

Workshop Report¹

Future Directions in Tensor-Based Computation and Modeling

May 1, 2009

Evrin Acar	Robert J. Harrison	Frank Olken
Orly Alter	Manal Helal	Larsson Omberg
Brett Bader	Anthony Kennedy	Haesun Park
Zhaojun Bai	Dongmin Kim	Robert Plemmons
Gregory Beylkin	Tamara Kolda	Stefan Ragnarsson
Lieven DeLathauwer	Julien Langou	Sri Priya Ponnappalli
Inderjit Dhillon	Lek-Heng Lim	J. Ram Ramanujam
Chris Ding	Michael Mahoney	James Raynolds
Lars Eldén	Carla Martin	Phillip Regalia
Petros Drineas	Martin Mohlenkamp	P. Saday Sadayappan
Christos Faloutsos	Jason Morton	Berkant Savas
Shmuel Friedland	Lenore Mullin	Charles Van Loan

Executive Summary

High-dimensional modeling is becoming ubiquitous across the sciences and engineering because of advances in sensor technology and storage technology. Computationally-oriented researchers no longer have to avoid what were once intractably large, tensor-structured data sets. The current NSF promotion of “computational thinking” is timely: we need a focused international effort to oversee the transition from matrix-based to tensor-based computational thinking. The successful problem-solving tools provided by the numerical linear algebra community need to be broadened and generalized. However, tensor-based research is not just matrix-based research with additional subscripts. Tensors are data objects in their own right and there is much to learn about their geometry and their connections to statistics and operator theory. This requires full participation of researchers from engineering, the natural sciences, and the information sciences, together with statisticians, mathematicians, numerical analysts, and software/language designers. Representatives from these disciplines participated in the Workshop. We believe that the NSF can help ensure the vitality of “big N” engineering and science by systematically supporting research in tensor-based computation and modeling.

¹Prepared by Charles Van Loan (Workshop Organizer) with editorial assistance from Sri Priya Ponnappalli and Stefan Ragnarsson (Workshop Scribes) and financial assistance from Lenore Mullin and Frank Olken (NSF Program Managers). Award number 0908059. The Workshop was held in Arlington, Virginia at the National Science Foundation, February 20-21, 2009. The Workshop web site is at <http://www.cs.cornell.edu/cv/TenWork/Home.htm>.

PART I. INTRODUCTION

1. What Are Tensors?

A *tensor* is an element of a tensor product of vector spaces. Up to a choice of bases, this can be represented as a multidimensional array of numerical values upon which algebraic operations generalizing matrix operations can be performed. In this representation, the entries in a k -th order tensor are identified by a k -tuple of subscripts, e.g., $A(i_1, i_2, i_3, i_4)$. A matrix is a second-order tensor. A vector is a first-order tensor. A scalar is a tensor of order zero.

The discretization of a continuous multivariate function on a grid yields a tensor, e.g., $A(i, j, k, \ell)$ might house the value of $f(w, x, y, z)$ at $(w, x, y, z) = (w_i, x_j, y_k, z_\ell)$. In other settings, a tensor might capture an n -way interaction, e.g., $A(i, j, k, \ell)$ is a value that captures an interaction between four variables/factors.

Tensors have been around since the middle of the 1800s and with varying levels of abstraction, play an important role in physics, engineering, and mathematics. For example, Einstein's whole theory of relativity was written in tensor format. Our use of the term will be specific and concrete: a tensor is an n -way array of real (or complex) numbers. The manipulation of such objects involves *multilinear algebra*. For an abstract, non-numerical treatment of that subject, see Greub (1978).

2. Tensor-Based Computation Is Not New

Over the last four decades the fields of chemometrics and psychometrics developed an infrastructure for tensor-based computation, see Tucker (1966). Understanding this research and its intersection with numerical linear algebra is essential. The multiway analysis texts by Smilde, Bro, and Geladi (2004) and Kroonenberg (2008) and the survey article by Kolda and Bader (2009) are excellent for this purpose with their many references. See also the expository papers by Bro (1997) and Bro (2006). A separate literature concerned with large-scale tensor computations has developed within the quantum chemistry and electronic structures communities. See White et al (1999), Head-Gordon et al (2003), Hirata (2003), and Chan et al (2007). Each of these research threads brings something unique to the table. The different camps must intermingle so as to avoid reinvention of the wheel.

3. It Is Increasingly About Big Data and High Dimension

In his acceptance speech for the Innovations Award at KDD 2007, U. Fayyad mentioned that Yahoo! Inc. has approximately five petabytes of data obtained via crawls that involve billions web pages. Another “big data” framework arises in the analysis of large social networks, where there are hundreds of millions of nodes with billions of conversations. See Leskovec and Horvitz (2008). Collecting and storing large datasets of sensor data, social network data, fMRI medical data is easier than ever with commodity, terabyte disks. This data explosion creates deep research challenges that require scalable, *tensor-based* algorithms.

The “volume” of a tensor is the product of the component dimensions $n_1, n_2, \dots, n_d$ and therein lies the *curse of dimensionality*. In many applications $N = n_1 n_2 \cdots n_d$ is big primarily because d is big. And d is getting bigger because researchers are interested in developing more

sophisticated models that capture multiple interactions instead of idealized, overly-simplistic pairwise interactions.

The development of tensor-based methods in the numerical optimization community illustrates this point. Research in this area started with tensor methods for nonlinear equations, where the Newton iteration is augmented with a low-rank approximation to the next term in the Taylor series after the Jacobian. The technique was later extended to an optimization strategy by including low-rank approximations to the third- and fourth-order tensors beyond the Hessian matrix. See Schnabel and Frank (1884), Schnabel and Chow (1991), Chow, Eskow and Schnabel (1994), Feng and Schnabel (1996), Bouaricha and Schnabel (1997), Feng and Pulliam (1997), Bader and Schnabel (2003), and Bader (2005).

In statistical settings there has been a comparable effort to “use more” terms in the power series expansion for a multivariate function $f(x_1, x_2, \dots, x_n)$. Truncated versions of the expansion provide a framework for modeling and computation. Typically, the higher-order cumulants in the expansion are ignored, e.g., the 3rd cumulant skewness and 4th cumulant Kurtosis both of which are tensors. These tensors describe higher-order dependence of random variables and in financial modeling situations can be used to estimate higher-order portfolio statistics. Just as Principal Components Analysis (PCA) identifies components that account for variation in covariance, Principal Cumulant Components Analysis (PCCA) identifies components that take into account variation in all higher-order cumulants simultaneously. See Morton and Lim (2009). An argument can be made that the current economic crisis was prompted, in part, by the adoption of crude, tensorless, approximations of risk. See Nocera (2009).

4. A Research Community Is Building


The workshop/conference profile of tensor-based computation has steadily increased during the past five years. We mention the Tensor Decomposition Workshop (2004, American Institute of Mathematics), the Workshop on Tensor Decompositions and Applications (2005, Luminy, France), and various minisymposia held at MMDS 2006, ICIAM 2007, MMDS 2008, 2008 SIGKDD, and the 2008 SIAM National Meeting. This upward trajectory of interest will continue; the 2009 SIAM National Meeting will have several NSF-supported minisymposia relating to tensor computation.

The articles by Kilmer and Martin (2004), Drineas, Golub, Lim, and Mahoney (2006), and Mahoney, Lim, and Carlsson (2009) nicely track the research coalitions that are forming to address tensor-related issues and large-scale computation. *SIAM J. Matrix Analysis and Applications* had a special tensor issue in 2008. On the educational front, we note that the recent introductory textbook by Eldén (2007) on data mining and pattern recognition includes a chapter on tensor computation. Numerical multilinear algebra is becoming more and more fundamental and is finding its way into mainstream scientific computing curriculum.

5. A New Chapter is About to be Written

A tensor can be regarded as a higher-order matrix. Conversely, a matrix with nested block structure can be regarded as a tensor. For example, $A(1:n_1, 1:n_2, \dots, 1:n_6)$ is an n_1 -by- n_2 block matrix whose entries are n_3 -by- n_4 block matrices whose entries are n_5 -by- n_6 matrices of real numbers. Given these point-of-view options, it is no surprise that the resurgence of

tensor-based scientific computing heralds the beginning of a new chapter in the field of matrix computations, a field that seems to “kick up” its level of thinking about every twenty years:



Computational thinking in the linear algebra area has a unique way of spreading from the specific to the general. The numerical PDE community prompted the development of the first sparse matrix solvers during the 1950s and 1960s, a technology that now permeates the sciences and engineering. Similar application-driven, coming-of-age stories apply to orthogonal matrix computations (statistics, the 1970s), structured matrix computations (control engineering, the 1980s), parallel matrix computations (real-time signal processing, the 1980s), and most recently, sample-based matrix computations (information science, the 2000s).

The development of tensor-based methods is well underway in a host of application areas. It is important to identify these areas and the common ground between them from the standpoint of algorithms, analysis, and software.

6. From Matrix to Tensor: A Complex Extrapolation

While it is important to tap into the traditions of numerical linear algebra, we must be mindful that tensor-based research is not just matrix-based research with additional subscripts. The extrapolation is complex and worthy of multiple perspectives. Tensors must be appreciated as data objects in their own right. They come equipped with their own geometry, their own statistics, and a measure of nonlinearity—attributes that are frequently lost when a tensor problem is reshaped into a matrix problem for the sake of computational convenience.

It is particularly important to appreciate the algorithmic and statistical perspectives when developing a tensor-based method. The reason we represent data in matrix form is to take advantage of the nice properties of vector spaces: structural properties (SVD), algorithmic properties, and statistical properties (PCA). These provide a clean framework. They are natural mathematical structures that provide algorithmic and statistical benefits. Tensor structures provide greater descriptive flexibility, but with this comes a computational cost. It is essential to appreciate this fact.

To understand high dimensions, it is sometimes necessary to study functions and operators, and not just vectors, matrices and tensors. This perspective can shed light on intrinsic issues and yield more concise proofs. It also permits the computational scientist to focus on dimensionality issues without worrying about indices. See Beylkin and Mohlenkamp (2002, 2005).

PART II. SOME APPLICATION DOMAINS

7. Tensor-Based Data Analysis and Mining

The surging interest in tensors is prompted, in part, by a proliferation of data-intensive applications. The scientific and engineering communities are awash in a sea of high-dimensional, multi-indexed data sets. Tensor methods are required to expose the underlying patterns.

Data in many disciplines inherently has more than two axes of variation and can be arranged as tensors (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 upon the initial set-up, i.e., number of components and the initialization of the component matrices. Also, models handling missing data and supervised learning are still being developed.

The 3-dimensional case is particularly important in situations that concern sequences of related images (matrices). Traditional numerical linear algebra is not enough in these data mining applications; a comprehensive set of numerical multilinear algebra tools will be required to spot patterns in n -way data. See Drineas, Golub, Lim, and Mahoney (2006), Sun (2007), Savas (2008), Kolda and Sun (2008), Sun, Tsourakakis, Hoke, Faloutsos, Eliassio-Rad (2008), and Acar and Yener (2009). Examples of tensor-based data analysis in the medical setting include Acar, Bingol, Bingol, Bro, and Yener (2007) and De Vos, Vergult, De Lathauwer, De Clercq, Van Huffel, Dupont, Palmini, and Van Paesschen (2007).

Huang, Ding, Luo, and Li (2008) identify important connections between well-known clustering strategies and certain tensor factorizations and algorithms. Hyperspectral data analysis is also rich in tensor computations, see Zhang, Wang, Plemmons and Pauca (2008).

The act of spotting patterns and redundancies in high-dimensional data sets requires a refinement of matrix-based techniques. Paying attention to tensor geometry and making use of cumulants is important. For non-Gaussian data not derived from independent factors, tensor decomposition techniques for factor analysis such as Principal Component Analysis (PCA) and Independent Component Analysis (ICA) 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 called Principal Cumulant Component Analysis (PCCA). Estimation can be achieved by maximizing over a Grassmannian. See Morton and Lim (2009).

8. Tensor-Based Information Science

An important research thread in the evolving field of information science involves the development of mathematical tools that can be used to describe the properties and behavior of networks and graphs. Linear algebra has long played a central role. However, with increased frequency, researchers are developing models that are multilinear in nature. Consequently, there is a need to develop an infrastructure that supports tensor computation in this area. See Kolda, Bader, and Kenny (2005), Dunlavy, Kolda, Kegelmeyer (2006), and Leskovec and Faloutsos (2007). Tensors are finding application in text analysis, analysis of semantic graphs and social networks, and multilingual information retrieval. See Bader, Berry, and Browne (2007), Bader, Harshman, and Kolda (2007), and Chew, Bader, Kolda, and Abdelali (2007).

We mention that tensor factorizations were a topic of interest at the 2009 Conference on Geometrical Models of Natural Language Semantics held in Athens.

9. Tensor-Based Signal/Image Processing

Numerical multilinear techniques are permeating the field of signal processing. For a nice introduction to these developments, see De Lathauwer and De Moor(1998) and Kofidis and Regalia(2001). Tensor models have been used for the analysis of EEG/ERP signals since mid-1980s. Various computational neuroscience applications are discussed in Acar and Yener (2009). Video compression, handwritten digit classification, and blind identification have tensor formulations. See Ding, Huang, and Luo (2008), Savas and Eldén (2007), and Comon (2004).

The “tensor faces” work of Vasilescu and Terzopoulos (2002) demonstrated the advantage of tensor-based representations when dealing with ensembles of facial images. As an analytical tool, the N-mode singular value decomposition is a natural choice when there are N modes to the data set, e.g., facial geometries, expressions, head poses, lighting conditions, etc.

Independent component analysis is in the process of being generalized through the application of some well-known tensor decompositions. Matrix-displacement theory, which has prompted the development of fast signal processing algorithms in the past, has been extended to the higher-order case using tensor-displacement structures. See Grigorasu and Regalia (1999). The overall role of cumulants and higher order statistics in signal processing is growing thereby creating a demand for fast, tensor-based algorithms.

10. Tensor-Based Computational Biology

As biology becomes more combinatoric, discrete, and data-intensive, tensor analysis will grow in importance. See Alter and Golub (2005), Omberg, Golub, and Alter (2007), and Acar, et al. (2007). There is great promise for the application-driven development of new tensor mathematics and algorithms. Future algorithms for integration and comparison of different large-scale data will come from the mathematical modeling of DNA microarray data. An example of this is the development of a higher-order tensor decomposition based on the generalized singular value decomposition. Models have been created using matrix and tensor computations, where the 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. These models predict previously unknown biological as well as physical principles.

It is worth remembering that the spread of tensor-based methods in any application area depends upon the availability of high-quality, easy-to-use software. This is particularly true in biology since the mathematization of the subject is relatively new compared to chemistry and physics.

11. Tensor-Based Simulation in Chemistry and Physics

Tensor-based models and computations arise in many areas of simulation in chemistry and physics, yet there are still many fundamental and technical challenges that inhibit the production of practical and efficient simulation tools. Some of the largest matrix problems in

all of computational science and engineering arise in this area making it an excellent venue for the development of tensor methods that scale. See Grotendorst (2000), Lechner, Alic, and Husa (2004), and Hachmann and Chan (2006).

Tensor-based, quantum-mechanical simulations that involve upwards of 10000 electrons are being developed. See Bai, Chen, Scalettar, and Yamazaki (2007).

12. Tomorrow

Our tour through application areas is not meant to be exhaustive. For example, there are interesting tensor threads in the area of quantum computing. Climate modeling has huge datasets that are naturally multi-dimensional. Several areas of engineering such as control engineering, where there is a great tradition of advanced matrix computations, appear to be on the verge of “going multilinear.” The overall point we are making is that the spread of tensor-based modeling and computation is pervasive and bound to continue.

13. The Babel Factor

It is inherently more difficult to describe a tensor computation than a matrix computation. There are vectors of subscripts and an exponential number of ways to carry out operations such as transposition. Loop nestings are deeper in tensor computations with many more possible rearrangements. To cope with these challenges, each application area tends to have its own notation. For example, the Einstein and Dirac notations find favor in physics and chemistry because they offer a certain economy. The downside is that those notations inhibit the flow of algorithmic ideas to other application communities.

It is pointless to recommend the adoption of a universal tensor notation. However, it would be extremely useful to have a “thesaurus” that shows how to articulate typical calculations in the different notations and a companion collection of Latex macros to facilitate the writing of tensor-related documents. For the sake of new-idea propagation, researchers should respect the diversity of tensor notations and broaden their exposition accordingly.

PART III. SOFTWARE/LIBRARY ISSUES

14. The Software Challenge

We need to make tensor computations as “easy” as matrix computations. New libraries will be required if the latest and greatest algorithmic ideas are to be accessible to researchers in the sciences and engineering. It is unclear whether there is enough of a “common denominator” across tensor applications to warrant the development of a tensor LAPACK; packages may have to be specialized to domain areas such as biology and information science.

Portability, reusability, reliability, correctness, and modularity needs to be reconciled with the computational scientists’ need for efficiency, especially on massively parallel multi-core architectures. Tensor-based algorithms are ideal for these computing environments and algorithmic research should take advantage of these systems. Zhang, Berry, Lamb and Samuel (2009) discuss parallelizing the nonnegative tensor factorization on multi-core architectures. Additional high-performance contributions have come from the lattice field theory community

as part of the DOE SciDAC initiative. However, 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 quantum chromodynamics, physicists are reluctant to accept small degradations in performance in exchange for the long-term benefits of portable software. Thus, the development of software frameworks that addresses this particular tension is a priority.

From a software engineer’s perspective, we need to make use of high-level languages and garner respect for the importance of maintainability. A library of test problems supported by the community would facilitate the assessment of new algorithms and implementations. This was the case for the field of numerical optimization, particularly during the 1980s and 1990s. See Averick, Carter, More, and Xue (1992).

15. Matlab

The MATLAB environment is excellent for the rapid prototyping of algorithmic ideas and its widespread use makes it possible to reach application communities relatively easily. In support of tensor computation, there is the N-Way toolbox of Anderson and Bro (2000) and the tensor class work of Bader and Kolda (2006). There is also support for sparse and factored tensors. See Bader and Kolda (2007). MATLAB is frequently the environment of choice for experimentalists and theoreticians who are venturing into computation for the first time. Thus, it is important for the system to be tensor-friendly if tensor-based thinking is to propagate throughout the sciences.

16. Towards a LAPACK for Tensors

As we mentioned, it is probably too early for a formal, comprehensive LAPACK-style initiative. However, it is very important for the tensor computation community to track research developments associated with high-performance numerical linear algebra. Topics include the exploitation of the multicore architectures, the use of self-tuning linear algebra code-generators, the exploitation of recursive data structures, and the extrapolation of the BLAS philosophy. See Dongarra, Gannon, Fox, and Kennedy (2007), Demmel et al (2005), Kågström, Ling, and Van Loan (1999), and Elmroth, Gustavson, Jonsson, and Kågström (2005).

Of course, it is not enough for an implementation to be fast—it must be numerically reliable. The issue of floating point stability in tensor computations has hardly been considered. Long chains of matrix manipulations are frequently involved in multilinear settings and it is absolutely essential that we build a numerically sound infrastructure upon which the CSE community can depend. In other words, respect for the shortcomings of floating arithmetic must find its way into the domain of tensor computations. A tensor LAPACK cannot be wholly successful until this happens.

17. Domain-Specific Software Tools for Tensor Computation

Can we define convenient and useful abstractions for the domain of tensor computations that can be automatically transformed for efficient parallel execution? The task of developing software for high-performance scientific computing is becoming increasingly difficult due

to deepening memory hierarchy 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 Tensor Contraction Engine is a very successful, domain-specific compiler for a class of tensor expressions arising in quantum chemistry. See Auer, et al (2006), Baumgartner, et al (2005), Bibireata, et al (2004), and Hartano, et al (2005). In this context, there is interest in developing compilers that can automatically parallelize and optimize for data locality in codes that are rich in affine loop nests. See Lam et al (1997), Lu et al (2005), Bondhugula, Hartono, Ramanujam, and Sadayappan (2008), and Bondhugula, Baskaran, Krishnamoorthy, Ramanujam, Rountev, and Sadayappan (2008).

18. The FFT Message

Multidimensional FFT computations have a lot in common with tensor contractions in that both are (a) rich in matrix-vector products, (b) highly parallelizable, and (c) plagued with all kinds of data-locality obstacles. Moreover, in both computational settings the Kronecker product has a prominent role to play. See Van Loan (1992). For these reasons it makes sense for the tensor community to be aware of high-performance FFT research. Indeed, FFT computations can be structured using the language of tensors and this has resulted in a design framework for high-performance codes. See Granata, Conner, and Tolimieri (1992a, 1992b) and Huang, Johnson, and Johnson (1991). There may also be a lesson to learn from the wildly successful FFTW system that automates the generation of optimal implementations. See Frigo and Johnson (2005).

19. Programming Languages and Array Theory

Currently, the design of high-performance tensor software requires (a) an ability to reason at the index-level about the constituent contractions and the order of their evaluation and (b) an ability to reason at the block matrix level in order to expose fast, underlying Kronecker product-like operations. Progress in numerical multilinear algebra will be inhibited without the development of languages and systems that provide high-level support for this type of computational thinking. Notation is a major challenge in multilinear settings where vectors of subscripts and recursion rule the day.

Theorizing about indices and compositions of indices is essential and critical to the analysis and optimization of high-dimensional array computing. Determining if there are boundary conditions (and eliminating them) in the software realm is important to the development of scalable programs. Since the watershed design of APL in the 1970s, there has been a longstanding interest in the theory of arrays within the programming language community. This literature and its ramifications for tensor-based computation needs to be understood and built upon. See Jenkins and Mullin (1990), Rosenkrantz, Mullin, and Hunt (2006), and Hunt, Mullin, Rosenkrantz, and Reynolds (2008).

Of particular interest is the Psi calculus, a calculus of indexing with shapes that provides optimal, verifiable, reproducible, scalable, and portable implementations of both hardware and software. It uses normal forms composed of multilinear operations on Cartesian coordinates that are transformed into simple abstract machines: starts, stops, strides, count, up and down

the processor/memory hierarchy. All of this can be automated. The Psi reduction allows one to prove that two tensor expressions are equivalent not just at the algorithmic level but at the level where the algorithm is “lifted” to a higher dimension to reflect processor/memory levels.

The Psi Calculus tensor notation was implemented in a C++ library and successfully tested with basic image and video processing applications. See Helal (2001). The work in Helal, El-Gindy, Mullin, and Gaeta (2008) and Helal, Mullin, Gaeta, and El-Gindy (2007) further developed the Psi Calculus tensor library to solve a higher-dimensional problem of Multiple Sequence Alignment in computational biology, using the dynamic programming algorithm. To handle the computational complexity using the multicore hardware and the clusters of computing nodes, a master/slave cubical partitioning model and a peer-to-peer diagonal partitioning model were developed. Both partitioning models keep the same tensor structure and are not based on lower dimension decompositions. To further reduce the complexity, a search space reduction technique has been developed. See Helal and Sintchenko (2009).

PART IV. ALGORITHMIC CHALLENGES

20. Coping With the Curse of Dimensionality

There are several approaches to the curse of dimensionality that attend the manipulation of high-order tensors. Approximation and separability are of paramount importance. By representing functions of many variables as sums of separable functions, one obtains a method to bypass the curse of dimensionality. See Beylkin and Mohlenkamp (2002, 2005) and also Hackbush and Khoromskij (2007) and Hackbush, Khoromskij, and Tyrtshnikov (2005). Research in this direction should torque software development. 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.

Tensor networks are another vehicle for representing huge vectors that arise in the context of solving Hamiltonian eigenvalue problems in quantum chemistry. A tensor network is a way of representing a very high-order tensor by connecting many low-order tensors through contractions. In this way, vectors of order $n = 2^{100}$ (for example) can be successfully approximated with many fewer than n numbers. Data-sparse representations of vectors through tensor networks promises to be important in problems with exponentially large dimension. See Chan (2004).

21. To and From the Numerical Linear Algebra Framework

Ongoing threads of research in the field of matrix computations include (a) the search for new computable matrix decompositions that broaden the set of solvable problems, (b) the exploitation of special structures such as sparsity and symmetry, and (c) the careful framing of numerical rank and conditioning issues through the singular value decomposition. See Golub and Van Loan (1996).

Generalizations of these interconnected pursuits are evident at the tensor level. However, multilinear complexities abound making it apparent that you can only run so far with the

classical numerical linear algebra paradigm. For matrices, it is obvious what a particular decomposition reveals. This is not always true with tensors. For matrices, our interest in data-sparse representations grows with n . For tensors, it is more likely to grow with the order d and that requires very different strategies. For matrices, the concept of rank is crystal clear while for tensors it is fuzzy and ambiguous.

Research in the emerging field of numerical multilinear algebra requires a different mind set. Numerical linear algebra has a lot to offer, but it is not the whole story. The issue of flattening is a case in point. A tensor is flattened (or “matricized” or “unfolded”) by cutting it into “matrix-thin” pieces and then assembling the slices into a block matrix, e.g.,

$$A(1:n_1, 1:n_2, 1:n_3) \rightarrow [A(:, 1, :) \ A(:, 2, :) \ \cdots \ A(:, n_2, :)].$$

This kind of reshaping of a tensor opens the door to matrix computations. For example, the higher order SVD amounts to a set of flattened tensor SVDs. See De Lathauwer, De Moor, and Vandewalle (2000a). This is extremely useful in certain multiway analysis applications. However, in other settings there is a loss of information when the tensor elements are scrambled together through matrix multiplication updates of the flattening. Caveat emptor. It would be handy to have a consumer’s guide to flattening.

We should also be receptive to the fact that the “tensorization” of matrix problems is sometimes constructive. For example, a block matrix is a flattened 4-th-order tensor with the (k, ℓ) entry of block (i, j) corresponding to tensor entry (i, j, k, ℓ) . A tensor decomposition of the data set might make more sense than a matrix decomposition. We need to be able to spot tensor problems that are disguised as matrix problems.

22. The Decomposition Paradigm

The classical PARAFAC/CANDECOMP and Tucker tensor decompositions are discussed in Kolda and Bader (2009) together with several variants. See also Comon (2001). Choosing the “right” decomposition depends on the underlying application. For example, the 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. See Bader, Berry, and Browne (2007), Bader, Harshman, and Kolda (2007), and Chew, Bader, Kolda, and Abdelali (2007).

Instead of looking for the single magic decomposition, it might make “data analysis sense” to apply a range of decompositions to a given problem and draw inferences from the union of the insights that they each provide. In the mean time, the search continues for tensor-level generalizations of the QR factorization and various eigenvalue decompositions. See Kilmer, Martin, and Perrone (2008). There are also variational approaches that can be used to extend the notion of eigenvalues and singular values to tensors. See Lim (2005) and Qi (2005).

New factorizations and techniques are being invented that can extract critical information from ultra-large, tensor-structured matrices through sampling. See de La Vega, Kannan, Karpinski, and Vempala (2005), Drineas, Kannan, and Mahoney (2006), Frommer, Mahoney, and Szyld (2007), and Mahoney, Maggioni, and Drineas (2008). The importance of sampling techniques for tensors is bound to increase as problem size grows.

23. Tensor Rank

Tensor rank is a much trickier issue than matrix rank. This complicates the problem of computing the distance of a given tensor to a specified set of rank-deficient tensors. For example, a random 2-by-2-by-2 tensor has rank three with probability .79 and rank two with probability .21. This kind of a split between full rank and reduced rank does not occur with matrices. Rank-related nearness questions for tensors become ill-posed. See Kolda (2003), de Silva and Lim (2008), and Friedland (2008). More can be said when the underlying tensor has certain symmetries as shown in Kofidis and Regalia (2002) and Comon, Golub, Lim, and Mourrain (2008).

An alternative approach to tensor rank is to regard the rank of a d -th order tensor as a d -tuple of integers, see De Lathauwer, De Moor, Vandewalle (2000b) and Eldén and Savas (2009). Seeking a low-rank representation of a tensor is a way of compressing the data. A Jacobi-like procedure that embodies the compression idea is given in Martin and Van Loan (2008). We mention that various complexity issues in computer science can be posed as tensor rank problems, see Landsberg (2005).

The well-known Eckart-Young theorem can be used to express the approximation error of matrices by their SVD. Although the exact errors can not be expressed using singular values of reformulated/folded matrices of tensors, error bounds for tensor decompositions have recently been derived in Ding Huang, and Luo (2008).

24. Sparsity and Symmetry

Sparse tensors occur frequently in the information sciences (§8). Krylov methods, well known and highly successful in matrix setting, can be generalized to compute low rank approximations to sparse tensors. See Savas and Eldén (2009).

Many tensors can have interesting and complicated symmetries. However, outside of the supersymmetric case, there appears to be little algorithmic work that addresses this issue.

25. Nonnegativity

The nonnegative factorization problem for matrices involves the approximation of a nonnegative matrix A with a low rank product FG^T where F and G are themselves nonnegative. Interest in such factorizations has exploded with the growth of information science.

Nonnegative tensor decompositions are useful in a wide variety of applications ranging from document analysis to image processing to bioinformatics. They can be used for spectral unmixing in material identification with hyperspectral data, and to analyze massive global multivariate climate datasets. See Chicocki, Zdunek, Choi, Plemmons, Amari (2007). Imposing a nonnegativity constraint in a tensor factorization in image rendering allows for editing, see Lawrence (2006). Bader, Berry, and Browne (2007) use a nonnegative tensor factorization in an application that involves tracking email messages over time.

Nonnegative versions of the PARAFAC model have been developed that require the solution of a succession of nonnegative least square problems. Sparsity can be exploited. See Kim and Park and Eldén (2007), Park and Kim (2008), and Lim and Comon (2009). Improved

Newton-type algorithms for the problem are currently being developed that overcome many computational deficiencies of existing methods. See Kim, Sra, and Dhillon (2008).

26. Multilinear Optimization

Most of the optimization procedures used to compute compressed tensor representations follow the alternating least squares (ALS) approach. The attractiveness is obvious. By freezing all but one vector component, a multilinear computation becomes a linear computation. 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. For examples of this strategy, see De Lathauwer, De Moor, and Vandewalle (2000), and Martin and Van Loan (2008). Nonlinear least squares methods have been proposed for fitting the CANDECOMP/PARAFAC model. See Paatero (1999) and Tomasi and Bro (2006).

It is important for the tensor computation community to explore new ways for handling the structured nonlinearity that is the hallmark of multilinear problems. New optimization approaches for tensor computations are being developed that solve for all variables simultaneously. Although this leads to a complex nonlinear optimization problem, initial results indicate that all-at-once optimization for fitting CANDECOMP/PARAFAC is competitive with ALS. See Acar, Kolda, Dunlavy (2009). In some important exact data cases, CANDECOMP/PARAFAC just amounts to computing a matrix eigenvalue decomposition. See De Lathauwer (2006) and De Lathauwer, De Moor, and Vandewalle (2004). In the case of noisy data, these algorithms can be used for starting values.

Other approaches to compute a Tucker model are Newton and quasi-Newton methods defined on a product of Grassmann manifolds, see Eldén and Savas (2009) and Savas and Lim (2008) and the related work in Ishteva, De Lathauwer, Absil, and Van Huffel (2009).

Fitting a tensor decomposition is, after all, a non-linear optimization problem. Alternating least squares is from the 1970s and with faster machines, it is about time we had new and improved optimization strategies. Collaboration between optimization experts and tensors experts will be essential.

27. Conclusions

It would be a mistake to conclude from the outline-form of this report that there are crisply-defined subdivisions within the tensor research community. Algorithmic advances will be driven by the rich interconnections that exist between applications and the themes that we have identified. See Figure 1. Undoing the “curse of dimensionality” in any particular area of applications will undoubtedly lead to new technologies and have far-reaching implications throughout mathematics and computational science and engineering.

In many respects, the “tensor” grand challenge is to enable solutions to the grand challenge problems confronting data-deluged researchers in other fields. Funding initiatives for tensor-related research should be considered a priority given current levels of support for information technology, biotechnology, climate modeling, and other critical areas that require sophisticated modeling and the analysis of large, multidimensional datasets.

By regarding tensors as tools for describing mathematical objects in high dimensions, it is clear that the development of computational multilinear algebra should parallel the development of analytical tools for spaces of high dimension. In fact, the distinction is rather artificial




Figure 1: *The “Geometry” of the Tensor Research Community*

since nonlinear approximation is the key tool that underpins both areas of research. In a very practical sense, multilinear algebra and an appropriate approximation theory are critical to the advancement of mathematics in applications where the curse of dimensionality is the main obstacle. The workshop highlighted the breadth of these problem areas even though there was only enough time to focus on a subset of pertinent issues.

References

- E. Acar and B. Yener (2009). “Unsupervised Multiway Data Analysis: A Literature Survey,” *IEEE Transactions on Knowledge and Data Engineering*, 21(1), 6–20.
- E. Acar, C.A. Bingol, H. Bingol, R. Bro and B. Yener (2007). “Multiway Analysis of Epilepsy Tensors,” *Bioinformatics*, 23(13), i10–i18.
- E. Acar, T.G. Kolda and D.M. Dunlavy (2009). “An Optimization Approach for Fitting Canonical Tensor Decompositions,” Technical Report Number SAND2009-0857, Sandia National Laboratories, Albuquerque, NM and Livermore, CA.
- O. Alter and G.H. Golub (2005). “Reconstructing the Pathways of a Cellular System from Genome-Scale Signals Using Matrix and Tensor Computations,” *Proceedings of the National Academy of Sciences*, 102(49), 17559–17564.
- C.A. Anderson and R. Bro (2000). “The N-Way Toolbox for MATLAB,” *Chemometrics and Intelligent Laboratory Systems*, 52, 1–4.
- A. Auer, G. Baumgartner, D. Bernholdt, A. Bibireata, V. Choppella, D. Cociorva, X. Gao, R. Harrison, S. Krishnamoorthy, S. Krishnan, C. Lam, Q. Lu, M. Nooijen, R. Pitzer, J. Ramanujam, P. Sadayappan, and A. Sibiryakov (2006). “Automatic Code Generation for Many-Body Electronic Structure Methods: The Tensor Contraction Engine,” *Molecular Physics*, 104, no. 2, 211–228.
- B.M. Averick, R.G. Carter, J.J. More, and G.-L. Xue (1992). “The Minpack-2 Test problem Collection,” Argonne National Laboratory, Report ANL/MCS-TM-150.

- B.W. Bader (2005). “Tensor-Krylov Methods for Solving Large-Scale Systems of Nonlinear Equations,” *SIAM J. Numerical Analysis*, 43, 1321–1347.
- B.W. Bader, M.W. Berry, and M. Browne (2007). “Discussion Tracking in Enron Email using PARAFAC,” in *Survey of Text Mining: Clustering, Classification, and Retrieval*, Second Ed., M.W. Berry and M. Castellanos, eds., Springer-Verlag.
- B.W. Bader, R.A. Harshman, and T.G. Kolda (2007). “Temporal Analysis of Semantic Graphs Using ASALSAN,” in *ICDM 2007: Proceedings of the 7th IEEE International Conference on Data Mining*, 33–42.
- B.W. Bader and T.G. Kolda (2006). “Algorithm 862: MATLAB Tensor Classes for Fast Algorithm Prototyping,” *ACM Transactions on Mathematical Software*, 32, 635–653.
- B.W. Bader and T.G. Kolda (2007). “Efficient MATLAB Computations with Sparse and Factored Tensors,” *SIAM J. Scientific Computing*, 30, 205–231.
- B.W. Bader and R.B. Schnabel (2003). “Curvilinear LineSearch for Tensor Methods,” *SIAM J. Scientific Computing*, 25, 604–622.
- Z. Bai, W. Chen, R. Scalettar, I. Yamazaki (2007). “Robust and Efficient Numerical Linear Algebra Solvers and Applications in Quantum Mechanical Simulations,” *Proceedings of the 4th International Congress of Chinese Mathematician (ICCM)*, Edited by Lizhan Ji, Kefeng Liu, Lo Yang, Shing-Tung Yau, Vol. III, 253–268, Higher Education Press, China.
- G. Baumgartner, A. Auer, D. Bernholdt, A. Bibireata, V. Choppella, D. Cociorva, X. Gao, R. Harrison, S. Hirata, S. Krishnamoorthy, S. Krishnan, C. Lam, Q. Lu, M. Nooijen, R. Pitzer, J. Ramanujam, P. Sadayappan, and A. Sibiryakov (2005). “Synthesis of High-Performance Parallel Programs for a Class of ab initio Quantum Chemistry Models,” *Proceedings of the IEEE*, 93, no. 2, 276–292.
- G. Beylkin and M.J. Mohlenkamp (2002). “Numerical Operator Calculus in Higher Dimensions,” *Proceedings of the National Academy of Sciences*, 99(16), 10246–10251.
- G. Beylkin and M. Mohlenkamp (2005). “Algorithms for Numerical Analysis in High Dimensions,” *SIAM J. Scientific Computing*, 26, 2133–2159.
- A. Bibireata, S. Krishnan, G. Baumgartner, D. Cociorva, C. Lam, P. Sadayappan, J. Ramanujam, D. Bernholdt, and V. Choppella (2004). “Memory-Constrained Data Locality Optimization for Tensor Contractions,” in *Languages and Compilers for Parallel Computing*, (L. Rauchwerger et al Eds.), Lecture Notes in Computer Science, Vol. 2958, 93–108, Springer-Verlag.
- U. Bondhugula, A. Hartono, J. Ramanujam, and P. Sadayappan (2008). “A Practical and Automatic Polyhedral Program Optimization System,” *Proc. ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation (PLDI 08)*, Tucson.
- U. Bondhugula, M. Baskaran, S. Krishnamoorthy, J. Ramanujam, A. Rountev, and P. Sadayappan (2008). “Automatic Transformations for Communication-Minimized Parallelization and Locality Optimization in the Polyhedral Model,” in *Proc. CC 2008 -International Conference on Compiler Construction*, Budapest.
- A. Bouaricha and R.B. Schnabel (1997). “Algorithm 768: TENSOLVE: A Software Package for Solving Systems of Nonlinear Equations and Nonlinear Least-Squares Problems Using Tensor Methods,” *ACM Transactions of Mathematical Software*, 23, 174–195.
- R. Bro (1997). “PARAFAC: Tutorial and applications,” *Chemometrics and Intelligent Laboratory Systems*, 38, 149–171.
- R. Bro (2006). “Review on Multiway Analysis in Chemistry 2000–2005,” *Critical Reviews in Analytical Chemistry*, 36, 279–293.
- G. K-L Chan (2004). “An Algorithm for Large Scale Density Matrix Renormalization Group Calculations,” *J. Chem. Phys.*, 120 (7), 3172.
- G.K.-L. Chan, J. Dorando, D. Ghosh, J. Hachmann, E. Neuscamman, H. Wang, and T. Yanai (2007). “An Introduction to the Density Matrix Renormalization Group Ansatz in Quantum Chemistry,” arXiv:cond-mat, vol. 0711.1398v1.

- P.A. Chew, B.W. Bader, T.G. Kolda, and A. Abdelali (2007). “Cross-language Information Retrieval Using PARAFAC2,” *KDD 2007: Proceedings of the 13th ACM SIGKDD International Conference on Data Mining*.
- T.-T. Chow and E. Eskow and R.B. Schnabel (1994). “A Software Package for Unconstrained Optimization Using Tensor Methods,” *ACM Transactions on Mathematical Software*, 20, 518–530.
- A. Cichocki, R. Zdunek, S. Choi, R. Plemmons, and S. Amari (2007). “Nonnegative Tensor Factorization Using Alpha and Beta Divergencies,” in *Proc. of the 32nd International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Honolulu.
- P. Comon (2001). “Tensor Decompositions: State of the Art and Applications,” in *Mathematics in Signal Processing V*, J. G. McWhirter and I. K. Proudler, eds., Oxford University Press, 1–24.
- P. Comon (2004). “Blind Identification and Source Separation in 2-x-3 Underdetermined Mixtures,” *IEEE Transactions on Signal Processing*, 52, 1–13.
- P. Comon, G.H. Golub, L.-H. Lim, and B. Mourrain (2008). “Symmetric Tensors and Symmetric Tensor Rank,” *SIAM Journal on Matrix Analysis Appl.*, to appear.
- M. De Vos, A. Vergult, L. De Lathauwer, W. De Clercq, S. Van Huffel, P. Dupont, A. Palmi, W. Van Paesschen (2007). “Canonical Decomposition of Ictal EEG Reliably Detects the Seizure Onset Zone”, *Neuroimage*, 37, No. 3, 844–854.
- L. De Lathauwer and B. De Moor (1998). “From Matrix to Tensor: Multilinear Algebra and Signal Processing,” in *Mathematics in Signal Processing IV*, J. McWhirter and I. Proudler, eds., Clarendon Press, Oxford, 1–15.
- L. De Lathauwer (2006). “A Link between the Canonical Decomposition in Multilinear Algebra and Simultaneous Matrix Diagonalization”, *SIAM J. Matrix Anal. Appl.*, 28, 642–666.
- L. De Lathauwer, B. De Moor, J. Vandewalle (2000a). “A Multilinear Singular Value Decomposition”, *SIAM J. Matrix Anal. Appl.*, 21, 1253–1278.
- L. De Lathauwer, B. De Moor, and J. Vandewalle (2000b). “On the Best Rank-1 and Rank-($r_1, r_2, \dots, r_N$) Approximation of Higher-Order Tensors,” *SIAM J. Matrix Analysis and Applications*, 21, 1324–1342.
- L. De Lathauwer and J. Vandewalle (2004). “Dimensionality Reduction in Higher-Order Signal Processing and Rank-($R_1; R_2; \dots; R_N$) Reduction in Multilinear Algebra,” *Linear Algebra and its Applications*, 391, 31–55.
- L. De Lathauwer, B. De Moor, and J. Vandewalle (2004). “Computation of the Canonical Decomposition by Means of a Simultaneous Generalized Schur Decomposition”, *SIAM J. Matrix Anal. Appl.*, 26, 295–327.
- W. de la Vega, R. Kannan, M. Karpinski, and S. Vempala (2005). “Tensor Decomposition and Approximation Algorithms for Constraint Satisfaction Problems,” in *Proceedings of the 39th Symposium on Theory of Computing (STOC)*, 747–754.
- J. Demmel, J. Dongarra, V. Eijkhout, E. Fuentes, A. Petit, R. Vuduc, R.C. Whaley, and K. Yelick (2005). “Self-Adapting Linear Algebra Algorithms and Software,” *Proc. IEEE*, 93, 293–312.
- V. de Silva and L.-H. Lim (2008). “Tensor rank and the Ill-Posedness of the Best Low-Rank Approximation Problem,” *SIAM J. Matrix Anal. Appl.*, 30, 1084–1127.
- C. Ding, H. Huang, and D. Luo (2008). “Tensor Reduction Error Analysis – Applications to Video Compression and Classification,” *IEEE Conference on Computer Vision and Pattern Recognition*.
- J. Dongarra, D. Gannon, G. Fox, and K. Kennedy (2007). “The Impact of Multicore on Computational Science Software,” *CTWatch Quarterly*, vol. 3.
- P. Drineas and M. Mahoney (2007). “A Randomized Algorithm for a Tensor-Based Generalization of the Singular Value Decomposition,” *Linear Algebra and Its Applications*, 420, 553–571.
- P. Drineas, R. Kannan, and M.W. Mahoney (2006). “Fast Monte Carlo Algorithms for Matrices III: Computing an Efficient Approximate Decomposition of a Matrix,” *SIAM J. Computing*, 36, 184–206.
- P. Drineas, G. Golub, L.-H. Lim, and M. Mahoney (2006). “Bridging the Gap Between Numerical Linear Algebra, Theoretical Computer Science, and Data Applications” *SIAM News*, 39, no. 8.

- D.M. Dunlavy, T.G. Kolda and W.P. Kegelmeyer (2006). “Multilinear Algebra for Analyzing Data with Multiple Linkages,” Technical Report Number SAND2006-2079, Sandia National Laboratories, Albuquerque, NM and Livermore, CA.
- L. Eldén (2007). *Matrix Methods in Data Mining and Pattern Recognition*, SIAM Publications, Philadelphia, PA.
- L. Eldén and B. Savas (2009). “A Newton-Grassmann Method for Computing the Best Multi-linear Rank-(r_1 ; r_2 ; r_3) Approximation of a Tensor,” *SIAM J. Matrix Analysis and Applications*, 31, 248–271.
- E. Elmroth, F. Gustavson, I. Jonsson, and B. Kågström (2005). “Recursive Blocked Algorithms and Hybrid Data Structures for Dense matrix Library Software,” *SIAM Review*, 46, 3–45.
- D. Feng and T.H. Pulliam (1997). “Tensor-GMRES Method for Large Systems of Nonlinear Equations,” *SIAM J. Optim.*, 7, 757–779.
- S. Friedland (2008). “On the Generic Rank of 3-Tensors,” arXiv 0805.3777v2.
- M. Frigo and S.G. Johnson (2005). “The Design and Implementation of FFTW3,” *Proc IEEE*, 93, 216–231.
- A. Frommer, M.W. Mahoney, and D.B. Szyld (2007). *Web Information Retrieval and Linear Algebra Algorithms*, Proc. of Dagstuhl Seminar 07071.
- G.H. Golub and C. Van Loan (1996). *Matrix Computations, 3rd Edition*, Johns Hopkins University Press, Baltimore, MD.
- J. Granata, M. Conner, and R. Tolimieri (1992). “The Tensor Product: A Mathematical Programming Language for FFTs and Other Fast DSP Operations,” *IEEE SP Magazine*, January, 40–48.
- J. Granata, M. Conner, and R. Tolimieri (1992). “Recursive Fast Algorithms and the Role of the Tensor Product,” *IEEE Transactions on Signal Processing*, 40, 2921–2930.
- W. Greub (1978). *Multilinear Algebra, 2nd ed*, Springer-Verlag, New York.
- V.S. Grigorascu and P.A. Regalia (1999). “Tensor Displacement Structures and Polyspectral Matching,” in *Fast Reliable Algorithms for Matrices with Structure*, T. Kailath and A.H. Sayed (eds), SIAM Publications, Philadelphia.
- J. Grotendorst (2000). *Modern Methods and Algorithms of Quantum Chemistry*, John von Neumann Institute for Computing, Julich, NIC Series, Vol. 1, ISBN 3-00-005618-1.
- J. Hachmann, W. Cardoenand, G. Chan (2006). “Multireference Correlation in Long Molecules with the Quadratic Scaling Density Matrix Renormalization Group”, *J. Chem. Phys.*, 125.
- W. Hackbusch and B.N. Khoromskij (2007). “Tensor-product Approximation to Operators and Functions in High Dimensions,” *Journal of Complexity*, 23, 697–714.
- W. Hackbusch, B.N. Khoromskij, and E.E. Tyrtysnikov (2005). “Hierarchical Kronecker tensor-product approximations,” *Journal of Numerical Mathematics*, 13, 119–156.
- A. Hartono, A. Sibiryakov, M. Nooijen, G. Baumgartner, D.E. Bernholdt, S. Hirata, C. Lam, R. Pitzer, J. Ramanujam, and P. Sadayappan (2005). “Automated Operation Minimization of Tensor Contraction Expressions in Electronic Structure Calculations,” in *Proc. International Conference on Computational Science 2005*, Atlanta.
- M. Head-Gordon, Y. Shao, C. Saravanan, and C. White (2003). “Curvy Steps for Density Matrix Based Energy Minimization: Tensor Formulation and Toy Applications,” *Molecular Physics*, 101, 37–43.
- M. Helal (2001). *Mathematics of Arrays - The Implementation and the Application*, Masters Thesis, American University in Cairo.
- M. Helal, H. El-Gindy, L.M. Mullin, and B. Gaeta (2008). “Parallelizing Optimal Multiple Sequence Alignment by Dynamic Programming,” *Proceedings of the International Symposium on Advances in Parallel and Distributed Computing Techniques (APDCT-08)*, 669–674.

- M. Helal, L.M. Mullin, B. Gaeta, and H. El-Gindy (2007). “Multiple Sequence Alignment Using Massively Parallel Mathematics of Arrays,” *Proceedings of the International Conference on High Performance Computing, Networking and Communication Systems (HPCNCS- 07)*, Orlando, FL., 120–127.
- M. Helal and V. Sintchenko (2009). “Dynamic Programming Algorithms for Discovery of Antibiotic Resistance in Microbial Genomes,” *Proceedings of Health Informatics Conference (HIC-09)*, to be held in 19-21 August, 2009, Canberra, Australia.
- S. Hirata (2003). “Tensor Contraction Engine: Abstraction and Automatic Parallel Implementation of Configuration Interaction, Coupled-Cluster, and Many-Body Perturbation Theories,” *J. Phys. Chem. A.*, 107, 9887–9897.
- H. Huang, C. Ding, D. Luo, and T. Li (2008). “Simultaneous Tensor Subspace Selection and Clustering: The Equivalence of High Order SVD and K-Means Clustering,” *Proc. ACM Int’l Conf. on Knowledge Discovery and Data Mining (KDD)*, 322–335.
- C-H Huang, J.R. Johnson, and R.W. Johnson (1991). “Multilinear Algebra and Parallel Programming,” *J. Supercomputing*, 5, 189–217.
- H.B. Hunt, L.R. Mullin, D.J. Rosenkrantz, and J.E. Raynolds (2008). “A Transformation–Based Approach for the Design of Parallel/Distributed Scientific Software: the FFT,” arXiv:0811.2535.
- M. Ishteva, L. De Lathauwer, P.-A. Absil, S. Van Huffel (2009). “Differential-Geometric Newton Algorithm for the Best Rank- (R_1, R_2, R_3) Approximation of Tensors”, *Numerical Algorithms*, to appear.
- M.A. Jenkins and L.R. Mullin (1990). “A Comparison of Array Theory and a Mathematics of Arrays”, in *Arrays, Functional Languages and parallelism, Proceedings of the Montreal Workshop*, Kluwer Academic Publishers.
- B. Kågström, P. Ling, and C. Van Loan (1998). “GEMM-Based Level 3 BLAS: Algorithms for the Model Implementations,” *ACM Transactions on Mathematical Software*, 24, 268–302.
- M.E. Kilmer and C.D. Moravitz Martin (2004). “Decomposing a Tensor,” *SIAM News*, 37, Number 9.
- M.E. Kilmer, C.D. Martin, and L. Perrone (2008). “A Third-Order Generalization of the Matrix SVD as a Product of Third-Order Tensors,” Technical Report Number TR-2008-4, Department of Computer Science, Tufts University, Medford, MA.
- H. Kim, H. Park, and L. Eldén (2007). “Non-negative Tensor Factorization Based on Alternating Large-scale Non-negativity-constrained Least Squares,” *Proceedings of the 7th IEEE International Conference on Bioinformatics & Bioengineering (IEEE BIBE 2007)*, Vol.II, 1147-1151.
- D. Kim, S. Sra, and I. S. Dhillon (2008). “Fast Projection-Based Methods for the Least Squares Nonnegative Matrix Approximation Problem,” *Statistical Analysis and Data Mining*, vol. 1:1, 38–51
- E. Kofidis and P. Regalia (2001). “Tensor Approximation and Signal Processing Applications,” in *Structured Matrices in Mathematics, Computer Science, and Engineering I*, V. Olshevsky (ed), American Mathematical Society, Providence, RI, 103–133.
- E. Kofidis and P.A. Regalia (2002). “On the Best Rank-1 Approximation of Higher-Order Super-Symmetric Tensors,” *SIAM J. Matrix Analysis and Applications*, 23, 863-884.
- T.G. Kolda (2001). “Orthogonal Tensor Decompositions,” *SIAM Journal on Matrix Analysis and Applications*, 23, 243–255.
- T.G. Kolda (2003). “A Counterexample to the Possibility of an Extension of the Eckart-Young Low-Rank approximation Theorem for the Orthogonal Rank Tensor Decomposition,” *SIAM Journal on Matrix Analysis and Applications*, 24, 762–767.
- T.G. Kolda and B.W. Bader (2009). “Tensor Decompositions and Applications,” *SIAM Review* 51, (to appear).
- T.G. Kolda, B.W. Bader, and J. P. Kenny (2005). “Higher-Order Web Link Analysis Using Multi- Linear Algebra,” in *ICDM 2005: Proceedings of the 5th IEEE International Conference on Data Mining, IEEE Computer Society*, 242–249.

- T.G. Kolda and J. Sun (2008). “Scalable Tensor Decompositions for Multi-aspect Data Mining,” *ICDM 2008: Proceedings of the 8th IEEE International Conference on Data Mining*, 363–372.
- P.M. Kroonenberg (2008). *Applied Multiway Data Analysis*, Wiley.
- C-C. Lam, P. Sadayappan, R. Wenger (1997). “Optimal Reordering and Mapping of a Class of Nested-Loops for Parallel Execution,” in *Languages and Compilers for Parallel Computing*, (D. Sehret et al Eds.), Lecture Notes in Computer Science, Vol. 1239, 315–329, Springer-Verlag.
- W. Landry (2003). “Implementing a High Performance Tensor Library,” *Scientific Programming*, 11, 273–290.
- J.M. Landsberg (2005). “The Border Rank of the Multiplication of 2-by-2 Matrices is Seven,” *Journal of the American Mathematical Society*, 19, 447–459.
- J.M. Landsberg and J. Morton (2009). *The Geometry of Tensors: Applications to Complexity, Statistics and Engineering*, to appear.
- J. Lawrence (2006). “Inverse Shade Trees for Non-parametric Material Representation and Editing,” *ACM Trans. Graph.*, 25, 735–745.
- C. Lechner, D. Alic, and S. Husa (2004). “From Tensor Equations to Numerical Code,” *Computer Algebra Tools for Numerical Relativity*, vol. 0411063, 2004.
- J. Leskovec and C. Faloutsos (2007). “Scalable Modeling of Real Graphs Using Kronecker Multiplication,” *International Conference on Machine Learning (ICML)*.
- J. Leskovec and E. Horvitz (2008). “Worldwide Buzz: Planetary-Scale Views on an Instant-Messaging Network,” *WWW2008*.
- L.-H. Lim (2005). “Singular Values and Eigenvalues of Tensors: a Variational Approach,” *Proc. IEEE Int. Workshop on Computational Advances in Multi-Sensor Adaptive Process*, 129–132.
- L.-H. Lim and P. Comon, “Nonnegative Approximations of Nonnegative Tensors,” *Journal of Chemometrics*, (to appear).
- N. Liu, B. Zhang, J. Yan, Z. Chen, W. Liu, F. Bai, and L. Chien (2005). “Text representation: From Vector to Tensor,” in *ICDM 2005: Proceedings of the 5th IEEE International Conference on Data Mining*, IEEE Computer Society, 725–728.
- Q. Lu, X. Gao, S. Krishnamoorthy, G. Baumgartner, J. Ramanujam, and P. Sadayappan (2005). “Empirical Performance-Model Driven Data Layout Optimization,” *Languages and Compilers for Parallel Computing*, (R. Eigenmann et al Eds.), Lecture Notes in Computer Science, Springer-Verlag.
- M.W. Mahoney, M. Maggioni, and P. Drineas (2008). “Tensor-CUR Decompositions For Tensor-Based Data”, *SIAM J. Matrix Analysis and Applications*, 30, 957–987.
- M.W. Mahoney, L-H Lim, and G. Carlsson (2009). “Algorithmic and Statistical Challenges in Modern Large-Scale Data Analysis are the Focus in MMDS 2008,” *SIGKDD Explorations*, 10(2), 57–60.
- C.D. Martin and C. Van Loan (2008). “A Jacobi-Type Method for Computing Orthogonal Tensor Decompositions,” *SIAM J. Matrix Analysis and Applications*, 29, 184–198.
- J. Morton and L.-H. Lim (2009). “Principal Cumulant Component Analysis,” Technical report, Stanford University.
- M. Moslashrup, L. Hansen, S. Arnfred, L.-H. Lim, and K. Madsen (2008). “Shift Invariant Multilinear Decomposition of Neuroimaging Data, *NeuroImage*, no. 4, 1439–1450.
- J. Nocera (2009). “Risk Mismanagement,” *The New York Time Magazine*, January 4, 2009.
- L. Omberg, G. Golub, and O. Alter (2007). “A Tensor Higher-Order Singular Value Decomposition for Integrative Analysis of Dna Microarray Data from Different Studies,” *Proc. National Academy of Science*, 107, 18371–18376.

- P. Paatero (1999). “The Multilinear Engine: A Table-Driven, Least Squares Program for Solving Multilinear Problems, Including the n-way Parallel Factor Analysis Model,” *Journal of Computational and Graphical Statistics*, 8, 854–888.
- H. Park and J. Kim (2008). “Toward Faster Nonnegative Matrix Factorization: A New Algorithm and Comparisons,” *The Proceedings of the IEEE International Conference on Data Mining*, to appear.
- L. Qi (2005). “Eigenvalues of a Real Supersymmetric Tensor,” *Journal of Symbolic Computation*, 40, 1302–1324.
- D.J. Rosenkrantz, L.R. Mullin, and H.B. Hunt III (2006). “On Minimizing Materializations of Array-Valued Temporaries”, *ACM Trans. Program. Lang. Syst.*, 28, 1145–1177
- B. Savas (2008). *Algorithms in Data Mining using Matrix and Tensor Methods*, PhD. Thesis, Linköping University, Sweden.
- B. Savas and L. Eldén (2007). “Handwritten Digit Classification Using Higher Order Singular Value Decomposition,” *Pattern Recognition*, 40, 993–1003.
- B. Savas and L. Eldén (2009). “Krylov Subspace Methods for Tensor Computations,” Report LITH-MAT-R-2009-02-SE, Linköping University, Sweden.
- B. Savas and L.-H. Lim (2008). “Best Multilinear Rank Approximation of Tensors with Quasi-Newton Methods on Grassmannians,” Department of Mathematics, Linköping University, Report LiTH-MAT-R-2008-01-SE.
- R.B. Schnabel and P.D. Frank (1984). “Tensor Methods for Nonlinear Equations,” *SIAM J. Numer. Anal.*, 21, 815–843.
- R.B. Schnabel and T.-T. Chow (1991). “Tensor Methods for Unconstrained Optimization Using Second Derivatives,” *SIAM J. Optim.*, 1, 293–315.
- A. Smilde, R. Bro, and P. Geladi (2004). *Multi-Way Analysis: Applications in the Chemical Sciences*, Wiley, West Sussex, England.
- J. Sun (2007). *Incremental Pattern Discovery on Streams, Graphs and Tensors*, PhD thesis, Carnegie Mellon University.
- J. Sun, C.E. Tsourakakis, E. Hoke, C. Faloutsos and T. Eliassi-Rad (2008). “Two Heads Better than One: Pattern Discovery in Time-Evolving Multi-Aspect Data,” *Journal Data Mining and Knowledge Discovery*, 17, 1, 111-128.
- G. Tomasi and R. Bro (2006). “A Comparison of Algorithms for Fitting the PARAFAC Model,” *Computational Statistics and Data Analysis*, 50, 1700–1734
- L.R. Tucker (1966). “Some Mathematical Notes on Three-Mode Factor Analysis,” *Psychometrika*, 31, 279–311.
- C. Van Loan (1992). *Computational Frameworks for the Fast Fourier Transform*, SIAM publications, Philadelphia, PA.
- C. Van Loan (2000). The Ubiquitous Kronecker Product, *Journal of Computational and Applied Mathematics*, 123, 85–100.
- M.A.O. Vasilescu and D. Terzopoulos (2002). “Multilinear Analysis of Image Ensembles: TensorFaces,” *Proceedings of the 7th European Conference on Computer Vision*, Vol. 2350, 447–460.
- S.R. White and R.L. Martin (1999). “Ab Initio Quantum Chemistry Using the Density Matrix Renormalization Group,” *J. Chem. Phys.*, 110, no. 9, p. 4127.
- Q. Zhang, M.W. Berry, B. Lamb and T. Samuel (2009). “A Parallel Nonnegative Tensor Factorization Algorithm for Mining Global Climate Data,” *Proceeding of LNCS*.
- Q. Zhang, H. Wang, R. Plemmons and P. Pauca (2008). “Tensor Methods for Hyperspectral Data Analysis: A Space Object Material Identification Study,” *J. Optical Soc. Amer. A*, 25, 3001–3012.
- T. Zhang and G. H. Golub (2001). “Rank-one Approximation to High order Tensors,” *SIAM Journal on Matrix Analysis and Applications*, 23, 534–550.