Tensor Approximation Properties for Multiresolution and Multiscale Volume Visualization

Susanne K. Suter* Renato Pajarola
Visualization and MultiMedia Lab, University of Zürich, Switzerland

Abstract

Interactive visualization and analysis of large and complex volume data is still a big challenge. Compression-domain volume rendering methods have shown that mathematical tools to represent and compress large data are very successful. We use a new framework that is widely used for data approximation – *tensor approximation* (TA). Specific properties of the TA bases are elaborated in the context of multiresolution and multiscale volume visualization.

Keywords: Tensor Approximation, Multiresolution Volume Visualization, Factor Matrix Properties, Tensor Rank Reduction.

1 Introduction

Todays volume datasets are not only large, but, in particular, exhibit an increasing complexity of internal structure at different spatial scales. To cope with these challenging volume datasets and show variable feature scales and spatial resolutions, we describe the advantageous properties and features of the TA framework [3]. TA is a compression-domain method that works with computed data-specific bases, while other compression-domain approaches like wavelets base on pre-defined filters. [8, 6] have shown that TA is more compact and is able to capture non-axis aligned features at different scales. Wavelets are advantageous to reconstruct the overall statistical distribution of a dataset at coarser resolutions, whereas TA extracts specific features based on statistical properties like the major direction. We aim at the latter.

For state-of-the-art out-of-core volume visualization [4, 2, 1] two key points are: (1) to decompose the data into subvolumes, and (2) to store the data at multiple resolutions in order to load the appropriate data portions on demand. Both key issues are supported intrinsically within the TA framework. Along the spatial axis (rows) of the TA factor matrices, the spatial distribution of the data is represented along the given data direction. It is possible to select subranges of the TA factor matrix rows in order to either (1) reconstruct only a specific subvolume or (2) reconstruct it at a lower resolution by mipmapping [7]. Another TA factor matrices property along the rank axis (columns) is that a column subrange approximates the data at a coarser scale of the visible structural features [6].

TA has successfully been applied to interactive large volume visualization in [5]. However, specific TA matrix properties have not yet been exploited. Here, we describe these TA features that can be utilized for multiresolution and multiscale volume visualization.

2 Background: Tensor Approximation (TA)

TA is a SVD/PCA-like approximation tool that works for higher order data tensors. A tensor is a generalization of a multidimensional array, e.g., a matrix is a second-order tensor. In the context of volume visualization, we represent a volume as a third-order tensor $\mathscr{A} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ with voxels $a_{i_1 i_2 i_3}$ and apply a tensor decomposition as described in [3]. We use the so called Tucker model, which consists of one factor matrix per mode $\mathbf{U}^{(\mathbf{n})} \in \mathbb{R}^{I_n \times R_n}$ and one core tensor, e.g., $\mathscr{B} \in \mathbb{R}^{R_1 \times R_2 \times R_3}$ for a volume. The core tensor $\mathscr{B}$ is in effect a projection of the original data $\mathscr{A}$ onto the basis of the factor matrices $\mathbf{U}^{(\mathbf{n})}$. In case of a volume, the Tucker model has three modes as illustrated in Fig. 1, and defines an approximation $\widetilde{\mathscr{A}} = \mathscr{B} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \times_3 \mathbf{U}^{(3)}$ of the original volume $\mathscr{A}$ (using n-mode products $\times_n$). The voxel-wise Tucker reconstruction of $\widetilde{\mathscr{A}}$ is defined as $\widetilde{a}_{i_1 i_2 i_3} = \sum_{r_1} \sum_{r_2} \sum_{r_3} b_{r_1 r_2 r_3} \cdot u^{(1)}_{i_1 r_1} \cdot u^{(2)}_{i_2 r_2} \cdot u^{(3)}_{i_3 r_3}$, with factor matrix and core tensor entries $u^{(n)}_{i_n r}$ and $b_{r_1 r_2 r_3}$.

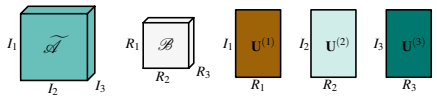

Figure 1: Tucker tensor approximation: $\widetilde{\mathscr{A}} = \mathscr{B} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \times_3 \mathbf{U}^{(3)}$

3 TA Factor Matrix Properties

The TA factor matrices $\mathbf{U}^{(n)}$ resemble a frequency pattern (as e.g. known from DCT-bases), illustrated for the hazelnut dataset in Fig. 2. As can be observed, the matrices reveal two different axis. The vertical axis (rows) shows the spatial distribution of the data, whereas the horizontal axis (columns) refines (frequency) details of the dataset. Seen in the context of a volume, the rows correspond to the spatial dimension along each mode n, and the columns correspond to the approximation quality and the feature scale space.

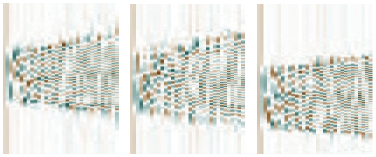

(a) $\mathbf{U}^{(1)}$ (b) $\mathbf{U}^{(2)}$ (c) $\mathbf{U}^{(3)}$

Figure 2: Factor matrix pattern shown for initial matrices (stretched). Value encoding: brown (negative), white (zero), green (positive).

The differences of adjusting the parameters along the two axes are illustrated in Figs. 3 and 4. By selecting a subset of rows, we reconstruct a subvolume; by selection a subset of columns, we reconstruct the full volume, but at a coarser feature scale.

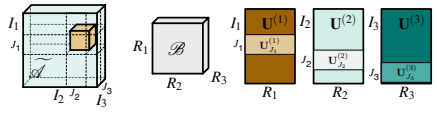

Figure 3: Spatial selectivity: A range of selected submatrix rows reconstructs a defined subvolume (in brown) of the original dataset.

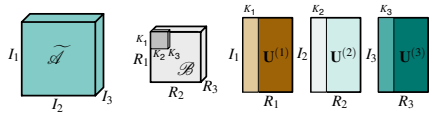

Figure 4: Rank reduced reconstruction: A reduced range of factor matrix columns with corresponding fewer core tensor entries reconstructs a lower quality approximation but at full resolution.

In summary, both axis represent parameter spaces for large volume visualization. We explain how these two properties can be exploited for multiresolution modeling (spatial selection and subsampling, Sec. 3.1) and multiscale approximation (tensor rank reduction, Sec. 3.2).

*e-mail: susuter@ifi.uzh.ch

3.1 Spatial Selectivity

For view-frustum culling and adaptive brick selection in interactive multiresolution volume visualization, efficient access to spatially restricted subvolumes is required. Since a TA factor matrix's rows directly correspond to its spatial dimension, we can reconstruct a subvolume directly from the TA factor matrices. That means, we can define row-index subranges $J_n \subseteq [0 \dots I_n]$ that reconstruct a defined spatial subvolume $J_1 \times J_2 \times J_3$ for the reduced index ranges $i_n \in J_n$ of the voxelwise reconstruction (Sec. 2, Fig. 3). Using these row-block submatrices $\mathbf{U}^{(n)}_{J_n}$ we formulate the subvolume reconstruction as $\widetilde{\mathscr{A}}_{J_1 \times J_2 \times J_3} = \mathscr{B} \times_1 \mathbf{U}^{(1)}_{J_1} \times_2 \mathbf{U}^{(2)}_{J_2} \times_3 \mathbf{U}^{(3)}_{J_3}$. In Fig. 5 an example of two selected subvolumes (1 and 2) is illustrated. For the two different subvolumes, we selected the factor matrix row vectors corresponding to the position of the subvolume in the input dataset.

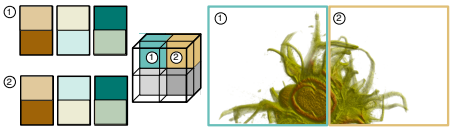

Figure 5: Spatial selectivity of factor matrices. Two selected bricks are reconstructed by the selection of row index subranges.

In multiresolution volume visualization, lower resolution subsampled representations of subvolume bricks are needed for view-dependent adaptive LOD rendering. Due to the direct spatial correspondence of factor-matrix rows to the spatial dimensions, we can apply the lower-resolution subsampling on the factor matrices before the brick reconstruction. Correspondingly, we construct a lower-resolution reconstruction in the n-th mode by averaging pairs of rows to get a downsampled matrix $\mathbf{U}^{(n)}_{\downarrow 1}$ (with $I_n/2$ rows). Downsampling and averaging pairs of rows corresponds to halving the reconstructed volume resolution (known as *mipmapping*). In Fig. 6, we show the factor matrix averaging as used for a multiresolution tensor representation and its effects on the visual reconstructions.

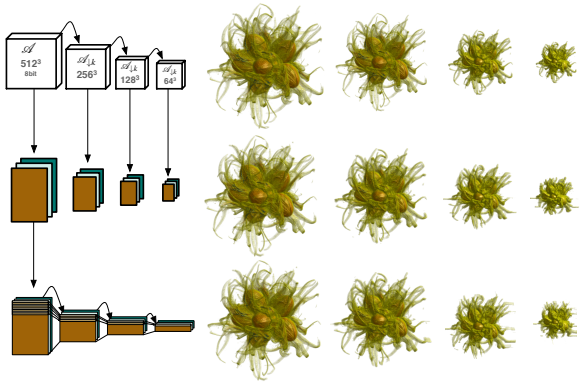

Figure 6: Factor matrix subsampling (bottom row) compared to direct TA (middle row) derived from original subsampled inputs $\mathscr{A}$ (top row).

3.2 Approximation and Rank Reduction

The Tucker model defines a rank-(R_1, R_2, R_3) approximation [3], whereas a small R_n corresponds to a low-rank approximation (many details removed) and a large R_n corresponds to an approximation closer matching the original. The rank R_n for the initial decomposition has to be explicitly given. However, further adaptive rank reductions can be applied after the initial decomposition (similar matrix SVD rank reduction). Even tough the ordering of the coefficients in the core tensor is not strictly decreasing, as in the matrix SVD case, in practice progressive tensor rank reduction in the Tucker model works for adaptive visualization of the data at different feature scales, Fig. 7. The tensor rank parameter R_n is responsible for the number of TA coefficients and bases that are used for the reconstruction and hence is responsible for the approximation level and the storage costs. Note that the reconstruction is still at the original spatial resolution as visualized with examples in Fig. 7, which compares the progressive rank reduction from an initial rank-$(256, 256, 256)$ TA (bottom) to a specific fixed rank-(R_1, R_2, R_3) decomposition (top) for a 512^3 input volume. Even low-rank approximations are able to capture the nuts as main features.

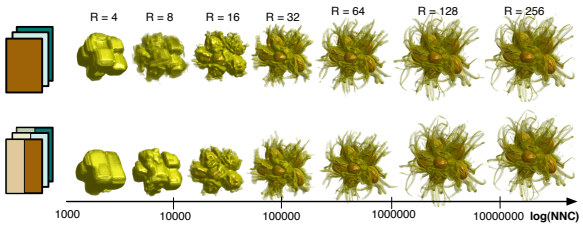

Figure 7: Tensor rank reduction (bottom row) compared to rank-(R, R, R) tensor approximations to specific rank R (top row).

4 Future Work

Interactive visual analysis of large and complex volume datasets is an ongoing and challenging problem, which was addressed e.g., with compression-domain volume rendering and out-of-core multiresolution volume rendering systems. For future work, we aim to improve on [5], which performs a single TA for every multiresolution brick, and exploit TA matrix properties to come up with a novel out-of-core volume rendering system that loads the data from global TA matrices to handle interactive visualization (multiresolution) and feature scale space reconstructions (multiscalability).

Acknowledgements

This work was supported by a Swiss National Science Foundation grant (project n°200021_132521) and the EU FP7 People Programme (Marie Curie Actions, REA Grant Agreement n°290227). Moreover, we acknowledge the CAP group and the VMMLab at University of Zürich for the hazelnut dataset acquisition.

References

[1] C. Crassin, F. Neyret, S. Lefebvre, and E. Eisemann. GigaVoxels: Ray-guided streaming for efficient and detailed voxel rendering. In *Proceedings ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pages 15–22, 2009.

[2] E. Gobbetti, F. Marton, and J. A. I. Guitiàn. A single-pass GPU ray casting framework for interactive out-of-core rendering of massive volumetric datasets. *The Visual Computer*, 24(7-9):797–806, Jul 2008.

[3] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, Sep 2009.

[4] P. Ljung, C. Winskog, A. Persson, C. Lundstrom, and A. Ynnerman. Full body virtual autopsies using a state-of-the-art volume rendering pipeline. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):869–876, Sep/Oct 2006.

[5] S. K. Suter, J. A. Iglesias Guitiàn, F. Marton, M. Agus, A. Elsener, C. P. Zollikofer, M. Gopi, E. Gobbetti, and R. Pajarola. Interactive multiscale tensor reconstruction for multiresolution volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2135–2143, Dec 2011.

[6] S. K. Suter, C. P. Zollikofer, and R. Pajarola. Application of tensor approximation to multiscale volume feature representations. In *Proceedings Vision, Modeling and Visualization*, pages 203–210, 2010.

[7] S. K. Suter, C. P. Zollikofer, and R. Pajarola. Multiscale tensor approximation for volume data. Technical Report IFI-2010.04, Department of Informatics, University of Zürich, Feb 2010.

[8] Q. Wu, T. Xia, C. Chen, H.-Y. S. Lin, H. Wang, and Y. Yu. Hierarchical tensor approximation of multidimensional visual data. *IEEE Transactions on Visualization and Computer Graphics*, 14(1):186–199, Jan/Feb 2008.