



Model reduction, data-based and advanced discretization in computational mechanics

Algebraic PGD for tensor separation and compression: An algorithmic approach



Pedro Díez*, Sergio Zlotnik, Alberto García-González, Antonio Huerta

Laboratori de Càlcul Numèric, E.T.S. de Ingenieros de Caminos, Canales y Puertos, Universitat Politècnica de Catalunya – (UPC BarcelonaTech),
Jordi Girona 1, 08034 Barcelona, Spain

ARTICLE INFO

Article history:

Received 7 December 2017
Accepted after revision 2 March 2018
Available online 7 May 2018

Keywords:

Tensor separation
Algebraic PGD
Least-squares approximation

ABSTRACT

Proper Generalized Decomposition (PGD) is devised as a computational method to solve high-dimensional boundary value problems (where many dimensions are associated with the space of parameters defining the problem). The PGD philosophy consists in providing a separated representation of the multidimensional solution using a greedy approach combined with an alternated directions scheme to obtain the successive rank-one terms. This paper presents an algorithmic approach to high-dimensional tensor separation based on solving the Least Squares approximation in a separable format of multidimensional tensor using PGD. This strategy is usually embedded in a standard PGD code in order to compress the solution (reduce the number of terms and optimize the available storage capacity), but it stands also as an alternative and highly competitive method for tensor separation.

© 2018 Académie des sciences. Published by Elsevier Masson SAS. All rights reserved.

1. Introduction

1.1. Framework and motivation

Data is often collected in terms of multidimensional arrays. The number of dimensions is denoted by d , and the object containing the information is a tensor \mathbf{F} of order d . A multi-index notation identifies each entry of the array as corresponding to specific values of d parameters. The size of the tensor in each dimension is denoted by n_i , for $i = 1, 2, \dots, d$, and \mathbf{F} is characterized by the generic term $[\mathbf{F}]_{j_1 j_2 \dots j_d}$ for $j_i = 1, 2, \dots, n_i$.

Tensor separation is a generalization of matrix diagonalization and allows representing the data collected in the tensor in a compact form. The separated form is expressed as the sum of M terms, each of them consisting in the tensorial product of d vectors of dimensions n_i , $i = 1, 2, \dots, d$. Note that only in terms of storage, the number of entries of \mathbf{F} is $\prod_{i=1}^d n_i$, and in the separable version it is described with $M \sum_{i=1}^d n_i$ scalar quantities. For *small* values of M , this represents a huge reduction of the information to be stored.

In the following, the methodology to perform this separation is presented first in the simple case of $d = 2$, where it can be done optimally. Then, the generalization to higher dimensions, in which there is no optimal strategy, is devised using the PGD philosophy.

* Corresponding author.

E-mail addresses: pedro.diez@upc.edu (P. Díez), sergio.zlotnik@upc.edu (S. Zlotnik), berto.garcia@upc.edu (A. García-González), antonio.huerta@upc.edu (A. Huerta).

1.2. SVD and matrix separation

For $d = 2$, \mathbf{F} is a matrix, and the Singular Value Decomposition (SVD) (and matrix diagonalization as a particular case for squared matrices) is the standard and optimal tool to obtain a reduced representation of a 2D array (a matrix or second-order tensor). The outcome of the SVD allows representing the matrix as a sum of tensor products (rank-one matrices) of the left and right eigenvectors, each weighted by the corresponding eigenvalue.

Namely, the SVD of $\mathbf{F} \in \mathbb{R}^{n_1 \times n_2}$ consists in finding square unit matrices $\mathbf{U} \in \mathbb{R}^{n_1 \times n_1}$ and $\mathbf{V} \in \mathbb{R}^{n_2 \times n_2}$ such that

$$\mathbf{F} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T \tag{1}$$

being $\mathbf{\Sigma}$ a diagonal matrix in $\mathbb{R}^{n_1 \times n_2}$ containing the singular values of \mathbf{F} . That is, \mathbf{U} and \mathbf{V} are such that $\mathbf{U}^T \mathbf{U} = \mathbb{I}_{n_1}$ and $\mathbf{V}^T \mathbf{V} = \mathbb{I}_{n_2}$, and $\mathbf{\Sigma}$ has the format

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_1 & & & & \\ & \sigma_2 & & & \\ & & \ddots & & \\ & & & \sigma_{n_2} & \\ 0 & 0 & 0 & 0 & \\ \vdots & \vdots & \vdots & \vdots & \\ 0 & 0 & 0 & 0 & \end{bmatrix} \text{ for } n_1 > n_2 \text{ or } \mathbf{\Sigma} = \begin{bmatrix} \sigma_1 & & 0 & \dots & 0 \\ & \sigma_2 & & & 0 \\ & & \ddots & & 0 \\ & & & \sigma_{n_1} & 0 \\ & & & & 0 \end{bmatrix} \text{ for } n_1 < n_2$$

It is assumed that the singular values are sorted in decreasing order, that is,

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(n_1, n_2)} \geq 0$$

The column vectors of matrices \mathbf{U} and \mathbf{V} are denoted by \mathbf{u}_j and \mathbf{v}_k , for $j = 1, 2, \dots, n_1$ and $k = 1, 2, \dots, n_2$, that is such that

$$\mathbf{U} = [\mathbf{u}_1 \ \mathbf{u}_2 \ \dots \ \mathbf{u}_{n_1}] \text{ and } \mathbf{V} = [\mathbf{v}_1 \ \mathbf{v}_2 \ \dots \ \mathbf{v}_{n_2}]$$

Thus, the SVD in (1) is rewritten as

$$\mathbf{F} = \sum_{j=1}^{\min(n_1, n_2)} \sigma_j \mathbf{u}_j \mathbf{v}_j^T = \sum_{j=1}^{\min(n_1, n_2)} \sigma_j \mathbf{u}_j \otimes \mathbf{v}_j \tag{2}$$

Note that the two notations in the equation above are equivalent because each term $\mathbf{u}_j \mathbf{v}_j^T$ is a rank-one $n_1 \times n_2$ matrix that can also be denoted as $\mathbf{u}_j \otimes \mathbf{v}_j$ using the tensor (or external) product. The latter notation is preferred in the following because it is easily extended to problems of dimension higher than two.

The representation displayed in (2) is particularly attractive to obtain a Least-Squares (LS) low-rank approximation of \mathbf{F} . The Frobenius matrix norm reads

$$\|\mathbf{F}\|^2 = \mathbf{F} : \mathbf{F} = [\mathbf{F}]_{j_1 j_2} [\mathbf{F}]_{j_1 j_2} \tag{3}$$

Using this *Euclidean-type* norm, the truncation of (2) to M terms (being $M < \min(n_1, n_2)$) results in the best LS rank- M approximation, that is,

$$\sum_{j=1}^M \sigma_j \mathbf{u}_j \otimes \mathbf{v}_j = \arg \min_{\mathbf{A} \in \mathbb{R}^{n_1 \times n_2} \text{ of rank } M} \|\mathbf{F} - \mathbf{A}\| \tag{4}$$

This LS optimality is associated with the orthogonality of the left and right eigenvectors.

In particular, the best rank-one approximation of \mathbf{F} is $\sigma_1 \mathbf{u}_1 \otimes \mathbf{v}_1$. This is easily shown by noting that for any rank-one matrix $\mathbf{w}^1 \otimes \mathbf{w}^2$ where $\mathbf{w}^1 = \sum_{k=1}^{n_1} \alpha_k \mathbf{u}_k$ and $\mathbf{w}^2 = \sum_{\ell=1}^{n_2} \beta_\ell \mathbf{v}_\ell$, LS discrepancy with \mathbf{F} reads

$$\|\mathbf{F} - \mathbf{w}^1 \otimes \mathbf{w}^2\|^2 = \|\mathbf{\Sigma} - \boldsymbol{\alpha} \otimes \boldsymbol{\beta}\|^2 \tag{5}$$

being $\boldsymbol{\alpha} \in \mathbb{R}^{n_1}$ and $\boldsymbol{\beta} \in \mathbb{R}^{n_2}$ the vectors representing \mathbf{w}^1 and \mathbf{w}^2 in the bases of left and right eigenvectors ($\{\mathbf{u}_1, \dots, \mathbf{u}_{n_1}\}$ and $\{\mathbf{v}_1, \dots, \mathbf{v}_{n_2}\}$, respectively). It is clear from the right-hand side of (5) that an optimal choice for $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ is $\alpha_k = \beta_k = \delta_{k1} \sqrt{\sigma_k}$.

Remark 1. The equality (5) is a direct consequence of observing that

$$\|\mathbf{F} - \mathbf{w}^1 \otimes \mathbf{w}^2\|^2 = \|\mathbf{F}\|^2 + \|\mathbf{w}^1 \otimes \mathbf{w}^2\|^2 - 2\mathbf{F} : \mathbf{w}^1 \otimes \mathbf{w}^2$$

$$\|\mathbf{F}\|^2 = \sum_{k=1}^{\min(n_1, n_2)} \sigma_k^2 = \|\boldsymbol{\Sigma}\|^2, \quad \|\mathbf{w}^1 \otimes \mathbf{w}^2\|^2 = \sum_{k=1}^{n_1} \sum_{\ell=1}^{n_2} \alpha_k^2 \beta_\ell^2 = \|\boldsymbol{\alpha} \otimes \boldsymbol{\beta}\|^2$$

and

$$\mathbf{F} : \mathbf{w}^1 \otimes \mathbf{w}^2 = \sum_{k=1}^{n_1} \sum_{\ell=1}^{n_2} \alpha_k \beta_\ell \mathbf{F} : (\mathbf{u}_k \otimes \mathbf{v}_\ell) = \sum_{k=1}^{\min(n_1, n_2)} \alpha_k \beta_k \sigma_k = \boldsymbol{\Sigma} : \boldsymbol{\alpha} \otimes \boldsymbol{\beta}$$

These properties are due to the orthonormality of the bases $\{\mathbf{u}_j\}_{j=1, \dots, n_1}$ and $\{\mathbf{v}_j\}_{j=1, \dots, n_2}$, which has as a direct consequence

$$\mathbf{F} : (\mathbf{u}_k \otimes \mathbf{v}_\ell) = \left[\sum_j \sigma_j \mathbf{u}_j \otimes \mathbf{v}_j \right] : (\mathbf{u}_k \otimes \mathbf{v}_\ell) = \sum_j \sigma_j (\mathbf{u}_j \cdot \mathbf{u}_k) (\mathbf{v}_j \cdot \mathbf{v}_\ell) = \delta_{k\ell} \sigma_k$$

The generalization of these tools to high-order tensor formats is the object of extensive research activities, and different strategies are labeled as High-Order SVD (HOSVD). Nevertheless, for a number of dimensions larger than two, the orthogonality among the eigenvectors and therefore the optimality of the separation is no longer guaranteed.

1.3. Problem statement: high-order separation

For $d > 2$, tensor $\mathbf{F} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ is expressed in a separable format if, for some integer value M , and for $i = 1, 2, \dots, d$ and $m = 1, 2, \dots, M$, there exists a set of vectors $\tilde{\mathbf{f}}_i^m \in \mathbb{R}^{n_i}$ such that

$$\mathbf{F} = \sum_{m=1}^M \tilde{\mathbf{f}}_1^m \otimes \tilde{\mathbf{f}}_2^m \otimes \dots \otimes \tilde{\mathbf{f}}_d^m \tag{6}$$

An alternative version of this format uses unit vectors \mathbf{f}_i^m and amplitudes σ_m , for $i = 1, \dots, d$ and $m = 1, \dots, M$

$$\mathbf{F} = \sum_{m=1}^M \sigma_m \mathbf{f}_1^m \otimes \mathbf{f}_2^m \otimes \dots \otimes \mathbf{f}_d^m \tag{7}$$

Note that (6) and (7) are readily shown to be equivalent by taking

$$\mathbf{f}_i^m = \frac{1}{\|\tilde{\mathbf{f}}_i^m\|} \tilde{\mathbf{f}}_i^m \text{ for } i = 1, \dots, d \text{ and } \sigma_m = \prod_{i=1}^d \|\tilde{\mathbf{f}}_i^m\| \text{ for } m = 1, \dots, M$$

1.4. Standard approaches

The problem of finding a separable expression with the form of (6) or (7) that approximates a tensor \mathbf{F} has a wide range of applications in many engineering and scientific fields. The pioneering works appeared in the field of psychometrics, see for example the work of Tucker [1] and Harshman [2], and in the field of chemometrics [3,4]. Later, the interest of tensor decompositions reached many different scientific communities such as neuroscience [5,6], computer vision [7,8], signal processing [9,10], and data mining [11]. Numerical analysis is no exception, and examples of applications on tensor decompositions include [12,13]. This list is by no means exhaustive; for a comprehensive review on tensor decomposition methods, see [14].

2. PGD strategy and rank-one algorithm

The PGD strategy aims at obtaining an approximation of some tensor \mathbf{F} in the separable form given by (7). The idea is extensively described in [15–18] and references therein, and consists in combining a *greedy* algorithm (that is, computing first for $m = 1$ and compute sequentially for any m when the previous $m - 1$ terms are already obtained) with an *alternated directions* scheme to solve the series of rank-one problems corresponding to term m .

2.1. Greedy scheme and rank-one terms

The first step in the greedy algorithm is to find a rank-one approximation of \mathbf{F} , namely find d unit vectors $\mathbf{f}_1^1, \dots, \mathbf{f}_d^1$ and the corresponding amplitude σ_1 such that

$$\mathbf{F} \approx \sigma_1 \mathbf{f}_1^1 \otimes \mathbf{f}_2^1 \otimes \dots \otimes \mathbf{f}_d^1$$

or, equivalently,

$$\mathbf{F} \approx \tilde{\mathbf{f}}_1^1 \otimes \tilde{\mathbf{f}}_2^1 \otimes \cdots \otimes \tilde{\mathbf{f}}_d^1 \tag{8}$$

Also at this stage, the LS criterion is used to qualify the best rank-one approximation, following the ideas draft in the appendix of [18] describing the PGD compression strategy. That is, the d vectors $\tilde{\mathbf{f}}_1^1, \dots, \tilde{\mathbf{f}}_d^1$ are sought such that they minimize the scalar functional $\mathcal{J}(\cdot)$ defined by

$$\mathcal{J}(\tilde{\mathbf{f}}_1^1, \dots, \tilde{\mathbf{f}}_d^1) = \|\mathbf{F} - \tilde{\mathbf{f}}_1^1 \otimes \tilde{\mathbf{f}}_2^1 \otimes \cdots \otimes \tilde{\mathbf{f}}_d^1\| \tag{9}$$

Note that, in this context, analogously to (3) the Frobenius-type norm for tensors reads

$$\|\mathbf{F}\|^2 = [\mathbf{F}]_{j_1 j_2 \dots j_d} [\mathbf{F}]_{j_1 j_2 \dots j_d} = \sum_{j_1=1}^{n_1} \sum_{j_2=1}^{n_2} \cdots \sum_{j_d=1}^{n_d} [\mathbf{F}]_{j_1 j_2 \dots j_d}^2 \tag{10}$$

That is,

$$(\tilde{\mathbf{f}}_1^1, \dots, \tilde{\mathbf{f}}_d^1) = \arg \min_{\mathbb{R}^{n_1} \times \cdots \times \mathbb{R}^{n_d}} \mathcal{J}(\tilde{\mathbf{f}}_1^1, \dots, \tilde{\mathbf{f}}_d^1) \tag{11}$$

Functional $\mathcal{J}(\cdot)$ is nonlinear and therefore also the problem given in (11) is nonlinear and requires devising an iterative solver. Note also that the set of possible solutions (constituted by all the rank-one tensors) is not provided with the structure of a linear vectorial space. Obviously, the sum of two rank-one tensors is, in general, a tensor of rank two.

The algorithm proposed to solve problem (11) is detailed in section 2.2. Once the solution to (11) is available, this is taken as the first term of the PGD approximation, that is,

$$\mathbf{F}_{\text{PGD}}^1 = \tilde{\mathbf{f}}_1^1 \otimes \tilde{\mathbf{f}}_2^1 \otimes \cdots \otimes \tilde{\mathbf{f}}_d^1 \tag{12}$$

From this point on, the algorithm is recursive and obtains the best approximation of the remainder part of \mathbf{F} . Namely, assuming that $\mathbf{F}_{\text{PGD}}^{M-1}$ is available

$$\mathbf{F}_{\text{PGD}}^{M-1} = \sum_{m=1}^{M-1} \tilde{\mathbf{f}}_1^m \otimes \tilde{\mathbf{f}}_2^m \otimes \cdots \otimes \tilde{\mathbf{f}}_d^m$$

the next term is obtained as

$$\mathbf{F}_{\text{PGD}}^M = \mathbf{F}_{\text{PGD}}^{M-1} + \tilde{\mathbf{f}}_1^M \otimes \tilde{\mathbf{f}}_2^M \otimes \cdots \otimes \tilde{\mathbf{f}}_d^M$$

and

$$(\tilde{\mathbf{f}}_1^M, \dots, \tilde{\mathbf{f}}_d^M) = \arg \min_{\mathbb{R}^{n_1} \times \cdots \times \mathbb{R}^{n_d}} \|\mathbf{F} - \mathbf{F}_{\text{PGD}}^{M-1} - \tilde{\mathbf{f}}_1^M \otimes \tilde{\mathbf{f}}_2^M \otimes \cdots \otimes \tilde{\mathbf{f}}_d^M\| \tag{13}$$

Note that (11) and (13) have exactly the same structure, they are both rank-one least-squares approximation problems and therefore the same iterative algorithm devised for (11) is going to be used for (13).

The stopping criteria used to decide whether the number of PGD terms, M is sufficiently large are mainly based on characterizing the relative importance of the last term added to the sum. An alternative approach is computing the residual $\|\mathbf{F} - \mathbf{F}_{\text{PGD}}^M\|$, but this is often discarded because it requires reconstructing $\mathbf{F}_{\text{PGD}}^M$ as a multidimensional tensor and this has a large computational cost.

Recall that using normalized vectors, the expression for the PGD solution reads

$$\mathbf{F}_{\text{PGD}}^M = \sum_{m=1}^M \sigma_m \mathbf{f}_1^m \otimes \mathbf{f}_2^m \otimes \cdots \otimes \mathbf{f}_d^m \tag{14}$$

Typically, M is considered to be large enough if, for some tolerance η^* , the following inequality holds

$$\sigma_M < \eta^* \sigma_1$$

That is, the greedy algorithm is stopped when the amplitude σ_M of the last term is significantly lower than the amplitude of the first one.

2.2. Alternated directions scheme: iterating in sectional problems

This section is devoted to describing the iterative algorithm devised to solve the rank-one problem (11) (or (13)). Thus, the goal is to compute $(\tilde{\mathbf{f}}_1^1, \dots, \tilde{\mathbf{f}}_d^1)$ minimizing the functional $\mathcal{J}(\tilde{\mathbf{f}}_1^1, \dots, \tilde{\mathbf{f}}_d^1)$ given by (9). The idea is to follow an alternated directions strategy, consisting in computing each of the sectional unknowns, say $\tilde{\mathbf{f}}_\gamma^1$ for $\gamma = 1, 2, \dots, d$, assuming that all the others ($\tilde{\mathbf{f}}_j^1$ for $j \neq \gamma$) are known. This is done for $\gamma = 1, 2, \dots, d$ and then iterated until convergence is reached.

The functional $\mathcal{J}(\cdot)$ is rewritten as:

$$\mathcal{J}(\tilde{\mathbf{f}}_1^1, \dots, \tilde{\mathbf{f}}_d^1) = \|\mathbf{F}\|^2 + \prod_{j=1}^d \tilde{\mathbf{f}}_j^{1\top} \tilde{\mathbf{f}}_j^1 - 2\mathbf{F} : \tilde{\mathbf{f}}_1^1 \otimes \dots \otimes \tilde{\mathbf{f}}_d^1 \tag{15}$$

where the symbol $:$ must be understood here as total tensor contraction (summing up in all the indices), and corresponds to the expression of the norm described in (10).

Thus, in order to compute an approximation to $\tilde{\mathbf{f}}_\gamma^1$, it is assumed that the other modes, $\tilde{\mathbf{f}}_j^1$ for $j \neq \gamma$ are known and the functional $\mathcal{J}(\cdot)$ is to be minimized with respect to $\tilde{\mathbf{f}}_\gamma^1$. Namely

$$\mathcal{J}(\tilde{\mathbf{f}}_\gamma^1) = \|\mathbf{F}\|^2 + \alpha \tilde{\mathbf{f}}_\gamma^{1\top} \tilde{\mathbf{f}}_\gamma^1 - 2\tilde{\mathbf{f}}_\gamma^{1\top} \mathbf{g} \tag{16}$$

where the computable quantities α and \mathbf{g} are

$$\alpha := \left(\prod_{j \neq \gamma}^d \tilde{\mathbf{f}}_j^{1\top} \tilde{\mathbf{f}}_j^1 \right) \tag{17}$$

and

$$\mathbf{g} := \mathbf{F} \dots \bigotimes_{j \neq \gamma}^d \tilde{\mathbf{f}}_j^1 \tag{18}$$

where symbol \dots indicates tensor contraction of all possible indices. In this case, provided that \mathbf{F} is a tensor of d dimensions and $\bigotimes_{j \neq \gamma}^d \tilde{\mathbf{f}}_j^1$ is a tensor of $d - 1$ dimensions, this means summing up in all indices i_j for $j = 1, \dots, d$ with $j \neq \gamma$. That is, (18) is equivalent to

$$[\mathbf{g}]_{i_\gamma} := \sum_{i_1=1}^{n_1} \dots \sum_{i_{\gamma-1}=1}^{n_{\gamma-1}} \sum_{i_{\gamma+1}=1}^{n_{\gamma+1}} \dots \sum_{i_d=1}^d [\mathbf{F}]_{i_1 \dots i_{\gamma-1} i_\gamma i_{\gamma+1} \dots i_d} \prod_{j \neq \gamma} [\tilde{\mathbf{f}}_j^1]_{i_j} \tag{19}$$

Thus, the operation represented in (18) (in compact form) and (19) (with the complete index notation) consists in contracting all the dimensions of tensor \mathbf{F} but one (the dimension γ , for $\gamma = 1, 2, \dots, d$) with the tensorial product of all the vectors $\tilde{\mathbf{f}}_j^1$, for $j \neq \gamma$.

The vector $\tilde{\mathbf{f}}_\gamma^1$ minimizing $\mathcal{J}(\tilde{\mathbf{f}}_\gamma^1)$ in (16) is precisely

$$\tilde{\mathbf{f}}_\gamma^1 = \frac{1}{\alpha} \mathbf{g} \tag{20}$$

This has to be done for all the dimensions, that is for $\gamma = 1, 2, \dots, d$, and iterated until the consecutive approximations of $\tilde{\mathbf{f}}_j^1$ reach stationarity. That is to say, assume each loop on γ is computing $\tilde{\mathbf{f}}_j^{\text{new}}$ as an approximation to $\tilde{\mathbf{f}}_j^1$ starting from initial guesses $\tilde{\mathbf{f}}_j^{\text{old}}$. Convergence is reached once, for some tolerance η , $\|\mathbf{f}_j^{\text{new}} - \mathbf{f}_j^{\text{old}}\| < \eta$ for all $j = 1, 2, \dots, d$. Note that the stationarity condition is expressed in terms of the normalized vectors \mathbf{f}_j to avoid the possibility that arbitrary constants with unit product may affect the d terms. Note that the sectional errors $\|\mathbf{f}_j^{\text{new}} - \mathbf{f}_j^{\text{old}}\|$ are already relative because all the vectors are normalized. Moreover, in the practical implementation, the global error taken into account is the product of all the sectional norms that stand for the Frobenius norm of the multidimensional error tensor, namely

$$\varepsilon = \|\mathbf{F}_{\text{PGD}}^{\text{new}} - \mathbf{F}_{\text{PGD}}^{\text{old}}\| = \left[\prod_{j=1}^d (\mathbf{f}_j^{\text{new}} - \mathbf{f}_j^{\text{old}})^\top (\mathbf{f}_j^{\text{new}} - \mathbf{f}_j^{\text{old}}) \right]^{1/2} \tag{21}$$

The strategy presented above for a rank-one approximation of some tensor \mathbf{F} is summarized in Algorithm 1.

Data: Tensor of order d to be approximated: \mathbf{F} with general term $[\mathbf{F}]_{i_1 i_2 \dots i_d}$, (for $i_\gamma = 1, \dots, n_\gamma$ and $\gamma = 1, \dots, d$)

Result: Rank-one approximation: $\mathbf{F}_{\text{PGD}} = \sigma \mathbf{f}_1 \otimes \mathbf{f}_2 \otimes \dots \otimes \mathbf{f}_d$

Initialize: Assign values to $\mathbf{f}_j^{\text{old}}$, for $j = 1, 2, \dots, d$; select a tolerance η

Compute $\sigma^{\text{old}} = \prod_{j=1}^d [\mathbf{f}_j^{\text{oldT}} \mathbf{f}_j^{\text{old}}]^{1/2}$;

Normalize: $\mathbf{f}_j^{\text{old}} \leftarrow [\mathbf{f}_j^{\text{oldT}} \mathbf{f}_j^{\text{old}}]^{-1/2} \mathbf{f}_j^{\text{old}}$, for $j = 1, 2, \dots, d$;

while $\varepsilon > \eta$ or $\varepsilon_\sigma > \eta$ **do**

for $\gamma = 1 \dots d$ **do**

 Compute $\alpha = \prod_{j \neq \gamma} \mathbf{f}_j^{\text{oldT}} \mathbf{f}_j^{\text{old}}$

 Compute \mathbf{g} such that

$$[\mathbf{g}]_{i_\gamma} = \sum_{i_1=1}^{n_1} \dots \sum_{i_{\gamma-1}=1}^{n_{\gamma-1}} \sum_{i_{\gamma+1}=1}^{n_{\gamma+1}} \dots \sum_{i_d=1}^d [\mathbf{F}]_{i_1 \dots i_{\gamma-1} i_\gamma i_{\gamma+1} \dots i_d} \prod_{j \neq \gamma} [\mathbf{f}_j^{\text{old}}]_{i_j}$$

 Compute $\mathbf{f}_\gamma^{\text{new}} = \frac{1}{\alpha} \mathbf{g}$

end

Compute $\sigma^{\text{new}} = \prod_{j=1}^d [\mathbf{f}_j^{\text{newT}} \mathbf{f}_j^{\text{new}}]^{1/2}$;

Normalize: $\mathbf{f}_j^{\text{new}} \leftarrow [\mathbf{f}_j^{\text{newT}} \mathbf{f}_j^{\text{new}}]^{-1/2} \mathbf{f}_j^{\text{new}}$, for $j = 1, 2, \dots, d$;

Compute modal error $\varepsilon = \left[\prod_{j=1}^d (\mathbf{f}_j^{\text{new}} - \mathbf{f}_j^{\text{old}})^T (\mathbf{f}_j^{\text{new}} - \mathbf{f}_j^{\text{old}}) \right]^{1/2}$;

Compute amplitude error $\varepsilon_\sigma = |\sigma^{\text{new}} - \sigma^{\text{old}}| / |\sigma^{\text{new}}|$;

Update $\mathbf{f}_j^{\text{old}} \leftarrow \mathbf{f}_j^{\text{new}}$, for $j = 1, 2, \dots, d$; $\sigma^{\text{old}} \leftarrow \sigma^{\text{new}}$;

end

Store $\mathbf{f}_j \leftarrow \mathbf{f}_j^{\text{new}}$, for $j = 1, 2, \dots, d$; $\sigma \leftarrow \sigma^{\text{new}}$;

Algorithm 1. Algebraic rank-one approximation for non-separable tensor \mathbf{F} .

2.3. Accounting for a separable input

Algorithm 1 is easily adapted to the case in which the input tensor is already provided in a separated format. That is, instead of having \mathbf{F} , we have Φ such that

$$\Phi = \sum_{\ell=1}^L \phi_1^\ell \otimes \phi_2^\ell \otimes \dots \otimes \phi_d^\ell \tag{22}$$

The LS approximation of Φ in the form of another separated tensor \mathbf{F}_{PGD} is meaningful, because it may significantly reduce the number of terms required to represent the same tensorial magnitude. That is, one would expect having $M \ll L$ and achieving a similar accuracy. Actually, this strategy is often used along the PGD computations because, when solving parametric boundary value problems, the PGD terms may exhibit redundancies (linear functional dependencies) from a LS viewpoint, see [18].

In practice, replacing the full tensor \mathbf{F} by the separated tensor Φ results only in a difference in the computation of the auxiliary vector \mathbf{g} in (18) or (19). Actually, in this case, (18) is readily replaced by

$$\mathbf{g} = \Phi \dots \bigotimes_{j \neq \gamma} \tilde{\mathbf{f}}_j^1 = \sum_{\ell=1}^L \left[\prod_{j \neq \gamma} \phi_j^{\ell T} \tilde{\mathbf{f}}_j^1 \right] \phi_\gamma^\ell \tag{23}$$

Thus, adapting Algorithm 1 to the case in which the input tensor is already expressed in separable format requires only replacing the line in which vector \mathbf{g} is computed, using (23) instead of (18). This is in practice the only difference between the algorithm providing rank-one tensor separation and rank-one tensor compression.

3. Complete LS PGD algorithm

The strategies described in the previous sections allow obtaining a LS PGD approximation $\mathbf{F}_{\text{PGD}}^M$ according to (14) of a d -dimensional tensor $\mathbf{F} \in \mathbb{R}^{n_1 \times \dots \times n_d}$. The same strategy works also for an already separated input tensor Φ as given in (22). Both the separation of \mathbf{F} and the compression of Φ are performed by algorithms having the same structure, and

Data: Tensor of order d to be approximated: \mathbf{F} with general term $[\mathbf{F}]_{i_1 i_2 \dots i_d}$,
 (for $i_\gamma = 1, \dots, n_\gamma$ and $\gamma = 1, \dots, d$)
 Tolerances $0 < \eta \ll 1$ for alternated directions and $0 < \eta^* \ll 1$ for greedy algorithm

Result: PGD approximation: $\mathbf{F}_{\text{PGD}} = \sum_{m=1}^M \sigma_m \mathbf{f}_1^m \otimes \mathbf{f}_2^m \otimes \dots \otimes \mathbf{f}_d^m$

Initialize counter of PGD terms $m = 1$ and starting value of amplitude $\sigma_1 = 1$

while $\sigma_m > \eta^* \sigma_1$ **do**

Initialize: Assign values to $\mathbf{f}_j^{\text{old}}$, for $j = 1, 2, \dots, d$;

Compute $\sigma^{\text{old}} = \prod_{j=1}^d [\mathbf{f}_j^{\text{oldT}} \mathbf{f}_j^{\text{old}}]^{1/2}$;

Normalize: $\mathbf{f}_j^{\text{old}} \leftarrow [\mathbf{f}_j^{\text{oldT}} \mathbf{f}_j^{\text{old}}]^{-1/2} \mathbf{f}_j^{\text{old}}$, for $j = 1, 2, \dots, d$;

while $\varepsilon > \eta$ or $\varepsilon_\sigma > \eta$ **do**

for $\gamma = 1 \dots d$ **do**

Compute $\alpha = \prod_{j \neq \gamma} \mathbf{f}_j^{\text{oldT}} \mathbf{f}_j^{\text{old}}$

Compute \mathbf{g} such that

$[\mathbf{g}]_{i_\gamma} = \sum_{i_1=1}^{n_1} \dots \sum_{i_{\gamma-1}=1}^{n_{\gamma-1}} \sum_{i_{\gamma+1}=1}^{n_{\gamma+1}} \dots \sum_{i_d=1}^d (\mathbf{F})_{i_1 \dots i_{\gamma-1} i_\gamma i_{\gamma+1} \dots i_d} \prod_{j \neq \gamma} [\mathbf{f}_j^{\text{old}}]_{i_j}$

$-\sum_{\tilde{m}=1}^{m-1} \sigma_{\tilde{m}} \left(\prod_{j \neq \gamma} \mathbf{f}_j^{\tilde{mT}} \mathbf{f}_j^{\text{old}} \right) [\mathbf{f}_\gamma^{\tilde{m}}]_{i_\gamma}$, for $i_\gamma = 1, 2, \dots, n_\gamma$;

Compute $\mathbf{f}_\gamma^{\text{new}} = \frac{1}{\alpha} \mathbf{g}$

end

Compute $\sigma^{\text{new}} = \prod_{j=1}^d [\mathbf{f}_j^{\text{newT}} \mathbf{f}_j^{\text{new}}]^{1/2}$;

Normalize: $\mathbf{f}_j^{\text{new}} \leftarrow [\mathbf{f}_j^{\text{newT}} \mathbf{f}_j^{\text{new}}]^{-1/2} \mathbf{f}_j^{\text{new}}$, for $j = 1, 2, \dots, d$;

Compute modal error $\varepsilon = \left[\prod_{j=1}^d (\mathbf{f}_j^{\text{new}} - \mathbf{f}_j^{\text{old}})^T (\mathbf{f}_j^{\text{new}} - \mathbf{f}_j^{\text{old}}) \right]^{1/2}$;

Compute amplitude error $\varepsilon_\sigma = |\sigma^{\text{new}} - \sigma^{\text{old}}| / |\sigma^{\text{new}}|$;

Update $\mathbf{f}_j^{\text{old}} \leftarrow \mathbf{f}_j^{\text{new}}$, for $j = 1, 2, \dots, d$; $\sigma^{\text{old}} \leftarrow \sigma^{\text{new}}$;

end

Store: $\sigma_m \leftarrow \sigma$; $\mathbf{f}_j^m \leftarrow \mathbf{f}_j^{\text{new}}$, for $i = 1, 2, \dots, d$;

Update $m \leftarrow m + 1$;

end

Store: $M \leftarrow m$;

Algorithm 2. Algebraic PGD approximation for multidimensional tensor \mathbf{F} .

the only difference is in the computation of vector \mathbf{g} , which is performed using (18) for separation and (23) for compression.

The present section is devoted to summarize the ideas of sections 2.1 and 2.2 in a compact algorithmic form.

As already indicated above, the main idea is to use the rank-one approximation of Algorithm 1 to compute the successive terms that conform \mathbf{F}_{PGD} , following the greedy approach described in section 2.1. The core of the algorithm is the rank-one updating from $\mathbf{F}_{\text{PGD}}^{m-1}$ to $\mathbf{F}_{\text{PGD}}^m$. This is essentially following the same rationale as in Algorithm 1, but replacing \mathbf{F} by $\mathbf{F} - \mathbf{F}_{\text{PGD}}^{m-1}$ as the tensor to be rank-one approximated.

Thus, the rank-one approximation inside the greedy loop aims at approximating $\mathbf{f}_1^m, \mathbf{f}_2^m, \dots, \mathbf{f}_d^m$ and σ_m , which provides the best rank-one update of $\mathbf{F}_{\text{PGD}}^{m-1}$. Here again, this results in a different expression to compute vector \mathbf{g} , alternative to (18), namely

$$\mathbf{g} = (\mathbf{F} - \mathbf{F}_{\text{PGD}}^{m-1}) \dots \otimes_{j \neq \gamma} \mathbf{f}_j^m = \mathbf{F} \dots \otimes_{j \neq \gamma} \mathbf{f}_j^m - \sum_{\tilde{m}=1}^{m-1} \sigma_{\tilde{m}} \left[\prod_{j \neq \gamma} \mathbf{f}_j^{\tilde{mT}} \mathbf{f}_j^m \right] \mathbf{f}_\gamma^{\tilde{m}} \quad (24)$$

Where the separated structure of \mathbf{F}_{PGD} is used in the last term, following exactly the same as in (23) with Φ .

Thus, the difference between the rank-one approximation described in Algorithm 1 and the complete PGD approximation summarized in Algorithm 2 lies in the addition of an outer loop on the number of terms (looping for m) and in the computation of \mathbf{g} , which is performed according to (24).

3.1. Algorithmic details: separated and complex inputs

Analogously as it discussed in Section 2.3, Algorithm 2 is easily adapted to accept a separable input Φ instead of a full multidimensional tensor F . Here again, the only difference between separation (of F) and compression (of Φ) lies in the expression to compute vector g , which for compression reads

$$g = (\Phi - F_{PGD}^{m-1}) \dots \bigotimes_{j \neq \gamma}^d f_j^m = \sum_{\ell=1}^L \left[\prod_{j \neq \gamma}^d \phi_j^{\ell T} f_j^m \right] \phi_\gamma^\ell - \sum_{\tilde{m}=1}^{m-1} \sigma_{\tilde{m}} \left[\prod_{j \neq \gamma}^d f_j^{\tilde{m} T} f_j^m \right] f_\gamma^{\tilde{m}} \tag{25}$$

In the case the input data is complex, that is either F or Φ lie in $\mathbb{C}^{n_1 \times \dots \times n_d}$ instead of $\mathbb{R}^{n_1 \times \dots \times n_d}$, the algorithm has to be slightly modified. In practice, all the modifications derive from the fact that the Frobenius-type norm for complex tensors differs from (10) in the sense that the first argument in the product has to be conjugated, and therefore it reads

$$\|F\|^2 = \bar{F} \dots F = [\bar{F}]_{j_1 j_2 \dots j_d} [F]_{j_1 j_2 \dots j_d} = \sum_{j_1=1}^{n_1} \sum_{j_2=1}^{n_2} \dots \sum_{j_d=1}^{n_d} |[F]_{j_1 j_2 \dots j_d}|^2 \tag{26}$$

where the bar ($\bar{\cdot}$) stands for the conjugate quantity. This affects the computation of the norms of the sectional vectors such that the scalar products require not only to transpose one of the vectors, but also to conjugate it. Moreover, in the tensorial multicontraction in, for instance, (25), what is located left to symbol \dots must be conjugated. This results in replacing (25) by

$$g = \overline{(\Phi - F_{PGD}^{m-1})} \dots \bigotimes_{j \neq \gamma}^d f_j^m = \sum_{\ell=1}^L \left[\prod_{j \neq \gamma}^d \bar{\phi}_j^{\ell T} f_j^m \right] \phi_\gamma^\ell - \sum_{\tilde{m}=1}^{m-1} \sigma_{\tilde{m}} \left[\prod_{j \neq \gamma}^d \bar{f}_j^{\tilde{m} T} f_j^m \right] \bar{f}_\gamma^{\tilde{m}} \tag{27}$$

3.2. Disambiguating normalization

In the stopping criterion of the alternated direction iterations, the comparison of the successive approximation to the sectional modes in Algorithms 1 and 2 is carried out in terms of the normalized vectors, see (21). Thus, in the algorithms, the normalization is readily indicated as $f_j^{new} \leftarrow [f_j^{new T} f_j^{new}]^{-1/2} f_j^{new}$, for $j = 1, 2, \dots, d$. However, this operation must be a little bit more sophisticated. In the real framework, the normalization must also disambiguate the sign. In other words, dividing by the norm, there could be ± 1 factors that may affect all the sectional modes (with an even number of -1 that globally cancel out). In this situation, the numerical stopping criterion is requiring more iterations, although, in practice, the stationarity in the alternated directions iterations has already been reached. The solution to this problem is simple, but it cannot be overlooked. One possibility is to disambiguate the sign by enforcing that the largest (in absolute value) component of f_j^{new} is positive.

If dealing with complex variables, the standard normalization would be $f_j^{new} \leftarrow [f_j^{new T} f_j^{new}]^{-1/2} f_j^{new}$, for $j = 1, 2, \dots, d$. In this case, the ambiguity is not only up to a sign, but up to any complex coefficient with unit module. The different modes f_j^{new} may be affected by coefficients z_j , $j = 1, 2, \dots, d$, with $|z_j| = 1$ and such that $\prod_{j=1}^d z_j = 1$ to cancel out globally. In order to disambiguate these factors, the normalization is carried out such that the largest (in module) component of f_j^{new} is real and positive.

4. Numerical examples

The methodologies presented above for PGD tensor separation and tensor compression are tested in this section for different numerical examples.

4.1. Example 1: separation of a 2D manufactured tensor

The first example uses a tensor F synthetically manufactured from a reduced number (six) of separated modes. Thus, the separation algorithm, whose input is the full reconstructed tensor is challenged to produce a number of modes similar to the number used to manufacture the input. The bidimensional ($d = 2$) tensor F is generated from a separated expression,

$$\Phi(x, y) = \sum_{\ell=1}^6 \phi_x^\ell \otimes \phi_y^\ell \tag{28}$$

where the modal vectors ϕ_x^ℓ and ϕ_y^ℓ , for $\ell = 1, \dots, 6$, are defined as the images of vector $\xi = [0 \ \frac{1}{99} \ \frac{2}{99} \ \dots \ \frac{98}{99} \ 1]^T \in \mathbb{R}^{100}$ by different continuous functions taking values in $[0, 1]$, namely $\phi_x^\ell = \phi_x^\ell(\xi)$ and $\phi_y^\ell = \phi_y^\ell(\xi)$, where

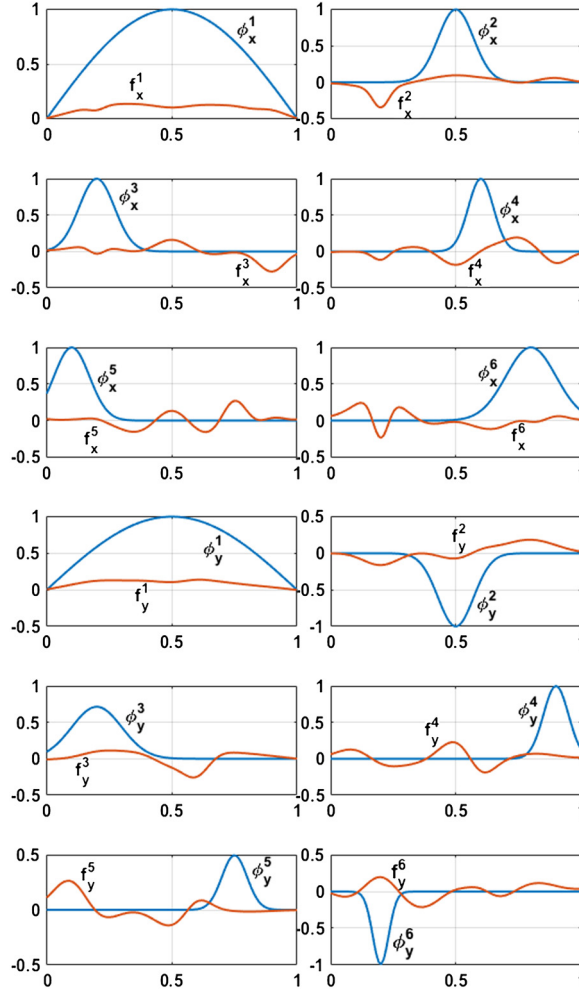


Fig. 1. Example 1: input modal functions ϕ_x^ℓ and ϕ_y^ℓ , for $\ell = 1, \dots, 6$ (in blue) used to manufacture \mathbf{F} and modes \mathbf{f}_x^j and \mathbf{f}_y^j , for $\ell = 1, \dots, M = 6$ (in red) of the separated representation of \mathbf{F}_{PGD} .

$$\phi_x^1(x) = \sin(\pi x) \qquad \phi_y^1(y) = \sin(\pi y) \qquad (29a)$$

$$\phi_x^2(x) = \exp\left(\frac{-(x-0.5)^2}{0.01}\right) \qquad \phi_y^2(y) = -\exp\left(\frac{-(y-0.5)^2}{0.01}\right) \qquad (29b)$$

$$\phi_x^3(x) = \exp\left(\frac{-(x-0.2)^2}{0.01}\right) \qquad \phi_y^3(y) = \exp\left(\frac{-(y-0.2)^2}{0.02}\right) \qquad (29c)$$

$$\phi_x^4(x) = \exp\left(\frac{-(x-0.6)^2}{0.005}\right) \qquad \phi_y^4(y) = \exp\left(\frac{-(y-0.9)^2}{0.005}\right) \qquad (29d)$$

$$\phi_x^5(x) = \exp\left(\frac{-(x-0.1)^2}{0.01}\right) \qquad \phi_y^5(y) = \exp\left(\frac{-(y-0.75)^2}{0.005}\right) \qquad (29e)$$

$$\phi_x^6(x) = \exp\left(\frac{-(x-0.8)^2}{0.02}\right) \qquad \phi_y^6(y) = -\exp\left(\frac{-(y-0.2)^2}{0.002}\right) \qquad (29f)$$

Fig. 1 shows the 1D modal functions ϕ_x^ℓ and ϕ_y^ℓ . The reconstructed 2D tensor \mathbf{F} is represented in Fig. 4a. The PGD separation described in Algorithm 2 is applied to \mathbf{F} , and provides

$$\mathbf{F}_{\text{PGD}}^M = \sum_{m=1}^M \sigma_m \mathbf{f}_x^m \otimes \mathbf{f}_y^m \qquad (30)$$

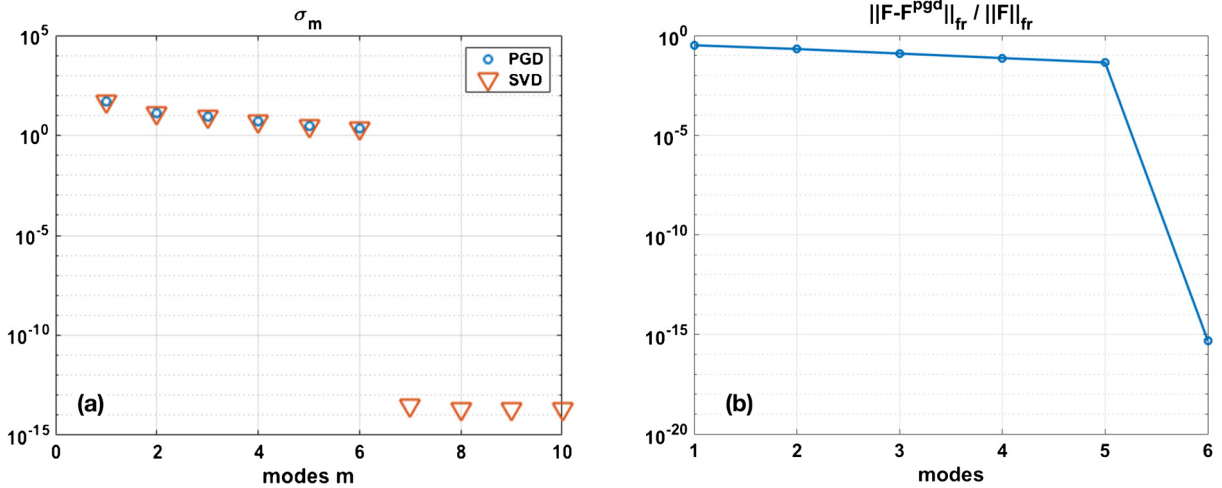


Fig. 2. Example 1: [panel a] modal amplitudes (σ_m) of the first six terms of the PGD separated tensor and comparison with the modal amplitudes provided by the Singular Value Decomposition method; [panel b] relative error of the PGD separated tensor in Frobenius norm as a function of the number of terms.

The functions corresponding to vectors f_x^m and f_y^m defining the first six modes (depicted in Fig. 1) account for the exact separated representation of F . The algorithm in two dimensions provides an optimal separated form coinciding with the results provided by a singular value decomposition method (see Fig. 2a). The modal error evolution is computed as the Frobenius norm of the difference of the tensors F and F_{PGD}^M . The exact solution is reached with six modes up to machine precision, as can be seen in Fig. 2b.

4.2. Example 2: separation of a higher-dimension tensor

In order to analyze the behavior of the algorithm in higher dimensions, the previous case has been extended to a seven-dimensional synthetic problem. A set of functions similar to (29) is used to build the higher-order tensor,

$$\begin{aligned} \phi_k^1(x_k) &= \sin(\pi x_k) \\ \phi_k^j(x_k) &= a_{jk} \exp\left(\frac{-(x_k - b_{jk})^2}{c_{jk}}\right) \end{aligned}$$

for $j = 2, \dots, 6$ and $k = 1, \dots, 7$, in the domain $[0, 1]^7$. The discretization of each dimension is done with 10 points, that is, $\xi = [0 \frac{1}{9} \frac{2}{9} \dots \frac{8}{9} 1]^T \in \mathbb{R}^{10}$, providing a full tensor with 10^7 real entries. Coefficients a , b and c are described in Appendix A.

The results of the seven-dimensional case show that the obtained separation does not recover the original six modes used to build the tensor. It actually requires 150 terms to obtain a separate representation with a relative error of 10^{-3} (Fig. 3). This is standard in higher dimensional problems, because none of the available separation procedures is able to provide optimal solutions for $d > 2$ [14]. Note, however, that the reduction in terms of storage is significant: the 10^7 real entries of the full tensor are reduced in the separated form to $150 \times 7 \times 10 \approx 10^4$. That is, the relative error of 10^{-3} is the (small) price to pay for saving three orders of magnitude in the storage.

4.3. Example 3: denoising data

An additional application of the proposed PGD algorithm is to filter noise in a data set. In this context, this example retakes the previous 2D case of Section 4.1 perturbing the input by adding a random noise with an amplitude of 2% of the maximum tensor value, see Fig. 4b. Thus, the 6 modes describing the tensor presented in Example 1 are duly identified and separated from the noise that is actually modeled (or represented) by the remaining modes. Fig. 5 shows the modal error evolution with respect to the original and noised tensors. As it can be seen, the first six modes provide an accurate approximation of the original tensor, while from mode 7 on, the algorithm is adding to the separated representation the information added by the noise. It can be observed in Fig. 5 that the discrepancies $\|F_{noised} - F_{PGD}\|$ and $\|F - F_{PGD}\|$ differ at the sixth mode and the following ones. This is associated with the fact that F_{PGD} is computed as an approximation of the tensor perturbed by the noise, F_{noised} , and not from the original one, F . This is essentially visible for the sixth mode, because it is where the relative amplitude of the mode (the accuracy) meets the magnitude of the noise.

This is the reason why the error computed against the original tensor increases after mode 6. Furthermore, in case of the error computed against the noised tensor, it decreases as the solution approaches the noised configuration. Note that from mode 6 on, such error decreases linearly with a very low convergence rate, as expected in the reproduction of a function with structure (noise).




Fig. 3. Example 2: [panel a] modal amplitudes (σ_m) of the first 150 terms of the PGD separated tensor; [panel b] relative error of the PGD separated tensor in Frobenius norm as a function of the number of terms.

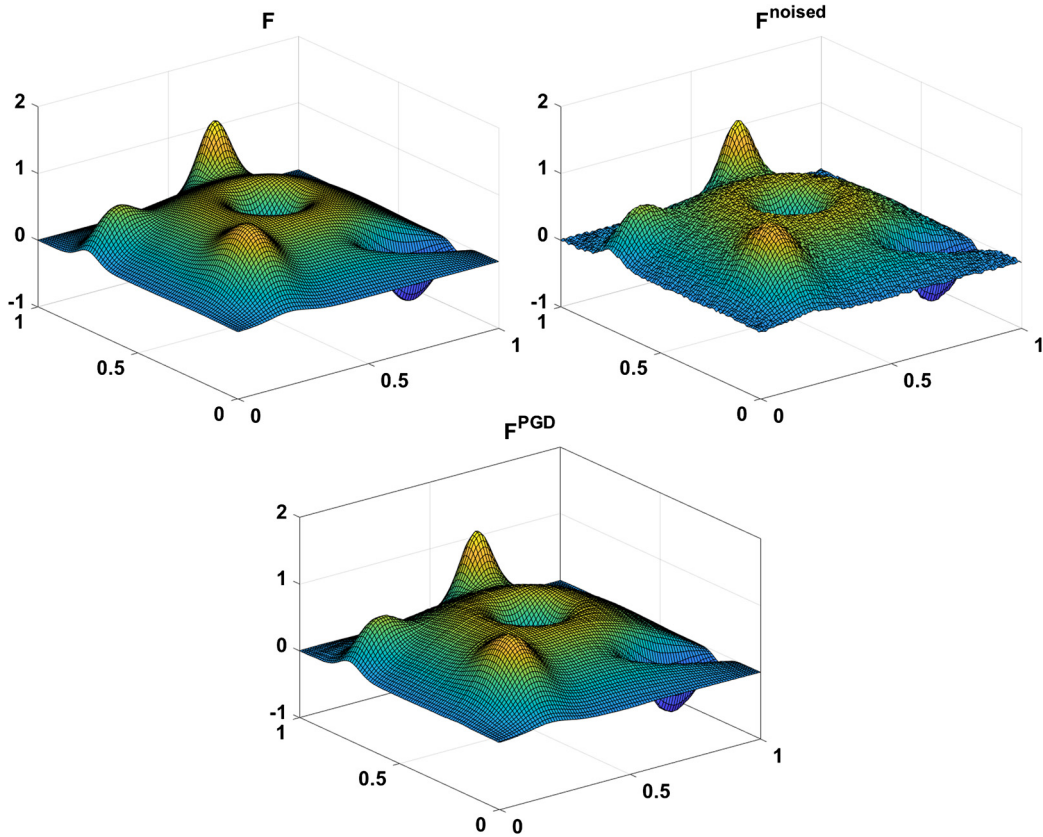


Fig. 4. Examples 1 and 3: [top left panel] Representation of the tensor F used in Example 1; [top right panel] noised version of F ; [bottom panel] denoised tensor via the PGD algorithm.

4.4. Example 4: compression of higher-order complex separated tensor

The PGD compression strategy presented above is here applied to a complex tensor of large size and dimension $d = 4$. The tensor, Φ , is already available in a separated format, as shown in (22), and contains the data to be compressed. It is the PGD solution to a parameterized sea wave propagation problem in harbor taken from [18], see Fig. 6 for an illustration. The four dimensions correspond to the space distribution of the (complex) wave height and the three parameters, which are: 1) the




Fig. 5. Example 3: evolution with the number of modes of the relative error of the result of the PGD algorithm, F_{PGD} , computed against i) the original tensor F , and ii) the noised version F_{noised} (used as input for PGD).

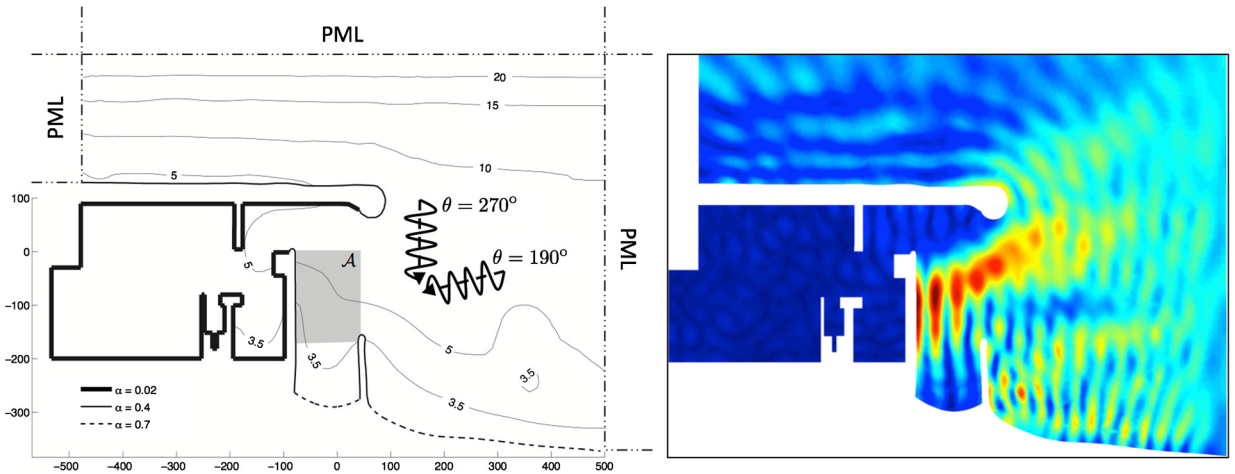


Fig. 6. Example 4: [left panel] domain, boundary conditions and parameters of the parameterized wave problem (from [18]); [right panel] illustration of one spatial solution for some set of parameters.

incoming wave direction, 2) the wave frequency, and 3) the reflectivity of the coastline. The original dataset has $L = 1500$ modes using a discretization of space, frequency, angle, and reflectivity of 15,757, 100, 50 and 10 degrees of freedom respectively. Note that the problem is stated in frequency domain and all dimensions are stored as complex numbers. The full tensor of this solution, therefore, occupies 12 GB of memory (corresponding to $15,757 \times 100 \times 50 \times 10 = 787.85 \times 10^6$ complex entries) The storage of the separated solution with 1500 modes (corresponding to $1500 \times (15,757 + 100 + 50 + 10) = 23.8755 \times 10^6$ complex entries) requires 364 MB. The goal of the compression is to approximate Φ by F_{PGD} with $M \ll L$ and to significantly reduce the storage requirements while maintaining the accuracy of the represented quantities. Here, PGD compression is carried out to obtain 200 modes (instead of using a tolerance η^* to stop the process). Note that with 200 modes, the storage requirement reduces to 48.5 MB. The objective is to check whether PGD compression produces a better description of the data when the number of modes is limited by the storage capacity. The results are shown in Fig. 7. On the left, the evolution of the modal amplitudes is represented, and, as expected, they globally show a decreasing trend until they get stabilized. More interestingly, in the right-hand-side curve, the actual error $\|\Phi - F_{PGD}^m\|$ is represented and shows a monotonic decreasing behavior. Moreover, also indicated as a reference value we display the error associated with the truncation of Φ to the first 200 terms, namely $\|\Phi - \sum_{\ell=1}^{200} \phi_1^\ell \otimes \phi_2^\ell \otimes \phi_3^\ell \otimes \phi_4^\ell\|$. This value is represented in the right panel of Fig. 7 as a red cross. It can be noted that, with respect to the first 200 terms of the original separation of Φ , the error associated with F_{PGD}^{200} is by one order of magnitude lower.

This example therefore demonstrates that the PGD compression is able to shorten the separable expressions and improve the accuracy of the overall representation at a limited storage capacity.




Fig. 7. Example 4: [left panel] evolution of modal amplitudes (σ_m) along the first 200 terms of PGD compression; [right panel] evolution of the relative error of the PGD compression (with respect to the full Φ tensor). The red cross marks the value of the norm of the difference between Φ and the tensor reconstructed with the first original 200 modes of Φ .

5. Concluding remarks

The PGD least-squares approximation is presented here as a computational tool to perform high-dimensional tensor separation at an affordable computational cost and with a limited coding complexity. The algorithms are discussed in detail, with special emphasis on the stopping criteria for both the greedy strategy and the alternated directions iteration scheme.

Moreover, the same idea is also used to compress separated approximations with a large number of terms and to reduce the storage requirements, while keeping the accuracy of the separated representation.

Appendix A

The functions used in the construction of the seven-dimensional tensor used in Section 4.2 are described next. The matrix \mathbf{A} , built with the coefficients a_{jk} reads:

$$\mathbf{A} = \begin{bmatrix} 1 & -1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}$$

Similarly, the coefficients b_{jk} and c_{jk} are:

$$\mathbf{B} = \begin{bmatrix} 0.5 & 0.5 & 0.5 & 0.5 & 0.5 & 0.5 & 0.5 \\ 0.2 & 0.2 & 0.2 & 0.2 & 0.2 & 0.2 & 0.2 \\ 0.6 & 0.9 & 0.6 & 0.6 & 0.6 & 0.6 & 0.6 \\ 0.1 & 0.75 & 0.1 & 0.1 & 0.1 & 0.1 & 0.1 \\ 0.8 & 0.2 & 0.8 & 0.8 & 0.8 & 0.8 & 0.8 \end{bmatrix} \quad \mathbf{C} = \begin{bmatrix} 0.01 & 0.01 & 0.01 & 0.01 & 0.01 & 0.01 & 0.01 \\ 0.01 & 0.02 & 0.01 & 0.02 & 0.02 & 0.02 & 0.02 \\ 0.01 & 0.01 & 0.01 & 0.01 & 0.01 & 0.01 & 0.01 \\ 0.01 & 0.01 & 0.01 & 0.01 & 0.01 & 0.01 & 0.01 \\ 0.01 & 0.02 & 0.01 & 0.02 & 0.02 & 0.02 & 0.02 \end{bmatrix}$$

References

- [1] L.R. Tucker, Some mathematical notes on three-mode factor analysis, *Psychometrika* 31 (1966) 279–311.
- [2] R.A. Harshman, Foundations of the parafac procedure: models and conditions for an explanatory multi-modal factor analysis, *UCLA Work. Pap. Phon.* 16 (1970) 1–84.
- [3] C.J. Appellof, E.R. Davidson, Strategies for analyzing data from video fluorometric monitoring of liquid chromatographic effluents, *Anal. Chem.* 53 (1981) 2053–2056.
- [4] R. Bro, PARAFAC. Tutorial and applications, *Chemom. Intell. Lab. Syst.* 38 (1997) 149–171.
- [5] C. Beckmann, S. Smith, Tensorial extensions of independent component analysis for multisubject FMRI analysis, *NeuroImage* 25 (2005) 294–311.
- [6] M. De Vos, A. Vergult, L. De Lathauwer, W. De Clercq, S. Van Huffel, P. Dupont, A. Palmi, W. Van Paesschen, Canonical decomposition of ictal scalp eeg reliably detects the seizure onset zone, *NeuroImage* 37 (2007) 844–854.
- [7] H. Wang, N. Ahuja, Facial expression decomposition, in: *ICCV 2003: Proceedings of the 9th IEEE International Conference on Computer Vision, Nice, France, 2003*, vol. 2, 2003, pp. 958–965.

- [8] A. Shashua, T. Hazan, Non-negative tensor factorization with applications to statistics and computer vision, in: ICML 2005: Proceedings of the 22nd International Conference on Machine Learning, Bonn, Germany, 2005, pp. 792–799.
- [9] B. Chen, A. Petropolu, L. De Lathauwer, Blind identification of convolutive mimo systems with 3 sources and 2 sensors, *EURASIP J. Appl. Signal Process.* 5 (2002) 487–496.
- [10] L. De Lathauwer, J. Vandewalle, Dimensionality reduction in higher-order signal processing and rank- (r_1, r_2, \dots, r_n) reduction in multilinear algebra, *Linear Algebra Appl.* 391 (2004) 31–55.
- [11] N. Liu, B. Zhang, J. Yan, Z. Chen, W. Liu, F. Bai, L. Chien, Text representation: from vector to tensor, in: Proceedings of the 5th IEEE International Conference on Data Mining, Houston, TX, USA, 2005, pp. 725–728.
- [12] I. Ibragimov, Application of the three-way decomposition for matrix compression, *Numer. Linear Algebra Appl.* 9 (2002) 551–565.
- [13] W. Hackbusch, B.N. Khoromskij, Tensor-product approximation to operators and functions in high dimensions, *J. Complex.* 23 (2007) 697–714.
- [14] T. Kolda, B. Bader, Tensor decompositions and applications, *SIAM Rev.* 51 (2009) 455–500.
- [15] F. Chinesta, E. Cueto, A. Huerta, PGD for solving multidimensional and parametric models, in: *Separated Representations and PGD-Based Model Reduction*, in: CISM Courses and Lectures, vol. 554, Springer, Vienna, 2014, pp. 27–89.
- [16] F. Chinesta, R. Keunings, A. Leygue, *The Proper Generalized Decomposition for Advanced Numerical Simulations. A Primer*, Springer Briefs in Applied Sciences and Technology, Springer, Cham, Switzerland, 2014.
- [17] S. Zlotnik, P. Díez, D. Modesto, A. Huerta, Proper generalized decomposition of a geometrically parametrized heat problem with geophysical applications, *Int. J. Numer. Methods Eng.* 103 (10) (2015) 737–758.
- [18] D. Modesto, S. Zlotnik, A. Huerta, Proper generalized decomposition for parameterized Helmholtz problems in heterogeneous and unbounded domains: application to harbor agitation, *Comput. Methods Appl. Mech. Eng.* 295 (2015) 127–149.