

Parallel Algorithms for CP and Tucker Decompositions


Grey Ballard

Low-rank Optimization and Applications Workshop
MPI Leipzig
April 4, 2019




WAKE FOREST
UNIVERSITY

Motivation: multidimensional data analysis requires scalable algorithms





Dynamic functional connectivity fMRI data

- measures correlation between regions of the brain over time
- experiments can include cognitive task
- study multiple subjects across groups




200 regions \times 200 regions \times 225 time steps \times 59 subjects
4 GB of data

CP decomposition discovers patterns of synchronization across brain networks



Motivation: Numerical simulations producing more data than we can handle




512 × 512 × 512 3D grid,
128 time steps, 64 variables:
8 terabytes of data
(double precision)


S3D MPI-based Combustion Code

- direct numerical simulation of engine combustion
- run on supercomputers
- single experiment produces terabytes of data
- storage resolution much less than computed resolution
- difficult to analyze or even transfer data

Tucker decomposition yields huge compression for combustion simulation data



Natural five-way multiway structure of scientific data




Compression rates as fidelity varies for 550GB simulation dataset

Parallel Computation of CP Decompositions with Nonnegativity Constraints

joint work with Srinivas Eswar¹, Koby Hayashi¹,
Ramakrishnan Kannan², and Haesun Park¹

¹ Georgia Tech

² Oak Ridge National Lab



$$\mathcal{X} \approx \mathbf{u}_1 \circ \mathbf{v}_1 \circ \mathbf{w}_1 + \cdots + \mathbf{u}_R \circ \mathbf{v}_R \circ \mathbf{w}_R, \quad \mathcal{X} \in \mathbb{R}^{I \times J \times K}$$


$$\mathcal{X} \approx [\![\mathbf{U}, \mathbf{V}, \mathbf{W}]\!], \quad \mathbf{U} \in \mathbb{R}^{I \times R}, \mathbf{V} \in \mathbb{R}^{J \times R}, \mathbf{W} \in \mathbb{R}^{K \times R}$$

are factor matrices

$$x_{ijk} \approx \sum_{r=1}^R u_{ir} v_{jr} w_{kr}, \quad 1 \leq i \leq I, 1 \leq j \leq J, 1 \leq k \leq K$$

Notation convention: scalar dimension N , index n with $1 \leq n \leq N$

NNCP Optimization Problem



For fixed rank R , we want to solve

$$\arg \min_{\mathbf{u}, \mathbf{v}, \mathbf{w} \geq 0} \left\| \mathcal{X} - \sum_{r=1}^R \mathbf{u}_r \circ \mathbf{v}_r \circ \mathbf{w}_r \right\|$$

Alternating Optimization (AO)

Fixing all but one factor matrix, we have a linear nonnegative least squares (NNLS) problem:

$$\arg \min_{\mathbf{V} \geq 0} \left\| \mathcal{X} - \sum_{r=1}^R \hat{\mathbf{U}}_r \circ \mathbf{V}_r \circ \hat{\mathbf{W}}_r \right\|$$

or equivalently

$$\arg \min_{\mathbf{V} \geq 0} \left\| \mathbf{X}_{(2)} - \mathbf{V}(\hat{\mathbf{W}} \odot \hat{\mathbf{U}})^T \right\|_F$$

\odot is the *Khatri-Rao* product, a column-wise Kronecker product

AO works by alternating over factor matrices, updating one at a time by solving the corresponding linear NNLS problem

Nonnegative (Linear) Least Squares

Generic problem:

$$\arg \min_{\mathbf{Y} \geq 0} \|\mathbf{A}\mathbf{Y} - \mathbf{B}\|_F$$

Many possible algorithms:

- Multiplicative Updates [LS99]
- Hierarchical Alternating Least Squares [CZPA09]
- Block Principal Pivoting [KP11]
- Alternating Direction Method of Multipliers [LS15]
- Nesterov-type Algorithm [LKL⁺17]

Nonnegative (Linear) Least Squares

Generic problem:

$$\arg \min_{\mathbf{Y} \geq 0} \|\mathbf{A}\mathbf{Y} - \mathbf{B}\|_F$$

Our problem:

$$\arg \min_{\mathbf{V} \geq 0} \left\| \mathbf{X}_{(2)} - \mathbf{V}(\hat{\mathbf{W}} \odot \hat{\mathbf{U}})^T \right\|_F$$

All algorithms compute $\mathbf{A}^T \mathbf{B}$ and $\mathbf{A}^T \mathbf{A}$, which for us are

$$\mathbf{A}^T \mathbf{B} \rightarrow \mathbf{X}_{(2)}(\hat{\mathbf{W}} \odot \hat{\mathbf{U}}) \quad \text{and} \quad \mathbf{A}^T \mathbf{A} \rightarrow (\hat{\mathbf{W}} \odot \hat{\mathbf{U}})^T (\hat{\mathbf{W}} \odot \hat{\mathbf{U}})$$

Nonnegative (Linear) Least Squares

Generic problem:

$$\arg \min_{\mathbf{Y} \geq 0} \|\mathbf{A}\mathbf{Y} - \mathbf{B}\|_F$$

Our problem:

$$\arg \min_{\mathbf{V} \geq 0} \left\| \mathbf{X}_{(2)} - \mathbf{V}(\hat{\mathbf{W}} \odot \hat{\mathbf{U}})^T \right\|_F$$

All algorithms compute $\mathbf{A}^T \mathbf{B}$ and $\mathbf{A}^T \mathbf{A}$, which for us are

$$\mathbf{A}^T \mathbf{B} \rightarrow \mathbf{X}_{(2)}(\hat{\mathbf{W}} \odot \hat{\mathbf{U}}) \quad \text{and} \quad \mathbf{A}^T \mathbf{A} \rightarrow (\hat{\mathbf{W}} \odot \hat{\mathbf{U}})^T (\hat{\mathbf{W}} \odot \hat{\mathbf{U}})$$

- $\mathbf{X}_{(2)}(\hat{\mathbf{W}} \odot \hat{\mathbf{U}})$ is called Matricized-Tensor Times Khatri-Rao Product (MTTKRP) and is expensive to compute
- $(\hat{\mathbf{W}} \odot \hat{\mathbf{U}})^T (\hat{\mathbf{W}} \odot \hat{\mathbf{U}})$ can be computed relatively cheaply as $\hat{\mathbf{W}}^T \hat{\mathbf{W}} * \hat{\mathbf{U}}^T \hat{\mathbf{U}}$, where $*$ is elementwise product

Our goal is to perform MTTKRP in parallel as fast as possible

- How do we distribute the tensor across processors?
- How do we distribute the matrices across processors?
- How do we divide up the computation?
- How much interprocessor communication will that require?

Parallel Communication Lower Bound

Theorem ([BKR18])

Any parallel MTTKRP algorithm involving a tensor with $I_k = I^{1/N}$ for all k and that evenly distributes one copy of the input and output performs at least

$$\Omega \left(\left(\frac{NIR}{P} \right)^{\frac{N}{2N-1}} + NR \left(\frac{I}{P} \right)^{1/N} \right)$$

sends and receives. (Second term will typically dominate.)

- N is the number of modes
- I is the number of tensor entries
- I_k is the dimension of the k th mode
- R is the rank of the CP model
- P is the number of processors

Parallelization Approach

- 1 Logically organize processors into N -D grid (match tensor)

$$P_1 \times P_2 \times \cdots \times P_N$$

- 2 Each processor is assigned a (cubical) subtensor of size


$$\frac{l_1}{P_1} \times \frac{l_2}{P_2} \times \cdots \times \frac{l_N}{P_N}$$

- 3 Each processor is assigned a subset of rows of each factor matrix, with dimensions

$$\frac{l_1}{P} \times R, \quad \frac{l_2}{P} \times R, \quad \cdots, \quad \frac{l_N}{P} \times R$$

- 4 Distribute computation based on where tensor data lives, communicating only the factor matrices as needed


Communication-Optimal Parallel Algorithm (3D)



Each processor

- 1 Starts with one subtensor and subset of rows of each input factor matrix


Communication-Optimal Parallel Algorithm (3D)



Each processor

- 1 Starts with one subtensor and subset of rows of each input factor matrix
- 2 All-Gathers all the rows needed from U


Communication-Optimal Parallel Algorithm (3D)



Each processor

- 1 Starts with one subtensor and subset of rows of each input factor matrix
- 2 All-Gathers all the rows needed from U
- 3 All-Gathers all the rows needed from W


Communication-Optimal Parallel Algorithm (3D)



Each processor

- 1 Starts with one subtensor and subset of rows of each input factor matrix
- 2 All-Gathers all the rows needed from **U**
- 3 All-Gathers all the rows needed from **W**
- 4 Computes its contribution to rows of **M** (local MTTKRP)


Communication-Optimal Parallel Algorithm (3D)



Each processor

- 1 Starts with one subtensor and subset of rows of each input factor matrix
- 2 All-Gathers all the rows needed from **U**
- 3 All-Gathers all the rows needed from **W**
- 4 Computes its contribution to rows of **M** (local MTTKRP)


Communication-Optimal Parallel Algorithm (3D)



Each processor

- 1 Starts with one subtensor and subset of rows of each input factor matrix
- 2 All-Gathers all the rows needed from U
- 3 All-Gathers all the rows needed from W
- 4 Computes its contribution to rows of M (local MTTKRP)
- 5 Reduce-Scatters to compute and distribute M evenly

Communication-Optimal Parallel Algorithm (3D)




Each processor

- 1 Starts with one subtensor and subset of rows of each input factor matrix
- 2 All-Gathers all the rows needed from U
- 3 All-Gathers all the rows needed from W
- 4 Computes its contribution to rows of M (local MTTKRP)
- 5 Reduce-Scatters to compute and distribute M evenly
- 6 Use M to solve NNLS problem for V

Rest of the Algorithm

- With correct processor grid, MTTKRP algorithm achieves communication lower bound
- Also need to compute $\mathbf{G} = \mathbf{U}^T \mathbf{U} * \mathbf{W}^T \mathbf{W}$
 - involves communication
 - generally lower order cost
- Lots of overlap across MTTKRP computations
 - save communication: keep temporary copies around
 - save computation: use dimension tree optimization
 - $O(N)$ savings, where N is the number of modes
- Can choose algorithm to compute \mathbf{V} from \mathbf{M} and \mathbf{G}
 - for some algorithms, this is all local computation
 - some algorithms require extra computation of global information, can add significant cost

Weak Scaling Results for 4D Synthetic Data



- local tensor is fixed at $128 \times 128 \times 128 \times 128$

Mouse Brain Data

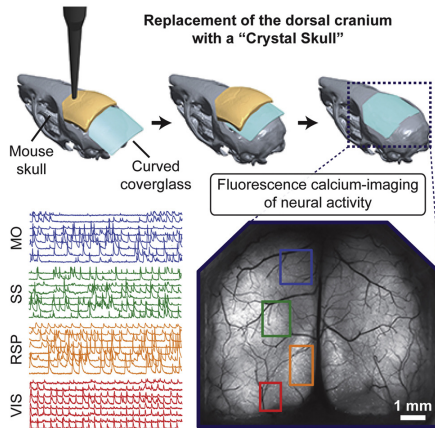




Image from [KZL⁺16]

- tensor is pixels \times time \times trial: $1.4M \times 69 \times 25$
- about 20 GB when stored in double precision

Strong Scaling Results for Mouse Brain Data



Convergence Results for Mouse Brain Data



Parallel Low-rank Approximations with Non-negativity Constraints



<https://github.com/ramkikannan/planc>

- Open source code for computing NMF and NNCP
- MPI/BLAS/LAPACK/C++11
- Designed for dense tensors and dense/sparse matrices
- Can offload computation to GPUs if available

Efficient Parallel Algorithm for Tucker Decompositions of Dense Tensors


joint work with Woody Austin⁴, Alicia Klinvex⁵, Tammy Kolda⁶, and Hemanth Kolla⁶

⁴ UT Austin

⁵ Bettis Atomic Power Laboratory

⁶ Sandia National Labs

Tucker Notation



$$\mathcal{X} \approx \mathcal{G} \times_1 \mathbf{U} \times_2 \mathbf{V} \times_3 \mathbf{W}$$

$$\mathcal{X} \in \mathbb{R}^{I \times J \times K}, \mathcal{G} \in \mathbb{R}^{P \times Q \times R}$$

is core tensor

$$\mathcal{X} \approx [\![\mathcal{G}; \mathbf{U}, \mathbf{V}, \mathbf{W}]\!],$$

$$\mathbf{U} \in \mathbb{R}^{I \times P}, \mathbf{V} \in \mathbb{R}^{J \times Q}, \mathbf{W} \in \mathbb{R}^{K \times R}$$

are factor matrices

$$x_{ijk} \approx \sum_{p=1}^P \sum_{q=1}^Q \sum_{r=1}^R g_{pqr} u_{ip} v_{jq} w_{kr}, \quad 1 \leq i \leq I, 1 \leq j \leq J, 1 \leq k \leq K$$

ST-HOSVD(\mathcal{X}, ε)

- 1 Compute \mathbf{U} with dimension $I \times P$
 - (a) Compute **Gram** matrix $\mathbf{X}_{(1)}\mathbf{X}_{(1)}^T$
 - (b) Use eigendecomposition to determine P and \mathbf{U}
 - (c) **TTM** to shrink to size $P \times J \times K$: $\mathcal{Y} = \mathcal{X} \times_1 \mathbf{U}^T$
- 2 Compute \mathbf{V} with dimension $J \times Q$
 - (a) Compute **Gram** matrix $\mathbf{Y}_{(2)}\mathbf{Y}_{(2)}^T$
 - (b) Use eigendecomposition to determine Q and \mathbf{V}
 - (c) **TTM** to shrink to size $P \times Q \times K$: $\mathcal{Z} = \mathcal{Y} \times_2 \mathbf{V}^T$
- 3 Compute \mathbf{W} with dimension $K \times R$
 - (a) Compute **Gram** matrix $\mathbf{Z}_{(3)}\mathbf{Z}_{(3)}^T$
 - (b) Use eigendecomposition to determine R and \mathbf{W}
 - (c) **TTM** to shrink to size $P \times Q \times R$: $\mathcal{G} = \mathcal{Z} \times_3 \mathbf{W}^T$


Key kernels of ST-HOSVD are

- **Gram**: short, fat matrix times its transpose ($\mathbf{X}_{(1)}\mathbf{X}_{(1)}^T$)
- **Evecs**: eigendecomposition of small symmetric matrix
- **TTM**: tensor times matrix to shrink problem ($\mathbf{U}^T\mathbf{X}_{(1)}$)

Our goal is to parallelize Gram and TTM efficiently

Tensor data distribution across processors

For N -way tensor, we use N -way processor grid with Cartesian block distribution (same as for CP)




Example: $P_1 \times P_2 \times P_3 = 3 \times 5 \times 2$


Local tensor size: $\frac{I}{P_1} \times \frac{J}{P_2} \times \frac{K}{P_3}$

Parallel matricization

Matricizing distributed tensor requires no data movement:
matricized tensor already in standard matrix distribution




Parallel Gram Computation




- each processor column redistributes its tensor data

Parallel Gram Computation



- each processor column redistributes its tensor data
- each processor computes local outer product
- sum across all processors via All-Reduce

Tensor Times Matrix




Tensor-times-matrix (TTM) is matrix multiplication
with matricized tensor

$$\mathcal{X} \times_1 \mathbf{U}' \Rightarrow_P \begin{matrix} I \\ \boxed{} \\ \mathbf{U}' \end{matrix} \begin{matrix} JK \\ \boxed{} \\ \mathbf{X}_{(1)} \end{matrix} = P \begin{matrix} JK \\ \boxed{} \end{matrix}$$

The diagram shows the TTM operation $\mathcal{X} \times_1 \mathbf{U}'$. The tensor \mathcal{X} is contracted with the matrix \mathbf{U}' along the first mode (dimension I). This is represented by the first part of the equation: a box labeled I and \mathbf{U}' is multiplied by a matrix $\mathbf{X}_{(1)}$ of size JK . The result is a matrix of size JK , shown in the second part of the equation. The operation is labeled with \Rightarrow_P and $= P$.


Parallel Tensor Times Matrix

Example: $P_1 \times P_2 \times P_3 = 3 \times 5 \times 2$



- Matrix \mathbf{V} distributed conformally to 2nd mode of tensor \mathcal{X}
- Matrix \mathbf{V} distributed redundantly on processor columns
- Local computation is matrix multiplication
- Communication pattern is reduce-scatter (MPI collective)

Time Breakdown of Parallel ST-HOSVD



Parallel running time example

- 5-way tensor of size 4.4 TB
- reduced to 10 GB (410X)
- 1100 processors (cores)
- 55 seconds total

Observations

- load-balanced execution
- cycle of Gram-Eig-TTM shrinks over time
- writing original tensor to disk is slower by 10X


Weak Scaling on Synthetic Data

Problem Setup


- local tensor fixed at $200 \times 200 \times 200 \times 200$
- local core fixed at $20 \times 20 \times 20 \times 20$

Result

- as problem size grows with number of processors, high efficiency maintained up to 31K cores



Combustion Simulation (S3D) Data



Stat-Planar dataset

- $500 \times 500 \times 500 \times 11 \times 400$
- 4.4 TB of total storage
- use 250 nodes to process


Two compression scenarios

- High: $1e-2$ error, 20,000X comp.
- Low: $1e-4$ error, 400X comp.

Three processor grids


- A: $1 \times 1 \times 40 \times 1 \times 100$
- B: $10 \times 8 \times 5 \times 1 \times 10$
- C: $40 \times 10 \times 1 \times 1 \times 10$

Gram Algorithm Comparison




- Old algorithm from our previous work [ABK16]

Processor Grid Comparison



Flame Surface Reconstruction



Flame surface at single time step.
Using temperature variable (iso-value is 2/3 of max).



<https://gitlab.com/tensors/TuckerMPI>

- Open source code for computing Tucker compression
- MPI/BLAS/LAPACK/C++11
- Designed for dense tensors

- CP and Tucker decompositions are useful and popular tools for multidimensional data analysis, can be applied to large datasets
- Parallel CP bottlenecked by MTTKRP
 - our algorithm's communication cost matches lower bound
 - we avoid redundant computation and communication across modes
- Parallel Tucker bottlenecked by SVD and TTM
 - need communication-efficient distributions and algorithms
 - we can tune the processor grid for efficiency
 - partial reconstruction for analysis can be done on laptop

Parallel Nonnegative CP Decomposition of Dense Tensors

Grey Ballard, Koby Hayashi, and Ramakrishnan Kannan

International Conference on High Performance Computing 2018

<https://ieeexplore.ieee.org/document/8638076>

[BHK18]

TuckerMPI: Efficient Parallel Software for Tucker Decompositions of Dense Tensors

Grey Ballard, Alicia Klinvex, and Tamara G. Kolda

arXiv 2019

<https://arxiv.org/abs/1901.06043>

[BKK19]


Local tensor data layout in memory

Local matricizations (with no data movement)
lead to matrices in funny layouts


Example: $2 \times 2 \times 2 \times 2$ tensor's matricization layouts




$X_{(1)}$



$X_{(2)}$




$X_{(3)}$




$X_{(4)}$



More eyeball norm comparisons (JICF)





Original



$\epsilon = 10^{-4}$
(110X)



$\epsilon = 10^{-2}$
(40000X)




More eyeball norm comparisons (HCCI)

Original


14X Compression

760X Compression




Compressibility depends on data (sing. value decay)

TJLR: 460 x 700 x 360 x 35 x 16



SP-50: 500 x 500 x 500 x 11 x 50



Partial reconstruction

Reconstruction requires as much space as the original data!

$$\hat{\mathbf{X}} = \mathcal{G} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \times_3 \mathbf{U}^{(3)} \times_4 \mathbf{U}^{(4)} \times_5 \mathbf{U}^{(5)}$$

$$N_1 \times N_2 \times N_3 \times N_4 \times N_5$$

But we can just reconstruct the portion that we need at the moment:

$$\tilde{\mathbf{X}} = \mathcal{G} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \times_3 \mathbf{C}^{(3)} \mathbf{U}^{(3)} \times_4 \mathbf{C}^{(4)} \mathbf{U}^{(4)} \times_5 \mathbf{C}^{(5)} \mathbf{U}^{(5)}$$

$$N_1 \times N_2 \times \frac{N_3}{2} \times 1 \times 1$$

$$\mathbf{C}^{(3)} = \begin{bmatrix} 1/2 & 0 & \cdots & 0 \\ 1/2 & 0 & \cdots & 0 \\ 0 & 1/2 & \cdots & 0 \\ 0 & 1/2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \end{bmatrix}$$


Downsample

$$\mathbf{C}^{(4)} = \begin{bmatrix} 0 \\ \vdots \\ 1 \\ 0 \\ \vdots \end{bmatrix}$$


Pick single variable

$$\mathbf{C}^{(5)} = \begin{bmatrix} 0 \\ \vdots \\ 1 \\ 0 \\ \vdots \end{bmatrix}$$

Pick single time step



Tucker compression




$$\underbrace{x}_{I \times J \times K} \approx \underbrace{g}_{P \times Q \times R} \times_1 \underbrace{u}_{I \times P} \times_2 \underbrace{v}_{J \times Q} \times_3 \underbrace{w}_{K \times R}$$

Compression ratio

$$C = \frac{IJK}{PQR + IP + JQ + KR} \approx \frac{IJK}{PQR}$$

Tucker approximation error



$$x_{ijk} \approx \tilde{x}_{ijk} = \sum_{p,q,r} g_{pqr} u_{ip} v_{jq} w_{kr}$$

Approximation error

$$\frac{\|\mathbf{x} - \tilde{\mathbf{x}}\|}{\|\mathbf{x}\|} = \frac{\left(\sum_{i,j,k} \left(x_{ijk} - \sum_{p,q,r} g_{pqr} u_{ip} v_{jq} w_{kr} \right)^2 \right)^{1/2}}{(\sum_{i,j,k} x_{ijk}^2)^{1/2}}$$

Strong scaling benchmark

Problem Setup


- $200 \times 200 \times 200 \times 200$ data tensor (12 GB)
- $20 \times 20 \times 20 \times 20$ core tensor
- $24 \cdot 2^k$ processors (cores)

Result

- small problem, but running time decreases with up to 6144 cores

Compute Platform

- Edison (NERSC), Cray XC30
- 24-core nodes





Woody Austin, Grey Ballard, and Tamara G. Kolda.

Parallel tensor compression for large-scale scientific data.

In *Proceedings of the 30th IEEE International Parallel and Distributed Processing Symposium*, pages 912–922, May 2016.



Grey Ballard, Koby Hayashi, and Ramakrishnan Kannan.

Parallel nonnegative CP decomposition of dense tensors.

In *25th IEEE International Conference on High Performance Computing (HiPC)*, pages 22–31, Dec 2018.



Grey Ballard, Alicia Klinvex, and Tamara G. Kolda.

TuckerMPI: Efficient parallel software for Tucker decompositions of dense tensors.

Technical Report 1901.06043, arXiv, 2019.



Grey Ballard, Nicholas Knight, and Kathryn Rouse.

Communication lower bounds for matricized tensor times Khatri-Rao product.

In Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium, pages 557–567, May 2018.



Andrzej Cichocki, Rafal Zdunek, Anh-Huy Phan, and Shun-ichi Amari.

Nonnegative matrix and tensor factorizations: applications to exploratory multi-way data analysis and blind source separation.

John Wiley & Sons, 2009.



Jingu Kim and Haesun Park.

Fast nonnegative matrix factorization: An active-set-like method and comparisons.

SIAM Journal on Scientific Computing, 33(6):3261–3281, 2011.



Tony Hyun Kim, Yanping Zhang, Jérôme Lecoq, Juergen C. Jung, Jane Li, Hongkui Zeng, Cristopher M. Niell, and Mark J. Schnitzer.

Long-term optical access to an estimated one million neurons in the live mouse cortex.

Cell Reports, 17(12):3385 – 3394, 2016.



A. P. Liavas, G. Kostoulas, G. Lourakis, K. Huang, and N. D. Sidiropoulos.

Nesterov-based alternating optimization for nonnegative tensor factorization: Algorithm and parallel implementation.

IEEE Transactions on Signal Processing, Nov 2017.



Daniel D Lee and H Sebastian Seung.

Learning the parts of objects by non-negative matrix factorization.

Nature, 401(6755):788, 1999.



A. P. Liavas and N. D. Sidiropoulos.

Parallel algorithms for constrained tensor factorization via alternating direction method of multipliers.

IEEE Transactions on Signal Processing, 63(20):5450–5463, Oct 2015.



Nick Vannieuwenhoven, Raf Vandebril, and Karl Meerbergen.

A new truncation strategy for the higher-order singular value decomposition.

SIAM Journal on Scientific Computing, 34(2):A1027–A1052, 2012.