Guarantees for well-posedness of canonical polyadic approximation and numerical linear algebra based estimation

Lieven De Lathauwer

Workshop on Tensor Theory and Methods
Huawei Paris, November 23, 2022

KU LEUVEN

STADIUS
Center for Dynamical Systems, Signal Processing and Data Analytics

EOS THE EXCELLENCE OF SCIENCE

erc European Research Council
Established by the European Commission

Overview

Introduction

Multiplicity based guarantees

Multiple pencil based computation

Positive definiteness based guarantees

Canonical polyadic decomposition

Definition: decomposition in minimal number of rank-1 terms [Harshman 1970; Carrol and Chang 1970]

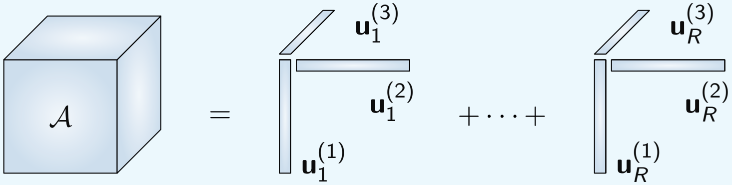

Surprising fact: unique under mild conditions on number of terms and differences between terms

Additional constraints such as orthogonality, triangularity, ... are not required, but may be imposed.

Basic tool for data analysis

row vector $\sim$ excitation spectrum
column vector $\sim$ emission spectrum
coefficients $\sim$ concentrations

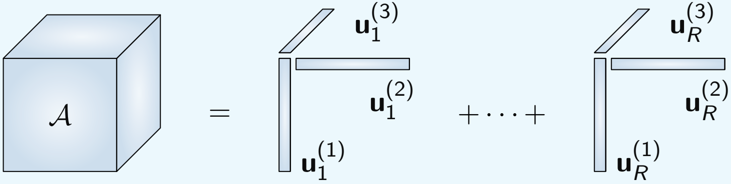

[Smilde, Bro, et al. 2004]

Orthogonality (often) undesired!

Existence of optimal CP approximation

Example:

$$\mathcal{T} = \mathbf{a_1} \otimes \mathbf{b_1} \otimes \mathbf{c_2} + \mathbf{a_1} \otimes \mathbf{b_2} \otimes \mathbf{c_1} + \mathbf{a_2} \otimes \mathbf{b_1} \otimes \mathbf{c_1} \qquad \text{(Rank 3)}$$

$$\mathcal{T}_n = n(\mathbf{a_1} + \frac{1}{n}\mathbf{a_2}) \otimes (\mathbf{b_1} + \frac{1}{n}\mathbf{b_2}) \otimes (\mathbf{c_1} + \frac{1}{n}\mathbf{c_2}) - n\mathbf{a_1} \otimes \mathbf{b_1} \otimes \mathbf{c_1} \qquad \text{(Border rank 2)}$$

$n \to \infty$: terms become large, almost proportional, opposite sign

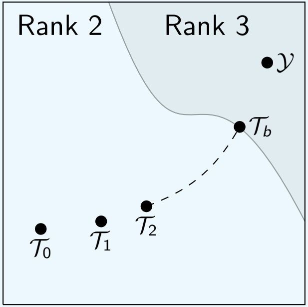

[Kruskal 1983; De Silva and Lim 2008; Qi, Michałek, et al. 2019]

Approximation problem ill-posed?

Matrices: the set $\{\mathbf{M} \in \mathbb{R}^{I \times J} \mid \text{rank}(\mathbf{M}) \leq R\}$ is closed for all R.

Tensor: the set $\{\mathcal{T} \in \mathbb{R}^{I \times J \times K} \mid \text{rank}(\mathcal{T}) \leq R\}$ is only closed for $R = 1$ and $R = R_{\text{max}}$.

Consequence: CP is sometimes ill-posed

Degeneracy: terms $\to \infty$ but partially cancel, fit improves

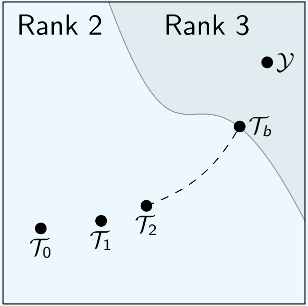

Engineering perspective

Example:

$$\mathcal{T} = \mathbf{a_1} \otimes \mathbf{b_1} \otimes \mathbf{c_2} + \mathbf{a_1} \otimes \mathbf{b_2} \otimes \mathbf{c_1} + \mathbf{a_2} \otimes \mathbf{b_1} \otimes \mathbf{c_1} \qquad \text{(Rank 3)}$$

$$\mathcal{T}_n = n(\mathbf{a_1} + \frac{1}{n}\mathbf{a_2}) \otimes (\mathbf{b_1} + \frac{1}{n}\mathbf{b_2}) \otimes (\mathbf{c_1} + \frac{1}{n}\mathbf{c_2}) - n\mathbf{a_1} \otimes \mathbf{b_1} \otimes \mathbf{c_1} \qquad \text{(Border rank 2)}$$

data tensor is not rank 2
$\leftrightarrow$ estimated tensor is rank 2

Mismatch between "data" and "model"

CPD: in practice often works well (but not always)

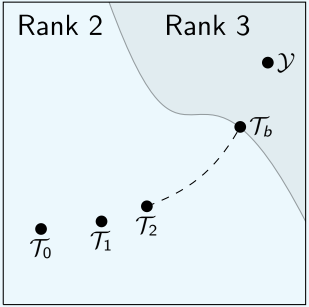

Engineering perspective (2)

Example: EVD

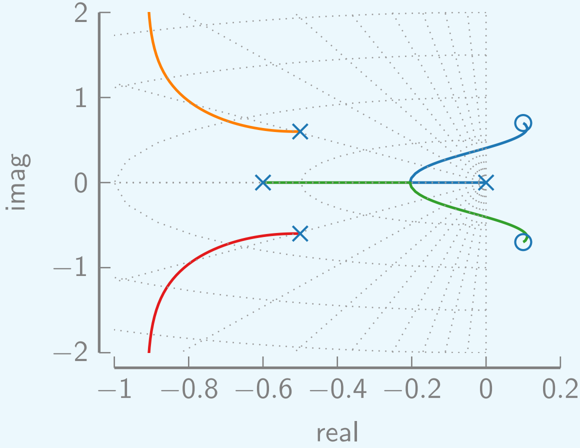

→ Perturbation bounds

Overview

CPD and (G)EVD: $(2 \times 2 \times 2)$ tensors

CPD: $$\mathcal{T} = \sum_{r=1}^{R} \mathbf{a}_r \otimes \mathbf{b}_r \otimes \mathbf{c}_r \quad \text{with } R = 2$$

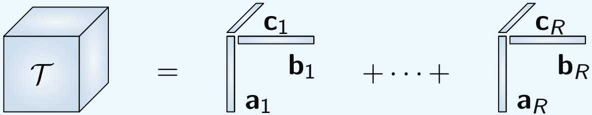

Slices: $$\mathbf{T}_{(:,:,1)} = \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 \end{bmatrix} \begin{bmatrix} c_{11} & \\ & c_{12} \end{bmatrix} \begin{bmatrix} \mathbf{b}_1 & \mathbf{b}_2 \end{bmatrix}^{\mathsf{T}}$$

$$\mathbf{T}_{(:,:,2)} = \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 \end{bmatrix} \begin{bmatrix} c_{21} & \\ & c_{22} \end{bmatrix} \begin{bmatrix} \mathbf{b}_1 & \mathbf{b}_2 \end{bmatrix}^{\mathsf{T}}$$

EVD: $$\mathbf{T}_{(:,:,1)} \cdot \mathbf{T}_{(:,:,2)}^{-1} = \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 \end{bmatrix} \begin{bmatrix} c_{11}/c_{21} & \\ & c_{12}/c_{22} \end{bmatrix} \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 \end{bmatrix}^{-1}$$

real-valued $\leftrightarrow$ complex-valued eigenvalues

Existence of optimal CP approximation: $(2 \times 2 \times 2)$ tensors

Boundary point 2 diverging components:

$$\left(\begin{array}{cc|cc} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{array}\right)$$

EVD $\rightarrow$ Jordan cell:

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}$$

(cf. GRAM)

[Kruskal 1983]

Story for $R \times R \times 2$ tensors is completely told by generalized eigenvalues

The generalized eigenvalues and generalized eigenvectors of $\mathcal{T}$ are (essentially) equal to the classical eigenvalues and eigenvectors of the matrix

$$\mathbf{T}_2^{-1}\mathbf{T}_1$$

Theorem: $\mathcal{T} \in \mathbb{R}^{R \times R \times 2}$ has rank R IFF $\mathcal{T}$ has a basis of generalized eigenvectors.

Idea: use perturbation theory for generalized eigenvalues to guarantee perturbation has distinct generalized eigenvalues.

An existence bound for $R \times R \times 2$ tensors

Theorem

Let $\mathcal{T}$ and $\hat{\mathcal{T}}$ be tensors of size $R \times R \times 2$. Assume that $\mathcal{T}$ has $\mathbb{R}$-rank R with CPD $[\![\mathbf{A}, \mathbf{B}, \mathbf{C}]\!]$. If

$$\|\mathcal{T} - \hat{\mathcal{T}}\|_{sp} < \frac{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B}) \min_{i \neq j} \chi(\mathbf{C}_i, \mathbf{C}_j)}{2},$$

then $\hat{\mathcal{T}}$ has $\mathbb{R}$-rank R and

$$md[\mathcal{T}, \hat{\mathcal{T}}] < \frac{\|\mathcal{T} - \hat{\mathcal{T}}\|_{sp}}{\sigma_{\min}(\mathbf{A})\sigma_{\min}(\mathbf{B})}.$$

Not limited to infinitesimal perturbations!

A multiple pencil based bound for existence

Theorem

Let $\mathcal{M} \in \mathbb{R}^{R \times R \times K}$ be any tensor. For each $i = 1, \ldots, \lfloor K/2 \rfloor$, let $\epsilon_i \geq 0$ be the bound computed using the $K = 2$ theorem for the pencil $(\mathbf{M}_{2i-1}, \mathbf{M}_{2i})$ and set $\epsilon = ||(\epsilon_1, \ldots, \epsilon_{\lfloor K/2 \rfloor})||_2$. If there exists some $\mathbb{R}$-rank R tensor $\mathcal{T}'$ such that

$$\|\mathcal{M} - \mathcal{T}'\|_F < \epsilon,$$

then $\mathcal{M}$ has a best $\mathbb{R}$-rank R approximation and any best $\mathbb{R}$-rank R approximation of $\mathcal{M}$ has a unique CPD.

Note: more pencils $\rightarrow$ relaxed bound

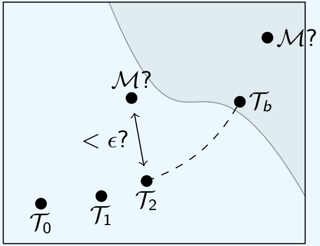

A multiple pencil based bound for existence (improved version)

Theorem

Let $\mathcal{M} \in \mathbb{R}^{R\times R\times K}$ be any tensor. Let $\mathbf{U} \in \mathbb{K}^{K\times K}$ be a unitary matrix and set $\mathcal{S} = \mathcal{M} \cdot_3 \mathbf{U}$. For each $i = 1, \ldots, \lfloor K/2 \rfloor$, let $\epsilon_i \geq 0$ the bound computed using the $K = 2$ theorem for the pencil

$$(\mathbf{S}_{2i-1}, \mathbf{S}_{2i})$$

and set $\epsilon = ||(\epsilon_1, \ldots, \epsilon_{\lfloor K/2 \rfloor})||_2$. If there exists some $\mathbb{R}$-rank R tensor $\mathcal{T}'$ such that

$$\|\mathcal{M} - \mathcal{T}'\|_F < \epsilon,$$

then $\mathcal{M}$ has a best $\mathbb{R}$-rank R approximation and any best $\mathbb{R}$-rank R approximation of $\mathcal{M}$ has a unique CPD.

SNR at which tensors of various sizes are guaranteed (in experiments) to have a best rank R approximation

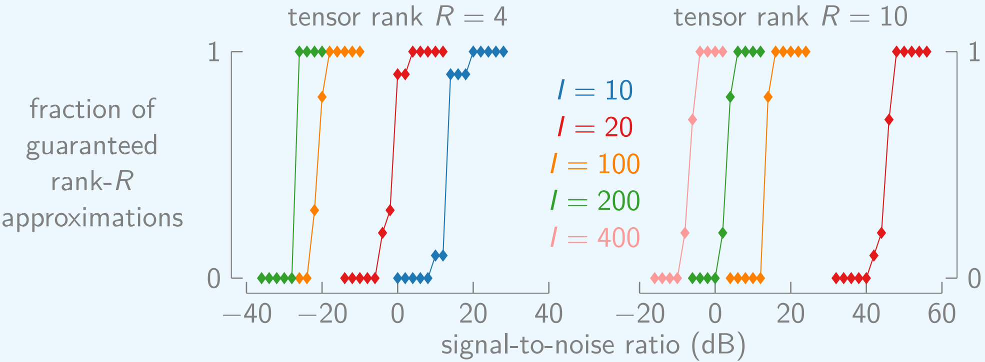

Proportion of $I \times I \times I$ tensors $\mathcal{T} + \mathcal{N}$ with truncated MLSVD guaranteed to have a best rank R approximation.

Confirms engineering practice!

Overview

Algorithm basics: CPD

CPD of a $9 \times 9 \times 9 \times 9 \times 9$ tensor of rank 11

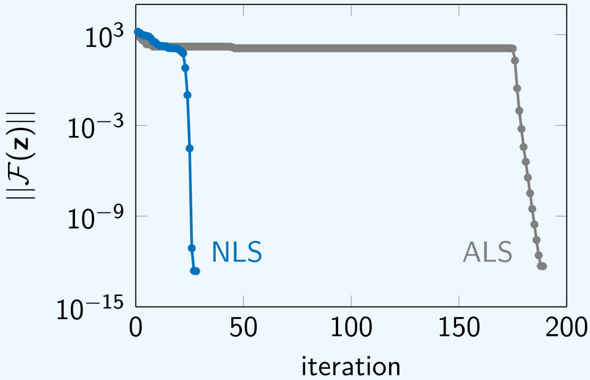

swamp

- init: EVD, random
- global $\leftrightarrow$ asymptotic
- asymptotic convergence: linear - superlinear - quadratic
- unconstrained decomposition $\leftrightarrow$ numerical challenges

Pencil-based computation: numerical implication

CPD:

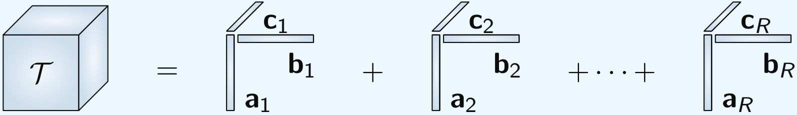

(G)EVD: $$\mathbf{T}_{(:,:,1)} \cdot \mathbf{T}_{(:,:,2)}^{-1} = \begin{bmatrix}\mathbf{a}_1 & \mathbf{a}_2 \cdots \mathbf{a}_R\end{bmatrix} \begin{bmatrix} c_{11}/c_{21} & & & \\ & c_{12}/c_{22} & & \\ & & \ddots & \\ & & & c_{1R}/c_{2R} \end{bmatrix} \begin{bmatrix}\mathbf{a}_1 & \mathbf{a}_2 \cdots \mathbf{a}_R\end{bmatrix}^{-1}$$

Algebraically equivalent but computational differences

- init optimization algorithm
- quantization noise $\rightarrow$ condition number [Beltrán, Breiding, et al. 2019]

CPD structure is collapsed into matrix pencil

Small eigenvalue gaps lead to inaccuracy

Gen. eigenvalues of $(\mathbf{T}_k, \mathbf{T}_\ell)$ are interpreted as points on the unit circle. The pencil $(\mathbf{T}_k, \mathbf{T}_\ell)$ has R generalized eigenvalues.

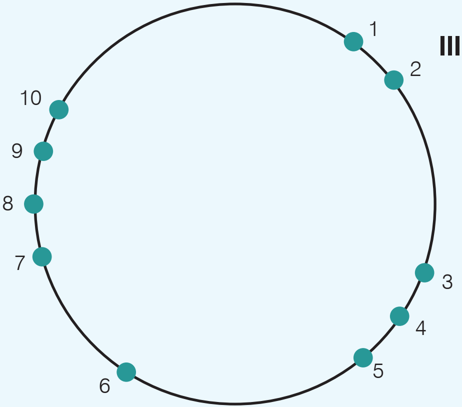

Illustration of generalized eigenvalues of $(\mathbf{T}_k, \mathbf{T}_\ell)$

● = generalized eigenvalue of $(\mathbf{T}_k, \mathbf{T}_\ell)$.

The small gap between generalized eigenvalues 1 and 2 leads to instability in computing the generalized eigenvectors $\mathbf{v}_1$ and $\mathbf{v}_2$.

Similar issues occur in the other clusters of generalized eigenvalues.

Generalized EigenSpace Decomp: Improve accuracy by computing eigenspaces corresponding to well separated eigenvalue clusters.

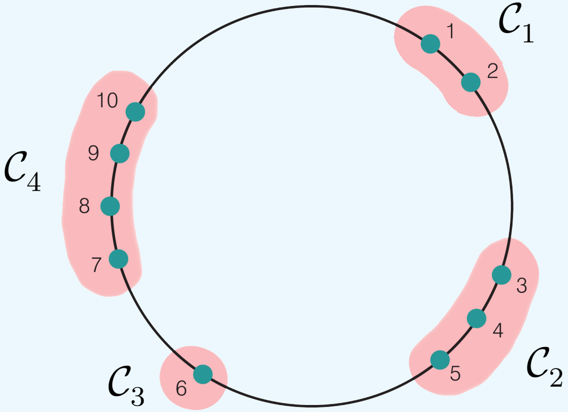

Clusters $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \mathcal{C}_4$ are well separated so can improve accuracy by only computing the corresponding eigenspaces $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3, \mathcal{E}_4$.

Use a new pencil to split eigenspaces!

Consider a new subpencil $(\mathbf{T}_m, \mathbf{T}_n)$. The eigenvectors of this pencil are the same as those of $(\mathbf{T}_k, \mathbf{T}_\ell)$, but the corresponding eigenvalues will lie in new positions on the unit circle.

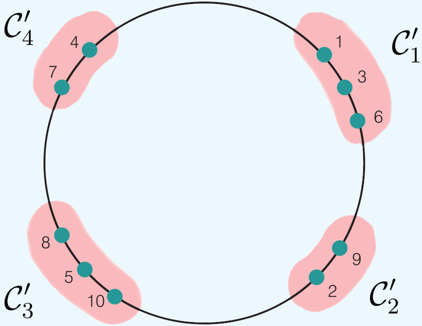

The clusters $\mathcal{C}'_1, \mathcal{C}'_2, \mathcal{C}'_3, \mathcal{C}'_4$ are well separated, so can compute the eigenspaces $\mathcal{E}'_1, \mathcal{E}'_2, \mathcal{E}'_3, \mathcal{E}'_4$.

Observe $\mathcal{E}_1 = \text{span}\{\mathbf{v}_1, \mathbf{v}_2\}$ and $\mathcal{E}'_1 = \text{span}\{\mathbf{v}_1, \mathbf{v}_3, \mathbf{v}_6\}$. Thus $\mathbf{v}_1 = \mathcal{E}_1 \cap \mathcal{E}'_1$.

GESD recursively deflates tensor rank.

In our implementation, GESD recursively writes $\mathcal{T}$ as a sum of tensors of reduced rank.

In the example, GESD would use $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3, \mathcal{E}_4$ to write the rank 10 tensor $\mathcal{T}$ as

$$\mathcal{T} = \mathcal{T}^1 + \mathcal{T}^2 + \mathcal{T}^3 + \mathcal{T}^4$$

where $\mathcal{T}^1, \mathcal{T}^2, \mathcal{T}^3$ and $\mathcal{T}^4$ have ranks $2, 3, 1$ and 4, respectively. $\mathcal{T}^1$ can then be decomposed into a sum of rank 1 tensors using the pencil $(\mathcal{T}^1_m, \mathcal{T}^1_n)$, etc.

Variations in GESD are possible. E.g. one could compute intersections of eigenspaces as described above rather than working recursively.

GESD vs synthetic data

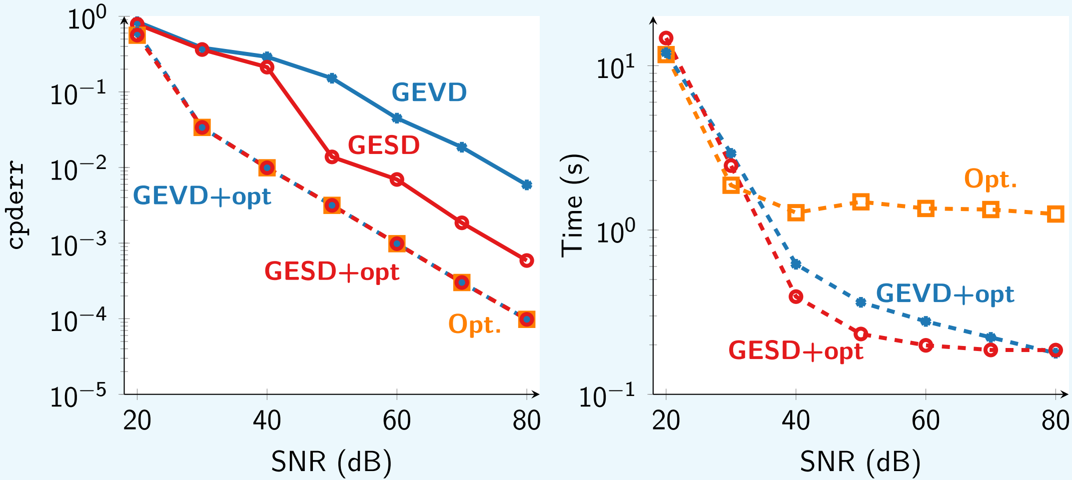

Accuracy and speed against Rank 10 tensors of size $100 \times 100 \times 100$ with highly correlated factor matrix columns.

Overview

Tensors with symmetric frontal slices

Say $\mathcal{T} \in \mathbb{R}^{R \times R \times K}$ has symmetric frontal slices (SFS) if $\mathbf{T}_k$ is symmetric for $k = 1, \ldots, K$.

Any SFS tensor $\mathcal{T}$ can be decomposed as

$$\mathcal{T} = \textstyle\sum_{\ell=1}^{L} \mathbf{a}_\ell \otimes \mathbf{a}_\ell \otimes \mathbf{c}_\ell =$$

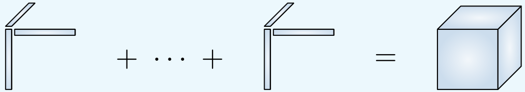

If L is as small as possible, then L is the SFS rank of $\mathcal{T}$.

Rank is not necessarily equal to SFS rank! (Shitov)

Very common in Latent Variable Analysis/Blind Source Separation: slices are statistics (symmetric)

Positive definite matrices and the spectral norm

Say $\mathbf{T} \in \mathbb{R}^{R \times R}$ is **positive definite** if $\mathbf{T}$ is symmetric and all eigenvalues of $\mathbf{T}$ are positive. I.e.

$$\mathbf{v}^\mathsf{T}\mathbf{T}\mathbf{v} = \mathbf{T} \cdot_1 \mathbf{v}^\mathsf{T} \cdot_2 \mathbf{v}^\mathsf{T} > 0 \qquad \text{for all} \;\; \mathbf{0} \neq \mathbf{v} \in \mathbb{R}^R$$

The spectral norm of $\mathcal{T} \in \mathbb{R}^{R \times R \times K}$ is

$$\|\mathcal{T}\|_{sp} = \max_{\|\mathbf{v}_i\|=1} \mathcal{T} \cdot_1 \mathbf{v}_1^\mathsf{T} \cdot_2 \mathbf{v}_2^\mathsf{T} \cdot_3 \mathbf{v}_3^\mathsf{T}$$

Low rank tensors with a positive definite slice mix are relatively closed

Theorem [Evert-De Lathauwer]

Let $\mathcal{T} \in \mathbb{R}^{R \times R \times K}$ be a tensor with border SFS rank at most R. If there is a vector $\mathbf{w} \in \mathbb{R}^K$ such that

$$\mathcal{T} \cdot_3 \mathbf{w}^T \succ 0.$$

Then $\mathcal{T}$ has rank and SFS rank equal to R.

Note: $\mathcal{T} \cdot_3 \mathbf{w}^{\mathsf{T}} = \sum_{k=1}^{K} w_k \mathbf{T}_k$.

Idea: PD property can be used in similar manner as eigenvalue multiplicity!

Spectral norm bound guaranteeing existence of best low rank approximation

Theorem [Evert-De Lathauwer]

Let $\mathcal{T}, \mathcal{N} \in \mathbb{R}^{R \times R \times K}$ and assume $\mathcal{T}$ has SFS rank R. If

$$\|\mathcal{N}\|_{sp} < \max_{\|\mathbf{w}\|=1} \min_{\|\mathbf{v}\|=1} (\mathcal{T} + \mathcal{N}) \cdot_1 \mathbf{v}^T \cdot_2 \mathbf{v}^T \cdot_3 \mathbf{w}^T$$

then $\mathcal{T} + \mathcal{N}$ has a best rank R approximation.

Consequence: Suppose you have some noisy rank R tensor $\mathcal{T} + \mathcal{N} \in \mathbb{R}^{R \times R \times K}$, and let $\hat{\mathcal{T}}$ be any rank R approximation to $\mathcal{T} + \mathcal{N}$. If

$$\|\mathcal{T} + \mathcal{N} - \hat{\mathcal{T}}\|_{sp} < \min_{\|\mathbf{v}\|=1} (\mathcal{T} + \mathcal{N}) \cdot_1 \mathbf{v}^{\mathrm{T}} \cdot_2 \mathbf{v}^{\mathrm{T}} \cdot_3 \mathbf{w}^{\mathrm{T}}$$

then $\mathcal{T} + \mathcal{N}$ has a best rank R approximation.

Computing our bound

Theorem [Evert-De Lathauwer]

Let $\mathcal{T} \in \mathbb{R}^{R\times R\times K}$ and assume $\mathcal{T}$ has SFS. The quantity

$$\lambda_* = \max_{\|\mathbf{w}\|=1} \min_{\|\mathbf{v}\|=1} \mathcal{T} \cdot_1 \mathbf{v}^T \cdot_2 \mathbf{v}^T \cdot_3 \mathbf{w}^T$$

is computable via semidefinite programming

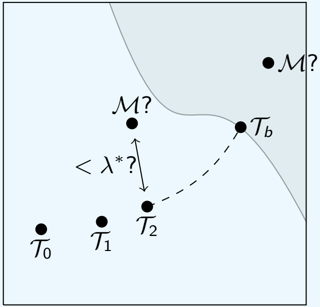

Sharpness of the bound

Theorem [Evert-De Lathauwer]

Let $\mathcal{T} \in \mathbb{R}^{R\times R\times K}$ and assume $\mathcal{T}$ has SFS rank R. Set

$$\lambda_* = \max_{\|\mathbf{w}\|=1} \min_{\|\mathbf{v}\|=1} \mathcal{T} \cdot_1 \mathbf{v}^T \cdot_2 \mathbf{v}^T \cdot_3 \mathbf{w}^T$$

and assume $\lambda_* \geq 0$. Then there exists a tensor $\mathcal{N}_* \in \mathbb{R}^{R\times R\times K}$ with $\|\mathcal{N}_*\|_{sp} = \lambda_*$ such that no linear combination of frontal slices of $\mathcal{T} + \mathcal{N}_*$ is positive definite.

Furthermore, if $K = 2$, then any open set containing $\mathcal{T} + \mathcal{N}_*$ contains a tensor which does not have a best rank R approximation.

Numerical experiments: Second order blind identification

Spectral distance to approximation (dB)

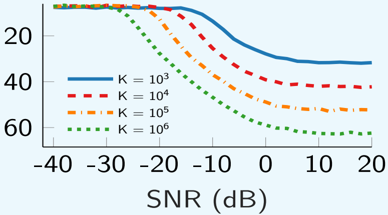

Mixing Matrix Error (dB)

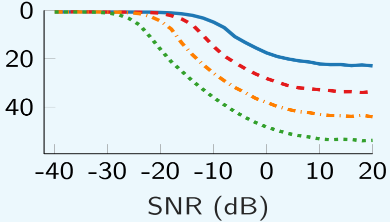

Bound Magnitude

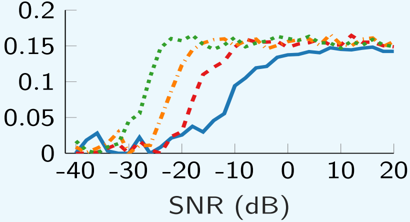

Proportion of tensors guaranteed to have best approximation

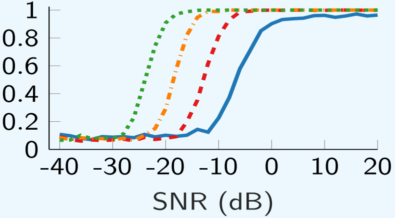

Bound sharpness vs. $4 \times 4 \times 2$ tensors. Approximations of $\mathcal{T} + \alpha\mathcal{N}_*$

Min column angle quantiles

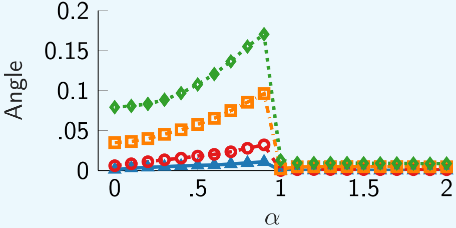

Factor matrix error quantiles

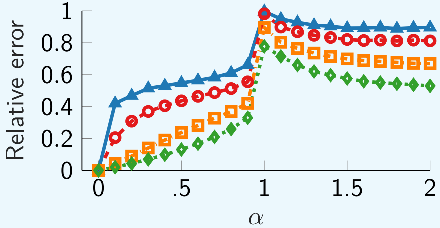

Max column norm quantiles

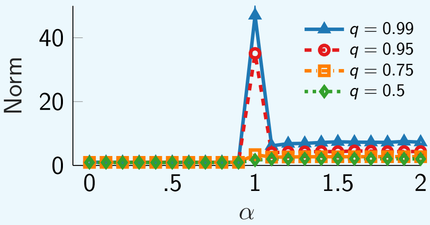

Proportion of MLSVD cores with repeated or complex generalized eigenvalues

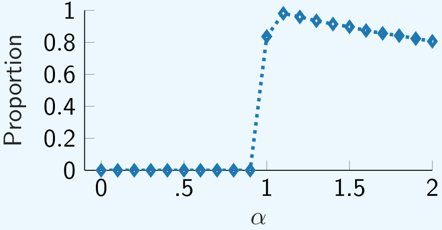

Bound sharpness vs. $4 \times 4 \times 4$ tensors. Approximations of $\mathcal{T} + \alpha\mathcal{N}_*$

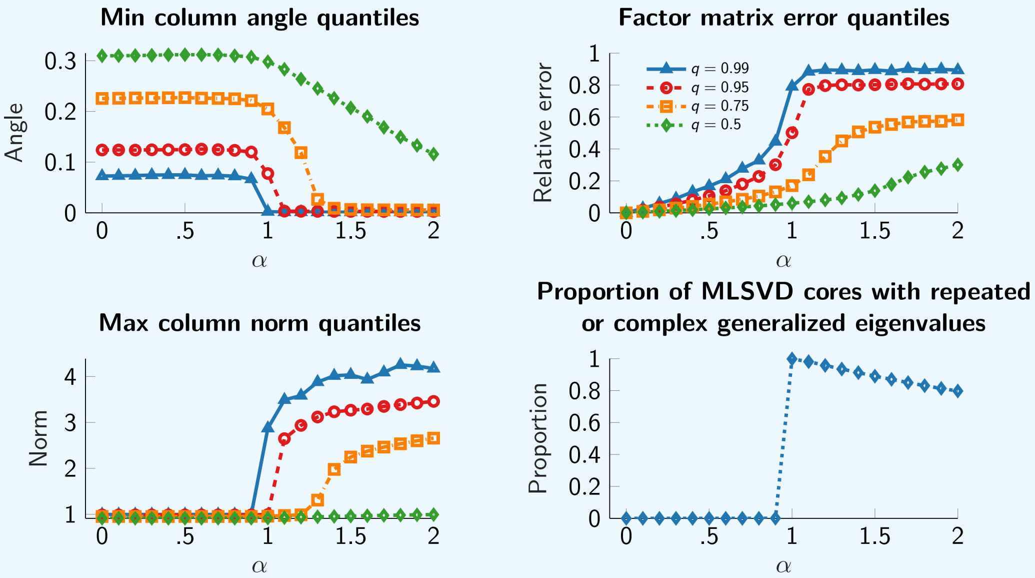

Guarantees for well-posedness of canonical polyadic approximation and numerical linear algebra based estimation

Lieven De Lathauwer

Workshop on Tensor Theory and Methods
Huawei Paris, November 23, 2022

KU LEUVEN STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics EOS THE EXCELLENCE OF SCIENCE erc European Research Council Established by the European Commission

References I

- Beltrán, C., P. Breiding, and N. Vannieuwenhoven (Jan. 2019). “Pencil-Based Algorithms for Tensor Rank Decomposition are not Stable”. In: *SIAM Journal on Matrix Analysis and Applications* 40.2, pp. 739–773.
- Carrol, J. D. and J. J. Chang (Sept. 1970). “Analysis of individual differences in multidimensional scaling via an *n*-way generalization of “Eckart–Young” decomposition”. In: *Psychometrika* 35.3, pp. 283–319.
- De Silva, V. and L.-H. Lim (Sept. 2008). “Tensor Rank and the Ill-Posedness of the Best Low-Rank Approximation Problem”. In: *SIAM Journal on Matrix Analysis and Applications* 30.3, pp. 1084–1127.
- Evert, E. and L. De Lathauwer (Mar. 2022a). “Guarantees for existence of a best canonical polyadic approximation of a noisy low-rank tensor”. In: *SIAM Journal On Matrix Analysis And Applications* 43.1, pp. 328–369.

References II

- Evert, E. and L. De Lathauwer (2022b). "On best low rank approximation of positive definite tensors". In: *Technical Report 22-45, ESAT-STADIUS, KU Leuven (Leuven, Belgium)*.
- Evert, E., M. Vandecappelle, and L. De Lathauwer (Feb. 2022a). "A recursive eigenspace computation for the canonical polyadic decomposition". In: *SIAM Journal On Matrix Analysis And Applications* 43.1, pp. 274–300.
- — (Mar. 2022b). "Canonical Polyadic Decomposition via the generalized Schur decomposition". In: *IEEE Signal Processing Letters* 29, pp. 937–941.
- Harshman, R. A. (1970). "Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multi-modal factor analysis". In: *UCLA Working Papers in Phonetics* 16, pp. 1–84.
- Kruskal, J. B. (1983). "Statement of some current results about three-way arrays". In: *Unpublished manuscript AT&T Bell Laboratories, Murray Hill, NJ*.

References III

Qi, Y., M. Michałek, and L.-H. Lim (2019). "Complex best r-term approximations almost always exist in finite dimensions". In: *Applied and Computational Harmonic Analysis.*

Smilde, A. K. et al. (2004). *Multi-way analysis with applications in the chemical sciences*. Wiley Chichester, UK.