

IEEE Vis 2013 Tutorial: Tensor Approximation in Scientific Visualization

Susanne K. Suter, *Student Member, IEEE*, and Rafael Ballester-Ripoll, and Renato Pajarola, *Member, IEEE*

Abstract— In this course, we will introduce the basic concepts of tensor approximation (TA) – a higher-order generalization of the SVD and PCA methods – as well as its applications to visual data representation, analysis and visualization, and bring the TA framework closer to visualization researchers and practitioners. The course will cover the theoretical background of TA methods, their properties and how to compute tensor decompositions, as well as practical applications of TA methods in visualization and visual computing. In a first theoretical part, the attendees will be instructed on the necessary mathematical background of tensor decomposition and approximation methods to learn the basics skills of using and applying these new tools in the context of the representation of large multidimensional visual data. Specific and very noteworthy features of the TA framework are highlighted which can effectively be exploited for spatio-temporal multidimensional data representation and visualization purposes. In a practical implementation and an application oriented session, compact TA data representation in scientific visualization and visual computing as well as decomposition and reconstruction algorithms will be demonstrated. At the end of the course, the participants will have a good basic knowledge of TA methods along with a practical understanding of implementation issues and its potential application in visualization.

Index Terms—Tensor decompositions, tensor approximations, Tucker model, CANDECOMP/PARAFAC model, hybrid block-diagonal TA models, compact visual data representation, higher-order SVD methods, data reduction, interactive volume visualization, multiresolution and multiscale modeling, clustered tensor decomposition.

1 ORGANIZATION

Organizers	Susanne K. Suter, Rafael Ballester-Ripoll, Prof. Dr. Renato Pajarola, University of Zürich, Switzerland
Lecturers	Susanne K. Suter Rafael Ballester-Ripoll Renato Pajarola
Duration	Half-day
Level	Intermediate
History	Tutorial parts were held at Eurographics 2013
Supplemental material	Theory document, example slides (selected parts), slides on vmmlib TA classes, vmmlib TA classes play project

As outlined in the abstract and further elaborated on below, this tutorial will bring the concepts of tensor approximation (TA) closer to the visualization communities. TA methods have already shown to be quite useful in visualization and graphics in a number of specific papers. This tutorial will review both, the underlying theory as well as some of the recent applications in this context.

We strongly believe that the TA framework is a powerful toolbox that has a large potential for a strong and lasting impact on large data representation, analysis and visualization solutions. This topic, tensor approximations and its applications in visual computing, has not received a significant and broad treatment in the past. A first TA focused tutorial has been given at Eurographics 2013. This tutorial was very successful, attracted a good audience, and feedback indicated increasing interest in the visual computing community. Furthermore, many people newly exposed to this topic, e.g. as during the VisWeek'12 poster session [SP12], have shown a strong interest in the general theory and possible applications of tensor approximation.

As applicable, we aim to provide practical insight into the implementation and usage of TA tools, e.g., using MATLAB or C++ code

examples. To that extent we will provide sample tutorial code for participants to exploit.

2 DESCRIPTION

2.1 Overview

The SVD and PCA approaches, which work effectively for matrix-based data compression (2nd-order tensors) cannot directly be extended in a straight-forward way to higher-dimensional data (higher-order tensors) and lose some of their unique properties. Nevertheless, a number of common visual datasets naturally lend themselves to a representation as higher-order tensors: volume data (3rd-order), spatio-temporal volume and FMRI data (4th-order), image stacks and video (3rd-order), BRDF/BTF illumination sample data (mostly 3rd-order or 4th order), as well as general image and sample collections (k -th order). In this tutorial, we briefly introduce the tensor approximation (TA) framework as an extension of the SVD and PCA approaches to higher-order tensor dimensionality, and describe TA with its special properties as a numerical linear algebra toolbox to process, analyze and represent complex visual data in novel ways.

The targeted audience consists of visualization researchers and practitioners since the presented methods and techniques exhibit direct applicability to represent and manage large multidimensional visual datasets. Previous and recent research work has already demonstrated the high potential of TA methods, and corresponding examples in visualization and visual computing in general will be reviewed as part of this tutorial.

At the end of the course, participants should understand the main concepts, properties and features of the TA framework and in particular the key differences between a Tucker model and a CANDECOMP/PARAFAC (CP) tensor model. Furthermore, he or she should be able to apply a TA and its reconstruction to simple data models, and be in a strong position to thoroughly understand the specific and advanced approaches presented from the recent research literature. In particular, the course will be augmented with practical examples.

Note that we propose a concise half-day course in order to keep the tutorial compact; however, if desired, we could see an even more detailed extension of the tutorial to a full day.

2.2 Target Audience

The targeted audience includes data visualization, visual computing and computer graphics experts, researchers as well as practitioners

• S.K. Suter, R. Ballester-Ripoll, and R. Pajarola are with the Visualization and MultiMedia Lab, University of Zürich, Switzerland
Email: susuter@ifi.uzh.ch, rballester@ifi.uzh.ch, and pajarola@acm.org.

with a solid background in linear numerical algebra who have to store, access and visualize large and complex multidimensional visual datasets. The presented TA methods have a high degree of direct applicability to the analysis and visualization of multidimensional visual data and to the compact representation and storage of large visual data in visualization, as well as to other closely related fields like multimedia and image or video processing.

2.3 Previous Courses and Tutorials

To the best of our knowledge, there has not been a similar course or tutorial on data tensors and applications of tensor approximations in visualization and visual computing at any visualization oriented conferences in the past. One related course on tensor methods in visualization that we could identify was the *Tensors in Visualization* course at the VisWeek 2010 [KST*10]. However, that course largely focused on tensor fields and their visualization, and only included data tensors briefly at the end of the tutorial. The proposed tutorial exclusively focuses on data tensors (general multi-way arrays) and not tensor fields. The main focus of this tutorial lies on tensor decomposition and tensor reconstruction of data tensors and the application of tensor approximation methods in scientific data visualization. Thus, the outlined tutorial is fundamentally different from the course at VisWeek 2010.

A first tutorial on *Tensor Approximation in Visualization and Graphics* in the context of graphics, visual computing and visualization was held at Eurographics 2013 [PSR13] (<http://vmml.ifi.uzh.ch/links/TutorTensorAprox.html>). This tutorial is closely related and shares parts of the structure and content. The main differences of this proposed tutorial compared to the one from Eurographics 2013 are its increased focus on scientific visualization, volume data representation, comparative analysis and practical details. In particular, the following major changes and differences will be integrated:

- The theory and properties of TA methods will be revised, integrating novel block-based tensor models and further examples exploiting certain TA properties.
- A new analysis will include substantial performance analysis of the various TA models as well as comparative studies of wavelet transform and TA based methods.
- In the applications we will focus more on scientific visualization, and specifically discuss the integration of concepts such as out-of-core data access, and hierarchical or clustered TA models.
- Extended implementation details we will include practical interactive demonstrations of the TA methods and their actual realization within C++ and/or MATLAB programs.

2.4 Syllabus

This tutorial addresses the application of advanced numerical linear algebra tools to compact data representations and interactive visualization of large multidimensional datasets. These datasets arise in many applications in scientific visualization and computer graphics, such as visualization of volumetric data or spatio-temporal simulation data, storage of reflectance data, motion synthesis or precomputed radiance transfer. TA methods have recently attracted increasing interest from the visual computing community, and a number of authors have shown that the TA framework is a viable tool for the compact representation of these multidimensional dataset. The idea of this tutorial is to give an introduction of TA methods and how they can be applied in visualization and graphics. Notably, we aim at making the successful TA application strategies available to the scientific audience.

The overall tutorial is structured into four main parts consisting of the general theoretical background and properties of TA methods, followed by practical applications of TA methods in scientific visualization and visual computing:

Part 1 Introduction of the TA framework

- Tucker and CP tensor decompositions

- Block-diagonal TA models
- Rank-reduced tensor approximations
- Useful TA properties and features for data visualization

Part 2 Comparative analysis

- Frequency analysis and DCT equivalence
- Performance analysis using MATLAB and vmmlib (C++)
- Multiscale feature expressiveness (wavelets vs. TA)

Part 3 TA applications in scientific visualization

- Out-of-core, level-of-detail and clustering concepts for large data tensor handling
- Multiscale and multiresolution hierarchy modeling for TA-based volume visualization

Part 4 Implementation and practice

- Live demo of multiscale and multiresolution TA
- Code examples with MATLAB and vmmlib (C++)
- GPU-based tensor reconstruction solutions

The sessions are designed in more detail as follows:

Opening (5min)

Presentation of the structure of this tutorial course and schedule of topics, introduction of speakers.

Introduction to the TA Framework (45min)

In this first session, we introduce the basic definition of a tensor approximation model, which is the decomposition of a higher-order tensor into a multilinear combination of bases and weighting coefficients. We introduce a wide range of TA models. The generalization of the matrix SVD is defined with two main models: the CANDECOMP/PARAFAC (CP) model and the Tucker tensor model. However, we also elaborate on various hybrid variants of the CP and Tucker approaches (so-called block-diagonal TA models, see e.g. [BRSEP13]), which we will present in this tutorial. Moreover, we explain the definitions of multilinear ranks as used for TA models and we elaborate on how a truncated tensor decomposition defines an approximation of the original.

In the context of data approximation, the tensor models are reviewed and analyzed as being a data point in a high-dimensional approximation space. Consequently, some specific properties and features of these approximation spaces, e.g., such as uniqueness, factor matrix orthonormality, all-orthogonality of core tensors or space-rank selectivity, are discussed as well as their effects on truncated tensor reconstructions (see e.g. [SMP13]). Finally, we show direct or incremental approaches for tensor decomposition algorithms.

Comparative Analysis (30min)

In this comparative part, we aim to highlight similarities of the TA framework to frequency domain analysis and compare the implementation performance of different TA toolboxes, as evaluated in [BRSEP13]. With respect to TA bases, we show the equivalence to frequency domain transformations for example by using discrete cosine transform (DCT) vectors as TA bases. Moreover, we will show performance results between the TA implementations of the MATLAB tensor toolbox and vmmlib (C++), which are described in more detail in Section 4.

Additionally, we compare the feature expressiveness and compact data representation capabilities of wavelet and TA based approaches [WXC*08, SZP10a].

TA Applications in Scientific Visualization (30min)

The goal of this part is to show various TA applications in the domain of scientific visualization. We start by reviewing compact data representation approaches (TA and wavelets [WXC*08, SZP10a]) and showing tensor approximation applications for interactive multiscale feature and multiresolution level-of-detail scientific volume visualization [SMP13].

In scientific visualization and visual computing applications, different tensor-based hierarchies have been proposed and will be compared in this tutorial, such as the progressively refined TA hierarchy of [WXC*08], as well as full multiresolution tensor data representation hierarchies (using local TA bases [SIGM*11] or global TA bases [SMP13]). We will show that most of these hierarchies support good data compression rates and can exploit out-of-core data management or decomposition methods. Moreover, it is shown in this tutorial how the various TA hierarchies exploit data redundancy between subvolumes and how typical brick-artifacts during volume rendering can be reduced effectively. Similarly, clustered TA models exploit redundancy along one specific tensor mode [TS06, TS12].

Implementation and Practice (45min)

In interactive scientific visualization, the tensor decomposition is usually carried out as an offline preprocessing routine which is less time critical, while the reconstruction process has to be performed at interactive rates. In this tutorial, the basic tensor data structures and tensor decomposition algorithms are explained by the examples of two available tensor libraries: the MATLAB Tensor Toolbox and the vmmllib template framework (see also Section 4). For the tensor reconstruction, we show a fast tensor reconstruction implementation on the GPU (see [SIGM*11]), which supports interactive large data visualization.

To demonstrate the applicability of TA methods in interactive visualization we intend to give a live demo of our tensor-based multiresolution and multiscale direct volume renderer [SMP13] (see also <http://github.com/VMML/Equalizer>). Within the tutorial, we show performance timings for the preprocessing step (generation of multiscale multiresolution TA data structure) and the real-time reconstruction for datasets of different sizes, such that a potential user gets a good impression about the offline and run-time performance for different typical dataset sizes. For the tutorial, documented test datasets and example routines are provided.

Closing (5min)

Finally, we will summarize the TA framework and its application in visualization and visual computing and provide a brief outlook on future challenges in the field.

3 DOCUMENTATION

In addition to the full tutorial slide sets, additional documentation will be made available to the attendees in form of summaries of related papers including dedicated links to electronic online versions (see Appendix B).

The presented tutorial is based on a number of articles and papers on tensor approximation methods and their applications. The basic theory of tensor decomposition and approximation methods are described in [dLdMV00a, dLdMV00b] and [KB09], of which we follow the latter on notation and formalism. Other main resources are the PhD theses of Tsai [Tsa09] and Suter [Sut13]. The attached theory part summarizes the used TA notation, the main tensor decomposition models and some decomposition properties, which can be exploited for scientific visualization. This theory part will be extended with block-diagonal tensor models as described in [BRSEP13]. Moreover, we review in this tutorial a number of key applications of tensor methods in visualization and visual computing [TS06, WWS*05, WXC*08, SZP10a, SIGM*11, TS12, SMP13, BRSEP13].

The slides from the previous Eurographics 2013 tutorial (see <http://vmmllib.ifl.uzh.ch/links/TutorTensorApprox.html>) will be extended with respect to the changes outlined in

Section 2.3 for block-diagonal TA models, more details on recent hierarchical TA models [TS12, SMP13], and including new comparative analysis (C++ vs. MATLAB, wavelets vs. TA).

Based on our extensive practical experience and work with tensor decomposition methods, we will provide a number of basic MATLAB examples (see https://files.ifl.uzh.ch/vmmllib/ta_tutorial/vmmllib_ta_classes.pdf and project example on <http://vmmllib.github.io/vmmllib/>). These results are shown based on our own development [vmm] (<https://github.com/VMML/vmmllib>) and the MATLAB tensor toolbox [BK06] (<http://www.sandia.gov/~tgkolda/TensorToolbox>). The practical examples, including test datasets will be made available to the attendees. Our test volume datasets are already available online (see <http://vmmllib.ifl.uzh.ch/research/datasets.html>) and will be provided for the tutorial as well.

4 AVAILABLE SOFTWARE

We consider the four presented toolboxes as the most convenient ones for tensor approximation applications; however, there is more tensor software available, as summarized in [KB09]. In the actual tutorial, we mostly refer to vmmllib (C++) and the MATLAB tensor toolbox.

4.1 vmmllib

The tensor classes that extend the *vector and matrix math library* vmmllib [vmm] are a C++ implementation with templates, and support tensors up to the 4th order. They provide methods for basic tensor manipulation (including foldings, unfoldings and product with vectors or matrices), as well as more advanced algorithms for decomposing and reconstructing (Tucker model). In particular, the 3rd-order version can also deal with the CP model, as well as with a number of hybrid block-based TA decompositions, as detailed in [BRSEP13]. Vmmllib also supports tensor memory-mapping in order to process large input tensors. So far, only compact tensor decompositions were considered, sparse implementations will need to be added on top. Desired extensions can be integrated by any developer since it is an open-source project.

<https://github.com/VMML/vmmllib>

4.2 MATLAB Tensor Toolbox

The MATLAB Tensor Toolbox [BK*12] is a comprehensive toolbox for tensor approximation applications. It offers optimized decomposition algorithms for the Tucker model as well as the CP model. The MATLAB tensor toolbox is a generic implementation for any N^{th} -order tensor decomposition and includes a well-documented help manual. The toolbox supports compact and sparse tensors, tensor unfoldings into matrices, and tensor multiplications. The main algorithms are published in [BK06]. This toolbox is a comprehensive tensor environment that is easy to use and extend in MATLAB.

<http://www.sandia.gov/~tgkolda/TensorToolbox>

4.3 MATLAB N-Way Toolbox

The N-Way Toolbox [AB00] is a MATLAB toolbox which provides functions for the computation of Tucker and CP approximations of a tensor. The implementations of these algorithms in the N-Way toolbox are very flexible and provide the user with a large number of options. Several initialization methods can be used, it is possible to specify orthogonality and non-negativity constraints for each of the modes individually and the imputation of missing values is supported. Furthermore, the computation of weighted CP approximations is possible.

<http://www.models.life.ku.dk/nwaytoolbox>

4.4 Tensorlab – MATLAB Toolbox

Just recently, a new MATLAB toolbox – Tensorlab [SVBDL13] – has been released. The toolbox offers algorithms for tensor decompositions, complex optimization (e.g., quasi-Newton), global minimization of bivariate polynomials and rational functions, and other features such as tensor visualization, estimating a tensor's rank or multilinear rank. The tensor decompositions include a large variety of algorithms such

as canonical polyadic decomposition (CPD), CPD with structured factor matrices and (partial) symmetry, multilinear singular value decomposition (MLSVD), block term decompositions (BTD) and low multilinear rank approximation (LMLRA).

<http://www.esat.kuleuven.be/sista/tensorlab/>

5 LECTURERS EXPERTISE

The tutorial is given by three experts on TA methods in visualization and computer graphics (two young and an experienced researcher). In the following, the lecturers' backgrounds and specializations are summarized.

Susanne K. Suter

Research Assistant, Post-Doc
Visualization and MultiMedia Lab (VMML)
University of Zürich, Switzerland
susuter@ifi.uzh.ch
<http://vmml.ifi.uzh.ch/people/current-staff/suter.html>

Susanne Suter is a post-doctoral research assistant at the University of Zürich, Switzerland. Her scientific expertise lies in data reduction and data compression, feature extraction, automation, real-time interactive visualization, and linear algebra in visualization.

Susanne K. Suter recently graduated from her PhD, where the core topic of her thesis [Sut13] matches the presented VisWeek tutorial. Her main focus in the area is interactive visualization of tensor approximated data from large micro-computed tomography or phase-contrast synchrotron datasets, where the main challenge lies in finding a mathematical framework to perform all tasks with one tool. That is (a) to reduce the actual amount of data, (b) to extract relevant features, and (c) to visualize from the decomposed data in real-time.

Susanne Suter contributed to the field by showing that TA is practical for multiscale volume visualization [SZP10b, SZP10a] and confirming that the online hardware-accelerated reconstruction for interactive rendering is fast enough [SIGM*11]. Moreover, she showed that multiscale feature visualization and state-of-the-art multiresolution DVR can be modeled directly into global tensor factor matrices and visualized with a single feature scale parameter [SMP13]. Susanne Suter explored other tensor decomposition approaches [FMPS13], suitable models for visualization purposes [BRSEP13], and the development of tensor approximation algorithms [vmm]. Furthermore, she recently took an active part on sharing her knowledge with the community [BRGI*13, PSR13].

During her PhD, Susanne Suter was granted a two-year fellowship from the University of Zurich. Just recently she received a one-year Swiss National Science foundation fellowship to pursue her research at the University of Florida, Gainesville, USA. She is a member of ACM, IEEE and IEEE VGTC.

Rafael Ballester-Ripoll

Research Assistant, PhD Candidate
Visualization and MultiMedia Lab (VMML)
University of Zürich, Switzerland
rballester@ifi.uzh.ch
<http://vmml.ifi.uzh.ch/people/current-staff/ballester.html>

Rafael Ballester-Ripoll is a doctoral candidate at the University of Zürich (UZH), Switzerland, since 2012. Previously, he obtained Diplomas in Mathematics and Informatic Engineering from the Technical University of Catalonia-BarcelonaTech, both in 2012. His research interests include volume visualization, multidimensional data processing and tensor-based compression, and real-time interactive visualization. At the UZH, he currently develops and applies tensor-approximation algorithms for volume visualization [BRSEP13, vmm]. His PhD position is funded by the Swiss National Science Foundation.

Professor, Dr. sc.-techn. ETH
Visualization and MultiMedia Lab (VMML)
University of Zürich, Switzerland
pajarola@acm.org
<http://vmml.ifi.uzh.ch/people/current-staff/pajarola.html>

Renato Pajarola

Renato Pajarola received his Dipl. Inf-Ing ETH and Dr. sc. techn. degrees in computer science from the Swiss Federal Institute of Technology (ETH) Zürich in 1994 and 1998 respectively. Subsequently he was a post-doctoral researcher and lecturer in the Graphics, Visualization & Usability (GVU) Center at Georgia Tech. In 1999 he joined the the University of California Irvine (UCI) as an Assistant Professor where he founded the Computer Graphics Lab. Since 2005 he has been leading the Visualization and MultiMedia Lab (VMML) at the University of Zürich (UZH) as Professor in the Department of Informatics. He is a member of ACM, ACM SIGGRAPH, IEEE and Eurographics.

Dr. Pajarola's research interests include real-time 3D graphics, multiresolution modeling, point based graphics, interactive large-scale scientific visualization, remote and parallel rendering, volume visualization and compression. He has published a wide range of internationally peer-reviewed research articles in top journals and conferences. He regularly serves on program committees, such as for example the IEEE Visualization Conference (2004-06,09-11), Eurographics (2010-11, 2013), Pacific Graphics (2002-03,07-08), IEEE Pacific Visualization (2008-10) or EuroVis (2001,2006-10, 2013). He chaired the 2010 EG Symposium on Parallel Graphics and Visualization and was papers co-chair in 2011, as well as papers co-chair of the 2007 and 2008 IEEE/EG Symposium on Point-Based Computer Graphics. He received a Eurographics Best Paper Award in 2005, an IADIS Best Paper Award in 2007 and a SPIE Best Paper Award in 2013.

Dr. Pajarola has previously participated in four quite successful and well received tutorials, at IEEE Visualization and ACM SIGGRAPH Asia on out-of-core, interactive massive model and parallel rendering methods [CESL*03, DGM*08, YMK*09], and at Eurographics on tensor approximation in visualization and graphics [PSR13].

Our intensive research activities on large scale multiresolution data representation, data reduction and interactive visualization, in particular volume rendering, has led us to the field of tensor approximation methods which are the central topic of this tutorial. Experiences from our own research on tensor approximations used in volume visualization [SZP10b, SZP10a, SIGM*11, SMP13, BRSEP13, BRGI*13, PSR13, FMPS13] as well as in-depth reviews of other work on compact visual data representation has triggered the proposal of this tutorial. Our current and future areas of specialization in tensor approximation methods is in the general context of novel multiresolution, hierarchical and out-of-core tensor decomposition models for large scale volume data representation, multi-scale feature extraction and interactive visualization.

ACKNOWLEDGMENTS

This work was supported in parts by the Forschungskredit of the University of Zürich, a Swiss National Science Foundation (SNSF) grant (project n°200021_132521), the EU FP7 People Programme (Marie Curie Actions) under REA Grant Agreement n°290227, as well as the German Science Foundation (DFG) under research grant KL 1142/4-1. Furthermore, we would like to acknowledge the Computer-Assisted Paleanthropology group and the VMML at the University of Zürich for the acquisition of the test datasets.

REFERENCES

- [AB00] ANDERSSON C. A., BRO R.: The n -way toolbox for MATLAB. *Chemometrics and Intelligent Laboratory Systems* 52, 1 (2000), 1–4.
- [BK06] BADER B. W., KOLDA T. G.: Algorithm 862: MATLAB tensor classes for fast algorithm prototyping. *ACM Transactions on Mathematical Software* 32, 4 (December 2006), 635–653.
- [BK*12] BADER B. W., KOLDA T. G., ET AL.: MATLAB tensor toolbox version 2.5. Available online, January 2012.

- [BRGI*13] Balsa Rodriguez M., Gobbetti E., Iglesias Guitián J. A., Makhinya M., Marton F., Pajarola R., Suter S. K.: A survey of compressed GPU direct volume rendering. Eurographics State of The Art Report (STAR), May 2013.
- [BRSEP13] Ballester-Ripoll R., Suter S. K., Elsen A., Pajarola R.: The use of tensor approximation in interactive volume visualization. *submitted to IEEE Transactions on Visualization and Computer Graphics* (2013).
- [CESL*03] Chiang Y.-J., El-Sana J., Lindstrom P., Pajarola R., Silva C. T.: Out-of-core algorithms for scientific visualization and computer graphics. In *IEEE Visualization Tutorial, Course Notes 4* (2003).
- [DGM*08] Dietrich A., Gobbetti E., Manocha D., Marton F., Pajarola R., Slusallek P., Yoon S.-E.: Interactive massive model rendering. In *ACM SIGGRAPH Asia Course Notes* (2008).
- [dLdMV00a] De Lathauwer L., De Moor B., Vandewalle J.: A multilinear singular value decomposition. *SIAM Journal of Matrix Analysis and Applications* 21, 4 (2000), 1253–1278.
- [dLdMV00b] De Lathauwer L., De Moor B., Vandewalle J.: On the best rank-1 and rank- (R_1, R_2, \dots, R_N) approximation of higher-order tensors. *SIAM Journal of Matrix Analysis and Applications* 21, 4 (2000), 1324–1342.
- [FMPS13] Friedland S., Mehrmann V., Pajarola R., Suter S. K.: On best rank one approximation of tensors. *Numerical Linear Algebra with Applications* (2013).
- [KB09] Kolda T. G., Bader B. W.: Tensor decompositions and applications. *SIAM Review* 51, 3 (September 2009), 455–500.
- [KST*10] Kindlmann G. L., Schultz T., Tricoche X., Vasilescu M. A. O., Vilanova A., Zhang E.: Tensors in visualization. In *IEEE Visualization Tutorial* (2010).
- [PSR13] Pajarola R., Suter S. K., Ruiters R.: Tensor approximation in visualization and graphics. Eurographics Tutorial, May 2013.
- [SIGM*11] Suter S. K., Iglesias Guitián J. A., Marton F., Agus M., Elsen A., Zollikofer C. P., Göpi M., Gobbetti E., Pajarola R.: Interactive multiscale tensor reconstruction for multiresolution volume visualization. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (December 2011), 2135–2143.
- [SMP13] Suter S. K., Makhinya M., Pajarola R.: TAMRESH: Tensor approximation multiresolution hierarchy for interactive volume visualization. *Computer Graphics Forum* (2013).
- [SP12] Suter S. K., Pajarola R.: Tensor approximation properties for multiresolution and multiscale volume visualization. Posters IEEE Visualization Conference, October 2012.
- [Sut13] Suter S. K.: *Interactive Multiresolution and Multiscale Visualization of Large Volume Data*. PhD thesis, University of Zurich, Switzerland, April 2013.
- [SVBDL13] Sorber L., Van Barel M., De Lathauwer L.: Tensorlab v1.0. <http://esat.kuleuven.be/sista/tensorlab/>, February 2013.
- [SZP10a] Suter S. K., Zollikofer C. P., Pajarola R.: Application of tensor approximation to multiscale volume feature representations. In *Proceedings Vision, Modeling and Visualization* (2010), pp. 203–210.
- [SZP10b] Suter S. K., Zollikofer C. P., Pajarola R.: *Multiscale Tensor Approximation for Volume Data*. Tech. Rep. IFI-2010.04, Department of Informatics, University of Zürich, February 2010.
- [TS06] Tsai Y.-T., Shih Z.-C.: All-frequency precomputed radiance transfer using spherical radial basis functions and clustered tensor approximation. *ACM Transactions on Graphics* 25, 3 (2006), 967–976.
- [TS12] Tsai Y.-T., Shih Z.-C.: K-clustered tensor approximation: A sparse multilinear model for real-time rendering. *ACM Transactions on Graphics* 31, 3 (May 2012), 19:1–19:17.
- [Tsa09] Tsai Y.-T.: *Parametric Representations and Tensor Approximation Algorithms for Real-Time Data-Driven Rendering*. PhD thesis, National Chiao Tung University, May 2009.
- [vmm] vmmllib: A vector and matrix math library. <http://vmmllib.sf.net>.
- [WWS*05] Wang H., Wu Q., Shi L., Yu Y., Ahuja N.: Out-of-core tensor approximation of multi-dimensional matrices of visual data. *ACM Transactions on Graphics* 24, 3 (Jul 2005), 527–535.
- [WXC*08] Wu Q., Xia T., Chen C., Lin H.-Y. S., Wang H., Yu Y.: Hierarchical tensor approximation of multidimensional visual data. *IEEE Transactions on Visualization and Computer Graphics* 14, 1 (Jan/Feb 2008), 186–199.
- [YMK*09] Yoon S.-E., Manocha D., Kasik D., Gobbetti E., Pajarola R., Slusallek P.: Interactive massive model rendering. In *IEEE Visualization Tutorial* (2009).

A ACRONYMS

ALS Alternating least-squares algorithm

BTF Bidirectional texture function

BRT Biscalar radiance transfer

BRDF Bidirectional reflectance distribution functions

CP CANDECOMP/PARAFAC

CTA Clustered tensor approximation

DCT Discrete cosine transform

DVR Direct volume rendering

HOEIGS Higher-order symmetric eigenvalue decomposition

HOOI Higher-order orthogonal iteration

HOPM Higher-order power method

HOSVD Higher-order singular value decomposition

K-CTA K-clustered tensor approximation

OLS Ordinary least squares

PCA Principal component analysis

PSNR Peak signal-to-noise ratio

RMSE Root mean square error

SVD Singular value decomposition

SH Spherical harmonics

TA Tensor approximation

TTM tensor times matrix

TTM1 tensor times matrix multiplication along mode 1

vmmllib Vector and matrix math library

VOTF View-dependent occlusion texture functions

WT Wavelet transform

B SUMMARY OF RELATED PAPERS

In the following, a selection of related papers are briefly summarized. Besides the major contributions of the papers, it is highlighted what tensor models are used in what context. The first few articles represent good background literature in order to get started with tensor approximation; then we added scientific visualization papers.

B.1 Kolda and Bader, 2009

Kolda and Bader [KB09] present an in-depth survey on various available higher-order tensor decomposition approaches. Besides the well-known Tucker model and CP model, they mention many other (hybrid) decomposition approaches. Hence, this survey is a great introduction to the theory and notations of tensor decompositions. It mentions most of the relevant related background works and gives a summary on the origins and development of tensor approximation. It gives also an overview what different terms are used for the same decomposition approaches, which were developed in parallel for a similar purpose. Furthermore, the main tensor decomposition algorithms are outlined.

With respect to applications, they mention several areas, where TA was applied, however, they do not provide results of own applications. Finally, they give an overview of software for tensor computing that was available before 2009.

[KB09] <http://epubs.siam.org/doi/pdf/10.1137/07070111X>

B.2 De Lathauwer et al, 2000a+b

De Lathauwer et al., introduce in [dLdMV00a] a generalization of many previously mentioned TA-like approaches. Since tensor approximation originated in applied sciences and in various areas in parallel, there was no clear general notation and definition of the tensor approximation concepts available for quite some time. De Lathauwer et al., name the extension of the singular value decomposition (SVD) to higher-orders the *multilinear singular value decomposition*. They give a clear overview on how the SVD can be extended to higher orders and what properties can be maintained with what model. The paper presents many definitions and notation for the multilinear SVD. Furthermore, the basic algorithm to perform an SVD in higher orders is generalized. This is the so-called higher order singular value decomposition or in short the HOSVD. In that context, they explain also the relationship between the SVD computation and the symmetric eigenvalue decomposition, which can be used to replace the SVD under certain constraints. If you look for mathematical definition around computing with TA including mathematical proofs, this is the paper to look at. The concepts are mainly explained with the higher-order extension of the Tucker model. However, they briefly mention the links between the Tucker model and a some other models.

In [dLdMV00b], De Lathauwer et al, present the generalization of the two main tensor decomposition algorithms: the higher-order orthogonal iteration (HOOI) and the higher order power-method (HOPM). Both algorithms belong to the family of alternating least squares (ALS) algorithms, which are applied to find a “best” approximation with a tensor decomposition for given rank conditions. The HOOI is applied to arrive at the Tucker model, the HOPM is applied to reach the CP model. Based on the concept of the matrix rank and the tensor rank, a rank- $(R_1, R_2 \dots R_N)$ approximation is defined for the Tucker model and a rank-R approximation for the CP model. Besides the generalization of the best rank-R and rank- $(R_1, R_2 \dots R_N)$ approximation, they given an overview on the ALS TA contributions that were performed previously. Finally, they explain the limitations of the truncation of tensor decompositions of higher orders.

[dLdMV00a] <http://epubs.siam.org/doi/pdf/10.1137/S0895479896305696>

[dLdMV00b] <http://epubs.siam.org/doi/pdf/10.1137/S0895479898346995>

B.3 Bader and Kolda, 2006

A good toolbox for computing with tensors was provided with the MATLAB tensor toolbox by Bader and Kolda. In [BK06], the main algorithms and their implementations are elaborated for the MATLAB tensor classes. This article provides helpful examples on how to compute with tensors in higher orders. For example, they explain how to multiply with tensors or how to rearrange a tensor into a matrix – both being elementary operations when working with tensor decompositions.

[BK06] <http://delivery.acm.org/10.1145/1190000/1186794/p635-bader.pdf>

B.4 Tsai, 2009

The PhD thesis of Tsai [Tsa09] introduced two novel compression algorithms, notably clustered tensor approximation (CTA) and K-clustered tensor approximation (K-CTA). The main applications are SRBFs and real-time data-driven rendering. The dissertation gives a detailed explanation on how the CTA and K-CTA work and how they are implemented. The development of the new TA algorithms was triggered by the fact that previous TA approaches are not compact enough for efficient reconstruction on the GPU. Therefore, the focus here is to introduce sparse representations and clustering to multi-linear models such as TA. An improved compression ratio with good image quality was achieved. Especially, K-CTA helps to improve smoother boundaries between subtensors by exploiting inter-cluster coherence. CTA and K-CTA seem to have some similarities with other matrix factorizations (e.g., two-stage SVD that exploits inter-block coherence); however, previous approaches did not cover sparse representations.

[Tsa09] http://www.cg.cse.yzu.edu.tw/research/phd/prof/Prof_Tsai.pdf

B.5 Suter, 2013

In the PhD thesis of Suter [Sut13], tensor approximation was chosen as the unique framework in scientific visualization (a) to reduce the actual amount of data, (b) to extract relevant features from the dataset, (c) to visualize the data directly from the mathematical frameworks’ coefficients for compression-domain multiresolution direct volume rendering (DVR). Particularly, the Tucker model was used to represent and compress 3D volume datasets. However, there is an overview of different TA models as well as TA notation and general formulations, too.

The inherent TA bases properties such as spatial selectivity and spatial subsampling were used to model multiresolution data structures. Furthermore, it was shown that the tensor rank can be used to steer feature visualization at different scales (multiscalability). In fact, the tensor rank is a parameter that adjusts (a) the amount of data used for the reconstruction, and (b) the scale of the features visualized in a certain reconstruction. Using more ranks adds details as well as finer scale features to a visualization, using only a few ranks visualizes the most prominent data structure (main statistical direction of the data distribution). Finally, the multiscalability available through TA has been successfully combined with the above mentioned multiresolution TA DVR models.

Moreover, this thesis includes a tensor specific quantization scheme [SIGM*11], which reduces the storage costs of one of the selected multiresolution models to 15 percent of the original data elements. In order to achieve interactive frame rates, a parallel GPU-based tensor reconstruction was developed [SIGM*11]. In fact, it could be shown that the tensor reconstruction overhead is marginal compared to the overall rendering costs. The developed algorithms were applied to large volume datasets up to 68GB (floating point values).

The theory part of the thesis on TA is available in the tutorial notes.

B.6 Wang et al., 2005

Wang et al. [WWS*05] focus in their paper on a tensor decomposition algorithm, which works for input tensors that do not fit into the main memory. They develop a so-called out-of-core ALS and perform experiments for initialization methods of the ALS. Since the computation of the HOSVD, which is used in the ALS algorithms, is expensive, they develop a block-based algorithm to perform a rank- $(R_1, R_2 \dots R_N)$ tensor decomposition. With respect to the ALS initialization, they observed that a random initialization results in the same decomposition as a HOSVD initialization; however, the random initialization was much cheaper. In their experiments, they decompose datasets larger than 10GB on a PC with 1GB memory. Particular applications are BTFs (4th-order tensor with the dimensions: row, column, illumination, and view direction), time-varying BTFs (5th-order tensor) and a 4D volume simulation sequence (4th-order tensor with the spatial dimensions X,Y,Z and time). The compression ratio is analyzed in terms of rate-distortion error based on PSNR.

[WWS*05] <http://delivery.acm.org/10.1145/1080000/1073224/p527-wang.pdf>

B.7 Tsai and Shih, 2006

Tsai and Shih [TS06] present a new data representation and compression approach for precomputed radiance transfer (PRT) based on spherical radial basis functions (SRBFs). They show experiments with clustered principal component analysis and clustered tensor approximation (CTA). They organize the PRTs into clusters of multi-dimensional arrays, which are iteratively updated in order to search for locally optimal solutions. The CTA algorithm has three phases: (1) initialization (obtain initial assignment of cluster members), (2) clustering (iteratively re-classify vertices with the minimum approximation error and repeat until convergence), and (3) approximation (extract optimal basis matrices). Their tensor is organized with the number of views of the BRDFs, the number of SRBF light transfer functions and the number of vertices and is based on the Tucker model. They use the block-based TA approach, as presented in [WWS*05]. In their experiments they compare their own results with OLS projection on SH bases and wavelets.

[TS06] <http://delivery.acm.org/10.1145/1150000/1141981/p967-tsai.pdf>

B.8 Wu et al., 2008

Wu et al. [WXC*08] present a hierarchical tensor approximation approach with so-called tensor ensembles. At each hierarchy level N subtensors of the current level are put into an $(N+1)^{th}$ -order tensor. Then the tensor decomposition is performed collectively in order to exploit more redundancy. They receive one set of factor matrices and one core tensor per hierarchy level, that is a sort of a hierarchical Tucker model. The hierarchy is created by applying rank-reduced TAs to the original tensor ensemble. The residual (error to original), is then further tensor decomposed in the next hierarchy level. Each next hierarchy level is divided into residual subtensors. The multilinear tensor rank $(R_1, R_2 \dots R_N)$ is given per hierarchy, where every next hierarchy level uses half of the rank of the current level. Similar to multiresolution analysis with wavelets, low-frequency components are represented at higher hierarchy levels and high-frequency components are at lower levels. High-frequency components have a smaller spatial support and can therefore be approximated with shorter bases vectors (that is why subdivision of hierarchy levels is performed). With that procedure a progressive reconstruction over the hierarchies is possible. Furthermore, they apply a thresholding of core tensor coefficients and perform a uniform core tensor and factor matrices quantization.

In their experiments, Wu et al. compare their hierarchical TA with wavelets, packet wavelets and single-level TA. The experiments are applied to medical and scientific multidimensional datasets, data-driven rendering (e.g., BTFs) and texture synthesis. The experiments are tested in terms of rate-distortion error based on PSNR.

[WXC*08] <http://ieeexplore.ieee.org/ielx5/2945/4384585/04359486.pdf>

B.9 Suter et al., 2010

In Suter et al. [SZP10a] tensor approximation was applied to direct volume visualization. A volume is represented as 3^{rd} -order Tucker tensor. The main idea of the paper is to use TA to compress data and to extract relevant features. For this, different rank-reduced (truncated) tensor reconstructions are compared. The features that can be visualized from tensor decompositions differ from other feature preserving decomposition approaches such as wavelets. While wavelets preserve an overall data distribution (or an averaged and down-sampled version of the original), TA reveals the major data directions differently. One observation was that TA could reveal features with lower number of coefficients, a second observation was that TA can preserve non-axis-aligned features better than wavelets, and a third observation was that TA makes it possible to show features at multiple spatial scales via truncation. One application of the TA was growth structures in dental material, both, simulated samples and phase-contrast synchrotron

tomography scanned samples. The data reduction levels are analyzed visually and in terms of rate-distortion error based on the RMSE.

[SZP10a] <http://diglib.eg.org/EG/DL/PE/VMV/VMV10/203-210.pdf>

B.10 Suter et al., 2011

In Suter et al. [SIGM*11] the basic observation that TA is a viable tool for multiscale volume feature visualization [SZP10a], was extended to large volumes. The main contributions are a tensor-specific quantization approach of the tensor decomposition coefficients, a GPU-based tensor reconstruction scheme, and the application of feature-preserving volume visualization to large multiresolution datasets. The Tucker model was used within a multiresolution direct volume rendering setup where each octree node was represented as a single tensor decomposition (each original subvolume or brick of size 32^3 is represented with a rank-(16,16,16) tensor decomposition). The results show a real-time interactive rendering system of large volumes (largest input tensor is of size 2048^3). Thanks to the GPU tensor reconstruction scheme, the tensor reconstruction overhead became marginal compared to the overall rendering costs. The observed multiscale feature visualization property of TA observed in [SZP10a] could be further confirmed with examples.

[SIGM*11] <http://ieeexplore.ieee.org/ielx5/2945/6064926/06064978.pdf>

B.11 Tsai and Shih, 2012

Tsai and Shih extend in their work the clustered TA (CTA) [TS06] to K-clustered TA [TS12]. The CTA is extended by introducing inter-cluster coherence and working with compact and sparse clustered TA. The inter-cluster coherence is exploited by assigning each subtensor to more than one clusters, notably to exactly K_m clusters. This approach controls the sparsity of the representation. The inter-cluster coherence helped to improve the boundaries between clusters. The K-CTA algorithm can be seen as a sparse extension of the CTA and a multilinear generalization of the K-SVD. The applied tensor approximation is based on the Tucker model. Since the K-CTA algorithm affects some orthogonality properties in the Tucker bases, an additional SVD is applied to each cluster factor matrix. These post-processed matrices are merged into a single global factor matrix for each mode. The experiments show applications with BTFs, BRTs, and VOTFs.

[TS12] <http://delivery.acm.org/10.1145/2170000/2167077/a19-tsai.pdf>

B.12 Suter et al., 2013

Suter et al. [SMP13] present a novel TA-based multiresolution hierarchy, which uses global bases instead of local brick-based tensor decompositions as in [SIGM*11]. The new hierarchy makes use of properties along the two TA factor matrix axes in order to model the multiresolution hierarchy through spatial subsampling and spatial selection, and to model multiscale volume feature visualization through tensor rank truncation. The key point is that all those multiresolution and multiscale features are modeled in one set of global TA bases, which exploit redundancies available in the multiresolution hierarchy and thus use less storage space. The modeling in one single TA hierarchy is possible thanks to a re-orthogonalization process applied to the global bases along the spatial brick positions (similar to [TS12]). A second contribution of the paper is the developed feature scale parameter. This feature scale parameter is linked to the direct volume rendering system, which allows a user to adaptively visualize volume data sets with his preferences on a feature scale. A high feature scale refers to many details in a data set, a low feature scale refers to only the most basic structures. This makes a flexible and user-guided multiscale and multiresolution direct volume rendering possible.

[SMP13]