Workshop

Low-rank Optimization and Applications

Hausdorff Center for Mathematics
University of Bonn

June 8-12, 2015

Organizers

Pierre-Antoine Absil (UC Louvain)
Daniel Kressner (EPF Lausanne)
André Uschmajew (Universität Bonn)

Low-rank Optimization and Applications

The problem of how to tackle problems of "high dimensionality" has become a research field of major importance in applied mathematics, statistics, signal processing, machine learning, complexity theory, and other fields. To a large extent this development is stimulated by the rapidly growing capacities of modern computers to collect and process data. One similarity between different communities are model reducing assumptions that involve notions of low-rank, sparsity, or other kinds of separability. They lead to deep theoretical questions at the intersection of geometry and approximation theory, require novel optimization methods, and pose challenging tasks on the level of implementation.

The topic of this workshop is low-rank approximation techniques with an emphasis on tensors, optimization theory, and applications in data processing and scientific computing. The goal is to foster collaboration and exchange knowledge from the different areas.

Program Overview

Time	Monday	Tuesday	Wednesday	Thursday	Friday
09:00 - 09:20	Opening				
09:20 - 10:00	Van Huffel	Rauhut	Landsberg	Saad	Verstraete
10:00 - 10:40	De Lathauwer	Usevich	Mohlenkamp	Dhillon	Nouy
10:40 - 11:10	Coffee break				
11:10 - 11:50	Oseledets	Schneider	Hackbusch	Fazel	Tyrtyshnikov
11:50 - 12:30	Lim	Grasedyck	Kollath	Gillis	Lee
12:30 - 14:30	Lunch break		Excursion	Lunch break	
14:30 - 15:10	Bachmayr	Zhang		Mishra	
15:10 - 15:50	Schwab	Nakatsukasa		Vandereycken	
15:50 - 16:20	Coffee break			Coffee break	
16:20 - 17:00	Ballani	Markovsky		Boumal	
⋮		Poster session			
20:00 -			Dinner		

Program

Monday, June 8

9:00 - 9:20 Opening

9:20 - 10:00 Sabine Van Huffel
The Power of Low Rank Matrix and Tensor Approximations in Smart Diagnostics

10:00 - 10:40 Lieven De Lathauwer
From tensor decomposition to coupled matrix/tensor decompositions

10:40 - 11:10 *Coffee Break*

11:10 - 11:50 Ivan Oseledets
Low-rank approximation of matrices and tensors: new application areas?

11:50 - 12:30 Lek-Heng Lim
Structured Matrix Computations via Tensor Decompositions

12:30 - 14:30 *Lunch*

14:30 - 15:10 Markus Bachmayr
Adaptive Low-Rank Solvers for High-Dimensional Operator Equations

15:10 - 15:50 Christoph Schwab
Quantized Tensor FEM for Elliptic Boundary Value Problems

15:50 - 16:20 *Coffee break*

16:20 - 17:00 Jonas Ballani
Hierarchical tensor approximation of parameter-dependent PDEs

Tuesday, June 9

9:20 - 10:00 Holger Rauhut
Low rank matrix and tensor recovery

10:00 - 10:40 Konstatin Usevich
Quasi-Hankel low-rank matrix completion: a convex relaxation

10:40 - 11:10 *Coffee Break*

11:10 - 11:50 Reinhold Schneider
Convex optimization for hierarchical tensor representations and tensor network states

11:50 - 12:30 Lars Grasedyck
Parallel Tensor Sampling in Hierarchical Low Rank Formats

12:30 - 14:30 *Lunch*

14:30 - 15:10 Shuzhong Zhang
On CP-Rank Approximation and Low CP-Rank Tensor Completion

15:10 - 15:50 Yuji Nakatsukasa
Finding a low-rank basis in a matrix subspace

15:50 - 16:20 *Coffee break*

16:20 - 17:00 Ivan Markovsky
System identification in the behavioral setting: A structured low-rank approximation approach

17:00 – *Poster Session,* *see page 10.*

Wednesday, June 10

9:20 - 10:00 Joseph Landsberg
Abelian tensors

10:00 - 10:40 Martin J. Mohlenkamp
The Optimization Landscape for Fitting a Rank-2 Tensor

10:40 - 11:10 *Coffee Break*

11:10 - 11:50 Wolfgang Hackbusch
Interconnection of tensor spectra and tensor ranks

11:50 - 12:30 Corinna Kollath
Using the matrix product state formalism to investigate non-equilibrium situation in cold atomic gases

12:30 - 20:00 *Lunch & Excursion*

20:00 – *Dinner*

Thursday, June 11

9:20 - 10:00 Yousef Saad
Computing Approximate Spectral Densities with Applications

10:00 - 10:40 Inderjit S. Dhillon
Bilinear Prediction using Low Rank Models

10:40 - 11:10 *Coffee Break*

11:10 - 11:50 Maryam Fazel
Variational Gram functions: convex analysis and optimization

11:50 - 12:30 Nicolas Gillis
Geometric Aspects of Nonnegative Matrix Factorization

12:30 - 14:30 *Lunch*

14:30 - 15:10 Bamdev Mishra
Riemannian preconditioning

15:10 - 15:50 Bart Vandereycken
Preconditioned Low-Rank Riemannian Optimization for Operators with Tensor Product Structure

15:50 - 16:20 *Coffee break*

16:20 - 17:00 Nicolas Boumal
A tale of convex and Riemannian geometry for low-rank optimization

Friday, June 12

9:20 - 10:00 Frank Verstraete
Quantum tensor networks

10:00 - 10:40 Anthony Nouy
Greedy algorithms for low-rank tensor approximation in hierarchical Tucker formats

10:40 - 11:10 *Coffee Break*

11:10 - 11:50 Eugene Tyrtyshnikov
Tensor Trains for the construction of multivariate optimization algorithms

11:50 - 12:30 Namgil Lee
Low-Rank Tensor Networks for Large-Scale Optimization Problems: Future Perspective and Challenges

Posters

Sergey Dolgov
Alternating iteration for low-rank solution of high-dimensional equations

Philippe Dreesen
Decoupling nonlinear functions using tensor decompositions

Mike Espig
On the Convergence of Alternating Least Squares Optimisation in Tensor Format Representation

Sebastian Garreis
Low-rank Tensor Methods in PDE-Constrained Optimization with Uncertainties

Christian Grussler
On optimal low-rank approximation of non-negative matrices

Wen Huang
A Riemannian Optimization Technique for Rank Inequality Constraints

Vladimir Kazeev
Approximation of Singularities by Quantized-Tensor Finite Elements

Sebastian Kraemer
A Variant of Alternating Least Squares Tensor Completion in TT-format

Dana Lahat
Multi-Set Data Analysis and Simultaneous Matrix Block Diagonalization

Namgil Lee
Singular Value Decomposition for Big Data Matrix Based on Tensor Train Networks

Robert Luce
A Projected Gradient Algorithm for Near Separable NMF

Max Pfeffer
Different Metrics for Riemannian Optimization on Tensor Train Component Manifolds

Maxim Rakhuba
Grid-based electronic structure calculations: the low-rank approach

Curt Da Silva
Irregular grid tensor completion

Michael Steinlechner
Riemannian Optimization for High-Dimensional Problems

Darya Sushnikova
Fast block-low rank direct solvers for sparse matrices

Arnaud Vandaele
Heuristics for Exact Nonnegative Matrix Factorizations

Tatsuya Yokota
Smooth PARAFAC Decomposition for Tensor Completion

Abstracts

Monday, 9:20 - 10:00

The Power of Low Rank Matrix and Tensor Approximations in Smart Diagnostics

Sabine Van Huffel
KU Leuven, Belgium

An overview of applications in Medical Diagnostics including low rank approximations into their core is presented. Accurate and automated extraction of clinically relevant information from patient recordings requires an ingenious combination of adequate pretreatment of the data (e.g. artefact removal), feature selection, pattern recognition, decision support, up to their embedding into user-friendly user interfaces.

The underlying computational problems can be solved by making use of low rank matrix and tensor approximations as building blocks of higher-level signal processing algorithms. A major challenge here is how to make the mathematical decompositions "interpretable" such that they reveal the underlying clinically relevant information and improve medical diagnosis. The addition of relevant constraints can help to achieve this.

The application of these approximations and their benefits will be illustrated in a variety of case studies, including neonatal brain monitoring and brain tumour tissue differentiation.

Monday, 10:00 - 10:40

From tensor decomposition to coupled matrix/tensor decompositions

Lieven De Lathauwer
KU Leuven, Belgium

Decompositions of higher-order tensors are becoming more and more important in signal processing, data analysis, machine learning, scientific computing, optimization and many other fields. As a current trend, coupled matrix/tensor decompositions (i.e., decompositions of multiple matrices and/or tensors with one or more factors in common) are now emerging. Applications can be found in various fields and include recommender systems, advanced array processing systems, multimodal biomedical data analysis and data completion. We give a short overview and discuss the state-of-the-art in the generalization of results for tensor decompositions to coupled matrix/tensor decompositions. We briefly discuss the remarkable uniqueness properties, which make these decompositions important tools for signal separation. Factor properties (such as orthogonality, but also nonnegativity, exponential structure, etc.) may be imposed when useful but are not required for uniqueness per se. Also remarkable, in the exact case the decompositions may under mild conditions be computed using only tools from standard linear algebra. We touch upon the computation of inexact decompositions via numerical optimization. We illustrate some of the ideas using Tensorlab, a Matlab toolbox for tensors and tensor computations that we have recently released, and

of which version 2 provides a comprehensive framework for the computation of (possibly constrained) coupled matrix/tensor decompositions.

Monday, 11:10 - 11:50

Low-rank approximation of matrices and tensors: new application areas?

Ivan Oseledets
Skolkovo Institute of Science and Technology, Russia

The low-rank approximation of matrices and tensors play a crucial role in different areas, but what is surprising, many results have been obtained independently in different communities.

In this talk I will focus on our recent result in the application of low-rank formats for machine learning problems, including completion, all-subsets regression and classification problems. There, the low-rank structure acts as a regularization, and the resulting mathematical problem is the optimization with low-rank constraints, which has to be solved using the machinery of manifold optimization.

Monday, 11:50 - 12:30

Structured Matrix Computations via Tensor Decompositions

Lek-Heng Lim
University of Chicago, USA

It is well-known that the asymptotic complexity of matrix-matrix product and matrix inversion is given by the rank of a 3-tensor, recently shown to be at most $O(n^{2.3728639})$ by Le Gall. This approach is attractive as a rank decomposition of that 3-tensor gives an explicit algorithm that is guaranteed to be the fastest possible and its tensor nuclear norm (according to Grothendieck and Schatten, not the mean of matrix nuclear norms obtained from flattening the tensor into matrices) quantifies the optimal numerical stability. There is also an alternative approach due to Cohn and Umans that relies on embedding matrices into group algebras. We will see that the tensor decomposition and group algebra approaches, when combined, allow one to systematically discover fast(est) algorithms. We will determine the exact tensor ranks, and correspondingly the fastest algorithms, for Circulant, Toeplitz, Hankel, Symmetric, and other structured matrix computations. This is joint work with Ke Ye (Chicago).

Monday, 14:30 - 15:10

Adaptive Low-Rank Solvers for High-Dimensional Operator Equations

Markus Bachmayr
UPM Paris 6, France

When low-rank tensor approximations to solutions of high-dimensional problems are computed by iterations that involve rank truncations, a crucial problem is to ensure that while convergence remains guaranteed, the arising ranks cannot become too large. In this talk, we consider two construction principles for iterations which provably yield such rank bounds, and how a posteriori error control and adaptivity can be realized in this context.

(Based on joint work with Wolfgang Dahmen and with Reinhold Schneider.)

Monday, 15:10 - 15:50

Quantized Tensor FEM for Elliptic Boundary Value Problems

Christoph Schwab
ETH Zurich, Switzerland

Solutions of elliptic boundary value problems in polygonal/ polyhedral domains are known to be piecewise analytic, with point or line singularities at corners (and edges in three space dimensions) of the domain.

We prove that, generically, solutions of second order, elliptic boundary value problems in polygonal/ polyhedral domains admit exponentially convergent low-rank approximations in the tensor train (TT) format. This implies exponential convergence of first order FEM with TT compression.

We report numerical examples of tensor structured Galerkin Finite Element Methods with tensor structured DMRG and AmEN system solvers which realize these rates. Work and memory scale poly-logarithmically in the target accuracy.

Joint work with Vladimir Kazeev (SAM, ETH). Work supported in part by the European Research Council (ERC) under AdG 247277.

Monday, 16:20 - 17:00

Hierarchical tensor approximation of parameter-dependent PDEs

Jonas Ballani
EPF Lausanne, Switzerland

Parametric PDEs appear in a large number of applications, as e.g. in uncertainty quantification or optimisation.

In a high-dimensional parameter regime, special numerical techniques are needed to avoid an exponential growth in the computational complexity with respect to the parameter dimension. In this talk, we will discuss low-rank tensor techniques that allow to represent approximate solutions with linear complexity in the parameter dimension.

In particular, our aim is to adaptively construct an approximation of the solution in the hierarchical tensor format from a relatively small set of data samples. Once this approximation from an offline computation is available, the evaluation of quantities of interest becomes a cheap online task. Moreover, the explicit tensor representation can be used to compute stochastic properties of the solution in a straightforward way. The potential of this approach is illustrated by numerical examples.

This is joint work with Lars Grasedyck (RWTH Aachen) and Daniel Kressner (EPF Lausanne).

Tuesday, 9:20 - 10:00

Low rank matrix and tensor recovery

Holger Rauhut
RWTH Aachen, Germany

Low rank recovery is an extension of compressive sensing which considers the reconstruction of matrices of (approximately) low rank from incomplete linear measurements via efficient algorithms such as convex relaxation or iterative methods. Rigorous guarantees on the number of measurements that are sufficient for accurate recovery are commonly derived for certain random measurement maps. In the first part of the talk I will discuss the situation where the measurements are inner products of the unknown matrix with random rank-one matrices. This setup arises as special case in the PhaseLift approach for the phase retrieval problem where the matrix to be recovered is of rank one. The general rank-r case appears in quantum tomography. In particular, we present result on robust uniform recovery from rank one-measurements which are taken with respect to approximate 4-designs. These results are derived via Mendelson's small ball method.

In the second part of the talk I will discuss the recovery of low rank tensors from incomplete linear measurements. Unfortunately, many tasks become hard both from a computational and a theoretical perspective when passing from matrices to tensors. In fact, no complete rigorous low rank tensor recovery results are yet available for tractable algorithms. I will discuss approaches which work well in practice and provide corresponding partial theoretical results.

Tuesday, 10:00 - 10:40

Quasi-Hankel low-rank matrix completion: a convex relaxation

Konstatin Usevich
GIPSA-lab Grenoble, France

The completion of matrices with missing values under the rank constraint is a non-convex optimization problem. A popular convex relaxation is based on minimization of the nuclear norm (sum of singular values) of the matrix. For this relaxation, an important question is when the two optimization problems lead to the same solution. This question was addressed in the literature mostly in the case of random positions of missing elements and random known elements. In this contribution, we analyze the case of structured matrices with fixed pattern of missing values, in particular, the case of Hankel and quasi-Hankel matrix completion, which appears as a subproblem in the computation of symmetric tensor canonical polyadic decomposition. We extend existing results on completion of rank-one real Hankel matrices to completion of rank-r complex Hankel and quasi-Hankel matrices.

Joint work with Pierre Comon. This work is supported by ERC AdG-2013-320594 DECODA.

Tuesday, 11:10 - 11:50

Convex optimization for hierarchical tensor representations and tensor network states

Reinhold Schneider
TU Berlin, Germany

The hierarchical Tucker tensor format (Hackbusch) and a particular case, Tensor Trains (TT) (Oseledets/Tyrtyshnikov), have been introduced for high dimensional problems.The parametrization has been already known in quantum physics as tree tensor network states. There are several ways to cast an approximate numerical solution into a variational framework. The Ritz Galerkin ansatz leads to an optimization problem on Riemannian manifolds. This provides very efficient techniques. But with these techniques, or in simplyfied form with a one-site DMRG or ALS algorithm, one can be easily trapped into local minima. In the spirit of adaptive approximation we pursue a variational framework motivated by concepts of compressive sensing. We introduce a *soft shrinkage iteration* scheme based on a Hierarchical SVD (HSVD) (or Vidal decomposition). We show that the iterates converge to the unique minimum of a convex optimization problem, even if the problem is ill-conditioned. With best n-term techniques we can show a quasi-optimal error rate for the solution (compare with talk of M. Bachmayr).

Tuesday, 11:50 - 12:30

Parallel Tensor Sampling in Hierarchical Low Rank Formats

Lars Grasedyck
RWTH Aachen, Germany

We consider the problem of tensor completion or recovery from a few samples, exploiting an underlying low rank property of the tensor under consideration. The low rank condition enables us to break the curse of dimension and represent tensors by a number of parameters that depends only linearly on the dimension. First we prove complete recovery in the noiseless case under a certain full rank assumption, which is fulfiled with overwhelming probability for quasi-random tensors. Second, we consider the noisy case where the singular values are assumed to decay. Third, we prove that the recovery algorithm has optimal complexity when the underlying tensor is only characterized by the hierarchical low rank property. Finally, we give a new variant of the sampling scheme that is fully parallelizable in the sense that it scales like $1/p$ on p processors.

Tuesday, 14:30 - 15:10

On CP-Rank Approximation and Low CP-Rank Tensor Completion

Shuzhong Zhang
University of Minnesota, USA

Computing the CP (CANDECOMP/PARAFAC) rank for a given tensor is already a difficult task, let alone the low CP-rank tensor completion problem which involves optimization with the CP-rank as the objective function. In this talk we present a new matricization approach, resulting in lower and upper approximations for the CP-rank. Theoretical properties of the new ranks (to be called the M-ranks) will be discussed. The M-ranks can be applied to solve the low CP-rank tensor completion problem. The numerical results show that the new approach has clear advantages over the existing ones for some classes of the problem.

Joint work with: Bo Jiang, and Shiqian Ma

Tuesday, 15:10 - 15:50

Finding a low-rank basis in a matrix subspace

Yuji Nakatsukasa
University of Tokyo, Japan

We consider the problem of finding a basis of possibly low rank in a subspace of matrices. It turns out that if a basis of rank-one matrices exists, then it can be obtained using a tensor CP decomposition. For bases involving matrices of rank higher than one, the situation is not as straightforward. We suggest an algorithm based on a greedy process. It finds linearly independent matrices one by another, by first estimating a possible rank on the basis of soft singular value thresholding, and then computing a matrix with that rank using the method of alternating projections. We provide local convergence results for this method. Applications include data compression, accurate computation of eigenvectors of a near-multiple eigenvalue and image separation.

Joint work with Tasuku Soma (Tokyo) and André Uschmajew (Bonn).

Tuesday, 16:20 - 17:00

System identification in the behavioral setting: A structured low-rank approximation approach

Ivan Markovsky
VU Brussel, Belgium

System identification is a fast growing research area that encompasses a broad range of problems and solution methods. The behavioral setting unifies them, however till recently, it lacked supporting numerical solution methods. In the last 10 years, algorithms for structured low-rank approximation were used to fulfill this gap. In this talk, we summarize recent progress on methods for system identification in the behavioral setting and pose some open problems. First, we show that errors-in-variables and output error system identification problems are equivalent to Hankel structured low-rank approximation. Then, we outline three generic solution approaches: 1) methods based on local optimization, 2) methods based on convex relaxations, and 3) subspace methods. In order to achieve the desired unification, the classical ARMAX identification problem should also be formulated as a structured low-rank approximation problem. This is an outstanding open problem.

Wednesday, 9:20 - 10:00
Abelian tensors

JOSEPH LANDSBERG
Texas A & M University, USA

The first "non-classical" equations for tensors are due to Strassen. We call the tensors that satisfy Strassen's equations plus a mild genericity condition "abelian tensors" because of their connection with abelian algebras. I will discuss the ranks, border ranks and other geometry of these tensors. This is joint work with Mateusz Michalek, available at `http://arxiv.org/abs/1504.03732`.

Wednesday, 10:00 - 10:40
The Optimization Landscape for Fitting a Rank-2 Tensor

MARTIN J. MOHLENKAMP
Ohio University, USA

The development of effective methods for approximating a tensor with a sum of separable tensors is greatly hindered by a lack of understanding of the approximation problem they are trying to solve. As a test case, we have been studying the problem of fitting a rank-2 tensor. We consider approximations with too few terms (rank 1), the correct number of terms (rank 2), and too many terms (rank 3). I will present analysis and graphical illustrations of the resulting optimization landscapes.

Wednesday, 11:10 - 11:50
Interconnection of tensor spectra and tensor ranks

WOLFGANG HACKBUSCH
CAU Kiel & MPI Leipzig, Germany

In Hilbert tensor spaces of order d, each tensor has a SVD spectrum associated to any subset α of $\{1, ..., d\}$. For instance it is well known that, in general, the d spectra of the Tucker tensors corresponding to $\alpha = \{j\}$ are different. This leads to the question whether any combination of spectra can be obtained by suitable tensors. The ranks are the number of nonzero singular values. In the case of the hierarchical format there are subsets α and β with $\alpha \subset \beta$. The question of how the corresponding ranks are related plays a role when tensors are transferred from the TT format into the hierarchical format or the other way.

Joint work with André Uschmajew (Bonn).

Wednesday, 11:50 - 12:30

Using the matrix product state formalism to investigate non-equilibrium situation in cold atomic gases

Corinna Kollath
Universität Bonn

Using the matrix product state formalism to investigate non-equilibrium situation in cold atomic gases

Experimentally, atomic gases cooled to Nanokelvin temperatures are a new exciting tool to study a broad range of quantum phenomena. The outstanding tunability of cold gases allows to rapidly change the system parameters and to observe the subsequent quantum evolution. This poses new challenges for the understanding of quantum dynamics in correlated many-body systems to a very high accuracy.

On the theoretical side numerical methods as different variants of the matrix product state methods give the possibility to predict theoretically the dynamics of cold atomic gases. We will show how the time-dependent version can be applied to an isolated system dynamics described by the Schroedinger equation and to open quantum systems described by Markovian Master equations.

Thursday, 9:20 - 10:00

Computing Approximate Spectral Densities with Applications

Yousef Saad
University of Minnesota, USA

We will consider two types of computations that are important in various areas of physics and scientific computing in general. First is the problem of computing the spectral density of a Hermitian matrix, also known as the Density Of States (DOS) in solid state physics, which represents the probability of finding an eigenvalue at any point of the real interval. The second is the problem of counting the number of eigenvalues in a given interval. The two problems are related but the methods used are somewhat different because of the different contexts. However, virtually all methods used for the two problems rely on estimating the trace of an operator by means of random sampling. Applications of these two calculations are numerous. One of these is to determine the approximate rank of a matrix and this will be the focus of the second part of the talk.

Thursday, 10:00 - 10:40

Bilinear Prediction using Low Rank Models

Inderjit S. Dhillon
University of Texas at Austin, USA

Linear prediction methods, such as linear regression and classification, form the bread-and-butter of modern machine learning. The classical scenario is the presence of data with multiple features and a single target variable. However, there are many recent scenarios, where there are multiple target variables. For example, predicting bid words for a web page (where each bid word acts as a target variable), or predicting diseases linked to a gene. In many of these scenarios, the target variables might themselves be associated with features. In these scenarios, we propose the use of bilinear prediction with low-rank models. The low-rank models serve a dual purpose: (i) they enable tractable computation even in the face of millions of data points as well as target variables, and (ii) they exploit correlations among the target variables, even when there are many missing observations. We illustrate our methodology on two modern machine learning problems: multi-label learning and inductive matrix completion, and show results on two applications: predicting Wikipedia labels, and predicting gene-disease relationships.

This is joint work with Prateek Jain, Nagarajan Natarajan, Hsiang-Fu Yu and Kai Zhong.

Thursday, 11:10 - 11:50

Variational Gram functions: convex analysis and optimization

Maryam Fazel
University of Washington, USA

We propose a new class of penalties, called variational Gram functions, that can promote pairwise relations such as orthogonality among a set of vectors in a vector space. Applications include hierarchical classification, multitask learning, and estimation of vectors with disjoint supports. We give conditions under which the penalties are convex, show how to characterize the subdifferential and compute the proximal operator, and discuss efficient optimization algorithms. Numerical experiments on a hierarchical classification problem are also presented.

Thursday, 11:50 - 12:30

Geometric Aspects of Nonnegative Matrix Factorization

Nicolas Gillis
Université de Mons, Belgium

Given an m-by-n nonnegative matrix M and a rank r, nonnegative matrix factorization (NMF) is the problem of finding two nonnegative matrices, U with r columns and V with r rows, such that the product UV approximates M. In the exact case, that is, for $M = UV$, finding such matrices U and V is equivalent to finding a polytope with r vertices (corresponding to the convex hull of the columns of U) nested between two given polytopes (corresponding to the convex hull of the columns of M and the nonnegative orthant). In this talk, we discuss two cases where this geometric interpretation is particularly meaningful: (i) The input matrix M has a special structure, referred to as separability, where its columns are contained in the convex hull of a small subset of its columns (up to the noise level). We describe a simple algorithm to compute NMF's under this assumption (the successive projection algorithm), discuss its robustness to noise, and apply it to hyperspectral images. (ii) Computing exact NMF's is equivalent to computing extended formulations of a polytope (that is, a higher dimensional polytope that projects onto the original one) with applications in combinatorial optimization. We focus on the regular n-gons and describe how we compute new minimum size extended formulations based on the computation of exact NMF's.

This is joint work with F. Glineur (UCLouvain), W.-K. Ma (Chinese U. of Hong Kong), A. Vandaele (UMons) and S. Vavasis (U. of Waterloo).

Thursday, 14:30 - 15:10

Riemannian preconditioning

Bamdev Mishra
University of Cambridge, U.K.

We exploit a basic connection between sequential quadratic programming and Riemannian gradient optimization to address the general question of selecting a metric in Riemannian optimization, in particular when the Riemannian structure is sought on a quotient manifold. The proposed method is shown to be particularly insightful and efficient in quadratic optimization with orthogonality and/or rank constraints, which covers most current applications of Riemannian optimization in matrix manifolds.

Joint work with Rodolphe Sepulchre.

Thursday, 15:10 - 15:50

Preconditioned Low-Rank Riemannian Optimization for Operators with Tensor Product Structure

Bart Vandereycken
Univeristé de Genève, Switzerland

Solving partial differential equations on high-dimensional domains leads to very large linear systems. In these cases, the degrees of freedom in the linear system grow exponentially with the number of dimensions, making classic approaches unfeasible. Approximation of the solution by low-rank tensor formats often allows us to avoid this *curse of dimensionality* by exploiting the underlying structure of the linear operator.

We propose a truncated Newton method on the manifold of tensors of fixed rank, in particular, *Tensor Train* (TT) / *Matrix Product States* (MPS) tensors. We demonstrate the flexibility of our algorithm by comparing different approximations of the Riemannian Hessian as preconditioners for the gradient directions. Finally, we compare the efficiency of our algorithm with other tensor-based approaches such as the *Alternating Linear Scheme* (ALS).

Thursday, 16:20 - 17:00

A tale of convex and Riemannian geometry for low-rank optimization

Nicolas Boumal
INRIA Paris, France

Motivated by applications involving orthogonal matrices, I will discuss optimization over the set of positive semidefinite matrices whose diagonal blocks are small identity matrices. This is a convex set. Interestingly, the subsets of feasible matrices of bounded rank admit a nice Riemannian geometry, making it relatively simple to optimize over them directly. This is especially interesting given that the optimizers of the full problem often have low-rank (a phenomenon we will study). I will show how, combining insights from both convex and Riemannian geometry, we can attain a better understanding of this class of problems (chiefly relating critical points of the rank-restricted problem to KKT points of the full problem) and develop better algorithms.

Friday, 9:00 - 9:40

Quantum tensor networks

Frank Verstraete
Universiteit Gent, Belgium

Quantum tensor networks are revolutionizing the way in which quantum many body physics can be understood and simulated. An overview will be given of the achievements and mathematical challenges encountered in this context.

Friday, 9:40 - 10:20

Greedy algorithms for low-rank tensor approximation in hierarchical Tucker formats

Anthony Nouy
Ecole Centrale de Nantes, France

We present an algorithm for the approximation of high-order tensors in hierarchical tensor formats. The algorithm generates a nested sequence of tensor spaces based on successive greedy enrichments. A heuristic error indicator is used for the selection of the dimensions to enrich, therefore resulting in an anisotropic construction of tensor spaces. Projections in the generated tensor spaces are again approximated using low-rank formats and are computed using a DMRG-like algorithm allowing for an automatic rank adaptation.

Joint work with Loic Giraldi.

Friday, 11:10 - 11:50

Tensor Trains for the construction of multivariate optimization algorithms

Eugene Tyrtyshnikov
INM RAS Moscow, Russia

We present a novel application of tensor trains for the construction of optimization algorithms that are competitive with some widely used heuristic methods of global optimization.

Friday, 11:50 - 12:30

Low-Rank Tensor Networks for Large-Scale Optimization Problems: Future Perspective and Challenges

Namgil Lee
RIKEN Tokyo, Japan

Back Compose Reply Reply all Forward Delete Print Mark More In this talk we discuss how tensor networks can be used to solve a wide class of big data optimization problems (that are far from tractable by classical numerical methods) by applying concept of tensorization and by performing all operations using relatively small size core tensors and applying iteratively optimized tensor network contractions. We discuss application of tensor networks, especially low-rank tensor train (TT) decompositions for some very specific constrained optimization problems related to solving large-scale SVD, EVD, sparse PCA, CCA (Canonical Correlation Analysis), and linear systems for structured non-square huge matrices. The basic idea is to convert intractable large scale constrained optimization problems to set of much smaller scale linked optimization problems via contraction of suitably designed tensor networks. We discus various tensor network models and associated algorithms to solve approximately a wide class of optimization problems under conditions that given big data structured matrices and vectors admit good low rank TT approximations. The experimental results will be provided to illustrate effectiveness and some limitations of the proposed method, e.g., to estimate principal components or a few extremal singular values and corresponding singular vectors. The simulation results will be compared with existing state-of the art algorithms. Open problems will be also addressed.

Joint work with Andrzej Cichocki.

Area map

Workshop location
Mathematik-Zentrum, Endenicher Allee 60, 53115 Bonn.

Cafeteria
The closest and easiest option for lunch is the campus cafeteria (in German "Mensa"), currently located in a tent behind the main building of Mathematik-Zentrum (due to construction work at the cafeteria building). You can only pay there using a plastic "mensa card", which has to be charged at corresponding machines. Please ask for a "guest mensa card" at the cash point upon your first visit. You will have to pay a deposit of 5.10 EUR, which you will get back (together with remaining credit) after handing the card back (at the cash point) at your last visit.

Lunch in restaurants
The suggestions on the map are within 10–15 minutes walking distance. The area in Poppelsdorf ("Various bars and restaurants") offers many different options for lunch. For dinner you may explore the area around the city center.