IEEE Vis 2016 Tutorial: Tensor Decomposition Methods in Visual Computing

Rafael Ballester-Ripoll, *Student Member, IEEE* and Renato Pajarola, *Member, IEEE*

Abstract— Initially proposed as an extension of the concept of matrix decomposition for three and more dimensions, tensor decompositions have found numerous applications in visualization and visual computing. They constitute a powerful mathematical framework for compactly representing and manipulating dense data fields, especially in many dimensions. This course will introduce the most popular decomposition models and showcase emerging tensor methods for compression, interactive visualization, texture synthesis, denoising, and multidimensional inpainting. Multidimensional visual data types of interest include image and geometry ensembles, hyperspectral images, volumes and corresponding time-varying data.

Index Terms—Tensor decompositions, high-dimensional data, compact visual data representation, higher-order singular value decomposition, data reduction, interactive volume visualization, volume compression, multiresolution and multiscale modeling

1 Organization

Organizers	Rafael Ballester-Ripoll, Prof. Dr. Renato Pajarola, University of Zürich, Switzerland
Lecturers	Rafael Ballester-Ripoll Renato Pajarola
Duration	Half-day
Level	Intermediate
History	A more specialized tutorial was held at Eurographics 2013

We will first review the fundamental building blocks of the most relevant tensor models in an introductory manner. We will emphasize learned (data-dependent) bases and their importance in data reduction, as opposed to predefined bases (such as the discrete cosine, Fourier or wavelet transforms). In the context of learned bases, we will showcase the advantages of tensor models over other representations such as vector quantization or dictionary encodings. The typical application pipeline for most tensor methods consists of a) decomposition (data reduction); b) manipulation in the resulting tensor-compressed format (transformation and learning); and c) reconstruction (for final display). We will present use cases from the literature; specific decomposition algorithms will be given less weight. When applicable, we will provide MATLAB or C++ sample code within our presentation slides and point to supplementary material in order to further emphasize on the practical aspects of the tutorial.

Throughout the course we will present higher-order tensor models by highlighting the main differences with their 2D counterparts (which are often more familiar to the visualization and visual computing community). We believe that delivering tensor material in such a comparative way will contribute to the overall clarity of the tutorial. By the end of the course participants will understand how to manipulate high-dimensional spatial and multimodal data using tensor decomposition (in particular the higher-order singular value decomposition, HOSVD [dLdMV00]) and will be familiar with recent applications in the field of visual computing.

To the best of our knowledge, this is the first tutorial with this particular scope and aims. A related tutorial (*Tensor Approximation in Visualization and Graphics*) was held at Eurographics 2013 [PSR13]. That course, however, focused largely on compression applications for rendering. Instead, the present course targets a wider scope by incorporating elements from visual data recovery, synthesis and multilinear learning. The past tutorial received a significant attendance, particularly among those interested in volume visualization and/or linear algebra applications. Previously, a course on tensor methods *Tensors in Visualization* was held at the VisWeek 2010 [KST*10]. It had a different scope, since it covered predominantly visualization of tensor fields, and only treated scalar field tensors briefly at the very end. Our tutorial will be more specific and delve into decompositions of multidimensional scalar fields in detail.

2 Structure

This tutorial is structured in three sections: an introduction to the foundations of the framework, followed by examples of practical visualization applications in a) cases where the goal is to manage and visualize large data (*data reduction*); and b) cases where *learning* is the main challenge (synthesizing new data, imputing missing values, visualizing and gaining understanding of high-dimensional fields).

2.1 Introduction to Tensor Decomposition

Several data set types can be naturally represented as higher-order tensors (multidimensional arrays): volume data (3D), spatio-temporal volume and FMRI data (4D), image stacks and video (3D), BRDF/BTF illumination sample data (5D and more), and collections thereof.

Numerous examples in the literature have established the Tucker model and the closely-related higher-order singular value decomposition (HOSVD [dLdMV00]) as one of the most successful tensor approximations in graphics and visualization. Their cornerstone is the concept of multiway projection (also known as tensor-times-matrix). We will also review the connections between these models and better-known tools such as the SVD, the Fourier and cosine transforms, and wavelets. To finalize this introduction we will provide a short comparison between HOSVD and the more recent tensor train (TT) model [Ose11], whose size grows only linearly with the number of dimensions and is thus better suited for data sets with a higher dimensionality.

2.2 Spatial Data Compression

Tensor decompositions have found use in numerous applications that handle visual dense spatial data on 3D regular Cartesian grids, such as X-ray tomography scans and computer simulations. Sometimes, such data sets are 4-dimensional as they vary with respect to time.

Two main goals in such data-intensive applications are a) to simplify the complex initial input to alleviate computational bottlenecks, while b) aiming for a fast reconstruction as faithful as pos-

• *R. Ballester-Ripoll and R. Pajarola are with the Visualization and MultiMedia Lab, University of Zürich, Switzerland Email: {rballester,pajarola}@ifi.uzh.ch.*

sible. This is the case when memory or time restrictions are an issue, particularly in interactive visualization. Early compression approaches for visual spatial data were proposed in [WA04, WWS*05] and [WXY07, WXC*08]. Progressive tensor rank reduction has been shown to reveal features and structural details at different scales also in volume data [SZP10a, SZP10b]. Further recent efforts in the context of tensor compression include [Tsa09, TS12, BRSP15, BRP15, Tsa15], [SIGM*11, SMP13, BRGI*14] for interactive volume visualization, and [WLHR12] for 3D displays. Multilinear bases resulting from tensor decomposition can be manipulated to reconstruct individual interpolated elements [KTW07], blocks [SMP13] or different resolution levels [SIGM*11]. We will also review tensor compression of bidirectional texture functions [WXC*08, RK09] and bidirectional reflectance distribution functions [BÖK11, RSK12], which are very attractive compression targets due to their intrinsic high dimensionality.

2.3 Tensor Methods for Multilinear Learning

In other settings, the purpose of a rank-reduced tensor is to capture non-obvious structural information from complex (and often large) data sets so that a *derivation* of the data can be newly generated, rather than just an approximation of the input. This includes denoising, inpainting/data recovery, 2D and 3D texture synthesis, etc. This forms the third and last section of our tutorial.

The problem of tensor completion arises often in signal processing and machine learning, and is a useful tool to handle missing values in visual data (e.g. corrupted regions in images or volumes). Advanced multidimensional completion algorithms include [OST08, KSV13, CHL14] and [FJ15]. We will not give in-depth mathematical details on these algorithms, and will focus on application examples and visual results instead.

Methods for image [RRB13] and volume [ZXJ*15] denoising have been proposed using the HOSVD model: essentially, window patches are stacked and filtered by low-rank truncation of the HOSVD core. As for texture synthesis, it has been demonstrated using HOSVD in [CSS06, CSS08] and [WXC*08]. A related topic is face transfer and synthesis [VT02, VT04, VBPP05, VT07]. We will highlight the important role of tensor basis manipulation in many of the aforementioned applications: by computing linear combinations of the basis elements (e.g. factor matrix rows, in the HOSVD case), one can interpolate entries from a tensor of arbitrary dimension.

3 SYLLABUS

Part 1 Introduction (30 minutes)

- Motivation: examples of common multidimensional visual data; limitations of the traditional SVD.
- Connections to frequency transforms and wavelets
- Spatial data manipulation in the tensor-compressed domain (spatial selectivity; tensor rank reduction; convolution)

Part 2 Tensor Compression of Large Data (60 minutes)

- Motivation: bandwidth as the bottleneck of the visualization pipeline
- Tensor compression performance (accuracy, compression rates, reconstruction speed)
- Applications: large scale volume compression; parallel reconstruction; out-of-core solutions; compression of BRDFs and BTFs

Part 3 Tensor Methods for Multilinear Learning (60 minutes)

- Handling missing values: low-rank completion and inpainting
- Multidimensional learning applications: synthesis, denoising and recognition

4 INSTRUCTORS' BACKGROUND

The tutorial is given by two researchers on TA methods in visualization and computer graphics (one young and one experienced). 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, BRGI*14, PSR13, FMPS13, BRGI*14, BRSP15, BRP15] as well as in-depth reviews of other work on compact visual data representation have 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.

In the following, the lecturers' backgrounds and specializations are summarized.

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 Computer Science 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 and high-dimensional spatial data analysis [BRSP15, BRP15] and is maintaining vmmlib, a C++ library for tensor manipulation and decomposition [vmm]. He is a member of IEEE.

Renato Pajarola

Professor, Dr. sc.-techn. ETH
Visualization and MultiMedia Lab (VMML)
University of Zürich, Switzerland
pajarola@ifi.uzh.ch
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 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].

References

[BÖK11] BILGILI A., ÖZTÜRK A., KURT M.: A general BRDF representation based on tensor decomposition. *Computer Graphics Forum 30*, 8 (December 2011), 2427–2439.

[BRGI*14] BALSA RODRÍGUEZ M., GOBBETTI E., IGLESIAS GUITIÁN J. A., MAKHINYA M., MARTON F., PAJAROLA R., SUTER S. K.: State-of-the-art in compressed GPU-based direct volume rendering. *Computer Graphics Forum 33*, 6 (2014), 77–100.

[BRP15] BALLESTER-RIPOLL R., PAJAROLA R.: Lossy volume compression using Tucker truncation and thresholding. *The Visual Computer* (2015), 1–14.

[BRSP15] BALLESTER-RIPOLL R., SUTER S. K., PAJAROLA R.: Analysis of tensor approximation for compression-domain volume visualization. *Computers & Graphics 47* (April 2015), 34–47.

[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).

[CHL14] CHEN Y.-L., HSU C.-T., LIAO H.-Y. M.: Simultaneous tensor decomposition and completion using factor priors. *IEEE Transactions on Pattern Analysis and Machine Intelligence 36*, 3 (2014), 577–591.

[CSS06] COSTANTINI R., SBAIZ L., SÜSSTRUNK S.: Dynamic texture analysis and synthesis using tensor decomposition. In *Advances in Visual Computing*, vol. 4292 of *Lecture Notes in Computer Science*. Springer-Verlag, 2006, pp. 245–254.

[CSS08] COSTANTINI R., SBAIZ L., SÜSSTRUNK S.: Higher order SVD analysis for dynamic texture synthesis. *IEEE Transactions on Image Processing 17*, 1 (2008), 42–52.

[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).

[dLdMV00] DE LATHAUWER L., DE MOOR B., VANDEWALLE J.: On the best rank-1 and rank-$(R_1, R_2, ..., R_N)$ approximation of higher-order tensors. *SIAM Journal of Matrix Analysis and Applications 21*, 4 (2000), 1324–1342.

[FJ15] FILIPOVIĆ M., JUKIĆ A.: Tucker factorization with missing data with application to low-n-rank tensor completion. *Multidimensional Syst. Signal Process. 26*, 3 (July 2015), 677–692.

[FMPS13] FRIEDLAND S., MEHRMANN V., PAJAROLA R., SUTER S. K.: On best rank one approximation of tensors. *Numerical Linear Algebra with Applications* (2013).

[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).

[KSV13] KRESSNER D., STEINLECHNER M., VANDEREYCKEN B.: Low-rank tensor completion by riemannian optimization. *BIT Numerical Mathematics 54*, 2 (2013), 447–468.

[KTW07] KRÜGER B., TAUTGES J., WEBER A.: Multi-mode representation of motion data. In *International Conference on Computer Graphics Theory and Applications* (2007), INSTICC - Institute for Systems and Technologies of Information, Control and Communication, pp. 21–29.

[Ose11] OSELEDETS I. V.: Tensor-train decomposition. *SIAM Journal on Scientific Computing 33*, 5 (September 2011), 2295–2317.

[OST08] OSELEDETS I. V., SAVOSTIANOV D. V., TYRTYSHNIKOV E. E.: Tucker dimensionality reduction of three-dimensional arrays in linear time. *SIAM Journal on Matrix Analysis and Applications 30*, 3 (2008), 939–956.

[PSR13] PAJAROLA R., SUTER S. K., RUITERS R.: Tensor approximation in visualization and graphics. Eurographics Tutorial, May 2013.

[RK09] RUITERS R., KLEIN R.: BTF compression via sparse tensor decomposition. *Computer Graphics Forum 28*, 4 (July 2009), 1181–1188.

[RRB13] RAJWADE A., RANGARAJAN A., BANERJEE A.: Image denoising using the higher order singular value decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence 35*, 4 (April 2013), 849–862.

[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.

[SIGM*11] SUTER S. K., IGLESIAS GUITIÀN 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 (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).

[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.

[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.

[Tsa15] TSAI Y.-T.: Multiway k-clustered tensor approximation: Toward high-performance photorealistic data-driven rendering. *ACM Transactions on Graphics 34*, 5 (October 2015), 157:1–15.

[VBPP05] VLASIC D., BRAND M., PFISTER H., POPOVIĆ J.: Face transfer with multilinear models. *ACM Transactions on Graphics 24*, 3 (2005), 426–433.

[vmm] vmmlib: A vector and matrix math library. http://vmmlib.sf.net.

[VT02] VASILESCU M. A. O., TERZOPOULOS D.: Multilinear analysis of image ensembles: TensorFaces. In *Procedings in European Conference on Computer Vision* (May 2002), vol. 2350, pp. 447–460.

[VT04] VASILESCU M. A. O., TERZOPOULOS D.: TensorTextures: Multilinear image-based rendering. *ACM Transactions on Graphics 23*, 3 (2004), 336–342.

[VT07] VASILESCU M. A. O., TERZOPOULOS D.: Multilinear projection for appearance-based recognition in the tensor framework. In *IEEE International Conference on Computer Vision* (Oct 2007), pp. 1–8.

[WA04] WANG H., AHUJA N.: Compact representation of multidimensional data using tensor rank-one decomposition. In *Proceedings Pattern Recognition Conference* (2004), pp. 44–47.

[WLHR12] WETZSTEIN G., LANMAN D., HIRSCH M., RASKAR R.: Tensor displays: Compressive light field synthesis using multilayer displays with directional backlighting. *ACM Transactions on Graphics 31*, 4 (July 2012), 80:1–80:11.

[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.

[WXY07] WU Q., XIA T., YU Y.: Hierarchical tensor approximation of multidimensional images. In *Proceedings IEEE International Conference in Image Processing* (2007), vol. 4, pp. IV–49–IV–52.

[YMK*09] YOON S.-E., MANOCHA D., KASIK D., GOBBETTI E., PAJAROLA R., SLUSALLEK P.: Interactive massive model rendering. In *IEEE Visualization Tutorial* (2009).

[ZXJ*15] ZHANG X., XU Z., JIA N., YANG W., FENG Q., CHEN W., FENG Y.: Denoising of 3d magnetic resonance images by using higher-order singular value decomposition. *Medical Image Analysis 19*, 1 (2015), 75–86.