A Multi-resolution Low-rank Tensor Decomposition

Sergio Rozada
King Juan Carlos University - Madrid, Spain

Joint work with:
A. G. Marques

icassp 2022 Singapore

2022 IEEE International Conference on Acoustics, Speech and Signal Processing - Singapore - May 22 - 27, 2022

- The world is becoming increasingly connected
- Every device is a factory of information
  - ⇒ Smartphones, sensors, etc.
  - ⇒ Enormous datasets - Big Data

- The world is becoming increasingly connected
- Every device is a factory of information
  - ⇒ Smartphones, sensors, etc.
  - ⇒ Enormous datasets - Big Data

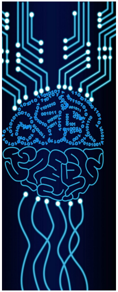

- Parsimony is key to enhance efficiency in multi-dimensional domains (N_s^D)
  - ⇒ E.g. chemometrics, and psychometrics [Bro97]
- Matrix models have been traditionally used to model these datasets
- Tensor models are breaking through [Papalexakis16]

Tensor Decomposition

Universidad Rey Juan Carlos

- Tensor decomposition tries to estimate a set of latent factors that summarize the tensor [Kolda2009]

Tensor Decomposition

- Tensor decomposition tries to estimate a set of latent factors that summarize the tensor [Kolda2009]
- The PARAFAC decomposition generalizes the SVD $\mathbf{X} = \sum_{r=1}^{R} \sigma_r \mathbf{u}_r \mathbf{v}_r^T$
  - ⇒ Consider the Ith order tensor $\underline{\mathbf{X}}$
  - ⇒ Consider the matrices $\mathbf{F}_i \in \mathbb{R}^{N_i \times R}$ for $i = 1, ..., I$
  - ⇒ Then, $\underline{\mathbf{X}}$ is said to have rank R if it can be written as

$$\underline{\mathbf{X}} = \sum_{r=1}^{R} [\mathbf{F}_1]_r \circledcirc [\mathbf{F}_2]_r \circledcirc ... \circledcirc [\mathbf{F}_I]_r \quad (1)$$

  - ⇒ The matrices $\mathbf{F}_i$ are obtained via ALS schemes

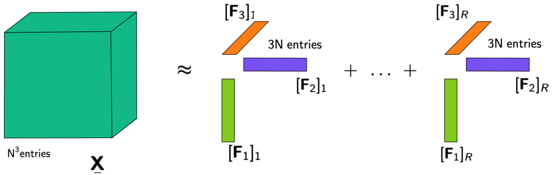

Vectorization and matrization I

- Tensors can be unfolded into lower order structures

Vectorization and matrization I

- Tensors can be unfolded into lower order structures
- Unfoldings are key to handle tensors
  $\Rightarrow$ There are different types of unfoldings

Universidad Rey Juan Carlos

- Tensors can be unfolded into lower order structures
- Unfoldings are key to handle tensors
  - ⇒ There are different types of unfoldings
- Tensors can be unfolded into vectors (vectorization)

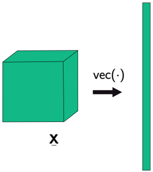

Vectorization and matrization I

- Tensors can be unfolded into lower order structures
- Unfoldings are key to handle tensors
  - ⇒ There are different types of unfoldings
- Tensors can be unfolded into vectors (vectorization)

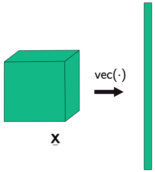

- Parsimony can be leveraged via tensor decomposition
  - ⇒ The Kronecker product factorizes a vectorization

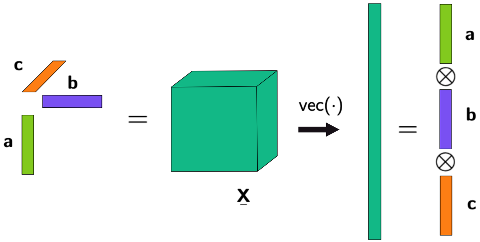

Vectorization and matrization II

- Tensor to matrix unfolding (matrization) key in tensor decomposition

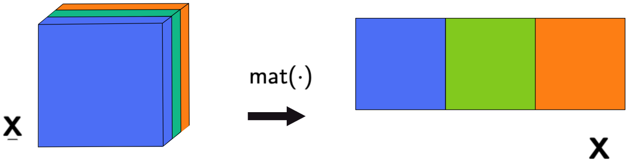

Vectorization and matrization II

- Tensor to matrix unfolding (matrization) key in tensor decomposition

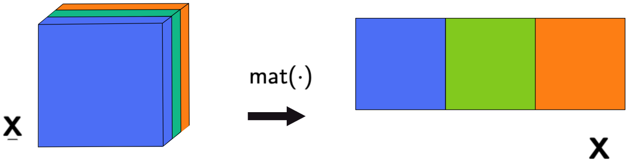

- Dimensions are rearranged in matrization
  - ⇒ One dimension remains fixed
  - ⇒ The other are grouped into one dimension

Vectorization and matrization II

- Tensor to matrix unfolding (matrization) key in tensor decomposition

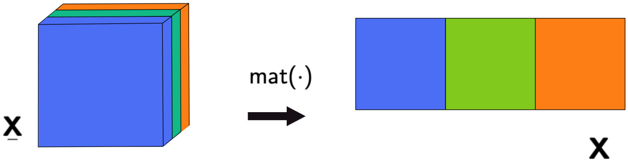

- Dimensions are rearranged in matrization
  - ⇒ One dimension remains fixed
  - ⇒ The other are grouped into one dimension
- Parsimony can be leveraged via low rank
  - ⇒ Low-rank tensor with rank R implies matrix with rank R
  - ⇒ The opposite is not necessarily true but...
  - ⇒ provides alternative parsimonious description (dimension dependent)

- Matrization and vectorization can be generalized

Tensorization and partitions

- Matrization and vectorization can be generalized
- Tensors can be unfolded into arbitrary low-order tensors
  - ⇒ Matrization is an special case, the lower-order tensor is a matrix
  - ⇒ Vectorization is another, the lower-order tensor is a vector

Tensorization and partitions

Universidad Rey Juan Carlos

- Matrization and vectorization can be generalized
- Tensors can be unfolded into arbitrary low-order tensors
  - $\Rightarrow$ Matrization is an special case, the lower-order tensor is a matrix
  - $\Rightarrow$ Vectorization is another, the lower-order tensor is a vector
- Partitions define how to unfold a tensor to a lower-order tensor:
  - $\Rightarrow$ Partitions tells how to group dimensions

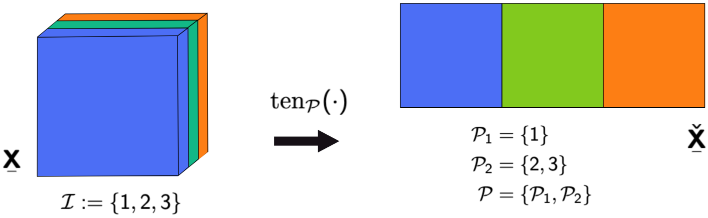

Formalizing partitions

- Consider a tensor $\underline{\mathbf{X}}$ of order I:
  - $\Rightarrow$ Let $\mathcal{I} := \{1, 2, ..., I\}$ denote the set containing all indexes

Definition: Partition

The ordered set $\mathcal{P} = \{\mathcal{P}_1, ..., \mathcal{P}_P\}$ is a partition of the set $\mathcal{I}$ if it holds that: $\mathcal{P}_p \neq \emptyset$ for all p, $\mathcal{P}_p \cap \mathcal{P}_{p'} = \emptyset$ for all $p' \neq p$, and $\bigcup_{p=1}^{P} \mathcal{P}_p = \mathcal{I}$.

- Consider a tensor $\underline{\mathbf{X}}$ of order I:
  - ⇒ Let $\mathcal{I} := \{1, 2, ..., I\}$ denote the set containing all indexes

Definition: Partition

The ordered set $\mathcal{P} = \{\mathcal{P}_1, ..., \mathcal{P}_P\}$ is a partition of the set $\mathcal{I}$ if it holds that: $\mathcal{P}_p \neq \emptyset$ for all p, $\mathcal{P}_p \cap \mathcal{P}_{p'} = \emptyset$ for all $p' \neq p$, and $\bigcup_{p=1}^{P} \mathcal{P}_p = \mathcal{I}$.

- Partitions are a mechanism to unfold tensors
  - ⇒ They can be used to decompose a tensor $\underline{\mathbf{X}}$ as the sum:

$$\underline{\mathbf{X}} = \sum_{l=1}^{L} \underline{\mathbf{Z}}_l \tag{2}$$

  - ⇒ Where each $\underline{\mathbf{Z}}_l$ is a low-order unfolding defined by partition $\mathcal{P}^{(l)}$

Contribution I

- A tensor can be unfolded into various lower-order tensors
  - $\Rightarrow$ Low rank can be imposed in each low-order tensor
  - $\Rightarrow$ These unfoldings can be mixed

Contribution I

- A tensor can be unfolded into various lower-order tensors
  - ⇒ Low rank can be imposed in each low-order tensor
  - ⇒ These unfoldings can be mixed
- We propose a Multi-Resolution Low Rank (MRLR) decomposition:

Definition: MRLR decomposition

1. Consider a *collection* of tensor lower-order representations
2. Postulate a low-rank decomposition for each representation
3. Map each representation back to the original tensor domain
4. Model the tensor as the sum of the low-rank representations

Contribution I

- A tensor can be unfolded into various lower-order tensors
  - ⇒ Low rank can be imposed in each low-order tensor
  - ⇒ These unfoldings can be mixed
- We propose a Multi-Resolution Low Rank (MRLR) decomposition:

Definition: MRLR decomposition

1. Consider a *collection* of tensor lower-order representations
2. Postulate a low-rank decomposition for each representation
3. Map each representation back to the original tensor domain
4. Model the tensor as the sum of the low-rank representations

- The MRLR decomposition
  - ⇒ exploits information at different resolutions
  - ⇒ leverages parsimony

Contribution II

- Formally, consider a collection of partitions $\mathcal{P}^{(1)},\ldots,\mathcal{P}^{(L)}$
  - ⇒ with $|\mathcal{P}^{(l)}| \leq |\mathcal{P}^{(l')}|$ for $l < l'$
  - ⇒ Given the Ith order tensor $\underline{\mathbf{X}}$, we propose

$$\underline{\mathbf{X}} = \sum_{l=1}^{L} \underline{\mathbf{Z}}_l, \quad \text{with} \quad \text{rank}\big(\text{ten}_{\mathcal{P}^{(l)}}(\underline{\mathbf{Z}}_l)\big) \leq R_l, \tag{3}$$

- Formally, consider a collection of partitions $\mathcal{P}^{(1)},\ldots,\mathcal{P}^{(L)}$
  - $\Rightarrow$ with $|\mathcal{P}^{(l)}| \leq |\mathcal{P}^{(l')}|$ for $l < l'$
  - $\Rightarrow$ Given the Ith order tensor $\underline{\mathbf{X}}$, we propose

$$\underline{\mathbf{X}} = \sum_{l=1}^{L} \underline{\mathbf{Z}}_l, \quad \text{with} \quad \text{rank}\big(\text{ten}_{\mathcal{P}^{(l)}}(\underline{\mathbf{Z}}_l)\big) \leq R_l, \tag{3}$$

- The tensorization operator $\text{ten}_{\mathcal{P}}(\underline{\mathbf{Z}})$ yields the $|\mathcal{P}|$th order tensor $\check{\underline{\mathbf{Z}}}$

$$\begin{aligned}
\check{\underline{\mathbf{Z}}} &= \text{ten}_{\mathcal{P}}(\underline{\mathbf{Z}}) \in \mathbb{R}^{\prod_{j=1}^{|\mathcal{P}_1|} |\mathcal{P}_1(j)| \times \ldots \times \prod_{j=1}^{|\mathcal{P}_P|} |\mathcal{P}_P(j)|} \\
[\check{\underline{\mathbf{Z}}}]_{k_1,\ldots,k_{|\mathcal{P}|}} &= [\underline{\mathbf{Z}}]_{n_1,\ldots,n_I} \quad \text{and} \\
k_p &= n_{\mathcal{P}_p(1)} \quad \text{if} \quad |\mathcal{P}_p| = 1 \\
k_p &= n_{\mathcal{P}_p(1)} + \sum_{i=2}^{|\mathcal{P}_p|} (n_{\mathcal{P}_p(i)} - 1) \prod_{j=1}^{i-1} N_{\mathcal{P}_p(j)} \quad \text{if} \quad |\mathcal{P}_p| > 1
\end{aligned} \tag{4}$$

MRLR decomposition

- Given a three-dimensional tensor $\underline{\mathbf{X}}$:
  - $\Rightarrow$ The partition $\mathcal{P}^{(1)} = \{\{1\},\{2,3\}\}$ defines the matrix $\check{\underline{\mathbf{X}}}_1$
  - $\Rightarrow$ The partition $\mathcal{P}^{(2)} = \{\{1,2,3\}\}$ defines the auto-map $\check{\underline{\mathbf{X}}}_2$
  - $\Rightarrow$ The low-rank decompositions $\check{\underline{\mathbf{Z}}}_1$ and $\check{\underline{\mathbf{Z}}}_2$ are estimated
  - $\Rightarrow$ $\check{\underline{\mathbf{Z}}}_1$ and $\check{\underline{\mathbf{Z}}}_2$ are mapped back, giving $\underline{\mathbf{Z}}_1$ and $\underline{\mathbf{Z}}_2$
  - $\Rightarrow$ Model the tensor as the sum of $\underline{\mathbf{Z}}_1$ and $\underline{\mathbf{Z}}_2$

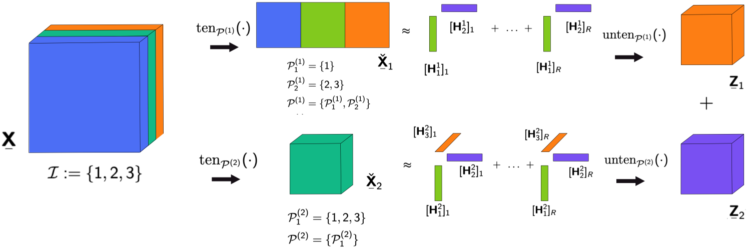

- The MRLR decomposition can be obtained via solving:

$$\min_{\underline{\mathbf{Z}}_1 \dots \underline{\mathbf{Z}}_L} \left\| \underline{\mathbf{X}} - \sum_{l=1}^{L} \underline{\mathbf{Z}}_l \right\|_F \quad (5)$$

$$\text{s. t. } \operatorname{rank}\big(\operatorname{ten}_{\mathcal{P}^{(l)}}(\underline{\mathbf{Z}}_l)\big) \leq R_l.$$

- The MRLR decomposition can be obtained via solving:

$$\min_{\underline{\mathbf{Z}}_1 \dots \underline{\mathbf{Z}}_L} \left\| \underline{\mathbf{X}} - \sum_{l=1}^{L} \underline{\mathbf{Z}}_l \right\|_F \qquad (5)$$

$$\text{s. t. } \operatorname{rank}\big(\operatorname{ten}_{\mathcal{P}^{(l)}}(\underline{\mathbf{Z}}_l)\big) \leq R_l.$$

- The expression in (5) can be worked out
  - ⇒ Sequentially solving for each $\underline{\mathbf{Z}}_l$
  - ⇒ Leveraging low rank in each $\underline{\mathbf{Z}}_l$ via the PARAFAC decomposition

$$\min_{\mathbf{H}_1^i, \dots, \mathbf{H}_{J_i}^i} \left\| \underline{\mathbf{X}} - \sum_{l \neq i}^{L} \underline{\mathbf{Z}}_l - \sum_{j=1}^{R_i} [\mathbf{H}_1^i]_j \circledcirc \dots \circledcirc [\mathbf{H}_{J_i}^i]_j \right\|_F . \qquad (6)$$

Numerical studies

- MRLR-dec has been tested in three scenarios (left to right):
  - $\Rightarrow$ The aminoacids dataset, a three-mode tensor $(5 \times 201 \times 61) \rightarrow I = 3$
  - $\Rightarrow$ An RGB video signal of 173 frames $(3 \times 173 \times 1080 \times 720) \rightarrow I = 4$
  - $\Rightarrow$ A discretized three-dimensional function $f(x_1, x_2, x_3) = \frac{x_1^2 + x_2^2}{e^{|x_2 + x_3|}} \rightarrow I = 3$

- MRLR-dec has been tested in three scenarios (left to right):
  - $\Rightarrow$ The aminoacids dataset, a three-mode tensor $(5 \times 201 \times 61) \rightarrow I = 3$
  - $\Rightarrow$ An RGB video signal of 173 frames $(3 \times 173 \times 1080 \times 720) \rightarrow I = 4$
  - $\Rightarrow$ A discretized three-dimensional function $f(x_1, x_2, x_3) = \frac{x_1^2 + x_2^2}{e^{|x_2 + x_3|}} \rightarrow I = 3$
- Performance measured in terms of Normalized Frobenius Error

$$\mathrm{NFE} = ||\underline{\mathbf{X}} - \check{\underline{\mathbf{X}}}||_F / ||\underline{\mathbf{X}}||_F. \tag{7}$$

Numerical studies

- MRLR-dec has been tested in three scenarios (left to right):
  - ⇒ The aminoacids dataset, a three-mode tensor ($5 \times 201 \times 61$) $\rightarrow I = 3$
  - ⇒ An RGB video signal of 173 frames ($3 \times 173 \times 1080 \times 720$) $\rightarrow I = 4$
  - ⇒ A discretized three-dimensional function $f(x_1, x_2, x_3) = \frac{x_1^2 + x_2^2}{e^{|x_2 + x_3|}} \rightarrow I = 3$
- Performance measured in terms of Normalized Frobenius Error

$$\mathrm{NFE} = ||\underline{\mathbf{X}} - \check{\underline{\mathbf{X}}}||_F / ||\underline{\mathbf{X}}||_F. \quad (7)$$

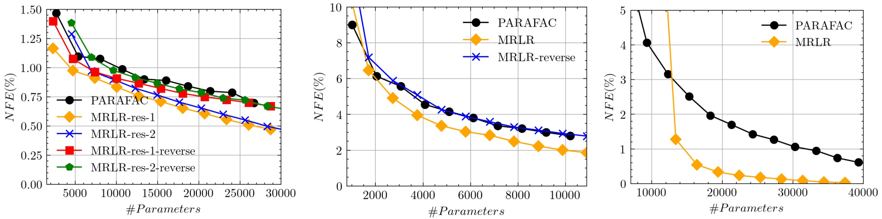

Closing remarks

- Tensor decompositions are parsimonious high-dimensional models

Closing remarks

Universidad Rey Juan Carlos

- Tensor decompositions are parsimonious high-dimensional models
- A new (MRLR) tensor decomposition has been proposed
  - ⇒ to leverage structure at different resolution levels

Closing remarks

- Tensor decompositions are parsimonious high-dimensional models
- A new (MRLR) tensor decomposition has been proposed
  - ⇒ to leverage structure at different resolution levels
- The MRLR decomposition can be algorithmic implemented
  - ⇒ obtaining each lower-order representation sequentially
  - ⇒ imposing low-rank in each lower-order representation

Closing remarks

Universidad Rey Juan Carlos

- Tensor decompositions are parsimonious high-dimensional models
- A new (MRLR) tensor decomposition has been proposed
  - ⇒ to leverage structure at different resolution levels
- The MRLR decomposition can be algorithmic implemented
  - ⇒ obtaining each lower-order representation sequentially
  - ⇒ imposing low-rank in each lower-order representation
- Practical cases where MRLR outperforms PARAFAC

Closing remarks

Universidad Rey Juan Carlos

- Tensor decompositions are parsimonious high-dimensional models
- A new (MRLR) tensor decomposition has been proposed
  - ⇒ to leverage structure at different resolution levels
- The MRLR decomposition can be algorithmic implemented
  - ⇒ obtaining each lower-order representation sequentially
  - ⇒ imposing low-rank in each lower-order representation
- Practical cases where MRLR outperforms PARAFAC
- Future research directions
  - ⇒ Identifiability/ambiguity results
  - ⇒ Intelligent ways of creating the partitions
  - ⇒ Alternative algorithms