Pencil-based algorithms for tensor rank decomposition are unstable

September 13, 2018

Nick Vannieuwenhoven
(FWO / KU Leuven)

Overview

Hitchcock (1927) introduced the **tensor rank decomposition**:

$$\mathcal{A} = \sum_{i=1}^{r} \mathbf{a}_i^1 \otimes \mathbf{a}_i^2 \cdots \otimes \mathbf{a}_i^d$$

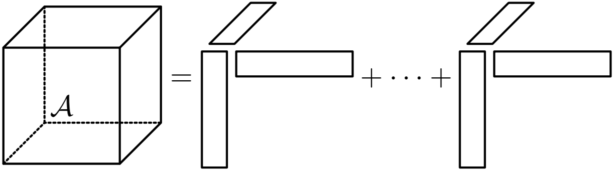

The **rank** of a tensor is the minimum number of rank-1 tensors of which it is a linear combination.

A tensor rank decomposition is also called a canonical polyadic decomposition (CPD).

If the set of rank-1 tensors $\{\mathcal{A}_1, \ldots, \mathcal{A}_r\}$ is uniquely determined given the rank-r tensor $\mathcal{A} = \mathcal{A}_1 + \cdots + \mathcal{A}_r$, then we call $\mathcal{A}$ an r-**identifiable** tensor.

Note that matrices are never r-identifiable, because

$$M = \sum_{i=1}^{r} \mathbf{a}_i \otimes \mathbf{b}_i = AB^T = (AX^{-1})(BX^T)^T$$

for every invertible X. In general, these factorizations are different.

Kruskal (1977) gave a famous sufficient condition for proving the r-identifiability of third-order tensors.

More recently r-identifiability was studied in algebraic geometry. This is a natural framework because the set of rank-1 tensors

$$\mathcal{S} := \left\{ \mathbf{a}^1 \otimes \mathbf{a}^2 \otimes \cdots \otimes \mathbf{a}^d \mid \mathbf{a}^k \in \mathbb{C}^{n_k} \setminus \{0\} \right\}$$

is the smooth projective **Segre variety**.

The set of tensors of rank bounded by r,

$$\sigma_r^0(\mathcal{S}) := \left\{ \mathcal{A}_1 + \cdots + \mathcal{A}_r \mid \mathcal{A}_i \in \mathcal{S} \right\},$$

is the Zariski-open constructible part of the projective r-**secant variety** of the Segre variety.

The number of distinct CPDs is an upper-semicontinuous function on $\overline{\sigma_r^0(\mathcal{S})}$, and its minimum value is called the r-**secant order** s_r, which was initially studied by Chiantini and Ciliberto (2001, 2006).

More precisely, there exists a Zariski-open subset of $\overline{\sigma_r^0(\mathcal{S})}$ where the number of distinct CPDs equals s_r.

If the r-secant order $s_r = 1$ then $\sigma_r(\mathcal{S})$ is called **generically** r-**identifiable**.

Generic r-identifiability of the tensors in $\mathbb{C}^{n_1} \otimes \cdots \otimes \mathbb{C}^{n_d}$,

$$\mathcal{A} = \sum_{i=1}^{r} \mathbf{a}_i^1 \otimes \cdots \otimes \mathbf{a}_i^d \quad \text{with } \mathbf{a}_i^k \in \mathbb{C}^{n_k},$$

is conjecturally understood because of

1. Strassen (1983) for $d = 3$ (partial result);
2. Bocci and Chiantini for $n_1 = \cdots = n_d = 2$;
3. Bocci, Chiantini, and Ottaviani (2013) for unbalanced cases;
4. Chiantini, Ottaviani, and V (2014) for $n_1 \cdots n_d \leq 15000$;
5. Abo, Ottaviani, and Peterson (2009); Chiantini and Ottaviani (2012); Chiantini, Mella, and Ottaviani (2014); etc.

Let $n_1 \geq \cdots \geq n_d$, $r_{\text{cr}} = \frac{n_1 \cdots n_d}{1+\sum_{i=1}^{d}(n_i - 1)}$, $r_{\text{ub}} = n_2 \cdots n_d - \sum_{k=2}^{d}(n_k - 1)$.

Conjectured general rule:

if $r \geq r_{\text{cr}}$ or $d = 2$	$\rightarrow$	not generically r-identifiable
if $n_1 > r_{\text{ub}}$ and $r \geq r_{\text{ub}}$	$\rightarrow$	not generically r-identifiable
if none of foregoing and $r < r_{\text{cr}}$	$\rightarrow$	generically r-identifiable

The real case is more involved because now

$$\sigma_r(\mathcal{S}_\mathbb{R}) := \left\{ \mathcal{A}_1 + \cdots + \mathcal{A}_r \mid \mathcal{A}_i \in \mathcal{S}(\mathbb{R}) \right\},$$

is only a **semi-algebraic set**.

Qi, Comon, and Lim (2016) showed that if $\sigma_r(\mathcal{S})$ is generically r-identifiable, then it follows that the set of real rank-r tensors with multiple complex CPDs is contained in a proper Zariski-closed subset of $\sigma_r(\mathcal{S}_\mathbb{R})$. In this sense, $\sigma_r(\mathcal{S}_\mathbb{R})$ is thus also generically r-identifiable.

See Angelini (2017) and Angelini, Bocci, Chiantini (2017) for more results on complex versus real identifiability.

Overview

Sensitivity

In numerical computations, the **sensitivity** of the output of a computation to **perturbations** at the input is very important, because representation and roundoff errors will corrupt any mathematical inputs.

Consider the matrix

$$A = \frac{1}{177147}\begin{bmatrix} 88574 & 88574 & 2 \\ 88574 & 88574 & 2 \\ 2 & 2 & 177146 \end{bmatrix}$$

Computing the singular value decomposition $\widehat{U}\widehat{S}\widehat{V}^T$ of the floating-point representation $\widetilde{A}$ of A numerically using Matlab, we find $\|A - \widehat{U}\widehat{S}\widehat{V}^T\| \approx 5.66 \cdot 10^{-16}$.

The singular values are

numerical	exact
0.0000000000000000098..	0
0.9999830649121916	$0.99998306491219156971328 8\ldots = 1 - 3^{-10}$
1.000016935087808	$1.00001693508780843028671 1\ldots = 1 + 3^{-10}$

In all cases, we found 16 correct digits of the exact solution.

However, when comparing the computed left singular vector corresponding to $\sigma_1 = 1 + 3^{-10}$ to the exact solution, we get

numerical	exact
0.577350269183747	$\frac{1}{\sqrt{3}}$
0.577350269183748	$\frac{1}{\sqrt{3}}$
0.577350269921281	$\frac{1}{\sqrt{3}}$

We have only recovered 11 digits correctly, even though the matrix $\widehat{U}\widehat{S}\widehat{V}^T$ contains at least 15 correct digits of each entry.

How is this possible?

We say that the problem of computing the singular values has a different sensitivity to perturbations than the computational problem of computing the left singular vectors.

Assuming the singular values are distinct, these problems can be modeled as functions

$$f_1 : \mathbb{F}^{m\times n} \to \mathbb{F}^{\min\{m,n\}}, \text{ respectively } f_2 : \mathbb{F}^{m\times n} \to \mathbb{F}^{m\times \min\{m,n\}}.$$

What we have observed above is that

$$0.4 \approx \frac{\|f_1(x) - f_1(x+\delta x)\|}{\|\delta x\|} \ll \frac{\|f_2(x) - f_2(x+\delta x)\|}{\|\delta x\|} \approx 800$$

at least $x = A$ and $\delta x = \widetilde{A} - A$ (with $\|\delta x\| \approx 5 \cdot 10^{-16}$).

Condition numbers

The **condition number** quantifies the **worst-case sensitivity** of f to perturbations of the input.

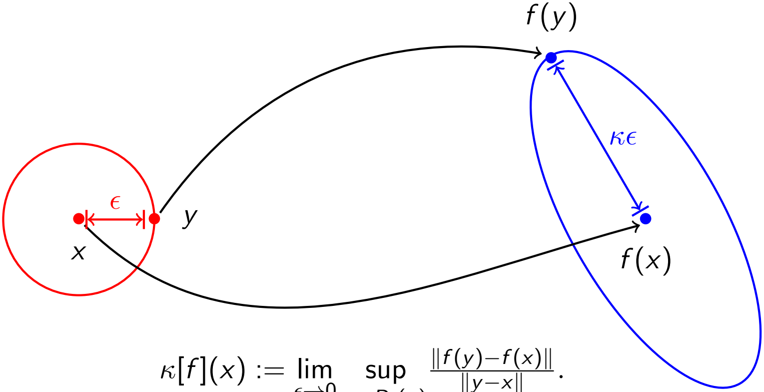

$$\kappa[f](x) := \lim_{\epsilon \to 0} \sup_{y \in B_\epsilon(x)} \frac{\|f(y)-f(x)\|}{\|y-x\|}.$$

If $f : \mathbb{F}^m \supset X \to Y \subset \mathbb{F}^n$ is a differentiable function, then the condition number is fully determined by the first-order approximation of f.

Indeed, in this case we have

$$f(\mathbf{x} + \mathbf{\Delta}) = f(\mathbf{x}) + J\mathbf{\Delta} + o(\|\mathbf{\Delta}\|),$$

where J is the Jacobian matrix containing all first-order partial derivatives. Then,

$$\begin{aligned}\kappa &= \lim_{\epsilon \to 0} \sup_{\|\mathbf{\Delta}\| \le \epsilon} \frac{\|f(\mathbf{x}) + J\mathbf{\Delta} + o(\|\mathbf{\Delta}\|) - f(\mathbf{x})\|}{\|\mathbf{\Delta}\|} \\ &= \max_{\|\mathbf{\Delta}\|=1} \frac{\|J\mathbf{\Delta}\|}{\|\mathbf{\Delta}\|} = \|J\|_2.\end{aligned}$$

More generally, for manifolds, we can apply Rice's (1966) **geometric framework of conditioning**:[1]

Proposition (Rice, 1966)

Let $\mathcal{X} \subset \mathbb{F}^m$ be a manifold of inputs and $\mathcal{Y} \subset \mathbb{F}^n$ a manifold of outputs with $\dim \mathcal{X} = \dim \mathcal{Y}$. *Then, the condition number of $F : \mathcal{X} \to \mathcal{Y}$ at $x_0 \in \mathcal{X}$ is*

$$\kappa[F](x_0) = \|\mathrm{d}_{x_0} F\| = \sup_{\|x\|=1} \|\mathrm{d}_{x_0} F(x)\|,$$

where $\mathrm{d}_{x_0} F : \mathrm{T}_{x_0}\mathcal{X} \to \mathrm{T}_{F(x_0)}\mathcal{Y}$ is the **derivative**.

[1] See, e.g., Blum, Cucker, Shub, and Smale (1998) or Bürgisser and Cucker (2013) for a more modern treatment.

The tensor decomposition problem

The condition number of the problem of computing CPDs was only recently investigated by Breiding and V (2018), after an initial study of a related problem involving CPDs in V (2017). I discuss the strategy we detailed in Beltrán, Breiding, and V (2018).

In the remainder, $\mathcal{S} = \mathcal{S}(\mathbb{R})$. To compute the condition number, we analyze the **addition map**:

$$\begin{aligned} \Phi_r : \mathcal{S} \times \cdots \times \mathcal{S} &\to \mathbb{R}^{n_1 \times \cdots \times n_d} \\ (\mathcal{A}_1, \ldots, \mathcal{A}_r) &\mapsto \mathcal{A}_1 + \cdots + \mathcal{A}_r \end{aligned}$$

Note that the domain and codomain are **smooth manifolds**.

For simplicity, we restrict the domain of Φ_r to a Zariski-open smooth submanifold such that Φ_r restricts to a **diffeomorphism** onto its image.

A set of vectors $\mathbf{p}_1, \ldots, \mathbf{p}_r \in \mathbb{R}^n$ is in **general linear position** (GLP) iff every subset of $\min\{r, n\}$ vectors is linearly independent.

A set of rank-1 tensors $\{\mathbf{a}_i^1 \otimes \cdots \otimes \mathbf{a}_i^d\}_i$ is in **super GLP** iff for every $1 \leq s \leq d$ and every subset $\mathbf{h} \subset \{1, \ldots, d\}$ of cardinality s, the set $\{\mathbf{a}_i^{h_1} \otimes \cdots \otimes \mathbf{a}_i^{h_s}\}_i$ is in GLP.

Let $\mathbf{n} = (n_1, \ldots, n_d)$. Let $\mathcal{M}_{r;\mathbf{n}} \subset \mathcal{S}^{\times r}$ be the set of tuples of $n_1 \times \cdots \times n_d$ rank-1 tensors $\mathfrak{a} = (\mathcal{A}_1, \ldots, \mathcal{A}_r)$ that satisfy:

1. $\Phi_r(\mathfrak{a})$ is a **smooth point** of the semi-algebraic set $\sigma_r^0(\mathcal{S})$;
2. $\Phi_r(\mathfrak{a})$ is r-**identifiable**;
3. the **derivative** $\mathrm{d}_{\mathfrak{a}}\Phi_r$ **is injective**;
4. $\mathfrak{a}$ is in **super GLP**;
5. for all $i = 1, \ldots, r$, $(\mathcal{A}_i)_{1,\ldots,1} \neq 0$.

Definition

The set of r-**nice tensors** is

$$\mathcal{N}_{r;\mathbf{n}} := \Phi_r(\mathcal{M}_{r;\mathbf{n}}).$$

One can prove the following results:

Proposition

Let $\mathcal{S}$ be generically r-identifiable. Then, $\widehat{\mathcal{M}}_{r;\mathbf{n}} := \mathcal{M}_{r;\mathbf{n}}/\mathfrak{S}_r$ is a manifold and the projection is a local diffeomorphism.

Proposition

Let $\mathcal{S}$ be generically r-identifiable. Then,

$$\Phi_r : \widehat{\mathcal{M}}_{r;\mathbf{n}} \to \mathcal{N}_{r;\mathbf{n}}, \{\mathcal{A}_1, \ldots, \mathcal{A}_r\} \to \mathcal{A}_1 + \cdots + \mathcal{A}_r$$

is a diffeomorphism. Moreover, $\mathcal{N}_{r,\mathbf{n}}$ is an open dense submanifold of $\sigma_r^0(\mathcal{S})$.

The inverse of Φ_r, restricted to the manifold of r-nice tensors, is

$$\tau_{r;\mathbf{n}} : \mathcal{N}_{r;\mathbf{n}} \to \widehat{\mathcal{M}}_{r;\mathbf{n}},\ \mathcal{A}_1 + \cdots + \mathcal{A}_r \to \{\mathcal{A}_1, \ldots, \mathcal{A}_r\},$$

which we call the **tensor rank decomposition map**.

As $\tau_{r;\mathbf{n}}$ is a smooth map between manifolds we can apply the standard framework. Since $\tau_{r;\mathbf{n}} \circ \Phi_r = \mathrm{Id}_{\mathcal{N}_{r;\mathbf{n}}}$ we have at $\mathcal{A} \in \mathcal{N}_{r;\mathbf{n}}$ that $\mathrm{d}_{\mathcal{A}}\tau_{r;\mathbf{n}} \circ \mathrm{d}_{\mathfrak{a}}\Phi_r = \mathrm{Id}_{\mathrm{T}_{\mathfrak{a}}\mathcal{N}_{r;\mathbf{n}}}$, so that

$$\kappa[\tau_{r;\mathbf{n}}](\mathcal{A}) = \|\mathrm{d}_{\mathcal{A}}\tau_{r;\mathbf{n}}\|_2 = \|(\mathrm{d}_{\mathfrak{a}}\Phi_r)^{-1}\|_2.$$

The derivative $\mathrm{d}_{\mathfrak{a}}\Phi$ is seen to be the map

$$\begin{aligned} \mathrm{d}_{\mathfrak{a}}\Phi : \mathrm{T}_{\mathcal{A}_1}\mathcal{S} \times \cdots \times \mathrm{T}_{\mathcal{A}_r}\mathcal{S} &\to \mathrm{T}_{\mathcal{A}}\mathbb{R}^{n_1\times\cdots\times n_d} \\ (\dot{\mathcal{A}}_1, \ldots, \dot{\mathcal{A}}_r) &\mapsto \dot{\mathcal{A}}_1 + \cdots + \dot{\mathcal{A}}_r. \end{aligned}$$

Hence, if U_i is an orthonormal basis of $\mathrm{T}_{\mathcal{A}_i}\mathcal{S} \subset \mathrm{T}_{\mathcal{A}_i}\mathbb{R}^{n_1\times\cdots\times n_d}$, then the map is represented in coordinates as the matrix

$$U = \begin{bmatrix} U_1 & U_2 & \cdots & U_r \end{bmatrix} \in \mathbb{R}^{n_1\cdots n_d \times r \dim \mathcal{S}}$$

Summarizing, if we are given a CPD $\mathfrak{a}$ of $\mathcal{A}$, then the condition number of computing this CPD may be computed as the inverse of the smallest singular value of U.

Interpretation

If

$$\mathcal{A} = \mathcal{A}_1 + \cdots + \mathcal{A}_r = \sum_{i=1}^{r} \mathbf{a}_i^1 \otimes \cdots \otimes \mathbf{a}_i^d$$

$$\mathcal{B} = \mathcal{B}_1 + \cdots + \mathcal{B}_r = \sum_{i=1}^{r} \mathbf{b}_i^1 \otimes \cdots \otimes \mathbf{b}_i^d$$

are tensors in $\mathbb{R}^{n_1 \times \cdots \times n_d}$, then for $\|\mathcal{A} - \mathcal{B}\|_F \approx 0$ we have the **asymptotically sharp bound**

$$\underbrace{\min_{\pi \in \mathfrak{S}_r} \sqrt{\sum_{i=1}^{r} \|\mathcal{A}_i - \mathcal{B}_{\pi_i}\|_F^2}}_{\text{forward error}} \lesssim \underbrace{\kappa[\tau_{r;\mathbf{n}}](\mathcal{A})}_{\text{condition number}} \cdot \underbrace{\|\mathcal{A} - \mathcal{B}\|_F}_{\text{backward error}}$$

Overview

An algebraic algorithm

In some cases, the CPD of third-order tensors can be computed via a **generalized eigendecomposition (GEVD)**.

For simplicity, assume that $\mathcal{A} \in \mathbb{R}^{n \times n \times n}$ is of rank n. Say

$$\mathcal{A} = \sum_{i=1}^{n} \mathbf{a}_i \otimes \mathbf{b}_i \otimes \mathbf{c}_i.$$

The steps are as follows.

1. Compute the **multilinear multiplication**

$$\mathcal{X} = (I, I, Q^T) \cdot \mathcal{A} := \sum_{i=1}^{n} \mathbf{a}_i \otimes \mathbf{b}_i \otimes (Q^T \mathbf{c}_i) \quad \in \mathbb{R}^{n \times n \times 2}$$

where $Q \in \mathbb{R}^{n \times 2}$ is a fixed matrix with orthonormal columns.

2. The two 3-**slices** X_1 and X_2 of $\mathcal{X}$ are

$$X_j = \sum_{i=1}^{n} \langle \mathbf{q}_j, \mathbf{c}_i \rangle \mathbf{a}_i \otimes \mathbf{b}_i = A \,\mathrm{diag}(\mathbf{q}_j^T C) B^T$$

where $A = [\mathbf{a}_i] \in \mathbb{R}^{n \times n}$ and likewise for B and C.

Hence, $X_1 X_2^{-1}$ has the following eigenvalue decomposition:

$$X_1 X_2^{-1} = A \,\mathrm{diag}(\mathbf{q}_1^T C) \,\mathrm{diag}(\mathbf{q}_2^T C)^{-1} A^{-1}$$

from which A can be found as the matrix of eigenvectors.

3. By a **1-flattening** we find

$$\mathcal{A}_{(1)} := \sum_{i=1}^{n} \mathbf{a}_i (\mathbf{b}_i \otimes \mathbf{c}_i)^T = A(B \odot C)^T,$$

where $B \odot C := [\mathbf{b}_i \otimes \mathbf{c}_i]_i \in \mathbb{R}^{n^2 \times n}$. Computing

$$A \odot (A^{-1}\mathcal{A}_{(1)})^T = A \odot (B \odot C) = [\mathbf{a}_i \otimes \mathbf{b}_i \otimes \mathbf{c}_i]_i,$$

solves the (ordered) tensor decomposition problem.

Let's perform an experiment in Tensorlab v3.0 with this decomposition algorithm.

Create the first tensor that comes to mind, a rank-25 random tensor of size $25 \times 25 \times 25$:

```
>> Ut{1} = randn(25,25);
>> Ut{2} = randn(25,25);
>> Ut{3} = randn(25,25);
>> A = cpdgen(Ut);
```

Compute $\mathcal{A}$'s decomposition and compare its distance to the input decomposition, relative to the machine precision $\epsilon \approx 2 \cdot 10^{-16}$:

```
>> Ur = cpd_gevd(A, 25);
>> E = kr(Ut) - kr(Ur);
>> norm( E(:), 2 ) / eps
ans =
    8.6249e+04
```

Of course, this can happen because of a high condition number. However,

```
>> kappa = condition_number( Ut )
ans =
    2.134
```

The only explanation is that there is something wrong with the algorithm.

Beltrán, Breiding, and V (2018) show that algorithms based on a reduction to $\mathbb{R}^{n_1 \times n_2 \times 2}$ are **numerically unstable**: the forward error produced by the algorithm divided by the backward error is "much" larger than the condition number, for some inputs.

Distribution of the condition number

We conceived the existence of this problem after seeing the distribution of the condition number of random rank-1 tensors

$$\mathcal{A}_i = \alpha_i \mathbf{a}_i \otimes \mathbf{b}_i \otimes \mathbf{c}_i \in \mathbb{R}^{n_1 \times n_2 \times n_3}$$

where

- $\alpha_i \in \mathbb{R}_+$, $\mathbf{a}_i \in \mathbb{S}^{n_1 - 1}$ and $\mathbf{b}_i \in \mathbb{S}^{n_2 - 1}$ are **arbitrary**, and
- the $\mathbf{c}_i$'s are random vectors i.i.d. on the unit sphere $\mathbb{S}^{n_3 - 1}$.

Informally, we showed, based on Cai, Fan, and Jiang (2013), that

$$\mathrm{P}[\kappa \geq \alpha] \geq \mathrm{P}\left[\max_{1 \leq i \neq j \leq r} \frac{1}{\sqrt{1 - \langle \mathbf{c}_i, \mathbf{c}_j \rangle}} \geq \alpha \right] \to 1 - e^{-K r^2 \alpha^{1 - n_3}},$$

as $r \to \infty$; herein, K is a constant depending only on n_3.

Distribution of the condition number

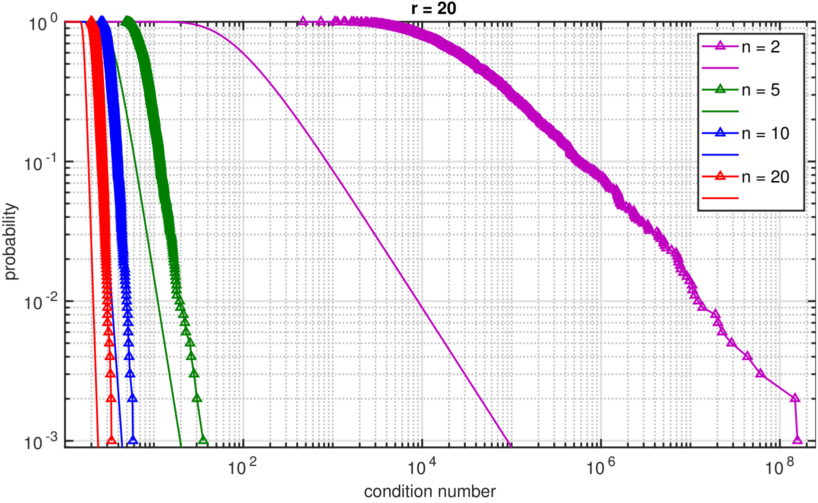

$1,000$ trials for $20 \times 20 \times n$ tensors.

Pencil-based algorithms

A **pencil-based algorithm** (PBA) is an algorithm that computes the CPD of

$$\mathcal{A} = \sum_{i=1}^{r} \mathbf{a}_i \otimes \mathbf{b}_i \otimes \mathbf{c}_i \in \mathcal{N}^* \subset \mathbb{R}^{n_1 \times n_2 \times n_d}$$

in a particular way, where $\mathcal{N}^* \subset \mathcal{N}_{r;\mathbf{n}}$ is some unspecified Zariski-open submanifold.[2]

[2] See Beltrán, Breiding, and V (2018) for the precise definition. The definition of a PBA is also more general than the one that I will present next.

Choose a fixed $Q \in \mathbb{R}^{n_3 \times 2}$ with orthonormal columns.

A PBA performs the following computations:

S1. $\mathcal{B} \leftarrow (I, I, Q^T) \cdot \mathcal{A}$;

S2. $\{\mathbf{a}_1, \ldots, \mathbf{a}_r\} \leftarrow \widehat{\theta}(\mathcal{B})$;

S3.a Choose an order $A := (\mathbf{a}_1, \ldots, \mathbf{a}_r)$;

S3.b $(\mathbf{b}_1 \otimes \mathbf{c}_1, \ldots, \mathbf{b}_r \otimes \mathbf{c}_r) \leftarrow (A^\dagger \mathcal{A}_{(1)})^T$;

S4. $\mathtt{output} \leftarrow \pi\big(\odot\big((\mathbf{a}_1, \ldots, \mathbf{a}_r), (\mathbf{b}_1 \otimes \mathbf{c}_1, \ldots, \mathbf{b}_r \otimes \mathbf{c}_r)\big)\big)$.

Herein, $\pi : \mathcal{S}^{\times r} \to (\mathcal{S}^{\times r} / \mathfrak{S}_r)$ and $\odot$ is the Khatri–Rao product: $\odot(A, B) := (\mathbf{a}_i \otimes \mathbf{b}_i)_i$.

Choose a fixed $Q \in \mathbb{R}^{n_3 \times 2}$ with orthonormal columns.

A PBA performs the following computations:

S1. $\mathcal{B} \leftarrow (I, I, Q^T) \cdot \mathcal{A}$;

S2. $\{\mathbf{a}_1, \ldots, \mathbf{a}_r\} \leftarrow \widehat{\theta}(\mathcal{B})$;

S3.a Choose an order $A := (\mathbf{a}_1, \ldots, \mathbf{a}_r)$;

S3.b $(\mathbf{b}_1 \otimes \mathbf{c}_1, \ldots, \mathbf{b}_r \otimes \mathbf{c}_r) \leftarrow (A^\dagger \mathcal{A}_{(1)})^T$;

S4. $\mathtt{output} \leftarrow \pi\Big(\odot\big((\mathbf{a}_1, \ldots, \mathbf{a}_r), (\mathbf{b}_1 \otimes \mathbf{c}_1, \ldots, \mathbf{b}_r \otimes \mathbf{c}_r)\big)\Big)$.

Herein, $\pi : \mathcal{S}^{\times r} \to (\mathcal{S}^{\times r} / \mathfrak{S}_r)$ and $\odot$ is the Khatri–Rao product: $\odot(A, B) := (\mathbf{a}_i \otimes \mathbf{b}_i)_i$.

The magic map $\widehat{\theta}$ needs to recover the vectors from the first factor matrix **when restricted to** $\mathcal{N}_{r;n_1,n_2,2}$:

$$\begin{aligned}\widehat{\theta}|_{\mathcal{N}_{r;n_1,n_2,2}} : \qquad \mathcal{N}_{r;n_1,n_2,2} &\longrightarrow (\mathbb{S}^{n_1-1})^{\times r}/\mathfrak{S}_r\\ \mathcal{B} = \sum_{i=1}^{r} \mathbf{a}_i \otimes \mathbf{b}_i \otimes \mathbf{z}_i &\longmapsto \{\mathbf{a}_1, \ldots, \mathbf{a}_r\}\end{aligned}$$

Since the input to $\widehat{\theta}$ will be the result of a previous numerical computation, the domain of definition of $\widehat{\theta}$ should also encompass a sufficiently large neighborhood of $\mathcal{N}_{r;n_1,n_2,2}$!

For proving instability, it does not matter what $\widehat{\theta}$ computes outside of $\mathcal{N}_{r;n_1,n_2,2}$.

For a valid input $\mathcal{A} \in \mathcal{N}^*$, let $\{\widetilde{\mathcal{A}}_1, \ldots, \widetilde{\mathcal{A}}_r\}$ be the CPD (in floating-point representation) returned by the PBA.

Our proof strategy consists of showing that for every $\epsilon > 0$ there exists an open neighborhood $\mathcal{O}_\epsilon \subset \mathcal{N}^*$ of r-nice tensors such that the **excess factor**

$$\omega(\mathcal{A}) = \frac{\text{observed forward error due to algorithm}}{\text{maximum forward error due to problem}}$$

$$:= \frac{\min_{\pi \in \mathfrak{S}_r} \sqrt{\sum_{i=1}^{r} \|\mathcal{A}_i - \widetilde{\mathcal{A}}_i\|^2}}{\kappa[\tau_{r;n_1,n_2,n_3}](\mathcal{A}) \cdot \|\mathcal{A} - \mathrm{fl}(\mathcal{A})\|_F}$$

behaves like a constant times ϵ^{-1}.

The maximum forward error of the problem is governed by the condition number of $\mathcal{A}$.

In an algorithm the error **can accumulate** in successive steps.

A PBA performs the following computations:

OK $\mathcal{B} \leftarrow (I, I, Q^T) \cdot \mathcal{A}$;

BAD $\{\mathbf{a}_1, \ldots, \mathbf{a}_r\} \leftarrow \widehat{\theta}(\mathcal{B})$;

OK Choose an order $A := (\mathbf{a}_1, \ldots, \mathbf{a}_r)$;

OK $(\mathbf{b}_1 \otimes \mathbf{c}_1, \ldots, \mathbf{b}_r \otimes \mathbf{c}_r) \leftarrow (A^\dagger \mathcal{A}_{(1)})^T$;

OK $\texttt{output} \leftarrow \pi\big(\odot\big((\mathbf{a}_1, \ldots, \mathbf{a}_r), (\mathbf{b}_1 \otimes \mathbf{c}_1, \ldots, \mathbf{b}_r \otimes \mathbf{c}_r)\big)\big)$.

The main intuition underpinning our proof is the fact that the condition number of $\widehat{\theta}$ can be very large even when the tensor decomposition problem has a small condition number.

For brevity, let's drop the rank r and dimensions $n_1 \times n_2 \times 2$ from the notation.

Consider the following diagram:

$$\begin{tikzcd}
\mathcal{N} \arrow[d, "\mathrm{Id}_{\mathcal{N}} \times \widehat{\theta}"'] \arrow[drr, "\tau"] & & \\
\mathcal{N} \times (\mathbb{S}^{n_1-1})^{\times r}/\mathfrak{S}_r \arrow[rr, "\widehat{\eta}"'] & & \widehat{\mathcal{M}}
\end{tikzcd}$$

Herein, $\widehat{\eta}$ is any map so that $\tau = \widehat{\eta} \circ (\mathrm{Id}_{\mathcal{N}} \times \widehat{\theta})$.

Since all involved domains and codomains are manifolds, we have

$$\begin{aligned}
\kappa[\tau](\mathcal{B}) = \|\mathrm{d}_{\mathcal{B}}\tau\|_2 &\le \|\mathrm{d}_{\mathcal{B}}\, \mathrm{Id}_{\mathcal{N}} \times \widehat{\theta}\|_2 \|\mathrm{d}_{(\mathcal{B},\widehat{\theta}(\mathcal{B}))}\widehat{\eta}\|_2 \\
&= \kappa[\mathrm{Id}_{\mathcal{N}} \times \widehat{\theta}](\mathcal{B}) \cdot \kappa[\widehat{\eta}](\mathcal{B}, \widehat{\theta}(\mathcal{B}))
\end{aligned}$$

Using foregoing idea, we are able to prove a lower on the condition number of $\widehat{\theta}$ at $\mathcal{B}$ in terms of $\kappa[\tau](\mathcal{B})$:

$$\kappa[\widehat{\theta}|_{\mathcal{N}}](\mathcal{B}) \geq \frac{\kappa[\tau](\mathcal{B})}{10r} - 1$$

But we know the right-hand side has a very bad distribution!

We then show that the neighborhood of the following orthogonally decomposable (**odeco**) tensor is problematic:

$$\mathcal{O} = \sum_{i=1}^{r} \mathbf{a}_i' \otimes \mathbf{b}_i' \otimes \mathbf{c}_i',$$

where $\mathbf{a}_i'$ (resp $\mathbf{b}_i'$) is an arbitrary orthonormal set of vectors and

$$C' = \frac{2}{n_3} \begin{bmatrix} Q^\perp & Q \end{bmatrix} \begin{bmatrix} \frac{n_3}{2} - 1 & 1 & 1 & \\ -1 & 1 - \frac{n_3}{2} & 1 & \\ -1 & 1 & 1 - \frac{n_3}{2} & \\ -1 & 1 & 1 & \cdots \\ \vdots & \vdots & \vdots & \\ -1 & 1 & 1 & \end{bmatrix}$$

where $[Q^\perp \quad Q]$ is an orthogonal matrix.

Odeco tensors like $\mathcal{O}$ have the **lowest sensitivity** to perturbations; their condition number is always 1.

It is a very bad omen that $\mathcal{O}$ is not a valid input for PBAs! Indeed, the projected tensor is

$$(I, I, Q^T) \cdot \mathcal{O} = -\frac{2}{n_3}\mathbf{a}_1' \otimes \mathbf{b}_1' \otimes \begin{bmatrix}1\\1\end{bmatrix} + \frac{2}{n_3}\sum_{i=2}^{r} \mathbf{a}_i' \otimes \mathbf{b}_i' \otimes \begin{bmatrix}1\\1\end{bmatrix},$$

which has a **positive-dimensional family** of decompositions, resulting in a non-unique first factor matrix. This is causes $\widehat{\theta}$ to have condition number ∞.

With some effort it can be shown that the remaining steps of the PBA cannot sufficiently reduce the sustained error.

Formally, we showed the following result:

Theorem (Beltrán, Breiding, and V (2018), Theorem 6.1)

There exist a constant $k > 0$ and a tensor $\mathcal{O} \in \mathcal{N}_{r;n_1,n_2,n_3}$ with the following properties: For all sufficiently small $\epsilon > 0$, there exists an open neighborhood $\mathcal{O}_\epsilon$ of $\mathcal{O}$, such that for all tensors $\mathcal{A} \in \mathcal{O}_\epsilon$ we have

1. *$\mathcal{A} \in \mathcal{N}^*$ is a valid input for a PBA, and*
2. *$\omega(\mathcal{A}) \geq k\epsilon^{-1}$.*

In other words, the forward error produced by a PBA can be larger than the maximum forward error expected from the tensor decomposition problem by an arbitrarily large factor.

The instability of the algorithm leads to an excess factor ω on top of the condition number of the computational problem:

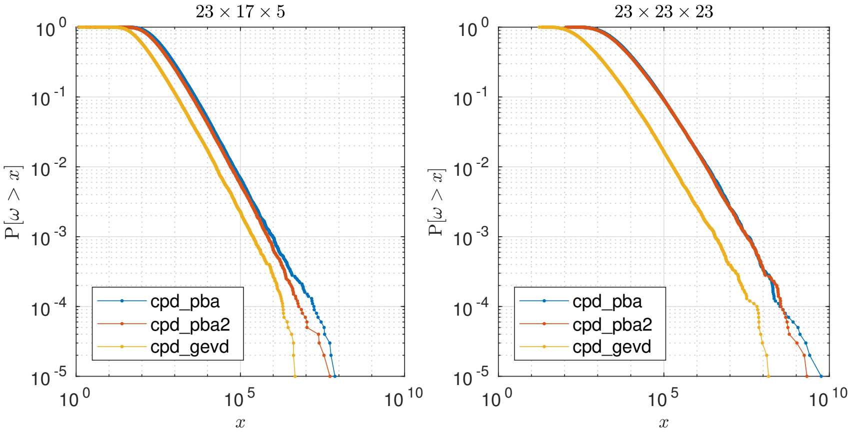

Test with random rank-1 tuples.

The excess factor ω in the neighborhood of the bad odeco tensor from our proof behaves exactly as the theory predicts:

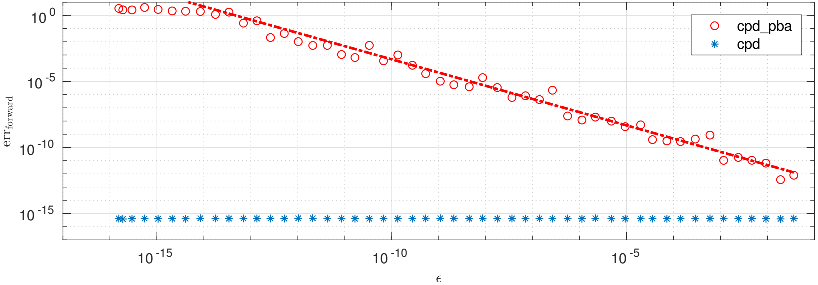

Test with a $85 \times 11 \times 29$ tensor.

Overview

Conclusions

Take-away story:

1. Tensors are conjectured to be generically r-identifiable for all strictly subgeneric r.
2. The condition number of the CPD measures the stability of the unique rank-1 tensors.
3. Reduction to a matrix pencil yields numerically unstable algorithms for computing CPDs.

Further reading

- Beltrán, Breiding, and V, *Pencil-based algorithms for tensor rank decomposition are not stable*, arXiv:1807.04159, 2018.
- Beltrán, Breiding, and V, *On the average condition number of the tensor rank decomposition*, In preparation.
- Breiding and V, *The condition number of join decompositions*, SIMAX, 2018.
- Breiding and V, *On the average condition number of tensor rank decompositions*, arXiv:1801.01673, 2018.

Grazie per l'attenzione!

References

Introduction

- Hitchcock, *The expression of a tensor or a polyadic as a sum of products*, J. Math. Phys., 1927.
- Kruskal, *Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics*, Lin. Alg. Appl., 1977.

Generic identifiability

- Angelini, *On complex and real identifiability of tensors*, Rev. Mat. Uni. Parma, 2017.
- Angelini, Bocci, Chiantini, *Real identifiability vs. complex identifiability*, Linear Multilinear Algebra, 2017.
- Ballico, *On the weak non-defectivity of veronese embeddings of projective spaces* , Centr. Eur. J. Math., 2005.
- Bocci, Chiantini, and Ottaviani, *Refined methods for the identifiability of tensors*, Ann. Mat. Pura Appl., 2013.
- Chiantini and Ottaviani, *On generic identifiability of 3-tensors of small rank*, SIAM J. Matrix Anal. Appl., 2013.
- Chiantini, Ottaviani, and Vannieuwenhoven, *An algorithm for generic and low-rank specific identifiability of complex tensors*, SIMAX, 2014.
- Chiantini, Ottaviani, and Vannieuwenhoven, *On generic identifiability of symmetric tensors of subgeneric rank*, Trans. Amer. Math. Soc., 2017.
- Chiantini, Ottaviani, and Vannieuwenhoven, *Effective criteria for specific identifiability of tensors and forms*, SIAM J. Matrix Anal. Appl., 2017.
- Galuppi and Mella, *Identifiability of homogeneous polynomials and Cremona Transformations*, J. Reine Angew. Math., 2017.
- Qi, Comon, and Lim, *Semialgebraic geometry of nonnegative tensor rank*, SIMAX, 2016.

Sensitivity

- Blum, Cucker, Shub, and Smale, *Complexity and Real Computation*, 1998.
- Beltrán, Breiding, and Vannieuwenhoven, *Pencil-based algorithms for tensor rank decomposition are not stable*, arXiv:1807.04159, 2018.
- Breiding and Vannieuwenhoven, *The condition number of join decompositions*, SIAM J. Matrix Anal. Appl., 2018.
- Bürgisser and Cucker, *Condition: The Geometry of Numerical Algorithms*, Springer, 2013.
- Lee, *Introduction to Smooth Manifolds*, 2013.
- Rice, *A theory of condition*, SIAM J. Numer. Anal., 1966.
- Vannieuwenhoven, *A condition number for the tensor rank decomposition*, Linear Algebra Appl., 2017.