Abstracts

Workshop on Future Directions in Tensor-Based Computation and Modeling
National Science Foundation, Arlington, Virginia, February 20-21, 2009.

Orly Alter
University of Texas (Austin)

Mathematical Modeling of DNA Microarray Data: Discovery of Biological Mechanisms with Tensor Decompositions, and Definitions of Novel Tensor Decompositions from Biological Applications

Future discovery and control in biology and medicine will come from the mathematical modeling of DNA microarray data. I will describe models that were recently created using matrix and tensor computations, where the mathematical variables, patterns uncovered in the data, correlate with activities of cellular elements. The operations, such as data reconstruction in subspaces of selected patterns, simulate experimental observation of the correlations and possibly also causal coordination of these activities. I will demonstrate the ability of these models to predict previously unknown biological as well as physical principles, and will show recent experimental results that verify one such computationally predicted cellular mechanism [1].

Future algorithms for integration and comparison of different large-scale data will come from the mathematical modeling of DNA microarray data. I will define a novel tensor decomposition and show that this decomposition extends to higher order most of the mathematical properties of the generalized singular value decomposition (GSVD) [2]. The GSVD enabled comparative modeling of DNA microarray data from yeast and human, the first comparative modeling of data from two different organisms and the only one to date that is not limited to two datasets with a one-to-one mapping between their rows, i.e., their genes. I will describe preliminary results where this novel higher-order GSVD enables comparative modeling of data from complete genomes of more than two organisms.

[1] Omberg, Meyerson, Kobayashi, Drury, Diffley & Alter, submitted.
[2] Ponnapalli, Saunders, Golub & Alter, submitted.

NOTES:

Evrim Acar Ataman
Sandia National Laboratory

Challenges in Tensor Mining

Data in many disciplines inherently has more than two axes of variation and can be arranged as tensors (or multiway arrays). Tensor decompositions have proven to be successful in extracting the underlying structure in such datasets. However, analyzing tensors is still challenging. Algorithms fitting tensor models depend heavily on the initial set-up, e.g., number of components and initialization of the component matrices. Also, models handling missing data and supervised learning are still to be developed. We discuss these challenges with example applications from computational neuroscience and social network analysis.

NOTES:

Brett W. Bader
Sandia National Laboratories

Unusual Tensor Decompositions for Informatics Applications

This presentation will introduce two tensor decompositions that are not well known but have proven useful in informatics applications. Three-way DEDICOM (decomposition into directional components) is an algebraic model with similarities to multidimensional scaling for the analysis of asymmetric 3-way arrays. PARAFAC2 is a modification of the popular PARAFAC (parallel factors) model that is less constrained and allows for different objects in one mode. Applications of both models to informatics problems will be shown.

NOTES:

Zhaojun Bai
University of California (Davis)

Multi-Length Scale Matrix Computations and Applications in Quantum Mechanical Simulations

Multi-length scale matrices refer to a class of matrices where dimensionality, eigenvalue and singular value distributions of the matrices can be characterized by a set of parameters. In the parameter space of interest, the matrices could be extremely large in dimensionality and ill-conditioned. In this talk, we will first describe a range of our synergistic activities in recent years on development of robust and efficient linear algebra solvers and high-performance software for multi-length scale matrix computation problems arising from quantum mechanical simulations of materials. We will also discuss the feasibility, through the use of tensor-based computation, of 10000-electron simulations on regular desktop PCs to allow physicists and chemists addressing important questions concerning properties of materials, such as magnetism and superconductivity.

NOTES:

Gregory Beylkin
University of Colorado (Boulder)

Nonlinear Approximations, Multi-linear Tools and Algorithms with Finite but Arbitrary Accuracy.

There are numerous instances of algorithms implemented repeatedly (with some variations) in different applications by different users. On the other hand, algorithms like the FFT or those of dense linear algebra (LAPACK) enjoy multiple users and relatively few professional implementers.

In order to have multiple users of the same software across various applications, we need adaptive algorithms that assure accuracy and map well onto some "standard" data structure. This, in turn, requires a systematic method of approximating and representing operators, in particular, those of mathematical physics.

This talk will discuss several successful examples of finding efficient approximations of operators and outline issues on the road ahead.

NOTES:

Lieven De Lathauwer
K.U. Leuven

An Introduction to Tensor-Based Independent Component Analysis

Independent component analysis aims at the blind separation of a mixture of statistically independent signals. Algebraic solutions are based on the computation of tensor decompositions. The actual source separation involves variants of the candecomp/parafac decomposition, with or without orthogonality constraints. Dimensionality reduction can be achieved by means of Tucker compression.

NOTES:

Inderjit S. Dhillon
University of Texas (Austin)

Newton-type Methods for Non-negative Tensor Approximation

Nonnegative Tensor Approximation is a decomposition technique that is useful in a wide variety of applications ranging from document analysis and image processing to bioinformatics. There are existing algorithms for nonnegative tensor approximation (NNTA), for example, Lee & Seung's multiplicative updates, alternating least squares, and certain gradient-descent based procedures. However, most existing procedures suffer from slow convergence or numerical instability. In this talk, I will present improved Newton-type algorithms for the NNTA problem, which overcome many computational deficiencies of existing methods. In particular, our methods use non-diagonal gradient scaling for faster convergence. These methods provide numerical results that are superior to both Lee & Seung's method as well to the alternating least squares (ALS) heuristic, which is known to work well in some situations but has no theoretical guarantees. I will present experimental results on both synthetic and real-world datasets to demonstrate the effectiveness of the new methods.

This is joint work with Dongmin Kim and Suvrit Sra.

NOTES:

Chris Ding
University of Texas (Arlington)

Tensor Clustering and Error Bounds

We show that widely used tensor decompositions such as HOSVD and ParaFac have clustering capabilities. More precisely, we prove that HOSVD/ParaFac objective functions are equivalent to relaxed K-means clustering, in the same framework that nonnegative matrix factors and PCA relate to K-means clustering. Error analysis provides insights to tensor decompositions. We derive tight error bounds for both HOSVD and ParaFac, generalizing Eckart-Young Theorem for SVD to these tensor decompositions. We present experiments on several real-life datasets to demonstrate the usefulness of these new theoretical results.

NOTES:

Petros Drineas
Rensselaer Polytechnic Institute

Tensor Decompositions from a Theoretical Computer Science Perspective

A large and particularly interesting class of NP-hard problems (the so-called Max-SNP problems) can be formulated using tensors, and approximation algorithms for such problems can be designed using tensor decompositions. Recent work in this area provided some provably accurate approximation algorithms for simple tensor decompositions that (unlike most other work) do not operate by reshaping (``flattening'') the tensors in order to produce matrices. In this talk, we will discuss such decompositions and their potential applications.

NOTES:

Lars Eldén
Linkoping University

Krylov Methods for Tensors I

Sparse tensors occur frequently in information sciences. We investigate different approaches for generalizing Krylov methods to tensors. We also discuss the applicability of Krylov-Schur methods. The purpose is to compute low-rank Tucker models of huge tensors. In addition, perturbation theory for the Tucker model is outlined.

NOTES:

Christos Faloutsos
Carnegie Melon University

Mining Graphs and Tensors

How do real graphs look like? How about time-evolving graphs, modeled as tensors? We present several patterns that graphs and tensors obey, and we show how to generate realistic tensors that follow such patterns. Finally, we present some results of mining real graphs using the 'hadoop' architecture (open source map/reduce).

NOTES:

Shmuel Friedland
University of Illinois (Chicago)

Results and Problems for 3-tensors

In this talk I state some results and open problems for 3-tensors that I worked on in the last couple of years.

1. The conjectured value of the generic, (border), rank of 3-tensor and its numerical verification over **C**. What about over **R**? Very partial results.
2. Best rank one approximation.
   - Restatement of L2 rank one approximation as a maximal problem. The corresponding generalization of singular vectors and singular values. Open problem: How many critical points, i.e. singular values, are in the generic case?
   - The analogous maximal problem for any Lp, p\in (1,\infty)$ norm. The corresponding p singular vectors and singular values.
   - Surprise: Existence of Perron-Frobenius theorem for irreducible nonnegative tensors for $p=3$, and more generally for $p\ge 3$, and nonexistence of Perron-Frobenius theorem for positive tensors for $p=2$, probably for $p<3$.
3. Analogs of SVD decomposition of 3-tensors.
   - The maximal number of zero entries in 3-tensor under the orthogonal conjugation in each of 3-modes.
   - The expected limit form of the tensor under the iteration an analog of QR algorithm.
   - An analog of Kogbetliantz's Algorithm

NOTES:

Robert J. Harrison
Oak Ridge National Laboratory

From Math to Peta-app: Challenges in Practical Computation with Tensor-Based Algorithms

Tensor-based models and computations arise in many areas of simulation in chemistry and physics yet there are still many fundamental and technical challenges to turning these formulations into practical and efficient simulation tools. I will discuss several such applications emphasizing research opportunities in math and C/S as well as the dearth of reliable, efficient software.

NOTES:

Manal Helal
University of New South Wales

Lenore Mullin
National Science Foundation

Tensor Computation on the High Performance Machines

High dimensional computational problems suffer from the dimensionality curse, leading the research towards dimensionality reduction and heuristics. However, some problems deviate significantly from the optimal solution when these methods are applied. On the other hand, mapping some problems to higher dimensions increases its chances for classification usling methods like Support Vector Machines. For example, Multiple Sequence Alignments (MSA) in computational biology is optimally solved using dynamic programming algorithm. However, exponential growth in the scoring tensor limits the number and lengths of sequences to a very small number. No dimensionality reduction method can be applied to this kind of scoring problems, while other heuristic methods introduce bias in directing the scoring through to a final solution. The advances in storage technology and parallel processing through high performance computers (HPC), clusters of computers, grids, and clouds, using the peta-scale technology, offer opportunities to address the dimensionality curse with new computational methods. This work implemented a tensor partitioning scheme for parallel processing and applied it on the dynamic programming MSA.

NOTES:

Anthony Kennedy
University of Edinburgh

Software Challenges in Computational Science

How can we write scientific programs that satisfy the computer scientists' goals of portability, reusable, reliability, correctness, and modularity together with the computational scientists' need for efficiency (especially on massively parallel multi-core architectures). There has been a lot of work on this in the lattice field theory community as part of the DOE SciDAC initiative, but it would clearly make sense for there to be a more general framework for expressing numerical algorithms. There are several significant issues, including abstraction of mathematical structures (such as linear spaces), memory management, data layout, and data exchange (XML schema and the like). Given the effort spent in designing and building supercomputers for QCD, physicists have been loath the lose even a few percent efficiency in exchange for the manifest advantages of portable software, and thus they have been highly motivated to look for software frameworks that achieve both.

In the UK we are setting up a new project with the goal of bridging the gap between new algorithms created by Numerical Analysts and efficient practical implementations of them on modern architectures. There are two principal hurdles to be overcome in order to solve this problem, a technical one of designing the right structures/interfaces/annotations that capture the information necessary for compilers to generate efficient codes on a wide range of architecture, and a "sociological" one of getting these accepted by users and implemented on a wide range of platforms so that they are used in real applications. This will require the joint efforts of Numerical Analysts, Computer Scientists, application experts, and supercomputer vendors, and support from funding agencies to overcome the extra risks of using new software approaches in leading-edge computational projects. Furthermore international collaboration in this effort is of vital importance.

NOTES:

Tamara G. Kolda
Sandia National Laboratories

Optimization Approaches for Solving Tensor Decompositions Models and Tensor Eigenvalue Problems

Tensor decompositions (e.g., higher-order analogues of matrix decompositions) are powerful tools for data analysis, especially as the size and complexity of data grow. Thus far, the most popular methods for fitting tensor decompositions rely on alternating optimization – typically alternating least squares (ALS). This can be thought of as a nonlinear generalization of block Gauss-Seidel because the idea is to solve for one block of variables at a time while holding all the others fixed. My team has been studying optimization approaches that solve for all variables simultaneously. Although this leads to a complex nonlinear optimization problem, our initial results indicate that all-at-once optimization for fitting CANDECOMP/PARAFAC is competitive with ALS. Others have achieved promising results for Tucker; see, e.g., Eldén and Savas (2007). I propose to discuss optimization approaches and their implications to the domain of tensor decompositions. *Joint work with Evrim Acar and Daniel Dunlavy, Sandia National Laboratories.*

NOTES:

Julien Langou
University of Colorado (Denver)

Communication Avoiding and Tiled Algorithm for "2D" Linear Algebra.

We will summarize our recent algorithms that accelerates dense "2D" linear algebra computation on multicore processors and parallel distributed machines. Both algorithms are somehow related to well-known updating factorization techniques.

NOTES:

Lek-Heng Lim
University of California (Berkeley)

Alternating Tensors and Rank Learning

Consider a collection of voters (e.g. viewers) each having rated a small fraction of a collection of alternatives (e.g. movies). Pairwise and triplewise comparisons of the ratings encode essentially all relevant information required to measure consensus and inconsistencies. These are naturally modelled by alternating tensors (skew-symmetric matrices and hypermatrices). The Hodge decomposition may then be applied to obtain a global ranking and an accompanying `certificate of reliability'. This is joint work with Y. Yao.

NOTES:

Michael W. Mahoney
Stanford University

Three Right Directions and Three Wrong Directions for Tensor Research

Much of the recent work in tensors has been motivated by large-scale scientific and Internet data analysis applications.   I will offer a few perspectives on large-scale scientific and Internet data analysis as well as my biased impression of three right directions and three wrong directions for tensor research in the upcoming years.

NOTES:

Carla Martin
James Madison University

Developing Tensor Operations with an Underlying Group Structure

Tensor computations frequently involve factoring or decomposing a tensor into a sum of rank-1 tensors (CANDECOMP-PARAFAC, HOSVD, etc.). These decompositions are often considered as different higher-order extensions of the matrix SVD. The HOSVD can be described using the n-mode product, which describes multiplication between a higher-order tensor and a matrix. Generalizing this multiplication leads to the contracted product, a type of multiplication between two higher-order tensors. However, the contracted product does not preserve the order of a tensor and it is therefore difficult to extend other concepts of linear algebra to tensors. We describe a type of tensor-tensor multiplication that has an underlying algebraic structure, in that the set of $n\times n\times n$ tensors forms an algebraic group. Another possible higher-order extension of the SVD emerges, but more importantly new algorithms for data compression result that utilize known efficient computational methods. Furthermore, this new framework provides a way to talk about higher-order extensions of QR factorization and eigenvalue decompositions. Ideally, these ideas attempt to bridge the gap between orthogonal decompositions and rank-revealing decompositions.

NOTES:

Martin Mohlenkamp
Ohio University

Computing with Sums of Separable Functions, with Applications in Quantum Mechanics

By representing functions of many variables as sums of separable functions, one obtains a method to bypass the curse of dimensionality. I will discuss efforts to develop, understand, and use this method, both in a general context and for applications in quantum mechanics.

NOTES:

Jason Morton
Stanford University

Algebraic Models for Multilinear Dependence

We discuss a new statistical technique inspired by research in tensor geometry and making use of cumulants, the higher order tensor analogs of the covariance matrix. For non-Gaussian data not derived from independent factors, tensor decomposition techniques for factor analysis such as Principal Component Analysis and Independent Component Analysis are inadequate. Seeking a small, closed space of models which is computable and captures higher-order dependence leads to a proposed extension of PCA and ICA, Principal Cumulant Component Analysis (PCCA). Estimation is performed by maximization over a Grassmannian. Joint work with L.-H. Lim.

NOTES:

Lenore M. Mullin
NSF

James E. Raynolds
SUNY Albany

Tensors and n-d Arrays: A Mathematics of Arrays (MoA}, Psi Calculus and the Composition of Tensor and Array Operations

Our presentation will discuss the outer product/tensor product and a special case of the tensor product, the Kronecker Product: the algorithms, their origin, and optimal implementation when composed, and mapped to complex processor/memory hierarchies. We discuss how the use of MoA and the Psi Calculus, a calculus of indexing with shapes, provides optimal, verifiable, reproducible, scalable, and portable implementations of both hardware and software. This is due to the fact that we are using normal forms composed of multi-linear operations on Cartesian coordinates which are transformed into simple abstract machines: starts, stops, strides, count, up and down the processor/memory hierarchy.

NOTES:

Haesun Park
Georgia Institute of Technology

Toward Faster Nonnegative Tensor Factorization: A New Activeset type Algorithm and Comparisons

A common formulation for nonnegative matrix factorization (NMF) appears as a non-convex optimization problem and various types of algorithms have been devised to solve the problem. The alternating nonnegative least squares (ANLS) framework is a block coordinate descent approach for solving NMF, which was recently shown to be theoretically sound and empirically efficient.

In this talk, we present novel algorithms for NMF and nonnegative tensor factorization (NTR) based on the ANLS framework. Our new algorithm for NMF is built upon the block principal pivoting method for the nonnegativity constrained least squares problem that overcomes some limitations of active set methods. We show that the proposed NMF algorithm can naturally extended to obtain highly efficient NTF in the form of the PARAFAC (PARAllel FACtor) model. Our algorithms inherit the convergence theory of the ANLS framework and can easily be extended to other NMF formulations such as sparse NMF and NTF with L1 norm constraints. Comparisons of algorithms using various data sets show that the proposed new algorithms outperform existing ones in computational speed for computing NMF and NTF.

This is a joint work with Krishnakumar Balabusramanian, Hyunsoo Kim, and Jingu Kim.

NOTES:

Robert Plemmons
Wake Forest University

Some Applications of Nonnegative Tensor Factorizations to Mining Hyperspectral and Global Climate Data

Nonnegativity constraints on solutions, or approximate solutions, to numerical problems are pervasive throughout computational science and engineering. In this talk it is shown how nonnegative tensor factorizations can be used for spectral unmixing in material identification with hyperspectral data, and to analyze massive global multivariate climate datasets. Some speculations and open questions will also be discussed.

NOTES:

J. (Ram) Ramanujam
Louisiana State University

P. Sadayappan
Ohio State University

Domain-Specific Abstractions for High-Productivity, High-Performance Scientific Computing

The task of developing software for high-performance scientific computing is becoming increasingly difficult due to deepening memory hierrachy and architectural heterogeneity (e.g. multicore CPU's and GPU's). A promising approach to addressing this problem is the development of domain-specific abstractions that are both convenient for application developers as well as amenable to automated compiler transformation for efficient execution on different targets.

The talk will briefly describe two efforts in this regard: 1) Tensor Contraction Engine, a domain-specific compiler for a class of tensor expressions arising in quantum chemistry, and 2) Pluto, a source-to-source compiler for automatic parallelization and data-locality optimization of affine loop nests.

A main question of interest is: Can we define convenient and useful abstractions for the domain of tensor computations that can be automatically transformed for efficient parallel execution?

NOTES:

Phillip Regalia
Catholic University of America

Cumulant Signal Processing, Tensors, and Some Recurring Problems

We revisit some signal processing problems involving cumulants and tensor structures. The first concerns attempts to extend interpolation problems from matrix displacement theory into higher-order analogs, using tensor displacement structures. This problem is relevant whenever one seeks to fit a linear model to a multidimensional data set, yet has met with only sparse attention. We address also convergence of higher-order symmetric power methods applied to tensors. This problem arises in blind signal separation, which is solved when the signal mixture is a square invertible matrix, but remains problematic in more realistic cases where a linear demixer can at best only partially separate signals of interest. We review how certain "sign definite" properties of tensors can ensure monotonic convergence, along with some counterexamples when the sign-definite property is not satisfied, thus identifying a further open problem.

NOTES:

Berkant Savas
Linkoping University

Krylov Methods for Tensors II

Sparse tensors occur frequently in information sciences. We investigate different approaches for generalizing Krylov methods to tensors. We also discuss the applicability of Krylov-Schur methods. The purpose is to compute low-rank Tucker models of huge tensors. In addition, perturbation theory for the Tucker model is outlined.

NOTES:

Charles Van Loan
Cornell University

Numerical Linear Algebra Issues Underlying a Tensor Network Computation

A Jacobi type method is presented for computing the smallest eigenvalue and associated eigenvector of a very large Hamiltonian that is a highly structured summation of Kronecker products. The vector iterates are represented in compressed form using tensor networks. Updates involve a succession of orthogonal matrix manipulations and projections. The application is an occasion to think about the message of numerical linear algebra when it is applied in a tensor setting.

NOTES: