(I, k)\)
        \(r=A(:, k)-A(:, 1: k-1) c\)
        \(I=[I, \operatorname{argmax}|r|]\)
```

Remark 2.2. Note that not all matrices can be well approximated by low-rank matrices. It might, however, be possible to approximate their off-diagonal blocks by low-rank matrices. Storing these approximations instead of the full blocks reduces the required storage. The resulting so-called hierarchical matrices [31, 155] are, in particular, useful for the approximation of invertible matrices, since solutions of linear systems involving hierarchical matrices can be computed efficiently.

### 2.1.2 Low-rank approximation of tensors

Let $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{d}}$ denote a rank-1 tensor represented by vectors $u^{(\ell)} \in \mathbb{R}^{n_{\ell}}, \ell=$ $1, \ldots, n$ as in (1.1). Using the multiple outer product, we can write this tensor as $\mathcal{T}=u^{(1)} \circ u^{(2)} \circ \cdots \circ u^{(d)}$. In general, the tensor-rank $r$ of the tensor $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{d}}$ is defined as the minimum $r \in \mathbb{N}$ such that $\mathcal{T}$ can be written as sum of $r$ rank-1 tensors [168]. A tensor $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times \ldots n_{d}}$ with rank at most $r$ is said to be represented in CANDECOMP/PARAFAC (CP)-format [190] if it is given in terms of vectors $u^{(\ell, k)} \in \mathbb{R}^{n_{\ell}}$ for $\ell=1, \ldots, d, k=1, \ldots, r$ as

$$
\begin{equation*}
\mathcal{T}=\sum_{k=1}^{r} u^{(1, k)} \circ u^{(2, k)} \circ \cdots \circ u^{(d, k)} . \tag{2.4}
\end{equation*}
$$

Note that storing these vectors only requires $\mathcal{O}(d r n)$ storage, whereas storing the entries of $\mathcal{T}$ directly requires $\mathcal{O}\left(n^{d}\right)$ storage, where we assume that $n_{1}=n_{2}=\cdots=n_{d}=n$.

In this work, we are primarily interested in using low-rank approximations of tensors to accelerate computations. As in the matrix case discussed in Section 2.1.1, we do not need to solve the hard problem [252, 292] of finding the best rank-r approximation with respect to a tensor norm such as the Frobenius norm $\|\cdot\|_{F}$, the entrywise $\ell_{1}$-norm $\|\cdot\|_{1}$ or the maximum norm $\|\cdot\|_{\infty}$. Instead any reasonably accurate low-rank approximation suffices.

The alternating least squares (ALS) algorithm [60, 161] is a commonly used alternating linear scheme [39, 145, 250, 274] to compute approximations in the CP-format (2.4). Let $\tilde{\ell} \in\{1, \ldots, d\}$. We observe that the problem of minimizing the approximation error

$$
\left\|\mathcal{T}-\sum_{k=1}^{r} u^{(1, k)} \circ u^{(2, k)} \circ \cdots \circ u^{(d, k)}\right\|_{F}
$$

is a linear least squares problem when we fix $u^{(\ell, k)}$ for $k=1, \ldots, n, \ell=1, \ldots, d, \ell \neq \tilde{\ell}$ and only optimize with respect to $u^{(\tilde{\ell}, 1)}, \ldots, u^{(\tilde{\ell}, r)}$. This allows us to improve a given approximation by updating $u^{(\tilde{\ell}, 1)}, \ldots, u^{(\tilde{\ell}, r)}$ based on the solution of this least squares problem. Performing such updates repeatedly for alternating values of $\tilde{\ell}$ leads to the ALS algorithm. Note that applying ALS to a randomly initialized approximation might not converge [198]. Still, this approach yields useful low-rank approximations in practice.

Remark 2.3. The ALS algorithm requires a priori knowledge of a suitable rank-r for the CP approximation (2.4). Note that the sum of two tensors in CP-format can again be represented in CP-format. This allows us to refine a given approximation by adding an approximation of the residual of the current approximation in CP-format. Performing this process successively leads to a rank-adaptive approximation algorithm [11, 145, 242].

Remark 2.4. We would like to point out that additional constraints can be imposed on the vectors $u^{(\ell, k)}$ in (2.4). For smoothness, nonnegativity and sparseness constraints we refer to [131]. For imposed linear dependencies we refer to the overview [119].

### 2.1.3 Low-rank tensor formats with additional structure

Prescribing additional structure to the sum (2.4) leads to different approximation formats. In the following, we introduce two commonly used formats.

Let $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{d}}$. A vector $v \in \mathbb{R}^{n_{\ell}}$ is called mode- $\ell$ fiber of $\mathcal{T}$ if there exist indices $i_{k} \in\left\{1, \ldots, n_{d}\right\}$ for $k=1, \ldots, d, k \neq \ell$ such that $\mathcal{T}_{i_{1}, \ldots, i_{\ell}, \ldots, i_{d}}=v_{i_{\ell}}$ for all $i_{\ell}=1, \ldots, n_{\ell}$. We denote by $\mathcal{T}^{\{\ell\}} \in \mathbb{R}^{n_{\ell} \times\left(n_{1} \cdots n_{\ell-1} \cdot n_{\ell+1} \cdots n_{d}\right)}$ the matrix containing all mode- $\ell$ fibers of $\mathcal{T}$ in its columns, i.e.

$$
\mathcal{T}_{i_{\ell}, i_{\neq \ell}}^{\{\ell\}}=\mathcal{T}_{i_{1}, i_{2}, \ldots, i_{d}}, \quad \text { where } i_{\neq \ell}=1+\sum_{\substack{j=1 \\ j \neq \ell}}^{d}\left(\left(i_{j}-1\right) \prod_{\substack{k=1 \\ k \neq \ell}}^{j-1} n_{k}\right),
$$

for all $i_{\ell}=1, \ldots, n_{\ell}, \ell=1, \ldots, d$. This matrix is known as mode- $\ell$ matricization of $\mathcal{T}$. The multilinear rank $\left(r_{1}, r_{2}, \ldots, r_{d}\right) \in \mathbb{N}^{d}$ of $\mathcal{T}$ is defined based on these matricizations as $r_{\ell}=\operatorname{rank}\left(\mathcal{T}^{\{\ell\}}\right)$ for $\ell=1, \ldots, d$.

A tensor $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{d}}$ with multilinear rank at most $\left(r_{1}, r_{2}, \ldots, r_{d}\right)$ is said to be represented in Tucker format [327] if it is given in terms of a so-called core tensor $\mathcal{C} \in \mathbb{R}^{r_{1} \times r_{2} \times \cdots \times r_{d}}$ and so-called factor matrices $U^{(\ell)} \in \mathbb{R}^{n_{\ell} \times r_{\ell}}$ for $\ell=1, \ldots, d$ as

$$
\mathcal{T}_{i_{1}, i_{2}, \cdots, i_{d}}=\sum_{k_{1}=1}^{r_{1}} \sum_{k_{2}=1}^{r_{2}} \ldots \sum_{k_{d}=1}^{r_{d}} \mathcal{C}_{k_{1}, k_{2}, \ldots, k_{d}} U_{i_{1}, k_{1}}^{(1)} U_{i_{2}, k_{2}}^{(2)} \cdots U_{i_{d}, k_{d}}^{(d)},
$$

for $i_{\ell}=1, \ldots, n_{\ell}, \ell=1, \ldots, d$. This can be written as

$$
\begin{equation*}
\mathcal{T}=\mathcal{C} \times_{1} U^{(1)} \times_{2} U^{(2)} \times_{3} \cdots \times_{d} U^{(d)}, \tag{2.5}
\end{equation*}
$$

where $\mathrm{x}_{\ell}$ denotes the mode- $\ell$ multiplication. For a tensor $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{d}}$ and a matrix $M \in \mathbb{R}^{m \times n_{\ell}}$ it is defined as the multiplication of every mode- $\ell$ fiber of $\mathcal{T}$ with $M$, i.e.

$$
(\mathcal{T} \times \neq M)_{i_{1}, \ldots, i_{d}}=\sum_{j_{\ell}=0}^{n_{\ell}} \mathcal{T}_{i_{1}, \ldots, i_{\ell-1}, j_{\ell}, i_{\ell+1}, \cdots, i_{d}} M_{i_{\ell}, j_{\ell}}, \quad \text { for } i_{\ell}=1, \ldots, n_{\ell}, \ell=1, \ldots, d .
$$

This is related to the matrizisation in the following way $\left(\mathcal{T} \times_{\ell} M\right)^{\{\ell\}}=M \mathcal{T}{ }^{\{\ell\}}$. Storing a tensor in Tucker format requires $\mathcal{O}\left(r^{d}+d n r\right)$ storage under the assumption that $n_{1}=n_{2}=\cdots=n_{d}=n$ and $r_{1}=r_{2}=\cdots=r_{d}=r$.

Remark 2.5. A special type of Tucker decompositions are so-called tensor CUR decompositions [53], which generalize matrix CUR decompositions (2.2) to tensors by requiring that the columns of the matrices $\mathcal{U}^{(\ell)}$ contain mode- $\ell$ fibers of the tensor $\mathcal{T}$. It is also possible to impose nonnegativity constraints on either the whole Tucker approximation [180, 309] or the individual factor matrices [191, 349] when $\mathcal{T}$ is nonnegative.

The second low-rank approximation format of interest are so-called tensor train (TT) decompositions [249]. These are also known as matrix product states [256, 337] in the Physics community.

Let $\mathcal{T}^{<\ell>}$ denote the tensor obtained by reshaping $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{d}}$ into $\mathbb{R}^{\left(n_{1} \cdot n_{2} \cdots n_{\ell}\right) \times\left(n_{\ell+1} \cdots n_{d}\right)}$ for $\ell=1, \ldots, d-1$. The TT ranks $\left(R_{1}, \ldots, R_{d-1}\right) \in \mathbb{N}^{d-1}$ of $\mathcal{T}$ are defined as $R_{\ell}=\operatorname{rank}\left(\mathcal{T}^{<\ell>}\right)$ for $\ell=1, \ldots, d-1[170]$. A tensor $\mathcal{T}$ with TT ranks at most ( $R_{1}, \ldots, R_{d-1}$ ) is said to be represented in the TT format when it is given in terms of so-called TT cores $\mathcal{G}^{(1)} \in \mathbb{R}^{n_{1} \times R_{1}}, \mathcal{G}^{(d)} \in \mathbb{R}^{R_{d-1} \times n_{d}}$ and $\mathcal{G}^{(\ell)} \in \mathbb{R}^{R_{\ell-1} \times n_{\ell} \times R_{\ell}}$ for $\ell=2, \ldots, d-1$ as

$$
\begin{equation*}
\mathcal{T}_{i_{1}, \ldots, i_{d}}=\sum_{k_{1}=1}^{R_{1}} \sum_{k_{1}=1}^{R_{1}} \ldots \sum_{k_{d-1}=1}^{R_{d-1}} \mathcal{G}_{i_{1}, k_{1}}^{(1)} \mathcal{G}_{k_{1}, i_{1}, k_{2}}^{(2)} \cdots \mathcal{G}_{k_{d-2}, i_{d-1}, k_{d-1}}^{(d-1)} \mathcal{G}_{k_{d-1}, i_{d}}^{(d)}, \tag{2.6}
\end{equation*}
$$

for $i_{\ell}=1, \ldots, n_{\ell}, \ell=1, \ldots, d$. Storing the TT cores requires $\mathcal{O}\left((d-2) n R^{2}+2 n R\right)$ storage under the assumption that $n_{1}=n_{2}=\cdots=n_{d}=n$ and $R_{1}=R_{2}=\cdots=R_{d-1}=R$. This results in linear (instead of exponential) growth with respect $d$ under the (strong) assumption that $R$ remains constant as $d$ increases.

Remark 2.6. Any tensor with multilinear ranks $\left(r_{1}, \ldots, r_{n}\right)$ can be represented in the Tucker format (2.5). However, not every tensor that can be represented in Tucker format (2.5) with core tensor $\mathcal{C} \in \mathbb{R}^{r_{1} \times \cdots \times r_{d}}$ has multilinear ranks $\left(r_{1}, \ldots, r_{n}\right)$. The same applies for the TT ranks $\left(R_{1}, \ldots, R_{d-1}\right)$ and the TT format (2.6).

Remark 2.7. There exist many other low-rank tensor formats in literature including the hierarchical Tucker format [152], two-level approximations [186] and tensor networks [73, 247, 255]. For an in-depth overview we refer to the book by Hackbusch [154].

### 2.1.4 Approximation algorithms

A classical algorithm to compute approximations of a given tensor $\mathcal{T}$ in the Tucker format (2.5) is the so-called higher order SVD (HOSVD) [88]. Its main idea is to compute the truncated SVDs of the matricizations $\mathcal{T}^{(\ell)}$ for $\ell=1, \ldots, d$. The columns of the factor matrices $U^{(\ell)}$ are set to the leading left singular vectors of the corresponding SVD. Using orthogonal projections we obtain the core tensor analogous to the CUR decomposition (2.2) as

$$
\begin{equation*}
\mathcal{C}=\mathcal{T} \times 1\left(U^{(1)}\right)^{T} \times_{2}\left(U^{(2)}\right)^{T} \times \cdots \times\left(U^{(d)}\right)^{T} . \tag{2.7}
\end{equation*}
$$

Approximations in the TT format (2.6) can be computed using the so-called TT-SVD [249] by applying a several SVDs to $\mathcal{T}$. However, both the HOSVD and TT-SVD require the evaluation of the full tensor $\mathcal{T}$ to compute the approximation. This might not be feasible for large tensors. We refer to the survey [145] for an overview of algorithms to compute low-rank approximations of $\mathcal{T}$. In the following, we introduce how TT approximations can
be computed efficiently using so-called TT-cross approximations [251] without evaluating every single entry of $\mathcal{T}$. For the efficient computation of Tucker approximations we refer to Chapter 3.

Cross approximation (2.1) can be generalized to tensors $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times \cdots \times n_{d}}$ in the following way. Let $\mathcal{I}^{\leq \ell} \subset\left\{\left(i_{1}, \ldots, i_{\ell}\right) \mid 1 \leq i_{k} \leq n_{k}, 1 \leq k \leq \ell\right\}$ and $\mathcal{I}^{>\ell} \subset\left\{\left(i_{\ell+1}, \ldots, i_{d}\right) \mid 1 \leq i_{k} \leq\right.$ $\left.n_{k}, \ell+1 \leq k \leq d\right\}$ denote index sets of cardinality $R_{\ell}$ for $\ell=1, \ldots, d-1$. We obtain the approximation

$$
\begin{equation*}
\mathcal{T}_{i_{1}, \ldots, i_{d}} \approx \sum_{s_{1}=1}^{R_{1}} \sum_{t_{1}=1}^{R_{1}} \cdots \sum_{s_{d-1}=1}^{R_{d-1}} \sum_{t_{d-1}=1}^{R_{d-1}} \prod_{\ell=1}^{d} \mathcal{T}\left(\mathcal{I}_{s_{\ell-1}}^{\leq \ell-1}, i_{\ell}, \mathcal{I}_{t_{\ell}}^{>\ell}\right)\left(T\left(\mathcal{I}_{s_{\ell}}^{\leq \ell}, \mathcal{I}_{t_{\ell}}^{>\ell}\right)\right)^{-1} \tag{2.8}
\end{equation*}
$$

for $i_{\ell}=1, \ldots, n_{\ell}, \ell=1, \ldots, d$, where we use $\mathcal{I}^{\leq 0}=\mathcal{I}^{\leq d}=\mathcal{I}^{>d}=\{ \}$. Note that Equation (2.8) is an approximation in TT format (2.6) [282]. Its cores can be computed by evaluating $\mathcal{O}\left(d n R^{2}\right)$ entries of $\mathcal{T}$ when $n_{1}=\cdots=n_{d}=n, R_{1}=\cdots=R_{d-1}=R$.

In order to derive suitable index sets, we present the rank-adaptive greedy restricted cross interpolation algorithm developed by Savostyanov [282]. We initialize the index sets $\mathcal{I}^{\leq \ell}, \mathcal{I}^{>\ell}$ randomly such that they each contain a single element and such that they are nested in the sense that $\left(i_{1}, \ldots, i_{\ell}\right) \in \mathcal{I}^{\leq \ell} \Rightarrow\left(i_{1}, \ldots, i_{\ell-1}\right) \in \mathcal{I}^{\leq \ell-1}$ and $\left(i_{\ell+1}, \ldots, i_{d}\right) \in$ $\mathcal{I}^{>\ell} \Rightarrow\left(i_{\ell+2}, \ldots, i_{d}\right) \in \mathcal{I}^{>\ell+1}$. We now update these index sets in alternating order until the approximation (2.8) approximates $\mathcal{T}$ well. For this purpose so-called DMRG supercores [281] are formed. These are defined as subtensors $\mathcal{T}\left(\mathcal{I}^{\leq \ell-1},:,:, \mathcal{I}^{>\ell+1}\right)$ which are reshaped into tensors in $\mathbb{R}^{R_{\ell-1} \times n_{\ell} \times n_{\ell+1} \times R_{\ell+1}}$. The rank-adaptive algorithm proposed by Savostyanov [282, Algorithm 2] computes cross approximations (2.1) of the supercores of the form

$$
\mathcal{T}\left(\mathcal{I}^{\leq \ell-1},:,:, \mathcal{I}^{>\ell+1}\right) \approx \sum_{s_{\ell}=1}^{R_{\ell}} \sum_{t_{\ell}=1}^{R_{\ell}} \mathcal{T}\left(\mathcal{I}^{\leq \ell-1},:, \mathcal{J}_{t_{\ell}}^{>\ell}\right)\left(\mathcal{T}\left(\mathcal{J}_{s_{\ell}}^{\leq \ell}, \mathcal{J}_{t_{\ell}}^{>\ell}\right)\right)^{-1} \mathcal{T}\left(\mathcal{J}_{s_{\ell}}^{\leq \ell},:, \mathcal{I}^{>\ell+1}\right),
$$

where the index sets $\mathcal{J} \leq \ell, \mathcal{J} \geq^{\geq \ell}$ of cardinality at most $R_{\ell}+1$ are constructed such that $\mathcal{I}^{\leq \ell} \subset \mathcal{J}^{\leq \ell}$ and $\mathcal{I}^{>\ell} \subset \mathcal{J}^{>\ell}$. We can now enhance the approximation (2.8) by replacing the index sets $\mathcal{I}^{\leq \ell}, \mathcal{I}^{>\ell}$ by $\mathcal{J}^{\leq \ell}, \mathcal{J}^{>\ell}$. In practice, this is achieved by determining the entry with the largest absolute value of the matrix

$$
\begin{equation*}
\left(\mathcal{T}\left(\mathcal{I}^{\leq \ell-1},:,:, \mathcal{I}^{>\ell+1}\right)-\sum_{s_{\ell}=1}^{R_{\ell}} \sum_{t_{\ell}=1}^{R_{\ell}} \mathcal{T}\left(\mathcal{I}^{\leq \ell-1},:, \mathcal{I}_{t_{\ell}}^{>\ell}\right)\left(\mathcal{T}\left(\mathcal{I}_{s_{\ell}}^{\leq \ell}, \mathcal{I}_{t_{\ell}}^{>\ell}\right)\right)^{-1} \mathcal{T}\left(\mathcal{I}_{s_{\ell}}^{\leq \ell},:, \mathcal{I}^{>\ell+1}\right)\right)^{<2>} \tag{2.9}
\end{equation*}
$$

using sampling as described in Remark 2.1. When the absolute value of this entry is above a prescribed tolerance, we obtain $\mathcal{J}^{\leq \ell}, \mathcal{J}^{\geq \ell}$ by adding the index corresponding to this entry to $\mathcal{I} \leq \ell, \mathcal{I} \geq \ell$. Note that this procedure adaptively increases the $\ell$ th entry of the

TT rank of the approximation (2.8). Otherwise we set $\mathcal{J} \leq \ell, \mathcal{J} \geq \ell$ to $\mathcal{I}^{\leq \ell}, \mathcal{I} \geq \ell$, i.e. we do not change the approximation. This process is repeatedly performed for $\ell=1, \ldots, d-1$ and stopped once the error of the approximation (2.8) is sufficiently small at sample points. We formalize this procedure in Algorithm 3.

```
Algorithm 3 TT-cross
    Input: procedure to evaluate entries of the tensor \(\mathcal{T} \in \mathbb{R}^{n_{1} \times \cdots \times n_{d}}\), tolerance \(\varepsilon\)
    Output: TT approximation \((2.6)\) of \(\mathcal{T}\) with core tensors \(\mathcal{H}^{(\ell)}\)
    Initialize \(\mathcal{I}^{\leq \ell}, \mathcal{I}^{>\ell}\) randomly with only one element each for \(\ell=1, \ldots, d-1\)
    while Error of the approximation (2.8) of \(\mathcal{T}\) is larger than \(\varepsilon\) at sample points.
        for \(\ell=1, \cdots, d-1\)
            Find a large entry of the matrix (2.9) and update \(\mathcal{I}^{\leq \ell}, \mathcal{I}^{>\ell}\) as described in
    Section 2.1.4.
    Compute the TT cores \(\mathcal{G}^{(\ell)}\) corresponding to the approximation (2.8).
```


### 2.2 Functional low-rank approximation

The concept of low-rank approximations can be generalized from tensor products of finite dimensional vector spaces to tensor products of infinite dimensional vector spaces [189]. Instead of matrices or tensors we are now interested in multivariate functions defined on the tensor product domain $[-1,1]^{d}$. In this continuous setting, the pendant to rank-1 tensors are separable functions; see Equation (1.2). The rank $r$ of a function $f:[-1,1]^{d} \rightarrow \mathbb{R}$ is the minimal $r \in \mathbb{N} \cup\{\infty\}$ for which $f$ can be written as a sum of separable functions, i.e.

$$
\begin{equation*}
f\left(x_{1}, x_{2}, \ldots, x_{d}\right)=\sum_{k=1}^{r} g_{k}^{(1)}\left(x_{1}\right) g_{k}^{(2)}\left(x_{2}\right) \cdots g_{k}^{(d)}\left(x_{d}\right), \quad \text { for all } x_{\ell} \in[-1,1], \ell=1, \ldots, \ell, \tag{2.10}
\end{equation*}
$$

where $g_{k}^{(\ell)}:[-1,1] \rightarrow \mathbb{R}$ for $k=1, \ldots, r, \ell=1, \ldots, d$.
We are again interested in approximating $f:[-1,1]^{d} \rightarrow \mathbb{R}$ by a function $h:[-1,1]^{d} \rightarrow \mathbb{R}$ subject to the constraint that the rank of $h$ is at most $r$ for a fixed $r \in \mathbb{N}$. The function $h$ is called (functional) low-rank approximation of $f$. Note that the original problem of finding separable decompositions of functions is intimately connected to low-rank decompositions of matrices and tensors [154, Chapter 7]. Throughout this work we typically measure the approximation error using the uniform norm $\|\cdot\|_{\infty}$ or the $L^{p}$-norm $\|\cdot\|_{L^{p}}$ defined on $[-1,1]^{d}$. We want to emphasize that it is reasonable to search for lowrank approximations of $f$ even when the rank of $f$ is infinity. For example, truncating the Schmidt decomposition [285] of an arbitrary function in the Lebesgue space $L^{2}\left([-1,1]^{2}\right)$ yields the best rank- $r$ approximation with respect to the $L^{2}$-norm [147].

Analogous to Section 2.1.3, we can impose additional structure on the sum of separable
functions (2.10). In the following, we generalize the Tucker format (2.5) and the TT format (2.6) from tensors to the continuous setting. This leads to functional low-rank approximations in the so-called functional Tucker [100, 163, 221, 270] and functional TT format [41, 139, 141].

Let $f:[-1,1]^{d} \rightarrow \mathbb{R}$. The multilinear rank $\left(r_{1}, \ldots, r_{d}\right) \in(\mathbb{N} \cup\{\infty\})^{d}$ of $f$ is defined as $r_{\ell}=\operatorname{rank}_{\{\ell\}}(f)$ for $\ell=1, \ldots, d$, where $\operatorname{rank}_{\{\ell\}}(f)$ denotes the minimum $r \in \mathbb{N} \cup\{\infty\}$ such that $f$ can be written as in terms of functions $g_{k}:[-1,1] \rightarrow \mathbb{R}$ and $h_{k}:[-1,1]^{d-1} \rightarrow \mathbb{R}$ for $k=1, \ldots, r$ as

$$
f\left(x_{1}, \ldots, x_{d}\right)=\sum_{k=1}^{r} g_{k}\left(x_{\ell}\right) h_{k}\left(x_{1}, \ldots, x_{\ell-1}, x_{\ell+1}, \ldots, x_{d}\right)
$$

A function $f:[-1,1]^{d}$ with multilinear rank at most $\left(r_{1}, \ldots, r_{d}\right)$ is said to be represented in functional Tucker format when it is given in terms of a so-called a core tensor $\mathcal{C} \in$ $\mathbb{R}^{r_{1} \times r_{2} \times \cdots \times r_{d}}$ and univariate functions $u_{k}^{(\ell)}:[-1,1] \rightarrow \mathbb{R}$ for $k=1, \ldots, r_{\ell}, \ell=1, \ldots, d$ as

$$
\begin{equation*}
f\left(x_{1}, \ldots, x_{d}\right)=\sum_{k_{1}=1}^{r_{1}} \sum_{k_{2}=1}^{r_{2}} \ldots \sum_{k_{d}=1}^{r_{d}} \mathcal{C}_{k_{1}, k_{2}, \ldots, k_{d}} u_{k_{1}}^{(1)}\left(x_{1}\right) u_{k_{2}}^{(2)}\left(x_{2}\right) \cdots u_{k_{d}}^{(d)}\left(x_{d}\right) . \tag{2.11}
\end{equation*}
$$

The TT $\operatorname{rank}\left(R_{1}, \ldots, R_{d-1}\right) \in(\mathbb{N} \cup\{\infty\})^{d-1}$ of $f$ is defined as $R_{\ell}=\operatorname{rank}_{<\ell>}(f)$ for $\ell=1, \ldots, d-1$, where $\operatorname{rank}_{<\ell>}(f)$ denotes the minimum $R \in \mathbb{N} \cup\{\infty\}$ such that $f$ can be written as in terms of functions $g_{k}:[-1,1]^{\ell} \rightarrow \mathbb{R}$ and $h_{k}:[-1,1]^{d-\ell} \rightarrow \mathbb{R}$ for $k=1, \ldots, R$ as

$$
f\left(x_{1}, \ldots, x_{d}\right)=\sum_{k=1}^{R} g_{k}\left(x_{1}, \ldots, x_{\ell}\right) h_{k}\left(x_{\ell+1}, \ldots, x_{d}\right)
$$

A function $f:[-1,1]^{d}$ with TT ranks at most $\left(r_{1}, \ldots, r_{d-1}\right)$ is said to be represented in functional TT format when it is given in terms of univariate functions $g_{\alpha_{1}}^{(1)}:[-1,1] \rightarrow \mathbb{R}$ for $\alpha_{1}=1, \ldots, R_{1}, g_{\alpha_{d-1}}^{(d)}:[-1,1] \rightarrow \mathbb{R}$ for $\alpha_{d-1}=1, \ldots, R_{d-1}$ and $g_{\alpha_{\ell-1}, \alpha_{\ell}}^{(\ell)}:[-1,1] \rightarrow \mathbb{R}$ for $\alpha_{\ell}=1, \ldots, R_{\ell}, \ell=2, \ldots, d-1$ as

$$
\begin{equation*}
f\left(x_{1}, \ldots, x_{d}\right)=\sum_{\alpha_{1}=1}^{R_{1}} \ldots, \sum_{\alpha_{d-1}=1}^{R_{d-1}} g_{\alpha_{1}}^{(1)}\left(x_{1}\right) g_{\alpha_{1}, \alpha_{2}}^{(2)}\left(x_{2}\right) \cdots g_{\alpha_{d-2}, \alpha_{d-1}}^{(d-1)}\left(x_{d-1}\right) g_{\alpha_{d-1}}^{(d)}\left(x_{d}\right) . \tag{2.12}
\end{equation*}
$$

Remark 2.8. We want to emphasize that many other low-rank approximation formats can be generalized analogously from tensors to functions. For instance, generalizations of the hierarchical Tucker format are used in [21, 91, 117, 241, 242, 286].

### 2.3 Chebyshev interpolation

### 2.3.1 Chebyshev interpolation for univariate functions.

In the following, we recall the fundamental ideas of Chebyshev interpolation.
Let $f:[-1,1] \rightarrow \mathbb{R}$ and $n \in \mathbb{N}$. The Chebyshev interpolant $p$ of $f$ with degree $n$ is the unique polynomial of degree $n$ that interpolates $f$ in the Chebyshev points (of the second kind) $x_{k}=\cos (k \pi / n), k=0, \ldots, n$, i.e. $p\left(x_{k}\right)=f\left(x_{k}\right)$ for $k=0, \ldots, n$. Using Chebyshev polynomials $T_{k}(x)=\cos \left(k \cos ^{-1}(x)\right), k=0, \ldots, n[67]$ as basis for the vector space of polynomials of degree at most $n$, we can write the interpolant as

$$
\begin{equation*}
p(x)=\sum_{i=0}^{n} a_{i} T_{i}(x), \tag{2.13}
\end{equation*}
$$

where $a_{i} \in \mathbb{R}, i=0, \ldots, n$ denotes the so-called Chebyshev coefficients. These coefficients are uniquely determined by the interpolation property

$$
f\left(x_{k}\right)=a_{0} T_{0}\left(x_{k}\right)+a_{1} T_{1}\left(x_{k}\right)+\cdots+a_{n} T_{n}\left(x_{k}\right), \quad \text { for } k=0, \ldots, n .
$$

These $n+1$ equations form a linear system, which can be inverted. By using the following property of the Chebyshev polynomials

$$
\frac{1}{2}\left(T_{i}\left(x_{0}\right) T_{j}\left(x_{0}\right)+T_{i}\left(x_{n}\right) T_{j}\left(x_{n}\right)\right)+\sum_{k=1}^{n-1} T_{i}\left(x_{k}\right) T_{j}\left(x_{k}\right)= \begin{cases}0 & i \neq j, \\ \frac{n}{2} & 0<i=j<n \\ n & i=j=0 \text { or } i=j=n\end{cases}
$$

we can construct the matrix $F \in \mathbb{R}^{(n+1) \times(n+1)}$ defined as

$$
F=\frac{2}{n}\left(\begin{array}{ccccc}
\frac{1}{4} T_{0}\left(x_{0}\right) & \frac{1}{2} T_{0}\left(x_{1}\right) & \frac{1}{2} T_{0}\left(x_{2}\right) & \ldots & \frac{1}{4} T_{0}\left(x_{n}\right)  \tag{2.14}\\
\frac{1}{2} T_{1}\left(x_{0}\right) & T_{1}\left(x_{1}\right) & T_{1}\left(x_{2}\right) & \ldots & \frac{1}{2} T_{1}\left(x_{n}\right) \\
\frac{1}{2} T_{2}\left(x_{0}\right) & T_{2}\left(x_{1}\right) & T_{2}\left(x_{2}\right) & \ldots & \frac{1}{2} T_{2}\left(x_{n}\right) \\
\vdots & \vdots & \vdots & \ddots & \vdots \\
\frac{1}{4} T_{n}\left(x_{0}\right) & \frac{1}{2} T_{n}\left(x_{1}\right) & \frac{1}{2} T_{n}\left(x_{2}\right) & \ldots & \frac{1}{4} T_{n}\left(x_{n}\right)
\end{array}\right),
$$

which maps the vector of function evaluations $\left(f\left(x_{0}\right), f\left(x_{1}\right), \ldots, f\left(x_{n}\right)\right)$ onto the vector of Chebyshev coefficients $\left(a_{0}, a_{1}, \ldots, a_{n}\right)$.

The following theorem summarizes the error bounds in [322, Theorem 8.2] and [322, Theorem 7.1]. Note that the error bound decays exponentially with respect to $n$ in part a) and geometrically in part b).

Theorem 2.1. a) Suppose that the continuous function $f:[-1,1] \rightarrow \mathbb{R}$ can be extended to an analytic function $f^{*}$ on $E_{\rho}$ with $\rho>1$, where $E_{\rho}$ denotes the Bernstein ellipse, i.e. the closed ellipse with foci at $\pm 1$ and the sum of major and minor semi-axes equal
to $\rho$. Then the Chebyshev interpolant $p$ of $f$ with degree $n$ satisfies

$$
\|f-p\|_{\infty} \leq \frac{4 \rho^{-n}}{\rho-1} \max _{z \in E_{\rho}}\left|f^{*}(z)\right|
$$

b) Assume that the function $f:[-1,1] \rightarrow \mathbb{R}$ has $(k-1)$ absolutely continuous derivatives and assume that the total variation of its $k$ th derivative is bounded by $V$. Then the Chebyshev interpolant $p$ of $f$ with degree $n>k$ satisfies

$$
\|f-p\|_{\infty} \leq \frac{4 V}{\pi k(n-k)^{k}}
$$

Remark 2.9. Throughout this work, we follow the standard notation established in [240], where the Chebyshev points, polynomials and coefficients are indexed starting from 0. Other vectors, matrices and tensors are indexed starting from 1.

Remark 2.10. Given a function $f:[-1,1] \rightarrow \mathbb{R}$, the cosine transform computes the values

$$
\tilde{a}_{k}=f\left(x_{0}\right) T_{k}\left(x_{0}\right)+f\left(x_{n}\right) T_{k}\left(x_{n}\right)+2 \sum_{i=1}^{n-1} f\left(x_{i}\right) T\left(x_{i}\right), \quad \text { for } k=0, \ldots, n .
$$

This can be written equivalently in terms of a discrete Fourier transform [135] as

$$
\tilde{a}_{k}=\sum_{j=0}^{2 n-1} f\left(\cos \left(\frac{\pi j}{n}\right)\right) \exp \left(\frac{2 \pi i j n}{2 n}\right), \quad \text { for } k=0, \ldots, n,
$$

where $i$ denotes the imaginary unity. Note that $a_{0}=1 / 2 \cdot \tilde{a}_{0}, a_{n}=1 / 2 \cdot \tilde{a}_{n}$ and $a_{k}=\tilde{a}_{k}$ for $k=1, \ldots, n-1$. Thus, we can use the fast Fourier transform to compute the coefficient vector in a fast and stable way.

Remark 2.11. Let $m \in \mathbb{N}$. Note that the Chebyshev points for $n=m$ are nested within the Chebyshev points for $n=2 m$. This implies, that given the function evaluations to compute a degree $m$ interpolant, we only need to evaluate the function at $m$ additional points to compute the degree $2 m$ interpolant. Moreover, there exist heuristics such as chop in Chebfun [16] to estimate whether a given Chebyshev interpolant is an accurate approximation of $f$. We can efficiently determine a sufficiently large degree for an accurate interpolant by initially starting with a small degree and by doubling the degree until the heuristic states that the interpolation is accurate.

### 2.3.2 Tensorized Chebyshev interpolation for multivariate functions

Chebyshev interpolation can be extended to multivariate functions using tensorization [237]. Let $f:[-1,1]^{d} \rightarrow \mathbb{R}$ and $\left(n_{1}, n_{2}, \ldots, n_{d}\right) \in \mathbb{N}^{d}$. The Chebyshev interpolant $p$ of $f$ with degree $\left(n_{1}, n_{2}, \ldots, n_{d}\right)$ is the unique multivariate polynomial of degree
$\left(n_{1}, n_{2}, \ldots, n_{d}\right)$, which satisfies

$$
\begin{equation*}
p\left(x_{i_{1}}^{(1)}, x_{i_{2}}^{(2)}, \ldots, x_{i_{d}}^{(d)}\right)=f\left(x_{i_{1}}^{(1)}, x_{i_{2}}^{(2)}, \ldots, x_{i_{d}}^{(d)}\right), \quad \text { for } i_{\ell}=0, \ldots, n_{\ell}, \ell=1, \ldots, d \tag{2.15}
\end{equation*}
$$

where the Chebyshev points are denoted by $x_{k}^{(\ell)}=\cos \left(k \pi / n_{\ell}\right), k=0, \ldots, n_{\ell}$. We can write the interpolant using a basis containing tensorized Chebyshev polynomial as

$$
\begin{equation*}
p\left(x_{1}, x_{2}, \ldots, x_{d}\right)=\sum_{i_{1}=0}^{n_{1}} \sum_{i_{2}=0}^{n_{2}} \ldots \sum_{i_{d}=0}^{n_{d}} \mathcal{A}_{i_{1}, i_{2}, \ldots, i_{d}} T_{i_{1}}\left(x_{1}\right) T_{i_{2}}\left(x_{2}\right) \ldots T_{i_{d}}\left(x_{d}\right) \tag{2.16}
\end{equation*}
$$

where $\mathcal{A} \in \mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times \cdots \times\left(n_{d}+1\right)}$ denotes the so-called coefficient tensor. Analogously to the case of univariate interpolation, the coefficient tensor is uniquely determined by the interpolation property (2.15), which can be written as

$$
\begin{equation*}
\mathcal{T}_{j_{1}, j_{2}, \ldots, j_{d}}=\sum_{i_{1}=0}^{n_{1}} \sum_{i_{2}=0}^{n_{2}} \ldots \sum_{i_{d}=0}^{n_{d}} \mathcal{A}_{i_{1}, i_{2}, \ldots, i_{d}} T_{i_{1}}\left(x_{j_{1}}^{(1)}\right) T_{i_{2}}\left(x_{j_{2}}^{(2)}\right) \ldots T_{i_{d}}\left(x_{j_{d}}^{(d)}\right) \tag{2.17}
\end{equation*}
$$

for all $j_{\ell}=0, \ldots, n_{\ell}, \quad \ell=1, \ldots, d$, where the evaluation tensor $\mathcal{T} \in \mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times \cdots \times\left(n_{d}+1\right)}$ is defined elementwise as

$$
\begin{equation*}
\mathcal{T}_{j_{1}, j_{2}, \ldots, j_{d}}=f\left(x_{j_{1}}^{(1)}, x_{j_{2}}^{(2)}, \ldots, x_{j_{d}}^{(d)}\right) \tag{2.18}
\end{equation*}
$$

Let $F^{(\ell)} \in \mathbb{R}^{\left(n_{\ell}+1\right) \times\left(n_{\ell}+1\right)}$ denote the matrix $F$ defined in (2.14) with $n$ replaced by $n_{\ell}$. Multiplying the system (2.17) with $F^{(1)}$ in the first mode yields

$$
\sum_{i_{1}=0}^{n_{1}} F_{j_{1}, i_{1}}^{(1)} \mathcal{T}_{i_{1}, j_{2}, \ldots, j_{d}}=\sum_{i_{2}=0}^{n_{2}} \sum_{i_{3}=0}^{n_{3}} \ldots \sum_{i_{d}=0}^{n_{d}} \mathcal{A}_{j_{1}, i_{2} \ldots, i_{d}} T_{i_{2}}\left(x_{j_{2}}^{(2)}\right) T_{i_{3}}\left(x_{j_{3}}^{(3)}\right) \ldots T_{i_{d}}\left(x_{j_{d}}^{(d)}\right)
$$

for all $j_{\ell}=0, \ldots, n_{\ell}, \ell=1, \ldots, d$. Repeating this procedure in each mode yields the explicit expression for the coefficient tensor

$$
\begin{equation*}
\mathcal{A}=\mathcal{T} \times_{1} F^{(1)} \times_{2} F^{(2)} \times_{3} \cdots \times_{d} F^{(d)} \tag{2.19}
\end{equation*}
$$

We would like to point out that the mode- $\ell$ product of the matrix $F^{(\ell)}$ and $\mathcal{T}$ can be computed efficiently by using the fast Fourier transform described in Remark 2.9. Note that computation of the coefficient tensor (2.19) requires the evaluation of $\left(n_{1}+1\right) \cdot\left(n_{2}+\right.$ 1) $\cdots\left(n_{d}+1\right)$ points of $f$.

The following theorem restates the error bound in [280, Lemma 7.3.3.]. Note that the approximation error (2.20) decays exponentially with respect to $\min \left\{n_{1}, \ldots, n_{d}\right\}$ for analytic functions $f$.

Theorem 2.2. Suppose that the continuous function $f:[-1,1]^{d} \rightarrow \mathbb{R}$ can be extended to an analytic function $f^{*}$ on $\mathcal{E}_{\rho}=E_{\rho_{1}} \times E_{\rho_{2}} \times \cdots \times E_{\rho_{d}}$ with $\rho_{\ell}>1$ for $\ell=1, \ldots, d$, where $E_{\rho}$ again denotes the Bernstein ellipse defined in Theorem 2.1. Then the Chebyshev
interpolant $p$ of $f$ with degree $\left(n_{1}, \ldots, n_{d}\right)$ satisfies

$$
\begin{equation*}
\|f-p\|_{\infty} \leq 2^{1+d / 2} \sqrt{d} \rho_{\min }^{-n_{\min }}\left(1-\rho_{\min }^{-2}\right)^{-d / 2} \max _{z \in \mathcal{E}_{\rho}}\left|f^{*}(z)\right| \tag{2.20}
\end{equation*}
$$

where $\rho_{\min }=\min \left\{\rho_{1}, \rho_{2}, \ldots, \rho_{d}\right\}$ and $n_{\min }=\min \left\{n_{1}, n_{2}, \ldots, n_{d}\right\}$.

### 2.3.3 Coefficient approximation

Due to the curse of dimensionality, it might not be feasible to compute the coefficient tensor (2.19) for large $d$. In the following lemma, we study how approximations of the evaluation tensor (2.18) affect the interpolant (2.16). Similar results with respect to the $L^{2}$-norm can be found in [41].

Lemma 2.2. Let $p$ be defined as in (2.16). For $\hat{\mathcal{T}} \in \mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times \cdots \times\left(n_{d}+1\right)}$, we define the polynomial

$$
\begin{equation*}
\hat{p}\left(x_{1}, \ldots, x_{n}\right)=\sum_{i_{1}=0}^{n_{1}} \ldots \sum_{i_{d}=0}^{n_{d}} \hat{\mathcal{A}}_{i_{1}, \ldots, i_{d}} T_{i_{1}}\left(x_{1}\right) \cdots T_{i_{d}}\left(x_{d}\right) \tag{2.21}
\end{equation*}
$$

where

$$
\begin{equation*}
\hat{\mathcal{A}}=\hat{\mathcal{T}} \times_{1} F^{(1)} \times_{2} F^{(2)} \times_{3} \cdots \times_{d} F^{(d)} \tag{2.22}
\end{equation*}
$$

Then

$$
\begin{equation*}
\|f-\hat{p}\|_{\infty} \leq\|f-p\|_{\infty}+\prod_{\ell=1}^{d}\left(\frac{2}{\pi} \log \left(n_{\ell}\right)+1\right)\|\mathcal{T}-\hat{\mathcal{T}}\|_{\infty} \tag{2.23}
\end{equation*}
$$

Proof. By applying the triangle inequality we obtain

$$
\|f-\hat{p}\|_{\infty} \leq\|f-p\|_{\infty}+\|p-\hat{p}\|_{\infty}
$$

The function $\tilde{p}=p-\hat{p}$ is the unique polynomial of degree $n$ satisfying

$$
\tilde{p}\left(x_{i_{1}}^{(1)}, x_{i_{2}}^{(2)}, \ldots, x_{i_{d}}^{(d)}\right)=(\mathcal{T}-\hat{\mathcal{T}})_{i_{1} i_{2} \ldots i_{d}}, \quad \text { for } i_{\ell}=0, \ldots, n_{\ell}, \ell=1, \ldots, n
$$

Let $n \in \mathbb{N}$. The Lebesgue constant $\Lambda_{n}$ bounds the ratio of the uniform norm approximation error for univariate interpolation and the maximum absolute value at the $n+1$ interpolation nodes. For univariate Chebyshev interpolation we have $\Lambda_{n} \leq(2 / \pi) \log (n)+1$ [322]. In [223], it is shown that the product of the Lebesgue constants for univariate interpolation bounds this ratio for multivariate interpolation, i.e.

$$
\|p-\hat{p}\|_{\infty}=\|\tilde{p}\|_{\infty} \leq \Lambda_{n_{1}} \Lambda_{n_{2}} \cdots \Lambda_{n_{d}}\|\mathcal{T}-\hat{\mathcal{T}}\|_{\infty}
$$

Note that the term $\|f-p\|_{\infty}$ in the error bound (2.23) does not depend on the approximation $\hat{\mathcal{T}}$. For analytic $f$, Theorem 2.2 states that we can choose a sufficiently large degree $\left(n_{1}, n_{2}, \ldots, n_{d}\right)$ such that $\|f-p\|_{\infty}$ is arbitrarily small.

Remark 2.12. In Lemma 2.2, we bound the error of $\|p-\hat{p}\|_{\infty}$ in terms of the Lebesgue constants and $\|\mathcal{T}-\hat{\mathcal{T}}\|_{\infty}$. In the following, we briefly discuss how $\|p-\hat{p}\|_{L^{1}}$ is related to a weighted norm of $\mathcal{T}-\hat{\mathcal{T}}$.

The tensorized Clenshaw-Curtis quadrature formula [74, 86] yields

$$
\begin{aligned}
\|p-\hat{p}\|_{L^{1}} & =\int_{-1}^{1} \int_{-1}^{1} \cdots \int_{-1}^{1}\left|p\left(x_{1}, \ldots, x_{d}\right)-\hat{p}\left(x_{1}, \ldots, x_{d}\right)\right| d x_{1} \ldots d x_{d} \\
& =\sum_{i-1=0}^{n_{1}} \sum_{i_{2}=0}^{n_{2}} \cdots \sum_{i_{d}=0}^{n_{d}} \mathcal{W}_{i_{1}, i_{2}, \ldots, i_{d}}\left|T_{i_{1}, i_{2}, \ldots, i_{d}}-\hat{T}_{i_{1} i_{2} \ldots i_{d}}\right|
\end{aligned}
$$

where $\mathcal{W}_{i_{1}, i_{2}, \ldots, i_{d}}=w_{i_{1}}^{(1)} w_{i_{2}}^{(2)} \cdots w_{i_{d}}^{(d)}$ with nonnegative weights $w_{i_{\ell}}^{(\ell)}$ are defined for $i_{\ell}=$ $0, \ldots, n_{\ell}, \ell=1, \ldots, d$ as in [334]

$$
\begin{aligned}
w_{i_{\ell}}^{(\ell)} & =\frac{c_{\ell}^{(\ell)}}{\left(n_{\ell}\right)}\left(1-\sum_{j=1}^{\left\lfloor n_{\ell} / 2\right\rfloor} \frac{b_{j}^{(\ell)}}{4 j^{2}-1} \cos \left(\frac{2 j i_{\ell} \pi}{n_{\ell}}\right)\right) \\
c_{i_{\ell}}^{(\ell)} & = \begin{cases}1 & i_{\ell}=0 \text { or } i_{\ell}=n_{\ell} \\
2 & \text { otherwise },\end{cases} \\
b_{i_{\ell}}^{(\ell)} & = \begin{cases}1 & i_{\ell}=\frac{n_{\ell}}{2} \\
2 & \text { otherwise }\end{cases}
\end{aligned}
$$

Based on these weights, we define the weighted tensor seminorm

$$
\|\mathcal{T}\|_{\mathcal{W}}=\sum_{i-1=0}^{n_{1}} \sum_{i_{2}=0}^{n_{2}} \cdots \sum_{i_{d}=0}^{n_{d}} \mathcal{W}_{i_{1}, i_{2}, \ldots, i_{d}}\left|T_{i_{1}, i_{2}, \ldots, i_{d}}\right|
$$

This yields the equality $\|p-\hat{p}\|_{L^{1}}=\|\mathcal{T}-\hat{\mathcal{T}}\|_{\mathcal{W}}$. Note that best approximations $\mathcal{T}_{\mathcal{W}}^{*}$ of $\mathcal{T}$ with respect to the seminorm $\|\cdot\|_{\mathcal{W}}$ are related to best approximations $\mathcal{T}_{1}^{*}$ of $\mathcal{W} * \mathcal{T}$ in the $\ell^{1}$-norm, where $*$ denotes the elementwise/Hadamard product for tensors. It holds

$$
\left(\mathcal{T}_{\mathcal{W}}^{*}\right)_{i_{1}, \ldots, i_{d}}= \begin{cases}\left(\mathcal{T}_{1}^{*}\right)_{i_{1}, \ldots, i_{d}} / \mathcal{W}_{i_{1}, \ldots, i_{d}} & \mathcal{W}_{i_{1}, \ldots, i_{d}} \neq 0 \\ \left(\mathcal{T}_{1}^{*}\right)_{i_{1}, \ldots, i_{d}} & \mathcal{W}_{i_{1}, \ldots, i_{d}}=0\end{cases}
$$

for $i_{\ell}=0, \ldots, n_{\ell}, \ell=1, \ldots, d$.

### 2.3.4 Connection to functional low-rank approximations

Assume that $\hat{\mathcal{T}}$ in (2.21) is given in Tucker format (2.5) with multilinear rank $\left(r_{1}, \ldots, r_{d}\right)$ as

$$
\hat{\mathcal{T}}=\mathcal{C} \times_{1} U^{(1)} \times_{2} U^{(2)} \times_{3} \cdots \times_{d} U^{(d)} .
$$

In this case the coefficient tensor $\hat{A}$ defined in (2.22) can be represented in the same low-rank format. This leads to a polynomial approximation (2.21) of the form

$$
\begin{align*}
\hat{p}\left(x_{1}, \ldots, x_{d}\right)= & \sum_{i_{1}=1}^{r_{1}} \ldots \sum_{i_{d}=1}^{r_{d}} \mathcal{C}_{i_{1}, \ldots, i_{d}}\left(\sum_{k_{1}=0}^{n_{1}} \sum_{j_{1}=0}^{n_{1}} F_{k_{1}, j_{1}}^{(1)} U_{j_{1}, i_{1}}^{(1)} T_{k_{1}}\left(x_{1}\right)\right) \ldots \\
& \left(\sum_{k_{d}=0}^{n_{d}} \sum_{j_{d}=0}^{n_{d}} F_{k_{d}, j_{d}}^{(d)} U_{j_{d}, i_{d}}^{(d)} T_{k_{d}}\left(x_{d}\right)\right) . \tag{2.24}
\end{align*}
$$

Note that this is an approximation in the functional Tucker format (2.11) with univariate functions

$$
u_{i}^{(\ell)}\left(x_{\ell}\right)=\sum_{k=0}^{n_{\ell}} \sum_{j=0}^{n_{\ell}} F_{k, j}^{(\ell)} U_{j, i}^{(\ell)} T_{k}\left(x_{\ell}\right), \quad \text { for } i=1, \ldots, r_{\ell}, \ell=1, \ldots, d
$$

The function $u_{i}^{(\ell)}\left(x_{\ell}\right)$ is the univariate Chebyshev interpolant based on the values stored in the $i$ th colum of $U^{(\ell)}$.

Analogously, we can transform a tensor train approximation (2.6) of $\hat{\mathcal{T}}$ with TT cores $\mathcal{G}^{(\ell)}$ into a functional tensor train approximation (2.12) of $f$ by defining the univariate functions

$$
\begin{aligned}
g_{\alpha_{1}}^{(1)}\left(x_{1}\right) & =\sum_{j=0}^{n_{1}} \sum_{k=0}^{n_{1}} F_{j, k}^{(1)} \mathcal{G}_{k, \alpha_{1}}^{(1)} T_{j}\left(x_{1}\right), \\
g_{\alpha_{d-1}}^{(d)}\left(x_{d}\right) & =\sum_{j=0}^{n_{d}} \sum_{k=0}^{n_{d}} F_{j, k}^{(d)} \mathcal{G}_{\alpha_{d-1}, k}^{(d)} T_{j}\left(x_{d}\right), \\
g_{\alpha_{\ell-1}, \alpha_{\ell}}^{(\ell)}\left(x_{\ell}\right) & =\sum_{j=0}^{n_{\ell}} \sum_{k=0}^{n_{\ell}} F_{j, k}^{(\ell)} \mathcal{G}_{\alpha_{\ell-1}, k, \alpha_{\ell}}^{(\ell)} T_{j}\left(x_{\ell}\right) .
\end{aligned}
$$

These can again be seen as univariate Chebyshev interpolants based on the values stored in the TT cores.

### 2.3.5 Proof of concept

The error bound (2.23) indicates that we ideally choose a Tucker approximation $\hat{\mathcal{T}}$ such that $\|\mathcal{T}-\hat{\mathcal{T}}\|_{\infty}$ is of the same order of magnitude as $\|f-p\|_{\infty}$. In the following, we
demonstrate that there exist functions for which we can balance the contributions of these errors with $r_{\ell} \ll n_{\ell}$ for $\ell=1, \ldots, d$. In this case the approximation (2.24) is almost as accurate as the interpolant (2.16), but it requires significantly less storage.

We consider the function

$$
f_{\varepsilon}(x, y, z)=\frac{1}{x+y+z+3+\varepsilon}
$$

on $[-1,1]^{3}$ with parameter $\varepsilon>0$. Let $\tau \geq 0$. In the following paragraphs, we show that the Chebyshev interpolant $p_{\varepsilon}$ defined in (2.16) requires polynomial degrees $n_{\ell}-1=$ $\mathcal{O}(1 / \log (1+\sqrt{\varepsilon}))$ to achieve an accuracy of $\left\|p_{\varepsilon}-f_{\varepsilon}\right\|_{\infty} \leq \tau$. However, one can achieve $\|\mathcal{T}-\hat{\mathcal{T}}\|_{\infty} \leq \tau$ with multilinear ranks $r_{\ell} \leq \mathcal{O}(|\log (\varepsilon)|)$, which grows much slower than $\mathcal{O}(1 / \log (1+\sqrt{\varepsilon})) \approx \mathcal{O}\left(\varepsilon^{-1 / 2}\right)$ for $\varepsilon \rightarrow 0$. For small values of $\varepsilon$ this implies that we can balance the error contributions with $r_{\ell} \ll n_{\ell}$ for $\ell=1, \ldots, d$.

Polynomial Degree For the degree $(n, n, n)$ Chebyshev interpolant $p_{\varepsilon}$ defined in (2.16) we require $\left\|p_{\varepsilon}-f_{\varepsilon}\right\|_{\infty} \leq \tau$, which is equivalent to $\left\|\varepsilon\left(p_{\varepsilon}-f_{\varepsilon}\right)\right\|_{\infty} \leq \varepsilon \tau$. By Theorem 2.2,

$$
\left\|\varepsilon\left(p_{\varepsilon}-f_{\varepsilon}\right)\right\|_{\infty} \leq \mathcal{O}\left(\rho_{\min }^{-n} \cdot \max _{(x, y, z) \in \mathcal{E}_{\rho}}\left|\varepsilon f_{\varepsilon}^{*}(x, y, z)\right|\right)
$$

We set $\rho_{1}=\rho_{2}=\rho_{3}=1+\varepsilon / 6+\sqrt{(1+\varepsilon / 6)^{2}-1}$ and extend $\varepsilon f_{\varepsilon}$ analytically to $\varepsilon f_{\varepsilon}^{*}(x, y, z)=\varepsilon(x+y+z+3+\varepsilon)^{-1}$ on $\mathcal{E}_{\rho}$. By construction $\max _{(x, y, z) \in \mathcal{E}_{\rho}}\left|\varepsilon f_{\varepsilon}^{*}(x, y, z)\right|=2$ is assumed for $x=y=z=-1-\varepsilon / 6$, where $|x+y+z+3+\varepsilon|$ is minimized. Hence, we can choose $n=\mathcal{O}(1 / \log (1+\sqrt{\varepsilon}))$ to obtain the desired accuracy. Although this is only an upper bound for the polynomial degree required, numerical experiments reported below indicate that it is tight.

Multilinear Rank An a priori approximation with exponential sums is used to obtain a bound on the multilinear rank for a tensor containing function values of $f_{\varepsilon}$; see [153]. Given $R>1$ and $r \in \mathbb{N}$, Braess and Hackbusch [47] showed that there exist coefficients $a_{i}$ and $b_{i}$ such that

$$
\begin{equation*}
\left|\frac{1}{x}-\sum_{i=1}^{r} a_{i} \exp \left(-b_{i} x\right)\right| \leq 16 \exp \left(-\frac{r \pi^{2}}{\log (8 R)}\right), \quad \forall x \in[1, R] \tag{2.25}
\end{equation*}
$$

Trivially, we have $\varepsilon f_{\varepsilon}(x, y, z)=1 / \omega$ for the substitution $\omega=(x+y+z+3+\varepsilon) / \varepsilon$ with $\omega \in[1,1+6 / \varepsilon]$. Applying (2.25) yields that there exist $a_{i}$ and $b_{i}$ such that $\left|1 / \omega-\sum_{i=1}^{r} a_{i} \exp \left(-b_{i} \omega\right)\right| \leq \tau \varepsilon$ or, equivalently,

$$
\begin{equation*}
\left\|f_{\varepsilon}(x, y, z)-g_{\varepsilon}(x, y, z)\right\|_{\infty} \leq \tau \tag{2.26}
\end{equation*}
$$

for every $x, y, z \in[-1,1]$ when

$$
r \geq \frac{-\log \left(8\left(1+\frac{6}{\varepsilon}\right)\right) \log \left(\frac{\tau}{16}\right)}{\pi^{2}}=\mathcal{O}(|\log (\varepsilon)|)
$$

where

$$
g_{\varepsilon}(x, y, z)=\sum_{i=1}^{r} \frac{a_{i}}{\varepsilon} \cdot \exp \left(-\frac{b_{i}}{\varepsilon} x\right) \cdot \exp \left(-\frac{b_{i}}{\varepsilon} y\right) \cdot \exp \left(-\frac{b_{i}}{\varepsilon} z\right) \cdot \exp \left(-\frac{b_{i}}{\varepsilon}(3+\varepsilon)\right)
$$

The approximation $g_{\varepsilon}$ in (2.26) has multilinear rank $(r, r, r)$. In turn, the tensor $\hat{\mathcal{T}}_{i, j, k}=$ $g_{\varepsilon}\left(x_{i}^{(1)}, x_{j}^{(2)}, x_{k}^{(3)}\right)$ has multilinear rank at most $(r, r, r)$ and satisfies $\|\mathcal{T}-\hat{\mathcal{T}}\|_{\infty} \leq \tau$.

Comparison In Figure 2.2, we estimate the maximal polynomial degree required to compute a Chebyshev interpolant with accuracy $\tau=10^{-10}$ for selected fibers of $f_{\varepsilon}$, which is a lower bound for the required polynomial degrees $n_{\ell}$. It perfectly matches the asymptotic behavior of the derived upper bound $\mathcal{O}(1 / \log (1+\sqrt{\varepsilon}))$. In Figure 2.2, we also plot the multilinear ranks from the truncated HOSVD [88] with tolerance $\tau$ applied to the tensor containing the evaluation of $f_{\varepsilon}$ on a $150 \times 150 \times 150$ Chebyshev grid. This estimate serves as a lower bound for the multilinear rank required to approximate $f_{\varepsilon}$. Due to the limited grid size, this estimate does not fully match the asymptotic behavior $|\log (\varepsilon)|$, but nonetheless it clearly reflects that the multilinear ranks can be much smaller than the polynomial degrees, as predicted by $|\log (\varepsilon)| \ll 1 / \log (1+\sqrt{\varepsilon})$, for sufficiently small $\varepsilon$.


Figure 2.2 - Comparison of the theoretical upper bounds $\mathcal{O}(|\log (\varepsilon)|)$ for the multilinear rank and $\mathcal{O}(1 / \log (1+\sqrt{\varepsilon}))$ for the polynomial degree for varying $\varepsilon$, the maximal polynomial degree, used by Chebfun to approximate selected functions fibers of $f_{\varepsilon}$ up to accuracy $10^{-10}$, and the maximal multilinear rank of the truncated HOSVD with tolerance $10^{-10}$ of a sample tensor on a $150 \times 150 \times 150$ Chebyshev grid. The constants in the bounds are chosen to result in curves close to the data.

## 3 Functional Tucker approximation

In this chapter, we focus on developing an efficient algorithm to approximate trivariate functions (that is, functions depending on three variables) defined on the tensor-product domain $[-1,1]^{3}$ in the functional Tucker format (2.11). We proceed by combining tensorized Chebyshev interpolation (2.16) with a Tucker decomposition (2.5) of the coefficient tensor (2.19). This leads to a functional Tucker approximation of the form (2.24).

Hashemi and Trefethen followed this approach to develop an approximation algorithm for Chebfun3 [163]. Their algorithm is adaptive in both the multilinear rank and the polynomial degree of the computed approximation, but the construction of the approximation proceeds indirectly, via a so-called slice-Tucker decomposition of the evaluation tensor (2.18). As a consequence, Chebfun3 sometimes uses unnecessarily many function evaluations and does not fully benefit from the potential of the Tucker decomposition to reduce, sometimes dramatically, the computational cost.

We propose a novel algorithm that utilizes univariate fibers instead of bivariate slices to construct a Tucker decomposition (2.5) of the evaluation tensor (2.18) directly. We call our approach Chebfun3F to emphasize that it is based on selecting the Fibers of the evaluation tensor to construct the initial factor matrices. In a second step, we replace orthogonal projections (2.7) by oblique projections (2.3) based on DEIM to construct the core tensor. We combine this approach with heuristics similar to the ones used in Chebfun3 for choosing the polynomial degree and for the accuracy verification. Chebfun3F reduces the cost for the approximation in terms of the number of function evaluations for nearly all functions considered, typically by $75 \%$, and sometimes by over $98 \%$.

This chapter is based on the article [100]. Its remainder is structured as follows. In Section 3.1, we briefly recall the approximation algorithm currently used in Chebfun3. Section 3.2 introduces our novel algorithm Chebfun3F. In Section 3.3, we perform numerical experiments to compare Chebfun3, Chebfun3F and sparse grid interpolation. We then analyse the coefficient decay of the resulting approximations and discuss how Chebfun3F can be incorporated into Chebfun in Section 3.4. In Section 3.5, we analyze
further properties of functional Tucker decompositions computed based on tensorized Chebyshev interpolation.

### 3.1 Existing algorithm: Chebfun3

In this section, we recall how an approximation of the form (2.11) is computed in Chebfun3 [163]. As discussed in Section 2.3.5, there are often situations in which approximations (2.24) require a much smaller $\left(r_{1}, \ldots, r_{d}\right)$ compared to the polynomial degree $\left(n_{1}, \ldots, n_{d}\right)$. Chebfun3 benefits from such a situation by first using a coarse evaluation tensor $\mathcal{T}_{c}$ to identify the fibers needed for the low-rank approximation. This allows to construct the actual approximation from a finer sample tensor $\mathcal{T}_{f}$ by only evaluating these fibers instead of the whole tensor.

Chebfun3 consists of three phases: preparation of the approximation by identifying fibers for a so-called block term decomposition [87] of $\mathcal{T}_{c}$, refinement of the fibers, conversion and compression of the refined block term decomposition of $\mathcal{T}_{f}$ into functional Tucker format (2.11).

### 3.1.1 Phase 1: Block term decomposition

In Chebfun3, the coarse evaluation tensor $\mathcal{T}_{c} \in \mathbb{R}^{\left(n_{1}^{(c)}+1\right) \times\left(n_{2}^{(c)}+1\right) \times\left(n_{3}^{(c)}+1\right)}$ is initially obtained by sampling $f$ on a $17 \times 17 \times 17$ grid of Chebyshev points. A block term decomposition of $\mathcal{T}_{c}$ is obtained by applying ACA (see Algorithm 1) recursively. In the first step, ACA is applied to a matricization of $\mathcal{T}_{c}$, say, the mode- 1 matricization $\mathcal{T}_{c}^{\{1\}}$. This results in index sets $I, J$ such that

$$
\begin{equation*}
\mathcal{T}_{c}^{\{1\}} \approx \mathcal{T}_{c}^{\{1\}}(:, J)\left(\mathcal{T}_{c}^{\{1\}}(I, J)\right)^{-1} \mathcal{T}_{c}^{\{1\}}(I,:) \tag{3.1}
\end{equation*}
$$

where $\mathcal{T}_{c}{ }^{\{1\}}(:, J)$ contains mode- 1 fibers of $\mathcal{T}_{c}$ and $\mathcal{T}_{c}{ }^{\{1\}}(I,:)$ contains mode- $(2,3)$ slices of $\mathcal{T}_{c}$. For each $i \in I$, such a slice $\mathcal{T}_{c}^{\{1\}}(i,:)$ is reshaped into a matrix $S_{i}=\mathcal{T}_{c}(i,:,:) \in$ $\mathbb{R}^{\left(n_{2}^{(c)}+1\right) \times\left(n_{3}^{(c)}+1\right)}$ and, in the second step, approximated by again applying ACA:

$$
\begin{equation*}
S_{i} \approx S_{i}\left(:, L_{i}\right)\left(S_{i}\left(K_{i}, L_{i}\right)\right)^{-1} S_{i}\left(K_{i},:\right) \tag{3.2}
\end{equation*}
$$

where $S_{i}\left(:, L_{i}\right)$ and $S_{i}\left(K_{i},:\right)$ contain mode-2 and mode-3 fibers of $\mathcal{T}_{c}$, respectively. Combining (3.1) and (3.2) yields the approximation

$$
\mathcal{T}_{c}^{\{1\}} \approx \mathcal{T}_{c}^{\{1\}}(:, J)\left(\mathcal{T}_{c}^{\{1\}}(I, J)\right)^{-1}\left(\begin{array}{c}
\operatorname{vec}\left(S_{1}\left(:, L_{1}\right)\left(S_{1}\left(K_{1}, L_{1}\right)\right)^{-1} S_{1}\left(K_{1},:\right)\right)  \tag{3.3}\\
\operatorname{vec}\left(S_{2}\left(:, L_{2}\right)\left(S_{2}\left(K_{2}, L_{2}\right)\right)^{-1} S_{2}\left(K_{1},:\right)\right) \\
\vdots
\end{array}\right)
$$

where vec denotes vectorization, i.e. the reshaping of a matrix into a vector. Reshaping this approximation into a tensor can be viewed as a block term decomposition in the sense of [87, Definition 2.2.]. This approximation scheme is visualized in Figure 3.1.


Figure 3.1 - Visualization of a block term decomposition (3.3).
If the ratios of $|I| /\left(n_{1}^{(c)}+1\right),\left|K_{i}\right| /\left(n_{2}^{(c)}+1\right)$ and $\left|L_{i}\right| /\left(n_{3}^{(c)}+1\right)$ are larger than the heuristic threshold $(2 \sqrt{2})^{-1}$ the coarse grid resolution $\left(n_{1}^{(c)}+1, n_{2}^{(c)}+1, n_{3}^{(c)}+1\right)$ is deemed insufficient to identify fibers. If this is the case, $n_{\ell}^{(c)}$ is increased to $\left\lfloor\sqrt{2}^{\left\lfloor 2 \log _{2}\left(n_{\ell}^{(c)}+1\right)+1\right\rfloor}\right\rfloor$ and Phase 1 is repeated.

### 3.1.2 Phase 2: Refinement

The block term decomposition (3.3) is composed of fibers of $\mathcal{T}_{c}$. Such a fiber $\mathcal{T}_{c}(:, j, k)$ corresponds to the evaluation of a univariate function $f(\cdot, y, z)$ for certain fixed $y, z$. We apply the Chebfun heuristic described in Remark 2.11 to decide whether the function values in $\mathcal{T}_{c}(:, j, k)$ suffice to yield an accurate interpolation of $f(\cdot, y, z)$ [16]. If this is not the case, we add nested evaluations of $f(\cdot, y, z)$ to refine the approximation.

In Chebfun3 this heuristic is applied to all mode- $\ell$ fibers contained in (3.3) simultaneously. Let $n_{\ell}^{(f)}$ denote the polynomial degree for which the heuristic states that the interpolation of is accurate. The refined mode- $\ell$ fibers are now vectors of length $n_{\ell}^{(f)}+1$. Replacing all fibers in (3.3) by their refined counterparts yields an approximation of the tensor $\mathcal{T}_{f}$, which contains the evaluations of $f$ on a $\left(n_{1}^{(f)}+1\right) \times\left(n_{2}^{(f)}+1\right) \times\left(n_{3}^{(f)}+1\right)$ Chebyshev grid. Note that $\mathcal{T}_{f}$ might be very large and is never computed explicitly. We visualize the refinement idea in Figure 3.2.

### 3.1.3 Phase 3: Compression

In the third phase of the Chebfun3 constructor, the refined block term decomposition is converted and compressed to the desired Tucker format (2.11), where the interpolants


Figure 3.2 - Visualization of the refinement step described in Section 3.1.2. In the refinement step, we replace the initial fiber $\mathcal{T}_{c}(:, j, 0)$ by the fiber $\mathcal{T}_{f}(:, j, 0)$. Here we depict tensors for $\left(n_{1}^{(c)}, n_{2}^{(c)}, n_{3}^{(c)}\right)=(11,11,11)$ and $\left(n_{1}^{(f)}, n_{2}^{(f)}, n_{3}^{(f)}\right)=(22,11,11)$.
$u_{i}(x)$ are stored as Chebfun objects [27]; see [163] for details. Lemma 2.2 guarantees a good approximation $\hat{f}$ when the polynomial degree $\left(n_{1}^{(f)}, n_{2}^{(f)}, n_{3}^{(f)}\right)$ is sufficiently large and when $\mathcal{T}_{f}$ is well approximated by the underlying Tucker approximation $\hat{\mathcal{T}}$. Neither of these properties can be guaranteed in Phases 1 and 2 alone. Therefore in a final step, Chebfun3 verifies the accuracy by comparing $f$ and the computed approximation at Halton points [239]. If the estimated error is too large, the whole algorithm is restarted on a finer coarse grid from Phase 1.

### 3.1.4 Disadvantages

The Chebfun3 algorithm often requires unnecessarily many function evaluations. As we illustrate in the following, this is due to redundancy among the mode- 2 and mode- 3 fibers. For this purpose we collect all (refined) mode-2 fibers $S_{i}\left(:, L_{i}\right)$ in the block term decomposition (3.3) into the columns of a big matrix $V_{m}^{\text {BTD }}=\left[\begin{array}{lll}S_{1}\left(:, L_{1}\right) & \cdots & S_{m}\left(:, L_{m}\right)\end{array}\right]$, where $m$ is the number of steps of the outer ACA (3.1). As will be demonstrated with an example below, the matrix $V_{m}^{\text {BTD }}$ is often observed to have low numerical rank, which in turn allows to represent its column space by much fewer columns, that is, much fewer mode-2 fibers. As the accuracy of the column space determines the accuracy of the Tucker decomposition after the compression, this implies that the other mode-2 fibers in $V_{m}^{\text {BTD }}$ are redundant.

Let us now consider the block term decomposition* (3.3) for the function

$$
f(x, y, z)=\frac{1}{1+25 \sqrt{x^{2}+y^{2}+z^{2}}} .
$$

In Figure 3.3 the numerical rank and the number of columns of $V_{m}^{\text {BTD }}$ are compared. For $m=10$ the approximation of the slices $S_{i_{1}}$ to $S_{i_{10}}$ leads to a total of 153 mode-2 fibers,

[^0]the sum of the corresponding red and blue bars in Figure 3.3. In contrast, their numerical rank (blue bar) is only 19. Thus, the red bar can be interpreted as number of redundant mode-2 fibers. This happens since nearby slices tend to be similar. The total block term


Figure 3.3 - Numerical rank and number of redundant columns ( $=$ total number of columns - numerical rank) of the matrix $V_{m}^{\text {BTD }}$, whose columns are given by the refined version of the mode-2 fibers determined after $m$ steps of the outer ACA (3.1) in Chebfun3.
decomposition contains 18 slices and is compressed into a Tucker decomposition with multilinear rank ( $17,19,19$ ). It contains 242 redundant fibers, the refinement requires 192 function evaluations for each of them. Note that the asymmetry in the rank of the Tucker decomposition is caused by the asymmetry of the block term decomposition.

Another disadvantage is that Chebfun3 always requires the full evaluation of $\mathcal{T}_{c}$ in Phase 1. This becomes expensive when a large size $\left(n_{1}^{(c)}+1\right) \times\left(n_{2}^{(c)}+1\right) \times\left(n_{3}^{(c)}+1\right)$ is needed in order to properly identify suitable fibers.

### 3.2 Novel algorithm: Chebfun3F

In this section, we describe our novel algorithm Chebfun3F to compute an approximation of the form (2.11). The goal of Chebfun3F is to the avoid the redundant function evaluations observed in Chebfun3. While the structure of Chebfun3F is similar to Chebfun3, consisting of 3 phases to identify/refine fibers and compute a Tucker decomposition, there is a major difference in Phase 1. Instead of proceeding via slices, we directly identify mode- $\ell$ fibers of $\mathcal{T}_{c}$ for building factor matrices. The core tensor is constructed in Phase 3.

### 3.2.1 Phase 1: Fiber indices and factor matrices

As in Chebfun3, the coarse tensor $\mathcal{T}_{c} \in \mathbb{R}^{\left(n_{1}^{(c)}+1\right) \times\left(n_{2}^{(c)}+1\right) \times\left(n_{3}^{(c)}+1\right)}$ is initially defined as evaluation tensor (2.18) with $\left(n_{1}, n_{2}, n_{3}\right)=(16,16,16)$. We seek to compute full rank
factor matrices $U_{c} \in \mathbb{R}^{\left(n_{1}^{(c)}+1\right) \times r_{1}}, V_{c} \in \mathbb{R}^{\left(n_{2}^{(c)}+1\right) \times r_{2}}$ and $W_{c} \in \mathbb{R}^{\left(n_{3}^{(c)}+1\right) \times r_{3}}$ such that the orthogonal projection of $\mathcal{T}_{c}$ onto the span of the factor matrices is an accurate approximation of $\mathcal{T}_{c}$, i.e.

$$
\begin{equation*}
\mathcal{T}_{c} \approx \mathcal{T}_{c} \times_{1} U_{c}\left(U_{c}^{T} U_{c}\right)^{-1} U_{c}^{T} \times_{2} V_{c}\left(V_{c}^{T} V_{c}\right)^{-1} V_{c}^{T} \times_{3} W_{c}\left(W_{c}^{T} W_{c}\right)^{-1} W_{c}^{T} . \tag{3.4}
\end{equation*}
$$

Additionally, we require that the columns in $U_{c}, V_{c}, W_{c}$ contain fibers of $\mathcal{T}_{c}$.
In the existing literature, algorithms to compute such factor matrices include the Higher Order Interpolatory Decomposition [277], which is based on a rank revealing QR decomposition, and the Fiber Sampling Tensor Decomposition [54], which is a generalization of the CUR decomposition. We propose a novel algorithm, which in contrast to the existing algorithms does not require the evaluation of the full tensor $\mathcal{T}_{c}$. We follow the ideas of TT-cross [251, 281] and its variants such as the Schur-Cross3D [271] and the ALS-cross [101]. After the publication of [100] several algorithms for this purpose have been suggested in [3]. The Tucker format with factor matrices containing fibers has been called generalized tensor CUR decomposition in [157], mode-wise tensor decomposition in [53] and tensor fiber CUR decomposition in [52].

Initially, we randomly choose index sets $\tilde{I}, \tilde{J}, \tilde{K}$ by partitioning $\{0, \ldots, 16\}$ into 6 subsets and sampling one index from each subset for each index set. In the first step, we apply Algorithm 1 to $\left(\mathcal{T}_{c}(:, \tilde{J}, \tilde{K})\right)^{\{1\}}$. Note that this needs only $36 \cdot\left(n_{1}^{(c)}+1\right)$ evaluations of the function $f$, in contrast to $\left(n_{1}^{(c)}+1\right)\left(n_{2}^{(c)}+1\right)\left(n_{3}^{(c)}+1\right)$ values in the whole tensor $\mathcal{T}_{c}$. The selected $r_{1}$ columns serve as a first candidate for the factor matrix $U_{c}$. The index set $\tilde{I}$ is set to the row indices selected by Algorithm 1 (see Figure 3.4). We use the updated index set and apply Algorithm 1 to $\left(\mathcal{T}_{c}(\tilde{I},: \tilde{K})\right)^{\{2\}}$ analogously, which yields $V_{c}$ and an updated $\tilde{J}$. From $\left(\mathcal{T}_{c}(\tilde{I}, \tilde{J},:)\right)^{\{3\}}$ we obtain $W_{c}$ and $\tilde{K}$. We repeat this process in an alternating fashion with the updated index sets, which leads to potentially improved factor matrices. Following the ideas of Chebfun3, we check after each iteration whether the ratios $r_{1} /\left(n_{1}^{(c)}+1\right), r_{2} /\left(n_{2}^{(c)}+1\right)$ and $r_{3} /\left(n_{3}^{(c)}+1\right)$ surpass the heuristic threshold $(2 \sqrt{2})^{-1}$. If this is the case, we increase the size of the coarse tensor $n_{\ell}^{(c)}$ to $\left.\left\lfloor\sqrt{2}^{\left\lfloor\log _{2}\left(n_{\ell}^{(c)}+1\right)+1\right.}\right\rfloor\right\rfloor$ and restart the whole process by reinitializing $\tilde{I}, \tilde{J}, \tilde{K}$ with $r_{1}, r_{2}, r_{3}$ random indices respectively.

It is not clear a priori how many iterations are needed to attain an approximation (3.4) that yields a Tucker approximation (2.11) which passes the accuracy verification in Phase 3. In numerical experiments, it has usually proven to be sufficient to stop after the second iteration, during which the coarse grid has not been refined, or when $|\tilde{I}| \leq 1$, $|\tilde{J}| \leq 1$ or $|\tilde{K}| \leq 1$. This is formalized in Algorithm 4. Note that $U_{c}, V_{c}, W_{c}$ are full rank by construction, since Algorithm 1 stops based on the tolerance $\varepsilon \geq 0$. In many cases, we found that the numbers of columns in the factor matrices are equal to the multilinear rank of the truncated HOSVD [88] of $\mathcal{T}_{c}$ with the same tolerance.


Figure 3.4 - Visualization of applying ACA (Algorithm 1) to a matricization of a subtensor.

```
Algorithm 4 Factor Matrix Computation
    Input: \(f,\left(n_{1}^{(c)}, n_{2}^{(c)}, n_{3}^{(c)}\right),\left(r_{1}, r_{2}, r_{3}\right)\)
    Output: \(U_{c}, V_{c}, W_{c},\left(n_{1}^{(c)}, n_{2}^{(c)}, n_{3}^{(c)}\right),\left(r_{1}, r_{2}, r_{3}\right)\)
    Let \(\mathcal{T}_{c}(i, j, k)=f\left(x_{i}^{(1)}, x_{j}^{(2)}, x_{k}^{(3)}\right)\) be a function that evaluates the values of \(f\) on a
    \(\left(n_{1}^{(c)}+1\right) \times\left(n_{2}^{(c)}+1\right) \times\left(n_{3}^{(c)}+1\right)\) Chebyshev grid on demand.
    initialize \(\tilde{J}, \tilde{K}\) with \(r_{2}, r_{3}\) randomly chosen indices in \(\left\{0, \ldots, n_{1}^{(c)}\right\},\left\{0, \ldots, n_{2}^{(c)}\right\}\)
    for iterations \(=1: 2\)
        compute ACA of \(\mathcal{T}_{c}(:, \tilde{J}, \tilde{K})^{\{1\}} \rightarrow U_{c} \in \mathbb{R}^{\left(n_{1}^{(c)}+1\right) \times r_{1}}=\) selected columns, \(\tilde{I}=\)
    selected row indices
        compute ACA of \(\mathcal{T}_{c}(\tilde{I},:, \tilde{K})^{\{2\}} \rightarrow V_{c} \in \mathbb{R}^{\left(n_{2}^{(c)}+1\right) \times r_{2}}=\) selected columns, \(\tilde{J}=\)
    selected row indices
        compute ACA of \(\mathcal{T}_{c}(\tilde{I}, \tilde{J},:)^{\{3\}} \rightarrow W_{c} \in \mathbb{R}^{\left(n_{3}^{(c)}+1\right) \times r_{3}}=\) selected columns, \(\tilde{K}=\)
    selected row indices
        if the multilinear ranks get too large \(\rightarrow\) adjust \(\left(n_{1}^{(c)}, n_{2}^{(c)}, n_{3}^{(c)}\right)\) and go to line 3
        if \(r_{1} \leq 1\) or \(r_{2} \leq 1\) or \(r_{3} \leq 1 \rightarrow\) return
```


### 3.2.2 Phase 2: Refinement of the factors

In Phase 2, the fibers in $U_{c}, V_{c}, W_{c}$ are refined using the Chebfun heuristic [16] as in Chebfun3 (see Section 3.1.2). This leads to full rank factor matrices $U_{f} \in \mathbb{R}^{\left(n_{1}^{(f)}+1\right) \times r_{1}}$, $V_{f} \in \mathbb{R}^{\left(n_{2}^{(f)}+1\right) \times r_{2}}$ and $W_{f} \in \mathbb{R}^{\left(n_{3}^{(f)}+1\right) \times r_{3}}$ containing the refined fibers of $\mathcal{T}_{f}$, where $\mathcal{T}_{f}$ denotes the evaluation tensor (2.18) of size $\left(n_{1}^{(f)}+1\right) \times\left(n_{2}^{(f)}+1\right) \times\left(n_{3}^{(f)}+1\right)$. This phase needs only $\mathcal{O}\left(\sum_{\ell=1}^{3}\left(n_{\ell}^{(f)}+1\right) r_{\ell}\right)$ evaluations of $f$.

### 3.2.3 Phase 3: Reconstruction of the core tensor

In the final Phase of Chebfun3F, we compute the core tensor $\hat{\mathcal{C}}$ to obtain the approximation $\mathcal{T}_{f} \approx \hat{\mathcal{C}} \times_{1} U_{f} \times_{2} V_{f} \times_{3} W_{f}$.

In principle, the best approximation (with respect to the Frobenius norm) for fixed factor
matrices $U_{f}, V_{f}, W_{f}$ is obtained by orthogonal projections as in (2.7) by

$$
\mathcal{T}_{f} \approx \mathcal{T}_{f} \times_{1} U_{f}\left(U_{f}^{T} U_{f}\right)^{-1} U_{f}^{T} \times{ }_{2} V_{f}\left(V_{f}^{T} V_{f}\right)^{-1} V_{f}^{T} \times_{3} W_{f}\left(W_{f}^{T} W_{f}\right)^{-1} W_{f}^{T} .
$$

Such an approach comes with the major disadvantage that the full evaluation of $\mathcal{T}_{f}$ is required. This can be circumvented by instead using oblique projections. Applying oblique projections in all three modes analogous to (2.3) yields

$$
\begin{aligned}
\mathcal{T}_{f} \approx \hat{\mathcal{T}}= & \mathcal{T}_{f} \times_{1} U_{f}\left(\Phi_{I}^{T} U_{f}\right)^{-1} \Phi_{I}^{T} \times_{2} V_{f}\left(\Phi_{J}^{T} V_{f}\right)^{-1} \Phi_{J}^{T} \times_{3} W_{f}\left(\Phi_{K}^{T} W_{f}\right)^{-1} \Phi_{K}^{T} \\
= & \underbrace{(\underbrace{\left(\mathcal{T}_{f} \times{ }_{1} \Phi_{I}^{T} \times_{2} \Phi_{J}^{T} \times_{3} \Phi_{K}^{T}\right)}_{=\mathcal{T}_{f}(I, J, K)} \times_{1}\left(\Phi_{I}^{T} U_{f}\right)^{-1} \times_{2}\left(\Phi_{J}^{T} V_{f}\right)^{-1} \times_{3}\left(\Phi_{K}^{T} W_{f}\right)^{-1})}_{=\hat{\mathcal{C}}} \\
& \times{ }_{1} U_{f} \times_{2} V_{f} \times{ }_{3} W_{f},
\end{aligned}
$$

for index sets $I, J, K$ and matrices $\Phi_{I}, \Phi_{J}, \Phi_{K}$ defined analogously as in (2.3). Note that the computation of the $\mathcal{T}_{f}$ only requires $r_{1} \cdot r_{2} \cdot r_{3}$ additional evaluations of $f$. To avoid the potentially ill-conditioned matrices $\left(\Phi_{I}^{T} U_{f}\right)^{-1},\left(\Phi_{J}^{T} V_{f}\right)^{-1},\left(\Phi_{K}^{T} W_{f}\right)^{-1}$, we compute the orthogonal matrices $Q_{U}, Q_{V}, Q_{W}$ in the QR decompositions of $U_{f}, V_{f}, W_{f}$. Note that $U_{f}\left(\Phi_{I}^{T} U_{f}\right)^{-1}=Q_{U}\left(\Phi_{I}^{T} Q_{U}\right)^{-1}$ and $\left(\mathcal{T}_{f}(I, J, K) \times_{1}\left(\Phi_{I}^{T} U_{f}\right)^{-1}\right) \times_{1} U_{f}=\mathcal{T}_{f}(I, J, K) \times_{1}$ $U_{f}\left(\Phi_{I}^{T} U_{f}\right)^{-1}$. To determine the index sets $I, J, K$ we apply DEIM [65], presented in Algorithm 2, to the matrices $Q_{U}, Q_{V}, Q_{W}$. This leads to the Tucker decomposition used in Chebfun3F

$$
\begin{align*}
\hat{\mathcal{T}} & =\mathcal{T}_{f} \times_{1} Q_{U}\left(\Phi_{I}^{T} Q_{U}\right)^{-1} \Phi_{I}^{T} \times_{2} Q_{V}\left(\Phi_{J}^{T} Q_{V}\right)^{-1} \Phi_{J}^{T} \times_{3} Q_{W}\left(\Phi_{K}^{T} Q_{W}\right)^{-1} \Phi_{K}^{T} \\
& =\mathcal{T}_{f}(I, J, K) \times_{1} Q_{U}\left(\Phi_{I}^{T} Q_{U}\right)^{-1} \times_{2} Q_{V}\left(\Phi_{J}^{T} Q_{V}\right)^{-1} \times_{3} Q_{W}\left(\Phi_{K}^{T} Q_{W}\right)^{-1} . \tag{3.5}
\end{align*}
$$

Note that in general neither $Q_{U}$ nor $Q_{U}\left(\Phi_{I}^{T} Q_{U}\right)^{-1}$ contain fibers of $\mathcal{T}_{f}$. This issue will be discussed in Section 3.4.

### 3.2.4 Chebfun3F algorithm

Having computed the Tucker factors $U_{f}, V_{f}$ and $W_{f}$ in Phase 2, and the core $\hat{\mathcal{C}}$ in Phase 3, we obtain the approximation (2.24) of the form (2.11). Following Chebfun3, we compute the Chebyshev interpolants of the columns of $U_{f}, V_{f}, W_{f}$ using Chebfun [27]. We also perform an accuracy verification for $\hat{f}$ by comparing its evaluations at Halton points to the original $f$. If the difference of the evaluations is too large, we restart the whole algorithm up to ten times using a finer coarse grid. Additionally, we modify the ranks such that if $r_{\ell_{1}} \leq 2$ we set $r_{\ell_{2}}=\max \left(6,2 r_{\ell_{2}}\right)$ and $r_{\ell_{1}}=3$ for $\ell_{1} \neq \ell_{2}$, and after the forth restart we set $r_{\ell}=2 r_{\ell}$ for $\ell=1,2,3$. This ensures that the multilinear ranks can grow in Phase 1. The overall Chebfun3F algorithm is formalized in Algorithm 5.

## Algorithm 5 Chebfun3F

Input: function $f(x, y, z)$, a procedure $\mathcal{T}_{f}(i, j, k)=f\left(x_{i}^{(1)}, x_{j}^{(2)}, x_{k}^{(3)}\right)$ that evaluates the values of $f$ on a $\left(n_{1}^{(f)}+1\right) \times\left(n_{2}^{(f)}+1\right) \times\left(n_{3}^{(f)}+1\right)$ Chebyshev grid on demand, a stopping tolerance $\varepsilon>0$.
Output: approximation $\hat{f}(x, y, z)=\sum_{i=1}^{r_{1}} \sum_{j=1}^{r_{2}} \sum_{k=1}^{r_{3}} \mathcal{C}_{i j k} u_{i}(x) v_{j}(y) w_{k}(z)$
Initialization: $\left(n_{1}^{(c)}, n_{c}^{(c)}, n_{3}^{(c)}\right)=(16,16,16),\left(r_{1}, r_{2}, r_{3}\right)=(6,6,6)$
Phase 1:
apply Algorithm 4 to compute the factor matrices $U_{c}, V_{c}, W_{c}$ and to update $\left(n_{1}^{(c)}, n_{2}^{(c)}, n_{3}^{(c)}\right),\left(r_{1}, r_{2}, r_{3}\right)$
Phase 2:
$\left(n_{1}^{(f)}, n_{2}^{(f)}, n_{3}^{(f)}\right)=\left(n_{1}^{(c)}, n_{2}^{(c)}, n_{3}^{(c)}\right), \quad U_{f}=U_{c}, V_{f}=V_{c}, W_{f}=W_{c}$
while Chebfun heuristic in [16] to decide if $U_{f}$ contains a sufficient number of entries is not satisfied $n_{1}^{(f)}+1=2 n_{1}^{(f)}, \quad$ refine $U_{f}$ such that its columns have length $n_{1}^{(f)}$ proceed analogously to obtain $V_{f}, W_{f}$
Phase 3:
$\left[Q_{U}, \sim\right]=\operatorname{qr}\left(U_{f}\right), \quad I=\operatorname{DEIM}\left(Q_{U}\right), \quad U=Q_{U} \cdot Q_{U}(I,:)^{-1}$
$\left[Q_{V}, \sim\right]=\operatorname{qr}\left(V_{f}\right), \quad J=\operatorname{DEIM}\left(Q_{V}\right), \quad V=Q_{V} \cdot Q_{V}(J,:)^{-1}$
$\left[Q_{W}, \sim\right]=\operatorname{qr}\left(W_{f}\right), \quad K=\operatorname{DEIM}\left(Q_{W}\right), \quad W=Q_{W} \cdot Q_{W}(K,:)^{-1}$
compute the Chebyshev interpolants $u_{i}(x), v_{j}(y), w_{k}(z)$ based on $U, V, W$ (see Section 2.3.4)
$\mathcal{C}=\mathcal{T}_{f}(I, J, K), \quad \hat{f}(x, y, z)=\sum_{i=1}^{r_{1}} \sum_{j=1}^{r_{2}} \sum_{k=1}^{r_{3}} \mathcal{C}_{i j k} u_{i}(x) v_{j}(y) w_{k}(z)$ if $|f(x, y, z)-\hat{f}(x, y, z)|>10 \varepsilon$ at Halton points $(x, y, z) \in[-1,1]^{3}$ modify the ranks $r_{\ell}$ if they are too small
restart from Phase 1 with $\left.n_{\ell}^{(c)}=\left\lfloor\sqrt{2}^{\left\lfloor 2 \log _{2}\left(n_{\ell}^{(c)}+1\right)+1\right.}\right\rfloor\right\rfloor$

Remark 3.1. In Chebfun3 upper bounds for the multilinear rank, polynomial degree and grid sizes are prescribed. The tolerances in the accuracy verification and in the ACA are initially set close to machine precision or provided by the user. Tolerance issues are avoided by relaxing these tolerances adaptively based on the computed function evaluations. In Chebfun3F, we handle these technicalities in the same manner.

### 3.2.5 Existence of a quasi-optimal Chebfun3F approximation

Due to the many heuristic ingredients in the Chebfun3F algorithm, it is difficult to analyze the convergence of the whole algorithm. Instead, we discuss the existence and error analysis of a specific Chebfun3F reconstruction. Lemma 2.2 shows how we can bound the approximation error depending on $\|\mathcal{T}-\hat{\mathcal{T}}\|_{\infty}$. Theorem 3.1 provides a bound for this error for a tensor approximation of the format (3.5) with specifically chosen fibers and index sets. The best approximation in the format will be at least as good. Although Chebfun3F is not guaranteed to return these specific fibers and index sets, it is hoped that its error is not much larger.

Theorem 3.1. Consider $\mathcal{T} \in \mathbb{R}^{n_{1} \times n_{2} \times n_{3}}$ of multilinear rank at least ( $r_{1}, r_{2}, r_{3}$ ). Let $Q_{U}, Q_{V}, Q_{W}$ denote orthonormal bases of $r_{1}, r_{2}, r_{3}$ selected mode-1,2,3 fibers of $\mathcal{T}$ respectively. Given index sets $I, J, K$ we consider a Tucker decomposition of the form

$$
\hat{\mathcal{T}}=\mathcal{T} \times_{1} Q_{U}\left(\Phi_{I}^{T} Q_{U}\right)^{-1} \Phi_{I}^{T} \times_{2} Q_{V}\left(\Phi_{J}^{T} Q_{V}\right)^{-1} \Phi_{J}^{T} \times_{3} Q_{W}\left(\Phi_{K}^{T} Q_{W}\right)^{-1} \Phi_{K}^{T} .
$$

There exists a choice of fibers and indices such that

$$
\begin{aligned}
\|\mathcal{T}-\hat{\mathcal{T}}\|_{\infty} \leq & \left(\sqrt{q\left(r_{1}, n_{1}\right) \cdot\left(r_{1}+1\right)}+\sqrt{q\left(r_{1}, n_{1}\right) \cdot q\left(r_{2}, n_{2}\right) \cdot\left(r_{2}+1\right)}\right. \\
& \left.+\sqrt{q\left(r_{1}, n_{1}\right) \cdot q\left(r_{2}, n_{2}\right) \cdot q\left(r_{3}, n_{3}\right) \cdot\left(r_{3}+1\right)}\right)\left\|\mathcal{T}-\hat{\mathcal{T}}_{\text {best }}\right\|_{F},
\end{aligned}
$$

where $q(r, n)=\sqrt{1+r(n-r)}$, and $\hat{\mathcal{T}}_{\text {best }}$ is the best Tucker approximation of $\mathcal{T}$ with multilinear rank at most $\left(r_{1}, r_{2}, r_{3}\right)$.

Proof. Using Frobenius norm properties and Lemma 2.1, we obtain

$$
\begin{align*}
\|\mathcal{T}-\hat{\mathcal{T}}\|_{\infty} \leq & \|\mathcal{T}-\hat{\mathcal{T}}\|_{F} \leq\left\|\left(\Phi_{I}^{T} Q_{U}\right)^{-1}\right\|_{2}\left\|\left(I-Q_{U} Q_{U}^{T}\right) \mathcal{T}^{\{1\}}\right\|_{F} \\
& +\left\|\left(\Phi_{I}^{T} Q_{U}\right)^{-1}\right\|_{2}\left\|\left(\Phi_{J}^{T} Q_{V}\right)^{-1}\right\|_{2}\left\|\left(I-Q_{V} Q_{V}^{T}\right) \mathcal{T}^{\{2\}}\right\|_{F} \\
& +\left\|\left(\Phi_{I}^{T} Q_{U}\right)^{-1}\right\|_{2}\left\|\left(\Phi_{J}^{T} Q_{V}\right)^{-1}\right\|_{2}\left\|\left(\Phi_{K}^{T} Q_{W}\right)^{-1}\right\|_{2}\left\|\left(I-Q_{W} Q_{W}^{T}\right) \mathcal{T}^{\{3\}}\right\|_{F} . \tag{3.6}
\end{align*}
$$

From [138, Lemma 2.1] it follows that there exists an index set $I$ such that

$$
\begin{equation*}
\left\|\left(\Phi_{I}^{T} Q_{U}\right)^{-1}\right\|_{2} \leq \sqrt{1+r_{1}\left(n_{1}-r_{1}\right)} \tag{3.7}
\end{equation*}
$$

From [93, Theorem 8] with the role of rows and columns interchanged it follows that we can select mode- 1 fibers $U$ of $\mathcal{T}$ such that

$$
\begin{equation*}
\left\|\left(I-Q_{U} Q_{U}^{T}\right) \mathcal{T}^{\{1\}}\right\|_{F}=\left\|\left(I-U\left(U^{T} U\right)^{-1} U^{T}\right) \mathcal{T}^{\{1\}}\right\|_{F} \leq \sqrt{r_{1}+1}\left\|\mathcal{T}-\hat{\mathcal{T}}_{\text {best }}\right\|_{F} \tag{3.8}
\end{equation*}
$$

Analogous bounds hold for $\left\|\left(\Phi_{J}^{T} Q_{V}\right)^{-1}\right\|_{2},\left\|\left(\Phi_{K}^{T} Q_{W}\right)^{-1}\right\|_{2},\left\|\left(I-Q_{V} Q_{V}^{T}\right) \mathcal{T}^{\{2\}}\right\|_{F}$ and $\left\|\left(I-Q_{W} Q_{W}^{T}\right) \mathcal{T}^{\{3\}}\right\|_{F}$. Applying the bounds (3.7) and (3.8) to the factors in (3.6) yields the claimed result.

Remark 3.2. If one uses orthogonal instead of oblique projections in Phase 3, Corollary 6 in [79] yields a bound similar to Theorem 3.1. Whilst the index sets obtained from DEIM [66] yield small errors in practice, their theoretical upper bounds for $\left\|\left(\Phi_{I}^{T} Q_{U}\right)^{-1}\right\|_{2}$ grow exponentially in $r$. In contrast, the strong rank-revealing $Q R$ decomposition [149] yields an index set for which a bound similar to inequality (3.7) is known [105, Lemma 2.1].

Remark 3.3. Note that the bound in Theorem 3.1 is the worst case bound for the optimal choice of index sets. This does not present the whole picture as even suboptimal index sets might yield a much better approximation in practice. To quantify the quality of the approximation in practice, we computed Chebfun3F approximations for the functions $f(x, y, z)=\log \left(1+x^{2}+y^{2}+z^{2}\right), f(x, y, z)=1 /\left(1+x^{2}+y^{2}+z^{2}\right), f(x, y, z)=\exp (x y z)$ and compared $\left\|\mathcal{T}_{f}-\hat{\mathcal{T}}\right\|_{\infty}$ and $\left\|\mathcal{T}_{f}-\mathcal{T}_{\text {HOSVD }}\right\|_{\infty}$, where $\hat{\mathcal{T}}$ is computed using Chebfun3F and $\mathcal{T}_{\text {HOSVD }}$ denotes the truncated HOSVD of $\mathcal{T}_{f}$ with multilinear ranks equal to those of $\hat{\mathcal{T}}$. We observed that both errors differ by at most a factor of 2. Even though the truncated HOSVD does not focus on $\|\cdot\|_{\infty}$, it still can serve as a good proxy. Hence by Lemma 2.2, the error in Chebfun3F is comparable to the error obtained from the truncated HOSVD.

### 3.2.6 Comparison of the theoretical cost

Assume $f$ can be approximated accurately in Tucker format (2.11) with multilinear rank $(r, r, r)$ and polynomial degrees $(n, n, n), n \geq r$. In a highly idealized setting Chebfun3 and Chebfun3F refine the coarse grid in Phase 1 until $\left(n_{\ell}^{(c)}+1\right)>(2 \sqrt{2}) r$ and identify fibers on this coarse grid. These fibers are refined until $n_{\ell}^{(f)} \geq n$ and lead to Tucker approximations which pass the accuracy check in Phase 3. Under these circumstances, both Chebfun3 and Chebfun3F use $\mathcal{O}\left(r^{3}\right)$ function evaluations in Phase 1 (see Section 2.2 in [163]). In total, Chebfun3 requires $\mathcal{O}\left(n r^{2}\right)$ function evaluations [163, Proposition 2.1]), whereas Chebfun3F only requires $\mathcal{O}\left(r^{3}+n r\right)$, since fewer fibers are refined in Phase 2. We want to emphasize that, in general, it is not guaranteed that this $n_{\ell}^{(c)}$ suffices to identify fibers leading to an accurate approximation. We summarize the breakdown of anticipated costs in Table 3.1.

|  | Chebfun3 | Chebfun3F |
| :---: | :---: | :---: |
| Phase 1 | $\mathcal{O}\left(r^{3}\right)$ | $\mathcal{O}\left(r^{3}\right)$ |
| Phase 2 | $\mathcal{O}\left(n r^{2}\right)$ | $\mathcal{O}(n r)$ |
| Phase 3 | 0 | $\mathcal{O}\left(r^{3}\right)$ |
| Total | $\mathcal{O}\left(r^{3}+n r^{2}\right)$ | $\mathcal{O}\left(r^{3}+n r\right)$ |

Table 3.1 - Anticipated numbers of functions evaluations in Chebfun3 and Chebfun3F.

### 3.3 Numerical results

In this section, we present numerical experiments ${ }^{\dagger}$ to compare Chebfun3F and Chebfun3. The main focus lies on the number of function evaluations required to compute the approximation in Tucker format (2.11). Unless mentioned otherwise the tolerance for the ACA and the accuracy check are initially set close to machine precision.

### 3.3.1 Chebfun3 vs. Chebfun3F

In Section 3.1.4 we illustrated that the function

$$
f(x, y, z)=\frac{1}{1+25 \sqrt{x^{2}+y^{2}+z^{2}}}
$$

leads to a lot of redundant fibers in Chebfun3. For this function, for both Chebfun3 and Chebfun3F, the accuracy check at the end of Phase 3 fails and the computation is restarted on a finer coarse grid, which leads to an approximation that passes the test. Overall 903364 function evaluations are required by Chebfun3, 421041 of them are used before the restart. In comparison, Chebfun3F only needs 222546 function evaluation in total and 100496 before the restart. In the final accuracy check, the estimated error for Chebfun3 is $7.8 \cdot 10^{-13}$ and $3.6 \cdot 10^{-13}$ for Chebfun3F, i.e. both approximations achieve around the same accuracy.

In Figure 3.5, the function evaluations per phase are juxtaposed. The figure shows, that Chebfun3 requires a large number of function evaluations in Phase 2. This is caused by the redundant fibers. In Phase 1, Chebfun3F requires fewer function evaluations, since $\mathcal{T}_{c}$ is not evaluated completely. Whilst Chebfun3 only requires function evaluations in Phase 3 for the accuracy check, Chebfun3F additionally needs to compute the core tensor, which only leads to a small number of function evaluations compared to the other phases.

Remark 3.4. The number of function evaluations required by Chebfun3F depends on the random initialization of the index sets $\tilde{I}, \tilde{J}, \tilde{K}$ in Algorithm 4. We computed the Chebfun3F approximation of $f$ for 1000 different random initializations and observed

[^1]

Figure 3.5 - Comparison of the function evaluations used by Chebfun3 and Chebfun3F to approximate $f(x, y, z)=\left(1+25 \sqrt{x^{2}+y^{2}+z^{2}}\right)^{-1}$. Both algorithms are restarted due to failing the accuracy check once. Evaluations before the restart are depicted in (a), evaluations after the restart in (b). The evaluations are subdivided into phases corresponding to the phases in Section 3.1 and Section 3.2 respectively.
numbers of function evaluations ranging from 213391 to 226073 with mean 221802.6 and variance $4.96 \cdot 10^{6}$. Similarly mild fluctuations have been observed for all other functions tested.

In Figure 3.6, the required function evaluations are depicted for four different functions. The corresponding computing times are depicted in Table 3.2. Again both algorithms lead to approximations of similar accuracy and Chebfun3F requires fewer function evaluations than Chebfun3. For

$$
\begin{equation*}
f_{1}(x, y, z)=\exp \left(-\sqrt{(x-1)^{2}+(y-1)^{2}+(z-1)^{2}}\right) \tag{3.9}
\end{equation*}
$$

Chebfun3 requires the refinement of a huge number of redundant fibers in Phase 2. The evaluations for

$$
\begin{equation*}
f_{2}(x, y, z)=[\cosh (3(x+y+z))]^{-2} \tag{3.10}
\end{equation*}
$$

differ most in Phase 1, where Chebfun3F benefits from not evaluating $\mathcal{T}_{c}$ completely. No additional refinement is required in Phase 2. For the function

$$
\begin{equation*}
f_{3}(x, y, z)=\frac{10^{5}}{1+10^{5}\left(x^{2}+y^{2}+z^{2}\right)} \tag{3.11}
\end{equation*}
$$

Chebfun3F reduces the number of required function evaluations by more than $98 \%$ to 1603693 from 109269332 required by Chebfun3.

In certain cases Chebfun3 outperforms Chebfun3F. This can happen in degenerated situations. For instance, the function $f(x, y, z)=\tanh (5(x+z)) \exp (y)$ requires a Tucker


Figure 3.6 - Total number of function evaluations per Phase in Chebfun3 and Chebfun3F for the functions: (a) $f_{1}$ (b) $f_{2}$, (c) $f_{3}$ as defined in (3.9)-(3.11) (d) the parametric PDE model (3.12)-(3.13) with prescribed tolerance $10^{-9}$.

|  | $f_{1}$ | $f_{2}$ | $f_{3}$ | PDE |
| :---: | :---: | :---: | :---: | :---: |
| Chebfun3 | 3.44 | 3.20 | 15.80 | 408.62 |
| Chebfun3F | 0.30 | 0.77 | 0.56 | 100.19 |

Table 3.2 - Computing times in MATLAB in seconds for Chebfun3 and Chebfun3F approximations of $f_{1}, f_{2}, f_{3}$ as defined in (3.9)-(3.11) and of the parametric PDE model (3.12)(3.13) with prescribed tolerance $10^{-9}$.
decomposition with rank $r=[71,1,71]$. In this case, Chebfun3F heavily relies on the heuristic to increase ( $r_{1}, r_{2}, r_{3}$ ) when restarting and requires 1641712 function evaluations compared to 1128061 in Chebfun3. Other functions that are difficult to approximate with Chebfun3F include (numerically) locally supported functions, such as a trivariate normal distribution with very small entries in the covariance matrix, for which identifying non-zero fibers in Phase 1 without fully evaluating $\mathcal{T}_{c}$ might be difficult.

## Application: Uncertainty quantification

Algorithms in uncertainty quantification, such as the Metropolis-Hastings method, often require the repeated evaluations of a parameter depended quantity of interest [307]. In many applications, the evaluation of this quantity requires the solution of a PDE depending on the parameters. To speed up computations, the mapping from the parameters to the quantity of interest is often replaced by a surrogate model [341]. In the context of models with three parameters (or after a dimension reduction to three parameters [77]), Chebfun3/Chebfun3F could be a suitable surrogate.

We consider the parametric elliptic PDE model problem on $\Omega=[1,3]^{2}$

$$
\begin{align*}
\nabla \cdot\left(\left(p_{1} f_{1}(x, y)+p_{2} f_{2}(x, y)+p_{3} f_{3}(x, y)\right) \nabla u(x, y)\right) & =1  \tag{3.12}\\
u(x, y) & =0 \quad(x, y) \in \Omega,  \tag{3.13}\\
& \in \partial \Omega,
\end{align*}
$$

with parameters $\left(p_{1}, p_{2}, p_{3}\right) \in[1,3]^{3}$ and functions $f_{1}(x, y)=\cos (x)+\sin (y)+2, f_{2}(x, y)=$ $\sin (x)+\cos (y)+2$, and $f_{3}(x, y)=\cos \left(x^{2}+y^{2}\right)+2$. The quantity of interest is defined as point evaluation $u(1.5,1.5)$. With Chebfun3F, we only need 3217 PDE solves to compute an approximation with prescribed accuracy $10^{-9}$, whereas Chebfun3 requires 7626 as depicted in Figure 3.6(d).

### 3.3.2 Comparison to sparse grids

Lastly, we study how efficient Chebfun3 approximations are compared to sparse grids [50]. Sparse grids are a method to interpolate functions by projecting them onto a particular space. This space is obtained by selecting the most beneficial elements from a hierarchical basis. In the following, we use dimension-adaptive sparse grids based on a Chebyshev-Gauss-Lobatto grid with polynomial basis functions as implemented in the sparse grid interpolation toolbox [195].

We compare how the approximation error decays compared to the number of function evaluations. Therefore, we prescribe varying tolerances to the algorithms. In Figure 3.7 the error decay is plotted for sparse grids, Chebfun3 and Chebfun3F. In (a), the function already studied in section 3.1.4 is depicted. We observe that both Chebfun3F and Chebfun3 require fewer function evaluations than sparse grids to achieve the same accuracy. The sparse grids perform poorly, since the function is smooth, but the norms of the second mixed derivatives are rather large. In contrast, the sum of 10 Gaussians depicted in (b) is well suited for sparse grids. In this case sparse grids require fewer function evaluations than Chebfun3F and Chebfun3 for the same accuracy. In (c) the Chebfun3F and sparse grids perform about equally well for

$$
\begin{equation*}
f_{4}(x, y, z)=\log (x+y z+\exp (x y z)+\cos (\sin (\exp (x y z)))) \tag{3.14}
\end{equation*}
$$



Figure 3.7 - Comparison of the number of function evaluations required to compute a Chebfun3, Chebfun3F and sparse grid approximation for the functions: (a) $f(x, y, z)=$ $\left(1+25 \sqrt{x^{2}+y^{2}+z^{2}}\right)^{-1}$, (b) sum of 10 Gaussians, (c) $f_{4}$ as defined in (3.14). The algorithms are initialized with varying tolerances. Their $\mathcal{L}^{2}$ error is estimated at 1000 sample points.

Note that the Chebfun3F heuristics only require a restart for a tolerance of around $10^{-7}$. This leads to the outlier in the Chebfun3F graph. For an arbitrary, black-box function it is not clear a priory whether a sparse grid interpolation or a Tucker decomposition (2.11) is the more efficient type of approximation. Note that when the Tucker decomposition is the better approximation format, we can expect that Chebfun3F requires fewer function evaluations compared to Chebfun3.

We conclude that trivariate functions defined on tensor product domains can be approximated efficiently by combining tensorized Chebyshev interpolation and a low-rank Tucker approximation of the evaluation tensor. Our numerical experiments demonstrate that Chebfun3F typically requires fewer function evaluations to compute such an approximation
of the same accuracy compared to Chebfun3.

### 3.4 Coefficient decay analysis

There is one major difference between the approximation computed using Chebfun3 and Chebfun3F as presented in Algorithm 5. The compression step in Phase 3 of Chebfun3 ensures that the Tucker approximation of the coefficient tensor (2.22) contains factor matrices with entries decaying to close to machine precision. This is not the case of Chebfun3F, where the corresponding factor matrices are of the form $F^{(1)} Q_{U}$ or ${ }^{(1)} Q_{U}\left(\Phi_{I}^{T} Q_{U}\right)^{-1}$ depending on if one treats the terms of the form $\left(\Phi_{I}^{T} Q_{U}\right)^{-1}$ in (3.5) as part of the factor matrix or as part of the core tensor. See Figure 3.8 for an example of this discrepancy.


Figure 3.8 - We plot the decay of the univariate Chebyshev interpolation coefficients (2.13) of the functions $u_{i}^{(1)}$ in (2.24) for the approximations of $f(x, y, z)=\sin (\exp (x+y+z))$ computed using Chebfun3 (left) and Chebfun3F (right). The different colors correspond to different values of $i \in\left\{1, \ldots, r_{1}\right\}$.

Many functions to perform computations with univariate functions in the Chebfun package assume that the function is represented in terms of coefficients (2.13) that decay until reaching values close to machine precision. In Chebfun3F the coefficient tensor (2.19) computed from (3.5) has decaying coefficients when studied as a whole, but the coefficients stored in the factor matrices do not decay to machine precision when the core tensor is not taken into account. This implies that applying operations such as e.g. computing a derivative by a mode- $\ell$ multiplication with a differentiation matrix (see Section7.2) are working as expected when the whole coefficient tensor is taken into account. However, when we only consider the multiplication of the $\ell$ th factor matrix with the differentiation matrix, heuristics [16] would state that the operation has not been performed accurately. Note that it might not be feasible to analyse the coefficient decay of the whole coefficient tensor in practice. In the following, we discuss how Chebfun3F approximations can be modified to circumvent this issue. Note that we can not simply replace $Q_{U}\left(\Phi_{I}^{T} Q_{U}\right)^{-1}$
in (3.5) by $U_{f}\left(\Phi_{I}^{T} U_{f}\right)^{-1}$, since this would lead to very ill-conditioned matrices $\Phi_{I}^{T} U_{f}$.

### 3.4.1 Observations in a two-dimensional setting

To study the aforementioned issue, we first analyze the problem in a two-dimensional setting using a fixed polynomial degree $(n, n)$. Throughout this section, we denote the evaluation matrix (2.18) by $T$ and the transformation matrix (2.14) by $F$. All numerical experiments are performed using the function $f(x, y)=1 /\left(1+25\left(x^{2}+y^{2}\right)\right)$ and $n=257$.

Gaussian elimination In Chebfun2 [320], the matrix $T$ is approximated using Gaussian elimination. The approximation is constructed iteratively using ACA. Instead of storing index sets $I, J$ the approximation (2.1) is stored as sum of the rank-1 matrices subtracted in each iteration in line 7 of Algorithm 1 . This can be written as $T \approx U_{g} C_{g} V_{g}^{T}$, where norms of the columns in $U_{g}, V_{g}$ decay rapidly and the matrix $C_{g} \in \mathbb{R}^{r \times r}$ is a diagonal matrix which contains the inverses of the pivot elements determined in line 5 . It is important to note that the columns in $U_{g}$ are not directly fibers of the original function. The matrix $C_{g}$ is very ill-conditioned, but the large entries in $C_{g}$ cancel out with the small norms in the corresponding columns of $U_{g}$ and $V_{g}$. The final approximation contains univariate interpolants (2.13) with coefficients given by the columns of $F U_{g}$, which decay as depicted in Figure 3.9(a).

Adaptive cross approximation If we were to apply ACA as presented in Algorithm 1 to $T$, we would obtain an approximation of the form $T \approx T(:, \tilde{J}) T(\tilde{I}, \tilde{J})^{-1} T(\tilde{I},:)$. Let $U=T(:, \tilde{J}), V=T(\tilde{I},:)^{T}$. This leads to a final approximation containing Chebfuns with coefficients given by the columns of $F U$, which decay as depicted in Figure 3.9(b). The final approximation can be evaluated with an accuracy up to machine precision by using the backslash operator of MATLAB to avoid inverting $T(\tilde{I}, \tilde{J})^{-1}$. However, we have to point out that the inverse $T(\tilde{I}, \tilde{J})^{-1}$ (computed using MATLAB) is an ill-conditioned matrix with condition number around $10^{15}$, which when used directly would lead to errors of order $10^{0}$ in point evaluation of the approximation. In fact, there we could not find any matrix $C \in \mathbb{R}^{r \times r}$ such that $\left\|T-U C V^{T}\right\|_{F}^{2} \leq 10^{-8}$ in double precision.

Note that neither Gaussian eliminations of the form $U_{g} C_{g} V_{g}^{T}$ nor the inversion using backslash can be generalized directly to the three-dimensional setting.

Two-dimensional Chebfun3F The two-dimensional equivalent of Chebfun3F would be to compute QR decompositions $U=Q_{U} R_{U}, V=Q_{V} R_{R}$ based on which we construct a decomposition of the form $T \approx Q_{U} \tilde{C} Q_{V}^{T}$, where $\tilde{C} \in \mathbb{R}^{r \times r}$. In contrast to the previous paragraph, where it was not possible to find a suitable matrix $C$, we can now find matrices $\tilde{C}$ that lead to accurate approximations of $T$. Using oblique projections based
on DEIM (2.3) we obtain $\tilde{C}=Q_{U}(I,:)^{-1} T(I, J) Q_{V}(J,:)^{-T}$, where $I, J$ are index sets obtained by applying DEIM to $Q_{U}$ and $Q_{V}$. This is the only approach so far that can be generalized to higher dimensions, but the interpolation coefficients in the columns of the factor matrix $F Q_{U}$ decay as depicted Figure 3.9(c) and would not be considered proper approximations of univariate functions in Chebfun.


Figure 3.9 - Decay of the coefficients stored in the columns of $F U_{g}, F U, F Q_{U}, F \bar{Q}_{U}$ defined in Sections 3.4.1 and 3.4.2.

### 3.4.2 Potential solutions in the two-dimensional setting

From the previous section, we conclude that we can not use the matrices $U$ and $V$ directly as factor matrices. At the same time, we should not use the matrices $Q_{U}$ and $Q_{V}$ directly due to the bad decay of the coefficients in the factor matrices. In the following, we propose two possible approaches to solve this issue, which lead to approximations as accurate as the approximation $Q_{U} \tilde{C} Q_{V}^{T}$. Both approaches can be generalized to higher
dimensional settings.

Storing the $R$ factors in addition to the core tensor The first approach is to compute the approximation $T \approx Q_{U} \tilde{C} Q_{V}^{T}$. We now insert the matrices $R_{U}, R_{V}$ to obtain an approximation of the form $T \approx U R_{U}^{-1} \tilde{C} R_{V}^{-T} V^{T}$, which contains the nicely decaying factor matrices $U, V$. This can be seen as approximation with core matrix $\hat{C}=R_{U}^{-1} \tilde{C} R_{V}^{-T}$, but we can not compute $\hat{C}$ explicitly due to the issues described in Section 3.4.1. We found that we can instead store $R_{U}, R_{V}$ separately. This allows us to accurately evaluate the approximation by using the following order of operations $\left((F U) R_{U}^{-1}\right) \tilde{C}\left(R_{V}^{-T}\left(V^{T} F^{T}\right)\right)$. At the same time, storing the additional terms $R_{U}, R_{V}$ is the main disadvantage of the approach since we can no longer explicitly form the approximation (2.24).

Rescaling of the $Q$ factors Starting from the approximation of the form $T \approx$ $Q_{U} \tilde{C} Q_{V}^{T}$, we can modify the coefficient decay by introducing diagonal scaling matrices $D_{U}, D_{V} \in \mathbb{R}^{r \times r}$. This leads to an approximation of the form $T \approx \bar{Q}_{U} \bar{C} \bar{Q}_{V}^{T}$, where $\bar{Q}_{U}=Q_{U} D_{U}, \bar{Q}_{V}=Q_{V} D_{V}$ and $\bar{C}=D_{U}^{-1} C D_{V}^{-1}$. We found that setting $\left(D_{U}\right)_{k, k}=\varepsilon / \max \left(\varepsilon, \min \left(\left|F Q_{U}(:, k)\right|\right)\right) \operatorname{and}\left(D_{V}\right)_{k, k}=\varepsilon / \max \left(\varepsilon, \min \left(\left|F Q_{V}(:, k)\right|\right)\right)$ for $k=1, \ldots, r$ leads to a final approximation containing Chebfuns with coefficients $F \bar{Q}_{U}$ and $F \bar{Q}_{V}$, which satisfy the Chebfun heuristic [16] and can, thus, be considered proper Chebfun objects; see Figure 3.9(d).

We would like to point out that this rescaling approach does not lead to a loss in accuracy when evaluating the resulting approximation. As in Chebfun2, we obtain an evaluation error of order $10^{-13}$. Also note that if we were to set all coefficients in $\bar{Q}_{U}$ with absolute value smaller than $10^{-13}$ to zero, we would still achieve an evaluation error of order $10^{-12}$.

### 3.4.3 Incorporation of Chebfun3F in Chebfun

In Chebfun3F we can avoid the issue, discussed in Section 3.4, that the coefficients of the factor matrices do not decay by generalizing the idea of rescaling the $Q$ factors described in Section 3.4.2. For this purpose, we modify Phase 3 in Algorithm 5 in the following way. We set $U=D_{U} Q_{U}, V=D_{V} Q_{V}, W=D_{W} Q_{W}$ and $\mathcal{C}=\mathcal{T}_{f}(I, J, K) \times_{1} D_{U} Q_{U}(I,:)^{-1} \times_{2}$ $D_{V} Q_{V}(J,:)^{-1} \times_{3} D_{w} Q_{W}(K,:)^{-1}$, where the diagonal matrices $D_{U}, D_{V}, D_{W}$ are defined as $\left(D_{U}\right)_{i, i}=\varepsilon / \max \left(\varepsilon, \min \left(\left|F^{(1)} Q_{U}(:, i)\right|\right)\right),\left(D_{V}\right)_{j, j}=\varepsilon / \max \left(\varepsilon, \min \left(\left|F^{(2)} Q_{V}(:, j)\right|\right)\right)$ and $\left(D_{W}\right)_{k, k}=\varepsilon / \max \left(\varepsilon, \min \left(\left|F^{(3)} Q_{W}(:, k)\right|\right)\right)$ for $i=1, \ldots, r_{1}, j=1, \ldots, r_{2}, k=1, \ldots, r_{3}$. Repeating the experiment presented in Figure 3.8 with this modified version of Chebfun3F yields the coefficient decay depicted in Figure 3.10. Observe that the coefficients now decay to almost machine precision as desired.

In collaboration with Nick Trefethen, Nick Hale, Behnam Hashemi and Alex Townsend, we incorporated this modified version of the Chebfun3F algorithm as alternative constructor


Figure 3.10 - We plot the decay of the univariate Chebyshev interpolation coefficients (2.13) of the functions $u_{i}^{(1)}$ in (2.24) for the approximation of $f(x, y, z)=\sin (\exp (x+y+z))$ computed using the modified version of Chebfun3F presented in Section 3.4.3.
for Chebfun3 objects into the official Chebfun package. Our alternative constructor can be called using chebfun3(f,'chebfun3f') and returns a standard Chebfun3 object.

Remark 3.5. As in Chebun3 [163], we included a coefficient truncation [16] in the modified version of Chebfun3F. Such a truncation is not included in the original version of Chebfun3F presented in Algorithm 5. This causes the difference in the polynomial degrees in Figure 3.8 and 3.10.

### 3.5 Additional insights

### 3.5.1 Approximations in Lebesgue spaces

In Remark 2.12, we discuss how an approximation of the evaluation tensor (2.18) affects the $L^{1}$-norm of the corresponding interpolants. In the following, we study this relationship for the $L^{2}$-norm.

Let $p:[-1,1]^{3} \rightarrow \mathbb{R}$ be defined based on the evaluation tensor $\mathcal{T} \in \mathbb{R}^{(n+1) \times(n+1) \times(n+1)}$ as in (2.16). Let $\hat{p}$ be defined as in (2.21) based on the tensor $\hat{\mathcal{T}} \in \mathbb{R}^{(n+1) \times(n+1) \times(n+1)}$. By using tensorized Clenshaw-Curtis quadrature as in Remark 2.12 we obtain

$$
\begin{aligned}
\|p-\hat{p}\|_{L^{2}}^{2} & =\int_{-1}^{1} \int_{-1}^{1} \int_{-1}^{1}(p(x, y, z)-\hat{p}(x, y, z))^{2} d x d y d z \\
& =\sum_{i=0}^{n} \sum_{j=0}^{n} \sum_{k=0}^{n} \mathcal{W}_{i, j, k}\left(T_{i, j, k}-\hat{T}_{i, j, k}\right)^{2}+\varepsilon_{\text {quad }}
\end{aligned}
$$

where $\varepsilon_{\text {quad }}$ denotes the quadrature error. The quadrature error occurs, since $(p-\hat{p})^{2}$ is a polynomial of degree up to $(2 n, 2 n, 2 n)$, but Clenshaw-Curtis quadrature is only exact
up to degree $(n, n, n)$ [325]. However, when the entries in the coefficients tensors $\mathcal{A}, \hat{\mathcal{A}}$ defined as in (2.19) based on $\mathcal{T}, \hat{\mathcal{T}}$ decay sufficiently quickly, we can assume that $\varepsilon_{\text {quad }}$ is negligible [27]. Under this strong assumption, we have

$$
\|p-\hat{p}\|_{L^{2}}^{2} \approx \sum_{i=0}^{n} \sum_{j=0}^{n} \sum_{k=0}^{n} \mathcal{W}_{i, j, k}\left((\mathcal{T}-\hat{\mathcal{T}})_{i, j, k}\right)^{2}=\|\sqrt{\mathcal{W}} *(\mathcal{T}-\hat{\mathcal{T}})\|_{F}^{2} .
$$

In order to obtain a good approximation of $p$ in the format (2.24) with $\mathcal{C} \in \mathbb{R}^{r \times r \times r}$, we try to minimize $\|\sqrt{\mathcal{W}} *(\mathcal{T}-\hat{\mathcal{T}})\|_{F}^{2}$. Analogous to Remark 2.12 , we can achieve this by computing an approximation of $\sqrt{\mathcal{W}} \mathcal{T}$ in Tucker format (2.5). Let $\tilde{\mathcal{T}}_{\text {HOSVD }}$ be obtained by applying a truncated HOSVD [88] to $\sqrt{\mathcal{W}} \mathcal{T}$. Let

$$
\left(\hat{\mathcal{T}}_{\text {HOSVD }}\right)_{i_{1}, \ldots, i_{d}}= \begin{cases}\left(\tilde{\mathcal{T}}_{\text {HOSVD }}\right)_{i_{1}, \ldots, i_{d}} / \sqrt{\mathcal{W}_{i_{1}, \ldots, i_{d}}} & \mathcal{W}_{i_{1}, \ldots, i_{d}} \neq 0 \\ \left(\tilde{\mathcal{T}}_{\text {HOSVD }}\right)_{i_{1}, \ldots, i_{d}} & \mathcal{W}_{i_{1}, \ldots, i_{d}}=0 .\end{cases}
$$

for $i_{\ell}=0, \ldots, n, \ell=1, \ldots, d$. Using the tensor $\hat{\mathcal{T}}_{\text {Hosvd }}$ to define $\hat{p}$ yields a close to optimal approximation of $p$ in the $L^{2}$-norm. Note that $\hat{\mathcal{T}}_{\text {HOSVD }}$ has multilinear rank at most $(r, r, r)$ and can thus be represented in Tucker format with with $\mathcal{C} \in \mathbb{R}^{r \times r \times r}$.

In Figure 3.11 we demonstrate that the singular values of $(\sqrt{\mathcal{W}} \mathcal{T})^{\{1\}}$ are closely related to the error $\|p-\hat{p}\|_{L^{2}}$. Surprisingly, both applying the HOSVD directly to $\mathcal{T}$ and using the tensor $\hat{\mathcal{T}}_{\text {HOSVD }}$ yield polynomials with almost identical errors. However, it is not possible to precisely determine the accuracy of $\|p-\hat{p}\|_{L^{2}}$ based on the singular values of $\hat{\mathcal{T}}^{\{1\}}$. We conclude that the we can accurately estimate the multilinear rank required for accurate approximations of the form (2.24) with respect to the $L^{2}$-norm by truncating the singular values of $\ell$-mode matricizations of $\sqrt{\mathcal{W}} \mathcal{T}$ for sufficiently large $n$. For computing an approximation with fixed multilinear rank it does not seem to be necessary to take the weight tensor $\sqrt{\mathcal{W}}$ into account.

### 3.5.2 Further compression

In this section, we want to gain additional insight into the impact of the size of the core tensor on approximations of the form (2.24). Moreover, we study if it is possible to further compress the Tucker approximation (2.5) of the evaluation tensor (2.18). This section is inspired by Rachel Minster's talk [229] on the paper [278] at SIAM-ALA 2021.

We are interested in evaluating a kernel function $k: \mathbb{R}^{2} \times \mathbb{R}^{2} \rightarrow \mathbb{R}$ at all pairs $\left(x^{(i)}, y^{(j)}\right)$ for $i=1, \ldots, M, j=1, \ldots, N$, where $x^{(i)} \in[a, b]^{2}, y^{(j)} \in[c, d]^{2}, a, b, c, d \in \mathbb{R}$ such that $[a, b] \cap[c, d]=\emptyset$. Let $K \in \mathbb{R}^{M \times N}$ denote the matrix of evaluations, i.e. $K_{i, j}=k\left(x^{(i)}, y^{(j)}\right)$. The direct computation of $K$ requires $M \cdot N$ kernel evaluations. In the following, we propose a method to approximate $K$ using fewer kernel evaluations.


Figure 3.11 - Let $\mathcal{T}$ denote the sample tensor of $f(x, y, z)=\cosh 3(x+y+z)^{-1}$ on a $75 \times 75 \times 75$ Chebyshev grid. We depict the singular value decay of $\mathcal{T}^{\{1\}}$ and of $(\sqrt{\mathcal{W}} \mathcal{T})^{\{1\}}$. The interpolant (2.16) based on $\mathcal{T}$ is denoted by $p$. The interpolant based $\hat{\mathcal{T}}_{\text {HOSVD }}$ with core tensor in $\mathbb{R}^{r \times r \times r}$ is denoted by $\hat{p}_{r}^{(\sqrt{\mathcal{W}} \mathcal{T})}$ and the interpolant obtained from a direct HOSVD of $\hat{T}$ with core tensor in $\mathbb{R}^{r \times r \times r}$ is denoted by $\hat{p}_{r}^{(\mathcal{T})}$. We plot estimations of $\left\|p-\hat{p}_{r}^{(\sqrt{\mathcal{W}} \mathcal{T})}\right\|_{L^{2}}$ and $\left\|p-\hat{p}_{r}^{(\mathcal{T})}\right\|_{L^{2}}$ computed using 10000 sample points for different values of $r$. Note that the estimated $L^{2}$-errors almost coincide for the same values of $r$.

We proceed by approximating $k$ as a function using tensorized interpolation (2.16) and approximations of the coefficient tensor (2.22) to obtain an approximation of the form

$$
k\left(x_{1}, x_{2}, y_{1}, y_{2}\right) \approx p\left(x_{1}, x_{2}, y_{1}, y_{2}\right)=\sum_{i=0}^{n_{1}} \sum_{j=0}^{n_{2}} \sum_{k=0}^{n_{3}} \sum_{l=0}^{n_{4}} \hat{\mathcal{A}}_{i, j, k, l} T_{i}\left(x_{1}\right) T_{j}\left(x_{2}\right) T_{k}\left(y_{1}\right) T_{l}\left(y_{2}\right)
$$

where $\hat{\mathcal{A}}$ is derived from an approximation $\hat{\mathcal{T}} \in \mathbb{R}^{\left(n_{1}+1\right) \times \cdots \times\left(n_{4}+1\right)}$ of the evaluation tensor (2.18). Given $\hat{\mathcal{T}}$ we can compute the approximation

$$
\begin{equation*}
\tilde{K}=E_{x}\left(F^{(2)} \otimes F^{(1)}\right) \hat{\mathcal{T}}^{<2>}\left(F^{(4)} \otimes F^{(3)}\right) E_{y}^{T} \tag{3.15}
\end{equation*}
$$

of the kernel matrix $K$, where $\otimes$ denotes the Kronecker product, $F^{(\ell)}$ is defined as in Equation (2.14), and $E_{x} \in \mathbb{R}^{M \times\left(n_{1}+1\right)\left(n_{2}+1\right)}$ and $E_{y} \in \mathbb{R}^{N \times\left(n_{3}+1\right)\left(n_{4}+1\right)}$ encode the evaluations of the Chebyshev polynomials at the points of interest, i.e.

$$
\begin{aligned}
& \left(E_{x}\right)_{i, j+k\left(n_{1}+1\right)}=T_{j}\left(x_{1}^{(i)}\right) \cdot T_{k}\left(x_{2}^{(i)}\right) \text { for } i=1, \ldots, M, j=0, \ldots, n_{1}, k=0, \ldots, n_{2} \\
& \left(E_{y}\right)_{i, j+k\left(n_{1}+1\right)}=T_{j}\left(y_{1}^{(i)}\right) \cdot T_{k}\left(y_{2}^{(i)}\right) \text { for } i=1, \ldots, N, j=0, \ldots, n_{3}, k=0, \ldots, n_{4}
\end{aligned}
$$

Note that multiplying the kernel approximation (3.15) with a vector requires $\mathcal{O}((M+$
N) $n^{2}+n^{4}$ ) operations when $n_{1}=n_{2}=n_{3}=n_{4}=n$. This is cheaper than multiplying $K$ directly with a vector when $(M+N) \gg n^{2}$.

Remark 3.6. Note that a generalization of these ideas to kernels $k: \mathbb{R}^{d} \times \mathbb{R}^{d} \rightarrow \mathbb{R}$ for arbitrary $d \in \mathbb{N}$ is straightforward, but generalizing Equation (3.15) is only beneficial for matrix vector products when $(M+N) \gg n^{d}$.

In [229], it is suggested to obtain $\hat{\mathcal{T}}$ in Tucker format (2.5) using a non-adaptive approach similar to Algorithm 5. First, the span of each matrizisation of $\mathcal{T}$ is approximated to obtain the factor matrices. For this purpose the randomized row interpolatory decomposition algorithm (RRID) [156] (see Algorithm 6) is applied to matricizations of subtensors of $\mathcal{T}$. The RRID algorithm immediately yields index sets that can be used to define the core tensor of the Tucker approximation using oblique projections. The overall kernel approximation algorithm is formalized in Algorithm 7.

```
Algorithm 6 Randomized row interpolatory decomposition
    Input: matrix \(M \in \mathbb{R}^{m \times n}\), target rank \(r\), oversampling parameter \(p\)
    Output: approximation \(M \approx U M(J,:)\)
    sample a Gaussian random matrix \(\Omega \in \mathbb{R}^{n \times(r+p)}\)
    compute a thin QR factorization \(Q R=M \Omega\)
    \(Q=Q(:, 1: r)\)
    apply a pivoted thin QR factorization to \(Q^{T}\) and let \(J\) denote the set of pivot elements
    \(U=Q(Q(J,:))^{-1}\)
```

```
Algorithm 7 Kernel approximation
    Input: kernel function \(k:[a, b]^{2} \times[c, d]^{2} \rightarrow \mathbb{R}\), polynomial degree \(\left(n_{1}, n_{2}, n_{3}, n_{4}\right)\),
    size of the core tensor \(r_{1} \times r_{2} \times r_{3} \times r_{4}\), index sets \(I_{\ell} \subseteq\left\{0, \ldots, n_{\ell}\right\}\) for \(1 \leq \ell \leq 4\)
    Output: approximation \(\hat{\mathcal{T}}\) of the evaluation tensor (2.18) in Tucker format (2.5)
    Let \(\mathcal{T}\) denote the tensor containing the evaluations of \(k\) on a \(\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times\)
    \(\left(n_{3}+1\right) \times\left(n_{4}+1\right)\) Chebyshev grid mapped onto the domain \([a, b]^{2} \times[c, d]^{2}\).
    Let \(M_{1}=T\left(:, I_{2}, I_{3}, I_{4}\right)^{\{1\}}, M_{2}=T\left(I_{1},:, I_{3}, I_{4}\right)^{\{2\}}, M_{3}=T\left(I_{1}, I_{2},:, I_{4}\right)^{\{3\}}, M_{4}=\)
    \(T\left(I_{1}, I_{2}, I_{3},:\right)^{\{4\}}\).
    for \(\ell=1, \ldots, 4\)
        Apply RRID with target rank \(r_{\ell}\) to obtain the approximation \(M_{\ell} \approx U_{\ell} M_{\ell}\left(J_{\ell},:\right)\).
    \(\hat{\mathcal{T}}=\mathcal{T}\left(J_{1}, J_{2}, J_{3}, J_{4}\right) \times{ }_{1} U_{1} \times_{2} U_{2} \times{ }_{3} U_{3} \times_{4} U_{4}\)
```

In the following numerical experiment, we consider the kernel function $k(x, y)=\frac{1}{\|x-y\|_{2}}$. We sample $x_{i}$ uniformly from $[0,1]^{2}$ for $1 \leq i \leq M=200$ and $y_{j}$ uniformly from $[2,3]^{2}$ for $1 \leq j \leq N=150$. In Figure 3.12 (a) we study the approximation error $\|K-\tilde{K}\|_{2}$ for approximations obtained using Equation (3.15), where $\hat{\mathcal{T}}$ is computed using Algorithm 7 with different prescribed core tensor sizes. We observe that error of tensorized interpolation (2.16) with polynomial degree ( $r, r, r, r$ ) without any approximation of the evaluation tensor yields the same approximation error as Algorithm 7 with core tensor in $\mathbb{R}^{r \times r \times r \times r}$ for any polynomial degree $(n, n, n, n)$ with $n>r$. This implies that the error
is dominated by the chosen core tensor size for this kernel function. Figure 3.12(b) shows that the matrix $\tilde{\mathcal{T}}<2>$ used in (3.15) does not seem to be full rank. Truncating the SVD of $\tilde{\mathcal{T}}^{<2>}$ with rank larger than $2 r$ does not improve the approximation error when the core tensor is in $\mathbb{R}^{r \times r \times r \times r}$. Thus, it is possible to further compress the approximation (3.15). So far, we did not analyze the impact of the index sets $I_{1}, \ldots, I_{4}$ in Algorithm 7. In Figure 3.12(c), we observe that choosing $m>2$ indices in each of the index sets $I_{1}, \ldots, I_{4}$ does not improve the error for core tensors in $\mathbb{R}^{5 \times 5 \times 5 \times 5}$, i.e. the $m^{3}=8$ randomly selected fibers in each mode suffice to find suitable factor matrices. For core tensors in $\mathbb{R}^{10 \times 10 \times 10 \times 10}$ and $\mathbb{R}^{15 \times 15 \times 15 \times 15}$, the approximation error decreases for $m$ increasing from 2 to 4 and stagnates afterwards. This demonstrates, that we only need slightly more than $r_{\ell}$ fibers to identify good factor matrices for a given $r_{\ell}$. At the same time, it confirms that we do not need to fully evaluate $\mathcal{T}$ to compute suitable factor matrices.

We conclude this section by summarizing our observations for general functional Tucker approximations of the form (2.24). Note that these findings also apply to Chebfun3 and Chebfun3F. While it can be beneficial to use a smaller core tensor compared to the multilinear rank as demonstrated in Section 5.5.1, we showed that this is not necessarily the case for all functions as demonstrated in Figure 3.12(a). A functional Tucker approximation of the form (2.24) can potentially be compressed further to minimize the required storage. This motivates the algorithms developed in Chapter 4, where we compress approximations (2.24) further by computing TT approximations of the core tensor. Lastly, we observe in Figure 3.12(c) that randomly selected subtensors (as used in Algorithm 4) tend to be sufficient to compute factor matrices that lead to accurate Tucker decompositions.


Figure 3.12 - Comparison of $\|K-\tilde{K}\|_{2}$ for the setting in Section 3.5.2. Top Left: The matrix $\tilde{K}(3.15)$ is computed based on $\tilde{\mathcal{T}}$ obtained using Algorithm 7 with polynomial degree $(n, n, n, n)$ and core tensor in $\mathbb{R}^{r_{1} \times r_{2} \times r_{3} \times r_{4}}$ as shown in the legend. The index sets in Algorithm 7 are set to $I_{\ell}=\left\{0, \ldots, n_{\ell}\right\}$ for $\ell=1, \ldots, 4$. Additionally plot the error when $\tilde{\mathcal{T}}$ is computed using tensorized interpolation (2.16) without any approximation of the evaluation tensor (full interpolation). Top Right: The tensor $\tilde{\mathcal{T}}$ obtained using Algorithm 7 with polynomial degree $(20,20,20,20)$ and core tensor in $\mathbb{R}^{r_{1} \times r_{2} \times r_{3} \times r_{4}}$ as shown in the legend. The index sets in Algorithm 7 are set to $I_{\ell}=\{0, \ldots, 20\}$ for $\ell=1, \ldots, 4$. We then compute a truncated SVD of $\tilde{\mathcal{T}}{ }^{<2>}$ with truncation rank $r_{\text {SVD }}$. We compute the matrix $\tilde{K}(3.15)$ based on this compressed version of $\tilde{\mathcal{T}}<2>$. Bottom: The matrix $\tilde{K}$ is computed using Algorithm 7 with polynomial degree ( $20,20,20,20$ ) and core tensor in $\mathbb{R}^{r_{1} \times r_{2} \times r_{3} \times r_{4}}$ as shown in the legend. The index sets $I_{\ell}$ each contain $m$ indices uniformly sampled from $\{0, \ldots, 20\}$.

## 4 Extended functional tensor train approximation

This chapter is concerned with the notoriously difficult task of approximating multivariate functions on the tensor product domain $[-1,1]^{d}$ for large values of $d$. In principle, we could extend the algorithms presented in Chapter 3 to functions depending on more than three variables. However, the functional Tucker format (2.11) is not well suited for the high-order tensors arising from the evaluation of a function in many variables due to the size of the core tensor scaling exponentially with respect to $d$. Other formats, such as the functional tensor train (FTT) format (2.12) and the extended functional tensor train (EFTT) format proposed in this chapter, are better suited for this purpose and will require different construction algorithms.

As described in Section 2.3.4, approximations in the FTT format can be obtained by combining tensorized Chebyshev interpolation (2.16) with Algorithm 3 to compute a tensor train decomposition (2.6) of the coefficient tensor (2.19) [41]. A slightly different approach to obtain FTT approximations is proposed in [141], where a continuous version of the TT-cross algorithm [251] is applied directly to the function. All univariate functions occurring in this process are discretized using parameterizations such as basis expansions. In principle, every univariate function could be parameterized differently. In the special case that univariate interpolation with the same basis functions is used to discretize all univariate functions in the same mode, this approach is equivalent to a TT approximation of the coefficient tensor for tensorized interpolation.

The FTT format (2.12) requires the storage of $R_{\ell-1} \cdot R_{\ell}$ univariate functions for each $\ell=1, \ldots, d$. These functions (or their parametrizations) will often be linearly dependent. In this case, we can compress the format further by storing these functions in terms of linear combinations of $r_{\ell} \ll R_{\ell-1} \cdot R_{\ell}$ basis functions. The resulting approximation can be seen as a variant of the hierarchical Tucker format proposed in [152]. When the approximation is constructed from a low-rank approximation of the coefficient tensor, we can obtain such a compressed format by following the ideas of Khoromskij [186]. There a two-level Tucker format is proposed, which is obtained by first computing a

Tucker approximation [327] before approximating the core tensor in $\mathbb{R}^{r_{1} \times \cdots \times r_{d}}$ further. Computing a TT approximation of the core tensor leads to a so-called extended TT format [99, 108, 286], from which we can derive a functional low-rank approximation in the EFTT format. This EFTT approximation corresponds to a compressed FTT approximation.

In this chapter, we propose a novel algorithm to efficiently compute functional low-rank approximations using the EFTT format. Our algorithm is based on first obtaining suitable factor matrices for a Tucker approximation (2.5) of the coefficient tensor. The columns of these matrices are determined by applying a combination of adaptive cross approximation [30] with randomized pivoting to matricizations of the coefficient tensor. We then apply oblique projections based on discrete empirical interpolation as in Section 3.2 to implicitly construct a suitable core tensor for the Tucker approximation. In a second step, we compute a TT approximation of this core tensor using Algorithm 3. The combination of the Tucker and TT approximation yields the desired functional low-rank approximation. The main advantage of our approach compared to a direct TT approximation of the coefficient tensor is that the approximated core tensor is potentially much smaller than the coefficient tensor. This reduces the storage complexity of the approximation and it significantly reduces the number of function evaluations needed in the computation of the TT approximation, which we demonstrate in our numerical experiments.

This chapter is based on the article [306]. Its remainder is structured as follows. In Section 4.1, we define the EFTT format. Our novel algorithm to compute approximations in the EFTT format is presented in Section 4.2. Our numerical experiments in Section 4.3 demonstrate the advantages of our novel algorithm compared to a direct TT approximation of the coefficient tensor as in [41] and to the FTT algorithm in [141]. The test functions for our numerical experiments are defined in the Appendix 4.A.

### 4.1 Extended functional tensor train format

A function $f:[-1,1]^{d} \rightarrow \mathbb{R}$ is said to be represented in EFTT format with TT representation ranks $\left(R_{1}, \ldots, R_{d-1}\right) \in \mathbb{N}^{d-1}$ and multi-linear representation ranks $\left(r_{1}, \ldots, r_{d}\right) \in \mathbb{N}^{d}$ when it is given in terms of univariate functions $u_{j}^{(\ell)}:[-1,1] \rightarrow \mathbb{R}$ for $j=1, \ldots, r_{\ell}, \ell=$ $1, \ldots, d$ and TT cores $\mathcal{H}^{(\ell)} \in \mathbb{R}^{R_{\ell-1} \times r_{\ell} \times \mathbb{R}_{\ell}}$ for $\ell=1, \ldots, d$ as

$$
\begin{equation*}
f\left(x_{1}, \ldots, x_{d}\right) \approx \sum_{j_{1}=1}^{r_{1}} \ldots \sum_{j_{d}=1}^{r_{d}} \sum_{\alpha_{1}=1}^{R_{1}} \ldots \sum_{\alpha_{d-1}}^{R_{d-1}} \mathcal{H}_{1, j_{1}, \alpha_{1}}^{(1)} \cdots \mathcal{H}_{\alpha_{d-1}, j_{d}, 1}^{(d)} u_{j_{1}}^{(1)}\left(x_{1}\right) \cdots u_{j_{d}}^{(d)}\left(x_{d}\right) \tag{4.1}
\end{equation*}
$$

where we formally set $R_{0}=R_{d}=1$ to simplify notation.
Following the construction in Section 2.3 .4 we can derive approximations in the EFTT format (4.1) from low-rank approximations of the evaluation tensor $\mathcal{T}$ (2.18). We first
compute a Tucker decomposition $(2.5)$ of $\mathcal{T} \in \mathbb{R}^{\left(n_{1}+1\right) \times \cdots \times\left(n_{d}+1\right)}$ of the form

$$
\begin{equation*}
\mathcal{T} \approx \mathcal{C} \times_{1} U^{(1)} \times_{2} \cdots \times_{d} U^{(d)} \tag{4.2}
\end{equation*}
$$

with $\mathcal{C} \in \mathbb{R}^{r_{1} \times \cdots \times r_{d}}$ and $U^{(\ell)} \in \mathbb{R}^{\left(n_{\ell}+1\right) \times r_{\ell}}$. In a second step, the core tensor $\mathcal{C}$ is approximated in TT format (2.6) by

$$
\begin{equation*}
\hat{\mathcal{C}}_{i_{1}, \ldots, i_{d}}=\sum_{\alpha_{1}=1}^{R_{1}} \cdots \sum_{\alpha_{d-1}=1}^{R_{d-1}} \mathcal{H}_{1, i_{1}, \alpha_{1}}^{(1)} \mathcal{H}_{\alpha_{1}, i_{1}, \alpha_{2}}^{(2)} \cdots \mathcal{H}_{\alpha_{d-2}, i_{d-1}, \alpha_{d-1}}^{(d-1)} \mathcal{H}_{\alpha_{d-1}, i_{d}, 1}^{(d)} \tag{4.3}
\end{equation*}
$$

for $i_{\ell}=1, \ldots, r_{\ell}, \ell=1, \ldots, d$, with the TT cores $\mathcal{H}^{(\ell)} \in \mathbb{R}^{R_{\ell-1} \times r_{\ell} \times R_{\ell}}$. Inserted into (4.2), this yields an approximation of the evaluation tensor in the extended TT format [286] of the form

$$
\begin{equation*}
\hat{\mathcal{T}}_{i_{1}, \ldots, i_{d}}=\sum_{j_{1}=1}^{r_{1}} \ldots \sum_{j_{d}=1}^{r_{d}} \sum_{\alpha_{1}=1}^{R_{1}} \ldots \sum_{\alpha_{d-1}}^{R_{d-1}} \mathcal{H}_{1, j_{1}, \alpha_{1}}^{(1)} \cdots \mathcal{H}_{\alpha_{d-1}, j_{d}, 1}^{(d)} U_{i_{1}, j_{1}}^{(1)} \cdots U_{i_{d}, j_{d}}^{(d)} \tag{4.4}
\end{equation*}
$$

for $i_{\ell}=0, \ldots, n_{\ell}, \ell=1, \ldots, d$. This only requires $\mathcal{O}\left(d r R^{2}+d n r\right)$ storage, where $R=\max R_{\ell}, r=\max r_{\ell}, n=\max n_{\ell}$, which compares favorably with the $\mathcal{O}\left(d n R^{2}\right)$ storage needed by a direct TT approximation (2.6) of $\mathcal{T}$, especially when $r \ll n$. In Figure 4.1, we visualize the corresponding approximation of the coefficient tensor (2.19) given by

$$
\begin{align*}
& \hat{\mathcal{A}}_{i_{1}, \ldots, i_{d}}=\sum_{k_{1}=0}^{n_{1}} \ldots \sum_{k_{d}=0}^{n_{d}} \sum_{j_{1}=1}^{r_{1}} \ldots \sum_{j_{d}=1}^{r_{d}} \sum_{\alpha_{1}=1}^{R_{1}} \ldots \sum_{\alpha_{d-1}}^{R_{d-1}} \mathcal{X}_{k_{1}, \ldots, k_{d}, j_{1}, \ldots, j_{d}, \alpha_{1}, \ldots, \alpha_{d-1}}  \tag{4.5}\\
& \mathcal{X}_{k_{1}, \ldots, k_{d}, j_{1}, \ldots, j_{d}, \alpha_{1}, \ldots, \alpha_{d-1}}=\mathcal{H}_{1, j_{1}, \alpha_{1}}^{(1)} \cdots \mathcal{H}_{\alpha_{d-1}, j_{d}, 1}^{(d)} U_{i_{1}, j_{1}}^{(1)} \cdots U_{i_{d}, j_{d}}^{(d)} F_{k_{1}, i_{1}}^{(1)} \cdots F_{k_{d}, i_{d},}^{(d)}
\end{align*}
$$

for $i_{\ell}=0, \ldots, n_{\ell}, \quad \ell=1, \ldots, d$. Inserting $\hat{\mathcal{A}}$ into (2.21) yields the desired approximation (4.1) with TT cores $H^{(\ell)}$ and with

$$
u_{j}^{(\ell)}(x):=\sum_{i=0}^{n_{\ell}} \sum_{k=0}^{n_{\ell}} F_{k, i}^{(\ell)} U_{i, j}^{(\ell)} T_{k}(x)
$$

The resulting approximation error is bounded by Lemma 2.2. When we store the tensors $\mathcal{H}^{(\ell)}$ and the matrices $\left(F^{(\ell)} U^{(\ell)}\right) \in \mathbb{R}^{\left(n_{\ell}+1\right) \times r_{\ell}}$ separately, we can compute the point evaluation of $(4.1)$ at $\left(x_{1}, \ldots, x_{d}\right) \in[-1,1]^{d}$ by first contracting $F^{(\ell)} U^{(\ell)}$ with the vectors $\left(T_{0}\left(x_{\ell}\right), \ldots, T_{n_{\ell}}\left(x_{\ell}\right)\right)$ for $\ell=1, \ldots, d$ and by then contracting the resulting vectors with the TT cores $\mathcal{H}^{(\ell)}$ [247]. This requires $\mathcal{O}\left(d r n+d r R^{2}\right)$ operations, where $R=\max R_{\ell}$, $r=\max r_{\ell}, n=\max n_{\ell}$.

Remark 4.1. Note that EFTT approximations obtained from evaluation tensor approximations of the form (4.4) are closely related to FTT approximations. Given an approximation (4.4), we can define the TT cores $\mathcal{G}^{(\ell)}=\mathcal{H}^{(\ell)} \times{ }_{2} U^{(\ell)}$. The TT approxima-


Figure 4.1 - Tensor network representation [247] of the coefficient tensor $\hat{\mathcal{A}}$ in (4.5) corresponding to EFTT format (4.1). The colored boxes mark the subtensors corresponding to the approximation of the evaluation tensor $\hat{\mathcal{T}}$ and the approximation of the core tensor $\hat{\mathcal{C}}$.
tion (2.6) given by the cores $\mathcal{G}^{(\ell)}$ can be used to derive an equivalent FTT approximation using the construction in Section 2.3.4. Conversely, EFTT approximations can be seen as compression of such a FTT approximation, since the TT cores are stored in terms of $\mathcal{H}^{(\ell)}$ and $U^{(\ell)}$ requiring only $\mathcal{O}\left(\left(n_{\ell}+1\right) r_{\ell}+R_{\ell-1} r_{\ell} R_{\ell}\right)$ storage instead of $\mathcal{O}\left(R_{\ell-1}\left(n_{\ell}+1\right) R_{\ell}\right)$ for the uncompressed core $\mathcal{G}^{(\ell)}$.

### 4.2 Approximation algorithm

In the following, we develop a novel algorithm for computing approximations in the EFTT format (4.1). Our algorithm obtains the factor matrices $U^{(1)}, \ldots, U^{(d)}$ for the Tucker approximation (4.2) from fibers of the evaluation tensor $\mathcal{T}$ via a variant of column subset selection [79, 93]. Following [100], the core tensor $\mathcal{C}$ is obtained as a subtensor of $\mathcal{T}$ by applying discrete empirical interpolation. Thus, there is no need to form $\mathcal{C}$ explicitly and we apply a variant of the TT-cross algorithm [251] to compute the TT cores $\mathcal{H}^{(1)}, \ldots, \mathcal{H}^{(d)}$ for the TT approximation (4.3) from only some entries of $\mathcal{C}$.

Fiber identification. Applying cross approximation (2.1) to $\mathcal{T}\{\ell\}$ yields index sets $\hat{I}_{\ell}, \hat{J}_{\ell}$ that determine an approximation of the form

$$
\begin{equation*}
\mathcal{T}^{\{\ell\}} \approx \mathcal{T}^{\{\ell\}}\left(:, \hat{J}_{\ell}\right)\left(\mathcal{T}^{\{\ell\}}\left(\hat{I}_{\ell}, \hat{J}_{\ell}\right)\right)^{-1} \mathcal{T}^{\{\ell\}}\left(\hat{I}_{\ell},:\right) . \tag{4.6}
\end{equation*}
$$

Note that the matrix $\mathcal{T}^{\{\ell\}}\left(:, \hat{J}_{\ell}\right)$ contains mode- $\ell$ fibers of $\mathcal{T}$. If the approximation error of (4.6) is small, we choose $\hat{U}^{(\ell)}=\mathcal{T}^{\{\ell\}}\left(:, \hat{\zeta}_{\ell}\right)$ as an approximate basis of mode- $\ell$ fibers. Due to the large size of $\mathcal{T}^{(\ell)}$, it might not be feasible to apply either ACA with full or ACA with partial pivoting as described in Algorithm 1 and Remark 2.1. Inspired by the success of randomization in numerical linear algebra [222, 339], we propose to determine the next indices in line 5 of Algorithm 1 based on sampling. This leads to Algorithm 8,
which samples a fixed number $s$ of random entries and picks the entry of largest absolute value. These random entries are also used for stopping the algorithm and to determine $r_{\ell}$ adaptively. Note that the update in line 10 is only performed implicitly and the subtraction is carried out each time an entry of $B$ is evaluated. Overall, Algorithm 8 applied to $\mathcal{T}^{\{\ell\}}$ determines the index sets $\hat{I}_{\ell}, \hat{J}_{\ell}$ using only $\mathcal{O}\left(r_{\ell}^{3}+s r_{\ell}^{2}\right)$ entries of $\mathcal{T}$. Afterwards, we explicitly compute $\hat{U}^{(\ell)}=\mathcal{T}^{\{\ell\}}\left(:, \hat{J}_{\ell}\right)$ by evaluating $n_{\ell} r_{\ell}$ entries of $\mathcal{T}$. We perform the described procedure for every $\ell=1, \ldots, d$ to determine $\hat{U}^{(1)}, \ldots, \hat{U}^{(d)}$; see also lines 4 and 5 of Algorithm 9 below.

```
Algorithm 8 Adaptive cross approximation with randomized pivoting
    Input: procedure to evaluate entries of \(A \in \mathbb{R}^{m \times n}\), tolerance \(\varepsilon\), number of samples \(s\)
    Output: index sets \(I\) and \(J\) such that \(A \approx A(:, J) A(I, J)^{-1} A(I,:)\)
    \(I=\emptyset, J=\emptyset\)
    \(B=A\)
    while true
        Construct \(S \subset\{1, \ldots, n\} \times\{1, \ldots, m\}\) by uniformly sampling \(s\) index pairs \((i, j)\).
        if \(\max _{(i, j) \in S}\left|B_{i, j}\right| \leq \varepsilon\) return \(I, J\) end
        \(\left(i^{*}, j^{*}\right)=\underset{(i, j) \in S}{\arg \max }\left|B_{i, j}\right|\)
        \(I=I \cup\left\{i^{*}\right\}, J=J \cup\left\{j^{*}\right\}\)
        \(B=A-A(:, J) A(I, J)^{-1} A(I,:)\)
```

Tucker approximation To arrive at the Tucker approximation (4.2) we need to project the fibers of $\mathcal{T}$ onto the spans of $\hat{U}^{(1)}, \ldots, \hat{U}^{(d)}$. Analogous to Section 3.2.3, we first compute economic QR decompositions $\hat{U}^{(\ell)}=Q^{(\ell)} R^{(\ell)}$ and apply DEIM (see Algorithm 2) to identify index sets $I_{\ell}$. We then apply oblique projections (2.3) to each mode of $\mathcal{T}$ to compute the Tucker approximation

$$
\begin{aligned}
\mathcal{T} & \approx\left(\mathcal{T} \times_{1} \Phi_{I_{1}}^{T} \times_{2} \cdots \times_{d} \Phi_{I_{d}}^{T}\right) \times_{1} Q^{(1)}\left(\Phi_{I_{1}}^{T} Q^{(1)}\right)^{-1} \times_{2} \cdots \times_{d} Q^{(d)}\left(\Phi_{I_{d}}^{T} Q^{(d)}\right)^{-1} \\
& =\mathcal{C} \times_{1} U_{1} \times_{2} \cdots \times_{d} U_{d},
\end{aligned}
$$

with factor matrices $U^{(\ell)}=Q^{(\ell)}\left(\Phi_{I_{\ell}}^{T} Q^{(\ell)}\right)^{-1}$ and core tensor $\mathcal{C}=\mathcal{T}\left(I_{1}, \ldots, I_{d}\right)$. Note that the subtensor $\mathcal{C}$ of $\mathcal{T}$ will not be explicitly formed; we only need to evaluate some its entries when computing the TT approximation of $\mathcal{C}$. The whole process of computing $U^{(\ell)}$ and $\mathcal{C}$ from $\hat{U}^{(\ell)}$ is summarized in Algorithm 9. Alternative randomized algorithms to compute Tucker approximations have been proposed in [228, 230].

Remark 4.2. For simplicity, we have assumed that the values of $n_{1}, \ldots, n_{d}$, determining the size of the expansion (2.16) and of $\mathcal{T}$ are given as input. In practice, these values are usually not provided. To ensure that the approximation error of the final approximation (4.1) is small, we use the heuristic introduced in Remark 2.11 to determine $n_{1}, \ldots, n_{d}$ adaptively. We initially set $n_{1}=\cdots=n_{d}=16$. If applying the chopping heuristic [16] to

```
Algorithm 9 Tucker approximation
    Input: procedure to evaluate entries of \(\mathcal{T} \in \mathbb{R}^{\left(n_{1}+1\right) \times \cdots \times\left(n_{d}+1\right)}\), tolerance \(\varepsilon\), number
    of samples \(s\)
    Output: Matrices \(U^{\ell} \in \mathbb{R}^{\left(n_{\ell}+1\right) \times r_{\ell}}\) and a procedure to evaluate entries of \(\mathcal{C} \in\)
    \(\mathbb{R}^{r_{1} \times \cdots \times r_{d}}\), defining a Tucker approximation (4.2)
    for \(\ell=1, \ldots, d\)
        \(\hat{J}_{\ell} \leftarrow\) index set \(J\) returned by Algorithm 8 applied to \(\mathcal{T}^{\{\ell\}}\) with tolerance \(\varepsilon\) and \(s\)
    samples.
        \(\hat{U}^{(\ell)}=\mathcal{T}^{\{\ell\}}\left(:, \hat{J}_{\ell}\right)\)
        Compute economic QR decomposition \(\hat{U}^{(\ell)}=Q^{(\ell)} R^{(\ell)}\).
        \(I_{\ell}=\operatorname{DEIM}\left(Q^{(\ell)}\right)(\) see Algorithm 2)
        \(U^{(\ell)}=Q^{(\ell)}\left(Q^{(\ell)}\left(I_{\ell},:\right)\right)^{-1}\)
    \(\mathcal{C}=\mathcal{T}\left(I_{1}, \ldots, I_{d}\right)\)
```

$\hat{U}^{(\ell)}$ indicates that $n_{\ell}$ is not sufficiently large then one sets $n_{\ell}=2 n_{\ell}+1$, updates $\mathcal{T}$ and repeats the procedure from line 4 to compute a new index set $\hat{\zeta}_{\ell}$.

TT cores It remains to compute a TT approximation of the core tensor (4.3) to obtain the desired extended TT approximation (4.4). For this purpose, we use the TT toolbox* implementation of Algorithm 3. This algorithm is rank-adaptive and requires the evaluation of only $\mathcal{O}\left(d r R^{2}\right)$ entries of $\mathcal{C}$ where $R=\max R_{\ell}, r=\max r_{\ell}$.

EFTT approximation algorithm We formalize the overall procedure for computing approximations in the EFTT format (4.1) in Algorithm 10. Note that $r_{1}, \ldots, r_{d}$ and $R_{1}, \ldots, R_{d-1}$ are determined adaptively in lines 5 and 6 . Let $R=\max R_{\ell}, r=\max r_{\ell}, n=$ $\max n_{\ell}$. The total number of evaluations of $f$ in Algorithm 10 is $\mathcal{O}\left(d r^{3}+d n r+d s r^{2}+d r R^{2}\right)$. In contrast, applying Algorithm 3 directly to $\mathcal{T}$ requires $\mathcal{O}\left(d n R^{2}\right)$ evaluations and yields an approximation in FTT format (2.6).

### 4.3 Numerical experiments

In this section, we present numerical experiments ${ }^{\dagger}$ to study the performance of Algorithm 10. Unless mentioned otherwise the sample size in Algorithm 1 is set to $s=\min \{\bar{n} / 2,50\}$, where $\bar{n}=\left(\left(n_{1}+1\right)\left(n_{2}+1\right) \cdots\left(n_{d}+1\right)\right)^{1 / d}$, and all tolerances are set to $10^{-10}$. The approximations error is measured via a Monte Carlo estimation of the relative $L^{2}$-error based on function evaluations at 10000 sample points.

[^2]```
Algorithm 10 EFTT approximation
    Input: function \(f:[-1,1]^{d} \rightarrow \mathbb{R}\), tolerance \(\varepsilon\), number of samples \(s\)
    Output: TT cores \(\mathcal{H}^{(\ell)}\) and procedures for evaluating univariate functions \(u_{j}^{(\ell)}\)
    defining an approximation of \(f\) in the EFTT format (4.1)
    \(n_{1}=\cdots=n_{d}=16\)
    Define a procedure to evaluate entries of the evaluation tensor \(\mathcal{T}\) defined in (2.18).
    Apply Algorithm 9 with tolerance \(\varepsilon\) and \(s\) samples to \(\mathcal{T}\) to determine \(\hat{U}^{(\ell)}\). Simulta-
    neously, we update \(n_{1}, \ldots, n_{d}\) as well as \(\mathcal{T}\) as described in Remark 4.2.
    6: Compute the tensors \(\mathcal{H}^{(\ell)}\) by applying Algorithm 3 with tolerance \(\varepsilon\) to \(\mathcal{C}\).
    : Create procedures to evaluate univariate functions \(u_{j}^{(\ell)}(x)=\)
    \(\sum_{i=0}^{n_{\ell}} \sum_{k=0}^{n_{\ell}} F_{k, i}^{(\ell)} U_{i, j}^{(\ell)} T_{k}(x)\).
```


### 4.3.1 Comparison to a direct TT approximation

We first compare Algorithm 10 yielding an approximation in the EFTT format (4.1) to a direct TT approximation of the evaluation tensor $\mathcal{T} \in \mathbb{R}^{100 \times \cdots \times 100}$ (2.18) using Algorithm 3 as in [41]. The latter approach yields an approximation in the FTT format (2.12). Since a direct TT approximation is not adaptive with respect to the polynomial degree, we fix the polynomial degree throughout this comparison, i.e., we replace line 3 by $n_{1}=\cdots=n_{d}=100$ and do not update adaptively in line 5 of Algorithm 10.

Benchmark functions. For the set of benchmark functions defined in Appendix 4.A.1, we on average reduce the number of function evaluations required by $30.6 \%$ and the required storage by $41.6 \%$ when using Algorithm 10 compared to the direct TT approximation of the coefficient tensor. For the Ackley function the reduction is $88.8 \%$ in terms of function evaluations and $93 \%$ in terms of storage. For the Borehole function, we have $r_{\ell}=R_{\ell}^{2}$, i.e. we can not compress the FTT approximation further. In this case, the direct TT approximation is slightly more efficient. At the same time, this demonstrates that all other test functions, taken from a wide range of applications, can be stored more efficiently using our approach. Details can be found in Table 4.1.

Genz functions In order to quantify the impact of the dimension on the approximation error, we study the Genz [136] functions defined in Appendix 4.A.2. Our numerical results for the approximation using a fixed polynomial degree, displayed in Figure 4.2, demonstrate that Algorithm 10 requires fewer function evaluation compared to a direct TT-cross approximation. However, in very large dimensions $(d>300)$ the direct TT-cross approach leads to slightly more accurate approximations. This might be caused by the fact that, for Algorithm 10, the error of $d$ projections needs to be taken into account in addition to the TT-cross approximation error. Moreover, approximating the smaller tensor $\mathcal{C}$ instead of $\mathcal{T}$ directly, might lead to slightly worse conditioned inverse matrices

Chapter 4: Extended functional tensor train approximation

| Function | Algorithm | Error | \# evals | \# dofs | $\max _{\ell} R_{\ell}$ | $\max _{\ell} r_{\ell}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Ackley | EFTT | $1.84 \mathrm{e}-02$ | 63168 | 15965 | 15 | 10 |
|  | DirectTT | $1.84 \mathrm{e}-02$ | 570176 | 226555 | 18 |  |
| Alpine | EFTT |  | 4677 | 1448 | 2 | 2 |
|  | DirectTT | $5.80 \mathrm{e}-03$ | 6860 | 2400 | 2 |  |
| Dixon | EFTT | $5.72 \mathrm{e}-14$ | 11879 | 3549 | 3 | 5 |
|  | DirectTT | $2.66 \mathrm{e}-14$ | 13022 | 5100 | 3 |  |
| Exponential | EFTT | $2.10 \mathrm{e}-14$ | 2108 | 707 | 1 | 1 |
|  | DirectTT | $2.09 \mathrm{e}-14$ | 2585 | 700 | 1 |  |
| Griewank | EFTT | $1.69 \mathrm{e}-07$ | 8091 | 2252 | 3 | 3 |
|  | DirectTT | $1.54 \mathrm{e}-07$ | 13022 | 5100 | 3 |  |
| Michalewicz | EFTT | $4.05 \mathrm{e}-02$ | 4677 | 1448 | 2 | 2 |
|  | DirectTT | $4.05 \mathrm{e}-02$ | 6860 | 2400 | 2 |  |
| Piston | EFTT | $3.22 \mathrm{e}-09$ | 202876 | 73978 | 24 | 11 |
|  | DirectTT | $3.02 \mathrm{e}-09$ | 991669 | 412186 | 18 |  |
| Qing |  | $1.11 \mathrm{e}-13$ | 5482 | 2172 | 2 | 3 |
|  | DirectTT | $2.29 \mathrm{e}-14$ | 6860 | 2400 | 2 |  |
| Rastrigin | EFTT | $2.29 \mathrm{e}-14$ | 4677 | 1448 | 2 | 2 |
|  | DirectTT | $2.30 \mathrm{e}-14$ | 6860 | 2400 | 2 |  |
| Rosenbrock | EFTT | $3.02 \mathrm{e}-14$ | 10970 | 2798 | 3 | 4 |
|  | DirectTT | $2.64 \mathrm{e}-14$ | 13022 | 5100 | 3 |  |
| Schaffer | EFTT | $6.76 \mathrm{e}-02$ | 1063304 | 288006 | 39 | 40 |
|  | DirectTT | $6.73 \mathrm{e}-02$ | 1510513 | 766565 | 30 |  |
| Schwefel | EFTT | $6.58 \mathrm{e}-04$ | 4677 | 1448 | 2 | 2 |
|  | DirectTT | $6.58 \mathrm{e}-04$ | 6860 | 2400 | 2 |  |
| Borehole | EFTT | $3.95 \mathrm{e}-02$ | 14365 | 3270 | 2 | 4 |
|  | DirectTT | $3.95 \mathrm{e}-02$ | 10178 | 2336 | 2 |  |
| OTL Circuit | EFTT | $3.85 \mathrm{e}-11$ | 16064 | 3277 | 5 | 5 |
|  | DirectTT | $8.31 \mathrm{e}-12$ | 27764 | 8300 | 4 |  |
| Robot Arm | EFTT | $6.65 \mathrm{e}-02$ | 523301 | 113365 | 35 | 36 |
|  | DirectTT | $6.88 \mathrm{e}-02$ | 752815 | 392234 | 32 |  |
| Wing Weight | EFTT | $3.72 \mathrm{e}-14$ | 6692 | 2072 | 2 | 2 |
|  | DirectTT | $8.11 \mathrm{e}-14$ | 10439 | 3600 | 2 |  |
| Friedman | EFTT | $5.49 \mathrm{e}-10$ | 12312 | 2376 | 4 | 4 |
|  | DirectTT | $1.04 \mathrm{e}-11$ | 14630 | 3136 | 3 |  |
| G \& L | EFTT | $2.52 \mathrm{e}-05$ | $3278$ | 1034 | 2 | 2 |
|  | DirectTT | $2.52 \mathrm{e}-05$ | $6651$ | 1800 | 2 |  |
| G \& P 8D | EFTT | $3.07 \mathrm{e}-11$ | 39588 | 8121 | 7 | 7 |
|  | DirectTT | $2.68 \mathrm{e}-11$ | 74693 | 30160 | 5 |  |
| D \& P Exp | EFTT | $1.56 \mathrm{e}-14$ | 1990 | 616 | 2 | 2 |
|  | DirectTT | $1.55 \mathrm{e}-14$ | 2087 | 800 | 2 |  |

Table 4.1 - We apply Algorithm 10 (EFTT) and an approximation of the evaluation tensor using Algorithm 3 (DirectTT) to approximate the test functions defined in Appendix 4.A.1. For each function, we display the estimated $L^{2}$ error of the approximations, the number of function evaluations required to construct the approximation, the degrees of freedom in the approximation and the largest $R_{\ell}$ and $r_{\ell}$ of the resulting approximation. The numbers displayed are the mean (geometric mean for the error and arithmetic mean for all other quantities) after approximating each function 100 times.
in (2.8).


Figure 4.2 - We apply Algorithm 10 (EFTT) and an approximation of the evaluation tensor using Algorithm 3 (DirectTT) to approximate the Genz functions (see Appendix 4.A.2) with fixed polynomial degree for dimension $d \in\{20,50,100,200,300,400,500\}$. Left: We plot the ratio of $L^{2}$ error for EFFT / DirectTT, i.e. for the function 'corner peak' with $d=20$ the EFTT approximation is around halve an order less accurate than the DirectTT approximation. Right: We plot the ratio of the function evaluations required to compute the approximation for EFFT / DirectTT, i.e. for the function 'corner peak' with $d=20$ EFTT we use only around $20 \%$ of the number of evaluations required for DirectTT. Throughout this figure, we sample 30 different parameters for the Genz functions as in [41] and display the geometric mean of the error and the arithmetic mean for the number of function evaluations.

### 4.3.2 Comparison to the FTT approximation algorithm

In this section, we compare the performance of our proposed Algorithm 10 based on the EFTT format (4.1) to approximation algorithm proposed in [139, 141] and implemented in the c3py package $\ddagger$. The c3py algorithm uses a continuous variant of the TT-cross algorithm to compute approximations in the FTT format (2.12).

The c3py package is based on the Legendre polynomials $P_{0}(x)=1, P_{1}(x)=x,(k+$ 2) $P_{k+2}(x)=(2 k+3) x P_{k+1}(x)-(k+1) P_{k}(x)$ for $k=0, \ldots, m-2[197]$ with degree at most $m \in \mathbb{N}$ instead of Chebyshev polynomials. This leads to multivariate polynomial approximations with degree at most $\left(m_{1}, \ldots, m_{d}\right)$ of the form

$$
\begin{equation*}
\sum_{i_{1}=0}^{m_{1}} \cdots \sum_{i_{d}=0}^{m_{d}} \mathcal{B}_{i_{1}, \ldots, i_{d}} P_{i_{1}}\left(x_{1}\right) \cdots P_{i_{d}}\left(x_{d}\right), \tag{4.7}
\end{equation*}
$$

where $\mathcal{B} \in \mathbb{R}^{\left(m_{1}+1\right) \times \cdots \times\left(m_{d}+1\right)}$. Since the asymptotic convergence rates of Legendre and Chebyshev interpolation are different [46], a fair comparison requires that we adjust our approximation algorithm to also return an approximation of the form (4.7). An approximation of the form (4.7) can be obtained from the evaluation tensor $\mathcal{T} \in$

[^3]$\mathbb{R}^{\left(n_{1}+1\right) \times \cdots \times\left(n_{d}+1\right)}(2.18)$ as $\mathcal{B}=\mathcal{T} \times{ }_{1} E^{(1)} \times{ }_{2} \cdots \times_{d} E^{(d)}$, where $E^{(\ell)} \in \mathbb{R}^{\left(n_{\ell}+1\right) \times\left(m_{\ell}+1\right)}$ is defined as
\[

E^{(\ell)}=\left($$
\begin{array}{cccc}
\frac{1}{2} w_{0}^{(\ell)} P_{0}\left(x_{0}^{(\ell)}\right) & \frac{1}{2} w_{1}^{(\ell)} P_{0}\left(x_{1}^{(\ell)}\right) & \ldots & \frac{1}{2} w_{m_{\ell}}^{(\ell)} P_{0}\left(x_{m_{\ell}}^{(\ell)}\right) \\
\frac{3}{2} w_{0}^{(\ell)} P_{1}\left(x_{0}^{(\ell)}\right) & \frac{3}{2} w_{1}^{(\ell)} P_{1}\left(x_{1}^{(\ell)}\right) & \ldots & \frac{3}{2} w_{m_{\ell}}^{(\ell)} P_{1}\left(x_{m_{\ell}}^{(\ell)}\right. \\
\frac{5}{2} w_{0}^{(\ell)} P_{2}\left(x_{0}^{(\ell)}\right) & \frac{5}{2} w_{1}^{(\ell)} P_{2}\left(x_{1}^{(\ell)}\right) & \ldots & \frac{5}{2} w_{m_{\ell}}^{(\ell)} P_{2}\left(x_{m_{\ell}}^{(\ell)}\right) \\
\vdots & \vdots & \ddots & \vdots \\
\frac{2 n_{\ell}+1}{2} w_{0}^{(\ell)} P_{n_{\ell}}\left(x_{0}^{(\ell)}\right) & \frac{2 n_{\ell}+1}{2} w_{1}^{(\ell)} P_{n_{\ell}}\left(x_{1}^{(\ell)}\right) & \ldots & \frac{2 n_{\ell}+1}{2} w_{m_{\ell}}^{(\ell)} P_{n_{\ell}}\left(x_{m_{\ell}}^{(\ell)}\right)
\end{array}
$$\right)
\]

with nonnegative weights $w_{i_{\ell}}^{(\ell)}$ defined as in Remark 2.12. The matrix $E^{(\ell)}$ encodes Clenshaw-Curtis quadrature [74] in $n_{\ell}+1$ nodes to (approximately) compute the Legendre coefficients via $L^{2}$-projections. Using analogous constructions as in Section 7.1, we can transform an extended TT approximation of the evaluation tensor (4.4) into an EFTT approximation with univariate functions represented in terms of linear combinations of Legendre polynomials.

Remark 4.3. Gauss-Legendre quadrature in $n+1$ points integrates a polynomials of degree $2 n+1$ exactly, whereas Clenshaw-Curtis quadrature is only exact for polynomials of degree $n$. We could define the evaluation tensor $\mathcal{T}$ and the matrices $E^{\ell}$ differently to encode Gauss-Legendre quadrature instead of Clenshaw-Curtis quadrature. However, in practice, both quadrature rules typically perform well as pointed out in [325].

In the following experiments we set $n_{\ell}=2 m_{\ell}$ for $\ell=1, \ldots, d$ to ensure accurate quadrature. To determine the polynomial degrees $m_{\ell}$ adaptively (see Remark 4.2), we follow the fiber adaptation strategy of c3py (see [139, Section 3.6.1]): We progressively increase the degrees until four sequential Legendre coefficients are smaller than the tolerance $10^{-10}$ or until the maximum of $m_{\ell}=105$ has been reached. In the c3py algorithm, we set all tolerances to $10^{-10}$.

Benchmark functions. We apply both the c3py algorithm from [141] and our novel Algorithm 10 to approximate the benchmark functions defined in Appendix 4.A.1. Our numerical results displayed in Table 4.2 demonstrate that our novel algorithm typically requires fewer function evaluations and still achieves the same accuracy compared to c3py. For the Ackley function, our approach reduces the number of required function evaluations by more than $96 \%$ and the storage by more than $90 \%$. At the same time, our approach is slightly more accurate.

Integration of sin function In the following, we repeat the experiment from [139, Figure 3-6]. The function $f\left(x_{1}, \ldots, x_{d}\right)=\sin \left(x_{1}+x_{2}+\cdots+x_{d}\right)$ can be represented in FTT format with TT ranks $(2,2, \ldots, 2)$ and its integral over the domain $[0,1]^{d}$ is known analytically [139]. In Figure 4.3, we plot the error of the integral of the approximations

| Function | Algorithm | Error | \# evals | \# dofs | $\max _{\ell} m_{\ell}$ | $\max _{\ell} R_{\ell}$ | $\max _{\ell} r_{\ell}$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Ackley | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 1.22 \mathrm{e}-02 \\ & 1.81 \mathrm{e}-02 \end{aligned}$ | $\begin{array}{r} 67107 \\ 2232528 \end{array}$ | $\begin{array}{r} 18465 \\ 197760 \end{array}$ | $\begin{aligned} & 105 \\ & 105 \end{aligned}$ | $\begin{aligned} & 15 \\ & 16 \end{aligned}$ | 9 |
| Alpine | $\underset{\text { c3py }}{\text { EFTT }}$ | $\begin{aligned} & 4.08 \mathrm{e}-03 \\ & 6.43 \mathrm{e}-03 \end{aligned}$ | $\begin{array}{r} 5464 \\ 50656 \end{array}$ | $\begin{aligned} & 1518 \\ & 2520 \end{aligned}$ | $\begin{aligned} & 105 \\ & 105 \end{aligned}$ | $\begin{aligned} & 2 \\ & 2 \end{aligned}$ | 2 |
| Dixon | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 3.21 \mathrm{e}-14 \\ & 2.89 \mathrm{e}-14 \end{aligned}$ | $\begin{aligned} & 13756 \\ & 13752 \end{aligned}$ | $\begin{array}{r} 3714 \\ 435 \end{array}$ | $\begin{gathered} 105 \\ 21 \end{gathered}$ | $\begin{aligned} & 3 \\ & 3 \end{aligned}$ | 5 |
| Exponential | $\underset{\text { c3py }}{\text { EFTT }}$ | $\begin{aligned} & 4.82 \mathrm{e}-15 \\ & 1.71 \mathrm{e}-14 \end{aligned}$ | $\begin{array}{r} 946 \\ 10518 \end{array}$ | $\begin{aligned} & 196 \\ & 133 \end{aligned}$ | $\begin{aligned} & 27 \\ & 19 \end{aligned}$ | $\begin{aligned} & 1 \\ & 1 \end{aligned}$ | 1 |
| Griewank | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 8.21 \mathrm{e}-08 \\ & 3.52 \mathrm{e}-06 \end{aligned}$ | $\begin{array}{r} 9139 \\ 51466 \end{array}$ | $\begin{aligned} & 2358 \\ & 4459 \end{aligned}$ | $\begin{aligned} & 105 \\ & 105 \end{aligned}$ | $\begin{aligned} & 3 \\ & 3 \\ & 3 \end{aligned}$ | 3 |
| Michalewicz | $\underset{\text { c3py }}{\text { EFTT }}$ | $\begin{aligned} & 2.54 \mathrm{e}-02 \\ & 1.45 \mathrm{e}-01 \end{aligned}$ | $\begin{array}{r} 5464 \\ 48745 \end{array}$ | $\begin{aligned} & 1518 \\ & 2443 \end{aligned}$ | $\begin{aligned} & 105 \\ & 105 \end{aligned}$ | $\begin{aligned} & 2 \\ & 2 \end{aligned}$ | 2 |
| Piston | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 3.71 \mathrm{e}-09 \\ & 3.85 \mathrm{e}-05 \end{aligned}$ | $\begin{aligned} & 174188 \\ & 251760 \end{aligned}$ | $\begin{aligned} & 69019 \\ & 66080 \end{aligned}$ | $\begin{aligned} & 33 \\ & 35 \end{aligned}$ | $\begin{aligned} & 23 \\ & 24 \end{aligned}$ | 11 |
| Qing | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 5.54 \mathrm{e}-13 \\ & 2.86 \mathrm{e}-14 \end{aligned}$ | $\begin{array}{r} 6996 \\ 11776 \end{array}$ | $\begin{array}{r} 2277 \\ 136 \end{array}$ | $\begin{array}{r} 105 \\ 21 \end{array}$ | $\begin{aligned} & 2 \\ & 2 \end{aligned}$ | 3 |
| Rastrigin | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 1.91 \mathrm{e}-14 \\ & 1.86 \mathrm{e}-10 \end{aligned}$ | $\begin{array}{r} 5463 \\ 22288 \end{array}$ | $\begin{aligned} & 1518 \\ & 1342 \end{aligned}$ | $\begin{array}{r} 105 \\ 63 \end{array}$ | $\begin{aligned} & 2 \\ & 2 \end{aligned}$ | 2 |
| Rosenbrock | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 1.25 \mathrm{e}-14 \\ & 9.43 \mathrm{e}-14 \end{aligned}$ | $\begin{array}{r} 4106 \\ 11530 \end{array}$ | $\begin{aligned} & 835 \\ & 633 \end{aligned}$ | $\begin{aligned} & 27 \\ & 14 \end{aligned}$ | $\begin{aligned} & 3 \\ & 3 \\ & 3 \end{aligned}$ | 4 |
| Schaffer | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 7.19 \mathrm{e}-02 \\ & 1.22 \mathrm{e}-01 \end{aligned}$ | $\begin{array}{r} 173787 \\ 3465760 \end{array}$ | $\begin{array}{r} 35016 \\ 214200 \end{array}$ | $\begin{aligned} & 105 \\ & 105 \end{aligned}$ | $\begin{aligned} & 16 \\ & 20 \end{aligned}$ | 17 |
| Schwefel | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 4.00 \mathrm{e}-04 \\ & 5.45 \mathrm{e}-04 \end{aligned}$ | $\begin{array}{r} 5463 \\ 50656 \end{array}$ | $\begin{aligned} & 1518 \\ & 2496 \end{aligned}$ | $\begin{aligned} & 105 \\ & 104 \end{aligned}$ | $\begin{aligned} & 2 \\ & 2 \end{aligned}$ | 2 |
| Borehole | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 3.95 \mathrm{e}-02 \\ & 2.08 \mathrm{e}-03 \end{aligned}$ | $\begin{array}{r} 6552 \\ 14346 \end{array}$ | $\begin{array}{r} 1116 \\ 577 \end{array}$ | $\begin{aligned} & 32 \\ & 70 \end{aligned}$ | $\begin{aligned} & 2 \\ & 2 \end{aligned}$ | 4 |
| OTL Circuit | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 7.93 \mathrm{e}-11 \\ & 4.07 \mathrm{e}-08 \end{aligned}$ | $\begin{array}{r} 6670 \\ 15674 \end{array}$ | $\begin{aligned} & 1083 \\ & 1782 \end{aligned}$ | $\begin{aligned} & 27 \\ & 28 \end{aligned}$ | $\begin{aligned} & 5 \\ & 5 \end{aligned}$ | 5 |
| Robot Arm | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 8.12 \mathrm{e}-02 \\ & 3.85 \mathrm{e}-01 \end{aligned}$ | $\begin{array}{r} 499954 \\ 2018017 \end{array}$ | $\begin{array}{r} 54760 \\ 228439 \end{array}$ | $\begin{array}{r} 94 \\ 105 \end{array}$ | $\begin{aligned} & 12 \\ & 20 \end{aligned}$ | 27 |
| Wing Weight | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 2.83 \mathrm{e}-14 \\ & 2.15 \mathrm{e}-13 \end{aligned}$ | $\begin{array}{r} 2867 \\ 12224 \end{array}$ | $\begin{aligned} & 560 \\ & 554 \end{aligned}$ | $\begin{aligned} & 24 \\ & 19 \end{aligned}$ | $\begin{aligned} & 2 \\ & 2 \end{aligned}$ | 2 |
| Friedman | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 2.16 e-11 \\ & 8.08 e-05 \end{aligned}$ | $\begin{array}{r} 5238 \\ 12142 \end{array}$ | $\begin{aligned} & 404 \\ & 710 \end{aligned}$ | $\begin{aligned} & 19 \\ & 15 \end{aligned}$ | $\begin{aligned} & 3 \\ & 4 \end{aligned}$ | 4 |
| G \& L | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 4.95 \mathrm{e}-06 \\ & 3.51 \mathrm{e}-02 \end{aligned}$ | $\begin{array}{r} 1547 \\ 13928 \end{array}$ | $\begin{aligned} & 356 \\ & 374 \end{aligned}$ | $\begin{array}{r} 29 \\ 105 \end{array}$ | $\begin{aligned} & 2 \\ & 2 \end{aligned}$ | 2 |
| G \& P 8D | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 4.77 \mathrm{e}-11 \\ & 9.54 \mathrm{e}-10 \end{aligned}$ | $\begin{aligned} & 19527 \\ & 27336 \end{aligned}$ | $\begin{aligned} & 3902 \\ & 5136 \end{aligned}$ | $\begin{aligned} & 24 \\ & 21 \end{aligned}$ | $\begin{aligned} & 6 \\ & 7 \end{aligned}$ | 7 |
| D \& P Exp | $\begin{aligned} & \text { EFTT } \\ & \text { c3py } \end{aligned}$ | $\begin{aligned} & 1.13 \mathrm{e}-14 \\ & 4.78 \mathrm{e}-10 \end{aligned}$ | $\begin{array}{r} 2404 \\ 12162 \end{array}$ | $\begin{aligned} & 646 \\ & 336 \end{aligned}$ | $\begin{array}{r} 105 \\ 49 \end{array}$ | $\begin{aligned} & 2 \\ & 2 \end{aligned}$ | 2 |

Table 4.2 - We apply Algorithm 10 (EFTT) and the algorithm in the c3py package to approximate the test functions defined in Appendix 4.A. 1 by functional low-rank approximation with Legendre polynomial basis functions (4.7). For each function, we display the estimated $L^{2}$-error of the approximations, the number of function evaluations required to construct the approximation, the degrees of freedom in the approximation and the largest $m_{\ell}, R_{\ell}$ and $r_{\ell}$ of the resulting approximation.
computed via Algorithm 10 and the c3py algorithm. The figure shows, that our novel approach achieves a similar level of accuracy. The difference in the number of function evaluations is rather small, since the multilinear ranks and TT ranks of the function are small.


Figure 4.3 - We apply Algorithm 10 (EFTT) and the algorithm in the c3py package to approximate $f\left(x_{1}, \ldots, x_{d}\right)=\sin \left(x_{1}+x_{2}+\cdots+x_{d}\right)$ by functional low-rank approximation with Legendre polynomial basis functions (4.7) as in [139, Figure 3-6]. Left: We plot the relative error of the integral of the approximations. Right: We plot the number of function evaluations required to compute the approximation.

Application: Uncertainty quantification. A classical application of multivariate function approximation is the computation of surrogates for uncertainty quantification [308, 341]. For the approximation of the quantity of interest mapping defined in Appendix 4.A.3, we find in Table 4.3 that Algorithm 10 leads to approximations requiring less storage compared to c3py.

|  |  | error | \# evals | \# dofs |
| :--- | :--- | :--- | ---: | ---: |
| $d=4$ | c3py | $4.21 \cdot 10^{-5}$ | 360 | 168 |
|  | EFTT | $1.05 \cdot 10^{-3}$ | 337 | 60 |
| $d=9$ | c3py | $9.79 \cdot 10^{-5}$ | 960 | 448 |
|  | EFTT | $1.30 \cdot 10^{-3}$ | 757 | 153 |
| $d=16$ | c3py | $1.73 \cdot 10^{-4}$ | 1800 | 840 |
|  | EFTT | $7.62 \cdot 10^{-4}$ | 1345 | 240 |

Table 4.3 - We apply Algorithm 10 (EFTT) and the algorithm in the c3py package to approximate the quantity of interest map $Q:[-1,1]^{d} \rightarrow \mathbb{R}$ defined in Appendix 4.A.3 for $d \in\{4,9,16\}$. All tolerances are set to $10^{-3}$. The table displays the $L^{2}$-error, the required number of evaluations (evals) and the number of degrees of freedom (dofs) for both approximations.

## 4.A Appendix

## 4.A. 1 Benchmark functions

In Table 4.4 and Table 4.5, we define the benchmark functions for our numerical experiments. Note that the functions are defined on different tensor product domains. In our experiments, we map the domain of these functions onto $[-1,1]^{d}$ using an affine linear transformation.

| Function | $d$ | Domain | References |
| :---: | :---: | :---: | :---: |
| $\begin{aligned} f_{\text {Ackley }}(\mathbf{x})= & -20 \exp \left(-0.2 \sqrt{\frac{1}{7} \sum_{i=1}^{7}\left(x_{i}\right)^{2}}\right) \\ & -\exp \left(\frac{1}{7} \sum_{i=1}^{7} \cos \left(2 \pi\left(x_{i}\right)\right)\right)+20+e^{1} \end{aligned}$ | 7 | $[-32.768,32.768]^{7}$ | [177, 310] |
| $f_{\text {Alpine }}(\mathbf{x})=\sum_{i=1}^{7}\left\|x_{i} \sin \left(x_{i}\right)+0.1 x_{1}\right\|$ | 7 | $[-10,10]^{7}$ | [177, 269] |
| $f_{\text {Dixon }}(\mathbf{x})=\left(x_{1}-1\right)^{2}+\sum_{i=2}^{7} i \cdot\left(2 x_{i}^{2}-x_{i-1}\right)^{2}$ | 7 | $[-10,10]^{7}$ | [177, 310] |
| $f_{\text {Exponential }}(\mathbf{x})=-\exp \left(-\frac{1}{2} \sum^{7} x_{i}^{2}\right)$ | 7 | $[-1,1]^{7}$ | [177, 268] |
| $f_{\text {Griewank }}(\mathbf{x})=\sum_{i=1}^{7} \frac{x_{i}^{2}}{4000}-\prod_{i=1}^{d} \cos \left(\frac{x_{i}}{\sqrt{i}}\right)+1$ | 7 | $[-600,600]^{7}$ | [177, 310] |
| $f_{\text {Michalewicz }}(\mathbf{x})=-\sum_{i=1}^{7} \sin \left(x_{i}\right) \sin ^{20}\left(\frac{i x_{i}^{2}}{\pi}\right)$ | 7 | $[0, \pi]^{7}$ | [310, 329] |
| $f_{\text {Piston }}(\mathbf{x})$ | 7 | see text | [310, 344] |
| $f_{\text {Qing }}(\mathbf{x})=\sum_{i=1}^{7}\left(x_{i}^{2}-i\right)^{2}$ | 7 | $[0,500]^{7}$ | [177, 265] |
| $f_{\text {Rastrigin }}(\mathbf{x})=70+\sum_{i=1}^{7}\left(x_{i}^{2}-10 \cos \left(2 \pi \cdot x_{i}\right)\right)$ | 7 | $[-5.12,5.12]^{7}$ | [98, 310] |
| $f_{\text {Rosenbrock }}(\mathbf{x})=\sum_{i=1}^{6}\left(100 \cdot\left(x_{i+1}-x_{i}^{2}\right)^{2}+\left(1-x_{i}\right)^{2}\right)$ | 7 | $[-2.048,2.048]^{7}$ | [177, 310] |
| $f_{\text {Schaffer }}(\mathbf{x})=\sum_{i=1}^{6}\left(0.5+\frac{\sin ^{2}\left(\sqrt{x_{i}^{2}+x_{i+1}^{2}}\right)-0.5}{\left(1+0.001\left(x_{i}^{2}+x_{i+1}^{2}\right)\right)^{2}}\right)$ | 7 | $[-100,100]^{7}$ | [177, 310] |
| $f_{\text {Schwefel }}(\mathbf{x})=2932.8803-\sum_{i=1}^{7} x_{i} \cdot \sin \left(\sqrt{\left\|x_{i}\right\|}\right)$ | 7 | $[-500,500]^{7}$ | [98, 310] |

Table 4.4 - Test functions from [71, Table 1].

Chapter 4: Extended functional tensor train approximation

| Function | $d$ | Domain | References |
| :--- | :---: | :---: | :---: |
| $f_{\text {Borehole }}(\mathbf{x})$ | 8 | see text | $[12,310]$ |
| $f_{\text {OTL Circuit }}(\mathbf{x})$ | 6 | see text | $[233,310]$ |
| $f_{\text {Robot Arm }}(\mathbf{x})$ | 8 | see text | $[12,310]$ |
| $f_{\text {Wing Weight }}(\mathbf{x})$ | 10 | see text | $[123,310]$ |
| $f_{\text {Friedman }}(\mathbf{x})=10 \sin \left(\pi x_{1} x_{2}\right)+20\left(x_{3}-0.5\right)^{2}+10 x_{4}+5 x_{5}$ | 5 | $[0,1]^{5}$ | $[128,310]$ |
| $f_{\text {G\&L }}(\mathbf{x})=\exp \left[\sin \left(\left(0.9\left(x_{1}+0.48\right)\right)^{10}\right)\right]+x_{2} x_{3}+x_{4}$ | 6 | $[0,1]^{6}$ | $[144,310]$ |
| $f_{\text {D\&P 8D }}(\mathbf{x})=4\left(x_{1}-2+8 x_{2}-8 x_{2}^{2}\right)^{2}+\left(3-4 x_{2}\right)^{2}$ |  |  |  |
| $\quad+16 \sqrt{x_{3}+1}\left(2 x_{3}-1\right)^{2}+\sum_{i=4}^{8} i \ln \left(1+\sum_{j=3}^{i} x_{j}\right)$ | 8 | $[0,1]^{8}$ | $[94,310]$ |
|  |  |  |  |
| $f_{\text {D\&P Exp }}(\mathbf{x})=100\left(e^{-2 / x_{1}^{1.75}}+e^{-2 / x_{2}^{1.5}}+e^{-2 / x_{3}^{1.25}}\right)$ |  |  | $[0,1]^{3}$ |

Table 4.5 - Test functions from [139, Table 3.2].

Some of these functions do not fit into the format of the table. These are defined in the following:

$$
f_{\text {Piston }}\left(M, S, V_{0}, k, P_{0}, T_{a}, T_{0}\right)=2 \pi \sqrt{\frac{M}{k+S^{2} \frac{P_{0} V_{0}}{T_{0}} \frac{T_{a}}{V^{2}}}},
$$

where

$$
V=\frac{S}{2 k}\left(\sqrt{A^{2}+4 k \frac{P_{0} V_{0}}{T_{0}} T_{a}-A}\right) \text { and } A=P_{0} S+19.62 M-\frac{k V_{0}}{S},
$$

with $M \in[30,60], S \in[0.005,0.02], V_{0} \in[0.002,0.01], k=[1000,5000], P_{0} \in$ $[90000,110000], T_{a} \in[290,296], T_{0} \in[340,360]$.

$$
f_{\text {Borehole }}\left(r_{w}, r, T_{u}, H_{u}, T_{l}, H_{l}, L, K_{w}\right)=\frac{2 \pi T_{u}\left(H_{u}-H_{l}\right)}{\log \left(r / r_{w}\right)\left(1+\frac{2 L T_{u}}{\log \left(r / r_{w}\right) r_{w}^{2} K_{w}}+\frac{T_{u}}{T_{l}}\right)}
$$

with $r_{w} \in[0.05,0.15], r \in[100,50000], T_{u} \in[63070,115600], H_{u} \in[990,1110], T_{l} \in$ [63.1, 116], $H_{l} \in[700,820], L \in[1120,1680], K_{w} \in[9855,12045]$.

$$
\begin{aligned}
& f_{\text {OTL Circuit }}\left(b_{1}, b_{2}, f, c_{1}, c_{2}, \beta\right)= \\
& \quad \frac{\left(\frac{12 b_{2}}{b_{1}+b_{2}}+0.74\right) \beta\left(c_{2}+9\right)}{\beta\left(c_{2}+9\right)+f}+\frac{11.35 f}{\beta\left(c_{2}+9\right)+f}+\frac{0.74 f \beta\left(c_{2}+9\right)}{\left(\beta\left(c_{2}+9\right)+f\right) c_{1}},
\end{aligned}
$$

with $b_{1} \in[50,150], b_{2} \in[25,70], f \in[0.5,3], c_{1} \in[1.2,2.5], c_{2} \in[0.25,1.2], \beta \in[50,300]$.

$$
f_{\text {Robot } \operatorname{Arm}}\left(\theta_{1}, \theta_{2}, \theta_{3}, \theta_{4}, L_{1}, L_{2}, L_{3}, L_{4}\right)=\sqrt{u^{2}+v^{2}},
$$

where $u=\sum_{i=1}^{4} L_{i} \cos \left(\sum_{j=1}^{4} \theta_{i}\right), v=\sum_{i=1}^{4} L_{i} \sin \left(\sum_{j=1}^{4} \theta_{i}\right)$ and $\theta_{i} \in[0,2 \pi], L_{i} \in[0,1]$ for $i=1, \ldots, 4$.

$$
\begin{aligned}
& f_{\text {Wing Weight }}\left(S_{w}, W_{f}, A, \Delta, q, \lambda, t_{c}, N_{z}, W_{d}, W_{p}\right)= \\
& \quad 0.036 S_{w}^{0.758} W_{f}^{0.0035}\left(\frac{A}{\cos ^{2}(\Delta)}\right)^{0.6} q^{0.006} \lambda^{0.04}\left(\frac{100 t_{c}}{\cos (\Delta)}\right)^{-0.3}\left(N_{z} W_{d}\right)^{0.49}+S_{w} W_{p},
\end{aligned}
$$

with $S_{w} \in[150,200], W_{f} \in[220,300], A \in[6,10], \Delta \in[-10,10], q \in[16,45], \lambda \in$ $[0.5,1], t_{c} \in[0.08,0.18], N_{z} \in[2.5,6], W_{d} \in[1700,2500], W_{p} \in[0.025,0.08]$.

## 4.A. 2 Genz functions

In the following, we define the Genz functions [136], which are frequently used to evaluate function approximation and integration schemes. On the domain $[-1,1]^{d}$, we consider

$$
\begin{array}{ll}
f_{1}(\mathbf{x})=\cos \left(2 \pi w_{1}+\sum_{i=1}^{d} c_{i} \frac{x_{i}+1}{2}\right) & \text { (oscillatory) } \\
f_{2}(\mathbf{x})=\left(1+\sum_{i=1}^{d} c_{i} \frac{x_{i}+1}{2}\right)^{-(d+1)} & \text { (corner peak) } \\
f_{3}(\mathbf{x})=\exp \left(-\sum_{i=1}^{d} c_{i}^{2}\left|\frac{x_{i}+1}{2}-w_{i}\right|\right) & \text { (continuous) }
\end{array}
$$

The parameters $w_{i}$ and $c_{i}$ are drawn uniformly from $[0,1]$, where $w_{i}$ act as a shift for the functions while $c_{i}$ determines the approximation difficulty of the functions. We normalized $c_{i}$ such that

$$
\sum_{i=1}^{d}\left|c_{i}\right|=\frac{b}{d^{h}},
$$

where the scaling constants $h$ and $b$ are defined for each function as in [41, Table 1]. Note that $f_{1}$ can be represented in FTT format (2.12) with $\max R_{\ell}=2$. The function $f_{3}$ is separable. For $f_{2}$, we are not aware of any analytic FTT representation.

## 4.A. 3 Parametric PDE problem

In the following, we recall the example from [207, Section 4]. Assume $\sqrt{d} \in \mathbb{N}$. Let $\Omega=[0,1]^{2}$. We consider the parametric elliptic PDE

$$
\begin{equation*}
-\nabla \cdot(a(x, p) \nabla u(x, p))=1, \quad x \in \Omega \tag{4.8}
\end{equation*}
$$

with homogeneous Dirichlet boundary conditions and parameter $p \in[-1,1]^{d}$. We define the piecewise constant coefficient $a(x, p): \Omega \times \mathbb{R}^{d}$ as

$$
a(x, p)= \begin{cases}1.25+0.75 p_{\sqrt{d}(t-1)+s} & x \in \Omega_{s, t} \\ 1 & \text { otherwise }\end{cases}
$$

where we denote the disk with radius $\rho=1 /(4 \sqrt{d}+2)$ centered around $(\rho(4 s-1), \rho(4 t-1))$ by $\Omega_{s, t}$ for $s, t=1, \ldots, \sqrt{d}$. In our numerical experiments, we approximate the quantity of interest $Q:[-1,1]^{d} \rightarrow \mathbb{R}$ defined as

$$
Q(p)=\int_{0}^{1} \int_{0}^{1} u(x, p) \partial x_{1} \partial x_{2}
$$

where $u(x, p)$ denotes the solution of the $\operatorname{PDE}(4.8)$ for the given parameter $p \in[-1,1]^{d}$. For each value of $p$, we solve the resulting PDE using a discretization based on linear finite elements [48].

## 5 Multi-marginal optimal transport

In this chapter, we study the solution of multi-marginal optimal transport problems, i.e. we search the nonnegative transport plan tensor $\mathcal{P}$ of order $d$, which minimizes $\langle\mathcal{C}, \mathcal{P}\rangle$ subject to marginal constraints, where $\mathcal{C}$ denotes a given nonnegative cost tensor of order $d$. After entropic regularization [82] the problem is equivalent to a tensor scaling problem of $\exp (-\mathcal{C} / \eta)$ [37] and can be solved using the Sinkhorn algorithm [58]. The multi-marginal Sinkhorn algorithm crucially relies on the repeated evaluation of marginals of the rescaled tensor $\exp (-\mathcal{C} / \eta)$.

Using a novel approach, we combine the ideas of exploiting underlying graphical models [150] and using low-rank approximations [7, 10] to compute marginals more efficiently. When the structure of transport plans is specified by a graphical model, we observe that the dual of this model is a tensor network [273], which contains a tensor network representation of $\exp (-\mathcal{C} / \eta)$. At the same time low-rank approximations of $\exp (-\mathcal{C} / \eta)$ can also be represented as tensor networks [247]. Facilitating this point of view, marginals can in both cases be computed by contracting [272] the tensor network and the scaling parameters. Note that the belief propagation and the junction tree algorithm for graphical models correspond to a particular order of contracting this network [273]. For tensor networks derived from graphical models, we propose to potentially accelerate the computation of marginals further by replacing tensors in the network by low-rank approximations. This yields a modified tensor network and can be seen as an approximation of $\exp (-\mathcal{C} / \eta)$. In our numerical experiments, we provide an example illustrating that our approach to introduce low-rank approximations in the tensor network is more efficient than directly working with the graphical model and more accurate than a direct tensor train approximation [249] of $\exp (-\mathcal{C} / \eta)$. We also demonstrate that our approach offers the potential to greatly speed up the computation of color transfer between several images without altering the resulting image significantly.

Moreover, we provide theoretical bounds on the error caused by introducing such approximations. In Theorem 5.2, we state a bound for the impact of using an approximation of
$\exp (-\mathcal{C} / \eta)$ on the entropically regularized transport cost. This result is a generalization of the bound in [10] for classical optimal transport problems. In contrast to the asymptotic bound in [7], our result contains explicit constants and does not assume that $\mathcal{C}$ is low-rank. We also would like to point out that the bound in [7] relies on the fact the rank of $\exp (-\mathcal{C} / \eta)$ can be bounded based on the rank of $\mathcal{C}$. The practical usefulness of these results is impeded by the fact that the elementwise exponential tends to increase (approximate) ranks drastically. For example, to obtain a reasonably good approximation of the elementwise exponential of a random $1000 \times 1000$ rank- 5 matrix by truncating all singular values smaller than $10^{-10}$ one easily ends up with a matrix of rank 800 or larger. In Theorem 5.3, we state how the parameters and tolerances need to be selected to obtain an accurate approximation of the original problem without regularization. This generalizes previous results in [217] from the case of using the tensor $\exp (-\mathcal{C} / \eta)$ directly to our case of using an approximation instead. In Lemma 5.2, we link approximations of parts in the tensor network to approximations of $\exp (-\mathcal{C} / \eta)$.

This chapter is based on the article [305]. Its remainder is structured as follows. In Section 5.1, we define the multi-marginal optimal transport problem and summarize the convergence results for the multi-marginal Sinkhorn algorithm in [127, 217]. The impact of approximating $\exp (-\mathcal{C} / \eta)$ is analyzed in Section 5.2. In Section 5.3, we derive the tensor network structure of $\exp (-\mathcal{C} / \eta)$ for transport plans represented by graphical models. Approximations of parts of this network are connected to approximations of $\exp (-\mathcal{C} / \eta)$ in Section 5.4. Our numerical experiments in Section 5.5 demonstrate how the multimarginal Sinkhorn algorithm can be accelerated by introducing low-rank approximations into the tensor network representation of $\exp (-\mathcal{C} / \eta)$.

### 5.1 Multi-marginal optimal transport and the Sinkhorn algorithm

Throughout this chapter, we let $\Delta^{n}=\left\{x \in \mathbb{R}_{>0}^{n}:\|x\|_{1}=1\right\}$ denote the set of strictly positive probability vectors of length $n$, where $\mathbb{R}_{>0}$ denotes the strictly positive real numbers. Further, we let $\mathbb{R}_{+}$denote the nonnegative real numbers. For a tensor $\mathcal{X} \in \mathbb{R}_{+}^{n_{1} \times \cdots \times n_{d}}$ containing the joint probability distribution of $d$ discrete random variables, we denote the $\ell$ th marginal (distribution) in $\Delta^{n_{\ell}}$ by

$$
r_{\ell}(\mathcal{X})=\operatorname{vec}\left(\mathcal{X} \times_{1} \mathbf{1}_{n_{1}}^{T} \cdots \times_{\ell-1} \mathbf{1}_{n_{\ell-1}}^{T} \times_{\ell+1} \mathbf{1}_{n_{\ell+1}}^{T} \cdots \times_{d} \mathbf{1}_{n_{d}}^{T}\right),
$$

where $\mathbf{1}_{n} \in \mathbb{R}^{n}$ denotes the vector of all ones. Throughout this chapter, the operations $\log$ and $\exp$ are always applied elementwise.

### 5.1.1 Mathematical setting

Given $d \geq 2$ marginals $r_{\ell} \in \Delta^{n_{\ell}}, \ell=1, \ldots, d$ and a cost tensor $\mathcal{C} \in \mathbb{R}_{+}^{n_{1} \times n_{2} \times \cdots \times n_{d}}$, the discrete multi-marginal optimal transport problem [238] is given by

$$
\begin{equation*}
\min _{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)}\langle\mathcal{P}, \mathcal{C}\rangle=\min _{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)} \sum_{i_{1}=1}^{n_{1}} \ldots \sum_{i_{d}=1}^{n_{d}} \mathcal{P}_{i_{1}, \ldots, i_{d}} \mathcal{C}_{i_{1}, \ldots, i_{d}}, \tag{5.1}
\end{equation*}
$$

where the set of feasible transport plans is given by

$$
B\left(r_{1}, \ldots, r_{d}\right)=\left\{\mathcal{P} \in \mathbb{R}_{+}^{n_{1} \times \cdots \times n_{d}} \mid r_{\ell}(\mathcal{P})=r_{\ell} \text { for } \ell=1, \ldots, d\right\} .
$$

Note that (5.1) is a linear optimization problem with $n_{1} \cdot n_{2} \cdots n_{d}$ degrees of freedom.
To solve (5.1) efficiently, it is common to add entropic regularization [82]. In the multimarginal setting, the entropy takes the form

$$
H(\mathcal{P})=-\langle\mathcal{P}, \log (\mathcal{P})\rangle .
$$

Given a regularization parameter $\eta>0$, the regularized problem takes the form

$$
\begin{equation*}
\min _{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)} V_{\mathcal{C}}^{\eta}(\mathcal{P}), \tag{5.2}
\end{equation*}
$$

where

$$
V_{\mathcal{C}}^{\eta}(\mathcal{P}):=\langle\mathcal{P}, \mathcal{C}\rangle-\eta H(\mathcal{P})
$$

is called entropic transport cost. It is known that the regularized problem (5.2) has a unique minimizer [37], which takes the form

$$
\begin{equation*}
\mathcal{P}_{\eta}^{*}=\mathcal{K} \times_{1} \operatorname{diag}\left(\exp \left(\beta_{1}\right)\right) \cdots \times_{d} \operatorname{diag}\left(\exp \left(\beta_{d}\right)\right), \tag{5.3}
\end{equation*}
$$

where $\mathcal{K}:=\exp (-\mathcal{C} / \eta)$ is called Gibbs kernel [58] and $\operatorname{diag}\left(\beta_{\ell}\right)$ denotes the diagonal matrix containing the entries of the so-called scaling parameters $\beta_{\ell} \in \mathbb{R}^{n_{\ell}}$ on its diagonal. The solution of the regularized problem (5.2) converges to a solution of (5.1) when $\eta \rightarrow 0$; see [35, 211].

### 5.1.2 Multi-marginal Sinkhorn algorithm

The multi-marginal Sinkhorn algorithm for solving (5.2) proceeds by iteratively updating the scaling parameters in (5.3). Let

$$
\begin{equation*}
\mathcal{P}_{\eta}^{(t)}=\mathcal{K} \times_{1} \operatorname{diag}\left(\exp \left(\beta_{1}^{(t)}\right)\right) \cdots \times_{d} \operatorname{diag}\left(\exp \left(\beta_{d}^{(t)}\right)\right) \tag{5.4}
\end{equation*}
$$

denote the scaled tensor obtained after $t$ iterations. In each iteration one selects an index $\ell$ and updates the $\ell$ th scaling parameter via $\beta_{\ell}^{(t+1)}=\log \left(r_{\ell}\right)-\log \left(r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right)\right)+\beta_{\ell}^{(t)}$, while the other vectors remain unchanged. For the choice of indices it has been suggested to traverse them cyclically [37] or use a greedy heuristics [127, 217]. The multi-marginal Sinkhorn algorithm is terminated once the stopping criterion

$$
\begin{equation*}
\sum_{\ell=1}^{d}\left\|r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right)-r_{\ell}\right\|_{1} \leq \varepsilon_{\text {stop }} \tag{5.5}
\end{equation*}
$$

is satisfied, where $\varepsilon_{\text {stop }}>0$ denotes a prescribed tolerance. Algorithm 11 summarizes the described procedure.

```
Algorithm 11 Multi-marginal Sinkhorn algorithm
    Input: cost tensor \(\mathcal{C}\), marginals \(r_{1}, \ldots, r_{d}\), regularization parameter \(\eta\)
    Output: transport plan \(\mathcal{P}_{\eta}^{(t)}\)
    \(\beta_{\ell}^{(0)}=\mathbf{0} \in \mathbb{R}^{n_{i}}\) for \(\ell=1, \ldots, d\) and \(\mathcal{K}=\exp (-\mathcal{C} / \eta)\)
    for \(t=0,1, \ldots\) until stopping criterion (5.5) is satisfied
        \(\mathcal{P}_{\eta}^{(t)}=\mathcal{K} \times_{1} \operatorname{diag}\left(\exp \left(\beta_{1}^{(t)}\right)\right) \cdots \times_{d} \operatorname{diag}\left(\exp \left(\beta_{d}^{(t)}\right)\right)\)
        Let \(\ell_{\text {next }}\) denote the index of the scaling parameter that should be updated next.
        \(\beta_{\ell}^{(t+1)}= \begin{cases}\log \left(r_{\ell}\right)-\log \left(r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right)\right)+\beta_{\ell}^{(t)} & \ell=\ell_{\text {next }} \\ \beta_{\ell}^{(t)} & \ell \neq \ell_{\text {next }}\end{cases}\)
```

When using cyclic order, it follows from an interpretation as Bregman projections [37] that the multi-marginal Sinkhorn algorithm converges. For greedy strategies, Theorem 5.1 below summarizes the statements of [217, Theorem 4.3] and [127, Theorem 3.4], which bound the number of iterations until the stopping criterion is reached. Note that the bound in a) gives a better rate with respect to $\varepsilon_{\text {stop }}$, but the bound depends on $r_{\ell}$, whereas the bound in b) does not involve $r_{\ell}$.
Theorem 5.1. Let $\mathcal{C} \in \mathbb{R}_{+}^{n_{1} \times \cdots \times n_{d}}, r_{\ell} \in \Delta^{n_{\ell}}$ for $\ell=1, \ldots, d, 0<\eta<\frac{1}{2}$, and $\varepsilon_{\text {stop }}>0$.
a) Suppose that Algorithm 11 selects the index of the scaling parameter in iteration $t$ according to

$$
\ell_{\text {next }}=\underset{\ell \in\{1, \ldots, d\}}{\arg \max } \mathbf{1}^{T}\left(r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right)-r_{\ell}\right)+\mathbf{1}^{T}\left(r_{\ell} * \log \left(r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right) * r_{\ell}^{-1}\right)\right),
$$

where * denotes the elementwise product and $r_{\ell}^{-1}$ denotes the elementwise inverse. Then the number of iterations to reach the stopping criterion (5.5) is bounded by

$$
\begin{equation*}
t \leq 2+2 d^{2} \varepsilon_{\text {stop }}^{-1}\left(\eta^{-1}\|\mathcal{C}\|_{\infty}-\log \min _{1 \leq \ell \leq d} \min _{1 \leq i \leq n_{\ell}}\left(r_{\ell}\right)_{i}\right) \tag{5.6}
\end{equation*}
$$

b) Assume that $n=n_{1}=\cdots=n_{d}$ and suppose that Algorithm 11 normalizes $\mathcal{P}_{\eta}^{(0)}$ to
have $\ell^{1}$-norm 1 and selects the index of the scaling parameter in iteration $t$ according to

$$
\begin{equation*}
\ell_{\text {next }}=\underset{\ell \in\{1, \ldots, d\}}{\arg \max }\left\|r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right)-\frac{\left\langle r_{\ell}, r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right)\right\rangle}{\left\|r_{\ell}\right\|_{2}^{2}} r_{\ell}\right\|_{1} . \tag{5.7}
\end{equation*}
$$

Then the number of iterations needed to reach the stopping criterion

$$
\begin{equation*}
\max _{\ell \in\{1, \ldots, d\}}\left\|r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right)-\frac{\left\langle r_{\ell}, r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right)\right\rangle}{\left\|r_{\ell}\right\|_{2}^{2}} r_{\ell}\right\|_{1}<\frac{\varepsilon_{\text {stop }}}{2 d} \tag{5.8}
\end{equation*}
$$

is bounded by

$$
t \leq 8 d^{2}(\sqrt{n}+1)^{2} \varepsilon_{\text {stop }}^{-2} \log \left(\eta^{-1}\|\exp (-\mathcal{C} / \eta)\|_{1}\right)
$$

When the alternative stopping criterion (5.8) is satisfied, the stopping criterion (5.5) is also satisfied.

A transport plan $\hat{\mathcal{P}} \in B\left(r_{1}, \ldots, r_{d}\right)$ is called $\varepsilon$-approximate solution of the original problem (5.1) if it satisfies

$$
\langle C, \hat{\mathcal{P}}\rangle \leq \min _{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)}\langle\mathcal{C}, \mathcal{P}\rangle+\varepsilon
$$

The transport plan $\mathcal{P}_{\eta}^{(t)}$ obtained from Algorithm 11 using either stopping criterion from Theorem 5.1 is, in general, not in $B\left(r_{1}, \ldots, r_{d}\right)$ because the marginal constraints $r_{\ell}=r_{\ell}\left(\mathcal{P}_{\varepsilon}^{(t)}\right)$ are satisfied simultaneously only in the limit $t \rightarrow \infty$. To fix this issue, rounding [9] can be applied to enforce the marginal constraints on $P_{\varepsilon}^{(t)}$ for finite $t$; see Algorithm 12. Note that for a tensor of the form (5.4), the operation in line 5 of Algorithm 12 can be phrased in terms of modifying $\beta_{\ell}^{(t)}$. Line 6 is a rank- 1 update. In [127, Lemma 3.6] and in [217, Theorem 4.4], the following property of the resulting tensor is proven.

Lemma 5.1. Let $\mathcal{A} \in \mathbb{R}_{>0}^{n_{1} \times \cdots \times n_{d}}$ and $r_{\ell} \in \Delta^{n_{\ell}}, \ell=1, \ldots, d$. Let $\mathcal{B}$ denote the output of Algorithm 12 applied to $\mathcal{A}$ and $r_{1}, \ldots, r_{d}$. Then $\mathcal{B} \in B\left(r_{1}, \ldots, r_{d}\right)$ and

$$
\|\mathcal{A}-\mathcal{B}\|_{1} \leq 2 \sum_{\ell=1}^{d}\left\|r_{\ell}-r_{\ell}(\mathcal{A})\right\|_{1}
$$

```
Algorithm 12 Rounding
    Input: tensor \(\mathcal{A} \in \mathbb{R}_{>0}^{n_{1} \times \cdots \times n_{d}}\), vectors \(r_{\ell} \in \Delta^{n_{\ell}}\) for \(\ell=1, \ldots, d\)
    Output: tensor \(\mathcal{B} \in \mathbb{R}^{n_{1} \times \cdots \times n_{d}}\)
    for \(\ell=1, \ldots, d\)
        \(v=\min \left(r_{\ell}(\mathcal{A})^{-1} * r_{\ell}, \mathbf{1}_{n_{\ell}}\right)\), where the min is taken elementwise
        \(\mathcal{A}=\mathcal{A} \times \ell \operatorname{diag}(v)\)
    \(\mathcal{B}=\mathcal{A}+\left\|r_{1}-r_{1}(\mathcal{A})\right\|_{1}^{-(d-1)} \bigcirc_{\ell=1}^{d}\left(r_{\ell}-r_{\ell}(\mathcal{A})\right)\), where \(\bigcirc\) denotes the outer product
```

By combining Algorithm 12 with Algorithm 11, using the index selection (5.7) and the stopping criterion (5.5), it follows [127, Corollary 3.8] that an $\varepsilon$-approximate solution can be computed in

$$
\mathcal{O}\left(\varepsilon^{-3} d^{4} n^{d+1} \log (n)(\max (\mathcal{C})-\min (\mathcal{C}))^{3}\right)
$$

operations, where $n_{1}=\cdots=n_{d}=n$. In practice, the computation can be accelerated by using slightly modified marginals and tolerances in the Sinkhorn algorithm [217], which ensure that the upper bound (5.6) for $t$ is not dominated by small entries in $r_{\ell}$.

### 5.2 Impact of approximating the Gibbs kernel

To accelerate the computation of marginals in the multi-marginal Sinkhorn algorithm, we will replace the Gibbs kernel $\mathcal{K}$ by an approximation $\tilde{\mathcal{K}}$; thus replacing Algorithm 11 by Algorithm 13. In this section, we will analyze the impact of this approximation on the transport cost of the computed transport plan. For $d=2$, such an analysis can be found in [10, Theorem 5]. In the following theorem, we generalize this result to the multi-marginal setting. Our proof closely follows the ideas in [10]. In contrast to the asymptotic result in [7, Theorem 7.4], we provide explicit bounds.

```
Algorithm 13 Multi-marginal Sinkhorn algorithm for a Gibbs kernel approximation
    Input: approximation \(\tilde{\mathcal{K}} \in \mathbb{R}_{>0}^{n_{1} \times n_{2} \cdots \times n_{d}}\) of Gibbs kernel \(\mathcal{K}=\exp (-\mathcal{C} / \eta)\), marginals
    \(r_{1}, \ldots, r_{d}\)
    Output: transport plan \(\tilde{\mathcal{P}}^{(t)}\)
    \(\beta_{\ell}^{(0)}=\mathbf{0} \in \mathbb{R}^{n_{\ell}}\) for \(\ell=1, \ldots, d\)
    for \(t=0,1, \ldots\) until stopping criterion (5.5) is satisfied
        \(\tilde{\mathcal{P}}^{(t)}=\tilde{\mathcal{K}} \times_{1} \operatorname{diag}\left(\exp \left(\beta_{1}^{(t)}\right)\right) \cdots \times_{d} \operatorname{diag}\left(\exp \left(\beta_{d}^{(t)}\right)\right)\)
        Let \(\ell_{\text {next }}\) denote the index of the scaling parameter that should be updated next.
        \(\beta_{\ell}^{(t+1)}= \begin{cases}\log \left(r_{\ell}\right)-\log \left(r_{\ell}\left(\tilde{\mathcal{P}}^{(t)}\right)\right)+\beta_{\ell}^{(t)} & \ell=\ell_{\text {next }} \\ \beta_{\ell}^{(t)} & \ell \neq \ell_{\text {next }}\end{cases}\)
```

Theorem 5.2. Let $\mathcal{K}=\exp (-\mathcal{C} / \eta)$ and assume that $\tilde{\mathcal{K}} \in \mathbb{R}_{>0}^{n_{1} \times n_{2} \cdots \times n_{d}}$ with $n_{\ell} \geq 2$ and
$d \geq 2$, satisfies $\|\log (\mathcal{K})-\log (\tilde{\mathcal{K}})\|_{\infty} \leq \varepsilon_{\log } \leq 1$. Let $\tilde{\mathcal{P}}$ denote the transport plan returned by Algorithm 13 with stopping criterion $\sum_{\ell=1}^{d}\left\|r_{\ell}(\tilde{\mathcal{P}})-r_{\ell}\right\|_{1} \leq \varepsilon_{\text {stop }}$. Then

$$
\left|V_{\mathcal{C}}^{\eta}\left(\mathcal{P}_{\eta}^{*}\right)-V_{\mathcal{C}}^{\eta}(\tilde{\mathcal{P}})\right| \leq \varepsilon_{V_{\mathcal{C}}^{\eta}}
$$

where $\mathcal{P}_{\eta}^{*}=\arg \min _{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)} V_{\mathcal{C}}^{\eta}(\mathcal{P})$ and

$$
\begin{align*}
\varepsilon_{V_{\mathcal{C}}^{\eta}}= & \eta\left(\varepsilon_{\log }\left(2+\log \left(\frac{2}{\varepsilon_{\log }}\right)\right)+\frac{\varepsilon_{\log }}{2} \log \left(\left(\prod_{\ell=1}^{d} n_{\ell}\right)-1\right)\right. \\
& \left.+2 \varepsilon_{\text {stop }} \log \left(\frac{1}{\varepsilon_{\text {stop }}}\left(\left(\prod_{\ell=1}^{d} n_{\ell}\right)-1\right)\right)\right)+\left(\varepsilon_{\text {log }}+2 \varepsilon_{\text {stop }}\right)\|\mathcal{C}\|_{\infty} \tag{5.9}
\end{align*}
$$

Proof. We denote by $\Pi^{S}$ the operator mapping a given tensor $\mathcal{T} \in \mathbb{R}_{>0}^{n_{1} \times \cdots \times n_{d}}$ to its unique [126] tensor scaling $\mathcal{U} \in B\left(r_{1}, \ldots, r_{d}\right)$ of the form $\mathcal{U}=\mathcal{T} \times_{1} \operatorname{diag}\left(\gamma_{1}\right) \times_{2} \cdots \times_{d}$ $\operatorname{diag}\left(\gamma_{d}\right)$ for some vectors $\gamma_{\ell} \in \mathbb{R}_{>0}^{n_{\ell}}$ for $\ell=1, \ldots, d$. Observe that $\mathcal{P}_{\eta}^{*}=\Pi^{S}(\mathcal{K})$. Using the triangle inequality, we decompose the error into

$$
\begin{align*}
\left|V_{\mathcal{C}}^{\eta}\left(\mathcal{P}_{\eta}^{*}\right)-V_{\mathcal{C}}^{\eta}(\tilde{\mathcal{P}})\right| \leq & \left|V_{\mathcal{C}}^{\eta}\left(\Pi^{S}(\mathcal{K})\right)-V_{\mathcal{C}}^{\eta}\left(\Pi^{S}(\tilde{\mathcal{K}})\right)\right|  \tag{5.10}\\
& +\left|V_{\mathcal{C}}^{\eta}\left(\Pi^{S}(\tilde{\mathcal{K}})\right)-V_{\tilde{\mathcal{C}}}^{\eta}\left(\Pi^{S}(\tilde{\mathcal{K}})\right)\right|  \tag{5.11}\\
& +\left|V_{\tilde{\mathcal{C}}}^{\eta}\left(\Pi^{S}(\tilde{\mathcal{K}})\right)-V_{\tilde{\mathcal{C}}}^{\eta}(\tilde{\mathcal{P}})\right|  \tag{5.12}\\
& +\left|V_{\tilde{\mathcal{C}}}^{\eta}(\tilde{\mathcal{P}})-V_{\mathcal{C}}^{\eta}(\tilde{\mathcal{P}})\right| \tag{5.13}
\end{align*}
$$

where $\tilde{\mathcal{C}}=-\eta \log (\tilde{\mathcal{K}})$. We derive bounds for each of these terms; their combination yields inequality (5.9).

Bound for (5.10): By definition of $V_{\mathcal{C}}^{\eta}$, we have

$$
\begin{aligned}
& \left|V_{\mathcal{C}}^{\eta}\left(\Pi^{S}(\mathcal{K})\right)-V_{\mathcal{C}}^{\eta}\left(\Pi^{S}(\tilde{\mathcal{K}})\right)\right| \leq \\
& \quad\left\|\Pi^{S}(\mathcal{K})-\Pi^{S}(\tilde{\mathcal{K}})\right\|_{1}\|\mathcal{C}\|_{\infty}+\eta\left|H\left(\Pi^{S}(\mathcal{K})\right)-H\left(\Pi^{S}(\tilde{\mathcal{K}})\right)\right| .
\end{aligned}
$$

Because of $\underset{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)}{\arg \min } V_{\mathcal{C}}^{\eta}(\mathcal{P})=\underset{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)}{\arg \min }\langle-\log (\mathcal{K}), \mathcal{P}\rangle-H(\mathcal{P})$, it follows that

$$
\begin{aligned}
& \left\|\Pi^{S}(\mathcal{K})-\Pi^{S}(\tilde{\mathcal{K}})\right\|_{1}= \\
& \quad\left\|\underset{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)}{\arg \min }\langle-\log (\mathcal{K}), \mathcal{P}\rangle-H(\mathcal{P})-\left(\underset{\tilde{\mathcal{P}} \in B\left(r_{1}, \ldots, r_{d}\right)}{\arg \min }\langle-\log (\tilde{\mathcal{K}}), \tilde{\mathcal{P}}\rangle-H(\tilde{\mathcal{P}})\right)\right\|_{1}
\end{aligned}
$$

Applying Lemma I in [10] to the right hand side of this expression yields $\| \Pi^{S}(\mathcal{K})-$ $\Pi^{S}(\tilde{\mathcal{K}})\left\|_{1} \leq\right\| \log \mathcal{K}-\log \tilde{\mathcal{K}} \|_{\infty} \leq \varepsilon_{\log }$. From Theorem 6 in [169] and Lemma D in [10] we
obtain

$$
\left|H\left(\Pi^{S}(\mathcal{K})\right)-H\left(\Pi^{S}(\tilde{\mathcal{K}})\right)\right| \leq \varepsilon_{\log } \log \left(\frac{2}{\varepsilon_{\log }}\right)+\frac{\varepsilon_{\log }}{2} \log \left(\left(\prod_{\ell=1}^{d} n_{\ell}\right)-1\right)
$$

Bound for (5.11) and (5.13): Using that $\left\|\Pi^{S}(\tilde{\mathcal{K}})\right\|_{1}=\|\tilde{\mathcal{P}}\|_{1}=1$ we obtain

$$
\begin{aligned}
\left|V_{\mathcal{C}}^{\eta}\left(\Pi^{S}(\tilde{\mathcal{K}})\right)-V_{\tilde{\mathcal{C}}}^{\eta}\left(\Pi^{S}(\tilde{\mathcal{K}})\right)\right| & \leq\left\langle\mathcal{C}, \Pi^{S}(\tilde{\mathcal{K}})\right\rangle-\left\langle\tilde{\mathcal{C}}, \Pi^{S}(\tilde{\mathcal{K}})\right\rangle \leq\|\mathcal{C}-\tilde{\mathcal{C}}\|_{\infty} \leq \eta \varepsilon_{\log } \\
\left|V_{\tilde{\mathcal{C}}}^{\eta}(\mathcal{\mathcal { P }})-V_{\mathcal{C}}^{\eta}(\tilde{\mathcal{P}})\right| & \leq\langle\mathcal{C}, \tilde{\mathcal{P}}\rangle-\langle\tilde{\mathcal{C}}, \tilde{\mathcal{P}}\rangle \leq\|\mathcal{C}-\tilde{\mathcal{C}}\|_{\infty} \leq \eta \varepsilon_{\log }
\end{aligned}
$$

Bound for (5.12): Using that the tensor $\tilde{\mathcal{P}}$ is the unique minimizer of $\underset{B\left(r_{1}(\tilde{\mathcal{P}}), \ldots, r_{d}(\tilde{\mathcal{P}})\right)}{\arg \min } V_{\tilde{\mathcal{C}}}^{\eta}(\mathcal{P})$, Lemma $H$ in [10] yields

$$
\left|V_{\tilde{\mathcal{C}}}^{\eta}\left(\Pi^{S}(\tilde{\mathcal{K}})\right)-V_{\tilde{\mathcal{C}}}^{\eta}(\tilde{\mathcal{P}})\right| \leq \omega\left(d_{H}\left(B\left(r_{1}(\tilde{\mathcal{P}}), \ldots, r_{d}(\tilde{\mathcal{P}})\right), B\left(r_{1}, \ldots, r_{d}\right)\right)\right)
$$

where $d_{H}(\cdot, \cdot)$ denotes the Hausdorff distance and

$$
\omega(x)=x\|\mathcal{C}\|_{\infty}+\eta\left(x \log \left(\frac{2}{x}\left(\left(\prod_{\ell=1}^{d} n_{\ell}\right)-1\right)\right)\right)
$$

We can bound $d_{H}\left(B\left(r_{1}(\tilde{\mathcal{P}}), \ldots, r_{d}(\tilde{\mathcal{P}})\right), B\left(r_{1}, \ldots, r_{d}\right)\right)$ by $2 \varepsilon_{\text {stop }}$, since Algorithm 12 maps any $\mathcal{A} \in B\left(r_{1}(\tilde{\mathcal{P}}), \ldots, r_{d}(\tilde{\mathcal{P}})\right)$ to $\mathcal{B} \in B\left(r_{1}, \ldots, r_{d}\right)$ with $\|\mathcal{A}-\mathcal{B}\|_{1} \leq 2 \varepsilon_{\text {stop }}$ as stated in Lemma 5.1. This implies

$$
\left|V_{\tilde{\mathcal{C}}}^{\eta}\left(\Pi^{S}(\tilde{\mathcal{K}})\right)-V_{\tilde{\mathcal{C}}}^{\eta}(\tilde{\mathcal{P}})\right| \leq 2 \varepsilon_{\text {stop }}\|\mathcal{C}\|_{\infty}+2 \eta \varepsilon_{\text {stop }} \log \left(\frac{1}{\varepsilon_{\text {stop }}}\left(\left(\prod_{\ell=1}^{d} n_{\ell}\right)-1\right)\right)
$$

The following theorem demonstrates how the previous result can be combined with Algorithm 12 to obtain $\varepsilon$-accurate solutions. The proof is inspired by [217, Theorem 4.5] and [127, Theorem 3.7], which state how $\varepsilon$-accurate solutions can be computed using Algorithm 11. Theorem 5.3 takes into account that we compute the transport plan using Algorithm 13 based on a perturbed Gibbs kernel.

Theorem 5.3. Let $\hat{\mathcal{P}}$ be the tensor obtained by applying Algorithm 12 to $\tilde{\mathcal{P}}$, where $\tilde{\mathcal{P}}$ is obtained from Algorithm 13 with $\tilde{\mathcal{K}}$ fulfilling the assumptions of Theorem 5.2. Let $\mathcal{P}^{*}=\arg \min _{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)}\langle\mathcal{C}, \mathcal{P}\rangle$. Then it holds
$\langle\mathcal{C}, \hat{\mathcal{P}}\rangle-\left\langle\mathcal{C}, \mathcal{P}^{*}\right\rangle \leq \varepsilon$,
where $\varepsilon=2 \eta \varepsilon_{\log }+2 \eta \sum_{\ell=1}^{d} \log \left(n_{\ell}\right)+4\|\mathcal{C}\|_{\infty} \varepsilon_{\text {stop }}$.

Proof. From $\mathcal{P}^{*} \in B\left(r_{1}, \ldots, r_{d}\right)$ and $\mathcal{P}_{\eta}^{*}=\arg \min _{\mathcal{P} \in B\left(r_{1}, \ldots, r_{d}\right)}\langle\mathcal{C}, \mathcal{P}\rangle-\eta H(\mathcal{P})$ follows $\left\langle\mathcal{C}, \mathcal{P}_{\eta}^{*}\right\rangle-\eta H\left(\mathcal{P}_{\eta}^{*}\right) \leq\left\langle\mathcal{C}, \mathcal{P}^{*}\right\rangle-\eta H\left(\mathcal{P}^{*}\right)$. Thus,

$$
\begin{equation*}
\left\langle\mathcal{C}, \mathcal{P}_{\eta}^{*}\right\rangle-\left\langle\mathcal{C}, \mathcal{P}^{*}\right\rangle \leq \eta H\left(\mathcal{P}_{\eta}^{*}\right)-\eta H\left(\mathcal{P}^{*}\right) \leq \eta \sum_{\ell=1}^{d} \log \left(n_{\ell}\right), \tag{5.14}
\end{equation*}
$$

where we use that $0 \leq H(\mathcal{X}) \leq \sum_{\ell=1}^{d} \log \left(n_{\ell}\right)$ for any tensor $\mathcal{X} \in \mathbb{R}_{+}^{n_{1} \cdots n_{d}}$ with $\|\mathcal{X}\|_{1}=$ 1 [81, Theorem 2.6.4].

Note that the marginals of $\tilde{\mathcal{P}}$ are, in general, not equal to $r_{1}, \ldots, r_{d}$. In order to compare $\tilde{\mathcal{P}}$ and $\mathcal{P}_{\eta}^{*}$, we construct $\mathcal{Q} \in B\left(r_{1}(\tilde{\mathcal{P}}), \ldots, r_{d}(\tilde{\mathcal{P}})\right)$ by applying Algorithm 12 to $\mathcal{P}_{\eta}^{*}$ with marginals $r_{1}(\tilde{\mathcal{P}}), \ldots, r_{d}(\tilde{\mathcal{P}})$. Since $\mathcal{P}_{\eta}^{*} \in B\left(r_{1}, \ldots, r_{d}\right)$, Lemma 5.1 implies that $\left\|\mathcal{Q}-\mathcal{P}_{\eta}^{*}\right\|_{1} \leq 2 \varepsilon_{\text {stop }}$. Hence,

$$
\begin{equation*}
\langle\mathcal{C}, Q\rangle-\left\langle\mathcal{C}, \mathcal{P}_{\eta}^{*}\right\rangle \leq\|\mathcal{C}\|_{\infty}\left\|\mathcal{Q}-\mathcal{P}_{\eta}^{*}\right\|_{1} \leq 2\|\mathcal{C}\|_{\infty} \varepsilon_{\text {stop }} \tag{5.15}
\end{equation*}
$$

The tensor $\tilde{\mathcal{P}}$ is the unique scaling of $\tilde{\mathcal{K}}$ with marginals $r_{1}(\tilde{\mathcal{P}}), \ldots, r_{d}(\tilde{\mathcal{P}})$ [26]. It is thus the unique minimizer [127] of $\arg \min _{\mathcal{P} \in B\left(r_{1}(\tilde{\mathcal{P}}), \ldots, r_{d}(\tilde{\mathcal{P}})\right)}\langle\tilde{\mathcal{C}}, \mathcal{P}\rangle-\eta H(\mathcal{P})$, where $\tilde{\mathcal{C}}=-\eta \log (\tilde{\mathcal{K}})$. Following the same argument as in (5.14), we obtain

$$
\begin{equation*}
\langle\tilde{\mathcal{C}}, \tilde{\mathcal{P}}\rangle-\langle\tilde{\mathcal{C}}, \mathcal{Q}\rangle \leq \eta H(\tilde{\mathcal{P}})-\eta H(\mathcal{Q}) \leq \eta \sum_{\ell=1}^{d} \log \left(n_{\ell}\right) . \tag{5.16}
\end{equation*}
$$

We further obtain

$$
\begin{align*}
& \langle\tilde{\mathcal{C}}, \mathcal{Q}\rangle-\langle\mathcal{C}, \mathcal{Q}\rangle \leq\|\tilde{\mathcal{C}}-\mathcal{C}\|_{\infty}\|\mathcal{Q}\|_{1}=\|\tilde{\mathcal{C}}-\mathcal{C}\|_{\infty} \leq \eta \varepsilon_{\log },  \tag{5.17}\\
& \langle\mathcal{C}, \tilde{\mathcal{P}}\rangle-\langle\tilde{\mathcal{C}}, \tilde{\mathcal{P}}\rangle \leq\|\tilde{\mathcal{C}}-\mathcal{C}\|_{\infty}\|\tilde{\mathcal{P}}\|_{1}=\|\tilde{\mathcal{C}}-\mathcal{C}\|_{\infty} \leq \eta \varepsilon_{\log } \tag{5.18}
\end{align*}
$$

where we use that $\|\tilde{\mathcal{P}}\|_{1}=\|\mathcal{Q}\|_{1}=1$. Additionally, Lemma 5.1 yields

$$
\begin{equation*}
\langle\mathcal{C}, \hat{\mathcal{P}}\rangle-\langle\mathcal{C}, \tilde{\mathcal{P}}\rangle \leq\|\mathcal{C}\|_{\infty}\|\hat{\mathcal{P}}-\tilde{\mathcal{P}}\|_{1} \leq 2\|\mathcal{C}\|_{\infty} \varepsilon_{\text {stop }} \tag{5.19}
\end{equation*}
$$

By adding the inequalities (5.14)-(5.19) we obtain

$$
\langle\mathcal{C}, \hat{\mathcal{P}}\rangle-\left\langle\mathcal{C}, \mathcal{P}^{*}\right\rangle \leq 2 \eta \varepsilon_{\log }+2 \eta \sum_{\ell=1}^{d} \log \left(n_{\ell}\right)+4\|\mathcal{C}\|_{\infty} \varepsilon_{\text {stop }}
$$

Remark 5.1. Note that the $\eta \varepsilon_{\log }$ only converges to zero in the limit $\eta \rightarrow 0$ when the approximation $\tilde{\mathcal{K}}$ is equal to the true Gibbs kernel $\mathcal{K}$.

Remark 5.2. Note that for any $\varepsilon>0$, we can find suitable $\eta, \varepsilon_{\text {log }}, \varepsilon_{\text {stop }}$ such that combining Algorithm 13 and Algorithm 12 yields an $\varepsilon$-accurate solution for the multimarginal optimal transport problem (5.1). Analogously, we can find $\eta, \varepsilon_{l o g}, \varepsilon_{\text {stop }}$ such that (5.9) is satisfied for any given $\varepsilon_{V_{c}^{\eta}}>0$. In particular, we can first select $\varepsilon_{\log }$ and $\varepsilon_{\text {stop }}$ based on $\|\mathcal{C}\|_{\infty}$. Afterwards, we can set $\eta$ sufficiently small.

Remark 5.3. It might seem counter-intuitive that $\varepsilon_{\log }$ and $\varepsilon_{\text {stop }}$ needs to be chosen inversely proportional to $\|\mathcal{C}\|_{\infty}$ in the previous remark. This is caused by the chosen objective function in (5.2). Let $c=\|C\|_{\infty}^{-1}$. Note that optimal solution of the regularized (5.2) and the set of optimal solutions of the original optimal transport problem (5.1) do not change when we replace both $\mathcal{C}$ by $\mathfrak{C C}$ and $\eta$ by $\mathrm{c} \mathrm{\eta}$. This normalization changes the optimal value of the objective functions in (5.2) and (5.1) by c, but it does not change the Gibbs kernel $\mathcal{K}$. Thus, we can transform the error bounds in Theorem 5.2 and 5.3 into bounds for this normalized problem, by multiplying $\varepsilon_{V_{c}^{\eta}}$ and $\varepsilon$ by $c$. To obtain a certain $c \varepsilon_{V_{c}^{\eta}}$ respectively $c \varepsilon$, we can select $\varepsilon_{\log }$ and $\varepsilon_{\text {stop }}$ independently from $\|\mathcal{C}\|_{\infty}$ before determining a suitable $\eta$.

### 5.3 Tensor networks and graphical models

In applications, the cost tensor $\mathcal{C} \in \mathbb{R}_{+}^{n_{1} \times \cdots \times n_{d}}$ usually carries additional structure. A broad class of structures leads to transport plans defined via graphical models [150]. In this case, the entries of $\mathcal{C}$ take the form

$$
\begin{equation*}
\mathcal{C}_{I}=\sum_{\alpha \in F} \mathcal{C}_{I_{\alpha}}^{\alpha} \quad \text { for every } I=\left(i_{1}, \ldots, i_{d}\right), \quad i_{\ell}=1, \ldots, n_{\ell}, \ell=1, \ldots, d, \tag{5.20}
\end{equation*}
$$

where the summation index tuples $\alpha=\left(\alpha_{1}, \ldots, \alpha_{M}\right)$ are contained in a fixed subset $F$ of

$$
\bigcup_{M=1}^{d}\left\{\left(\alpha_{1}, \ldots, \alpha_{M}\right) \in \mathbb{N}^{M} \mid 1 \leq \alpha_{1}<\cdots<\alpha_{M} \leq d\right\}
$$

$I_{\alpha}:=\left(i_{\alpha_{1}}, \ldots, i_{\alpha_{M}}\right)$ and $\mathcal{C}^{\alpha} \in \mathbb{R}_{+}^{n_{\alpha_{1}} \times \cdots \times n_{\alpha_{M}}}$. The corresponding Gibbs kernel is given by

$$
\begin{equation*}
\mathcal{K}_{I}=\prod_{\alpha \in F} \mathcal{K}_{I_{\alpha}}^{\alpha} \quad \text { for every } I=\left(i_{1}, \ldots, i_{d}\right), \quad i_{\ell}=1, \ldots, n_{\ell}, \ell=1, \ldots, d, \tag{5.21}
\end{equation*}
$$

where $\mathcal{K}^{\alpha}=\exp \left(-\mathcal{C}^{\alpha} / \eta\right)$. This Gibbs kernel can be represented in terms of a tensor network [247]. In general, a tensor network represents a high-order tensor that is constructed by contracting several low-order tensors. By contraction we refer to the sum over a joint index in two low-order tensors. A graph is used to describe precisely how the low-order tensors should to be contracted. Its vertices correspond to the low-order tensors. Each edge corresponds to a contraction of the two low-order tensors corresponding to the vertices connected by the edge.

In the following, we describe how to construct the particular network for $\mathcal{K}$. First, we add each tensor $\mathcal{K}^{\alpha}$ as a vertex. For every $\ell=1, \ldots, d$, we add an additional tensor $\mathcal{D}^{(\ell)}$ as vertex. For every $\ell=1, \ldots, d$ and $\alpha \in F$, we add edges from $\mathcal{K}^{\alpha}$ to $\mathcal{D}^{(\ell)}$ if $\ell$ is contained in $\alpha$. We then add one open edge to each $\mathcal{D}^{(\ell)}$ that corresponds to the index in the $\ell$ th mode of $\mathcal{K}$. The order $d_{\ell}$ of the tensors $\mathcal{D}^{(\ell)} \in \mathbb{R}^{n_{\ell} \times \cdots \times n_{\ell}}$ is equal to the number of connected edges. Their entries are given by $\mathcal{D}_{i_{1}, \ldots, i_{d_{\ell}}}^{(\ell)}=\delta_{i_{1}, i_{2}} \delta_{i_{2}, i_{3}} \cdots \delta_{i_{d_{\ell}-1}, i_{d_{\ell}}}$, where $\delta$ denotes the Kronecker delta. Each edge in the tensor network corresponds to the sum over the corresponding index in the connected vertices. See Figure 5.1 for an example.

(a) Direct construction

(b) After simplification

Figure 5.1 - Tensor network representation of $\mathcal{K}$ for a cost tensor of the form (5.20) constructed from $\mathcal{C}^{(1,2)}, \mathcal{C}^{(1,4)} \mathcal{C}^{(2,3,4)}, \mathcal{C}^{(4,5)}$. In (b), we slightly simplified the network by contracting the identity matrices $\mathcal{D}^{(3)}, \mathcal{D}^{(5)}$ with their connected open edges. Summing over all internal edges yields the elementwise representation

$$
\sum_{j_{1}=1}^{\mathcal{K}_{i_{1}, i_{2}, i_{3}, i_{4}, i_{5}}^{n_{1}}} \sum_{j_{2}=1}^{n_{1}} \sum_{j_{3}=1}^{n_{2}} \sum_{j_{4}=1}^{n_{2}} \sum_{j_{5}=1}^{n_{4}} \sum_{j_{6}=1}^{n_{4}} \sum_{j_{7}=1}^{n_{4}} \mathcal{D}_{i_{1}, j_{1}, j_{2}}^{(1)} \mathcal{D}_{i_{2}, j_{3}, j_{4}}^{(2)} \mathcal{D}_{i_{4}, j_{5}, j_{6}, j_{7}}^{(4)} \mathcal{K}_{j_{1}, j_{3}}^{(1,2)} \mathcal{K}_{j_{2}, j_{5}}^{(1,4)} \mathcal{K}_{j_{4}, i_{3}, j_{6}}^{(2,3,4)} \mathcal{K}_{j_{7}, i_{5}}^{(4,5)} .
$$

Remark 5.4. We want to emphasize that this tensor network is closely related to the dual tensor network of the graphical model representing the transport plan [273]. In the context of graphical models, the cost tensors $\mathcal{C}$ is given in the form of (5.20). The resulting transport plan $\mathcal{P}_{\eta}^{(t)}$ defined in (5.3) can be represented as tensor network by attaching the matrices $\operatorname{diag}\left(\exp \left(\beta_{\ell}\right)\right)$ to the corresponding open modes of the tensor network representation of $\mathcal{K}$. At the same time, $\mathcal{P}_{\eta}^{(t)}$ represents a discrete probability distribution, which can be represented by a graphical model [150]. This graphical model and the tensor network of $\mathcal{P}_{\eta}^{(t)}$ are duals of each other [273].

We can compute $\mathcal{K}$ from the tensor network by contracting each of the internal edges sequentially. The contraction of one internal edge corresponds to the merging the
connected vertices by evaluating of the sum over the index corresponding to the edge [247]. When the tensor network contains circles, multi-edges will occur, which can be contracted by summing over all the corresponding indices simultaneously. The order of contracting the internal edges determines the degree of the occurring intermediate vertices and the computational complexity. In our constructed tensor network, we can exploit the special structure of $\mathcal{D}^{(\ell)}$ to contract all connected edges simultaneously. The book [272] discusses several heuristics to optimize the order of contractions. Note that the optimal contraction sequence might still incur a large computational cost.

In Algorithm 11, we need to evaluate the marginals of $\mathcal{P}_{\eta}^{(t)}$ in each iteration. Let $\gamma_{\ell}^{(t)}=\exp \left(\beta_{\ell}^{(t)}\right)$ for $\ell=1, \ldots, d$. Given a tensor network representation of $\mathcal{K}$, the mode- $\ell$ marginals can be written as a contraction of the network after connecting the matrix $\operatorname{diag}\left(\gamma_{\ell}^{(t)}\right)$ to the open edge corresponding to mode $\ell$, and the vectors $\gamma_{\tilde{\ell}}^{(t)}$ for $\tilde{\ell} \neq \ell$ to their respective open edges. By contracting all inner edges of the resulting network, we obtain $r_{\ell}\left(\mathcal{P}_{\eta}^{(t)}\right)$. We refer to Figure 5.2 for examples. The depicted networks will again be used in the numerical experiments in Section 5.5.

(a)

(b)

Figure 5.2 - Examples for the tensor network diagram representation of $r_{3}\left(\mathcal{P}_{\eta}^{(t)}\right)$ based on different tensor network structures for the Gibbs kernel (5.21).

We give a brief example on how to efficiently contract the network depicted in Figure 5.2(a). For the complexity analysis we assume $n_{1}=n_{2}=n_{3}=n_{4}=n$. We first compute the vectors $v_{j}^{(1)}=\sum_{i=1}^{n}\left(\gamma_{1}^{(t)}\right)_{i} \mathcal{K}_{i, j}^{(1,2)}$ and $v_{j}^{(4)}=\sum_{i=1}^{n} \mathcal{K}_{j, i}^{(3,4)}\left(\gamma_{4}^{(t)}\right)_{i}$ for $j=1, \ldots, n$ in $n(2 n-1)$ operations each, where operations refers to the required number of additions and multiplications. This corresponds to contracting the edge between $\gamma_{1}^{(t)}$ and $\mathcal{K}^{(1,2)}$ as well as the edge between $\mathcal{K}^{(3,4)}$ and $\gamma_{4}^{(t)}$. In the next step, we contract the edges from $\mathcal{D}^{(2)}$ to $v^{(1)}, \gamma_{2}^{(t)}$ and $\mathcal{K}^{(2,3)}$ simultaneously. This corresponds to computing the vector $v_{j}^{(2)}=\sum_{i=1}^{n}\left(v^{(1)} * \gamma_{2}^{(t)}\right)_{i} \mathcal{K}_{i, j}^{(2,3)}$ for $j=1, \ldots, n$ in $n+n(2 n-1)=2 n^{2}$ operations, where we first compute the elementwise product before evaluating the matrix vector product. Finally, we contract the remaining edges around $\mathcal{D}^{(3)}$ to compute $r_{3}\left(\mathcal{P}_{\eta}^{(t)}\right)=v^{(2)} * \gamma_{3}^{(t)} * v^{(4)}$ in $2 n$ operations. In total, we need $6 n^{2}$ operations to compute the marginal $r_{3}\left(\mathcal{P}_{\eta}^{(t)}\right)$. Note that this contraction strategy corresponds to the following reordering of the sums in
the definition of the marginal

$$
\begin{aligned}
& r_{3}\left(P_{\eta}^{(t)}\right)_{i_{3}}=\sum_{i_{1}=1}^{n} \sum_{i_{2}=1}^{n} \sum_{i_{4}=1}^{n} \mathcal{K}_{i_{1}, i_{2}}^{(1,2)} \mathcal{K}_{i_{2}, i_{3}}^{(2,3)} \mathcal{K}_{i_{3}, i_{4}}^{(3,4)}\left(\gamma_{1}^{(t)}\right)_{i_{1}}\left(\gamma_{2}^{(t)}\right)_{i_{2}}\left(\gamma_{3}^{(t)}\right)_{i_{3}}\left(\gamma_{4}^{(t)}\right)_{i_{4}} \\
& \quad=\left(\gamma_{3}^{(t)}\right)_{i_{3}}\left(\sum_{i_{2}=1}^{n} \mathcal{K}_{i_{2}, i_{3}}^{(2,3)}\left(\left(\gamma_{2}^{(t)}\right)_{i_{2}}\left(\sum_{i_{1}=1}^{n} \mathcal{K}_{i_{1}, i_{2}}^{(1,2)}\left(\gamma_{1}^{(t)}\right)_{i_{1}}\right)\right)\right)\left(\sum_{i_{4}=1}^{n} \mathcal{K}_{i_{3}, i_{4}}^{(3,4)}\left(\gamma_{4}^{(t)}\right)_{i_{4}}\right) .
\end{aligned}
$$

We can reuse the intermediate terms $v^{(1)}, v^{(2)}, v^{(4)}$ in the computation of the other marginals. This allows us to compute all four marginals in $12 n^{2}+4 n$ operations, whereas computing the marginals based on the full tensor requires $\mathcal{O}\left(n^{4}\right)$ operations.

Remark 5.5. Instead of the tensor network based on an ordinary graph with special tensors $\mathcal{D}^{(\ell)}$, we could consider a tensor network based on a hypergraph as in [273]. The structure of $\mathcal{D}^{(\ell)}$ can be modeled by a single hyperedge containing all vertices connected to $\mathcal{D}^{(\ell)}$. The contraction of these hypergraph based networks is studied in [174].

Remark 5.6. Note that the structure of $\mathcal{K}$ can also be exploited to compute marginals when applying Algorithm 12 to $\mathcal{P}_{\eta}^{(t)}$. Storing the rank-1 update in line 6 separately in terms of its factors offers the potential to avoid storing the transport plan explicitly as a full tensor.

### 5.4 Low-rank approximations in tensor networks

Assuming that the Gibbs kernel is represented as in Equation (5.21), we obtain a tensor network containing the coefficient tensors $\mathcal{K}^{\alpha}$. The following lemma bounds the impact on the Gibbs kernel when replacing each $\mathcal{K}^{\alpha}$ by an approximation $\tilde{\mathcal{K}}^{\alpha}$.

Lemma 5.2. Let $\mathcal{K}$ be defined based on tensors $\mathcal{K}^{\alpha} \in \mathbb{R}_{>0}^{n_{\alpha_{1}} \times \cdots \times n_{\alpha_{M}}}$ for $\alpha \in F$ as in (5.21). Let $\tilde{\mathcal{K}}^{\alpha} \in \mathbb{R}_{>0}^{n_{\alpha_{1}} \times \cdots \times n_{\alpha_{M}}}$ for $\alpha \in F$. We define $\tilde{\mathcal{K}} \in \mathbb{R}_{>0}^{n_{1} \times \cdots \times n_{d}}$ elementwise as

$$
\begin{equation*}
\tilde{\mathcal{K}}_{I}=\prod_{\alpha \in F} \tilde{\mathcal{K}}_{I_{\alpha}}^{\alpha} \quad \text { for every } I=\left(i_{1}, \ldots, i_{d}\right), \quad i_{\ell}=1, \ldots, n_{\ell}, \quad \ell=1, \ldots, d . \tag{5.22}
\end{equation*}
$$

Then

$$
\|\log (\mathcal{K})-\log (\tilde{\mathcal{K}})\|_{\infty} \leq \sum_{\alpha \in F}\left\|\log \left(\mathcal{K}^{\alpha}\right)-\log \left(\tilde{\mathcal{K}}^{\alpha}\right)\right\|_{\infty}
$$

Proof.

$$
\begin{aligned}
\|\log (\mathcal{K})-\log (\tilde{\mathcal{K}})\|_{\infty} & =\max _{i_{1}, \ldots, i_{d}}\left|\sum_{\alpha \in F} \log \left(\mathcal{K}_{I_{\alpha}}^{\alpha}\right)-\sum_{\alpha \in F} \log \left(\tilde{\mathcal{K}}_{I_{\alpha}}^{\alpha}\right)\right| \\
& \leq \max _{i_{1}, \ldots, i_{d}} \sum_{\alpha \in F}\left|\log \left(\mathcal{K}_{I_{\alpha}}^{\alpha}\right)-\log \left(\tilde{\mathcal{K}}_{I_{\alpha}}^{\alpha}\right)\right| \\
& \leq \sum_{\alpha \in F}\left\|\log \left(\mathcal{K}^{\alpha}\right)-\log \left(\tilde{\mathcal{K}}^{\alpha}\right)\right\|_{\infty} .
\end{aligned}
$$

Combining Lemma 5.2 and Theorem 5.2 implies that Algorithm 13 with $\tilde{\mathcal{K}}$ defined as in (5.22) yields an accurate approximation of the optimal transport plan when $\| \log \left(\mathcal{K}^{\alpha}\right)-$ $\log \left(\tilde{\mathcal{K}}^{\alpha}\right) \|_{\infty}$ is sufficiently small for all $\alpha \in F$. This allows one to replace each $\mathcal{K}^{\alpha}$ by a low-rank approximation $\tilde{\mathcal{K}}^{\alpha}$, which in turn accelerates the computation of tensor network contractions. For the example at the end of Section 5.3, the number of operations reduces from $\mathcal{O}\left(n^{2}\right)$ to $\mathcal{O}(n r)$ when every $\mathcal{K}^{\alpha}$ is approximated by a rank- $r$ matrix of the form $\mathcal{K}^{\alpha}=U^{\alpha}\left(V^{\alpha}\right)^{T}$ with $U^{\alpha}, V^{\alpha} \in \mathbb{R}^{n \times r}$ as depicted in Figure 5.3(a).

### 5.5 Numerical experiments

All numerical experiments* in this section were performed in MATLAB R2018b on a Lenovo Thinkpad T480s with Intel Core i7-8650U CPU and 15.4 GiB RAM. In Algorithms 11 and 13 we select the next index to be updated using index selection (5.7) and we stop using stopping criterion (5.8) with $\varepsilon_{\text {stop }}=10^{-4}$.

### 5.5.1 Proof of concept

In the following, we study the impact of approximating the Gibbs kernel on the transport cost. We define a multi-marginal optimal transport problem, whose cost tensor is of the form studied in [110, Section 5.2]. Let $n=420$. For $\ell=1, \ldots, 4$, we generate random point sets $X^{(\ell)}=\left\{x_{1}^{(\ell)}, \ldots, x_{n}^{(\ell)}\right\}$ by sampling the points $x_{i}^{(\ell)} \in \mathbb{R}^{2}$ independently randomly from the uniform distribution on $[0,1]^{2}$ for $i=1, \ldots, n$. We define $n \times n$ matrices $C^{(1,2)}, C^{(2,3)}, C^{(3,4)}$ with entries

$$
\mathcal{C}_{i, j}^{(1,2)}=\left\|x_{i}^{(1)}-x_{j}^{(2)}\right\|_{2}^{2}, \mathcal{C}_{j, k}^{(2,3)}=\left\|x_{j}^{(2)}-x_{k}^{(3)}\right\|_{2}^{2}, \mathcal{C}_{k, l}^{(3,4)}=\left\|x_{k}^{(3)}-x_{l}^{(4)}\right\|_{2}^{2} \quad \text { for } 1 \leq i, j, k, l \leq n .
$$

We construct the cost tensor

$$
\mathcal{C}_{i, j, k, l}=\mathcal{C}_{i, j}^{(1,2)}+\mathcal{C}_{j, k}^{(2,3)}+\mathcal{C}_{k, l}^{(3,4)} \quad \text { for } 1 \leq i, j, k, l \leq n .
$$

[^4]Let $\mathcal{K}^{\alpha}=\exp \left(-\mathcal{C}^{\alpha}\right)$ for $\alpha \in\{(1,2),(2,3),(3,4)\}$. The Gibbs kernel $\mathcal{K}=\exp (-\mathcal{C})$ is represented by the tensor network shown in Figure 5.2(a).

Let $r \leq n$. We compare two different approximations of $\mathcal{K}$. For the first one, we replace the matrices $\mathcal{K}^{\alpha}$ by their rank-r best approximations $\tilde{\mathcal{K}}^{\alpha}$ using truncated singular value decompositions and define

$$
\left(\mathcal{K}_{\mathrm{SVDs}}\right)_{i, j, k, l}=\tilde{\mathcal{K}}_{i, j}^{(1,2)} \cdot \tilde{\mathcal{K}}_{j, k}^{(2,3)} \cdot \tilde{\mathcal{K}}_{k, l}^{(3,4)} \quad \text { for } 1 \leq i, j, k, l \leq n
$$

For the second approximation $\mathcal{K}_{\mathbf{T T}}$, we compute a TT approximation 2.6 of $\mathcal{K}$ with TT rank ( $r, r, r$ ) using the TT-DMRG-cross algorithm [281] ignoring the underlying graph structure. The tensor network representation of $\mathcal{K}_{\text {SVDs }}$ and $\mathcal{K}_{\mathrm{TT}}$ is depicted in Figure 5.3. All four marginals can be computed in $\mathcal{O}(n r)$ operations for $\mathcal{K}_{\text {SVDs }}$ and in $\mathcal{O}\left(n r^{2}\right)$ operations for $\mathcal{K}_{\mathrm{TT}}$ by contracting the tensor networks. In contrast, exploiting the graph structure in $\mathcal{K}$ without low-rank approximations requires $\mathcal{O}\left(n^{2}\right)$ operations.

(a) Based on $\mathcal{K}_{\text {sVDs }}$

(b) Based on $\mathcal{K}_{\mathbf{T T}}$

Figure 5.3 - Tensor networks for the computation of $r_{3}\left(\mathcal{P}_{\eta}^{(t)}\right)$ for the example in Section 5.5.1. We express the truncated SVDs in $\mathcal{K}_{\text {SVDs }}$ as $\tilde{\mathcal{K}}^{\alpha}=U^{\alpha}\left(V^{\alpha}\right)^{T}$ with $U^{\alpha}, V^{\alpha} \in R^{n \times r}$. The TT cores in $\mathcal{K}_{\mathrm{TT}}$ are denoted by $\mathcal{G}_{1} \in \mathbb{R}^{n \times r}, \mathcal{G}_{2}, \mathcal{G}_{3} \in \mathbb{R}^{r \times n \times r}$ and $\mathcal{G}_{4} \in \mathbb{R}^{r \times n}$.

Based on the tensors $\mathcal{K}, \mathcal{K}_{\mathrm{SVDs}}, \mathcal{K}_{\mathrm{TT}}$, we compute transport plans $\mathcal{P}, \mathcal{P}_{\mathrm{SVDs}}, \mathcal{P}_{\mathrm{TT}}$ by first applying Algorithm 13 with marginals $r_{\ell}=\frac{1}{n} \cdot \mathbf{1}_{n}$ and then rounding the resulting tensor using Algorithm 12. In Figure 5.4, we compare the different transport plans. Note that we can efficiently evaluate the transport cost $\langle\mathcal{C}, \mathcal{P}\rangle$ using tensor network contractions of $\mathcal{C}^{\alpha}$ and $\mathcal{P}$ without evaluating the full tensors. We observe that the difference in transport cost of $\mathcal{P}_{\mathrm{SVDs}}, \mathcal{P}_{\mathrm{TT}}$ and $\mathcal{P}$ is much smaller than the norm of the difference of the logarithms of $\mathcal{K}_{\text {SVDs }}, \mathcal{K}_{\mathrm{TT}}$ and $\mathcal{K}$. The approximation $\mathcal{P}_{\text {SVDs }}$ that exploits the graph structure leads to slightly better approximations compared to $\mathcal{P}_{\text {TT }}$. We want to emphasize that computing $\mathcal{P}_{\text {SVDs }}$ with $r=25$ is faster than using the graph structure of $\mathcal{K}$ directly and only leads to a difference in transport cost in the order of machine precision. Computing $\mathcal{P}_{\mathrm{TT}}$ is
faster than computing $\mathcal{P}$ for very small ranks. The different scaling in the number of operations required to compute marginals leads to larger computation times for $\mathcal{P}_{\mathrm{T} T}$ compared to $\mathcal{P}_{\text {SVDs }}$ for increasing values of $r$. We want to emphasize that storing a tensor in $\mathbb{R}^{n \times n \times n \times n}$ explicitly would require more than 240 GB of memory. Thus, it would not be feasible to solve this problem without exploiting either the underlying structure of $\mathcal{C}$ or the structure of the TT approximation.


Figure 5.4 - Difference in transport cost for the example in Section 5.5.1 for various ranks $r$. Left: We depict the difference in transport cost for $\mathcal{P}_{\mathrm{SVDs}}, \mathcal{P}_{\mathrm{TT}}$ and $\mathcal{P}$ and an estimation of the norm of the difference of the logarithms of $\mathcal{K}_{\text {SVDs }}, \mathcal{K}_{\mathrm{TT}}$ and $\mathcal{K}$ based on 1000 sample points. Right: Measured computation times for applying Algorithm 13 and 12 to compute the transport plans exploiting the structures of $\mathcal{K}, \mathcal{K}_{\text {SVDs }}, \mathcal{K}_{\mathrm{TT}}$. Note that this time does not include the computation of $\mathcal{K}^{\alpha}, \mathcal{K}_{\text {SVDs }}, \mathcal{K}_{\mathrm{TT}}$.

Remark 5.7. Figure 5.4 shows that larger ranks lead to smaller approximation errors, but at the same time larger ranks increase the computation time. This needs to be balanced in practice. In particular, the rank needs to be sufficiently large such that the approximation $\tilde{\mathcal{K}}$ of the Gibbs kernel is strictly positive, which implies that $\|\log (\mathcal{K})-\log (\tilde{\mathcal{K}})\|_{\infty}$ is bounded. This can be achieved by choosing an approximation such that $\|\mathcal{K}-\tilde{\mathcal{K}}\|_{\infty}$ is smaller than the smallest entry of $\mathcal{K}$.

Remark 5.8. The difference of the entropic cost for the tensors $\mathcal{P}_{S V D S}, \mathcal{P}_{T T}$ and the optimal transport plan $\mathcal{P}_{\eta}^{*}$ is bounded by Theorem 5.2. We can assume that $\mathcal{P}$ is a good approximation of $\mathcal{P}_{\eta}^{*}$. This would allow us to study the sharpness of the bound numerically. However, the evaluation of the entropic cost requires the explicit computation of the full tensors, which is not feasible for $n=420$ without a Monte Carlo approximation. Instead, we repeat the experiment in Section 5.5 .1 with a smaller $n$. The results are depicted in Figure 5.5. We find that the theoretical bound is much larger than the error observed in practice.


Figure 5.5 - We repeat the experiment described in Section 5.5 .1 with $n=10$ for various ranks $r$. We depict the difference in the entropic transport cost (5.2) of $\mathcal{P}_{\text {SVDs }}, \mathcal{P}_{\mathrm{TT}}$ and $\mathcal{P}$. Further, we depict the norm of the difference of the elementwise logarithm of $\mathcal{K}_{\text {SVDs }}, \mathcal{K}_{\mathrm{TT}}$ and $\mathcal{K}$. Based on these values we compute the value of $\varepsilon_{V_{\mathcal{C}}^{1}}$ as defined in (5.9).

### 5.5.2 Application: Color transfer from color barycenters

In the following, we apply our algorithms for color transfer as in [150]. We consider $d=4$ images containing $n=100^{2}$ pixels each. Let $0 \leq \lambda_{1}, \lambda_{2}, \lambda_{3}$ such that $\lambda_{1}+\lambda_{2}+\lambda_{3}=1$. In a first step, we compute an approximation of the Wasserstein barycenter [2] with weights $\lambda=\left(\lambda_{1}, \lambda_{2}, \lambda_{3}\right)$ of the color of the first three images by solving the multi-marginal optimal transport problem in [37, Section 4.2]. We then transfer the color from the approximation of the barycenter onto the fourth image by solving a two-marginal optimal transport problem [267].

Let $x_{i}^{(\ell)} \in[0,1]^{3}$ denote the color of pixel $i$ in image $\ell$, where we treat the RGB values as element in $\mathbb{R}^{3}$ and rescale to $[0,1]^{3}$. Let $x_{i}^{(B)}=\lambda_{1} x_{i}^{(1)}+\lambda_{2} x_{i}^{(2)}+\lambda_{3} x_{i}^{(3)}$ for $i=1, \ldots, n$. We choose these points to define a reference template [335] for computing an approximation of the barycenter as in [37]. Let $\mathcal{C}_{i, j}^{(\ell, 4)}=\left\|x_{i}^{(\ell)}-x_{j}^{(B)}\right\|_{2}^{2}$ for $1 \leq i, j \leq n$ and $\mathcal{K}^{(\ell, 4)}=\exp \left(-\mathcal{C}^{(\ell, 4)} / \eta\right)$ for $1 \leq \ell \leq 3$. We define the cost tensor

$$
\mathcal{C}_{i, j, k, l}=\lambda_{1} \mathcal{C}_{i, l}^{(1,4)}+\lambda_{2} \mathcal{C}_{j, l}^{(2,4)}+\lambda_{3} \mathcal{C}_{k, l}^{(3,4)} \quad \text { for } 1 \leq i, j, k, l \leq n,
$$

and the Gibbs kernel tensor $\mathcal{K}=\exp (-\mathcal{C} / \eta)$ for a given regularization parameter $\eta \geq 0$. Following the ideas in $[110,150]$, we compute $r_{B}=r_{4}\left(\mathcal{P}_{\eta}^{*}\right)$, where

$$
\begin{equation*}
\mathcal{P}_{\eta}^{*}=\mathcal{K} \times_{1} \operatorname{diag}\left(\exp \left(\beta_{1}\right)\right) \times_{2} \operatorname{diag}\left(\exp \left(\beta_{2}\right)\right) \times_{3} \operatorname{diag}\left(\exp \left(\beta_{3}\right)\right) \times_{4} \operatorname{diag}\left(\mathbf{1}_{n}\right), \tag{5.23}
\end{equation*}
$$

with scaling parameters $\beta_{1}, \beta_{2}, \beta_{3} \in \mathbb{R}^{n}$ chosen such that $r_{\ell}\left(\mathcal{P}_{\eta}^{*}\right)=\mathbf{1}$ for $1 \leq \ell \leq 3$. To compute an approximation $\mathcal{P}_{\eta}$ of $\mathcal{P}_{\eta}^{*}$ we run Algorithm 11 with cost tensor $\mathcal{C}$ and marginals $r_{1}=r_{2}=r_{3}=\mathbf{1}_{n}$ and index selection (5.7), which we modify to use arg max $\ell \in\{1,2,3\}$
instead of $\underset{\ell \in\{1,2,3,4\}}{\arg \max }$. This modification ensures that we only update the first three scaling parameters, which leads to an approximation of the form (5.23) [151, Theorem 3.5]. The Gibbs kernel corresponds to a star shaped graph structure as in Figure 5.2(b). The approximation of the color barycenter is now given by the points $x^{(B)}$ with masses $r_{B}$.

Remark 5.9. Proposition 3.4 in [151] states that multi-marginal optimal transport problems with star shaped graph structures can be decomposed into several independent two-marginal problems when all marginals are prescribed. This does not apply for the computation of (5.23), since $r_{4}\left(\mathcal{P}_{\eta}^{*}\right)$ is unknown.

We now want to transfer the color from the approximation of the barycenter to the fourth image. We define the cost matrix $C_{i, j}=\left\|x_{i}^{(B)}-x_{j}^{(4)}\right\|_{2}^{2}$ for $1 \leq i, j \leq n$ and Gibbs kernel matrix $K=\exp (-C / \eta)$. Let the matrix $P$ denote the approximate solution obtained from Algorithm 11 with cost matrix $C$ and marginals $r_{1}=r_{B}$ and $r_{2}=\mathbf{1}_{n}$. The color vector of the target image with transferred colors is now given by $x_{j}^{*}=\sum_{i=1}^{n} P_{i, j} x_{i}^{(B)}$ for $1 \leq j \leq n$. Note that we can transfer the color to several target images without recomputing the barycenter approximation.

To accelerate the computation of marginals, we replace the Gibbs kernel tensor $\mathcal{K}$ by an approximation $\tilde{\mathcal{K}}$ obtained by replacing $\mathcal{K}^{(1,4)}, \mathcal{K}^{(2,4)}, \mathcal{K}^{(3,4)}$ by rank-r approximations using the randomized SVD [156]. We also replace the Gibbs kernel matrix $K$ by a rank- $r$ approximation $\tilde{K}$. Marginals and the target color vector $x^{*}$ can now be computed in $\mathcal{O}(n r)$ operations.

In the following numerical experiments, we set $\eta=1 / 10$ and compute approximate transport plans $\tilde{\mathcal{P}}, \tilde{P}$ by applying Algorithm 13 to $\tilde{\mathcal{K}}$ and $\tilde{K}$. In Figure 5.6, we study the impact of $\lambda$ and $r$ onto the color transferred image for example images from the COCO data set [218]. We observe that small values of $r$ suffice to accurately approximate the desired image. The computation including the assembling of the matrices and the randomized SVD takes 0.25 seconds for $r=50$, whereas using the full matrices in $\mathbb{R}^{10000 \times 10000}$ directly in the tensor network takes 7.65 seconds, i.e. our low-rank approach reduces the computation time by over $96 \%$.

Remark 5.10. When we are only interested in transferring the color onto a single target image, we can alternatively solve the following multi-marginal optimal transport problem without computing an approximation of the Barycenter. Based on the matrices $\hat{\mathcal{C}}_{i, j}^{(\ell, 4)}=\left\|x_{i}^{(\ell)}-x_{j}^{(4)}\right\|_{2}^{2}$ for $\ell=1,2,3$, we define the cost tensor

$$
\hat{\mathcal{C}}_{i, j, k, l}=\lambda_{1} \hat{\mathcal{C}}_{i, l}^{(1,4)}+\lambda_{2} \hat{\mathcal{C}}_{j, l}^{(2,4)}+\lambda_{3} \hat{\mathcal{C}}_{k, l}^{(3,4)} \quad \text { for } 1 \leq i, j, k, l \leq n .
$$

Let $\hat{\mathcal{P}}$ denote the transport plan obtained by solving the corresponding optimal transport problem with marginals $r_{\ell}=\mathbf{1}_{n}$ for $\ell=1, \ldots, 4$. To accelerate the solution of this problem, we can again replace the matrices $\exp \left(-\hat{\mathcal{C}}^{(1,4)} / \eta\right), \exp \left(-\hat{\mathcal{C}}^{(2,4)} / \eta\right), \exp \left(-\hat{\mathcal{C}}^{(3,4)} / \eta\right)$


Figure 5.6 - Given the images displayed in the top row of (c), we use the method described in Section 5.5.2 to transfer their color to the bottom right image in (c). For fixed $\lambda=(1 / 3,1 / 3,1 / 3)$, we plot the resulting images for ranks $r=3, r=5, r=10, r=50$ from left to right in (a). The rightmost picture is obtained by using the full matrices $\mathcal{K}^{\alpha}$ and $K$ directly. In (b), we plot $\left\|\tilde{x}^{r}-\tilde{x}^{*}\right\|_{\infty}$ where $\tilde{x}^{r}$ denotes the resulting image vector for a given rank $r$ and $\tilde{x}^{*}$ is computed using the full matrices. Moreover, we display the resulting image for different values of $\lambda$ in (c): middle row left to right $\lambda=(1,0,0), \lambda=(0,1,0), \lambda=(0,0,1)$, bottom row left $\lambda=(1 / 3,2 / 3,0)$, bottom row middle $\lambda=(1 / 5,1 / 5,3 / 5)$.
by rank-r approximations. Note that it would also be possible to decompose the problem into several independent two-marginal problems using Proposition 3.4 in [151]. The color vector of the target image with transferred colors can now be defined as $x_{l}^{*}=$ $\sum_{i=1}^{n} \sum_{j=1}^{n} \sum_{k=1}^{n} P_{i, j, k, l}\left(\lambda_{1} x_{i}^{(1)}+\lambda_{2} x_{j}^{(2)}+\lambda_{3} x_{k}^{(3)}\right)$ for $1 \leq l \leq n$. In Figure 5.7, it is shown that the resulting target image seems to be similar to the image obtained after solving the Barycenter problem (see Figure 5.6). However, we actually solve different
problems yielding different solutions. This is demonstrated in our experiments, where the maximum norm of the difference of the color vector obtained using this approach to the color vector obtained using the barycenter approach described in Section 5.5.2 is around 94. Also note that the error introduced by using rank-r approximations using this alternative approach is much smaller compared to using rank-r approximations using the barycenter approach.


Figure 5.7 - We repeat the experiments used to generate Figure 5.6 using the alternative approach described in Remark 5.10.

### 5.5.3 A tensor network with circles

In the following, we describe an optimal transport problem arising in the context of Schrödinger bridges $[96,151]$. Let $m=5$ and $n=40^{2}$. We denote by $x^{(i)}$ for $1 \leq i \leq n$ the $i$ th grid point on the grid $\{1, \ldots, 40\}^{2}$. Let $C_{i, j}=\left\|x^{(i)}-x^{(j)}\right\|_{2}^{2}$ for $1 \leq i, j \leq n$ and
$K=\exp (-C / \eta)$. We consider the Gibbs kernel

$$
\begin{equation*}
\mathcal{K}_{i_{1}, i_{2}, i_{3}, i_{4}, i_{5}}=\prod_{\alpha \in F} K_{I_{\alpha}} \quad \text { for } 1 \leq i_{1}, i_{2}, i_{3}, i_{4}, i_{5} \leq n \tag{5.24}
\end{equation*}
$$

where $F=\{(1,2),(2,3),(3,4),(4,5)\}$. Note that this corresponds to a tensor network structure similar to Figure 5.2(a). Given $r_{1}, r_{5} \in \Delta^{n}$, we now consider the problem of finding scaling parameters $\beta_{1}, \beta_{5} \in \mathbb{R}^{n}$ such that

$$
\begin{equation*}
\mathcal{P}_{\eta}^{*}=\mathcal{K} \times_{1} \operatorname{diag}\left(\exp \left(\beta_{1}\right)\right) \times_{2} \operatorname{diag}(\mathbf{1}) \times_{3} \operatorname{diag}(\mathbf{1}) \times_{4} \operatorname{diag}(\mathbf{1}) \times_{5} \operatorname{diag}\left(\exp \left(\beta_{5}\right)\right) \tag{5.25}
\end{equation*}
$$

satisfies $r_{1}\left(\mathcal{P}_{\eta}^{*}\right)=r_{1}$ and $r_{5}\left(\mathcal{P}_{\eta}^{*}\right)=r_{5}$. In the context of Schrödinger bridges the marginals $r_{\ell}\left(\mathcal{P}_{\eta}^{*}\right)$ for $\ell=2,3,4$ describe how the initial distribution $r_{1}$ most likely evolved into $r_{5}$ [151]. As in Section 5.5.2, we can again compute approximations of $\mathcal{P}_{\eta}^{*}$ by modifying Algorithm 13 such that only $\beta_{1}$ and $\beta_{5}$ are updated based on the prescribed marginals $r_{1}, r_{5}$; see [150, Section III.A].


Figure 5.8 - Tensor network diagram representation of $\mathcal{K}$ as defined in (5.24). The black network is obtained for $F=\{(1,2),(2,3),(3,4),(4,5)\}$. The blue part depicts the additional nodes and vertices introduced by setting $F=$ $\{(1,2),(1,3),(2,3),(2,4),(3,4),(3,5),(4,5)\}$.

The Schrödinger bridge problem is based on a Markov chain model, in the sense that each distribution only depends on the previous distribution. We now introduce additional dependencies by setting $F=\{(1,2),(1,3),(2,3),(2,4),(3,4),(3,5),(4,5)\}$ in Equation (5.24). The tensor network structure of $\mathcal{K}$ is depicted in Figure 5.8. Note that this results in a graphical model for $\mathcal{P}_{\eta}^{*}$ with window graph structure as in [7, Figure 2]. In Figure 5.9, we depict the corresponding tensor network after replacing each matrix $K$ by a rank- $r$ approximation. Marginals of this network can be computed in $\mathcal{O}\left(n r^{4}\right)$ operations. For instance, to compute $r_{3}\left(\mathcal{P}_{\eta}^{(t)}\right)$ we first evaluate the colored tensors $\mathcal{T}^{[1]}, \ldots, \mathcal{T}^{[5]}$ in $\mathcal{O}\left(n r^{4}\right)$ operations by contracting the colored edges simultaneously using the structure of $\mathcal{D}^{(k)}$. We then compute $\mathcal{T}^{[1,2]} \in \mathbb{R}^{r \times r \times r}$ in $\mathcal{O}\left(r^{4}\right)$ operations by contracting $\mathcal{T}^{[1]}$ and $\mathcal{T}^{[2]}$ along their common edge. Analogously we compute $\mathcal{T}^{[4,5]}$. We then contract $\mathcal{T}^{[1,2]}$ and $\mathcal{T}{ }^{[4,5]}$ along their common edge in $\mathcal{O}\left(r^{5}\right)$ operations before contracting all edges of the resulting tensor simultaneously with $\mathcal{T}^{[3]}$ in $\mathcal{O}\left(n r^{4}\right)$ operations. We want to emphasize that contracting the network without replacing $K$ by low-rank approximations would not be feasible due to the required memory for storing the intermediate tensors.

In Figure 5.10, we study the impact of introducing these additional dependencies on the


Figure 5.9 - Tensor network diagram representation of $r_{3}\left(\mathcal{P}_{\eta}^{(t)}\right)$ corresponding to the graph structure (5.24) with $F=\{(1,2),(1,3),(2,3),(2,4),(3,4),(3,5),(4,5)\}$. Here, we replace each matrix $K$ by $U V^{T}$ with $U, V \in \mathbb{R}^{n \times r}$. The colors mark subtensors $\mathcal{T}^{[1]}, \ldots \mathcal{T}^{[5]}$ computed during the contraction.


Figure 5.10 - We consider the scaling problem (5.25) for $\eta=0.1$. We obtain an approximation $\tilde{\mathcal{K}}$ of the Gibbs kernel by replacing $\hat{\mathcal{K}}$ in (5.24) by a rank-10 approximation computed via the randomized SVD. We apply Algorithm 13 with index selection (5.7), where we take $\arg$ max only over the set $\{1,5\}$. Let $\mathcal{P}$ denote the resulting transport plan. In each row, we depict the prescribed $r_{1}, r_{5}$ as well as $r_{2}(\mathcal{P}), r_{3}(\mathcal{P}), r_{4}(\mathcal{P})$ reshaped into $\mathbb{R}^{40 \times 40}$. The top row is obtained using $F=\{(1,2),(2,3),(3,4),(4,5)\}$ as in the Schrödinger bridge setting. The bottom row is obtained using $F=\{(1,2),(1,3),(2,3),(2,4),(3,4),(3,5),(4,5)\}$ with additional dependencies.
marginals $r_{\ell}(\mathcal{P})$, where $\mathcal{P}$ denotes the computed approximation of (5.25). We observe that the additional dependencies lead to a more concentrated $r_{3}(\mathcal{P})$, in the sense that the mass is less spread out. Moreover, the additional dependencies lead to $r_{2}(\mathcal{P})$ and $r_{4}(\mathcal{P})$ being concentrated slightly closer to the center of the images.

Remark 5.11. In this experiment, the graphical model of the transport plan contains circles. Hence, we can not apply the belief propagation algorithm directly [118, 150]. There exist generalizations such as the loopy belief propagation algorithm [181], which does not always converge, and the junction tree algorithm [173], which applies the belief propagation on a so-called junction tree. This junction tree encodes the connection structure of the vertices. It can be used to determine the contraction order for the corresponding tensor network representation of the Gibbs kernel [273, Section 4.2]. In particular, the minimal cost to contract the tensor network is bounded from above by the cost of the junction tree
algorithm.

## 6 Self-diffusion matrix

This chapter is concerned with the computation of the self-diffusion matrix of a tagged particle process on a grid. In contrast to classical methods based on estimating long-time averages of empirical means of deviations of some stochastic processes on a finite-size supercell, we propose a novel numerical method based on solving the equivalent highdimensional optimization problem formulated in [42, 210, 266] using an alternating linear scheme [39, 171, 274]. Once we obtained an approximation of the optimal solution, we use a carefully tuned variance reduction method to evaluate the desired entries of the self-diffusion matrix. To illustrate the method, we compute the self-diffusion matrix of a two-dimensional tagged particle process defined on a Cartesian grid. Our numerical experiments demonstrate that our novel low-rank approach is much less subject to statistical noise and that it yields a very accurate approximation of the self-diffusion matrix.

In addition, we present so-called cross-diffusion systems [49, 51, 182, 183] as example for applications, in which self-diffusion matrices play an important role. These systems appear in various fields such as population dynamics [290], tumor growth in medical biology [176], the modeling of biological systems [313] and lattice gases [13, 14, 111, 288], and the simulation of diffusion processes within mixtures of chemical compounds in materials science [23, 225]. When these systems are derived as hydrodynamic limits of lattice-based stochastic hopping models at the microscopic level [111, 266], they involve self-diffusion matrices as coefficients. Computing these self-diffusion coefficients is a major challenge that needs to be overcome before we can solve the cross-diffusion system numerically using for instance finite volume [116] or finite element methods [48].

This chapter is based on the articles [84, 85]. Its remainder is structured as follows. In Section 6.1, we recall the two equivalent definitions of the self-diffusion matrix of the tagged particle process. In Section 6.2, we construct finite-dimensional approximations of both definitions. The low-rank approach we propose here in order to compute a numerical approximation is presented in Section 6.3. The efficiency of our approach is illustrated
through several numerical experiments presented in Section 6.4. Finally, in Section 6.5, we develop a cell-centered finite volume method to simulate a cross-diffusion system involving self-diffusion matrices as coefficients.

### 6.1 Infinite-dimensional definition

Let $d=1,2,3$ denote the physical dimension of the problem. We consider a symmetric tagged particle process [276] defined on the infinite grid $\mathbb{Z}^{d}$. Let $K \in \mathbb{N}$ denote the number of possible jump directions for the particles. Let $\left(v_{k}\right)_{1 \leq k \leq K} \subset \mathbb{Z}^{d} \backslash\{\mathbf{0}\}$ denote the set of possible jump directions. For all $1 \leq k \leq K$, the probability of jumping in the direction $v_{k}$ is denoted by $\left.\left.p_{k} \in\right] 0,1\right]$. This jumping scheme is assumed to be symmetric in the sense that if the jump in the direction $v_{k}$ occurs with probability $p_{k}$, then the jump in the direction $-v_{k}$ occurs with the same probability.

The self-diffusion matrix application

$$
D_{s}:\left\{\begin{array}{ccc}
{[0,1]} & \rightarrow & \mathbb{R}^{d \times d} \\
\rho & \mapsto & D_{s}(\rho):=\left(D_{s, i j}(\rho)\right)_{1 \leq i, j \leq d}
\end{array}\right.
$$

can be defined in the following two ways.

### 6.1.1 Definition as optimization problem

Let us first introduce some notation. Let $S:=\mathbb{Z}^{d} \backslash\{\mathbf{0}\}$. For all so-called environments $\eta:=\left(\eta_{s}\right)_{s \in S} \in\{0,1\}^{S}$ and all $y \neq z \in S$, we define by $\eta^{y, z}:=\left(\eta_{s}^{y, z}\right)_{s \in S}$ the element of $\{0,1\}^{S}$ such that

$$
\eta_{s}^{y, z}:= \begin{cases}\eta_{s} & \text { if } s \neq y, z \\ \eta_{y} & \text { if } s=z \\ \eta_{z} & \text { if } s=y\end{cases}
$$

Furthermore, for all $w \in S$, we define by $\eta^{\mathbf{0}, w}:=\left(\eta_{s}^{\mathbf{0}, w}\right)_{s \in S}$ the element of $\{0,1\}^{S}$ such that

$$
\eta_{s}^{\mathbf{0}, w}:= \begin{cases}\eta_{s+w} & \text { if } s \neq-w \\ 0 & \text { if } s=-w\end{cases}
$$

For a mean particle density $\rho \in[0,1]$, we denote by $\rho^{\otimes}$ the Bernoulli product measure on $\{0,1\}^{S}$ with constant marginals equal to $\rho$ in each mode. Let us also define the set of functions $H_{\rho}:=L_{\rho^{\otimes}}^{2}\left(\{0,1\}^{S}\right)$, where $L_{\rho \otimes}^{2}$ denotes the $L^{2}$ Lebesgue space with measure $\rho^{\otimes}$.

For a drift direction $u \in \mathbb{R}^{d}$ and mean particle density $\rho \in[0,1]$, the self-diffusion
coefficient $u^{T} D_{s}(\rho) u$ is given by:

$$
\begin{align*}
& u^{T} D_{s}(\rho) u:=2 \inf _{\Psi \in H_{\rho}} \mathbb{E}_{\rho^{*}}\left[\sum _ { k = 1 } ^ { K } p _ { k } \left(\left(1-\eta_{v_{k}}\right)\left(u \cdot v_{k}+\Psi\left(\eta^{\mathbf{0}, v_{k}}\right)-\Psi(\eta)\right)^{2}\right.\right. \\
&\left.\left.+\frac{1}{2} \sum_{\substack{y \in S \\
y+v_{k} \neq \mathbf{0}}}\left(\Psi\left(\eta^{y+v_{k}, y}\right)-\Psi(\eta)\right)^{2}\right)\right] \tag{6.1}
\end{align*}
$$

where the notation $\mathbb{E}_{\rho \otimes}$ refers to the fact that the expectation is taken over the product measure $\rho^{\otimes}[200,266]$, i.e. we need to consider the expectation over all possible environments $\eta$. Problem (6.1) thus reads as an infinite-dimensional optimization problem over the set $H_{\rho}$.

Remark 6.1. Naturally, for all $\rho \in[0,1]$, since $D_{s}(\rho)$ is a symmetric matrix, one can easily deduce the full matrix $D_{s}(\rho)$ from the knowledge of $u^{T} D_{s}(\rho) u$ for three vectors $u \in \mathbb{R}^{d}$.

### 6.1.2 Definition as long time mean square deviation

The quantity $u^{T} D_{s}(\rho) u$ can be equivalently expressed as the long time limit of the following expectation [266, Theorem 2.3]. Let us assume that, at time $t=0$, the tagged particle is located at position $\mathbf{0}$ and all other sites are occupied with probability $\rho$ following a Bernoulli distribution. The duration between two consecutive jumping events follows an exponential law with parameter 1. The jumping directions (respectively rates) are given by $\left\{v_{1}, \ldots, v_{K}\right\}$ (respectively $p_{1}, \ldots, p_{K}$ ). Jumps are not allowed if the final site of the jumping particle is already occupied. Let $w(t)$ denote the position of the tagged particle at time $t \geq 0$. It then holds that, for a drift direction $u \in \mathbb{R}^{d}$ and mean particle density $\rho \in[0,1]$, the quantity $u^{T} D_{s}(\rho) u$ can be equivalently formulated as the long-time limit mean square deviation of the tagged particle in the direction $u \in \mathbb{R}^{d}$, i.e.

$$
\begin{equation*}
u^{T} D_{s}(\rho) u=2 \lim _{t \rightarrow \infty} \frac{\mathbb{E}_{\rho^{\otimes}}\left[\langle u, w(t)\rangle^{2}\right]}{t} \tag{6.2}
\end{equation*}
$$

Remark 6.2. Starting from a Bernoulli-product measure and symmetric transition rates, it is a classical problem in probability theory to study the motion of a tagged particle on $\mathbb{Z}^{d}$ [192, 215, 276] and more recently also on trees [68, 132]. For $d \geq 2$ and symmetric nearest neighbor transition rates, the tagged particle is known to satisfy a central limit theorem with non-degenerate limiting variance. However, to our best knowledge, a general closed formula for the limiting variance is not available.

Notice that when starting from a Bernoulli-product measure with a tagged particle in the origin, the resulting environment process is stationary with respect to the Bernoulli-product measure conditioned to contain a particle in the origin, usually called the Palm measure. This suggests that the limiting variance for the tagged particle can be described in terms of
the sum of separable functions, i.e. in terms of a low-rank function. In turn, this indicates that the equivalent characterization of the limiting variance as in Equation (6.1) is also related to low-rank functions.

### 6.2 Finite-dimensional approximation

The quantity $u^{T} D_{s}(\rho) u$ can not be computed exactly in practice, neither using expression (6.1) nor expression (6.2). On the one hand, (6.1) reads as an infinite-dimensional optimization problem and has to be approximated by a finite-dimensional optimization problem in practice [178, 210]. On the other hand, (6.2) requires the computation of the long-time average of the stochastic process defined on the infinite lattice grid $\mathbb{Z}^{d}$. Equivalently, it has to be approximated in a finite-dimensional setting, as long-time average of the mean-square deviation of the tagged particle associated to a stochastic process defined on a finite-size grid with periodic boundary conditions. Both finite-dimensional approximations are presented in detail in the following.

### 6.2.1 Discretized minimization problems

Let $M \in \mathbb{N}$ denote a discretization parameter and introduce the finite grid $S_{M}:=$ $\{-M, \cdots, M\}^{d} \backslash\{\boldsymbol{0}\}$. For the sake of simplicity, we assume that $M \geq \max _{k \in\{1, \ldots, K\}}\left\|v_{k}\right\|_{1}$. For any $\eta \in\{0,1\}^{S_{M}}$, we can construct by periodicity an extension $\widetilde{\eta}:=\left(\widetilde{\eta}_{s}\right)_{s \in S} \in\{0,1\}^{S}$ by assuming with a slight abuse of notation that the site $\mathbf{0}$ is occupied, i.e. $\widetilde{\eta}_{s}=1$ for $s \in(2 M+1) \cdot \mathbb{Z}^{d} \backslash\{\mathbf{0}\}$. Using this notation, for all $\eta \in\{0,1\}^{S_{M}}$ and all $y, z, w \in S$, we define $\eta^{y, z} \in\{0,1\}^{S_{M}}$ and $\eta^{\mathbf{0}, w} \in\{0,1\}^{S_{M}}$ as

$$
\eta^{y, z}:=\left(\widetilde{\eta}_{s}^{y, z}\right)_{s \in S_{M}} \quad \text { and } \quad \eta^{\mathbf{0}, w}:=\left(\widetilde{\eta}_{s}^{\mathbf{0}, w}\right)_{s \in S_{M}} .
$$

Figure 6.1 shows an illustration of $\eta^{y, z}$ and $\eta^{\mathbf{0}, w}$.
Let $N:=(2 M+1)^{d}-1=\operatorname{Card}\left(S_{M}\right)$. For all $\ell \in\{0, \ldots, N\}$, let $C_{M, \ell}:=\{\eta \in$ $\left.\{0,1\}^{S_{M}} \mid \sum_{s \in S_{M}} \eta_{s}=\ell\right\}$ denote the set of all possible configurations of the particles on $S_{M}$ so that the total number of occupied sites is equal to $\ell$.

Let us define $H_{M}:=\left\{\Psi:\{0,1\}^{S_{M}} \rightarrow \mathbb{R}\right\}$. For every $u \in \mathbb{R}^{d}$ and $l \in\{0, \ldots, N\}$, we introduce the quadratic functional $A_{M, \ell}^{u}: H_{M} \rightarrow \mathbb{R}$ defined by

$$
\begin{align*}
& A_{M, \ell}^{u}(\Psi):=\frac{1}{\left|C_{M, \ell}\right|} \sum_{\eta \in C_{M, \ell}} \sum_{k=1}^{K} p_{k}\left[\left(1-\eta_{v_{k}}\right)\left(u \cdot v_{k}+\Psi\left(\eta^{\mathbf{0}, v_{k}}\right)-\Psi(\eta)\right)^{2}\right. \\
&\left.+\frac{1}{2} \sum_{\substack{y \in S_{M} \\
y+v_{k} \neq \mathbf{0}}}\left(\Psi\left(\eta^{y+v_{k}, y}\right)-\Psi(\eta)\right)^{2}\right] . \tag{6.3}
\end{align*}
$$



Figure 6.1 - Middle: Visualization of one particular $\eta \in S_{1}$ for $d=2$ with three occupied sites marked in blue, cyan and magenta. Additionally, we mark the tagged particle $\mathbf{0}$ in red. Left: Visualization of the occupied sites of $\eta^{\mathbf{0},(-1,0)}$. This can be seen as a jump of the imaginary red particle one step to the left, followed by immediately relabeling of the sites such that red particle remains at $\mathbf{0}$. By exploiting the periodicity, we obtain $\eta^{\mathbf{0},(-1,0)} \in S_{1}$. Right: Visualization of the occupied sites of $\eta^{(1,1)+(-1,0),(1,1)}$. This can be seen as jump of the cyan particle one step to the left.

Then, for some $\ell \in\{0, \ldots, N\}$, one can define [42] for all $u \in \mathbb{R}^{d}$,

$$
\begin{equation*}
u^{T} D_{s}^{M}\left(\frac{\ell}{N}\right) u:=2 \min _{\Psi \in H_{M}} A_{M, \ell}^{u}(\Psi) \tag{6.4}
\end{equation*}
$$

For any $\bar{\rho} \in[0,1]$, it is proved in [210] that $\lim _{M \rightarrow+\infty} u^{T} D_{s}^{M}\left(\frac{\ell}{N}\right) u=u^{T} D_{s}(\bar{\rho}) u$.

### 6.2.2 Combined minimization problem

The collection of sets $C_{M, 0}, \ldots, C_{M, N}$ forms a partition of the set $\{0,1\}^{S_{M}}$. Observe for a given $\Psi \in H_{M}, A_{M, \ell}^{u}(\Psi)$ only depends on the values of $\Psi(\eta)$ for $\eta \in C_{M, \ell}$, since $\eta \in C_{M, \ell}$ implies that also $\eta^{\mathbf{0}, v_{k}} \in C_{M, \ell}$ and $\eta^{y+v_{k}, y} \in C_{M, \ell}$ for all $k \in\{1, \ldots, K\}$. As a consequence, if $\Psi_{\mathrm{opt}}^{M, u} \in H_{M}$ is a minimizer of

$$
\begin{equation*}
\min _{\Psi \in H_{M}} A_{M}^{u}(\Psi) \tag{6.5}
\end{equation*}
$$

where

$$
\begin{gather*}
A_{M}^{u}(\Psi):=\sum_{\eta \in\{0,1\}^{S_{M}}} \sum_{k=1}^{K} p_{k}\left[\left(1-\eta_{v_{k}}\right)\left(u \cdot v_{k}+\Psi\left(\eta^{\mathbf{0}, v_{k}}\right)-\Psi(\eta)\right)^{2}+\right. \\
\left.\frac{1}{2} \sum_{\substack{y \in S_{M} \\
y+v_{k} \neq \mathbf{0}}}\left(\Psi\left(\eta^{y+v_{k}, y}\right)-\Psi(\eta)\right)^{2}\right] \tag{6.6}
\end{gather*}
$$

then it holds that $A_{M, \ell}^{u}\left(\Psi_{\text {opt }}^{M, u}\right)=\min _{\Psi \in H_{M}} A_{M, \ell}^{u}(\Psi)$ for all $\ell \in\{0, \ldots, N\}$. The knowledge of $\Psi_{\mathrm{opt}}^{M, u}$ then allows us to compute $u^{T} D_{s}^{M}\left(\frac{\ell}{N}\right) u$ for all $\ell=0, \ldots, N$ as $2 A_{M, \ell}^{u}\left(\Psi_{\mathrm{opt}}^{M, u}\right)$. Note that the minimization problem (6.5) is independent of $\ell$, in contrast to (6.4). Also note that the minimization problem (6.5) has $2^{N}$ degrees of freedom and is thus intractable for large values of $N$.

Our main contribution is to propose the low-rank tensor approximation algorithm to compute an approximation of $\Psi_{\text {opt }}^{M, u}$ in order to mitigate this issue.

### 6.2.3 Estimation of long-time mean square deviation

The quantity $u^{T} D_{s}^{M}\left(\frac{\ell}{N}\right) u$ can be equivalently approximated by discretizing the definition (6.2). This leads to the problem of computing the long-time mean square deviation of a tagged particle evolving in the periodic environment $S_{M}[210]$. The Monte Carlo algorithm is the standard method of choice to compute an approximation of $u^{T} D_{s}^{M}\left(\frac{\ell}{N}\right) u$ by means of empirical averages. The quality of the obtained approximations then depends on the choice of two numerical parameters, namely the number of Monte Carlo samples (i.e. stochastic realizations of the periodized tagged particle process) and the value of the chosen final time.

We want to compare the low-rank tensor approximation algorithm we propose in this chapter with Monte Carlo estimations of the long-time mean square deviation. To this aim, we detail our Monte Carlo algorithm in the following. Let $\ell \in\{1, \ldots, N\}$. An initial environment $\eta$ is obtained by sampling uniformly from $C_{M, \ell}$. We sample from exponential distributions to determine the next time a particle in the environment or the tagged particle performs a jump. Whenever a jump is performed the environment $\eta$ is updated accordingly. Let $w=\mathbf{0} \in \mathbb{Z}^{d}$ denote the initial position of the tagged particle. Throughout the simulation, we track the position of the tagged particle in $\mathbb{Z}^{d}$, i.e. whenever the tagged particle successfully jumps in direction $v_{k}$, we set $w=w+v_{k}$. To approximate the expectation of $w$, we repeat this simulation $N_{s}$ times with different samples. Each of these simulations is stopped at the same stopping time $T$ at which we compute the approximation of $u^{T} D_{s}(\rho) u$. This approach is formalized in Algorithm 14.

Remark 6.3. In Algorithm 14, we set $N_{s}=\left\lceil\hat{N}_{s} /(\ell+1)\right\rceil$, where $\hat{N}_{s}$ is a given input parameter. This ensures that the expected runtime is approximately equal for all choices of $\ell$. Additionally Figure 6.5 shows that the variance of the output is comparable independently of $\ell$. If we were to use the same $N_{s}$ for all $\ell$, we would observe a much larger variance for smaller $\ell$ and the runtime would increase for larger $\ell$.

```
Algorithm 14 Long-term Monte Carlo estimation
    Input: \(M \in \mathbb{N}, u \in \mathbb{R}^{d}\), number of occupied sites \(1 \leq \ell \leq N\) with \(N=(2 M+1)^{d}-1\),
    final time \(T>0, \hat{N}_{s} \in \mathbb{N}\)
    Output: approximation \(\alpha \approx \frac{1}{2} u^{T} D_{s}^{M}\left(\frac{\ell}{N}\right) u=A_{M, \ell}^{u}\left(\Psi_{\mathrm{opt}}^{M, u}\right)\)
    \(w=\mathbf{0}, \alpha=0, N_{s}=\left\lceil\hat{N}_{s} /(\ell+1)\right\rceil\)
    for \(i=1,2,3, \ldots, N_{s}\)
        randomly initialize \(\eta \in C_{M, \ell}\) and set \(t_{\text {тотад }}=0\),
        while true
            sample \(t_{\text {NEW }}\) from an exponential distribution with mean \(\frac{1}{\ell+1}\)
            set \(t_{\text {TOTAL }}=t_{\text {TOTAL }}+t_{\text {NEW }}\)
            if \(t_{\text {TOTAL }}>T\)
                break
                uniformly select either one occupied site \(y \in S_{M}\) with \(\eta_{y}=1\) or the tagged
    particle \(y=\mathbf{0}\)
            select a jump direction \(v_{k}\) with probability \(p_{k}\)
            if \(\tilde{\eta}_{y+v_{k}}=\mathbf{0}\), i.e. the target location is not occupied
                    when \(y \neq \mathbf{0} \Rightarrow\) set \(\eta=\eta^{y+v_{k}, y}\)
                when \(y=\mathbf{0} \Rightarrow\) set \(\eta=\eta^{\mathbf{0}, v_{k}}\) and update \(w=w+v_{k}\)
        update \(\alpha=\alpha+\langle u, w\rangle^{2}\)
    \(\alpha=\frac{\alpha}{T N_{s}}\)
```


### 6.3 Low-rank solutions for the optimization problem

The aim of this section is to present our novel low-rank method for the numerical resolution of the high-dimensional optimization problem (6.5). Instead of solving the problem for all functions in $H_{M}$, we restrict the solution space to functions of rank at most $r$ defined analogous to (2.4), i.e. to functions of the form $\Psi=R^{(1)}+\cdots+R^{(r)}$, with

$$
R(\eta)^{(k)}=\Pi_{s \in S_{M}} R_{s}\left(\eta_{s}\right)^{(k)}, \quad \forall \eta=\left(\eta_{s}\right)_{s \in S_{M}} \in\{0,1\}^{S_{M}}
$$

for some $R_{s}^{(k)}:\{0,1\} \rightarrow \mathbb{R}$ for $s \in S_{M}, k=1, \ldots, r$. We denote by $H_{M}^{r} \subset H_{M}$ the set of all rank at most $r$ functions. For all $r \in \mathbb{N}$, it holds

$$
\min _{\Psi \in H_{M}} A_{M}^{u}(\Psi)=\min _{\Psi \in H_{M}^{2^{N}}} A_{M}^{u}(\Psi) \leq \min _{\Psi \in H_{M}^{r+1}} A_{M}^{u}(\Psi) \leq \min _{\Psi \in H_{M}^{r}} A_{M}^{u}(\Psi)
$$

In the next sections, we derive an algorithm to approximate $\arg \min _{\Psi \in H_{M}^{r}} A_{M}^{u}(\Psi)$. We first introduce a fast and stable algorithm for the evaluation of $A_{M}^{u}$ for functions in $H_{M}^{r}$. We then develop a successive minimization scheme to compute low-rank solutions of (6.5). Each minimization step is performed using an alternating linear scheme, which relies on the fast and stable evaluations of $A_{M}^{u}$. Lastly, we discuss the evaluation of $A_{M, \ell}^{u}$ for the computation of the self-diffusion coefficient (6.4).

### 6.3.1 Fast and stable evaluation

In this section, we introduce a fast and stable method to evaluate $A_{M}^{u}(\Psi)$ for $\Psi \in H_{M}^{r}$. The naive direct evaluation would require to sum over $2^{N}$ terms, which is intractable for large values of $N$. In principle, the order of summation can be exchanged leading to terms of the form $\sum_{\eta \in\{0,1\}^{s_{M}}} \Psi(\eta)^{2}$, which can be evaluated efficiently for $\Psi \in H_{M}^{r}$. However, subtracting these terms leads to a lot of numerical cancellation for $M>1$. We circumvent these issues by treating the evaluation of $A_{M}^{u}(\Psi)$ as the computation of the Frobenius norms of certain tensors. These Frobenius norms can then be evaluated in a fast and stable way.

## Reformulation as tensor norm

For $\Psi \in H_{M}$, we consider the order $N$ tensor $\mathcal{T}^{(\Psi)} \in \mathbb{R}^{2 \times \cdots \times 2}$ with entries $\mathcal{T}_{i_{1}, \ldots, i_{N}}^{(\Psi)}=\Psi(\eta)$, where $\eta=\left(\eta_{s_{1}}, \ldots, \eta_{s_{N}}\right)$ for some enumeration $s_{1}, \ldots, s_{N}$ of the sites in $S_{M}$ and $\eta_{s_{j}}=i_{j}-1$ for $1 \leq j \leq N$. Note that $\Psi \in H_{M}^{r}$, as defined in Section 6.3, can be represented by a rank- $r$ tensor in CP-format (2.4) of the form

$$
\begin{equation*}
\mathcal{T}_{i_{1}, \ldots, i_{N}}^{(\Psi)}=\sum_{k=1}^{r} \prod_{j=1}^{N} a_{i_{j}}^{(j, k)}, \tag{6.7}
\end{equation*}
$$

where $a^{(j, k)} \in \mathbb{R}^{2}$ is defined as $a^{(j, k)}=\left(R_{s_{j}}^{(k)}(0), R_{s_{j}}^{(k)}(1)\right)$.
For $v, y \in S_{M}$, we analogously define the order $N$ tensor $\mathcal{T}^{\left(\Psi^{(0, v)}\right)} \in \mathbb{R}^{2 \times \cdots \times 2}$ with entries $\mathcal{T}_{i_{1}, \ldots,,_{N}}^{\left(\Psi^{\mathbf{0}, v)}\right)}=\Psi\left(\eta^{\mathbf{0}, v}\right)$. When $y+v \neq \mathbf{0}$ we additionally define the order $N$ tensor $\mathcal{T}^{\left(\Psi^{(y+v, y)}\right)} \in \mathbb{R}^{2 \times \cdots \times 2}$ with entries $\mathcal{T}_{i_{1}, \ldots, i_{N}}^{\left(\Psi^{(y+v, y)}\right)}=\Psi\left(\eta^{y+v, y}\right)$. We observe that these again are rank- $r$ tensors when $\Psi \in H_{M}^{r}$. Lastly, we define the order $N$, rank-1 tensor $\mathcal{T}^{(u \cdot v)} \in$ $\mathbb{R}^{2 \times \cdots \times 2}$ with entries $\mathcal{T}_{i_{1}, \ldots, i_{N}}^{(u \cdot v)}=u \cdot v$.

For an order $N$ tensor $\mathcal{T} \in \mathbb{R}^{2 \times \cdots \times 2}$ and a site $s$ we define the projection operators $\mathcal{P}_{s}: \mathbb{R}^{2 \times \cdots \times 2} \rightarrow \mathbb{R}^{2 \times \cdots \times 2}$ as

$$
\left(\mathcal{P}_{s}(\mathcal{T})\right)_{i_{1}, \ldots, i_{N}}:=\left\{\begin{array}{ll}
\mathcal{T}_{i_{1}, \ldots, i_{N}} & i_{s}=1 \\
0 & \text { otherwise }
\end{array}, \quad \text { for } 1 \leq i_{1}, \ldots, i_{N} \leq 2,\right.
$$

where $i_{s}$ denotes to the index assigned to site $s$. We can now rewrite (6.6) as

$$
\begin{gather*}
A_{M}^{u}(\Psi)=\sum_{k=1}^{K} p_{k}\left[\left\|P_{v_{k}}\left(\mathcal{T}^{\left(u \cdot v_{k}\right)}+\mathcal{T}^{\left(\Psi^{\left(0, v_{k}\right)}\right)}-\mathcal{T}^{(\Psi)}\right)\right\|_{F}^{2}+\right. \\
\left.\frac{1}{2} \sum_{\substack{y \in S_{M} \\
y+v_{k} \neq 0}}\left\|\mathcal{T}^{(\Psi)}-\mathcal{T}^{\left(\Psi^{\left(y+v_{k}, y\right)}\right)}\right\|_{F}^{2}\right] \tag{6.8}
\end{gather*}
$$

## Efficient evaluation of Frobenius norms for sums of low-rank tensors

Let $k \in\{1, \ldots, K\}$ and $y \in S_{M}$ such that $y+v_{k} \neq \mathbf{0}$ be fixed. We derive an algorithm inspired by TT orthogonalization [249] to evaluate $\left\|\mathcal{T}^{(\Psi)}-\mathcal{T}^{\left(\Psi^{\left(y+v_{k}, y\right)}\right)}\right\|_{F}^{2}$.

We start by rewriting Equation (6.7) in TT format (2.6) as

$$
\begin{equation*}
T_{i_{1}, \ldots, i_{N}}^{(\Psi)}=\sum_{k_{1}=1}^{r} \ldots \sum_{k_{N-1}=1}^{r} \mathcal{C}_{1, i_{1}, k_{1}}^{1} \mathcal{C}_{k_{1}, i_{2}, k_{2}}^{2} \mathcal{C}_{k_{2}, i_{3}, k_{3}}^{3} \cdots \mathcal{C}_{k_{N-1}, i_{N}, 1}^{N} \tag{6.9}
\end{equation*}
$$

with tensors $\mathcal{C}^{1} \in \mathbb{R}^{1 \times 2 \times r}, \mathcal{C}^{N} \in \mathbb{R}^{r \times 2 \times 1}$ and $\mathcal{C}^{j} \in \mathbb{R}^{r \times 2 \times r}$ for $1<j<N$, whose entries are given by

$$
\begin{aligned}
\mathcal{C}_{k_{1}, i, k_{2}}^{j} & =\left\{\begin{array}{ll}
a_{i}^{j, k_{1}} & k_{1}=k_{2} \\
0 & \text { otherwise }
\end{array},\right. & \text { for } 1<j<N, 1 \leq i \leq 2,1 \leq k_{1}, k_{2} \leq N \\
\mathcal{C}_{1, i, k}^{1} & =a_{i}^{1, k} & \text { for } 1 \leq i \leq 2,1 \leq k \leq N \\
\mathcal{C}_{k, i, 1}^{N} & =a_{i}^{N, k} & \text { for } 1 \leq i \leq 2,1 \leq k \leq N
\end{aligned}
$$

The representation (6.9) can be generalized to other low-rank tensors. Let $\mathcal{T}^{\left(\Psi^{\left(y+v_{k}, y\right)}\right)}$ be analogously represented by tensors $\mathcal{D}^{j}$. The tensor $\mathcal{T}^{(\Psi)}-\mathcal{T}^{\left(\Psi^{\left(y+v_{k}, y\right)}\right)}$ is at most rank- $2 r$. It's representation in the form of (2.6) with tensors $\mathcal{E}^{j}$ can be constructed from the tensors $\mathcal{C}^{j}, \mathcal{D}^{j}$, i.e. the tensors $\mathcal{E}^{1} \in \mathbb{R}^{1 \times 2 \times 2 r}, \mathcal{E}^{N} \in \mathbb{R}^{2 r \times 2 \times 1}$ and $\mathcal{E}^{j} \in \mathbb{R}^{2 r \times 2 \times 2 r}$ for $1<j<N$ are defined as $\mathcal{E}_{1,,, 1: r}^{1}=\mathcal{C}^{1}, \mathcal{E}_{1,, r+1: 2 r}^{1}=-\mathcal{D}^{1}, \mathcal{E}_{1: r,:, 1}^{N}=\mathcal{C}^{N}, \mathcal{E}_{r+1: 2 r, ;, 1}^{N}=\mathcal{D}^{N}$ and $\mathcal{E}_{1: r, ;, 1: r}^{j}=\mathcal{C}^{j}, \mathcal{E}_{r+1: 2 r, ; r+1: 2 r}^{j}=\mathcal{D}^{j}$. Note that for $1<j<N$ these tensors are sparse by construction.

We can compute the norm $\left\|\mathcal{T}^{(\Psi)}-\mathcal{T}^{\left(\Psi^{\left(y+v_{k}, y\right)}\right)}\right\|_{F}^{2}$ directly form the tensors $\mathcal{E}^{j}$ in a fast and stable way using TT orthogonalization [249] as defined in Algorithm 15. We want to emphasize that the sparsity structure of $\mathcal{E}^{j}$ is preserved in Algorithm 15. The evaluation of $\left\|\mathcal{E}^{N}\right\|_{F}^{2}$ in line 13 requires summing up $4 r$ terms, whereas a naive evaluation of $\left\|\mathcal{T}^{(\Psi)}-\mathcal{T}^{\left(\Psi^{\left(y+v_{k}, y\right)}\right)}\right\|_{F}^{2}$ involves $2^{N}$ terms.

Remark 6.4. The sum of several low-rank tensors can be represented in the format (2.6). The representation of the sum can again be constructed from the representations of the individual low-rank tensors. This allows us to apply TT orthogonalization to evaluate all Frobenius-norms in Equation (6.8) in a fast and stable way.

### 6.3.2 Successive minimization

Let $r \in \mathbb{N}$. In order to approximate $\underset{\Psi \in H_{M}^{r}}{\arg \min } A_{M}^{u}(\Psi)$, we decompose the problem into a sequence of successive minimization problems following the idea of Remark 2.3. We assume that $\Psi \in H_{M}^{r}$ is of the form $\Psi=R^{(1)}+\cdots+R^{(r)}$ with rank-1 functions $R^{(k)} \in H_{M}^{1}$.

```
Algorithm 15 Frobenius norm evaluation
    Input: tensors \(\mathcal{E}^{1} \in \mathbb{R}^{1 \times 2 \times 2 r}, \mathcal{E}^{N} \in \mathbb{R}^{2 r \times 2 \times 1}\) and \(\mathcal{E}^{j} \in \mathbb{R}^{2 r \times 2 \times 2 r}\) for \(1<j<N\)
    Output: \(\|\mathcal{T}\|_{F}^{2}\), where \(\mathcal{T} \in \mathbb{R}^{2^{N}}\) has entries \(T_{i_{1}, \ldots, i_{N}}=\)
    \(\sum_{k_{1}=1}^{r} \cdots \sum_{k_{N-1}=1}^{r} \mathcal{E}_{1, i_{1}, k_{1}}^{1} \mathcal{E}_{k_{1}, i_{2}, k_{2}}^{2} \cdots \mathcal{E}_{k_{N-1}, i_{N}, 1}^{N}\)
    Compute \(Q R\) decomposition of \(\mathcal{E}^{1}\) reshaped as element in \(\mathbb{R}^{2 \times 2 r}\).
    Set \(\mathcal{E}^{1}\) to \(Q\) reshaped as element in \(\mathbb{R}^{1 \times 2 \times 2 r}\).
    Update \(\mathcal{E}^{2}\) : reshape to \(\mathbb{R}^{2 r \times 2 \cdot 2 r}\), multiply with \(R\) from the left, reshape back to
    \(\mathbb{R}^{2 r \times 2 \times 2 r}\).
    for \(j=2, \ldots, N-1\)
        Compute \(Q R\) decomposition of \(\mathcal{E}^{j}\) reshaped as element in \(\mathbb{R}^{2 r \cdot 2 \times 2 r}\).
        Set \(\mathcal{E}^{j}\) to \(Q\) reshaped as element in \(\mathbb{R}^{2 r \times 2 \times 2 r}\).
        if \(j<N-1\)
            Update \(\mathcal{E}^{j+1}\) : reshape to \(\mathbb{R}^{2 r \times 2 \cdot 2 r}\), multiply with \(R\) from the left, reshape back
    to \(\mathbb{R}^{2 r \times 2 \times 2 r}\).
        else
            Update \(\mathcal{E}^{N}\) : reshape to \(\mathbb{R}^{2 r \times 2}\), multiply with \(R\) from the left, reshape back to
    \(\mathbb{R}^{2 r \times 2 \times 1}\).
    \(\|\mathcal{T}\|_{F}^{2}=\left\|\mathcal{E}^{N}\right\|_{F}^{2}\)
```

We first determine $R^{(1)}$ as solution of $\min _{R^{(1)} \in H_{M}^{1}} A_{M}^{u}\left(R^{(1)}\right)$. In a successive step $R^{(2)}$ is determined as solution of $\min _{R^{(2)} \in H_{M}^{1}} A_{M}^{u}\left(R^{(1)}+R^{(2)}\right)$. This is continued successively until $\Psi$ is determined. The idea is formalized in Algorithm 16.

```
Algorithm 16 Successive minimization
    Input: rank \(r\)
    Output: approximation \(\Psi \approx \arg \min _{\Psi \in H_{M}^{r}} A_{M}^{u}(\Psi)\)
    \(\Psi=0\)
    for \(k=1, \ldots, r\)
        \(R^{(k)}=\arg \min _{R \in H_{M}^{1}} A_{M}^{u}(\Psi+R)\)
        \(\Psi=\sum_{i=1}^{k} R^{(i)}\)
```


### 6.3.3 Alternating least squares

In Algorithm 16, we need to solve minimization problems of the form

$$
\begin{equation*}
\min _{R \in H_{M}^{1}} A_{M}^{u}(\Psi+R) \tag{6.10}
\end{equation*}
$$

for given $\Psi \in H_{M}^{r}$. In the following, we introduce an ALS algorithm (see Section 2.1.2) to solve such minimization problems. The main idea is to approximate the solution of (6.10) by an iterative scheme which amounts to solving a sequence of small-dimensional linear problems. We start from an initial $R(\eta):=\Pi_{s \in S_{M}} R_{s}\left(\eta_{s}\right)$. We first minimize $A^{u}(\Psi+R)$ only with respect to a selected $R_{s_{0}}:\{0,1\} \rightarrow \mathbb{R}$ for some $s_{0} \in S_{M}$ leaving the other $R_{s}$,
$s \neq s_{0}$ fixed. By partially evaluating $A^{u}(\Psi+R)$ for all terms not depending on $R_{s_{0}}$, we obtain that $\min _{R_{s_{0}} \in\{\{0,1\} \rightarrow \mathbb{R}\}} A^{u}(\Psi+R)$ with $R(\eta):=R_{s_{0}}\left(\eta_{s_{0}}\right) \Pi_{s \in S_{M} \backslash\left\{s_{0}\right\}} R_{s}\left(\eta_{s}\right)$ is equivalent to a quadratic optimization problem

$$
\begin{equation*}
\min _{\left.R_{s_{0}} \in\{0,1\} \rightarrow \mathbb{R}\right\}} \alpha_{1} R_{s_{0}}(1)^{2}+\alpha_{2} R_{s_{0}}(0)^{2}+\alpha_{3} R_{s_{0}}(1) R_{s_{0}}(0)+\alpha_{4} R_{s_{0}}(1)+\alpha_{5} R_{s_{0}}(0)+\alpha_{6} \tag{6.11}
\end{equation*}
$$

with constants $\alpha_{1}, \ldots, \alpha_{6} \in \mathbb{R}$ depending on the fixed $R_{s}, s \neq s_{0}$ and $\Psi$. This quadratic optimization problem always admits a unique optimal $R_{s_{0}}$, which is given by $R_{s_{0}}(1)=a$ and $R_{s_{0}}(0)=b$, where $a, b \in R$ are the solution of the linear system

$$
\left(\begin{array}{cc}
2 \alpha_{1} & \alpha_{3}  \tag{6.12}\\
\alpha_{3} & 2 \alpha_{2}
\end{array}\right)\binom{a}{b}=\binom{-\alpha_{4}}{-\alpha_{5}}
$$

This allows us to optimize $A_{M}^{u}(\Psi+R)$ with respect to individual $R_{s_{0}}$. By alternating the selected $s_{0} \in S_{M}$, we obtain the alternating least squares algorithm, which is formalized in Algorithm 17.

```
Algorithm 17 Alternating least squares
    Input: initial functions \(R_{s}^{0}:\{0,1\} \rightarrow \mathbb{R}\) for \(s \in S_{M}\), function \(\Psi \in H_{M}^{r}\), vector \(u\),
    tolerance \(\varepsilon\)
    Output: approximation \(R_{\text {opt }}(\eta)=\Pi_{s \in S_{M}} R_{s}\left(\eta_{s}\right)\) of \(\arg \min _{R \in H_{M}^{1}} A_{M}^{u}(\Psi+R)\)
    \(v_{\text {old }}=\infty, v_{\text {new }}=A_{M}^{u}\left(\Psi+R^{0}\right)\) with \(R^{0}(\eta):=\Pi_{s \in S_{M}} R_{s}^{0}\left(\eta_{s}\right)\)
    \(\forall s \in S_{M}, R_{s}:=R_{s}^{0}\).
    while \(\left|v_{\text {old }}-v_{\text {new }}\right|>\varepsilon\left|v_{\text {new }}\right|\).
        \(v_{\text {old }}=v_{\text {new }}\)
        for \(s_{0} \in S_{M}\)
            \(R_{s_{0}}=\arg \min _{\widetilde{R}_{s_{0}}:\{0,1\} \rightarrow \mathbb{R}} A^{u}(\Psi+\widetilde{R})\) where \(\widetilde{R}(\eta)=\widetilde{R}_{s_{0}}\left(\eta_{s_{0}}\right) \Pi_{s \in S_{M} \backslash\left\{s_{0}\right\}} R_{s}\left(\eta_{s}\right)\)
    for all \(\eta=\left(\eta_{s}\right)_{s \in S_{M}}\)
        \(v_{\text {new }}=A_{M}^{u}(\Psi+R)\)
```

Remark 6.5. To compute the constants $\alpha_{i}$, we can either explicitly implement the partial evaluations of $A_{M}^{u}(\Psi+R)$. Alternatively, we can treat $A_{M}^{u}(\Psi+R)$ as a function in $\mathbb{R}^{2} \rightarrow \mathbb{R}$ depending on the values $R_{s_{0}}(0)$ and $R_{s_{0}}(1)$. We know that this function is a multivariatepolynomial of the form $\alpha_{1} R_{s_{0}}(1)^{2}+\alpha_{2} R_{s_{0}}(0)^{2}+\alpha_{3} R_{s_{0}}(1) R_{j}(0)+\alpha_{4} R_{s_{0}}(1)+\alpha_{5} R_{s_{0}}(0)+\alpha_{6}$. The constants can be computed using multivariate-polynomial interpolation (2.16) in six points. This interpolation has the advantages that it is non-intrusive and that the evaluations of $A_{M}^{u}(\Psi+R)$ can be performed efficiently using the ideas of Section 6.3.1.

Remark 6.6. Throughout this chapter, we use approximations in the CP-format. Note that alternating algorithms can also be applied to TT tensors [145, 171, 312].

### 6.3.4 Monte Carlo methods

Let $\Psi \in H_{M}^{r}$ denote an approximation of the solution of (6.5). In the following, we discuss how to evaluate $A_{M, \ell}^{u}(\Psi)$. In the definition (6.3) of $A_{M, \ell}^{u}$ the function

$$
\begin{array}{r}
f(\eta):=\sum_{k=1}^{K} p_{k}\left[\left(1-\eta_{v_{k}}\right)\left(u \cdot v_{k}+\Psi\left(\eta^{\mathbf{0}, v_{k}}\right)-\Psi(\eta)\right)^{2}+\right. \\
\left.\frac{1}{2} \sum_{\substack{y \in S \backslash\{\mathbf{0}\} \\
y+v_{k} \neq \mathbf{0}}}\left(\Psi\left(\eta^{y+v_{k}, y}\right)-\Psi(\eta)\right)^{2}\right] \tag{6.13}
\end{array}
$$

is evaluated for all $\eta \in C_{M, \ell}$. For larger $N$ and most choices of $\ell$, evaluating $\left|C_{M, \ell}\right|=\binom{N}{\ell}$ terms is intractable. We thus propose to use a Monte Carlo method [227] to approximate $A_{M, \ell}^{u}(\Psi)$.

In a naive Monte Carlo method, we compute $N_{s}$ samples $\eta^{(i)} \in C_{M, \ell}$ for $i=1, \ldots, N_{s}$ and replace $\frac{1}{\left|C_{M, \ell}\right|} \sum_{\eta \in C_{M, \ell}}$ in (6.3) by $\frac{1}{N_{s}} \sum_{\eta \in\left\{\eta^{(1)}, \ldots, \eta^{\left(N_{s}\right)}\right\}}$.
In Algorithm 18, we describe a Monte Carlo method with additional variance reduction [158], which as demonstrated in Section 6.4 reduces the number of samples needed to obtain a given approximation accuracy. The main idea is to observe that the sites $v_{1}, \ldots, v_{K}$ play a special role since they are the most relevant for jumps of the tagged particle. Instead of sampling the $\eta$ uniformly in $C_{M, \ell}$, we now ensure that all possible states of the sites $v_{1}, \ldots, v_{K}$ are occurring equally often in the set of sample points. We obtain the environments used to approximate $A_{M, \ell}^{u}(\Psi)$ using the following construction. We first sample the states of all sites other than $v_{1}, \ldots, v_{K}$ randomly. We then combine these states with all possible states of the sites $v_{1}, \ldots, v_{K}$ to obtain $2^{K}$ different environments. After evaluating $f$ for each of these environments, we obtain the next set of environments by randomly sampling the states of all sites other than $v_{1}, \ldots, v_{K}$. This procedure is repeated until desired the total number of evaluations of $f$ has been reached.

Remark 6.7. In line 9 of Algorithm 18, we sample form the set set

$$
\left\{\vec{w}^{(2)} \in\{0,1\}^{S_{M} \backslash\left\{v_{1}, \ldots, v_{k}\right\}} \mid\left\|\vec{w}^{(2)}\right\|_{1}=n_{2}\right\}
$$

which contains $\binom{N-K}{n_{2}}$ elements. When the number of elements in this set is smaller than $\widetilde{N}_{s}$, we can compute $\tilde{S}$ based on all possible $\vec{w}^{(2)}$ instead of sampling $\widetilde{N}_{s}$ times. This decreases the variance of the approximation further and reduces the number of required evaluations of $f$.

### 6.3.5 Limitations of the approach

In this section, we briefly discuss the limitations of the proposed algorithm. These are

```
Algorithm 18 Monte Carlo method with variance reduction
    Input: function \(f: C_{M, \ell} \rightarrow \mathbb{R}\), parameter \(\tilde{N}_{s}\)
    Output: approximation \(S\) of \(\frac{1}{\left|C_{M, \ell}\right|} \sum_{\eta \in C_{M, \ell}} f(\eta)\)
    \(S=0\)
    for \(\vec{w}^{(1)} \in\{0,1\}^{K}\)
        \(n_{1}:=\|\vec{w}\|_{1}, n_{2}:=\ell-n_{1}\)
        if \(0 \leq n_{2} \leq N-\ell\)
            \(\tilde{S}=0\)
                for \(i=1, \ldots, \tilde{N}_{s}\)
                    Sample \(\vec{w}^{(2)} \in\{0,1\}^{S_{M} \backslash\left\{v_{1}, \ldots, v_{k}\right\}}\) such that \(\left\|\vec{w}^{(2)}\right\|_{1}=n_{2}\).
                    Construct \(\eta \in C_{M, \ell}\) such that \(\eta(s)=\left\{\begin{array}{ll}\vec{w}_{k}^{(1)} & s=v_{k} \\ \vec{w}_{s}^{(2)} & \text { otherwise }\end{array}\right.\).
                    \(\tilde{S}=\tilde{S}+f(\eta)\)
        \(S=S+\binom{N-K}{n_{2}} \cdot \tilde{S} / \widetilde{N}_{s}\)
    \(S=\frac{S}{\binom{N}{\ell}}\)
```

primarily related to the following observation. Since $A_{M, \ell}^{u}\left(\Psi_{\text {opt }}^{M, u}\right)$ defines the entries of the self-diffusion coefficient (6.4), we know that $A_{M, \ell}^{u}\left(\Psi_{\text {opt }}^{M, u}\right) \in[0,1]$. Note that $A_{M, \ell}^{u}$ contains the normalization constant $\left|C_{M, \ell}\right|^{-1}$. The objective function $A_{M}^{u}$ (6.6) does not include a normalization factor. This implies that the $A_{M}^{u}\left(\Psi_{\mathrm{opt}}^{M, u}\right)$ is of order $2^{N}$. Trying to minimize this objective function numerically using floating point arithmetic leads to issues caused by rounding errors for large $N$.

For larger values of the objective function, it becomes increasingly challenging to solve the individual minimization problems in line 8 of Algorithm 17. In particular, our approach of using polynomial interpolation to find the location of the minimum might encounter numerical rounding issues. We find that the positive eigenvalues of the $2 \times 2$ matrix in (6.12) tend to be many orders of magnitude smaller than the norm of the right hand side. This implies that small rounding errors in the constants $\alpha_{1}, \ldots, \alpha_{6}$ might lead to vastly different solutions. For $M>2$, we even find that rounding errors can lead to negative eigenvalues in the system, i.e. we can not find the minimum of the polynomial (6.11) at all. It might be necessary to use a different approach to solve the individual ALS minimization problems for larger $M$. Alternatively, one could develop an algorithm to minimize the function $\log \left(A_{M}^{u}\right)$.

Moreover, the number of ALS iterations needed to find a good rank-1 minimizer tends to increase when increasing the size of the domain. This is especially true when a random initialization is used in Algorithm 17. Note that it might be possible to circumvent this issue by initializing Algorithm 17 based on the solution computed for a smaller value of $M$. Overcoming these limitations will then require further investigation. A possible path could be to combine domain decomposition approaches together with the tensor-based
optimization algorithm we propose here, in order to obtain a global optimization procedure where only independent parallel local optimization problems on medium-sized cells are solved.

### 6.4 Numerical Experiments

In the following numerical experiments*, we consider a tagged particle process with $K=4$ displacement vectors $v_{1}=(1,0), v_{2}=(-1,0), v_{3}=(0,1)$, and $v_{4}=(0,-1)$ and with associated probability $p_{k}=1 / 4$ for $k=1,2,3,4$. All computing times are measured without parallelization in MATLAB R2018b on a Lenovo Thinkpad T480s with Intel Core i7-8650U CPU and 15.4 GiB RAM. Note that both the sampling Algorithm 14 and Algorithm 18 as well as the evaluation of the different Frobenius-norms in Equation (6.8) can be parallelized to speed up computations.

### 6.4.1 Solving the optimization problem

In the following, we analyze the numerical approximation of the self-diffusion coefficient by solving the optimization problem (6.4).

Low-rank approximation error First, we study how the rank $r$ affects the value of $\min _{\Psi \in H_{M}^{r}} A_{M}^{u}(\Psi)$ for $u=(1,0)$. Let $\Psi_{\text {ALS }}^{r, u} \in H_{M}^{r}$ denote the function obtained by using Algorithm 16 with rank $r$, where the minimization problem in each iteration is solved using Algorithm 17. The initial functions $R_{s}^{0}$ in Algorithm 17 are selected randomly by assigning random values drawn form the uniform distribution on $[0,1]$ to $R_{s}^{0}(0)$ and $R_{s}^{0}(1)$. In Algorithm 17, we set $\varepsilon=10^{-12}$ and additionally stop after line 8 has been executed 420 times.

We depict the error $\left\|A_{1}^{u}\left(\Psi_{\mathrm{ALS}}^{r, u}\right)-A_{1}^{u}\left(\Psi_{\mathrm{ALS}}^{10, u}\right)\right\| /\left\|A_{1}^{u}\left(\Psi_{\mathrm{ALS}}^{10, u}\right)\right\|$, evaluated using the ideas of Section 6.3.1, of $\Psi_{\text {ALS }}^{r, u}$ for various $r$ in Figure 6.2. The computation with $r=10$ takes 4 minutes for $M=1$ and 82 minutes for $M=2$. We observe that the error decays quickly with increasing rank. In particular, for $M=1$ an approximation with $r=1$ is less than $0.1 \%$ away from a direct least squares solution of the minimization problem (6.5) (see Remark 6.8).

Remark 6.8. Given the tensor $\mathcal{T}^{(\Psi)}$ as defined in (6.7), we can equivalently write the evaluation of $A_{M}^{u}(\Psi)$ as sum of $Q=N K \cdot 2^{N-1}$ quadratic terms, which can be expressed as $\left\|B \operatorname{vec}\left(\mathcal{T}^{(\Psi)}\right)+b\right\|_{2}^{2}$ for some matrix $B \in \mathbb{R}^{Q \times 2^{N}}$ and vector $b \in \mathbb{R}^{Q}$. Thus, (6.5) boils down to solving the least squares problem

$$
\begin{equation*}
\min _{\Psi \in H_{M}}\left\|B \operatorname{vec}\left(\mathcal{T}^{(\Psi)}\right)+b\right\|_{2}^{2} . \tag{6.14}
\end{equation*}
$$

[^5]

Figure 6.2 - Algorithm 16 yields successive approximations $\Psi_{\text {ALS }}^{r, u} \in H_{M}^{r}$. For $r \in$ $\{1, \ldots, 9\}$, we plot the relative error of $\Psi_{\text {ALS }}^{k, u}$ compared to $\Psi_{\text {ALS }}^{10, u}$. We repeat this experiment 12 times with different random initial functions in Algorithm 17. The resulting relative errors are displayed in different colors. Left: $M=1$. Right: $M=2$.

In our setting of $d=2, M=1$ and $K=4$, it holds that $N=8$ so that $2^{N}=256$ and $Q=4096$. This is small enough to solve (6.14) up to a very high precision using lsqr in MATLAB. For larger values of $M$ it is no longer tractable to solve this least squares problem.

Sampling-based evaluation Let $M=2$ and $u=(1,0)$. For a given approximate solution $\Psi_{\text {ALS }}^{3, u}$, we want to evaluate $A_{2, \ell}^{u}\left(\Psi_{\text {ALS }}^{3, u}\right)$. A single evaluation of (6.13) with $\Psi$ replaced by $\Psi_{\text {ALS }}^{3, u}$ requires on average around $5.1 \cdot 10^{-4}$ seconds. Computing $A_{2, \ell}^{u}\left(\Psi_{\text {ALS }}^{3, u}\right)$ for $0 \leq \ell \leq N$ directly would require $2^{24} \approx 1.6 \cdot 10^{7}$ evaluations of (6.13), i.e. around 2.4 hours. In Section 6.3.4, we proposed two Monte Carlo algorithms to obtain approximations of $A_{2, \ell}^{u}\left(\Psi_{\text {ALS }}^{3, u}\right)$. We visualize the variance in the approximation obtained by these algorithms in Figure 6.3. The studied quantity $\frac{2}{N+1} \sum_{\ell=0}^{N} u^{T} D_{s}\left(\frac{\ell}{N}\right) u$ approximates $\int_{0}^{1} \operatorname{Tr}\left(D_{s}(\bar{\rho})\right) d \bar{\rho}$, where $\operatorname{Tr}$ denotes the trace operator. Algorithm 18 clearly leads to a variance reduction. With only $10^{5}$ samples, which can be evaluated in about one minute, we can already reach a variance of $10^{-6}$.

### 6.4.2 Estimation of long-time mean square deviation

In the following, we study estimating the long-term limit using Algorithm 14.
Figure 6.4 motivates our choice of $T=300$ and $\hat{N}_{s}=30000$ for the following numerical experiments. The error caused by stopping with $T=300$ is negligible compared to the stochastic variance when stopping with $\hat{N}_{s}=30000$.

In Figure 6.5 we analyze how the parameter $\hat{N}_{s}$ affects the approximation of the quantity


Figure 6.3 - For $M=2$ and $u=(1,0)$ we compute $\Psi_{\text {ALS }}^{3, u}$ using Algorithm 16. We then approximate $u^{T} D_{s}\left(\frac{\ell}{N}\right) u=2 A_{2, \ell}^{u}\left(\Psi_{\text {ALS }}^{3, u}\right)$ for $0 \leq \ell \leq N$ using sampling as in Section 6.3.4. We sample $\eta \in C_{M, \ell}$ using two different approaches; from the uniform distribution (naive Monte Carlo) and using Algorithm 18 (with variance reduction). We use a total of 460800 evaluations of (6.13) for both sampling methods. After every 3072 evaluations, we evaluate $\frac{2}{N+1} \sum_{\ell=0}^{N} u^{T} D_{s}\left(\frac{\ell}{N}\right) u$ based on the already computed samples. The whole experiment is repeated 250 times. Left: Evolution of $\frac{2}{N+1} \sum_{\ell=0}^{N} u^{T} D_{s}\left(\frac{\ell}{N}\right) u$ with increasing number of samples. The shaded areas mark one standard deviation from the mean for the respective algorithm. Right: Evolution of the variance of $\frac{2}{N+1} \sum_{\ell=0}^{N} u^{T} D_{s}\left(\frac{\ell}{N}\right) u$ with increasing number of samples.


Figure 6.4 - We run Algorithm 14 with $M=2, u=(1,0), \ell=0$ and different values for $T$ and $\hat{N}_{s}$. For each setting, we plot how the approximation of the value $u^{T} D_{s}(0) u$, which is known to be equal to 1 , evolves over time for 12 random initializations visualized in different colors. Left: $\hat{N}_{s}=10000, T=100000$. Right: $\hat{N}_{s}=100000, T=1000$.
$\frac{2}{N+1} \sum_{\ell=0}^{N} u^{T} D_{s}\left(\frac{\ell}{N}\right) u$. We observe that the quantity of interest converges to the same value around 0.84 as in Figure 6.3. Note that a single run with $\hat{N}_{s}=30000$ requires 33 minutes and reaches a variance of $10^{-4}$. Compared to the results in Figure 6.3 the computation time needed to achieve the same variance with Algorithm 14 compared to Algorithm 18 is much higher. This is visualized in Figure 6.6.



Figure 6.5 - We run Algorithm 14 with $T=300, u=(1,0)$ for $\ell=0, \ldots, 24$ to approximate $D_{s}\left(\frac{\ell}{N}\right)$. We use $\hat{N}_{s}=30000$ and store intermediate values for $\hat{N}_{s}=$ $100,200,300, \ldots$. Based on these approximations we compute $\frac{2}{N+1} \sum_{\ell=0}^{N} u^{T} D_{s}\left(\frac{\ell}{N}\right) u$. This procedure is repeated 50 times with different random initializations. Left: Evolution of $\frac{2}{N+1} \sum_{\ell=0}^{N} u^{T} D_{s}\left(\frac{\ell}{N}\right) u$ with increasing values of $\hat{N}_{s}$. The shaded area marks one standard deviation from the mean. Right: Evolution of the variance of $\frac{2}{N+1} \sum_{\ell=0}^{N} u^{T} D_{s}\left(\frac{\ell}{N}\right) u$ with increasing values of $\hat{N}_{s}$. We additionally display the evolution of the variance of $u^{T} D_{s}\left(\frac{\ell}{N}\right) u$ for individual values of $\ell$.


Figure 6.6 - We plot the variances displayed on the right of Figure 6.3 and Figure 6.5. Instead of the number of samples, we display the computation time for evaluating the corresponding number of samples on the x-axis. Left: Computation times only for Algorithm 14, the naive Monte Carlo method and Algorithm 18. Right: We additionally include the computation time of 736 seconds for computing $\Psi_{\text {ALS }}^{3, u}$ using Algorithm 16.

### 6.4.3 Comparison of algorithms

We now compare the two approaches of approximating the self-diffusion coefficient. We evaluate $u^{T} D_{s}(\ell / N) u$ for $u \in\{(1,0),(0,1),(1,1)\}$ to recover all entries of the symmetric matrix $D_{s}(\ell / N)$. Interpolation of the entries yields an approximation of $D_{s}(\rho)$ for $\rho \in[0,1]$. Alternatively, interpolation in the space of symmetric positive definite matrices
could be used [232].
For the following experiments, we use Algorithm 14 with a slight modification to approximate $u^{T} D_{s}(\bar{\rho}) u$ for different values of $u$ simultaneously. This is achieved by keeping track of different $\alpha$ for different $u$ in lines 16 and 17 . The optimization problem (6.5) needs to be solved for every $u$ separately, but the solution can be evaluated for different $\ell$. In contrast, the sampling in Algorithm 14 needs to be recomputed from scratch for every $\ell$.

In Figure 6.7, we display the trace of $D_{s}(\rho)$ computed using different approaches. Computing the whole matrix $D_{s}(\ell / N)$ for all $0 \leq \ell \leq N$ took 3 minutes (37 minutes) for the approach based on solving the minimization problem, whereas using Algorithm 14 took 15 minutes ( 45 minutes) for $M=1(M=2)$. The much higher variance of Algorithm 14 for approximating the long-term limit leads to clearly visible changes in the graph of trace at the values $\rho=\frac{\ell}{N}$. The figure also contains results using a $6 \times 6$ grid $(N=35)$ for which the minimization based approach took 5 hours and 40 and the sampling of the long-term limit took 9 hours and 24 minutes. The sampled solution of the minimization problem changes less due to a smaller sampling variance in Algorithm 18. The variance for both approaches is depicted in Figure 6.8 and listed in Table 6.1. We want to emphasize that the approach based on solving the minimization problem is faster and simultaneously yields a lower variance in the entries of the self-diffusion compared to the estimation of the long-term limit.


Figure 6.7 - For $M=1(N=8)$ and $M=2(N=24)$ we solve the minimization problem (6.5) using Algorithm 16 with $r=3$. We evaluate $A_{M, \ell}^{u}$ directly for $M=1$, whereas for $M=2$ we use Algorithm 18 with $\widetilde{N}_{s}=50$. We compare to Algorithm 14 with $\hat{N}_{s}=30000$ and $T=300$ for $M=1$ and $M=2$. Repeating this for $u \in$ $\{(1,0),(1,1),(0,1)\}$ yields $D_{s}(\ell / N)$. Element-wise linear interpolation allows us to evaluate $D_{s}(\rho)$ for $\rho \in[0,1]$. We plot the trace of $D_{s}(\rho)$. In addition, we run both algorithms on a $4 \times 4$ and $6 \times 6$ periodic grid ( $N=15$ and $N=35$ ). For $N=15$ we use the same parameters as for $N=24$. For $N=35$ we use rank $10, \widetilde{N}_{s}=150$ and $\hat{N}_{s}=80000, T=600$. Note that the $4 \times 4$ grid lies in between the $3 \times 3$ grid for $M=1$ and the $5 \times 5$ grid for $M=2$. Right: Zoom on part of the graphs.


Figure 6.8 - We repeat the experiment presented in Figure 6.7100 times for $N=8$, $N=15$ and $N=24$. We plot a zoom on a part of the graphs. The shaded areas mark one standard deviation from the mean for each of the methods.

|  |  | $\max _{\ell \in\{0, \ldots, N\}} \operatorname{Var}\left(\operatorname{Tr}\left(D_{s}(\ell / N)\right)\right)$ | $\frac{1}{N+1} \sum_{\ell=0}^{N} \operatorname{Var}\left(\operatorname{Tr}\left(D_{s}(\ell / N)\right)\right)$ |
| :---: | :---: | :---: | :---: |
| $\mathrm{N}=8$ | long-term limit | $2.313 \cdot 10^{-4}$ | $1.072 \cdot 10^{-4}$ |
|  | minimization | $8.791 \cdot 10^{-9}$ | $4.231 \cdot 10^{-9}$ |
| $\mathrm{~N}=15$ | long-term limit | $3.433 \cdot 10^{-4}$ | $1.410 \cdot 10^{-4}$ |
|  | minimization | $2.548 \cdot 10^{-4}$ | $6.835 \cdot 10^{-5}$ |
| $\mathrm{~N}=24$ | long-term limit | $4.377 \cdot 10^{-4}$ | $1.989 \cdot 10^{-4}$ |
|  | minimization | $3.262 \cdot 10^{-4}$ | $8.001 \cdot 10^{-5}$ |

Table 6.1 - We display the mean and maximum of the variance for the 100 approximations of the self-diffusion coefficient computed in Figure 6.8. Note that we do not need to sample to evaluate $A_{M, \ell}^{u}$ for $N=8$. The variance for minimization with $N=8$ is solely caused by the random initializations of $R_{s}^{0}$ in Algorithm 17.

### 6.5 Application: Cross-diffusion system

In this section, we recall how self-diffusion matrices arise in the context of cross-diffusion systems. We then use our novel low-rank approach described in Section 6.3 to approximate the self-diffusion matrices needed to resolve a slightly simplified cross-diffusion system with a cell-centered finite volume method.

### 6.5.1 Hydrodynamic limit of a lattice-based stochastic hopping model

Lattice-based stochastic hopping models $[214,266,301]$ describe the evolution of a mixture of multi-species particles, the positions of which are clamped on a (periodic) lattice $T_{n}:=$ $\left\{0, \ldots, \frac{n-1}{n}\right\}^{d}, d=1,2,3$. After a certain amount of time, a single particle jumps to a nearby node according to a given Markovian random walk. Let $v_{k}:=\left(\left(v_{k}\right)_{i}\right)_{1 \leq i \leq d} \in \mathbb{Z}^{d} \backslash\{0\}$ for $k=1, \ldots, K$ be the displacement vectors for all possible jumps that one particle can


Figure 6.9 - The evolution of state of the particles on a lattice with $n=5, d=2$. Left: Initial state. Middle: After the blue particle at $s=\left(\frac{1}{5}, \frac{3}{5}\right)$ jumped with $\mathbf{v}=(1,0)$ to $s+\frac{1}{5} \mathbf{v}$. Right: After the red particle at position $\left(\frac{4}{5}, \frac{1}{5}\right)$ jumped with $\mathbf{v}=(0,-2)$. Since the lattice is periodic the particle ends up in position $\left(\frac{4}{5}, \frac{4}{5}\right)$.
make, i.e. a particle at position $x \in T_{n}$ lands at position $x+\frac{1}{n} v_{k} \in T_{n}$. The rate of jumping in the direction $\frac{1}{n} v_{k}$ is given by $\left.\left.p_{k} \in\right] 0,1\right]$. We assume that the particles can be of two possible species, either species blue or species red. We study how the distribution of the particles belonging to each species evolves over time in the limit when the number of lattice points and particles tends to infinity. See Figure 6.9 for a schematic illustration of this lattice-hopping process.

We are interested in simulating the hydrodynamic limit of this lattice-based hopping model, which was identified in [266]. In the limit $n \rightarrow+\infty$, this hydrodynamic limit reads as a cross-diffusion system defined on the $d$-dimensional torus $T^{d}:=(\mathbb{R} / \mathbb{Z})^{d}$. At this scale, we consider densities $\rho_{\text {red }}: T^{d} \times \mathbb{R}_{+} \rightarrow[0,1]$ and $\rho_{\text {blue }}: T^{d} \times \mathbb{R}_{+} \rightarrow[0,1]$, where $\rho_{\text {red }}(x, t)$ (respectively $\rho_{\text {blue }}(x, t)$ ) denotes the local volumic fraction of red (respectively blue) particles at point $x \in T^{d}$ and time $t \geq 0$. The local volumic fractions $\rho_{\text {red }}$ and $\rho_{\text {blue }}$ are shown to be solutions to the following cross-diffusion system [266]

$$
\frac{\partial}{\partial t}\binom{\rho_{\text {red }}}{\rho_{\text {blue }}}=\frac{1}{2} \nabla \cdot\left[\left(\begin{array}{cc}
\frac{\rho_{\text {blue }}}{\rho} D_{s}(\rho)+\frac{\rho_{\text {red }}}{\rho} D & \frac{\rho_{\text {red }}}{\rho}\left(D-D_{s}(\rho)\right)  \tag{6.15}\\
\frac{\rho_{\text {blue }}}{\rho}\left(D-D_{s}(\rho)\right) & \frac{\rho_{\text {red }}}{\rho} D_{s}(\rho)+\frac{\rho_{\text {blue }}}{\rho} D
\end{array}\right)\binom{\nabla \rho_{\text {red }}}{\nabla \rho_{\text {blue }}}\right],
$$

in $T^{d} \times \mathbb{R}_{+}$, where $\rho:=\rho_{\text {red }}+\rho_{\text {blue }}$. The matrix $D \in \mathbb{R}^{d \times d}$ is defined such that $D_{i j}:=\sum_{k=1}^{K} p_{k}\left(v_{k}\right)_{i}\left(v_{k}\right)_{j}$ for all $i, j=1, \ldots, d$. The matrix $D_{s}(\rho)$ denotes the selfdiffusion matrix application (6.1) of the tagged particle process with jump directions $v_{k}$ and jump probabilities $p_{k}$ for $k=1, \ldots, K$.

We complement system (6.15) by the initial conditions $\rho_{\text {red }}(0, x)=\rho_{\text {red }}^{0}(x), \rho_{\text {blue }}(0, x)=$ $\rho_{\text {blue }}^{0}(x)$, where $\rho_{\text {red }}^{0}, \rho_{\text {blue }}^{0} \in L^{\infty}\left(T^{d}\right)$ satisfy $0 \leq \rho_{\text {red }}^{0}+\rho_{\text {blue }}^{0} \leq 1$ almost everywhere in $T^{d}$. In this case, it can be shown [266] that there exists a unique solution $\left(\rho_{\text {red }}, \rho_{\text {blue }}\right) \in$ $L^{\infty}\left(T^{d} \times \mathbb{R}_{+}\right)^{2}$ to system (6.15). In addition, it holds that for almost all $t \geq 0$ and $x \in T^{d}$, $0 \leq \rho_{\text {red }}(x, t), 0 \leq \rho_{\text {blue }}(x, t)$ and $\rho_{\text {red }}(x, t)+\rho_{\text {blue }}(x, t) \leq 1$.

### 6.5.2 Deterministic resolution of a simplified cross-diffusion system

Since, the design and numerical analysis of a numerical scheme for the resolution of the cross-diffusion system (6.15) is still subject to future research, we present a numerical scheme to resolve a simplified system. In this section, we develop a cell-centred finite volume method $[115,116]$ to simulate this simplified cross-diffusion system, assuming that a numerical approximation of the self-diffusion matrix $D_{s}(\bar{\rho})$ can be computed for any $\bar{\rho} \in[0,1]$.

## Simplified cross-diffusion system

Let $\Omega \subset \mathbb{R}^{d}$ be a polyhedric bounded domain of $\mathbb{R}^{d}$. The local volumic fractions of red and blue particles are now given by functions $\rho_{\text {red }}: \Omega \times \mathbb{R}_{+} \rightarrow[0,1]$ and $\rho_{\text {blue }}: \Omega \times \mathbb{R}_{+} \rightarrow[0,1]$. These functions are assumed to satisfy the following simplified cross-diffusion system:

$$
\begin{align*}
& \frac{\partial}{\partial t}\binom{\rho_{\text {red }}}{\rho_{\text {blue }}}= \\
& \frac{1}{2} \nabla \cdot\left[\left(\begin{array}{cc}
\operatorname{Tr}\left[\frac{\rho_{\text {blue }}}{\rho} D_{s}(\rho)+\frac{\rho_{\text {red }}}{\rho} D\right] & \operatorname{Tr}\left[\frac{\rho_{\text {red }}}{\rho}\left(D-D_{s}(\rho)\right)\right] \\
\operatorname{Tr}\left[\frac{\rho_{\text {blue }}}{\rho}\left(D-D_{s}(\rho)\right)\right] & \operatorname{Tr}\left[\frac{\rho_{\text {red }}}{\rho} D_{s}(\rho)+\frac{\rho_{\text {blue }}}{\rho} D\right]
\end{array}\right)\binom{\nabla \rho_{\text {red }}}{\nabla \rho_{\text {blue }}}\right], \tag{6.16}
\end{align*}
$$

in $\Omega \times \mathbb{R}_{+}$, where $\rho:=\rho_{\text {red }}+\rho_{\text {blue }}$. We complement system (6.15) by the initial conditions $\rho_{\text {red }}(0, x)=\rho_{\text {red }}^{0}(x)$, $\rho_{\text {blue }}(0, x)=\rho_{\text {blue }}^{0}(x)$, where $\rho_{\text {red }}^{0}, \rho_{\text {blue }}^{0} \in L^{\infty}(\Omega)$ satisfying $0 \leq \rho_{\text {red }}^{0}+\rho_{\text {blue }}^{0} \leq 1$ almost everywhere in $\Omega$. We also assume that system (6.16) is complemented with Neumann boundary conditions instead of periodic boundary conditions.

## Time discretization

Let $T_{f}>0$ denote some final time. For the time discretization, we introduce a division of the interval $\left[0, T_{f}\right]$ into subintervals $I_{p}:=\left[t_{p-1}, t_{p}\right], 1 \leq p \leq P_{t}$ for some $P_{t} \in \mathbb{N}$ such that $0=t_{0}<t_{1}<\cdots<t_{P_{t}}=T_{f}$. The time steps are denoted by $\Delta t_{p}=t_{p}-t_{p-1}$, $p=1, \ldots, P_{t}$. For a given sequence of real numbers $\left(v^{p}\right)_{p \in \mathbb{N}}$, we define the approximation of the first-order time derivative thanks to the backward Euler scheme as follows:

$$
\partial_{t} v^{p}:=\frac{v^{p}-v^{p-1}}{\Delta t_{p}} \quad \forall 1 \leq p \leq P_{t}
$$

## Space discretization

For the space discretization, we consider a conforming simplicial mesh $\mathcal{T}_{h}$ of the domain $\Omega$, i.e. $\mathcal{T}_{h}$ is a set of elements $K$ verifying $\bigcup_{K \in \mathcal{T}_{h}} \bar{K}=\bar{\Omega}$, where the intersection of the closure of two elements of $\mathcal{T}_{h}$ is either an empty set, a vertex, or an $l$-dimensional face, $0 \leq l \leq d-1$. Denote by $h_{K}$ the diameter of the generic element $K \in \mathcal{T}_{h}$ and $h=\max _{K \in \mathcal{T}_{h}} h_{K}$. We denote by $\mathcal{E}_{h}$ the set of mesh faces. Boundary faces are collected in the set $\mathcal{E}_{h}^{\text {ext }}=\left\{\sigma \in \mathcal{E}_{h} ; \sigma \subset \partial \Omega\right\}$ and internal faces are collected in the set $\mathcal{E}_{h}^{\text {int }}=\mathcal{E}_{h} \backslash \mathcal{E}_{h}^{\text {ext }}$. To each face $\sigma \in \mathcal{E}_{h}$, we associate a unit normal vector $n_{\sigma}$; for $\sigma \in \mathcal{E}_{h}^{\text {int }}, \sigma=K \cap L, n_{\sigma}$ points from $K$ towards $L$ and for $\sigma \in \mathcal{E}_{h}^{\text {ext }}$ it coincides with the outward unit normal vector $n_{\Omega}$ of $\Omega$. We also denote by $N_{\mathrm{e}}$ the number of elements in the mesh $\mathcal{T}_{h}$. Furthermore, the notation $n_{K, \sigma}$ stands for the outward unit normal vector to the element $K$ on $\sigma$. We also assume that the family $\mathcal{T}_{h}$ is superadmissible in the sense that for all cells $K \in \mathcal{T}_{h}$ there exists a point $x_{K} \in K$ (the cell center) and for all edges $\sigma \in \mathcal{E}_{h}$ there exists a point $x_{\sigma} \in \partial K$ (the edge center) such that, for all edges $\sigma \in \mathcal{E}_{K}$, the line segment joining $x_{K}$ with $x_{\sigma}$ is orthogonal to $\sigma$ (see [116]). For an interior edge $\sigma \in \mathcal{E}_{h}^{\text {int }}$ shared by two elements $K$ and $L$ (denoted in the sequel by $\sigma=K \cap L$ ), we define the distance between these elements $d_{K L}:=\operatorname{dist}\left(x_{K}, x_{L}\right)$. Figure 6.10 provides a schematic illustration.


Figure 6.10 - Illustration of the notation for a mesh with two elements.

## The cell-centered finite volume method

In the context of the cell-centered finite volume method, the unknowns of the model are discretized using one value per cell: $\forall 1 \leq p \leq P_{t}$ we let

$$
U^{p}:=\left(U_{K}^{p}\right)_{K \in \mathcal{T}_{h}} \in \mathbb{R}^{2 N_{\mathrm{e}}}, \quad \text { with } \quad U_{K}^{p}:=\left(\rho_{\mathrm{red}, K}^{p}, \rho_{\mathrm{blue}, K}^{p}\right) \in \mathbb{R}^{2}
$$

where $\rho_{\text {red }, K}^{p}$ and $\rho_{\mathrm{blue}, K}^{p}$ are respectively the discrete elementwise unknowns approximating the values of $\rho_{\text {red }}\left(t_{p}\right)$ and $\rho_{\text {blue }}\left(t_{p}\right)$ in the element $K \in \mathcal{T}_{h}$. More precisely, $\rho_{\text {red }, K}^{p} \approx \frac{1}{|K|} \int_{K} \rho_{\text {red }}\left(t_{p}, x\right) \mathrm{d} x$ and $\rho_{\text {blue }, K}^{p} \approx \frac{1}{|K|} \int_{K} \rho_{\text {blue }}\left(t_{p}, x\right) \mathrm{d} x$. Integrating (6.15) over the element $K \in \mathcal{T}_{h}$, using next the Green formula and the Neumann boundary conditions, the cell-centered finite volume scheme we solve is the following: for a given
fixed value of $U^{p-1}$, find $U^{p} \in \mathbb{R}^{2 N_{\mathrm{e}}}$ such that

$$
\begin{align*}
& H_{1, K}\left(U^{p}\right)=|K| \partial_{t} \rho_{\mathrm{red}, K}^{p}-\frac{1}{2} \sum_{\sigma \in \mathcal{E}_{K}} \mathcal{F}_{1, K, \sigma}\left(U^{p}\right)=0 \quad \forall K \in \mathcal{T}_{h}  \tag{6.17}\\
& H_{2, K}\left(U^{p}\right)=|K| \partial_{t} \rho_{\mathrm{blue}, K}^{p}-\frac{1}{2} \sum_{\sigma \in \mathcal{E}_{K}} \mathcal{F}_{2, K, \sigma}\left(U^{p}\right)=0 \quad \forall K \in \mathcal{T}_{h}
\end{align*}
$$

The numerical fluxes $\mathcal{F}_{1, K, \sigma}\left(U^{p}\right)$ and $\mathcal{F}_{2, K, \sigma}\left(U^{p}\right)$ are respectively approximations of

$$
\begin{aligned}
& \mathcal{F}_{1, K, \sigma}\left(U^{p}\right) \approx\left(S^{11}\left(U^{p}\right) \nabla \rho_{\text {red }}^{p}+S^{12}\left(U^{p}\right) \nabla \rho_{\mathrm{blue}}^{p}\right) \cdot n_{K, \sigma} \\
& \mathcal{F}_{2, K, \sigma}\left(U^{p}\right) \approx\left(S^{21}\left(U^{p}\right) \nabla \rho_{\mathrm{red}}^{p}+S^{22}\left(U^{p}\right) \nabla \rho_{\mathrm{blue}}^{p}\right) \cdot n_{K, \sigma},
\end{aligned}
$$

where $S^{11}\left(U^{p}\right):=\left(\mathcal{S}_{K}^{11}\left(U^{p}\right)\right)_{K \in \mathcal{T}_{h}}, S^{12}\left(U^{p}\right):=\left(\mathcal{S}_{K}^{12}\left(U^{p}\right)\right)_{K \in \mathcal{T}_{h}}, S^{21}\left(U^{p}\right):=\left(\mathcal{S}_{K}^{21}\left(U^{p}\right)\right)_{K \in \mathcal{T}_{h}}$, $S^{22}\left(U^{p}\right):=\left(\mathcal{S}_{K}^{22}\left(U^{p}\right)\right)_{K \in \mathcal{T}_{h}}$ are collections of real numbers defined on each cell of the mesh as follows. For all $K \in \mathcal{T}_{h}$,

$$
\begin{aligned}
& \mathcal{S}_{K}^{11}\left(U^{p}\right):=\operatorname{Tr}\left[\frac{\rho_{\mathrm{blue}, K}^{p}}{\rho_{K}^{p}} D_{s}\left(\rho_{K}^{p}\right)+\frac{\rho_{\mathrm{red}, K}^{p}}{\rho_{K}^{p}} D\right] \\
& \mathcal{S}_{K}^{12}\left(U^{p}\right):=\operatorname{Tr}\left[\frac{\rho_{\mathrm{red}, K}^{p}}{\rho_{K}^{p}}\left(D-D_{s}\left(\rho_{K}^{p}\right)\right)\right] \\
& \mathcal{S}_{K}^{21}\left(U^{p}\right):=\operatorname{Tr}\left[\frac{\rho_{\mathrm{blue}, K}^{p}}{\rho_{K}^{p}}\left(D-D_{s}\left(\rho_{K}^{p}\right)\right)\right] \\
& \mathcal{S}_{K}^{22}\left(U^{p}\right):=\operatorname{Tr}\left[\frac{\rho_{\mathrm{red}, K}^{p}}{\rho_{K}^{p}} D_{s}\left(\rho_{K}^{p}\right)+\frac{\rho_{\mathrm{blue}, K}^{p}}{\rho_{K}^{p}} D\right]
\end{aligned}
$$

where $\rho_{K}^{p}:=\rho_{\text {red, } K}^{p}+\rho_{\mathrm{blue}, K}^{p}$. Now we are interested in computing an approximation of $\left(S^{i j}\left(U^{p}\right) \nabla \rho_{\sharp}^{p}\right) \cdot n_{K, \sigma}$ for all $1 \leq i, j \leq 2$ and $\sharp \in\{$ red, blue $\}$.

For any edge $\sigma \in \mathcal{E}_{h}^{\text {int }}$, denoting by $K, L \in \mathcal{T}_{h}$ the two cells sharing the edge $\sigma$, we employ the following harmonic averaging formula for all $1 \leq i, j \leq 2$,

$$
\left(S^{i j}\left(U^{p}\right) \nabla \rho_{\sharp}^{p}\right) \cdot n_{K, \sigma} \approx|\sigma| \frac{\mathcal{S}_{K}^{i j}\left(U^{p}\right) \mathcal{S}_{L}^{i j}\left(U^{p}\right)}{d_{K, \sigma} \mathcal{S}_{L}^{i j}\left(U^{p}\right)+d_{L, \sigma} \mathcal{S}_{K}^{i j}\left(U^{p}\right)}\left(\rho_{\sharp, L}^{p}-\rho_{\sharp, K}^{p}\right) .
$$

We refer to $[1,113,114]$ for more details. This averaging formula enables then to obtain the expression of the fluxes $\mathcal{F}_{1, K, \sigma}\left(U^{p}\right)$ and $\mathcal{F}_{2, K, \sigma}\left(U^{p}\right)$.

At each time step $p$, we thus have to solve the nonlinear problem: find $U^{p} \in \mathbb{R}^{2 N_{\mathrm{e}}}$ such that $\mathcal{H}^{p}\left(U^{p}\right)=0$, where $\mathcal{H}^{p}\left(U^{p}\right):=\left(H_{1, K}\left(U^{p}\right), H_{2, K}\left(U^{p}\right)\right)_{K \in \mathcal{T}_{h}} \in \mathbb{R}^{2 N_{\mathrm{e}}}$, where $H_{1, K}\left(U^{p}\right)$ and $H_{2, K}\left(U^{p}\right)$ are defined by (6.17).

## Resolution of the nonlinear problem by a Newton method

In this section, we present a Newton procedure for computing the approximate solution $U^{p}$ at each time step $p$. For $1 \leq p \leq P_{t}$ and $U^{p, 0} \in \mathbb{R}^{2 N_{e}}$ fixed (typically $U^{p, 0}=U^{p-1}$ ), the Newton algorithm generates a sequence $\left(U^{p, k}\right)_{k \geq 1}$, with $U^{p, k} \in \mathbb{R}^{2 N_{e}}$ given by the system of linear algebraic equations:

$$
\begin{equation*}
A^{p, k-1} U^{p, k}=B^{p, k-1} \tag{6.18}
\end{equation*}
$$

Here, the Jacobian matrix $A^{p, k-1} \in \mathbb{R}^{2 N_{e}, 2 N_{e}}$ and the right-hand side vector $B^{p, k-1} \in \mathbb{R}^{2 N_{e}}$ are defined by

$$
A^{p, k-1}:=J_{\mathcal{H}^{p}}\left(U^{p, k-1}\right) \quad \text { and } \quad B^{p, k-1}:=J_{\mathcal{H}^{p}}\left(U^{p, k-1}\right) U^{p, k-1}-\mathcal{H}^{p}\left(U^{p, k-1}\right) .
$$

Note that $J_{\mathcal{H}^{p}}\left(U^{p, k-1}\right)$ is the Jacobian matrix of the function $\mathcal{H}^{p}$ at point $U^{p, k-1}$. A classical stopping criterion for system (6.18) is for instance

$$
\begin{equation*}
\left\|\mathcal{H}^{p}\left(U^{p, k}\right)\right\|_{2} /\left\|\mathcal{H}^{p}\left(U^{p, 0}\right)\right\|_{2}<\varepsilon_{\operatorname{lin}} \tag{6.19}
\end{equation*}
$$

where $\varepsilon_{\text {lin }}>0$ is small enough. We refer to the book [185] for large descriptions on linearization techniques.

### 6.5.3 Numerical Experiment

In the following numerical experiment ${ }^{\dagger}$, we solve the PDE system (6.16) using the finite volume scheme described in this section. A approximation of the self-diffusion coefficient $D_{s}(\bar{\rho})$ is obtained from $\Psi_{\text {ALS }}^{1, u}$ as in Section 6.4. We set the domain $\Omega:=[0,1] \times[0,1]$ and consider a uniform spatial mesh ( $N_{\mathrm{e}}=312$ elements). We use a constant time step $\Delta t_{p}=\Delta t=10^{-3}$ seconds $\forall 1 \leq p \leq P_{t}$. The final time of simulation is $T_{f}=10$ seconds. We employ at each time step $1 \leq p \leq P_{t}$ a Newton solver as described in Section 6.5.2 with $\varepsilon_{\operatorname{lin}}=10^{-8}$. The initial values are defined by

$$
\rho_{\text {red }}^{0}(x, y):=0.25+0.25 \cos (\pi x) \cos (\pi y) \quad \text { and } \quad \rho_{\text {blue }}^{0}(x, y):=0.5-0.5 \cos (\pi x) \cos (\pi y) .
$$

Figure 6.11 displays the shape of the numerical solutions $\rho_{\text {red }}$ and $\rho_{\text {blue }}$ for several time steps. We observe that the local volumic fractions evolve over time to reach constant profiles (around 0.25 for $\rho_{\text {red }}$ and 0.5 for $\rho_{\text {blue }}$ ) in the long time limit. This indicates, that our scheme is stable in the sense that the densities lie in their physical ranges. In every time step, only 2 or 3 Newton iterations are required to reach the stopping criterion (6.19).

[^6]

Figure 6.11 - Solutions $\rho_{\text {red }}$ and $\rho_{\text {blue }}$ for several time steps. Top left $p=0$, top middle $p=10$, top right $p=30$, bottom left $p=50$, bottom middle $p=110$, bottom right $p=250$.

# 7 Fast global spectral method 

Global spectral methods [122, 143, 319, 326] offer the potential to solve PDEs to very high accuracy. These methods are based on studying the impact of the differential operator when applied to a finite set of basis functions such as tensorized Chebyshev polynomials with bounded degree (2.16). We then compute a linear combination of these basis functions that (approximately) satisfies the PDE. For one- and two-dimensional rectangular domains this approach led to the development of the solvers Chebop [102, 103, 246] and Chebop2 [319] contained in the Chebfun package [104]. Extensions to triangular domains and to disks have been derived in [244, 245, 338]. In the three-dimensional setting only specialised solvers for the Poissons's equation have been developed so far [125, 321, 346].

This, chapter is concerned with the numerical solution of general linear PDEs on cubes of the form

$$
\begin{equation*}
\mathcal{L} u=f \quad \text { on }[-1,1]^{3}, \tag{7.1}
\end{equation*}
$$

complemented with linear boundary conditions. Note that even though the domains of three-dimensional PDEs arising in medicine [166, 219, 311], engineering [28, 107, 263] and geosciences [175, 226, 347] are rarely cubes, they can often be mapped onto cubes [124]. Following the idea of global spectral methods, we approximate both the solution $u$ and the right-hand side $f$ by multivariate polynomials of the form (1.3). We approximate the operator $\mathcal{L}$ by a mapping mimicking the impact of applying $\mathcal{L}$ on the level of the coefficient tensors. In order to define this mapping, we express the three-dimensional differential operator as a combination of one-dimensional linear differential operators, which we determine using a CP decomposition (2.4). By representing $u$ in a Chebyshev basis (2.16) and $f$ in an ultraspherical basis, we can express the action of each of these one-dimensional operators on the coefficients in terms of sparse and well-conditioned differentiation and multiplication matrices [246]. The combination of applying these matrices yields the desired mapping. Solving the PDE corresponds to inverting this
mapping, which can be seen as solving a tensor-valued linear system. Discretizing the boundary conditions provides additional linear constraints under which inverting the linear system has a unique solution. Using substitution we obtain an unconstrained tensor-valued linear system uniquely determining a subtensor, from which we recover the full coefficient tensor of the solution. We would like to point out that our discretization approach can be applied to a wide variety of PDEs, but it does not necessarily preserve certain desirable properties of the differential operator.

It turns out that our discretization of the PDE (7.1) has a Kronecker structure that can often be exploited. For instance, Poisson's equation with homogeneous Dirichlet boundary conditions leads to a Laplace-like equation [24, 70, 206, 236]. Laplace-like equations can be solved efficiently using the blocked recursive algorithm in [69]. This algorithm is asymptotically slower than the nested alternating direction implicit method in [125], but our numerical experiments in Section 7.4 demonstrate that the recursive blocked algorithm is significantly faster in practice. Even if the PDE of interest does not immediately lead to a structure suitable for the fast solver, we can often apply the blocked recursive algorithm as a preconditioner for GMRES.

The algorithmic ideas presented in this chapter, can be extended from solving stationary linear PDEs, such as the Helmholtz equation and (convection) diffusion problems, to solving time-dependent PDEs and PDE eigenvalue problems of the form

$$
\frac{\partial}{\partial t} u=\mathcal{L} u \quad \text { and } \quad \mathcal{L} u=\lambda u
$$

which we solve using implicit Euler and inverse iteration methods, respectively. We expect to obtain accurate approximations of the solution, when both the PDE coefficients and the solution are sufficiently smooth to be well approximated by truncated expansions with tensorized polynomial basis functions. For non-smooth solutions, our global spectral method might lead to inaccurate approximations.

Remark 7.1. Our global spectral method can be used to compute solutions numerically to very high accuracy. It is not to be confused with so-called spectral element methods [124, 164, 184] and $p$ - and hp-finite element methods [18, 244, 342]. In those methods, $u$ is not approximated globally by a truncated series expansion. Instead, $u$ is written as sum of (locally supported) functions, each of which is approximated individually by a truncated series expansion. There also exist solvers relying on using domain decompositions in combination with truncated series expansions [160, 259]. We want to emphasize that the methods presented in this work can be used as local solver, when the elements/subdomains can be mapped to cubes.

This chapter is based on the article [304]. Its remainder is structured as follows. In Section 7.1, we define the linear differential operator and the approximation format. In Section 7.2, we derive the mapping on the level of the coefficient tensor. The discretization
of PDEs and the efficient solution of the resulting tensor-valued linear system is discussed in Section 7.3. In Section 7.4, we apply our global spectral method to solve stationary PDEs, parabolic PDEs and PDE eigenvalue problems numerically to very high accuracy.

### 7.1 Problem setting

### 7.1.1 Structure of a linear differential operator

A linear partial differential operator $\mathcal{L}$ on the domain $[-1,1] \times[-1,1] \times[-1,1]$ maps a sufficiently smooth function $u:[-1,1]^{3} \rightarrow \mathbb{R}$ to

$$
\begin{equation*}
\mathcal{L} u(x, y, z)=\sum_{a=0}^{N_{x}} \sum_{b=0}^{N_{y}} \sum_{c=0}^{N_{z}} \alpha_{a b c}(x, y, z) \frac{\partial^{a+b+c}}{\partial x^{a} \partial y^{b} \partial z^{c}} u(x, y, z), \tag{7.2}
\end{equation*}
$$

where $N_{x}, N_{y}, N_{z}$ are called differential order and $\alpha_{a b c}(x, y, z):[-1,1]^{3} \rightarrow \mathbb{R}$ are coefficient functions for $0 \leq a \leq N_{x}, 0 \leq b \leq N_{y}, 0 \leq c \leq N_{z}$. For the differential operator $\mathcal{L}$ we consider the linear PDE

$$
\begin{equation*}
\mathcal{L} u=f \tag{7.3}
\end{equation*}
$$

with right hand side $f:[-1,1]^{3} \rightarrow \mathbb{R}$. The PDE can be solved uniquely when the system is complemented with sufficient boundary conditions.

### 7.1.2 Approximation format

We approximate the solution $u:[-1,1]^{3} \rightarrow \mathbb{R}$ of the $\operatorname{PDE}(7.3)$ in the space $\mathbb{P}_{n_{1}, n_{2}, n_{3}}$ of trivariate polynomials of degree at most $\left(n_{1}, n_{2}, n_{3}\right)$. We express $u$ in terms of a tensorized basis of Chebyshev polynomials (see Section 2.3), which leads to a representation of the form

$$
\begin{equation*}
u(x, y, z) \approx \sum_{i=0}^{n_{1}} \sum_{j=0}^{n_{2}} \sum_{k=0}^{n_{3}} \mathcal{U}_{i j k} T_{i}(x) T_{j}(y) T_{k}(z) \tag{7.4}
\end{equation*}
$$

with coefficient tensor $\mathcal{U} \in \mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times\left(n_{3}+1\right)}$.

### 7.2 Operator discretization

In the following, we discretize the differential operator $\mathcal{L}$ for fixed coefficient functions $\alpha_{\text {abc }}$ by approximating $\mathcal{L}$ as mapping from $\mathbb{P}_{n_{1}, n_{2}, n_{3}}$ to $\mathbb{P}_{n_{1}, n_{2}, n_{3}}$. This is particularly easy for constant coefficients in the differential operator $\mathcal{L}$, i.e., $\alpha_{a b c}(x, y, z)=\alpha_{a b c} \in \mathbb{R}$ in Equation (7.2). Applying a linear differential operator with constant coefficients to a
polynomial does not increase the polynomial degree. For non-constant coefficients an additional truncation is needed to obtain a polynomial in $\mathbb{P}_{n_{1}, n_{2}, n_{3}}$. Thus, it is natural to see the application of the operator as a transformation of the coefficient tensor. We discuss the discretization for the constant case first before generalizing to the non-constant case in Section 7.2.3.

### 7.2.1 One-dimensional differential operators

We briefly recapitulate how to obtain the mapping describing the transformation of coefficients in a one-dimensional setting before returning to the three-dimensional setting. Let $u:[-1,1] \rightarrow \mathbb{R}$ be a polynomial of degree $n$ represented in the Chebyshev basis by $u(x)=\sum_{k=0}^{n} u_{k} T_{k}(x)$, with coefficients $u_{k} \in \mathbb{R}$. Applying a one-dimensional linear differential operator $\mathcal{L}$ of order $N$ with constant coefficients to $u$ can be written as

$$
\begin{equation*}
\mathcal{L} u(x)=\sum_{a=0}^{N} \alpha_{a} \frac{d^{a}}{d x^{a}} u(x), \tag{7.5}
\end{equation*}
$$

with coefficients $\alpha_{a} \in \mathbb{R}$. This is a linear combination of (higher order) derivatives of $u$. For every derivative of the polynomial $u$, there exists a so-called differentiation matrix which maps the coefficient vector $\mathbf{u}=\left(u_{0}, \ldots, u_{n}\right)$ to the coefficients of the derivative. The remainder of this section follows the ideas of [246] to represent the derivative using an ultraspherical basis instead of a Chebyshev basis, which leads to better conditioned and sparse differentiation matrices.

For the parameter $\lambda>0$ and $k=0,1, \ldots$, ultraspherical polynomials follow the recurrence relation

$$
\begin{aligned}
& 4 \lambda(k+\lambda+1)\left(1-x^{2}\right) C_{k}^{(\lambda+1)}(x)= \\
& \quad-(k+1)(k+2) C_{k+2}^{(\lambda)}(x)+(k+2 \lambda)(k+2 \lambda+1) C_{k}^{(\lambda)}(x),
\end{aligned}
$$

where $C_{k}^{(1)}(x)=\left(\sin (k+1) \cos ^{-1}(x)\right) / \sin \left(\cos ^{-1}(x)\right)$ [240]. Let $v(x)=\sum_{k=0}^{n} v_{k} C_{k}^{(\lambda)}(x)$ denote the $\lambda$ th derivative of $u$ represented in $C^{(\lambda)}$ basis, then the coefficient vector $\mathbf{v}=$ $\left(v_{0}, \ldots, v_{n}\right)$ satisfies $\mathbf{v}=D_{\lambda} \mathbf{u}$, where the sparse differentiation matrix $D_{\lambda} \in \mathbb{R}^{(n+1) \times(n+1)}$
is defined as

$$
D_{\lambda}=2^{\lambda-1}(\lambda-1)!\left(\begin{array}{ccccc}
\overbrace{0 \ldots 0}^{\lambda \text { times }} & & & & \\
\\
& & \lambda+1 & & \\
& & \ddots & \\
& & & n \\
& & & & 0 \\
& & & & 0
\end{array}\right) .
$$

Note that the ultraspherical basis is different for different $\lambda$. The sparse transformation matrices $S_{0} \in \mathbb{R}^{(n+1) \times(n+1)}$, mapping Chebyshev to $C^{(1)}$ coefficients, and $S_{\lambda} \in \mathbb{R}^{(n+1) \times(n+1)}$, mapping $C^{(\lambda)}$ to $C^{(\lambda+1)}$, are defined as

$$
\begin{aligned}
& S_{0}=\left(\begin{array}{cccccccc}
1 & 0 & -\frac{1}{2} & & & & \\
& \frac{1}{2} & 0 & -\frac{1}{2} & & & \\
& & \frac{1}{2} & 0 & -\frac{1}{2} & & \\
& & & \ddots & \ddots & \ddots & \\
& & & & \frac{1}{2} & 0 & -\frac{1}{2} \\
& & & & & \frac{1}{2} & 0 \\
& & & & & \frac{1}{2}
\end{array}\right), \\
& S_{\lambda}=\left(\begin{array}{ccccccc}
1 & 0 & -\frac{\lambda}{\lambda+2} & & & & \\
& \frac{\lambda}{\lambda+1} & 0 & -\frac{\lambda}{\lambda+3} & & & \\
& & \frac{\lambda}{\lambda+2} & 0 & -\frac{\lambda}{\lambda+4} & & \\
& & & & \ddots & \ddots & \ddots \\
\\
& & & & \frac{\lambda}{\lambda+n-2} & 0 & -\frac{\lambda}{\lambda+n} \\
& & & & & \frac{\lambda}{\lambda+n-1} & 0 \\
& & & & & & \frac{\lambda}{\lambda+n}
\end{array}\right) .
\end{aligned}
$$

Let $\mathcal{L} u(x)=\sum_{k=0}^{n} w_{k} C_{k}^{(N)}(x)$ be represented in $C^{(N)}$ ultraspherical basis with coefficient vector $\mathbf{w}=\left(w_{0}, \ldots, w_{n}\right)$, then

$$
\begin{equation*}
\mathbf{w}=(\underbrace{a_{N} D_{N}+a_{N-1} S_{N-1} D_{N-1}+\cdots+a_{1} S_{N-1} \cdots S_{1} D_{1}+a_{0} S_{N-1} \cdots S_{0}}_{=L}) \mathbf{u} . \tag{7.6}
\end{equation*}
$$

The matrix $L \in \mathbb{R}^{(n+1) \times(n+1)}$ describes how the one-dimensional differential operator $\mathcal{L}$ acts on Chebyshev coefficients of the solution.

### 7.2.2 Three-dimensional differential operators

In Chebop2 [319] it is suggested to split two-dimensional differential operators via an SVD into a sum of one-dimensional operators, for which the matrices $L$ can be computed as in Section 7.2.1. Such splittings can be generalized to higher dimensional settings via CP decompositions (2.4).

Let $\mathcal{A} \in \mathbb{R}^{\left(N_{x}+1\right) \times\left(N_{y}+1\right) \times\left(N_{z}+1\right)}$ denote the tensor with entries $\mathcal{A}_{a, b, c}$ given by the coefficients $\alpha_{a b c}$ in Equation (7.2) for $0 \leq a \leq N_{x}, 0 \leq b \leq N_{y}, 0 \leq c \leq N_{z}$. Given a CP decompositon (2.4) of the form $\mathcal{A}=\sum_{r=1}^{R} \mathbf{a}^{(r)} \circ \mathbf{b}^{(r)} \circ \mathbf{c}^{(r)}$, we can rewrite the application of $\mathcal{L}$ to $u$ as

$$
\begin{equation*}
\mathcal{L} u(x)=\sum_{r=1}^{R} \underbrace{\sum_{a=0}^{N_{x}} \mathbf{a}_{a}^{(r)} \frac{\partial^{a}}{\partial x^{a}}}_{=\mathcal{L}_{r}^{(x)}} \underbrace{\sum_{b=0}^{N_{y}} \mathbf{b}_{b}^{(r)} \frac{\partial^{b}}{\partial y^{b}}}_{=\mathcal{L}_{r}^{(y)}} \underbrace{\sum_{a=0}^{N_{z}} \mathbf{c}_{c}^{(r)} \frac{\partial^{c}}{\partial z^{c}}}_{=\mathcal{L}_{r}^{(z)}} u(x, y, z) \tag{7.7}
\end{equation*}
$$

in terms of one-dimensional differential operators $\mathcal{L}_{r}^{(x)}, \mathcal{L}_{r}^{(y)}, \mathcal{L}_{r}^{(z)}$.
Let $L_{r}^{(x)} \in \mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{1}+1\right)}, L_{r}^{(y)} \in \mathbb{R}^{\left(n_{2}+1\right) \times\left(n_{2}+1\right)}, L_{r}^{(z)} \in \mathbb{R}^{\left(n_{3}+1\right) \times\left(n_{3}+1\right)}$ denote the matrices $L$ associated with the operators $\mathcal{L}_{r}^{(x)}, \mathcal{L}_{r}^{(y)}, \mathcal{L}_{r}^{(z)}$ as defined in Equation (7.6). Let $u$ be a polynomial with coefficient tensor $\mathcal{U} \in \mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times\left(n_{3}+1\right)}$ satisfying Equation (7.4) with equality. Then

$$
\mathcal{L} u(x, y, z)=\sum_{i=0}^{n_{1}} \sum_{j=0}^{n_{2}} \sum_{k=0}^{n_{3}} \mathcal{V}_{i j k} C_{i}^{\left(N_{x}\right)}(x) C_{j}^{\left(N_{y}\right)}(y) C_{j}^{\left(N_{z}\right)}(z)
$$

with coefficient tensor

$$
\begin{equation*}
\mathcal{V}=\sum_{r=1}^{R} \mathcal{U} \times_{1} L_{r}^{(x)} \times_{2} L_{r}^{(y)} \times_{3} L_{r}^{(z)} \tag{7.8}
\end{equation*}
$$

### 7.2.3 Generalization to non-constant coefficients

## One-dimensional differential operators

We now consider the one-dimensional differential operator (7.5) with non-constant coefficients $\alpha_{a}:[-1,1] \rightarrow \mathbb{R}$. In this section, we define multiplication matrices to incorporate these coefficients in the discretization.

Let $u(x)$ be a polynomial of the form $u(x)=\sum_{k=0}^{n} u_{k} B_{k}(x)$, with basis functions $B_{k}(x)$ chosen as Chebyshev polynomials $T_{k}$ or ultraspherical polynomials $C_{k}^{(\lambda)}$. To multiply $u(x)$ by a function $v:[-1,1] \rightarrow \mathbb{R}$, we approximate $v(x) \approx \tilde{v}(x)=\sum_{k=0}^{n} v_{k} B_{k}(x)$ by a polynomial in the same basis using Chebyshev interpolation [322] (and basis transformations).

Further, we also approximate the product $\tilde{v}(x) u(x) \approx z(x)=\sum_{k=0}^{n} z_{k} B_{k}(x)$ by a polynomial in the same basis. There exist so-called multiplication matrices [124, 246] depending on the coefficients $\mathbf{v}=\left(v_{0}, \ldots, v_{n}\right)$, which map the coefficients $\mathbf{u}=\left(u_{0}, \ldots, u_{n}\right)$ to the coefficients $\mathbf{z}=\left(z_{0}, \ldots, z_{n}\right)$ such that $z(x)$ approximates $\tilde{v}(x) u(x)$. In the following, we summarize how these matrices are defined in [318] for both Chebyshev and ultraspherical bases. For the Chebyshev basis, we define $\mathbf{z}=M[\mathbf{v}] \mathbf{u}$ and the multiplication matrix $M[\mathbf{v}] \in \mathbb{R}^{(n+1) \times(n+1)}$ given by

$$
\begin{aligned}
M[\mathbf{v}]= & \frac{1}{2}\left(\begin{array}{cccccc}
2 v_{0} & v_{1} & v_{2} & \ldots & v_{n-1} & v_{n} \\
v_{1} & 2 v_{0} & v_{1} & \ldots & v_{n-2} & v_{n-1} \\
v_{2} & v_{1} & 2 v_{0} & \ldots & v_{n-3} & v_{n-2} \\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
v_{n-1} & v_{n-2} & v_{n-3} & \ldots & 2 v_{0} & v_{1} \\
v_{n} & v_{n-1} & v_{n-2} & \ldots & v_{1} & 2 v_{0}
\end{array}\right) \\
& +\frac{1}{2}\left(\begin{array}{ccccccc}
0 & 0 & 0 & \ldots & 0 & 0 & 0 \\
v_{1} & v_{2} & v_{3} & \ldots & v_{n-1} & v_{n} & 0 \\
v_{2} & v_{3} & v_{4} & \ldots & v_{n} & 0 & 0 \\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\
v_{n-1} & v_{n} & 0 & \ldots & 0 & 0 & 0 \\
v_{n} & 0 & 0 & \ldots & 0 & 0 & 0
\end{array}\right)
\end{aligned}
$$

For the $C^{(\lambda)}$ ultraspherical basis, we define $\mathbf{z}=M^{(\lambda)}[\mathbf{v}] \mathbf{u}$ and the multiplication matrix $M^{(\lambda)}[\mathbf{v}] \in \mathbb{R}^{(n+1) \times(n+1)}$ as $M^{(\lambda)}[\mathbf{v}]=\sum_{i=0}^{n} v_{i} M_{i}^{(\lambda)}$. The matrices $M_{i}^{(\lambda)} \in \mathbb{R}^{(n+1) \times(n+1)}$ are defined recursively for $i=0,1, \ldots$ by

$$
(i+2) M_{i+2}^{(\lambda)}=2(i+\lambda+1) N^{(\lambda)} M_{i+1}^{(\lambda)}-(i+2 \lambda) M_{i}^{(\lambda)}
$$

with $M_{0}^{(\lambda)}=I, M_{1}^{(\lambda)}=2 \lambda N^{(\lambda)}$ and $N^{(\lambda)} \in \mathbb{R}^{(n+1) \times(n+1)}$ defined as

$$
N^{(\lambda)}=\left(\begin{array}{ccccccc}
0 & \frac{2 \lambda}{2(\lambda+1)} & & & & & \\
\frac{1}{2 \lambda} & 0 & \frac{2 \lambda+1}{2(\lambda+2)} & & & & \\
0 & \frac{2}{2(\lambda+1)} & 0 & \frac{2 \lambda+2}{2(\lambda+3)} & & & \\
& & \frac{2}{2(\lambda+2)} & 0 & \frac{2 \lambda+3}{2(\lambda+4)} & & \\
& & & \ddots & \ddots & \ddots & \\
& & & & \frac{2}{2(\lambda+n-2)} & 0 & \frac{2 \lambda+n-1}{2(\lambda+n)} \\
& & & & & \frac{2}{2(\lambda+n-1)} & 0
\end{array}\right) .
$$

We now incorporate the non-constant coefficients $\alpha_{a}(x)$ into the coefficient mapping. Let $\mathbf{a}_{a} \in \mathbb{R}^{(n+1)}$ denote coefficients of polynomial approximations of $\alpha_{a}(x)$ in the Chebyshev basis for $a=0$ and in $C^{(a)}$ ultraspherical basis for $a=1,2, \ldots$ Analogous to

Equation (7.6), we approximate $\mathcal{L} u \approx \sum_{k=0}^{n} w_{k} T_{k}(x)$ by computing the coefficients $\mathbf{w}=\left(w_{0}, \ldots, w_{n}\right)$ defined as

$$
\begin{aligned}
\mathbf{w}=L \mathbf{u}:= & \left(M^{(N)}\left[\mathbf{a}_{N}\right] D_{N}+S_{N-1} M^{(N-1)}\left[\mathbf{a}_{N-1}\right] D_{N-1}+\ldots\right. \\
& \left.+S_{N-1} \cdots S_{1} M^{(1)}\left[\mathbf{a}_{1}\right] D_{1}+S_{N-1} \cdots S_{0} M\left[\mathbf{a}_{0}\right]\right) \mathbf{u}
\end{aligned}
$$

Remark 7.2. The interpolation of $v$ and the approximation of the product $\tilde{v} u$ introduce truncation errors. For sufficiently large $n$, these errors are close to machine precision.

## Three-dimensional differential operators

We now consider the three-dimensional differential operator $\mathcal{L}$ as defined in Equation (7.2) with non-constant coefficients. We proceed by splitting this operator into one-dimensional operators with non-constant coefficients.

Let the polynomial degree $\left(n_{1}, n_{2}, n_{3}\right)$ be chosen sufficiently large to accurately approximate the coefficient functions $\alpha_{a b c}(x, y, z)$ using tensorized Chebyshev interpolation in the form of Equation (7.4). Let $\mathcal{B}^{(a b c)} \in \mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times\left(n_{3}+1\right)}$ denote the coefficient tensors corresponding to $\alpha_{a b c}(x, y, z)$ for each multi-index $a b c$. We define the tensor $\mathcal{A} \in \mathbb{R}^{\left(N_{x}+1\right) \times\left(n_{1}+1\right) \times\left(N_{y}+1\right) \times\left(n_{2}+1\right) \times\left(N_{z}+1\right) \times\left(n_{3}+1\right)}$ with entries

$$
\mathcal{A}_{a i b j c k}=\mathcal{B}_{i j k}^{(a b c)}
$$

for $0 \leq a \leq N_{x}, 0 \leq b \leq N_{y}, 0 \leq c \leq N_{z}, 0 \leq i \leq n_{1}, 0 \leq j \leq n_{2}, 0 \leq k \leq n_{3}$. This approximates the differential operator in the sense that

$$
\begin{aligned}
& \mathcal{L} u(x, y, z) \approx \tilde{\mathcal{L}} u(x, y, z):= \\
& \quad \sum_{a=0}^{N_{x}} \sum_{i=0}^{n_{1}} \sum_{b=0}^{N_{y}} \sum_{j=0}^{n_{2}} \sum_{c=0}^{N_{z}} \sum_{k=0}^{n_{3}} \mathcal{A}_{a i b j c k} T_{i}(x) T_{j}(y) T_{k}(z) \frac{\partial^{a+b+c}}{\partial x^{a} \partial y^{b} \partial z^{c}} u(x, y, z)
\end{aligned}
$$

We now reshape $\mathcal{A}$ into a tensor of order 3 in $\mathbb{R}^{\left(N_{x}+1\right)\left(n_{1}+1\right) \times\left(N_{y}+1\right)\left(n_{2}+1\right) \times\left(N_{z}+1\right)\left(n_{3}+1\right)}$ and compute a CP decomposition of the form

$$
\mathcal{A}=\sum_{r=1}^{R} \mathbf{a}^{(r)} \circ \mathbf{b}^{(r)} \circ \mathbf{c}^{(r)}
$$

where $\mathbf{a}^{(r)} \in \mathbb{R}^{\left(N_{x}+1\right)\left(n_{1}+1\right)}, \mathbf{b}^{(r)} \in \mathbb{R}^{\left(N_{y}+1\right)\left(n_{2}+1\right)}, \mathbf{c}^{(r)} \in \mathbb{R}^{\left(N_{z}+1\right)\left(n_{3}+1\right)}$. We reshape the vectors $\mathbf{a}_{r}, \mathbf{b}_{r}, \mathbf{c}_{r}$ back into matrices in $\mathbf{A}^{(r)} \in \mathbb{R}^{\left(N_{x}+1\right) \times\left(n_{1}+1\right)}, \mathbf{B}^{(r)} \in \mathbb{R}^{\left(N_{y}+1\right) \times\left(n_{2}+1\right)}$ and $\mathbf{C}^{(r)} \in \mathbb{R}^{\left(N_{z}+1\right) \times\left(n_{3}+1\right)}$, respectively. Analogously to Equation (7.7), we now define
one-dimensional differential operators

$$
\begin{aligned}
\mathcal{L}_{r}^{(x)} u(x) & =\sum_{a=0}^{N_{x}}\left(\sum_{i=0}^{n_{1}} \mathbf{A}_{a i}^{(r)} T_{i}(x)\right) \frac{\partial^{a}}{\partial x^{a}} u(x), \\
\mathcal{L}_{r}^{(y)} u(y) & =\sum_{b=0}^{N_{y}}\left(\sum_{j=0}^{n_{2}} \mathbf{B}_{b j}^{(r)} T_{j}(y)\right) \frac{\partial^{b}}{\partial y^{b}} u(y), \\
\mathcal{L}_{r}^{(z)} u(z) & =\sum_{c=0}^{N_{z}}\left(\sum_{k=0}^{n_{3}} \mathbf{C}_{c k}^{(r)} T_{k}(z)\right) \frac{\partial^{c}}{\partial z^{c}} u(z),
\end{aligned}
$$

which satisfy by construction

$$
\begin{equation*}
\tilde{\mathcal{L}} u(x, y, z)=\sum_{r=1}^{R}\left(\mathcal{L}_{r}^{(x)} \otimes \mathcal{L}_{r}^{(y)} \otimes \mathcal{L}_{r}^{(z)}\right) u(x, y, z) \tag{7.9}
\end{equation*}
$$

where $\otimes$ denotes the tensor product for linear operators. Note that the terms of the form $\sum_{i=0}^{n_{1}} \mathbf{A}_{a i}^{(r)} T_{i}(x)$ are univaritate functions. So, each of the one-dimensional differential operators fits into the setting of Section 7.2.3 and we can obtain matrices $L_{r}^{(x)}, L_{r}^{(y)}, L_{r}^{(z)}$ as in Equation (7.8).

Remark 7.3. In our numerical experiments we use using tensorlab [332] to obtain the CP decomposition from the full tensor $\mathcal{A}$. Since $\mathcal{A}$ is potentially very large, it might sometimes be necessary to rely on heuristic algorithms that do not explicitly require the full tensor.

### 7.3 A spectral method for three-dimensional linear PDEs

In this section, we present how to compute approximate solutions for PDEs of the form (7.3). We again discretize the differential operator $\mathcal{L}$ as in Section 7.2. Additionally, we discretize the right hand side $f$ using a truncated expansion with ultraspherical basis functions of the form

$$
f(x, y, z) \approx \sum_{i=0}^{n_{1}} \sum_{j=0}^{n_{2}} \sum_{k=0}^{n_{3}} \mathcal{F}_{i j k} C_{i}^{\left(N_{x}\right)}(x) C_{j}^{\left(N_{y}\right)}(y) C_{k}^{\left(N_{z}\right)}(z),
$$

with coefficient tensor $\mathcal{F} \in \mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times\left(n_{3}+1\right)}$. The tensor $\mathcal{F}$ can be computed by applying suitable basis transformations to the coefficient tensor obtained from tensorized Chebyshev interpolation. The discretized PDE reads as tensor-valued linear system of the form

$$
\begin{equation*}
\sum_{r=1}^{R} \mathcal{U} \times_{1} L_{r}^{(x)} \times_{2} L_{r}^{(y)} \times_{3} L_{r}^{(z)}=\mathcal{F} \tag{7.10}
\end{equation*}
$$

It remains to incorporate the boundary conditions.

### 7.3.1 Boundary condition discretization

Following the ideas of [319], we can discretize commonly used boundary conditions for three-dimensional PDEs on cubes as constraints of the form

$$
\begin{equation*}
\mathcal{U} \times_{1} B_{1}=\mathcal{G}_{1}, \quad \mathcal{U} \times_{2} B_{2}=\mathcal{G}_{2}, \quad \mathcal{U} \times_{3} B_{3}=\mathcal{G}_{3}, \tag{7.11}
\end{equation*}
$$

where the matrices $B_{1} \in \mathbb{R}^{N_{x} \times\left(n_{1}+1\right)}, B_{2} \in \mathbb{R}^{N_{y} \times\left(n_{2}+1\right)}, B_{3} \in \mathbb{R}^{N_{z} \times\left(n_{3}+1\right)}$ have linearly independent rows and $\mathcal{G}_{1} \in \mathbb{R}^{N_{x} \times\left(n_{2}+1\right) \times\left(n_{3}+1\right)}, \mathcal{G}_{2} \in \mathbb{R}^{\left(n_{1}+1\right) \times N_{y} \times\left(n_{3}+1\right)}, \mathcal{G}_{3} \in$ $\mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times N_{z}}$. We present two examples of how constraints of the form (7.11) can be derived. We want to emphasize that the discretized boundary conditions need to satisfy compatibility constraints as in [318, Section 6.4].

## Dirichlet conditions

We consider the Dirichlet boundary condition $u(1, y, z)=h(y, z)$ for a given function $h:[-1,1] \rightarrow \mathbb{R}$. We approximate the function $h$ using bivariate Chebyshev interpolation in $\left(n_{2}+1\right) \times\left(n_{3}+1\right)$ points. Let $H \in \mathbb{R}^{\left(n_{2}+1\right) \times\left(n_{3}+1\right)}$ denote the corresponding coefficient matrix of $h$. We can enforce that the solution $u$ given in Chebyshev basis (7.4) coincides with the interpolant of $h$ by demanding that

$$
\begin{equation*}
u(1, y, z)=\sum_{i=0}^{n_{1}} \sum_{j=0}^{n_{2}} \sum_{k=0}^{n_{3}} \mathcal{U}_{i j k} T_{i}(1) T_{j}(y) T_{k}(z)=\sum_{j=0}^{n_{2}} \sum_{k=0}^{n_{3}} H_{j k} T_{j}(y) T_{k}(z) \tag{7.12}
\end{equation*}
$$

This can be equivalently written as $U \times_{1} B_{1}=\mathcal{G}_{1}$ with $B_{1}=\left(T_{0}(1), \ldots, T_{n_{1}}(1)\right)$ and $\left(G_{1}\right)_{0,0: n_{2}, 0: n_{3}}=H$.

We now consider an additional Dirichlet boundary condition $u(-1, y, z)=\tilde{h}$. Again, we compute the coefficient matrix $\tilde{H}$ corresponding to $\tilde{h}$. Analogously to Equation (7.12), we enforce that $u(-1, y, z)$ coincides with the interpolant of $\tilde{h}$. We can express both conditions simultaneously in the form of (7.11) by defining

$$
B_{1}=\left(\begin{array}{cccc}
T_{0}(-1) & T_{1}(-1) & \ldots & T_{n}(-1) \\
T_{0}(1) & T_{1}(1) & \ldots & T_{n}(1)
\end{array}\right)
$$

and $\left(G_{1}\right)_{0,0: n_{2}, 0: n_{3}}=\tilde{H},\left(G_{1}\right)_{1,0: n_{2}, 0: n_{3}}=H$.

## Mixed Dirchlet and Neumann conditions

We consider one example of mixed boundary conditions with Neumann boundary conditions on the right side of the cube $[-1,1]^{3}$ and Dirichlet boundary conditions on all other sides. The Neumann boundary condition is given by $\frac{\partial}{\partial x} u(1, y, z)=h(y, z)$ for a given function $h$ and $u \in \mathbb{P}_{n_{1}, n_{2}, n_{3}}$ represented in Chebyshev basis (7.4). As in the Dirichlet case, we use bivariate Chebyshev interpolation $h$ to obtain the coefficient matrix $H \in \mathbb{R}^{\left(n_{2}+1\right) \times\left(n_{3}+1\right)}$. We now demand that

$$
\begin{equation*}
\frac{\partial}{\partial x} u(1, y, z)=\sum_{i=0}^{n_{1}} \sum_{j=0}^{n_{2}} \sum_{k=0}^{n_{3}} \mathcal{U}_{i j k} T_{i}^{\prime}(1) T_{j}(y) T_{k}(z)=\sum_{j=0}^{n_{2}} \sum_{k=0}^{n_{3}} H_{j k} T_{j}(y) T_{k}(z), \tag{7.13}
\end{equation*}
$$

where $T_{i}^{\prime}(1)=i^{2}$ for $i=0, \ldots, n[322]$. Equation (7.13) can be expressed in the form of (7.11) with

$$
B_{1}=\left(T_{0}^{\prime}(1), \ldots, T_{n_{1}}^{\prime}(1)\right) \quad \text { and } \quad\left(G_{1}\right)_{0,0: n_{2}, 0: n_{3}}=H
$$

The Dirichlet boundary conditions on the left side can be included in $B_{1}, \mathcal{G}_{1}$ as in Section 7.3.1.

### 7.3.2 Incorporating the boundary conditions

We need to incorporate the discretized boundary conditions (7.11) into the discretized $\operatorname{PDE}(7.10)$ to obtain the unique solution $\mathcal{U}$. Following the ideas of [319, Section 6], we compute $\mathcal{U}$ by substituting (7.11) into (7.10).

Since $B_{1}, B_{2}, B_{3}$ have linearly independent rows, we can assume without loss of generality that
$\left(B_{1}\right)_{0:\left(N_{x}-1\right), 0:\left(N_{x}-1\right)}=I, \quad\left(B_{2}\right)_{0:\left(N_{y}-1\right), 0:\left(N_{y}-1\right)}=I, \quad\left(B_{3}\right)_{\left(0:\left(N_{z}-1\right), 0:\left(N_{z}-1\right)\right.}=I$.

For $r=1, \ldots, R$, we rewrite the boundary conditions as

$$
\begin{aligned}
& -\mathcal{U} \times_{1}\left(L_{r}^{(x)}\right)_{0: n_{1}, 0:\left(N_{x}-1\right)} B_{1}=-\mathcal{G}_{1} \times_{1}\left(L_{r}^{(x)}\right)_{0: n_{1}, 0:\left(N_{x}-1\right)}, \\
& -\mathcal{U} \times_{2}\left(L_{r}^{(y)}\right)_{0: n_{2}, 0:\left(N_{y}-1\right)} B_{2}=-\mathcal{G}_{2} \times_{2}\left(L_{r}^{(y)}\right)_{0: n_{2}, 0:\left(N_{y}-1\right)}, \\
& -\mathcal{U} \times_{3}\left(L_{r}^{(z)}\right)_{0: n_{3}, 0:\left(N_{z}-1\right)} B_{3}=-\mathcal{G}_{3} \times_{3}\left(L_{r}^{(z)}\right)_{0: n_{3}, 0:\left(N_{z}-1\right)} .
\end{aligned}
$$



Figure 7.1 - Visualization of the block decomposition of $\mathcal{U} \in \mathbb{R}^{\left(n_{1}+1\right) \times\left(n_{2}+1\right) \times\left(n_{3}+1\right)}$ with block $\mathcal{U}^{111} \in \mathbb{R}^{N_{x} \times N_{y} \times N_{z}}$.

Substituting these modified boundary conditions into the discretized PDE (7.10) leads to

$$
\begin{aligned}
& \sum_{r=1}^{R} \mathcal{U} \times \underbrace{\left(L_{r}^{(x)}-\left(L_{r}^{(x)}\right)_{0: n_{1}, 0:\left(N_{x}-1\right)} B_{1}\right)}_{=\tilde{L}_{r}^{(x)}} \times 2 \underbrace{\left(L_{r}^{(y)}-\left(L_{r}^{(y)}\right)_{0: n_{2}, 0:\left(N_{y}-1\right)} B_{2}\right)}_{=\tilde{L}_{r}^{(y)}} \\
& \times_{3} \underbrace{\left(L_{r}^{(z)}-\left(L_{r}^{(z)}\right)_{0: n_{3}, 0:\left(N_{z}-1\right)} B_{3}\right)}_{=\tilde{L}_{r}^{(z)}}= \\
& \mathcal{F}-\sum_{r=1}^{R} \mathcal{G}_{1} \times_{1}\left(L_{r}^{(x)}\right)_{\left(0: n_{1}, 0:\left(N_{x}-1\right)\right.} \times{ }_{2} L_{r}^{(y)} \times_{3} L_{r}^{(z)} \\
& -\sum_{r=1}^{R} \mathcal{G}_{2} \times_{1}\left(L_{r}^{(x)}-\left(L_{r}^{(x)}\right)_{0: n_{1}, 0:\left(N_{x}-1\right)} B_{1}\right) \times_{2}\left(L_{r}^{(y)}\right)_{0: n_{2}, 0:\left(N_{y}-1\right)} \times{ }_{3} L_{r}^{(z)} \\
& -\sum_{r=1}^{R} \mathcal{G}_{3} \times_{1}\left(L_{r}^{(x)}-\left(L_{r}^{(x)}\right)_{0: n_{1}, 0:\left(N_{x}-1\right)} B_{1}\right) \times_{2}\left(L_{r}^{(y)}-\left(L_{r}^{(y)}\right)_{0: n_{2}, 0:\left(N_{y}-1\right)} B_{2}\right) \\
& \times_{3}\left(L_{r}^{(z)}\right)_{0: n_{3}, 0:\left(N_{z}-1\right)},
\end{aligned}
$$

where we denote the right hand side by $\tilde{\mathcal{F}}$. Observe that the first $N_{x}, N_{y}, N_{z}$ columns of $\tilde{L}_{r}^{(x)}, \tilde{L}_{r}^{(y)}, \tilde{L}_{r}^{(z)}$ are zero due to assumption (7.14). Let $\mathcal{U}$ be decomposed into tensor blocks as described in Figure 7.1. The system after substitution uniquely determines the block $\mathcal{U}^{222}=\mathcal{U}_{N_{x}: n_{1}, N_{y}: n_{2}, N_{z}: n_{3}}$, which can be written as

$$
\begin{align*}
& \sum_{r=1}^{R} \mathcal{U}^{222} \times \underbrace{\left(\tilde{L}_{r}^{(x)}\right)_{0:\left(n_{1}-N_{x}\right), N_{x}: n_{1}}}_{=\hat{L}_{r}^{(x)}} \times \underbrace{}_{2} \underbrace{\left(\tilde{L}_{r}^{(y)}\right)_{0:\left(n_{2}-N_{y}\right), N_{y}: n_{2}}}_{=\hat{L}_{r}^{(y)}} \times 3 \underbrace{\left(\tilde{L}_{r}^{(z)}\right)_{0:\left(n_{3}-N_{z}\right), N_{z}: n_{3}}}_{=\hat{L}_{r}^{(z)}} \\
& \quad=\underbrace{\tilde{\mathcal{F}}_{0:\left(n_{1}-N_{x}\right), 0:\left(n_{2}-N_{y}\right), 0:\left(n_{3}-N_{z}\right)}} . \tag{7.15}
\end{align*}
$$

The remaining blocks can be reconstructed from the discretized boundary conditions (7.11) as

$$
\begin{array}{ll}
\mathcal{U}^{122}=\mathcal{G}_{1}-\mathcal{U}^{222} \times_{1} B_{1}, & \mathcal{U}^{212}=\mathcal{G}_{2}-\mathcal{U}^{222} \times_{2} B_{2}, \\
\mathcal{U}^{221}=\mathcal{G}_{3}-\mathcal{U}^{222} \times_{3} B_{3}, & \mathcal{U}^{112}=\mathcal{G}_{1}-\mathcal{U}^{212} \times_{1} B_{1}, \\
\mathcal{U}^{121}=\mathcal{G}_{1}-\mathcal{U}^{221} \times_{1} B_{1}, & \mathcal{U}^{211}=\mathcal{G}_{3}-\mathcal{U}^{212} \times_{3} B_{3}, \\
\mathcal{U}^{111}=\mathcal{G}_{1}-\mathcal{U}^{211} \times_{1} B_{1} . &
\end{array}
$$

### 7.3.3 Solving tensor-valued linear systems

The computation of $\mathcal{U}$ requires the solution of the unconstrained tensor-valued linear system (7.15) of the form

$$
\begin{equation*}
\sum_{r=1}^{R} \mathcal{U}^{222} \times_{1} \hat{L}_{r}^{(x)} \times_{2} \hat{L}_{r}^{(y)} \times_{3} \hat{L}_{r}^{(z)}=\hat{\mathcal{F}} . \tag{7.16}
\end{equation*}
$$

This system can be solved by reshaping the tensor-valued linear system into a vector-valued linear system of the form

$$
\begin{equation*}
\left(\sum_{r=1}^{R} \hat{L}_{r}^{(z)} \otimes \hat{L}_{r}^{(y)} \otimes \hat{L}_{r}^{(x)}\right) \operatorname{vec}\left(\mathcal{U}^{222}\right)=\operatorname{vec}(\hat{\mathcal{F}}) \tag{7.17}
\end{equation*}
$$

For certain PDEs we can transform the system (7.16) into a Laplace-like equation

$$
\begin{equation*}
\mathcal{U}^{222} \times_{1} U+\mathcal{U}^{222} \times_{2} V+\mathcal{U}^{222} \times_{3} W=\breve{\mathcal{F}}, \tag{7.18}
\end{equation*}
$$

with matrices $U \in \mathbb{R}^{\left(n_{1}+1-N_{x}\right) \times\left(n_{1}+1-N_{x}\right)}, \quad V \in \mathbb{R}^{\left(n_{2}+1-N_{y}\right) \times\left(n_{2}+1-N_{y}\right)}, W \in$ $\mathbb{R}^{\left(n_{3}+1-N_{z}\right) \times\left(n_{3}+1-N_{z}\right)}$ and tensor $\breve{\mathcal{F}} \in \mathbb{R}^{\left(n_{1}+1-N_{x}\right) \times\left(n_{2}+1-N_{y}\right) \times\left(n_{3}+1-N_{z}\right)}$. For instance, this can be achieved for $N_{x}=N_{y}=N_{z}, n_{1}=n_{2}=n_{3}$ and $B_{1}=B_{2}=B_{3}$, when there exists a CP decomposition of the tensor $\mathcal{A}$ defined in Section 7.2 with symmetry constraints [61] of the form

$$
\begin{equation*}
\mathcal{A}=\mathbf{a} \circ \mathbf{v} \circ \mathbf{w}+\mathbf{u} \circ \mathbf{b} \circ \mathbf{w}+\mathbf{u} \circ \mathbf{v} \circ \mathbf{c} . \tag{7.19}
\end{equation*}
$$

Then Equation (7.16) is equivalent to the Laplace-like equation (7.18) with

$$
\begin{align*}
U & =\left(\hat{L}_{1}^{(y)}\right)^{-1} \hat{L}_{1}^{(x)}, \quad V=\left(\hat{L}_{2}^{(x)}\right)^{-1} \hat{L}_{2}^{(y)}, \quad W=\left(\hat{L}_{3}^{(x)}\right)^{-1} \hat{L}_{3}^{(z)}, \\
\breve{\mathcal{F}} & =\hat{\mathcal{F}} \times{ }_{1}\left(\hat{L}_{1}^{(y)}\right)^{-1} \times_{2}\left(\hat{L}_{2}^{(x)}\right)^{-1} \times_{3}\left(\hat{L}_{3}^{(x)}\right)^{-1} . \tag{7.20}
\end{align*}
$$

To solve Laplace-like equations (7.18), we apply the recursive blocked algorithm developed
in [69]. It transforms the matrices $U, V, W$ into quasi-triangular form by computing Schur decompositions. Block decompositions for the quasi-triangular matrices reveal an equivalent system of Laplace-like equations with smaller matrices. We apply this observation recursively, until we can solve the small Laplace-like equations efficiently by reshaping. This yields the blocked recursive algorithm. For $n=n_{1}=n_{2}=n_{3}$ this approach has a theoretical runtime of $\mathcal{O}\left(n^{4}\right)$ operations. In Section 7.4.1, we demonstrate that this recursive blocked algorithm is much faster than directly reshaping the tensorvalued linear system.

Remark 7.4. For general PDEs with the same type of boundary conditions in each mode, we can obtain a CP decomposition of the form (7.19) if the differential operator does not contain mixed derivatives and if all coefficients $\alpha_{a b c}$ only depend on the variable corresponding to the mode in which their corresponding derivative acts. This includes, for instance, differential operators of the form

$$
\alpha_{1}(x)+\alpha_{2}(x) \frac{\partial}{\partial x}+\alpha_{3}(x) \frac{\partial^{2}}{\partial x^{2}}+\alpha_{4}(y) \frac{\partial}{\partial y}+\alpha_{5}(y) \frac{\partial^{2}}{\partial y^{2}}+\alpha_{6}(z) \frac{\partial}{\partial z}+\alpha_{7}(z) \frac{\partial^{2}}{\partial z^{2}},
$$

with univariate coefficient functions $\alpha_{i}:[-1,1] \rightarrow \mathbb{R}$.
Remark 7.5. So far, we introduced the global spectral method for fixed polynomial degrees $\left(n_{1}, n_{2}, n_{3}\right)$. In order to heuristically determine if the solution is accurate, we can analyze the residual of (7.10) and the decay of the coefficients in $\mathcal{U}$ [16]. This can be used to adaptively increase $\left(n_{1}, n_{2}, n_{3}\right)$ until the solution is accurate.
Remark 7.6. In general, it is not recommended to explicitly invert the matrices in (7.20). The inversion can be avoided completely by extending the recursive solver in [69] similar to how recursive solvers for Sylvester equations can be extended to solve generalized Sylvester equations.

### 7.4 Numerical results

All numerical experiments* in this section were performed in MATLAB R2018b on a Lenovo Thinkpad T480s with Intel Core i7-8650U CPU and 15.4 GiB RAM.

### 7.4.1 Runtime comparison

We consider Poissons's equation $\Delta u=f$ with homogeneous zero Dirichlet boundary conditions. In Table 7.1, we compare our global spectral method to the nested alternating direction implicit method (NADIM) proposed in [125, Section 5]. NADIM relies on solving two-dimensional Sylvester equations recursively on three levels. On each level an iterative algorithm is used, which leads to a large total number of iterations.

[^7]|  | $n=10$ |  | $n=30$ |  | $n=50$ |  | $n=150$ |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
|  | Time | Error | Time | Error | Time | Error | Time | Error |
| NADIM | 21.8 | $1.47 \cdot 10^{-5}$ | 742 | $2.80 \cdot 10^{-8}$ | - | - | - | - |
| reshape | 0.019 | $1.55 \cdot 10^{-5}$ | 2.40 | $1.22 \cdot 10^{-15}$ | 102 | $1.22 \cdot 10^{-15}$ | - | - |
| recursive | 0.054 | $1.55 \cdot 10^{-5}$ | 0.10 | $3.12 \cdot 10^{-13}$ | 0.41 | $8.20 \cdot 10^{-13}$ | 10.9 | $1.15 \cdot 10^{-10}$ |

Table 7.1 - Comparison of algorithms to solve Poisson's equation with known solution $u^{*}(x, y, z)=\sin (\pi x) \sin (\pi y) \sin (\pi z)$, from which the right hand side $f$ is explicitly computed. To compute solutions $u \in \mathbb{P}_{n, n, n}$ with our global spectral method, we compare reshaping (7.15) and solving (7.17) with backslash in MATLAB (reshape) to transforming (7.15) into a Laplace-like equation (7.18) and solving with the blocked recursive solver [69] (recursive). Additionally, we compare to NADIM [125]. For various n, we measure the runtime in seconds and we estimate $\left\|u-u^{*}\right\|_{\infty}$ by sampling 1000 random points.

We observe that even though NADIM has an asymptotic runtime of $\mathcal{O}\left(n^{3}(\log (n))^{3}\right)$ [125], it is the slowest algorithm in our setting and can not handle $n>30$ in a reasonable amount of time. The asymptotic runtime of the recursive algorithm is slower with $\mathcal{O}\left(n^{4}\right)$, but in our experiments it is the fastest method and it can handle values of up to $n=150$ in less than 11 seconds. For $n=50$ solving the reshaped system (7.17) with backslash leads to an error of order $10^{-15}$, whereas the recursive algorithm only achieves an error of order $10^{-13}$. While the recursive approach is able to solve much larger systems, it is slightly more sensitive to numerical rounding errors.

### 7.4.2 Stationary problems

## Helmholtz problems

We consider the Helmholtz equation on $[-1,1]^{3}$ with non-homogeneous Dirichlet boundary conditions as in [160, 319], which arises for instance in the context of three-dimensional wave equations in acoustics [336] and seismic-imaging [261]. It is defined as

$$
\Delta u(x, y, z)+\kappa^{2} u(x, y, z)=f(x, y, z),
$$

with coefficient $\kappa \in \mathbb{R}$. For this differential operator, we define the tensor $\mathcal{A}$ as in Section 7.2.2. We observe that a CP decomposition in the form (7.19) is given by

$$
\mathcal{A}=\left(\begin{array}{c}
\kappa^{2}  \tag{7.21}\\
0 \\
1
\end{array}\right) \circ\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right) \circ\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right)+\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right) \circ\left(\begin{array}{l}
0 \\
0 \\
1
\end{array}\right) \circ\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right)+\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right) \circ\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right) \circ\left(\begin{array}{l}
0 \\
0 \\
1
\end{array}\right) .
$$

Hence, we can derive a Laplace-like structure (7.18) for Equation (7.15) and employ the recursive solver.

In the following, we employ the global spectral method for the Helmholtz equation in [56, Section 5.3] given by

$$
\begin{equation*}
\Delta u(x, y, z)+\kappa(x)^{2} u(x, y, z)=f(x, y, z), \tag{7.22}
\end{equation*}
$$

with function $\kappa(x)=\gamma_{1}-\gamma_{2} \cos \left(\pi \gamma_{3} x / 2\right)$ and scalar coefficients $\gamma_{1}, \gamma_{2}, \gamma_{3} \in \mathbb{R}$. The right hand side $f$ and the Dirichlet boundary conditions are computed explicitly from the solution

$$
\begin{equation*}
u^{*}(x, y, z)=\exp \left(-\kappa(x) / \gamma_{3}\right) \cos \left(\pi \gamma_{1} y / 2\right) \cos \left(\pi \gamma_{2} z / 2\right) . \tag{7.23}
\end{equation*}
$$

In order to incorporate the coefficient function $\kappa(x)$, we compute one-dimensional differential operators as in Equation (7.7) from the CP-decomposition (7.21). We then set $\mathcal{L}_{1}^{(x)} u(x)=\kappa(x)^{2} u(x)+\frac{\partial^{2}}{\partial x^{2}} u(x)$ and discretize this operator as described in Section 7.2.3. The resulting equation (7.15) can still be transformed into a Laplace-like equation (7.18) by setting $U, V, W, \breve{\mathcal{F}}$ as defined in (7.20).

In Figure 7.2, we depict how well the solution $u \in \mathbb{P}_{n, n, n}$ of our global spectral method approximates $u^{*}$. We compare to the trivariate Chebyshev interpolant $\tilde{u}$ of $u^{*}$ that is close to the best approximation of $u$ in $\mathbb{P}_{n, n, n}$ and converges quickly [100, 322]. For $n<80$, the solution $u$ and the interpolant $\tilde{u}$ almost coincide. For $n>80$, the error $\left\|u-u^{*}\right\|_{\infty}$ stagnates, whilst the interpolation error $\left\|\tilde{u}-u^{*}\right\|_{\infty}$ continues to decrease further before reaching a plateau close to machine precision. This discrepancy is caused by the recursive blocked solver.


Figure 7.2 - Left: We plot $u^{*}(x, y, z)$ as defined in (7.23) for $\left(\gamma_{1}, \gamma_{2}, \gamma_{3}\right)=(5,3,5)$ and fixed $z=1 / 4$. Right: We solve the Helmholtz equation (7.22) for $\left(\gamma_{1}, \gamma_{2}, \gamma_{3}\right)=(5,3,5)$ with known solution $u^{*}(7.23)$ via our global spectral method to obtain the solution $u \in \mathbb{P}_{n, n, n}$ for various $n$. Further, we compute the trivariate Chebyshev interpolant $\tilde{u} \in \mathbb{P}_{n, n, n}$ of $u^{*}$. For each $n$, we estimate $\left\|u-u^{*}\right\|_{\infty}$ and $\left\|\tilde{u}-u^{*}\right\|_{\infty}$ using 1000 sample points (evaluation error).

Remark 7.7. The derivation of the Laplace-like system from the CP decomposition can be extended from the Helmholtz equation to convection diffusion problems of the form

$$
-\nu \triangle u+\xi^{T} \nabla u=f
$$

with $\nu \in R$ and $\xi \in \mathbb{R}^{3}$ as studied in [24, 70, 340]. For these problems we define the CP-decomposition

$$
\mathcal{A}=\left(\begin{array}{c}
0 \\
\xi_{1} \\
-\nu
\end{array}\right) \circ\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right) \circ\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right)+\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right) \circ\left(\begin{array}{c}
0 \\
\xi_{2} \\
-\nu
\end{array}\right) \circ\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right)+\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right) \circ\left(\begin{array}{l}
1 \\
0 \\
0
\end{array}\right) \circ\left(\begin{array}{c}
0 \\
\xi_{3} \\
-\nu
\end{array}\right) .
$$

## Diffusion problems with separable coefficient

Diffusion problems of the form

$$
\begin{equation*}
-\nabla \cdot(a(x, y, z) \nabla u(x, y, z))=f(x, y, z) \tag{7.24}
\end{equation*}
$$

with a separable coefficient $a(x, y, z)=a_{1}(x) a_{2}(y) a_{3}(z)$ defined by univariate functions $a_{1}, a_{2}, a_{3}:[-1,1] \rightarrow \mathbb{R}$, are not directly given as linear partial differential operator of the form (7.2). We can, however, decompose the three-dimensional differential operator into a sum of three differential operators similar to (7.7) as

$$
\begin{aligned}
& \nabla \cdot(a(x, y, z) \nabla u(x, y, z))= \\
& (\underbrace{\frac{\partial}{\partial x}\left(a(x, y, z) \frac{\partial}{\partial x}\right)}_{=\tilde{\mathcal{L}}_{1}^{(x)}}+\underbrace{\frac{\partial}{\partial y}\left(a(x, y, z) \frac{\partial}{\partial y}\right)}_{=\tilde{\mathcal{L}}_{2}^{(y)}}+\underbrace{\frac{\partial}{\partial z}\left(a(x, y, z) \frac{\partial}{\partial z}\right)}_{=\tilde{\mathcal{L}}_{3}^{(z)}}) u(x, y, z) .
\end{aligned}
$$

Applying the differential operator $\tilde{\mathcal{L}}_{1}^{(x)}$ to a polynomial $u \in \mathbb{P}_{n, n, n}$ can be written analogously to (7.8) as

$$
\mathcal{V}=\mathcal{U} \times{ }_{1} \underbrace{S_{1} D_{1} S_{0}^{-1} M^{(1)}\left[\mathbf{a}_{1}\right] D_{1}}_{=L_{1}^{(x)}} \times 2 \underbrace{M^{(2)}\left[\mathbf{a}_{2}\right] S_{1} S_{0}}_{=L_{1}^{(y)}} \times 3 \underbrace{M^{(2)}\left[\mathbf{a}_{3}\right] S_{1} S_{0}}_{=L_{1}^{(y)}}
$$

where $\mathbf{a}_{1} \in \mathbb{R}^{n_{1}+1}, \mathbf{a}_{2} \in \mathbb{R}^{n_{2}+1}, \mathbf{a}_{3} \in \mathbb{R}^{n_{3}+1}$ denote the coefficients for univariate Chebyshev interpolation of $a_{1}, a_{2}, a_{3}$ as in Section 7.2.3. The multiplication matrices $M[\mathbf{a}]$, differentiation matrices $D_{1}$ and transformation matrices $S_{0}, S_{1}$ are defined as in Section 7.2. We obtain analogous discretizations for $\tilde{\mathcal{L}}_{2}^{(y)}$ and $\tilde{\mathcal{L}}_{3}^{(z)}$. Observe that the discretization has the same symmetric structure as the CP decomposition (7.19). Thus, we can transform (7.15) to a Laplace-like equation (7.18) by defining $U, V, W, \breve{\mathcal{F}}$ as in (7.20) and use the recursive solver.

## Diffusion problems with higher rank coefficient

Most diffusion problems (7.24) arising in the study of groundwater flow and uncertainty quantification $[75,133,220,287,316,328]$ do not have a rank- 1 coefficient $a(x, y, z)$. The coefficient is often given by a truncated expansion as a sum of separable functions of the form

$$
a(x, y, z)=\sum_{r=1}^{R_{a}} a_{1}^{(r)}(x) a_{2}^{(r)}(y) a_{3}^{(r)}(z)
$$

with $R_{a}>1$ and univariate functions $a_{1}^{(r)}, a_{2}^{(r)}, a_{3}^{(r)}:[-1,1] \rightarrow \mathbb{R}$. The PDE (7.24) can now be written as

$$
\begin{equation*}
-\sum_{r=1}^{R_{a}} \nabla \cdot\left(a_{1}^{(r)}(x) a_{2}^{(r)}(y) a_{3}^{(r)}(z) \nabla u(x, y, z)\right)=f(x, y, z) . \tag{7.25}
\end{equation*}
$$

Following the ideas in Section 7.4.2, we can discretize (7.24) for each separable function $a_{1}^{(r)}(x) a_{2}^{(r)}(y) a_{3}^{(r)}(z)$. Adding these discretizations yields a discretization of the form (7.10) with $R=3 R_{a}$. Since $R>3$, we can not find a Laplace-like formulation of (7.15) and we can not use the recursive solver.

We can, however, use preconditioned GMRES [137] to compute solutions of (7.16) seen as tensor-valued linear system. Throughout this work, we restart GMRES every 15 iterations. As preconditoner we employ the recursive solver to solve Equation (7.15) for a discretization of the same diffusion problem (7.25) with the coefficient $a(x, y, z)$ replaced by a separable coefficient $b(x, y, z)$. This can be seen as effectively solving the system

$$
\begin{aligned}
\left(\sum_{r=1}^{3} \hat{L}_{r}^{(z)}[b] \otimes \hat{L}_{r}^{(y)}[b] \otimes \hat{L}_{r}^{(x)}[b]\right)^{-1}( & \left.\sum_{r=1}^{R} \hat{L}_{r}^{(z)}[a] \otimes \hat{L}_{r}^{(y)}[a] \otimes \hat{L}_{r}^{(x)}[a]\right) \operatorname{vec}\left(\mathcal{U}^{222}\right)= \\
& \left(\sum_{r=1}^{3} \hat{L}_{r}^{(z)}[b] \otimes \hat{L}_{r}^{(y)}[b] \otimes \hat{L}_{r}^{(x)}[b]\right)^{-1} \operatorname{vec}(\hat{\mathcal{F}}),
\end{aligned}
$$

where $\hat{L}_{r}^{(x)}[a], \hat{L}_{r}^{(y)}[a], \hat{L}_{r}^{(z)}[a]$ and $\hat{L}_{r}^{(x)}[b], \hat{L}_{r}^{(y)}[b], \hat{L}_{r}^{(z)}[b]$ denote the matrices in (7.16) based on discretizations of the PDE with coefficient $a$ and $b$ respectively. The application of the inverse of $\sum_{r=1}^{3} \hat{L}_{r}^{(z)}[b] \otimes \hat{L}_{r}^{(y)}[b] \otimes \hat{L}_{r}^{(x)}[b]$ can be computed by solving a Laplace-like equation (7.18).

From now on, we consider the rank-2 coefficient $a(x, y, z)=\left(1+x^{2}\right)\left(1+y^{2}\right)\left(1+z^{2}\right)+$ $\exp (x+y+z)$. We use our global spectral method with $n=n_{1}=n_{2}=n_{3}=30$ to solve the diffusion problem (7.25) with known solution $u^{*}(x, y, z)=\sin (\pi x) \sin (\pi y) \sin (\pi z)$, from which we explicitly compute the right hand side $f$ and Dirichlet boundary conditions. In Figure 7.3, we display the convergence rates for GMRES with preconditioners based on the constant coefficient $b_{1}(x, y, z)=\|a\|_{\mathcal{L}^{2}}$ and the separable coefficient $b_{2}(x, y, z)=$


Figure 7.3 - Convergence rate of preconditioned GMRES for solving Equation (7.15) for the Diffusion problem (7.25) with coefficient $a(x, y, z)=\left(1+x^{2}\right)\left(1+y^{2}\right)\left(1+z^{2}\right)+\exp (x+y+z)$. We compare no preconditioning and preconditioning with a separable coefficient $b(x, y, z)$ as described in Section 7.4.2.
$\left(1+x^{2}\right)\left(1+y^{2}\right)\left(1+z^{2}\right)$. Both preconditioners yield solutions $u$ such that the error $\left\|u-u^{*}\right\|_{\infty}$ is close to machine precision. Solving with $b_{1}$ takes 4.87 seconds. For the separable coefficient $b_{2}$, fewer iterations are necessary and the computation only takes 2.14 seconds. For comparison, reshaping as in Section 7.4.1 would take 107 seconds.

## Helmholtz equation with non-constant coefficients

Next, we study a variable coefficient Helmholtz equation as in [319, Example 2]. We consider the PDE

$$
\begin{equation*}
\Delta u(x, y, z)+\kappa(x, y, z) u(x, y, z)=f(x, y, z), \tag{7.26}
\end{equation*}
$$

with function $\kappa(x, y, z)=\sqrt{x+y+z+42}$. For $n=n_{1}=n_{2}=n_{3}=30$, we compute a discretization (7.10) by computing an approximate CP decomposition of $\mathcal{A}$ with $R=10$ as defined as in Section 7.2.3 using tensorlab [332]. In 2.08 seconds, we obtain a CP decomposition with error $2.04 \cdot 10^{-9}$ in the uniform norm. This leads to a discretized PDE, for which we solve Equation (7.15) with preconditioned GMRES with restarting as in Section (7.4.2). In the preconditioner we solve Equation (7.15) for the Helmholtz equation (7.22) with constant coefficient $\|\kappa\|_{\mathcal{L}^{2}}$.

When the right hand side and the Dirichlet boundary conditions are chosen to match the known solution $u^{*}(x, y, z)=\sin (\pi x) \sin (\pi y) \sin (\pi z)$, solving with preconditioned GMRES takes 1.04 seconds. The computed solution $u$ satisfies $\left\|u-u^{*}\right\|_{\infty} \approx 2.38 \cdot 10^{-10}$, where we estimate the uniform norm using 1000 sample points.

Remark 7.8. To obtain very high accuracy in the PDE solution, we would need a very accurate CP decomposition. The efficient computation of accurate $C P$ decompositions is


Figure 7.4 - Left: Let $u[n] \in \mathbb{P}_{n, n, n}$ denote the solution computed for the Helmholtz equation described in Section 7.4.2. We plot $u[45](x, y, 1 / 4)$. Right: For various $n$ we plot $\|u[n]-u[45]\|_{\infty}$ estimated at 1000 sample points. Additionally, we plot the combined residual defined as maximum of the residual of (7.10) and the residuals of (7.11) in the uniform norm. Further, we plot the error of the CP approximation of $\mathcal{A}$ in the uniform norm.
still subject to research [293]. Here, we could avoid computing a CP decompositoin of $\mathcal{A} \in \mathbb{R}^{93 \times 93 \times 93}$ by instead approximating the operators $\Delta$ and $\kappa(x, y, z) \mathcal{I}$ separately, where $\mathcal{I}$ denotes the identity operator. Discretizing $\kappa(x, y, z) \mathcal{I}$ only involves a CP decomposition of a tensor in $\mathbb{R}^{31 \times 31 \times 31}$, which can be computed for $R=7$ in only 0.09 seconds with error $1.93 \cdot 10^{-11}$ in the uniform norm. The two resulting discretizations in the form of (7.9) can be added to obtain a discretization of $\Delta+\kappa(x, y, z) \mathcal{I}$.

## Helmholtz equation with unknown solution

As a final stationary problem, we consider the Helmholtz equation (7.26) with nonconstant coefficient $\kappa(x, y, z)=\sqrt{x+y+z+42}$ and $f(x, y, z)=1$. We use mixed boundary conditions with zero Neumann boundary conditions on the right and zero Dirichlet boundary conditions on all other sides as in Section 7.3.1. As in Section 7.4.2, we compute a CP decomposition of $\mathcal{A}$ and solve using preconditioned GMRES. In Figure 7.4, we display the computed solution and the error decay for different polynomial degrees. We observe that the residual decays when the polynomial degree is increased, which indicates that the solutions become more accurate.

### 7.4.3 Time-dependent problems

In this section, we introduce an implicit Euler scheme to solve parabolic PDEs of the form

$$
\frac{\partial}{\partial t} u(x, y, z, t)+\mathcal{L} u(x, y, z, t)=0
$$

where $\mathcal{L}$ is a linear partial differential operator acting only on the spatial variables $x, y, z$. The system is complemented with boundary conditions. We are interested in the time evolution starting from a given initial function $u(x, y, z, 0)=u_{0}(x, y, z)$. We discretize the equation in time using the uniform step length $h$. For each $\tau=1,2, \ldots$, we compute $u_{\tau}(x, y, z) \approx u(x, y, z, \tau h)$ as solution of the stationary linear PDE

$$
\begin{equation*}
(\mathcal{I}-h \mathcal{L}) u_{\tau+1}=u_{\tau} \tag{7.27}
\end{equation*}
$$

We approximate each function $u_{\tau}$ by a polynomial of the form (7.4) represented by the coefficient tensor $\mathcal{U}_{\tau}$. The initial $\mathcal{U}_{0}$ is computed via tensorized Chebyshev interpolation. For $\tau=1,2, \ldots$, we obtain $\mathcal{U}_{\tau}$ by applying our spectral method to solve Equation (7.27). We discretize the operator $(\mathcal{I}-h \mathcal{L})$ directly like the Helmholtz equation in Section 7.4.2. This allows us to employ the recursive solver. Note that the right hand side is represented in terms of an ultraspherical basis. Hence, the computation of $\mathcal{U}_{\tau}$ requires multiplying $\mathcal{U}_{\tau-1}$ with appropriate basis transformation matrices.

We demonstrate this implicit Euler scheme for the parabolic PDE studied in [257, Section 6.1]. The function $u^{*}(x, y, z, t)=\exp \left(-3 \pi^{2} t\right) \sin (\pi x) \sin (\pi y) \sin (\pi z)$ satisfies the parabolic PDE

$$
\frac{\partial}{\partial t} u+\Delta u=0
$$

on the domain $[-1,1]^{3}$ with homogeneous Dirichlet boundary conditions. We use an implicit Euler scheme for $u_{0}(x, y, z)=\sin (\pi x) \sin (\pi y) \sin (\pi z)$ and compare our global spectral method to a finite difference method [303]. In each timestep of the finite difference method, we solve a linear system with a sparse Kronecker-structured finite difference matrix using the backslash operator in MATLAB.

The time evolution of the errors is displayed in Figure 7.5. We observe that for $n=20$ and $h=10^{-2}$ both approximations lead to very similar errors, but the computation time for 50 implicit Euler steps decreases from 6.22 seconds for the finite difference method to 0.93 seconds for our global spectral method. In this case the error is dominated by the implicit Euler scheme for both methods. In contrast, for $n=30$ and $h=10^{-4}$, the errors for the spectral method are smaller. For large $\tau$ the error of both approaches is dominated by the time discretization via the implicit Euler scheme. However, in the initial 750 time steps the error of the spatial discretization dominates for the finite difference scheme,
whereas the global spectral method is able to represent $u_{\tau}$ accurately.


Figure 7.5 - Error plots for the parabolic PDE example in Section 7.4.3. We apply an implicit Euler scheme to compute approximations $u_{\tau}(x, y, z) \approx u(x, y, z, \tau h)$ and compare the solutions for the global spectral method (GSM) with solutions in $\mathbb{P}_{n, n, n}$ and a finite difference discretization (FD) on a regular $(n+1) \times(n+1) \times(n+1)$ grid. The absolute error $\left\|u^{*}(\cdot, \cdot, \cdot, \tau h)-u_{\tau}(\cdot, \cdot, \cdot)\right\|_{\infty}$ is estimated in each timestep from evaluations at 100 sample points. The relative error is computed from the estimation of the absolute error.

Remark 7.9. Instead of using the implicit Euler scheme, we could generalize the ideas from [83] to derive a tensor-oriented exponential Euler scheme.

Remark 7.10. The methods presented in this work can also be used to solve two-dimensional parabolic PDEs on rectangles by treating time as third space variable in the discretization of the operator. We refer to [319] for more details on this approach.

### 7.4.4 Eigenvalue problems

The methods presented in this work can be extended to solve PDE eigenvalue problems, in which we search eigenvalues $\lambda$ and eigenfunctions $u$ satisfying the equation $\mathcal{L} u=\lambda u$ complemented with homogeneous Dirichlet boundary conditions. We are particularly interested in finding the eigenvalue with minimal absolute value. For this purpose we employ the inverse iteration algorithm [137]. Starting from an initial function $u_{0}$ we iteratively compute an approximation of the eigenfunction. For $s=1,2, \ldots$, we compute $u_{s}$ as solution of the PDE

$$
\begin{equation*}
\mathcal{L} u_{s}=\frac{u_{s-1}}{\left\|u_{s-1}\right\|_{\mathcal{L}^{2}}} \tag{7.28}
\end{equation*}
$$

We approximate the eigenvalue using the Rayleigh quotient $\frac{1}{\lambda} \approx \frac{\left\langle u_{s-1}, u_{s}\right\rangle}{\left\langle u_{s-1}, u_{s-1}\right\rangle}$, where $\langle\cdot, \cdot\rangle$ denotes the standard $\mathcal{L}^{2}$ scalar product. We again proceed by discretizing the differential
operator and the function $u$ to solve Equation (7.28) using the spectral method introduced in Section 7.3.

Let the functions $u, v \in \mathbb{P}_{n_{1}, n_{2}, n_{3}}$ be given in the form of (7.4). We evaluate the norm as $\|u\|_{\mathcal{L}^{2}}^{2}=\langle u, u\rangle$ and the scalar product $\langle u, v\rangle$ by interpolating the function $u v$ using tensorized Chebyshev polynomial basis functions as in [100, 224]. In a second step, we integrate the approximation of $u v$ using the exact values for the integrals of the basis functions [322, Theorem 19.2].

We test our method for the elliptic PDE eigenvalue problem

$$
\begin{equation*}
-\Delta u(x, y, z)+v(x, y, z) u(x, y, z)=\lambda u(x, y, z), \tag{7.29}
\end{equation*}
$$

with potential $v(x, y, z)=\sin (\pi / 2(x+1)) \sin (\pi / 2(y+1)) \sin (\pi / 2(z+1))$ as in [154, 208]. Due to the potential $v(x, y, z)$ the inverse iteration steps (7.28) are Helmholtz equations (7.22) with non-constant separable coefficients. We discretize these by adding a discretization of $v(x, y, z) \mathcal{I}$ with $R=1$ to a discretization of the Laplacian. The resulting discretized PDE (7.10) has $R=4$ and we apply preconditioned GMRES with restarting using a Helmholtz equation with constant coefficient $\|v\|_{\mathcal{L}^{2}}$ to compute solutions as in Section 7.4.2. We compare our global spectral method to a finite difference scheme on a regular grid for solving (7.28). In Figure 7.6 we observe that the eigenvalues computed as Rayleigh coefficients converge at a faster rate for the global spectral method than the finite difference approach. Already for $n=20$ the global spectral method reaches an error close to machine precision.


Figure 7.6 - Eigenvalue convergence rate for the PDE eigenvalue problem (7.29) computed as Rayleigh coefficients from $u_{50}$ in the inverse iteration method, where the solutions $u_{s}$ are computed via the global spectral method with $u_{s} \in \mathbb{P}_{n, n, n}$ and a finite difference method on a regular $(n+1) \times(n+1) \times(n+1)$ grid. We denote the eigenvalue obtained for $n$ by $\lambda[n]$. Since the exact $\lambda$ is not known, we plot $|\lambda[n]-\lambda[45]|$ for both algorithms, where $\lambda[45]$ is computed with the respective algorithm.

Remark 7.11. We would like to point out that our approach can be used to compute a
basis for the (continuous) Krylov subspace, which is used in the Arnoldi method [15, 137] and in the Least-squares spectral method in [162].

## 8 Conclusions and outlook

In this thesis, we demonstrated that three very different high-dimensional problems can be solved efficiently by identifying and exploiting underlying low-rank tensor structures to overcome the curse of dimensionality.

Functional low-rank approximation In Chapters 3 and 4, we developed two novel algorithms to compute functional low-rank approximations of multivariate functions based on selected function evaluations. The key point in deriving these algorithms has been that low-rank approximations of the evaluation tensor for tensorized interpolation lead to particular types of functional low-rank approximations. Compared to previously existing algorithms, we managed to significantly reduce the number of function evaluations needed to obtain an approximation with prescribed accuracy for most functions.

Throughout this thesis, our approximation algorithms have been derived based on expansions in terms of tensorized polynomials. Extensions to other basis functions are possible and might offer certain advantages. For example, Ballfun [43] avoids the introduction of artificial boundaries when approximating functions on the ball by combining a mix of Chebyshev and Fourier basis functions with additional symmetry constraints. So far, we have focused on approximations in the functional Tucker or the (extended) functional TT format. While we have demonstrated that many functions can be well approximated in these formats, it might be possible to compress certain functions even more when using other functional low-rank approximation formats [21]. To obtain such approximations efficiently, we will need to develop specialized approximation algorithms. Those algorithms potentially could reuse parts of the algorithms presented in this work, such as first computing a Tucker approximation using Algorithm 9 before approximating the core tensor in the desired format. Future work could additionally cover how to compute numerically with functions [323] once they are compressed in functional low-rank formats. For arithmetic operations, one could perform certain operations by directly manipulating the low-rank approximation of the evaluation and the coefficient tensor. For instance,
the multiplication of functions could be seen as Hadamard product of the corresponding low-rank approximations of the evaluation tensor [204].

In Chapter 7, we derived a global spectral method for solving three-dimensional linear PDEs on cubes with very high accuracy. Our numerical experiments demonstrated that using the blocked recursive solver [69] to either solve the resulting equations directly or as a preconditioner when the equations are not Laplace-like vastly outperforms all existing methods. We additionally presented the versatility of our method by the extension to eigenvalue and time-dependent problems.

In future work, additional features included in Chebop2 [319] such as adaptivity and input parsing could be incorporated into the current implementation of our global spectral method. Currently, the computational complexity of our approach is heavily influenced by the storage needed to store the full coefficient tensors $\mathcal{U}, \mathcal{F}$ and $\mathcal{G}_{\ell}$ for representing the solution, the right-hand side and the boundary conditions, respectively. Under the assumption that both the solution of the PDE and the right-hand side can be well approximated in functional low-rank formats, also these coefficient tensors can be well approximated in the corresponding low-rank tensor formats. Such approximations have the potential to drastically reduce the storage complexity. Exploiting this potential requires determining suitable approximation formats and using a specialized solver for the chosen formats. For instance, when choosing the Tucker or the TT format, we can use the ideas from [289] to solve Laplace-like equations directly in these low-rank formats. Alternatively, one could add the boundary conditions as penalty terms to the discretized PDE as in [162] and apply standard techniques such as Riemannian optimization [264] or alternating linear schemes [171] to solve the resulting system in the TT format. The computational complexity of time-dependent problems could be reduced using rank-adaptive dynamical low-rank approximations [62, 64, 90]. We also would like to point out that extending our global spectral method to higher dimensional PDEs on hypercubes is straightforward, but it might no longer be feasible to solve the resulting system numerically without exploiting additional low-rank structure of the coefficient tensors.

Multi-marginal optimal transport In Chapter 5, we analyzed the impact of approximating the Gibbs kernel on the Sinkhorn algorithm for solving entropically regularized multi-marginal optimal transport problems. We demonstrated that the introduction of low-rank approximations offers the potential to reduce the computational complexity of evaluating marginals by orders of magnitude. For transport plans defined via graphical models, we showed that constructing the low-rank approximation based on the dual graphical model can be faster and more accurate than the direct approximation of the Gibbs kernel. An application of our method is presented by the drastic reduction of the computation time for transferring colors from several images onto one target image.

So far, we computed the approximation $\tilde{\mathcal{K}}$ of the Gibbs kernel $\mathcal{K}$ using classical lowrank approximation algorithms, which are designed to minimize $\|\mathcal{K}-\tilde{\mathcal{K}}\|_{\infty}$ with respect to the uniform or the Frobenius norm. However, our error bounds indicate that we are actually interested in finding potentially very different approximations for which $\|\log (\mathcal{K})-\log (\tilde{\mathcal{K}})\|_{\infty}$ is minimized. Efficient algorithms to compute such approximations are still subject to future work. The design of these algorithms can potentially be inspired by [172, 330], where approximations in the CP format are computed such that non-standard objective functions are minimized. In [112], the related problem of finding low-rank approximations of elementwise exponentials is studied in the context of lognormal random fields. For multi-marginal optimal transport problems arising in the context of density functional theory [35], several obstacles need to be overcome before our approach can be put into practice. The Coulomb cost leads to zero entries on the diagonal of the Gibbs kernel. Any suitable approximation of the Gibbs kernel would need to approximate these entries exactly, which is usually not feasible without either adding sparse correction terms to the low-rank approximation as for instance in [213, 275] or generalizing the idea of using hierarchical matrices [234] to tensors. Moreover, the underlying graphical model leads to a fully connected tensor network that can not be contracted efficiently.

Self-diffusion matrix In Chapter 6, we introduced a novel approach to compute the self-diffusion matrix of a tagged particle process. We exploited the fact that the latter quantity can be expressed via the solution of a high-dimensional deterministic minimization problem to develop an alternating optimization scheme to build a low-rank approximation of this solution. Our numerical experiments showed that the variance in the computed self-diffusion matrix is much smaller when using our novel approach compared to using standard sampling based methods with the same amount of computational resources.

There still remain several interesting open questions that need to be investigated in future work. As already described in Section 6.3.5, the current implementation of our low-rank algorithm is limited in terms of the size of the computational domain for which the algorithm can be run due to round-off errors. We still believe that our approach is promising since the error linked to the truncation of the computational domain is small in comparison to statistical errors. In this thesis, all our experiments have been concerned with two-dimensional tagged particle systems. Extensions to higher dimensional-systems are possible, but the rounding issues become more and more pressing. So far, we only derived a finite volume scheme to solve a simplified cross-diffusion system involving the computed self-diffusion matrices. It still remains to develop a convergent entropy diminishing finite volume scheme for hydrodynamic limits of multi-species exclusion processes.

## Bibliography

[1] L. Agelas, R. Eymard, and R. Herbin. A nine-point finite volume scheme for the simulation of diffusion in heterogeneous media. C. R. Math. Acad. Sci. Paris 347.11-12 (2009), pp. 673-676.
[2] M. Agueh and G. Carlier. Barycenters in the Wasserstein space. SIAM J. Math. Anal. 43.2 (2011), pp. 904-924.
[3] S. Ahmadi-Asl et al. Cross tensor approximation methods for compression and dimensionality reduction. IEEE Access 9 (2021), pp. 150809-150838.
[4] M. Ali and A. Nouy. Approximation with tensor networks. Part I: Approximation spaces. arXiv e-prints (2020), arXiv:2007.00118.
[5] M. Ali and A. Nouy. Approximation with tensor networks. Part II: Approximation rates for smoothness classes. arXiv e-prints (2020), arXiv:2007.00128.
[6] M. Ali and A. Nouy. Approximation with tensor networks. Part III: Multivariate approximation. arXiv e-prints (2021), arXiv:2101.11932.
[7] J. M. Altschuler and E. Boix-Adsera. Polynomial-time algorithms for multimarginal optimal transport problems with structure. Math. Program. (2022), pp. 172.
[8] J. M. Altschuler and E. Boix-Adserà. Hardness results for multimarginal optimal transport problems. Discrete Optim. 42 (2021), Paper No. 100669, 21.
[9] J. Altschuler, J. Weed, and P. Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In: Adv. Neural Inf. Process. Syst. 30. 2017, pp. 1964-1974.
[10] J. Altschuler et al. Massively scalable Sinkhorn distances via the Nyström method. In: Adv. Neural Inf. Process. Syst. 32. 2019, pp. 4427-4437.
[11] A. Ammar et al. Proper generalized decomposition of time-multiscale models. Internat. J. Numer. Methods Eng. 90.5 (2012), pp. 569-596.
[12] J. An and A. Owen. Quasi-regression. J. Complexity 17.4 (2001), pp. 588-607.
[13] C. Arita, P. L. Krapivsky, and K. Mallick. Bulk diffusion in a kinetically constrained lattice gas. J. Phys. A 51 (2018), p. 125002.
[14] C. Arita, P. L. Krapivsky, and K. Mallick. Variational calculation of transport coefficients in diffusive lattice gases. Phys. Rev. E 95 (2017), p. 032121.
[15] W. E. Arnoldi. The principle of minimized iteration in the solution of the matrix eigenvalue problem. Quart. Appl. Math. 9 (1951), pp. 17-29.
[16] J. L. Aurentz and L. N. Trefethen. Chopping a Chebyshev series. ACM Trans. Math. Software 43.4 (2017), pp. 1-21.
[17] F. A. Ba and M. Quellmalz. Accelerating the Sinkhorn algorithm for sparse multi-marginal optimal transport via fast Fourier transforms. Algorithms 15.9 (2022).
[18] I. Babuška and M. Suri. The p-and h-p versions of the finite element method, an overview. Comput. Methods Appl. Mech. Eng. 80 (1990), pp. 5-26.
[19] F. BACH. Breaking the curse of dimensionality with convex neural networks. J. Mach. Learn. Res. 18 (2017), pp. 1-53.
[20] M. Bachmayr and A. Cohen. Kolmogorov widths and low-rank approximations of parametric elliptic PDEs. Math. Comp. 86.304 (2017), pp. 701-724.
[21] M. Bachmayr, A. Nouy, and R. Schneider. Approximation by tree tensor networks in high dimensions: Sobolev and compositional functions. arXiv e-prints (2021), arXiv:2112.01474.
[22] M. Bachmayr, R. Schneider, and A. Uschmajew. Tensor networks and hierarchical tensors for the solution of high-dimensional partial differential equations. Found. Comput. Math. 16.6 (2016), pp. 1423-1472.
[23] A. Bakhta and V. Ehrlacher. Cross-diffusion systems with non-zero flux and moving boundary conditions. ESAIM Math. Model. Numer. Anal. 52.4 (2018), pp. 1385-1415.
[24] J. Ballani and L. Grasedyck. A projection method to solve linear systems in tensor format. Numer. Linear Algebra Appl. 20.1 (2013), pp. 27-43.
[25] R. Ballester-Ripoll, E. G. Paredes, and R. Pajarola. Sobol tensor trains for global sensitivity analysis. Reliab. Eng. Syst. Saf. 183 (2019), pp. 311-322.
[26] R. Bapat. $D_{1} A D_{2}$ theorems for multidimensional matrices. Linear Algebra Appl. 48 (1982), pp. 437-442.
[27] Z. Battles and L. N. Trefethen. An extension of MATLAB to continuous functions and operators. SIAM J. Sci. Comput. 25.5 (2004), pp. 1743-1770.
[28] Y. Bazilevs et al. 3D simulation of wind turbine rotors at full scale. Part II: Fluid-structure interaction modeling with composite blades. Int. J. Numer. Meth. Fluids 65.1-3 (2010), pp. 236-253.
[29] M. Bebendorf and S. RJasanow. Adaptive low-rank approximation of collocation matrices. Computing 70.1 (2003), pp. 1-24.
[30] M. Bebendorf. Approximation of boundary element matrices. Numer. Math. 86.4 (2000), pp. 565-589.
[31] M. Bebendorf. Hierarchical Matrices. Vol. 63. Lect. Notes Comput. Sci. Eng. Springer, 2008.
[32] F. Beier et al. Unbalanced multi-marginal optimal transport. J. Math. Imaging Vis. (2021), pp. 1-20.
[33] R. Bellman. Dynamic Programming. Princeton University Press, 1957.
[34] J.-D. Benamou. Optimal transportation, modelling and numerical simulation. Acta Numer. 30 (2021), pp. 249-325.
[35] J.-D. Benamou, G. Carlier, and L. Nenna. A numerical method to solve multimarginal optimal transport problems with Coulomb cost. In: Splitting Methods in Communication, Imaging, Science, and Engineering. Springer, 2016, pp. 577-601.
[36] J.-D. Benamou, G. Carlier, and L. Nenna. Generalized incompressible flows, multi-marginal transport and Sinkhorn algorithm. Numer. Math. 142.1 (2019), pp. 33-54.
[37] J.-D. Benamou et al. Iterative Bregman projections for regularized transportation problems. SIAM J. Sci. Comput. 37.2 (2015), A1111-A1138.
[38] J.-P. Berrut and L. N. Trefethen. Barycentric Lagrange interpolation. SIAM Rev. 46.3 (2004), pp. 501-517.
[39] G. Beylkin and M. J. Mohlenkamp. Algorithms for numerical analysis in high dimensions. SIAM J. Sci. Comput. 26.6 (2005), pp. 2133-2159.
[40] G. Beylkin and M. J. Mohlenkamp. Numerical operator calculus in higher dimensions. Proc. Natl. Acad. Sci. USA 99.16 (2002), pp. 10246-10251.
[41] D. Bigoni, A. P. Engsig-Karup, and Y. M. Marzouk. Spectral tensor-train decomposition. SIAM J. Sci. Comput. 38.4 (2016), A2405-A2439.
[42] O. Blondel and C. Toninelli. Kinetically constrained lattice gases: Tagged particle diffusion. Ann. Inst. Henri Poincaré Probab. Stat. 54.4 (2018), pp. 23352348.
[43] N. Boullé and A. Townsend. Computing with functions in the ball. SIAM J. Sci. Comput. 42.4 (2020), pp. C169-C191.
[44] N. Boumal. An Introduction to Optimization on Smooth Manifolds. Cambridge University Press, 2023.
[45] C. Boutsidis and D. P. Woodruff. Optimal CUR matrix decompositions. SIAM J. Comput. 46.2 (2017), pp. 543-589.
[46] J. P. Boyd and R. Petschek. The relationships between Chebyshev, Legendre and Jacobi polynomials: The generic superiority of Chebyshev polynomials and three important exceptions. J. Sci. Comput. 59.1 (2014), pp. 1-27.
[47] D. Braess and W. Hackbusch. Approximation of $1 / x$ by exponential sums in $[1, \infty)$. IMA J. Numer. Anal. 25.4 (2005), pp. 685-697.
[48] S. C. Brenner and L. R. Scott. The Mathematical Theory of Finite Element Methods. 3rd. Vol. 15. Texts Appl. Math. Springer, 2008.
[49] M. Bruna and S. J. Chapman. Diffusion of multiple species with excluded-volume effects. J. Chem. Phys. 137.20 (2012), p. 204116.
[50] H.-J. Bungartz and M. Griebel. Sparse grids. Acta Numer. 13 (2004), pp. 147269.
[51] M. Burger et al. Nonlinear cross-diffusion with size exclusion. SIAM J. Math. Anal. 42.6 (2010), pp. 2842-2871.
[52] H. CAI et al. Fast robust tensor principal component analysis via fiber CUR decomposition. In: IEEE/CVF Int. Conf. Comput. Vis. Workshop. 2021, pp. 189197.
[53] H. Cai et al. Mode-wise tensor decompositions: Multi-dimensional generalizations of CUR decompositions. J. Mach. Learn. Res. 22 (2021), Paper No. 185, 36.
[54] C. F. Caiafa and A. Cichocki. Generalizing the column-row matrix decomposition to multi-way arrays. Linear Algebra Appl. 433.3 (2010), pp. 557-573.
[55] E. J. CANDÈS and T. TAO. The power of convex relaxation: Near-optimal matrix completion. IEEE Trans. Inform. Theory 56.5 (2010), pp. 2053-2080.
[56] L. F. Canino et al. Numerical solution of the Helmholtz equation in 2D and 3D using a high-order Nyström discretization. J. Comput. Phys. 146.2 (1998), pp. 627-663.
[57] J. Cao et al. Multi-marginal wasserstein GAN. In: Adv. Neural Inf. Process. Syst. 32. 2019, pp. 1776-1786.
[58] G. Carlier. On the linear convergence of the multimarginal Sinkhorn algorithm. SIAM J. Optim. 32.2 (2022), pp. 786-794.
[59] G. Carlier, A. Oberman, and E. Oudet. Numerical methods for matching for teams and Wasserstein barycenters. ESAIM Math. Model. Numer. Anal. 49.6 (2015), pp. 1621-1642.
[60] J. D. Carroll and J.-J. Chang. Analysis of individual differences in multidimensional scaling via an N-way generalization of "Eckart-Young" decomposition. Psychometrika 35.3 (1970), pp. 283-319.
[61] J. D. Carroll, S. Pruzansky, and J. B. Kruskal. Candelinc: A general approach to multidimensional analysis of many-way arrays with linear constraints on parameters. Psychometrika 45.1 (1980), pp. 3-24.
[62] G. Ceruti, J. Kusch, and C. Lubich. A rank-adaptive robust integrator for dynamical low-rank approximation. BIT 62.4 (2022), pp. 1149-1174.
[63] G. Ceruti and C. Lubich. An unconventional robust integrator for dynamical low-rank approximation. BIT 62.1 (2022), pp. 23-44.
[64] G. Ceruti, C. Lubich, and D. Sulz. Rank-adaptive time integration of tree tensor networks. arXiv e-prints (2022), arXiv:2201.10291.
[65] S. Chaturantabut and D. C. Sorensen. Discrete empirical interpolation for nonlinear model reduction. In: 48h IEEE Conf. Decis. Control and 28th Chin. Control Conf. CCC. 2009, pp. 4316-4321.
[66] S. Chaturantabut and D. C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. SIAM J. Sci. Comput. 32.5 (2010), pp. 2737-2764.
[67] P. L. Chebyshev. Théorie des Mécanismes Connus sous le Nom de Parallélogrammes. Imprimerie de l'Académie impériale des sciences, 1853.
[68] D. CHEN et al. Limit theorems for the tagged particle in exclusion processes on regular trees. Electron. Commun. Probab. 24 (2019), Paper No. 2, 10.
[69] M. Chen and D. Kressner. Recursive blocked algorithms for linear systems with Kronecker product structure. Numer. Algorithms 84.3 (2020), pp. 1199-1216.
[70] Z. Chen and L. Lu. A projection method and Kronecker product preconditioner for solving Sylvester tensor equations. Sci. China Math. 55.6 (2012), pp. 1281-1292.
[71] A. Chertkov, G. Ryzhakov, and I. Oseledets. Black box approximation in the tensor train format initialized by ANOVA decomposition. arXiv e-prints (2022), arXiv:2208.03380.
[72] A. Cichocki et al. Tensor decompositions for signal processing applications: From two-way to multiway component analysis. IEEE Signal Process. Mag. 32.2 (2015), pp. 145-163.
[73] A. Cichocki et al. Tensor networks for dimensionality reduction and large-scale optimization: Part 1 low-rank tensor decompositions. Found. Trends Mach. Learn. 9.4-5 (2016), pp. 249-429.
[74] C. W. Clenshaw and A. R. Curtis. A method for numerical integration on an automatic computer. Numer. Math. 2 (1960), pp. 197-205.
[75] A. Cohen, R. Devore, and C. Schwab. Analytic regularity and polynomial approximation of parametric and stochastic elliptic PDEs. Anal. Appl. (Singap.) 9.1 (2011), pp. 11-47.
[76] P. G. Constantine. Active Subspaces. Vol. 2. SIAM Spotlights. SIAM, 2015.
[77] P. G. Constantine, C. Kent, and T. Bui-Thanh. Accelerating Markov chain Monte Carlo with active subspaces. SIAM J. Sci. Comput. 38.5 (2016), A2779A2805.
[78] A. Cortinovis. Fast deterministic and randomized algorithms for low-rank approximation, matrix functions, and trace estimation. PhD thesis. EPFL, 2022.
[79] A. Cortinovis and D. Kressner. Low-rank approximation in the Frobenius norm by column and row subset selection. SIAM J. Matrix Anal. Appl. 41.4 (2020), pp. 1651-1673.
[80] A. Cortinovis, D. Kressner, and S. Massei. On maximum volume submatrices and cross approximation for symmetric semidefinite and diagonally dominant matrices. Linear Algebra Appl. 593 (2020), pp. 251-268.
[81] T. M. Cover and J. A. Thomas. Elements of Information Theory. 2nd. WileyInterscience, 2006.
[82] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In: Adv. Neural Inf. Process. Syst. 26. 2013, pp. 2292-2300.
[83] M. C. D'Autilia, I. Sgura, and V. Simoncini. Matrix-oriented discretization methods for reaction-diffusion PDEs: Comparisons and applications. Comput. Math. Appl. 79.7 (2020), pp. 2067-2085.
[84] J. Dabaghi, V. Ehrlacher, and C. Strössner. Computation of the self-diffusion coefficient with low-rank tensor methods: Application to the simulation of a crossdiffusion system. arXiv e-prints (2021), 2111.11349. To appear in ESAIM Proc. Surveys.
[85] J. Dabaghi, V. Ehrlacher, and C. Strössner. Tensor approximation of the selfdiffusion matrix of tagged particle processes. arXiv e-prints (2022), arXiv:2204.03943. To appear in J. Comput. Phys..
[86] P. J. Davis and P. Rabinowitz. Methods of Numerical Integration. 2nd. Comput. Sci. Appl. Math. Academic Press, 1984.
[87] L. De Lathauwer. Decompositions of a higher-order tensor in block terms. II. Definitions and uniqueness. SIAM J. Matrix Anal. Appl. 30.3 (2008), pp. 10331066.
[88] L. De Lathauwer, B. De Moor, and J. Vandewalle. A multilinear singular value decomposition. SIAM J. Matrix Anal. Appl. 21.4 (2000), pp. 1253-1278.
[89] A. De Masi and E. Presutti. Mathematical Methods for Hydrodynamic Limits. Vol. 1501. Lecture Notes in Math. Springer, 1991.
[90] A. Dektor, A. Rodgers, and D. Venturi. Rank-adaptive tensor methods for high-dimensional nonlinear PDEs. J. Sci. Comput. 88:36.36 (2021).
[91] A. Dektor and D. Venturi. Dynamically orthogonal tensor methods for highdimensional nonlinear PDEs. J. Comput. Phys. 404 (2020), p. 103501.
[92] A. Dektor and D. Venturi. Tensor rank reduction via coordinate flows. arXiv e-prints (2022), arXiv:2207.11955.
[93] A. Deshpande and L. Rademacher. Efficient volume sampling for row/column subset selection. In: 51st Annu. IEEE Symp. Found. Comput. Sci. FOCS. 2010, pp. 329-338.
[94] H. Dette and A. Pepelyshev. Generalized Latin hypercube design for computer experiments. Technometrics 52.4 (2010), pp. 421-429.
[95] H. P. Deutsch and K. Binder. Interdiffusion and self-diffusion in polymer mixtures: A Monte Carlo study. J. Chem. Phys. 94.3 (1991), pp. 2294-2304.
[96] S. Di Marino and A. Gerolin. An optimal transport approach for the Schrödinger bridge problem and convergence of Sinkhorn algorithm. J. Sci. Comput. 85.2 (2020), pp. 1-28.
[97] S. Di Marino, A. Gerolin, and L. Nenna. Optimal transportation theory with repulsive costs. In: Topological Optimization and Optimal Transport. Vol. 17. Radon Ser. Comput. Appl. Math. De Gruyter, 2017, pp. 204-256.
[98] J. Dieterich and B. Hartke. Empirical review of standard benchmark functions using evolutionary global optimization. Applied Math. 3.10 (2012).
[99] S. Dolgov and B. Khoromskij. Two-level QTT-Tucker format for optimized tensor calculus. SIAM J. Matrix Anal. Appl. 34.2 (2013), pp. 593-623.
[100] S. Dolgov, D. Kressner, and C. Strössner. Functional Tucker approximation using Chebyshev interpolation. SIAM J. Sci. Comput. 43.3 (2021), A2190-A2210.
[101] S. Dolgov and R. Scheichl. A hybrid alternating least squares-TT-cross algorithm for parametric PDEs. SIAM/ASA J. Uncertain. Quantif. 7.1 (2019), pp. 260291.
[102] T. A. Driscoll, F. Bornemann, and L. N. Trefethen. The chebop system for automatic solution of differential equations. BIT 48.4 (2008), pp. 701-723.
[103] T. A. Driscoll and N. Hale. Rectangular spectral collocation. IMA J. Numer. Anal. 36.1 (2016), pp. 108-132.
[104] T. A. Driscoll, N. Hale, and L. N. Trefethen. Chebfun Guide. Pafnuty Publications, 2014.
[105] Z. Drmač and A. K. Saibaba. The discrete empirical interpolation method: canonical structure and formulation in weighted inner product spaces. SIAM J. Matrix Anal. Appl. 39.3 (2018), pp. 1152-1180.
[106] P. Dvurechensky, A. Gasnikov, and A. Kroshnin. Computational optimal transport: Complexity by accelerated gradient descent is better than by Sinkhorn's algorithm. In: 35th Int. Conf. Mach. Learn. 2018, pp. 1367-1376.
[107] D. EbLing et al. Multiphysics simulation of thermoelectric systems for comparison with experimental device performance. J. Electron. Mater. 38.7 (2009), pp. 14561461.
[108] M. Eigel, R. Gruhlke, and M. Marschall. Low-rank tensor reconstruction of concentrated densities with application to Bayesian inversion. Stat. Comput. 32.2 (2022), Paper No. 27, 27.
[109] L. Einkemmer and C. Lubich. A low-rank projector-splitting integrator for the Vlasov-Poisson equation. SIAM J. Sci. Comput. 40.5 (2018), B1330-B1360.
[110] F. Elvander et al. Multi-marginal optimal transport using partial information with applications in robust localization and sensor fusion. Signal Process. 171 (2020), p. 107474.
[111] C. Erignoux. Limite hydrodynamique pour un gaz sur réseau de particules actives. PhD thesis. Université Paris Diderot, 2016.
[112] M. Espig et al. Efficient low-rank approximation of the stochastic Galerkin matrix in tensor formats. Comput. Math. Appl. 67.4 (2014), pp. 818-829.
[113] R. Eymard, T. Gallouët, and R. Herbin. A finite volume scheme for anisotropic diffusion problems. C. R. Math. Acad. Sci. Paris 339.4 (2004), pp. 299-302.
[114] R. Eymard, T. Gallouët, and R. Herbin. A new finite volume scheme for anisotropic diffusion problems on general grids: Convergence analysis. C. R. Math. Acad. Sci. Paris 344.6 (2007), pp. 403-406.
[115] R. Eymard, T. Gallouët, and R. Herbin. Convergence of finite volume schemes for semilinear convection diffusion equations. Numer. Math. 82.1 (1999), pp. 91116.
[116] R. Eymard, T. Gallouët, and R. Herbin. Finite volume methods. In: Handbook of Numerical Analysis, VII. North-Holland, 2000, pp. 713-1020.
[117] A. Falcó, W. Hackbusch, and A. Nouy. Tree-based tensor formats. SeMA Journal (2018).
[118] J. FAN et al. On the complexity of the optimal transport problem with graphstructured cost. In: 25th Int. Conf. Artif. Intell. Stat. 2022, pp. 9147-9165.
[119] G. FAVIER and A. L. de Almeida. Overview of constrained PARAFAC models. EURASIP J. Adv. Signal Process. 142 (2014), pp. 1-25.
[120] R. Ferrando and E. Scalas. Self-diffusion in a 2D lattice gas with lateral interactions. Surf. Sci. 281.1-2 (1993), pp. 178-190.
[121] J. FEYDY et al. Optimal transport for diffeomorphic registration. In: MICCAI. 2017, pp. 291-299.
[122] B. Fornberg. A Practical Guide to Pseudospectral Methods. Vol. 1. Cambridge Monogr. Appl. Comput. Math. Cambridge University Press, 1996.
[123] A. I. J. Forrester, A. Sóbester, and A. J. Keane. Engineering Design via Surrogate Modelling: A Practical Guide. John Wiley \& Sons, 2008.
[124] D. Fortunato, N. Hale, and A. Townsend. The ultraspherical spectral element method. J. Comput. Phys. 436 (2021), p. 110087.
[125] D. Fortunato and A. Townsend. Fast Poisson solvers for spectral methods. IMA J. Numer. Anal. 40.3 (2020), pp. 1994-2018.
[126] J. Franklin and J. Lorenz. On the scaling of multidimensional matrices. Linear Algebra Appl. 114/115 (1989), pp. 717-735.
[127] S. Friedland. Tensor optimal transport, distance between sets of measures and tensor scaling. arXiv e-prints (2020), arXiv:2005.00945.
[128] J. H. Friedman. Multivariate adaptive regression splines. Ann. Statist. 19.1 (1991), pp. 1-141.
[129] G. Friesecke and M. Penka. The GenCol algorithm for high-dimensional optimal transport: General formulation and application to barycenters and Wasserstein splines. arXiv e-prints (2022), arXiv:2209.09081.
[130] G. Friesecke, A. S. Schulz, and D. VöGler. Genetic column generation: Fast computation of high-dimensional multimarginal optimal transport problems. SIAM J. Sci. Comput. 44.3 (2022), A1632-A1654.
[131] X. Fu et al. Computing large-scale matrix and tensor decomposition with structured factors: A unified nonconvex optimization perspective. IEEE Signal Process. Mag. 37.5 (2020), pp. 78-94.
[132] N. Gantert and D. Schmid. The speed of the tagged particle in the exclusion process on Galton-Watson trees. Electron. J. Probab. 25 (2020), Paper No. 71, 27.
[133] S. Garreis and M. Ulbrich. Constrained optimization with low-rank tensors and applications to parametric problems with PDEs. SIAM J. Sci. Comput. 39.1 (2017), A25-A54.
[134] A. Genevay, G. Peyre, and M. Cuturi. Learning generative models with Sinkhorn divergences. In: Int. Conf. Artif. Intell. Stat. 2018, pp. 1608-1617.
[135] W. M. Gentleman. Algorithm 424: Clenshaw-Curtis quadrature [D1]. Commun. ACM 15.5 (1972), pp. 353-355.
[136] A. Genz. A package for testing multiple integration subroutines. In: Numerical Integration. NATO ASI Series. Springer, 1987, pp. 337-340.
[137] G. H. Golub and C. F. Van Loan. Matrix Computations. 4th. Johns Hopkins Stud. Math. Sci. Johns Hopkins University Press, 2013.
[138] S. A. Goreinov, E. E. Tyrtyshnikov, and N. L. Zamarashkin. A theory of pseudoskeleton approximations. Linear Algebra Appl. 261 (1997), pp. 1-21.
[139] A. Gorodetsky. Continuous low-rank tensor decompositions, with applications to stochastic optimal control and data assimilation. PhD thesis. MIT, 2017.
[140] A. A. Gorodetsky and J. D. Jakeman. Gradient-based optimization for regression in the functional tensor-train format. J. Comput. Phys. 374 (2018), pp. 12191238.
[141] A. Gorodetsky, S. Karaman, and Y. Marzouk. A continuous analogue of the tensor-train decomposition. Comput. Methods Appl. Mech. Eng. 347 (2019), pp. 59-84.
[142] A. Gorodetsky, S. Karaman, and Y. Marzouk. High-dimensional stochastic optimal control using continuous tensor decompositions. Int. J. Robot. Res. 37.2-3 (2018), pp. 340-377.
[143] D. Gottlieb and S. A. Orszag. Numerical Analysis of Spectral Methods: Theory and Applications. SIAM, 1977.
[144] R. B. Gramacy and H. K. H. Lee. Adaptive design and analysis of supercomputer experiments. Technometrics 51.2 (2009), pp. 130-145.
[145] L. Grasedyck, D. Kressner, and C. Tobler. A literature survey of low-rank tensor approximation techniques. GAMM-Mitt. 36.1 (2013), pp. 53-78.
[146] M. Griebel and H. Harbrecht. Analysis of tensor approximation schemes for continuous functions. Found. Comput. Math. (2021), pp. 1-22.
[147] M. Griebel and H. Harbrecht. Approximation of bi-variate functions: Singular value decomposition versus sparse grids. IMA J. Numer. Anal. 34.1 (2014), pp. 2854.
[148] M. Griebel, H. Harbrecht, and R. Schneider. Low-rank approximation of continuous functions in Sobolev spaces with dominating mixed smoothness. arXiv e-prints (2022), arXiv:2203.04100.
[149] M. Gu and S. C. Eisenstat. Efficient algorithms for computing a strong rankrevealing QR factorization. SIAM J. Sci. Comput. 17.4 (1996), pp. 848-869.
[150] I. HAASLER et al. Multi-marginal optimal transport and probabilistic graphical models. IEEE Trans. Inf. Theory 67.7 (2021), pp. 4647-4668.
[151] I. HaAsLer et al. Multimarginal optimal transport with a tree-structured cost and the Schrödinger bridge problem. SIAM J. Control Optim. 59.4 (2021), pp. 24282453.
[152] W. Hackbusch and S. Kühn. A new scheme for the tensor representation. J. Fourier Anal. Appl. 15.5 (2009), pp. 706-722.
[153] W. Hackbusch. Computation of best $L^{\infty}$ exponential sums for $1 / x$ by Remez' algorithm. Comput. Vis. Sci. 20.1-2 (2019), pp. 1-11.
[154] W. Hackbusch. Tensor Spaces and Numerical Tensor Calculus. Vol. 42. Springer Ser. Comput. Math. Springer, 2012.
[155] W. Hackbusch, L. Grasedyck, and S. Börm. An introduction to hierarchical matrices. In: EQUADIFF 10. Vol. 127. 2. 2002, pp. 229-241.
[156] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. SIAM Rev. 53.2 (2011), pp. 217-288.
[157] K. Hamm. Generalized pseudoskeleton decompositions. arXiv e-prints (2022), arXiv:2206.14905.
[158] J. M. Hammersley and D. C. Handscomb. Monte Carlo Methods. Monogr. Statist. Appl. Probab. Springer, 1964.
[159] S. Hao and D. S. Sholl. Self-diffusion and macroscopic diffusion of hydrogen in amorphous metals from first-principles calculations. J. Chem. Phys. 130.24 (2009), p. 244705.
[160] S. Hao and P.-G. Martinsson. A direct solver for elliptic PDEs in three dimensions based on hierarchical merging of Poincaré-Steklov operators. J. Comput. Appl. Math. 308 (2016), pp. 419-434.
[161] R. A. Harshman. Foundations of the PARAFAC Procedure: Models and Conditions for an Explanatory Multimodal factor Analysis. Working Papers in Phonetics 16. University of California at Los Angeles, 1970.
[162] B. Hashemi and Y. Nakatsukasa. Least-squares spectral methods for ODE eigenvalue problems. SIAM J. Sci. Comput. 44.5 (2022), A3244-A3264.
$[163]$ B. Hashemi and L. N. Trefethen. Chebfun in three dimensions. SiAm J. Sci. Comput. 39.5 (2017), pp. C341-C363.
$[164]$ J. S. Hesthaven, S. Gottlieb, and D. Gottlieb. Spectral Methods for TimeDependent Problems. Vol. 21. Cambridge Monogr. Appl. Comput. Math. Cambridge University Press, 2007.
[165] N. J. Higham. The numerical stability of barycentric Lagrange interpolation. IMA J. Numer. Anal. 24.4 (2004), pp. 547-556.
[166] T. Hillen and K. J. Painter. A user's guide to PDE models for chemotaxis. J. Math. Biol. 58.1-2 (2009), pp. 183-217.
$[167]$ G. E. Hinton and R. Zemel. Autoencoders, minimum description length and Helmholtz free energy. In: Adv. Neural Inf. Process. Syst. 6. Vol. 6. MorganKaufmann, 1993.
[168] F. L. Hitchсоск. Multiple invariants and generalized rank of a p-way matrix or tensor. J. Math. Phys. 7.1-4 (1928), pp. 39-79.
[169] S.-W. Ho and R. W. Yeung. The interplay between entropy and variational distance. IEEE Trans. Inform. Theory 56.12 (2010), pp. 5906-5929.
[170] S. Holtz, T. Rohwedder, and R. Schneider. On manifolds of tensors of fixed TT-rank. Numer. Math. 120.4 (2012), pp. 701-731.
[171] S. Holtz, T. Rohwedder, and R. Schneider. The alternating linear scheme for tensor optimization in the tensor train format. SIAM J. Sci. Comput. 34.2 (2012), A683-A713.
[172] D. Hong, T. G. Kolda, and J. A. Duersch. Generalized canonical polyadic tensor decomposition. SIAM Rev. 62.1 (2020), pp. 133-163.
[173] C. Huang and A. Darwiche. Inference in belief networks: A procedural guide. Internat. J. Approx. Reason. 15.3 (1996), pp. 225-263.
[174] C. Huang et al. Efficient parallelization of tensor network contraction for simulating quantum computation. Nat. Comput. Sci. 1.9 (2021), pp. 578-587.
[175] H. IGEL. Wave propagation in three-dimensional spherical sections by the Chebyshev spectral method. Geophys. J. Int. 136.3 (1999), pp. 559-566.
[176] T. L. Jackson and H. M. Byrne. A mechanical model of tumor encapsulation and transcapsular spread. Math. Biosci. 180 (2002), pp. 307-328.
[177] M. Jamil and X.-S. Yang. A literature survey of benchmark functions for global optimisation problems. Int. J. Math. Model. Numer. Optim. 4.2 (2013), pp. 150194.
[178] M. Jara. Finite-dimensional approximation for the diffusion coefficient in the simple exclusion process. Ann. Probab. 34.6 (2006), pp. 2365-2381.
[179] Y. Ji et al. A survey on tensor techniques and applications in machine learning. IEEE Access 7 (2019), pp. 162950-162990.
[180] T.-X. Jiang et al. Nonnegative low rank tensor approximation and its application to multi-dimensional images. Numer. Math. (2022), pp. 1-30.
[181] Y. W. Jonathan S Yedidia William Freeman. Generalized belief propagation. In: Adv. Neural Inf. Process. Syst. 13. 2000, pp. 689-695.
[182] A. JÜNGEL. Entropy Methods for Diffusive Partial Differential Equations. Springer Briefs Math. Springer, 2016.
[183] A. JÜNGEL. The boundedness-by-entropy method for cross-diffusion systems. Nonlinearity 28.6 (2015), pp. 1963-2001.
[184] G. E. Karniadakis and S. J. Sherwin. Spectral/hp Element Methods for Computational Fluid Dynamics. 2nd. Numer. Math. Sci. Comput. Oxford University Press, 2005.
[185] C. T. Kelley. Iterative Methods for Linear and Nonlinear Equations. Vol. 16. Front. Appl. Math. SIAM, 1995.
[186] B. N. KhoromskiJ. Structured rank- $\left(R_{1}, \ldots, R_{D}\right)$ decomposition of functionrelated tensors in $\mathbb{R}^{D}$. Comput. Methods Appl. Math. 6.2 (2006), pp. 194-220.
[187] B. N. Khoromskij. $O(d \log N)$-quantics approximation of $N-d$ tensors in highdimensional numerical modeling. Constr. Approx. 34.2 (2011), pp. 257-280.
[188] B. N. KhoromskiJ. Tensor numerical methods for multidimensional PDEs: Theoretical analysis and initial applications. ESAIM Proc. Surveys 48 (2015), pp. 128.
[189] B. N. KhoromskiJ. Tensor Numerical Methods in Scientific Computing. Vol. 19. Radon Ser. Comput. Appl. Math. De Gruyter, 2018.
[190] H. A. L. Kiers. Towards a standardized notation and terminology in multiway analysis. J. Chemom. 14.3 (2000), pp. 105-122.
[191] Y. Kim and S. Choi. Nonnegative Tucker decomposition. In: IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. 2007, pp. 1-8.
[192] C. Kipnis and S. R. S. Varadhan. Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions. Comm. Math. Phys. 104.1 (1986), pp. 1-19.
[193] C. Kipnis, C. Landim, and S. Olla. Hydrodynamical limit for a nongradient system: The generalized symmetric exclusion process. Commun. Pure Appl. Math. 47.11 (1994), pp. 1475-1545.
[194] N. Kishore Kumar and J. Schneider. Literature survey on low rank approximation of matrices. Linear Multilinear Algebra 65.11 (2017), pp. 2212-2244.
[195] A. Klimke. Sparse grid interpolation toolbox v5.1.1. 2008.
[196] O. Koch and C. Lubich. Dynamical tensor approximation. SIAM J. Matrix Anal. Appl. 31.5 (2010), pp. 2360-2375.
[197] W. Koepf. Hypergeometric Summation. Adv. Lect. Math. Friedr. Vieweg \& Sohn, 1998.
[198] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. SIAM Rev. 51.3 (2009), pp. 455-500.
[199] S. Kolouri et al. Optimal mass transport: Signal processing and machine-learning applications. IEEE Signal Process. Mag. 34.4 (2017), pp. 43-59.
[200] T. Komorowski, C. Landim, and S. Olla. Fluctuations in Markov Processes. Vol. 345. Grundlehren Math. Wiss. Springer, 2012.
[201] K. Konakli and B. Sudret. Global sensitivity analysis using low-rank tensor approximations. Reliab. Eng. Syst. Saf. 156 (2016), pp. 64-83.
[202] K. Konakli and B. Sudret. Polynomial meta-models with canonical low-rank approximations: Numerical insights and comparison to sparse polynomial chaos expansions. J. Comput. Phys. 321 (2016), pp. 1144-1169.
[203] K. Konakli and B. Sudret. Uncertainty quantification in high-dimensional spaces with low-rank tensor approximations. In: 1st Int. Conf. Uncertain. Quantif. Comput. Sci. Eng. 2015.
[204] D. Kressner and L. Periša. Recompression of Hadamard products of tensors in Tucker format. SIAM J. Sci. Comput. 39.5 (2017), A1879-A1902.
[205] D. Kressner, M. Steinlechner, and B. Vandereycken. Low-rank tensor completion by Riemannian optimization. BIT 54.2 (2014), pp. 447-468.
[206] D. Kressner and C. Tobler. Krylov subspace methods for linear systems with tensor product structure. SIAM J. Matrix Anal. Appl. 31.4 (2009), pp. 1688-1714.
[207] D. Kressner and C. Tobler. Low-rank tensor Krylov subspace methods for parametrized linear systems. SIAM J. Matrix Anal. Appl. 32.4 (2011), pp. 12881316.
[208] D. Kressner and C. Tobler. Preconditioned low-rank methods for high-dimensional elliptic PDE eigenvalue problems. Comput. Methods Appl. Math. 11.3 (2011), pp. 363-381.
[209] L. Kronecker. Über einige Interpolationsformeln für ganze Funktionen mehrerer Variablen. In: Monatsberichte der Königlichen Preuß. Akademie der Wissenschaften zu Berlin aus dem Jahre 1865. Buchdruckerei der Königl. Akademie der Wissenschaften, 1866, pp. 686-691.
[210] C. Landim, S. Olla, and S. R. S. Varadhan. Finite-dimensional approximation of the self-diffusion coefficient for the exclusion process. Ann. Probab. 30.2 (2002), pp. 483-508.
[211] C. LÉOnard. From the Schrödinger problem to the Monge-Kantorovich problem. J. Funct. Anal. 262.4 (2012), pp. 1879-1920.
[212] K.-C. Li. Sliced inverse regression for dimension reduction. J. Amer. Statist. Assoc. 86.414 (1991), pp. 316-342.
[213] S. Li et al. Low-rank tensor decomposition based anomaly detection for hyperspectral imagery. In: IEEE Int. Conf. Image Process. (ICIP). 2015, pp. 4525-4529.
[214] T. M. Liggett. Interacting Particle Systems. Vol. 276. Grundlehren Math. Wiss. Springer, 1985.
[215] T. M. Liggett. Stochastic Interacting Systems: Contact, Voter and Exclusion Processes. Vol. 324. Grundlehren Math. Wiss. Springer, 1999.
[216] T. Lin, N. Ho, and M. I. Jordan. On efficient optimal transport: An analysis of greedy and accelerated mirror descent algorithms. In: 36th Int. Conf. Mach. Learn. 2019, pp. 3982-3991.
[217] T. Lin et al. On the complexity of approximating multimarginal optimal transport. J. Mach. Learn. Res. 23.65 (2022), pp. 1-43.
[218] T.-Y. Lin et al. Microsoft COCO: Common objects in context. In: ECCV. 2014, pp. 740-755.
[219] G. Liu et al. Module-based multiscale simulation of angiogenesis in skeletal muscle. Theor. Biol. Med. Model. 8.6 (2011), pp. 1-21.
[220] G. J. Lord, C. E. Powell, and T. Shardlow. An Introduction to Computational Stochastic PDEs. Cambridge Texts Appl. Math. Cambridge University Press, 2014.
[221] T. H. LuU et al. A new method for reconstruction of cross-sections using Tucker decomposition. J. Comput. Phys. 345 (2017), pp. 189-206.
[222] P.-G. Martinsson and J. A. Tropp. Randomized numerical linear algebra: Foundations and algorithms. Acta Numer. 29 (2020), pp. 403-572.
[223] J. C. Mason. Near-best multivariate approximation by Fourier series, Chebyshev series and Chebyshev interpolation. J. Approx. Theory 28.4 (1980), pp. 349-358.
[224] J. C. Mason and D. C. Handscomb. Chebyshev Polynomials. Chapman and Hall/CRC, 2002.
[225] D. M. Mattox. Handbook of Physical Vapor Deposition (PVD) Processing. 2nd. William Andrew Publishing, 2010.
[226] D. McBride et al. Computational modelling of variably saturated flow in porous media with complex three-dimensional geometries. Internat. J. Numer. Methods Fluids 50.9 (2006), pp. 1085-1117.
[227] N. Metropolis and S. Ulam. The Monte Carlo method. J. Amer. Statist. Assoc. 44 (1949), pp. 335-341.
[228] R. Minster, Z. Li, and G. Ballard. Parallel randomized Tucker decomposition algorithms. arXiv e-prints (2022), arXiv:2211.13028.
[229] R. Minster, A. K. Saibaba, and M. E. Kilmer. Efficient tensor-based approximations to kernel interactions. SIAM-ALA presentation. 2021.
[230] R. Minster, A. K. Saibaba, and M. E. Kilmer. Randomized algorithms for low-rank tensor decompositions in the Tucker format. SIAM J. Math. Data Sci. 2.1 (2020), pp. 189-215.
[231] L. Mirsky. Symmetric gauge functions and unitarily invariant norms. Quart. J. Math. Oxford Ser. (2) 11 (1960), pp. 50-59.
[232] M. Moakher and P. G. Batchelor. Symmetric positive-definite matrices: From geometry to applications and visualization. In: Visualization and Processing of Tensor Fields. Math. Vis. Springer, 2006, pp. 285-298, 452.
[233] H. Moon, A. M. Dean, and T. J. Santner. Two-stage sensitivity-based group screening in computer experiments. Technometrics 54.4 (2012), pp. 376-387.
[234] M. Motamed. A hierarchically low-rank optimal transport dissimilarity measure for structured data. BIT (2022).
[235] M. Mozaffari et al. Optimal transport theory for power-efficient deployment of unmanned aerial vehicles. In: IEEE Int. Conf. Commun. 2016, pp. 1-6.
[236] E. A. Muravleva and I. V. Oseledets. Approximate solution of linear systems with Laplace-like operators via cross approximation in the frequency domain. Comput. Methods Appl. Math. 19.1 (2019), pp. 137-145.
[237] S. Narumi. Some formulas in the theory of interpolation of many independent variables. Tohoku Math. J. (1) 18 (1920), pp. 309-321.
[238] L. Nenna. Numerical methods for multi-marginal optimal transportation. PhD thesis. PSL Research University, 2016.
[239] H. Niederreiter. Random Number Generation and Quasi-Monte Carlo Methods. Vol. 63. CBMS-NSF Regional Conf. Ser. in Appl. Math. SIAM, 1992.
[240] NIST Handbook of Mathematical Functions. Cambridge University Press, 2010.
[241] A. Nouy. Higher-order principal component analysis for the approximation of tensors in tree-based low-rank formats. Numer. Math. 141.3 (2019), pp. 743-789.
[242] A. Nouy. Low-rank methods for high-dimensional approximation and model order reduction. In: Model Reduction and Approximation. Vol. 15. Comput. Sci. Eng. SIAM, 2017, pp. 171-226.
[243] A. Nouy. Low-rank tensor methods for model order reduction. In: Handbook of Uncertainty Quantification. Springer, 2017, pp. 857-882.
[244] S. Olver, R. M. Slevinsky, and A. Townsend. Fast algorithms using orthogonal polynomials. Acta Numer. 29 (2020), pp. 573-699.
[245] S. Olver, A. Townsend, and G. Vasil. A sparse spectral method on triangles. SIAM J. Sci. Comput. 41.6 (2019), A3728-A3756.
[246] S. Olver and A. Townsend. A fast and well-conditioned spectral method. SIAM Rev. 55.3 (2013), pp. 462-489.
[247] R. Orús. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. Ann. Physics 349 (2014), pp. 117-158.
[248] I. V. Oseledets. Approximation of matrices with logarithmic number of parameters. Doklady Math. 428.1 (2009), pp. 23-24.
[249] I. V. Oseledets. Tensor-train decomposition. SIAM J. Sci. Comput. 33.5 (2011), pp. 2295-2317.
[250] I. V. Oseledets, M. V. Rakhuba, and A. Uschmajew. Alternating least squares as moving subspace correction. SIAM J. Numer. Anal. 56.6 (2018), pp. 3459-3479.
[251] I. Oseledets and E. Tyrtyshnikov. TT-cross approximation for multidimensional arrays. Linear Algebra Appl. 432.1 (2010), pp. 70-88.
[252] P. PaAtERO. Construction and analysis of degenerate PARAFAC models. J. Chemom. 14.3 (2000), pp. 285-299.
[253] B. PASS. Multi-marginal optimal transport: Theory and applications. ESAIM: M2AN 49.6 (2015), pp. 1771-1790.
[254] R. Peng, J. Gray, and G. Kin-Lic Chan. Arithmetic circuit tensor networks, multivariable function representation, and high-dimensional integration. arXiv e-prints (2022), arXiv:2209.07410.
[255] R. Penrose. Applications of negative dimensional tensors. In: Combin. Math. Appl. Academic Press, London, 1971, pp. 221-244.
[256] D. PEREZ-GARCIA et al. Matrix product state representations. Quantum Info. Comput. 7.5 (2007), pp. 401-430.
[257] T. von Petersdorff and C. Schwab. Numerical solution of parabolic equations in high dimensions. M2AN Math. Model. Numer. Anal. 38.1 (2004), pp. 93-127.
[258] G. Peyré and M. Cuturi. Computational optimal transport: With applications to data science. Found. Trends Mach. Learn. 11.5-6 (2019), pp. 355-607.
[259] H. P. PFEIFFER et al. A multidomain spectral method for solving elliptic equations. Comput. Phys. Comm. 152.3 (2003), pp. 253-273.
[260] R. B. Platte and L. N. Trefethen. Chebfun: A new kind of numerical computing. In: Progress in Industrial Mathematics at ECMI 2008. Vol. 15. Math. Ind. Springer, 2010, pp. 69-87.
[261] R.-E. Plessix. A Helmholtz iterative solver for 3D seismic-imaging problems. Geophysics 72.5 (2007), SM185-SM194.
[262] T. Poggio et al. Why and when can deep-but not shallow-networks avoid the curse of dimensionality: A review. Int. J. Autom. Comput. 14.5 (2017).
[263] A. Popov and N. Y. ZHU. Modeling radio wave propagation in tunnels with a vectorial parabolic equation. IEEE Trans. Antennas Propag. 48.9 (2000), pp. 14031412.
[264] M. Psenka and N. Boumal. Second-order optimization for tensors with fixed tensor-train rank. arXiv e-prints (2020), arXiv:2011.13395.
[265] A. Qing. Dynamic differential evolution strategy and applications in electromagnetic inverse scattering problems. IEEE Trans. Geosci. Remote Sens. 44.1 (2006), pp. 116-125.
[266] J. Quastel. Diffusion of color in the simple exclusion process. Comm. Pure Appl. Math. 45.6 (1992), pp. 623-679.
[267] J. Rabin, S. Ferradans, and N. Papadakis. Adaptive color transfer with relaxed optimal transport. In: IEEE Int. Conf. Imag. Process. 2014, pp. 4852-4856.
[268] S. Rahnamayan, H. Tizhoosh, and M. Salama. Opposition-based differential evolution ( $O D E$ ) with variable jumping rate. In: IEEE Symp. Found. Comput. Intell. 2007, pp. 81-88.
[269] S. Rahnamayan, H. R. Tizhoosh, and M. M. A. Salama. A novel population initialization method for accelerating evolutionary algorithms. Comput. Math. Appl. 53.10 (2007), pp. 1605-1614.
[270] P. RAI et al. Randomized functional sparse Tucker tensor for compression and fast visualization of scientific data. arXiv e-prints (2019), arXiv:1907.05884.
[271] M. V. Rakhuba and I. V. Oseledets. Fast multidimensional convolution in low-rank tensor formats via cross approximation. SIAM J. Sci. Comput. 37.2 (2015), A565-A582.
[272] S.-J. RAN et al. Tensor Network Contractions: Methods and Applications to Quantum Many-Body Systems. Vol. 964. Lecture Notes in Phys. Springer, 2020.
[273] E. Robeva and A. Seigal. Duality of graphical models and tensor networks. Inf. Inference 8.2 (2019), pp. 273-288.
[274] T. Rohwedder and A. Uschmajew. On local convergence of alternating schemes for optimization of convex problems in the tensor train format. SIAM J. Numer. Anal. 51.2 (2013), pp. 1134-1162.
[275] S. F. Roohi et al. Multi-dimensional low rank plus sparse decomposition for reconstruction of under-sampled dynamic MRI. Pattern Recognit. 63 (2017), pp. 667679.
[276] E. SAADA. A limit theorem for the position of a tagged particle in a simple exclusion process. Ann. Probab. 15.1 (1987), pp. 375-381.
[277] A. K. Saibaba. HOID: Higher order interpolatory decomposition for tensors based on Tucker representation. SIAM J. Matrix Anal. Appl. 37.3 (2016), pp. 1223-1249.
[278] A. K. Saibaba, R. Minster, and M. E. Kilmer. Efficient randomized tensorbased algorithms for function approximation and low-rank kernel interactions. Adv. Comput. Math. 48 (2022).
[279] E. Sanz and D. Marenduzzo. Dynamic Monte Carlo versus Brownian dynamics: A comparison for self-diffusion and crystallization in colloidal fluids. J. Chem. Phys. 132.19 (2010), p. 194102.
[280] S. A. Sauter and C. Schwab. Boundary Element Methods. Vol. 39. Springer Ser. Comput. Math. Springer, 2011.
[281] D. Savostyanov and I. Oseledets. Fast adaptive interpolation of multi-dimensional arrays in tensor train format. In: 7th Int. Workshop Multidimens. (nD) Syst. 2011, pp. 1-8.
[282] D. V. SaVOSTYANOV. Quasioptimality of maximum-volume cross interpolation of tensors. Linear Algebra Appl. 458 (2014), pp. 217-244.
[283] M. Scetbon and M. Cuturi. Low-rank optimal transport: Approximation, statistics and debiasing. arXiv e-prints (2022), arXiv:2205.12365.
[284] M. Scetbon, M. Cuturi, and G. Peyré. Low-rank Sinkhorn factorization. In: Proc. 38th Int. Conf. Mach. Learn. 2021, pp. 9344-9354.
[285] E. Schmidt. Zur Theorie der linearen und nicht linearen Integralgleichungen Zweite Abhandlung. Math. Ann. 64.2 (1907), pp. 161-174.
[286] R. SChNEIDER and A. Uschmajew. Approximation rates for the hierarchical tensor format in periodic Sobolev spaces. J. Complexity 30.2 (2014), pp. 56-71.
[287] N. SchWENCK et al. Dimensionally reduced flow models in fractured porous media: Crossings and boundaries. Comput. Geosci. 19.6 (2015), pp. 1219-1230.
[288] A. Shapira. Hydrodynamic limit of the Kob-Andersen model. arXiv e-prints (2020), arXiv:2003.08495.
[289] T. Shi and A. Townsend. On the compressibility of tensors. SIAM J. Matrix Anal. Appl. 42.1 (2021), pp. 275-298.
[290] N. Shigesada, K. Kawasaki, and E. Teramoto. Spatial segregation of interacting species. J. Theoret. Biol. 79.1 (1979), pp. 83-99.
[291] N. D. Sidiropoulos et al. Tensor decomposition for signal processing and machine learning. IEEE Trans. Signal Process. 65.13 (2017), pp. 3551-3582.
[292] V. de Silva and L.-H. Lim. Tensor rank and the ill-posedness of the best low-rank approximation problem. SIAM J. Matrix Anal. Appl. 30.3 (2008), pp. 1084-1127.
[293] N. Singh et al. Comparison of accuracy and scalability of Gauss-Newton and alternating least squares for CANDECOMC/PARAFAC decomposition. SIAM J. Sci. Comput. 43.4 (2021), pp. C290-C311.
[294] R. Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. Ann. Math. Statist. 35 (1964), pp. 876-879.
[295] S. A. Smolyak. Quadrature and interpolation formulas for tensor products of certain classes of functions. Dokl. Akad. Nauk SSSR 148.5 (1963), pp. 1042-1045.
[296] I. M. Sobol'. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Math. Comput. Simulation 55.1-3 (2001), pp. 271280.
[297] I. M. Sobol'. Sensitivity estimates for nonlinear mathematical models. Math. Modeling Comput. Experiment 1.4 (1993), 407-414 (1995).
[298] M. B. Soley et al. Functional tensor-train Chebyshev method for multidimensional quantum dynamics simulations. J. Chem. Theory Comput. 18 (2022), pp. 25-36.
[299] Z. Song, D. P. Woodruff, and P. Zhong. Low rank approximation with entrywise $\ell_{1}$-norm error. In: 49th Annu. ACM Symp. Theory Comput. ACM, New York, 2017, pp. 688-701.
[300] D. C. Sorensen and M. Embree. A DEIM induced CUR factorization. SIAM J. Sci. Comput. 38.3 (2016), A1454-A1482.
[301] H. Spohn. Large Scale Dynamics of Interacting Particles. Texts Monogr. Phys. Springer, 1991.
[302] M. Stoll. A literature survey of matrix methods for data science. GAMM-Mitt. 43.3 (2020), e202000013, 26.
[303] J. C. Strikwerda. Finite Difference Schemes and Partial Differential Equations. 2nd. SIAM, 2004.
[304] C. StröSSNER and D. Kressner. Fast global spectral methods for three-dimensional partial differential equations. IMA J. Numer. Anal. (2022), pp. 1-24.
[305] C. StröSSner and D. Kressner. Low-rank tensor approximations for solving multimarginal optimal transport problems. SIAM J. Imaging Sci. 16.1 (2023), pp. 169-191.
[306] C. Strössner, B. Sun, and D. Kressner. Approximation in the extended functional tensor train format. arXiv e-prints (2022), arXiv:2211.11338.
[307] A. M. Stuart. Inverse problems: A Bayesian perspective. Acta Numer. 19 (2010), pp. 451-559.
[308] B. Sudret, S. Marelli, and J. Wiart. Surrogate models for uncertainty quantification: An overview. In: 17th Eur. Conf. Antennas Propag. 2017, pp. 793797.
[309] A. Sultonov, S. Matveev, and S. Budzinskiy. Low-rank nonnegative tensor approximation via alternating projections and sketching. arXiv e-prints (2022), arXiv:2209.02060.
[310] S. Surjanovic and D. Bingham. Virtual library of simulation experiments: Test functions and datasets. Retrieved November 14, 2022, from https://www.sfu.ca/ ~ssurjano/. 2013.
[311] K. R. Swanson et al. Virtual and real brain tumors: Using mathematical modeling to quantify glioma growth and invasion. J. Neurol. Scie. 216.1 (2003), pp. 1-10.
[312] S. Szalay et al. Tensor product methods and entanglement optimization for ab initio quantum chemistry. Int. J. Quantum Chem. 115.19 (2015), pp. 1342-1391.
[313] P. R. Taylor. Stochastic lattice-based models of diffusion in biological systems. PhD thesis. University of Oxford, 2016.
[314] A. Thibault et al. Overrelaxed Sinkhorn-Knopp algorithm for regularized optimal transport. Algorithms 14.5 (2021).
[315] A. L. Thorneywork et al. Effect of hydrodynamic interactions on self-diffusion of quasi-two-dimensional colloidal hard spheres. Phys. Rev. Lett. 115 (26 2015), p. 268301.
[316] R. A. Todor and C. Schwab. Convergence rates for sparse chaos approximations of elliptic problems with stochastic coefficients. IMA J. Numer. Anal. 27.2 (2007), pp. 232-261.
[317] C. Toninelli, G. Biroli, and D. S. Fisher. Spatial structures and dynamics of kinetically constrained models of glasses. Phys. Rev. Lett. 92 (18 2004), p. 185504.
[318] A. Townsend. Computing with functions in two dimensions. PhD thesis. University of Oxford, 2014.
[319] A. Townsend and S. Olver. The automatic solution of partial differential equations using a global spectral method. J. Comput. Phys. 299 (2015), pp. 106-123.
[320] A. Townsend and L. N. Trefethen. An extension of Chebfun to two dimensions. SIAM J. Sci. Comput. 35.6 (2013), pp. C495-C518.
[321] A. Townsend, H. Wilber, and G. B. Wright. Computing with functions in spherical and polar geometries I. The sphere. SIAM J. Sci. Comput. 38.4 (2016), pp. C403-C425.
[322] L. N. Trefethen. Approximation Theory and Approximation Practice. SIAM, 2013.
[323] L. N. Trefethen. Computing numerically with functions instead of numbers. Math. Comput. Sci. 1.1 (2007), pp. 9-19.
[324] L. N. Trefethen. Cubature, approximation, and isotropy in the hypercube. SIAM Rev. 59.3 (2017), pp. 469-491.
[325] L. N. Trefethen. Is Gauss quadrature better than Clenshaw-Curtis? SIAM Rev. 50.1 (2008), pp. 67-87.
[326] L. N. Trefethen. Spectral Methods in MATLAB. Vol. 10. Software Environ. Tools. SIAM, 2000.
[327] L. R. Tucker. Some mathematical notes on three-mode factor analysis. Psychometrika 31 (1966), pp. 279-311.
[328] E. Ullmann, H. C. Elman, and O. G. Ernst. Efficient iterative solvers for stochastic Galerkin discretizations of log-transformed random diffusion problems. SIAM J. Sci. Comput. 34.2 (2012), A659-A682.
[329] C. Vanaret et al. Certified global minima for a benchmark of difficult optimization problems. arXiv e-prints (2020), arXiv:2003.09867.
[330] M. Vandecappelle, N. Vervliet, and L. De Lathauwer. A second-order method for fitting the canonical polyadic decomposition with non-least-squares cost. IEEE Trans. Signal Process. 68 (2020), pp. 4454-4465.
[331] F. Verstraete, V. Murg, and J. Cirac. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. Adv. Phys. 57.2 (2008), pp. 143-224.
[332] N. VERVLIET et al. Tensorlab 3.0. 2016.
[333] C. Villani. Optimal Transport: Old and New. 1st. Grundlehren Math. Wiss. Springer, 2009.
[334] J. Waldvogel. Fast construction of the Fejér and Clenshaw-Curtis quadrature rules. BIT 46.1 (2006), pp. 195-202.
[335] W. WANG et al. A linear optimal transportation framework for quantifying and visualizing variations in sets of images. Int. J. Comput. Vis. 101.2 (2013), pp. 254269.
[336] Z. WANG and S. F. Wu. Helmholtz equation-least-squares method for reconstructing the acoustic pressure field. J. Acoust. Soc. Am. 102.4 (1997), pp. 2020-2032.
[337] S. R. White. Density matrix formulation for quantum renormalization groups. Phys. Rev. Lett. 69 (19 1992), pp. 2863-2866.
[338] H. Wilber, A. Townsend, and G. B. Wright. Computing with functions in spherical and polar geometries II. The disk. SIAM J. Sci. Comput. 39.3 (2017), pp. C238-C262.
[339] D. P. Woodruff. Sketching as a tool for numerical linear algebra. Found. Trends Theor. Comput. Sci. 10.1-2 (2014), pp. iv +157.
[340] H. Xiang and L. Grigori. Kronecker product approximation preconditioners for convection-diffusion model problems. Numer. Linear Algebra Appl. 17.4 (2010), pp. 691-712.
[341] D. Xiu. Stochastic collocation methods: A survey. In: Handbook of Uncertainty Quantification. Springer, 2017, pp. 699-716.
[342] H. Xu et al. Spectral/hp element methods: Recent developments, applications, and perspectives. J. Hydrodyn. 30.1 (2018), pp. 1-22.
[343] N. L. Zamarashkin, S. V. Morozov, and E. E. Tyrtyshnikov. On the best approximation algorithm by low-rank matrices in Chebyshev's norm. Comput. Math. Math. Phys. 62.5 (2022), pp. 701-718.
[344] V. P. Zankin, G. V. Ryzhakov, and I. V. Oseledets. Gradient descent-based D-optimal design for the least-squares polynomial approximation. arXiv e-prints (2018), arXiv:1806.06631.
[345] C. Zenger. Sparse grids. In: Parallel Algorithms for Partial Differential Equations. Vol. 31. Notes Numer. Fluid Mech. Friedr. Vieweg \& Sohn, 1991, pp. 241-251.
[346] S. Zhao and M. J. Yedlin. A new iterative Chebyshev spectral method for solving the elliptic equation $\nabla \cdot(\sigma \nabla u)=f$. J. Comput. Phys. 113.2 (1994), pp. 215-223.
[347] M. S. Zhdanov, S. K. Lee, and K. Yoshioka. Integral equation method for 3D modeling of electromagnetic fields in complex structures with inhomogeneous background conductivity. Geophysics 71.6 (2006), G333-G345.
[348] B. Zhou and M. Parno. Efficient and exact multimarginal optimal transport with pairwise costs. arXiv e-prints (2022), arXiv:2208.03025.
[349] G. ZHOU et al. Efficient nonnegative Tucker decompositions: Algorithms and uniqueness. IEEE Trans. Image Process. 24.12 (2015), pp. 4990-5003.
[350] Y. Zhou et al. Large-scale parallel collaborative filtering for the Netflix prize. In: Algorithmic Aspects in Information and Management. Springer, 2008, pp. 337-348.

# Curriculum Vitae 

## Personal Details

Christoph Max Strößner<br>born 22.09.1995 in Hof/Germany

## Education

| École Polytechique Fédérale de Lausanne | Switzerland <br> since $01 / 2019$ |
| :--- | ---: |
| PhD in Mathematics |  |
| Thesis: Low-Rank Tensor Methods for High-Dimensional Problems |  |
| Advisor: Prof. Daniel Kressner | Germany |
| Technical University of Munich | $10 / 2016-09 / 2018$ |
| Master of Science in Mathematics |  |
| Thesis: A Surrogate Model for Bayesian Inversion Based on Sparse Grids |  |
| Thesis advisor: Prof. Barbara Wohlmuth | Germany |
| Technical University of Munich | $10 / 2013-09 / 2016$ |
| Bachelor of Science in Mathematics |  |
| Thesis: An Optimization Algorithm for Finding Invariant Sets |  |
| Thesis advisor: Prof. Oliver Junge | South-Korea |
| KAIST | $08 / 2015-12 / 2015$ |

## Professional Experience

| École Polytechique Fédérale de Lausanne | Switzerland <br> since $01 / 2019$ |
| :--- | ---: |
| Research Assistant | Germany |
| Technical University of Munich | $09 / 2016-01 / 2018$ |

ESG Elektroniksystem- und Logistik-GmbH
Germany
Research and Development Intern

## Publications and Preprints

C. Strössner and D. Kressner, Low-rank tensor approximations for solving multimarginal optimal transport problems. SIAM J. Imaging Sci. 2023 16:1, 169-191. DOI. GitHub.
C. Strössner and D. Kressner, Fast global spectral methods for three-dimensional partial differential equations. IMA J. Numer. Anal. 2022. DOI. GitHub.
S. Dolgov, D. Kressner and C. Strössner, Functional Tucker approximation using Chebyshev interpolation, SIAM J. Sci. Comput. 2021 43:3, A2190-A2210. DOI. GitHub.
J. Dabaghi, V. Ehrlacher and C. Strössner, Computation of the self-diffusion coefficient with low-rank tensor methods: application to the simulation of a cross-diffusion system, To appear in ESAIM Proc. Surveys. arXiv. GitHub.
J. Dabaghi, V. Ehrlacher and C. Strössner, Tensor approximation of the self-diffusion matrix of tagged particle processes. To appear in J. Comput. Phys.. arXiv. GitHub.
C. Strössner, B. Sun and D.Kressner, Approximation in the extended functional tensor train format. Submitted. arXiv. GitHub.

## Awards

Dean's Award for Excellence in Teaching 2021-22

## Scholarships

Max Weber Program
KAIST Scholarship Award

## Memberships

SIAM member

## Conferences, Seminars, Summer Schools and Workshops

Workshop on Low-rank Models and Applications 22
Talk: Low-rank tensor approximations for solving multi-marginal optimal ..... 2022transport problems
Swiss Numerics Day 2022 ..... Zurich
Poster: Low-rank multi-marginal optimal transport ..... 2022
ApplMath22 BrijuniPoster: Low-rank multi-marginal optimal transport2022
MATHICSE Retreat Villars-sur-Ollon
Talk: Functional low-rank approximations based on tensorized Chebyshev interpolation
CERMICS Applied Mathematics SeminarMarne-la-Vallée
Invited Talk: Functional low-rank approximations for trivariate functions ..... 2021
based on tensorized Chebyshev polynomials
Joint Group Day 2021Lausanne
Talk: Functional Tucker approximations: Novel constructors and solvers ..... 2021
Swiss Numerics Day 2021 Lausanne2021
Matrix Equations and Tensor Techniques IX Perugia
Talk: Solving PDEs on hypercubes using Chebyshev interpolation ..... 2021
CEMRACS Research Project Luminy ..... 2021 with tensor methods.
CEMRACS 2021 Summer School Luminy ..... 2021
SIAM Conference on Applied Linear Algebra (LA21) ..... virtual2021
GAMM 2020@21 virtual
Invited Talk: Functional Tucker approximation using Chebyshev interpolation ..... 2021
Communications in NLA ..... virtual
Talk: Functional Tucker approximation using Chebyshev interpolation ..... 2021
Low-Rank Models 2020
Villars-sur-Ollon
Poster: Chebychev interpolation in Tucker format ..... 2020
Model Order Reduction Summer School Eindhoven2019
8th Workshop on High-Dimensional Approximation ..... Zürich2019
MATHICSE Retreat
Champéry
Talk: Interpolation based on low-rank tensors ..... 2019
Swiss Numerics Day 2019 Lugano2019
Winter School on Hierarchical Matrices ..... Kiel

## Teaching Assistantship at EPFL

## Analyse IV

Analysis 1
Fall 2021
Numerical Analysis
Spring 2021
Analysis 1
Probabilities and Statistics II
Fall 2020
Probabiries and Statistics II
Spring 2020
Discrete Mathematics
Fall 2019
Probabilités et Statistique
Spring 2019

## Co-supervised Semester Projects at EPFL

Algorithms for image deconvolution
Functional tensor train approximation using Chebyshev interpolation Nonnegative low-rank matrix approximation for nonnegative matrices Mixed precision algorithms for fast matrix factorizations
The Sinkhorn algorithm for optimal transport

Spring 2022
Fall 2021
Spring 2021
Fall 2020
Fall 2020


[^0]:    *Note that the accuracy verification in Phase 3 fails once for this function. Here we only consider to block term decomposition obtained after restarting the procedure.

[^1]:    ${ }^{\dagger}$ The MATLAB code to reproduce these results is available from https://github.com/cstroessner/ Chebfun3F.

[^2]:    *The TT toolbox is available from https://github.com/oseledets/TT-Toolbox.
    ${ }^{\dagger}$ The MATLAB and Python code to reproduce these results is available from https://github.com/ cstroessner/EFTT.

[^3]:    $\ddagger$ The c3py package
    Th
    mpressed-Continuous-Computation is available from https://github.com/goroda/

[^4]:    *The MATLAB code to reproduce these results is available from https://github.com/cstroessner/ Optimal-Transport.git.

[^5]:    ${ }^{*}$ The code to reproduce these results is available from https://github.com/cstroessner/SelfDiffusion

[^6]:    ${ }^{\dagger}$ The code to reproduce these results is available from https://github.com/cstroessner/ SelfDiffusionCoefficent.git

[^7]:    *The MATLAB code to reproduce these results is available from https://github.com/cstroessner/ SpectralMethod3D.

