

Eurographics 2013 Tutorial: Tensor Approximation in Visualization and Graphics

R. Pajarola¹ and S. K. Suter¹ and R. Ruiters²

¹Visualization and Multimedia Lab, University of Zurich, Switzerland

²Computer Graphics Group, University of Bonn, Germany

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 and computer graphics researchers and practitioners. The course will cover the theoretical background of TA methods, their properties and how to compute them, as well as practical applications of TA methods in visualization and computer graphics contexts. In a first theoretical part, the attendees will be instructed on the necessary mathematical background of TA 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 two application oriented sessions, compact TA data representation in scientific visualization and computer graphics 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 its potential application in visualization and graphics related projects.

Keywords: Tensor decompositions, tensor approximations, Tucker model, CANDECOMP/PARAFAC model, compact visual data representation, higher-order SVD methods, data reduction, interactive volume visualization, multi-resolution and multiscale modeling, clustered tensor decomposition, bidirectional reflectance distribution functions, bidirectional texture functions, precomputed radiance transfer.

1. Organization

Organizers	Prof. Dr. Renato Pajarola, Susanne K. Suter University of Zürich, Switzerland pajarola@acm.org, susuter@ifi.uzh.ch
Lecturers	Renato Pajarola Susanne K. Suter Roland Ruiters
Duration	half-day
Level	intermediate
History	no previous offerings

As outlined in the abstract and further elaborated on below, this tutorial will bring the concepts of tensor approximation (TA) closer to the computer graphics and 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 in fact 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 and it is the right time to give a first and practice oriented introduction into this topic to the graphics and visualization community.

As applicable, we aim to provide practical insight, e.g., using MATLAB, CUDA or C++ code examples.

2. Description

2.1. Overview

The SVD and PCA approaches, which work great for matrix-based data (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), BRD-F/BTF illumination data (mostly 3rd-order or 4-th 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 linear algebra tool to process, analyze and represent complex visual data in novel ways.

The targeted audience consists of graphics and visualization researchers and practitioners with a solid background in linear numerical algebra, graphics techniques and/or visualization methods. The presented methods and techniques exhibit a high 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 graphics and visualization 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 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.

2.2. Syllabus

This tutorial addresses the application of advanced numerical linear algebra tools to compact data representations and interactive visualization of large multidimensional dataset. These datasets arise in many applications in scientific visualization and computer graphics, such as visualization of volumetric datasets, storage of reflectance data, motion synthesis or precomputed radiance transfer. *Tensor approximation* (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 three main parts consisting of the general background and properties of TA methods, followed by practical applications of TA methods in scientific visualization and computer graphics:

Part 1 Introduction of the TA framework

- Tucker and CANDECOMP/PARAFAC (CP) tensor decompositions
- Rank-reduced tensor approximations, ALS methods
- Useful TA properties and features for data visualization
- Frequency analysis and DCT equivalence

Part 2 Applications of TA in scientific visualization

- Implementation details of tensor decomposition and tensor reconstruction algorithms
- Practical examples (MATLAB, vmmllib)
- TA-based volume visualization

Part 3 Applications of TA in rendering and graphics

- Examples for multidimensional datasets in rendering and graphics applications
- Influence of data organization, parametrization and error metric
- Clustering and sparsity
- Processing irregular and sparse input samples

The sessions are designed as follows:

Introduction (Pajarola, 10min)

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

Tensor Decomposition Models (Pajarola, 25min)

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. The two models introduced are the CANDECOMP/PARAFAC (CP) model and the Tucker tensor model. In particular, we elaborate on the tensor decomposition being the generalization of the SVD approach to higher order tensor data, and how a rank-reduced tensor decomposition defines an approximation of the original data.

Properties and Features (Pajarola, 25min)

In this segment, 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 rank-reduced tensor reconstructions. Variations to the standard methods, such as incremental construction of tensor decompositions, are described. Finally, the difference of pre-defined and learned bases is discussed and the equivalence to frequency domain transformations is reviewed.

Scientific Visualization Applications (Suter, 30min)

The goal of this segment is to show various TA applications in the domain of scientific visualization. We start by comparing compact data representation approaches (TA and wavelets [WXC*08, SZP10a]) while showing tensor approximation applications for interactive multiscale scientific volume visualization [SIGM*11]. The examples include hierarchical tensor approximation approaches, which are often used for multiresolution visual data representations.

Coffee Break

Implementation Examples in Scientific Visualization (Suter, 25min)

In scientific visualization, the tensor decomposition is usually carried out as a preprocessing routine, while the reconstruction process has to be performed in real-time. In this tutorial, the basic tensor data structures and tensor decomposition algorithms are explained by the example of two available tensor libraries: The MATLAB Tensor Toolbox [BK06]

and the *vmmlib* (vector and matrix math library) [vmm]. For the tensor reconstruction, we show a GPGPU tensor reconstruction which enables for interactive large data visualization. Within the tutorial, we show performance timings for the decomposition step and the real-time reconstruction for datasets of different sizes such that a potential user gets an impression about the offline/run-time performance of different dataset sizes. During the tutorial and its documentation test datasets and example routines are provided.

Graphics Applications (Ruiters, 30min)

In this segment, we will present applications of tensor approximations for datasets that commonly arise in graphics and rendering. The main focus will be on representations for reflectance data, such as bidirectional reflectance distribution functions (BRDFs) [SZC*07, SKB10, BÖK11] and bidirectional texture functions (BTFs) [FKIS02, VT04, WWS*05, WXC*08], but we will also shortly present some other applications such as precomputed radiance transfer (PRT) [TS06], view-dependent occlusion texture functions (VOTFs) [TS12] or storage and synthesis of motions [Vas02, VBPP05, KTMW08, MLC10]. We will discuss graphics related aspects such as the arrangement of the data in a tensor and the influence of the utilized parametrization. Furthermore, we will show how to modify the error metric during the tensor approximation via the use of per entry weights. This can be used to cope with the high dynamic range that reflectance data exhibit.

Clustering and Sparsity (Ruiters, 25min)

In this part, approaches combining tensor approximations with clustering [TS06, TS12] and approaches utilizing sparse tensor decompositions [RK09] are discussed. These representations are useful for rendering applications, since they offer both, good compression ratios and fast random data access. Additionally, we will discuss an approach [RSK12] to directly fit a tensor representation to sparse and irregular measurements without first computing a dense representation of the tensor. This way, it is possible to process larger tensors without having to store them explicitly. Furthermore, this approach allows to integrate additional regularization constraints such as smoothness into the computation of the tensor approximation.

Summary/Outlook (Pajarola, 10min)

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

2.3. 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 Sec. 3).

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. A number of key applications of tensor methods in visualization and graphics which we review in this tutorial have been presented in [SZP10a, SIGM*11], [RK09, RSK12, TS06, TS12] and [WXC*08]. Other resources are the PhD thesis of Tsai ([Tsa09]) and Suter ([Sut13]).

Based on our extensive practical experience and work with tensor methods, we will provide a number of basic MATLAB examples, based on our own development [vmm] (<https://github.com/VMML/vmmlib>) as well as the MATLAB tensor toolbox [BK06] (<http://www.sandia.gov/~tgkolda/TensorToolbox>) and the MATLAB N-way toolbox [AB00] (<http://www.models.life.ku.dk/nwaytoolbox>). These practical examples including test datasets will be made available to the attendees.

3. 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 a section with scientific visualization papers and a section with papers in computer graphics.

3.1. Tensor Approximation Background Literature

3.1.1. 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>

3.1.2. 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>

3.1.3. 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>

3.1.4. 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 sub-tensors 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

3.1.5. 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.

3.2. Scientific Visualization

3.2.1. 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>

3.2.2. 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>

3.3. Computer Graphics

3.3.1. Vasilescu and Terzopoulos, 2002

Vasilescu and Terzopoulos present a technique, called *TensorFaces* [VT02], which computes a multi-linear extension of the PCA for an image collections of human faces. They use a set of registered photographs depicting faces of 28 male subjects in five head poses under three illuminations and showing three different expressions. These images are stored in a 5^{th} -order tensor with the dimensions: subject, pose, illumination, expression and pixel number. They utilize the N-mode SVD, which computes a Tucker factorization of this tensor via the application of an SVD along each

of the tensors modes. This factorization provides a multi-linear extension of the PCA, which decomposes the dataset into Eigenmodes. In contrast to the often employed PCA, which only provides the global axes of variation for the whole dataset, this technique is able to characterize the axes of variation, ordered by their importance, along each of the modes independently. Furthermore, by multiplying with the corresponding factor matrices, it is possible to get the axes of variation for specific persons, expressions, or poses.

[VT02] http://link.springer.com/chapter/10.1007%2F3-540-47969-4_30

3.4. Furukawa et al., 2002

The paper [FKIS02] by Furukawa et. al is the first work to apply tensor approximation to the compression of BTFs. They evaluated two different tensor layouts, a 3^{rd} order tensor with the dimensions pixel index, view direction and light direction and a 4^{th} order tensor in which the light direction is stored in two modes according to its polar angles. They then utilize a CP decomposition to compress this tensor. In contrast to most other works on BTF compression, they compress the BTF for each triangle individually, resulting in tensors, with a rather small spatial dimension compared to other BTF compression approaches (the tensor sizes are $136 \times 36 \times 72$ and $136 \times 36 \times 12 \times 6$). Since the correlation between independent triangles is not exploited, the compact representation of the angular behavior is more important in this setting than for larger BTFs. Their evaluation shows that in this case the tensor based compression is superior to an SVD based approach and in particular that the 3^{rd} order representation is superior to a 4^{th} order representation.

[FKIS02] <http://diglib.eg.org/EG/DL/WS/EGWR/EGWR02/257-266.pdf>

3.5. Vasilescu et al., 2004

TensorTextures [VT04] by Vasilescu and Terzopoulos is an BTF compression technique based on the N-mode SVD. They represent the BTF as a 3^{rd} -order tensor with the dimensions pixel index and color, view direction and light direction. To compress this tensor, an N-mode SVD is applied and the resulting decomposition is then truncated along its view and light modes, whereas the spatial mode is not reduced allowing for reasonably fast random access. Though this representation does not achieve better compression ratios compared to the PCA, when measured under the RMSE, the authors show that it provides higher flexibility to the user. Since the truncated rank can be selected independently for view and illumination, it is possible to create results which preserve the view dependence well, neglecting illumination information, resulting in sharp images with few highlights and shadows and results which preserve the illumination well, at the cost of view dependent information, resulting in images which are unsharp due to the parallax effects but reproduce highlights more faithful. The PCA in contrast always offers the best compromise between these aspects under the RMSE giving the user no options to change the trade-off.

[VT02] <http://delivery.acm.org/10.1145/1020000/1015725/p336-vasilescu.pdf>

3.5.1. Vlasic et al., 2005

Vlasic et. al present a multi-linear model [VBPP05] for 3D face meshes and demonstrate its usefulness for the imputation of missing data, the synthesis of novel faces and expressions and the tracking of faces in videos. They utilize a database of 3D scans of human faces. After semi-automatically registering the scans to obtain accurate vertex correspondences, these are stored in a tensor. They investigate both a representation via a 3^{rd} -order tensor (vertex number, expression and subject) and a 4^{th} -order tensor (vertex number, viseme, expression and subject). A truncated N-mode SVD is then utilized to obtain multi-linear models of these datasets. In the following, they demonstrate several applications of this model. First, they impute missing data in the tensor and show that this enables them to predict an expression, which was removed from the data set for some of the subjects, based on the data for the other persons and expressions. They then demonstrate that it is possible to synthesize novel expression and persons by modifying the parameters of the multi-linear model. Finally, the multi-linear model is fitted to a video stream which enables the tracking of facial pose and expression parameters. This could be utilized for the editing of videos, e.g., by modifying the appearance, age or performance of an actor.

[VBPP05] <http://delivery.acm.org/10.1145/1080000/1073209/p426-vlasic.pdf>

3.5.2. 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 (4^{th} -order tensor with the dimensions: row, column, illumination, and view direction), time-varying BTFs (5^{th} -order tensor) and a 4D volume simulation sequence (4^{th} -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>

3.5.3. 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>

3.5.4. 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>

3.5.5. Ruiters and Klein, 2009

Ruiters and Klein propose a novel tensor compression technique, the *sparse tensor decomposition* [RK09], and demonstrate the use of this technique for the compression of BTFs.

They store the BTF in a 3^{rd} order tensor, with light direction and color channel together in the first mode, and then view and pixel position in the other two modes. To compress this tensor, they introduce a sparse tensor decomposition, which utilizes the K-SVD [AEB06] to subdivide a tensor along one of its modes into a smaller dictionary tensor and a sparse tensor. By repeatedly performing this decomposition along all of its modes, a tensor is finally decomposed into the product of a 2^{nd} order dictionary tensor and several sparse 3^{rd} tensors. They demonstrate that this decomposition achieves compression ratios for BTFs which are superior to PCA, N-mode SVD and per cluster factorization based techniques. Furthermore, they show that, especially at very high compression ratios, the sparsity of the tensors allows for fast random access into the compressed tensor which results for these cases in faster rendering timings than for PCA based compression in a CPU implementation.

[RK09] <http://diglib.org/EG/DL/CGF/volume28/issue2/PDF/v28i2pp513-522.pdf>

3.5.6. 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>

3.5.7. Ruiters et al., 2012

Ruiters et al. propose an approach [RSK12] to directly fit a spatially varying bidirection reflectance distribution function (SVBRDF) represented as a sum of separable functions to a sparse and irregular set of input samples, utilizing additional smoothness and non-local spatial regularization. When acquiring the spatially varying reflectance for an object with curved surfaces, the variation of the local coordinate system and occlusions result in an irregular and sparse sampling. From these samples, the authors reconstruct an isotropic SVBRDF in a Half/Diff parameterization at a high angular resolution. This results in a five dimensional function with the dimensions color, θ_h , θ_d , ϕ_d and position on the surface. Explicitly storing this function at the utilized high angular samplings ($90 \times 90 \times 180$) for the full spatial resolutions

(256×256 and 512×512) would be completely infeasible due to the enormous sizes of the resulting dataset. Instead a representation as a sum of separable functions, which can be regarded as the continuous analogue of a CP decomposition of the tabulated function, is fitted directly to the samples via an alternating least squares approach. This avoids the necessity to store the full dataset to enable a tensor factorization of the dataset. A regularization term is integrated into the fitting to obtain smooth reconstructions and take advantage of spatial self-similarity. In their evaluation, the authors demonstrate that this technique allows to reconstruct specular highlights which could not be represented in a BTF due to insufficient angular resolution and input sampling.

[RSK12] <http://diglib.org/EG/DL/CGF/volume31/issue2/PDF/v31i2pp315-324.pdf>

4. Available Software

We consider the three presented toolboxes as the most convenient ones for tensor approximation applications; however, there is more tensor software available, as summarized in [KB09].

4.1. 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.

[AB00] <http://www.models.life.ku.dk/nwaytoolbox>

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.

[BK*12] <http://www.sandia.gov/~tgkolda/TensorToolbox>

4.3. Vmmlib

The tensor classes that extend the *vector and matrix math library* vmmlib [vmm], are a C++ implementation with templates. The tensor classes are mainly based on the 3^{rd} -order

implementation of tensors and algorithms; however, 4th-order tensor algorithms are planned. The tensor classes implement tensor unfoldings, tensor multiplications and tensor decomposition approaches for the Tucker model, the CP model and some hybrid block-based TA models. The ALS can be applied using the HOSVD or the HOEIGS (higher-order symmetric eigenvalue decomposition based on the covariance matrix of the unfolded input tensor). VmmLib also supports tensor memory-mapping in order to process large input tensors. So far, only compact tensor decompositions were considered, sparse implementations would need to be added on top. Desired extensions can be integrated by any developer since it is an open-source project.

[vmm] (<https://github.com/VMML/vmmLib>)

5. Lecturers Resumes

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.

Renato Pajarola

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 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 (as co-author) in 2005, and an IADIS Best Paper Award in 2007.

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

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] 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.

Susanne K. Suter

Research Assistant, PhD Candidate
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 PhD candidate and research assistant at the University of Zürich, Switzerland. Her main scientific interest is data reduction and data compression, feature extraction, automation, real-time interactive visualization as well as linear algebra in visualization.

Susanne K. Suter's PhD thesis matches the core topic of the presented Eurographics tutorial. Her main focus in the area is interactive visualization of tensor approximated data from large micro-computed tomography or phase-contrast synchrotron datasets, whereas 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. She 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]. She is a student member of IEEE and ACM.

Roland Ruiters

Research Assistant, PhD Candidate
Computer Graphics Group
University of Bonn, Germany
ruiters@cs.uni-bonn.de
<http://cg.cs.uni-bonn.de/en/people/dipl-inform-roland-ruiters/>

Roland Ruiters is a PhD student and research assistant at the University of Bonn, Germany. His research focuses on

the acquisition, representation and editing of reflectance data in the form of BTFs.

As part of his doctoral studies, which he intends to finish during the next year, he worked on tensor approximations for SVBRDF and BTF data. He showed that a compact representation, which still allows for efficient random access during rendering, can be derived by decomposing a tensor into the product of several sparse tensors [RK09]. He also developed techniques to directly derive a tensor approximation from a sparse and irregular set of measured reflectance samples. This avoids the necessity to resample the measurements into a dense tensor representation prior to the approximation and thus enables the processing of large datasets which would be infeasible otherwise [RSK12].

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 Paleoanthropology 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. 3, 7
- [AEB06] AHARON M., ELAD M., BRUCKSTEIN A.: K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing* 54, 11 (November 2006), 4311–4322. 7
- [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. 2, 3, 4, 7
- [BK*12] BADER B. W., KOLDA T. G., ET AL.: MATLAB tensor toolbox version 2.5. Available online, January 2012. 7
- [BÖK11] BILGILI A., ÖZTÜRK A., KURT M.: A general BRDF representation based on tensor decomposition. *Computer Graphics Forum* 30, 8 (December 2011), 2427–2439. 3
- [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). 8
- [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). 8
- [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. 3, 4
- [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. 3, 4
- [FKIS02] FURUKAWA R., KAWASAKI H., IKEUCHI K., SAKAUCHI M.: Appearance based object modeling using texture database: acquisition, compression and rendering. In *Proceedings Eurographics Workshop on Rendering* (2002), pp. 257–266. 3, 5
- [KB09] KOLDA T. G., BADER B. W.: Tensor decompositions and applications. *SIAM Review* 51, 3 (September 2009), 455–500. 3, 7
- [KTMW08] KRÜGER B., TAUTGES J., MÜLLER M., WEBER A.: Multi-mode tensor representation of motion data. *Journal of Virtual Reality and Broadcasting* 5, 5 (July 2008). 3
- [MLC10] MIN J., LIU H., CHAI J.: Synthesis and editing of personalized stylistic human motion. In *Proceedings ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (2010), pp. 39–46. 3
- [RK09] RUITERS R., KLEIN R.: BTF compression via sparse tensor decomposition. *Computer Graphics Forum* 28, 4 (July 2009), 1181–1188. 3, 6, 7, 9
- [RSK12] RUITERS R., SCHWARTZ C., KLEIN R.: Data driven surface reflectance from sparse and irregular samples. *Computer Graphics Forum* 31, 2 (May 2012), 315–324. 3, 7, 9
- [SIGM*11] SUTER S. K., IGLESIAS GUITIÁN J. A., MARTON F., AGUS M., ELSENER A., ZOLLIKOFER C. P., GOPI M., GOBBETTI E., PAJAROLA R.: Interactive multiscale tensor reconstruction for multiresolution volume visualization. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (Dec 2011), 2135–2143. 2, 3, 4, 5, 8
- [SKB10] SCHWENK K., KUIJPER A., BOCKHOLT U.: Modeling wavelength-dependent brdfs as factored tensors for real-time spectral rendering. In *Proceedings of the International Conference on Computer Graphics Theory and Applications* (2010), pp. 165–172. 3
- [Sut13] SUTER S. K.: *Interactive Multiresolution and Multiscale Visualization of Large Volume Data*. PhD thesis, University of Zurich, Switzerland, April 2013. 3, 4
- [SZC*07] SUN X., ZHOU K., CHEN Y., LIN S., SHI J., GUO B.: Interactive relighting with dynamic BRDFs. *ACM Transactions on Graphics* 26, 3 (2007), 27. 3
- [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. 2, 3, 4, 5, 8
- [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. 8
- [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. 3, 6, 7
- [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. 3, 7
- [Tsa09] TSAI Y.-T.: *Parametric Representations and Tensor Approximation Algorithms for Real-Time Data-Driven Rendering*. PhD thesis, National Chiao Tung University, May 2009. 3, 4
- [Vas02] VASILESCU M. A. O.: Human motion signatures: Analysis, synthesis, recognition. In *Proceedings IEEE Conference on Pattern Recognition* (2002), vol. 3, pp. 456–460. 3
- [VBPP05] VLASIC D., BRAND M., PFISTER H., POPOVIĆ J.: Face transfer with multilinear models. *ACM Transactions on Graphics* 24, 3 (2005), 426–433. 3, 6

- [vmm] vmmllib: A vector and matrix math library.
<http://vmmllib.sf.net>. 3, 7, 8
- [VT02] VASILESCU M. A. O., TERZOPOULOS D.: Multilinear analysis of image ensembles: TensorFaces. In *Proceedings in European Conference on Computer Vision* (May 2002), vol. 2350, pp. 447–460. 5, 6
- [VT04] VASILESCU M. A. O., TERZOPOULOS D.: TensorTextures: Multilinear image-based rendering. *ACM Transactions on Graphics* 23, 3 (2004), 336–342. 3, 5
- [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. 3, 6
- [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. 2, 3, 6
- [YMK*09] YOON S.-E., MANOCHA D., KASIK D., GOBBETTI E., PAJAROLA R., SLUSALLEK P.: Interactive massive model rendering. In *IEEE Visualization Tutorial* (2009). 8

Appendix A: Acronyms

ALS Alternating least-squares algorithm
BTF Bidirectional texture function
BTF Byscale radiance transfer
BRDF Bidirectional reflectance distribution functions
CP CANDECOMP/PARAFAC
CTA Clustered tensor approximation
DCT Discrete cosine transform
DVR Direct volume rendering
FT Fourier transform
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
PRT Precomputed radiance transfer
PSNR Peak signal-to-noise ratio
RMSE Root mean square error
SVBRDF spatially varying bidirection reflectance distribution function
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
VQ Vector quantization
WT Wavelet transform