



HAL
open science

The M-Tensor Format: Optimality in High Dimensional Regression for Nonlinear Models with Scarce Data

Rémi Cloarec, Sebastian Rodriguez, Xavier Kestelyn, Francisco Chinesta

► To cite this version:

Rémi Cloarec, Sebastian Rodriguez, Xavier Kestelyn, Francisco Chinesta. The M-Tensor Format: Optimality in High Dimensional Regression for Nonlinear Models with Scarce Data. 2026. <hal-05500936>

HAL Id: hal-05500936

<https://hal.science/hal-05500936v1>

Preprint submitted on 9 Feb 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

The M-Tensor Format: Optimality in High Dimensional Regression for Nonlinear Models with Scarce Data *

Rémi Cloarec^{†‡} Sebastian Rodriguez[†] Xavier Kestelyn[‡]
Francisco Chinesta[†]

Monday 9th February, 2026

Abstract

We present a nonlinear regression framework based on tensor algebra tailored to high dimensional contexts where data is scarce. We exploit algebraic properties of a partial tensor product, namely the m-tensor product, to leverage structured equations with separated variables. The proposed method combines kernel properties along with tensor algebra to prevent the curse of dimensionality and tackle approximations up to hundreds of parameters while avoiding the fixed point strategy. This formalism allows us to provide different regularization techniques fit for low amount of data with a high number of parameters while preserving well-known matrix-based properties. We demonstrate complexity scaling on a general benchmark and dynamical systems to show robustness for engineering problems and ease of implementation.

Keywords Multilinear algebra, Tensor format, High-dimensional regression, Kernel-based regression

MSC codes 15A23, 15A69, 65D40

Contents

1	Introduction	2
2	Least squares approximation of a scalar function	4
3	Separable variables and tensor structure	5
3.1	Row-wise tensor product	6
3.2	Transposition	7
3.3	Contractions	7
3.4	Hadamard product	7
3.5	Inner product and norm	8
3.6	N-way contracted product	8

*Submitted to the editors DATE.

This work was funded by the ENSAM RTE Chair.

[†]ENSAM RTE Chair, PIMM, Arts et Metiers ParisTech, Paris, France (remi.cloarec@ensam.eu, sebastian.rodriguez_iturra@ensam.eu, francisco.chinesta@ensam.eu).

[‡]ENSAM RTE Chair, University of Lille, Arts et Metiers ParisTech, Centrale Lille, HEI, EA 2697-L2EP, Lille, France (xavier.kestelyn@ensam.eu).

4	M-Tensor least squares regression	8
5	Regularization of m-tensor regression	10
5.1	Tikhonov regularization	10
5.2	Spectral truncation	11
5.3	ALI regularization	12
6	Kernel regression interpretation	13
7	Numerical examples	14
7.1	Rosenbrock function approximation	14
7.2	Prediction of dynamical systems in time	16
7.2.1	3D Lorenz attractor	17
7.2.2	Kuramoto oscillators	18
7.3	A practical note	19
8	Discussion	20
9	Conclusion and perspectives	21
	Bibliography	22
A	Tensor operations complexity analysis	24
A.1	Rank-1 tensor contraction	24
A.2	Rank-1 tensors Hadamard product	24
A.3	M-product complexity	25
B	Polynomial toy problem	26
C	Algorithms to compute ALI decomposition	27
C.1	Greedy decomposition	28
C.2	Optimal decomposition	30

1 Introduction

Engineering problems are often made dependent on parameters. In certain cases, the number of those parameters might be reduced by identifying fewer latent variables as is done by model order reduction techniques. In certain scenarii the dependence has to be explicit. In that case even for problems with low latent dimensionality, making explicit use of those parameters becomes intractable. This issue is often found in design optimization problems for instance. Manipulating a large number of parameters often becomes numerically intractable because of a wall of complexity coined *curse of dimensionality* [4]. This curse translates the exponential growth of models (*i.e.* the number of coefficients) and that of the required amount of data with respect to the number of parameters. In contexts where data is scarce and costful, that curse makes models simply inaccessible.

Solutions proposed to circumvent this curse can rely on a modification of the grid required by the model, *e.g.* with a sparse grid [15], or with an *adaptive* selection of the pertinent points of that grid to retain [6]. A reduction of complexity can also be found by leveraging the structure of the equations. A common approach is to formulate the problem with separated variables. A prime example of this approach can be found in the works of Beylkin and Mohlenkamp [5]. The separated variables representation inherently relies on a tensorial factorization of the operators,

often based on the use of Kronecker product-based linear systems [2, 19]. This reliance on tensor representation can explain the recent development of plenty of methods based on tensor network formalisms, more specifically the Tensor-Train (TT-) format [23] and the overall gain of interest toward tensor methods [17, 12, 13].

A popular variable separation technique stemming from computational mechanics is the Proper Generalized Decomposition (PGD) [1] and its sparse data-based version, the s-PGD [14]. Those approaches rely on a Canonical Polyadic Decomposition (CPD) of the field that is learnt from one-dimensional modes, *i.e.* a rank-1 tensor is added at each loop on all axes. It resembles the CP-ALS method [17] and solves for each mode employing a fixed-point strategy. This fixed-point approach implies a sequential construction and may append modes late in the algorithm that serve only as correctives for early ones that were not optimal.

Many recent approaches are based on the TT-format and its good scaling with high-dimensional problems. A framework that inspired this paper is that of MANDy [10]. In this work, Gelß et al. adapt the popular SINDy framework [7] to high-dimensional dynamical system identification, which remains a limitation for SINDy. To achieve good scaling, the authors built sparse TT cores from a choice of one-dimensional basis functions. The same authors later expanded the same methodology to image classification [16] where they started mixing their tensorial framework with kernel interpretations to tackle efficiently higher dimensional problems without relying on the TT-SVD sequential algorithm which they demonstrated became a computational bottleneck. At the same time as Baddoo et al. with their fully kernel-based approach to dynamic model learning [3], Gelß et al. published the kernel Feature Space Approximation (kFSA) method that once again mixes TT-formulated feature maps to kernel methods to provide a high-dimensional regularization (and model reduction) technique. We further investigated the work of Engel et al. in signal processing [8] that develop a kernel recursive least squares algorithm for online learning of signals. In their work, they provide nice proofs, noticeably on the limitation of the number of retained signals (finite dictionary) and a bound with respect to ℓ_2 optimality provided by the kernel Principal Component Analysis (kPCA). The evaluation of distance in the feature space that defines their Almost Linearly Dependent (ALD) criterion is the same later used by Gelß and Baddoo in their respective works. We use the same approach in our framework.

In this work, we keep the philosophy of structuring the equations that build the regression operator from a separated variables perspective. Considering the lack of data as a structural constraint, we isolate it in the structure of the proposed operator and define a *partial* tensor product. We consider it an algebraically minimalistic approach to generalize matrix-based regression. It allows us to preserve properties of usual matrix-based linear systems while avoiding the curse of dimensionality. By circumventing the fixed point strategy usually encountered in high dimensional contexts, we can produce an optimal result in the ℓ_2 sense from available data at a very low cost. We present the approach as *algebraically minimalistic* for its conceptual simplicity compared to tensor networks or abstract kernel feature spaces for instance. Working with tensor algebra and *then* identifying the kernel from the dot products also allows us to avoid the usual limitation of kernel methods to a few well-known kernels while leveraging their properties. For instance we are able to use the properties of kPCA. Furthermore, the sought minimalism allows us to avoid the necessity for sparse operators in practice for computational efficiency in the case of TT-formulation. The structure of the solution we build relates to the state of the art as a rank-wise unconstrained version of the s-PGD [14].

The proposed approach is structured around a tensor formalism from which we derive nice algebraic properties and computational tricks. We call it the *m-tensor* format as it directly relates to the unmanageable *matrix* that would be built naively to solve the problem. This is how we retain many of the desirable properties of the matrix-based approach in a fully (tensor-)factorized manner. By preserving a structure of higher order, we enjoy the great efficiency

of tensor-based approaches. Those properties would otherwise be lost by the matrix-based structure. As it is tailored for regression operators, this formalism is not investigated as a potential data structure for high order problems. The tools used in this paper are available on GitHub¹ along with numerical examples.

The paper is organized as follows. In section 2, we set the problem of the least squares approximation of a scalar function in a general form. Section 3 builds the link between separable variables and tensor structure, this leads us to present the tensor operations we will use within our framework. The m-tensor format is presented in the case of least squares regression in section 4. Section 5 provides different regularization techniques for m-tensor regression that we relate to kernel regression in section 6. We apply our method to demonstrate scalability and robustness in section 7 on a polynomial benchmark with growing dimension and then on two dynamical systems: a Lorenz attractor and Kuramoto oscillators. We propose a short discussion comparing the m-tensor formalism to other techniques in section 8. Section 9 draws conclusions and presents perspectives for future works.

Notations

We present the notations adopted in the remainder of the paper. Operation symbols are given with their definition. Working with high dimensional objects relies on burdensome notations in order to disambiguate definitions, therefore we try to strictly stick to conventions given here. Uppercase bold symbols (*e.g.* \mathbf{T}) are used for tensors. Uppercase symbols (*e.g.* M) are used for matrices. Lowercase bold symbols (*e.g.* \mathbf{v}) are used for vectors. The notation $[\cdot]_{ijk\dots}$ denotes the $ijk\dots$ -th element of the tensor in brackets, each index evolving along an axis of the tensor. Due to the high order nature of the object of the paper, indices are more frequently subscripted with axis number, *i.e.* elements of a tensor \mathbf{T} of order n will be denoted as $[\mathbf{T}]_{i_1 i_2 \dots i_n}$.

2 Least squares approximation of a scalar function

In order to provide a general approach, let $\mathcal{D} \subset \mathbb{R}^n$ denote a n -dimensional domain, defined as a tensor product of one-dimensional intervals in \mathbb{R}

$$\mathcal{D} = \bigotimes_{i=1}^n [a_i, b_i], \quad [a_i, b_i] \subset \mathbb{R}.$$

The aim is to approximate a scalar function f defined over that domain. Let φ denote basis functions to be chosen, we define an approximation of f with the form

$$\begin{aligned} \hat{f}(\mathbf{x}) &= \sum_{i=1}^{\mathcal{N}} c_i \varphi_i(\mathbf{x}), \\ &= \boldsymbol{\varphi}^\top(\mathbf{x}) \mathbf{c}, \end{aligned} \tag{2.1}$$

where coefficients \mathbf{c} are the *best fit* (in the least squares sense) for the chosen bases functions. The number of terms \mathcal{N} depends on the choice of bases as well. This formulation is meant to accommodate for a large array of problems in a uniform manner. Let \mathcal{P} denote a set of m distinct samples in \mathcal{D} such that

$$\mathcal{P} = \{\mathbf{x}_k \in \mathcal{D}, \forall k, 1 \leq k \leq m\}.$$

Consider we are given the evaluation of f at each sample of \mathcal{P} . We denote $\mathbf{y} \in \mathbb{R}^m$ the vector defined as

$$[\mathbf{y}]_k = f(\mathbf{x}_k) \quad \forall k, 1 \leq k \leq m, \mathbf{x}_k \in \mathcal{P}.$$

¹GitHub repository for the toolbox and examples: <https://github.com/RemiC-OV/MTensor>.

Let $\Phi \in \mathbb{R}^{m \times \mathcal{N}}$ denote the evaluation of bases φ used in eq. (2.1) over the set \mathcal{P} . With $[\Phi]_k$ denoting the k -th row of Φ , we have

$$[\Phi]_k = \varphi^\top(\mathbf{x}_k) \quad \forall k, 1 \leq k \leq n, \mathbf{x}_k \in \mathcal{P}.$$

Then, the least squares solution $\hat{\mathbf{c}}$ is the one that, for any given bases φ , minimizes the squared ℓ_2 error with respect to \mathbf{y} :

$$\hat{\mathbf{c}} = \underset{\mathbf{c}}{\operatorname{argmin}} \left\| \sum_{k=1}^m \varphi^\top(\mathbf{x}_k) \mathbf{c} - [\mathbf{y}]_k \right\|_2^2. \quad (2.2)$$

This minimization problem admits an explicit solution based on whether Φ has linearly independent columns eq. (2.3) or rows eq. (2.4) that we build using the Moore-Penrose inverse formulation [24]:

$$\hat{\mathbf{c}} = (\Phi^\top \Phi)^{-1} \Phi^\top \mathbf{y}, \quad (2.3)$$

$$\hat{\mathbf{c}} = \Phi^\top (\Phi \Phi^\top)^{-1} \mathbf{y}. \quad (2.4)$$

Our choice is to use this widely used explicit formulation of least squares solution to determine the coefficients \mathbf{c} of the approximation. Without any prior knowledge of f , the aim is then to define basis functions φ such that f is well approximated. To that end, one can chose to tailor the functions in order to minimize \mathcal{N} or provide a rich enough set of bases so that all features of f can be captured. In what follows we chose the latter and show how our tensorial framework alleviates the curse of dimensionality within the regression approach.

3 Separable variables and tensor structure

A popular approach in high dimension is the separation of variables. In this setting, the bases functions φ are sought as products of one-dimensional functions. We denote those one-dimensional functions ψ . This leads us to define multidimensional bases φ as

$$\forall i, 1 \leq i \leq \mathcal{N}, \varphi_i(\mathbf{x}) = \prod_{j=1}^n \psi_{ij}([\mathbf{x}]_j).$$

Like all combinatorial problems, no standard approach would scale well in order to compute all combinations of one-dimensional basis functions. For that reason we demonstrate the use of a convenient tensor framework. Let $\boldsymbol{\psi} \in \mathbb{R}^p$ denote the vector of p one-dimensional basis functions deemed rich enough to represent the problem in one dimension given as

$$\boldsymbol{\psi}(x) = [\psi_1(x) \quad \psi_2(x) \quad \cdots \quad \psi_p(x)].$$

All combinations of one-dimensional bases would be found without redundancy in the elements of the tensor φ defined as

$$\varphi(\mathbf{x}) = \bigotimes_{j=1}^n \boldsymbol{\psi}([\mathbf{x}]_j).$$

We note that $\varphi \in \mathbb{R}^{p \times p \times \cdots \times p}$ is a rank-1 tensor of order n . Having the same set of basis functions along each axis is an arbitrary choice made to alleviate notations. Based on this representation, the formulation of the approximation of f in (2.1) can be written as the inner product

$$\hat{f}(\mathbf{x}) = \langle \varphi(\mathbf{x}), \mathbf{C} \rangle. \quad (3.1)$$

Here $\mathbf{C} \in \mathbb{R}^{p \times p \times \dots \times p}$ is an order n tensor of coefficients. The least squares coefficient tensor $\hat{\mathbf{C}}$ is the one minimizing the ℓ_2 error with respect to \mathbf{y} :

$$\hat{\mathbf{C}} = \underset{\mathbf{C}}{\operatorname{argmin}} \left\| \sum_{k=1}^m \langle \boldsymbol{\varphi}(\mathbf{x}_k), \mathbf{C} \rangle - [\mathbf{y}]_k \right\|_2^2. \quad (3.2)$$

If we were to unfold (vectorize) this tensor formulation, we would find the expression given for (2.1) with $\mathcal{N} = p^n$. Given this structure, we propose a tensor framework to efficiently build approximations in high dimensions.

We give a few reminders to work with tensors that allow us to introduce our framework. Tensors are high order objects and their use allows far more flexibility than the usual matrix/vector linear algebra. They allow to preserve structure that would be lost by flattening the problem in a matrix/vector formulation. The most important aspect is the ability to only use factorized formulations for problems that would otherwise be inaccessible memory- and computationally-wise. We justify the given complexities in appendix section A. For more details on tensor formats and decompositions see [12, 13, 17].

3.1 Row-wise tensor product

We first define a row-wise tensor product. Further in this paper we refer to it as the *m-tensor product*. It generalizes the face splitting product [28] (transposed Khatri-Rao) the same way the tensor product generalizes the matrix Kroenecker product.

Our approach uses a similar treatment of the indices as the face-splitting product but conserves its multidimensional structure, therefore avoiding the index mapping burden. The m-tensor product is computed from an arbitrary number of input matrices, provided they all share the same number of rows. For a given set of matrices $\{\Psi_i \in \mathbb{R}^{m \times p_i}\}_{i=1}^n$, we define the order $n + 1$ tensor $\Phi \in \mathbb{R}^{m \times p_1 \times p_2 \times \dots \times p_n}$ from the m-tensor product as

$$\begin{aligned} [\Phi]_{ki_1 i_2 \dots i_n} &= \left[\bigotimes_{j=1}^n \Psi_j \right]_{ki_1 i_2 \dots i_n}, \\ &= \prod_{j=1}^n [\Psi_j]_{ki_j}. \end{aligned}$$

We show further in this paper how well suited this tensor format is for multilinear regression. We refer to the matrices $\{\Psi_i\}_{i=1}^n$ as *cores* of the *m-tensor* Φ . We introduce the *r-dim* and *c-dim* of the resulting tensor as the dimensions of the *rows* and *columns* of this tensor respectively. In this setting we call *rows* the rank-1 tensors resulting from the row-wise tensor product and *columns* the mode-1 fibers of the tensor. Those are vectors in \mathbb{R}^m . The mode-1 unfolding of this tensor, denoted $\Phi_{(1)}$, results in a dense matrix in $\mathbb{R}^{m \times p_1 p_2 \dots p_n}$. An important note is to be made of the structure of this unfolding. It is *exactly* equal to the face-splitting product of its cores defined using notations introduced by Slyusar [28] as

$$\Phi_{(1)} = \Psi_1 \square \Psi_2 \square \dots \square \Psi_n.$$

This can easily be demonstrated element-wise from both definitions. From this observation, we access properties of an otherwise inaccessible matrix operator $\Phi_{(1)}$ based on tricks allowed by a separated variables tensor representation.

3.2 Transposition

As is the case for arbitrary tensors, the transposition results in a reordering of the tensor's axes. For example with $\Phi \in \mathbb{R}^{m \times p_1 \times p_2 \times \dots \times p_n}$ built as the m-tensor product of cores $\{\Psi_j\}_{j=1}^n$, its transpose is defined by

$$\Phi^\top \in \mathbb{R}^{p_n \times p_{n-1} \times \dots \times p_1 \times m}, \text{ with } [\Phi^\top]_{i_n i_{n-1} \dots i_1 k} = \prod_{j=1}^n [\Psi_j]_{k i_j}.$$

3.3 Contractions

Contractions are operations that consist in summing axes of tensors, reducing their order. Any kind of contraction defined in the general case can be applied to the case of m-tensors, for that definition see appendix section A. We specify here two kinds of contraction for m-tensors: the c-dim contraction and the r-dim contraction. We denote them by $\text{contract}_c(\cdot)$ and $\text{contract}_r(\cdot)$ respectively. Then we have for $\Phi \in \mathbb{R}^{m \times p_1 \times p_2 \times \dots \times p_n}$ the c-dim contraction defined as

$$\text{contract}_c(\Phi) \in \mathbb{R}^{p_1 \times p_2 \times \dots \times p_n}, \quad [\text{contract}_c(\Phi)]_{i_1 i_2 \dots i_n} = \sum_{k=1}^m [\Phi]_{k i_1 i_2 \dots i_n} \quad \forall i_j, 1 \leq i_j \leq p_j.$$

The r-dim contraction is defined as

$$\text{contract}_r(\Phi) \in \mathbb{R}^m, \quad [\text{contract}_r(\Phi)]_k = \sum_{i_1=1}^{p_1} \sum_{i_2=1}^{p_2} \dots \sum_{i_n=1}^{p_n} [\Phi]_{k i_1 i_2 \dots i_n} \quad \forall k, 1 \leq k \leq m.$$

The r-dim contraction acts on rank-1 tensors whose contraction is computed with complexity $\mathcal{O}(np)$ by denoting $\forall i, p_i = p$ for ease of notation, bringing the overall r-dim contraction to a complexity $\mathcal{O}(mnp)$. This property, demonstrated in appendix section A, makes operations involving m-tensor r-dim contractions desirable. In the contrary we avoid c-dim contractions whose complexity is $\mathcal{O}(mp^n)$. We denote $\text{contract}_{r,c}(\cdot)$ the full contraction of the tensor resulting in a scalar.

3.4 Hadamard product

The Hadamard product is a fundamental operation, otherwise known as the element-wise product. The Hadamard product of two m-tensors can be computed in a fully factorized manner, therefore the scaling of its computation will not suffer the curse of dimensionality. Let $\Phi_1, \Phi_2 \in \mathbb{R}^{m \times p_1 \times p_2 \times \dots \times p_n}$ be m-tensors respectively based on cores $\{\Psi_i^1\}_{1 \leq i \leq n}$ and $\{\Psi_i^2\}_{1 \leq i \leq n}$.

$$\Phi_1 = \bigotimes_{i=1}^n \Psi_i^1, \quad \Phi_2 = \bigotimes_{i=1}^n \Psi_i^2. \quad (3.3)$$

Elements of their Hadamard product are defined as

$$\begin{aligned} [\Phi_1 \odot \Phi_2]_{k i_1 i_2 \dots i_n} &= [\Phi_1]_{k i_1 i_2 \dots i_n} [\Phi_2]_{k i_1 i_2 \dots i_n}, \\ &= \prod_{j=1}^n [\Psi_j^1]_{k i_j} [\Psi_j^2]_{k i_j}. \end{aligned}$$

The complexity of the Hadamard product of two m-tensors is $\mathcal{O}(nmp)$ by letting $\forall i, p_i = p$ for ease of notation.

3.5 Inner product and norm

The inner product of two m-tensors leverages the efficiency of the r-dim contraction defined earlier. Let $\Phi_1, \Phi_2 \in \mathbb{R}^{m \times p_1 \times p_2 \times \dots \times p_n}$ be the m-tensors defined in eq. (3.3). Their inner product is defined as

$$\begin{aligned} \langle \Phi_1, \Phi_2 \rangle &= \sum_{k=1}^m \sum_{i_1=1}^{p_1} \sum_{i_2=1}^{p_2} \dots \sum_{i_n=1}^{p_n} [\Phi_1]_{ki_1i_2\dots i_n} [\Phi_2]_{ki_1i_2\dots i_n}, \\ &= \text{contract}_{r,c}(\Phi_1 \odot \Phi_2). \end{aligned}$$

The complexity for this inner product is $\mathcal{O}(mnp)$ considering $\forall i, p_i = p$ for ease of notation. The ℓ_2 norm of an m-tensor inherits the efficiency highlighted above. We define it for $\Phi \in \mathbb{R}^{m \times p_1 \times p_2 \times \dots \times p_n}$ as

$$\begin{aligned} \|\Phi\|_2 &= \sqrt{\langle \Phi, \Phi \rangle}, \\ &= \sqrt{\text{contract}_{r,c}(\Phi^{\odot 2})}. \end{aligned}$$

3.6 N-way contracted product

Operations with tensors allow to define the sums and product depending on the aim of the operation. Here we define a n -way contracted product that we call the m -product and denote \times . Let $\Phi_1 \in \mathbb{R}^{m_1 \times p_1 \times p_2 \times \dots \times p_n}$ and $\Phi_2 \in \mathbb{R}^{m_2 \times p_1 \times p_2 \times \dots \times p_n}$ be m-tensors respectively based on cores $\{\Psi_i^1\}_{1 \leq i \leq n}$ and $\{\Psi_i^2\}_{1 \leq i \leq n}$. Provided that both tensors have the same order and same row dimensions, we can build an element of $\Phi_1 \times (\Phi_2)^\top \in \mathbb{R}^{m_1 \times m_2}$ as the inner product of order $n-1$ rank-1 tensors. The overall computation results in the Hadamard product of matrix products of corresponding cores of the input m-tensors:

$$\Phi_1 \times \Phi_2^\top = \bigodot_{i=1}^n \Psi_i^1 (\Psi_i^2)^\top.$$

Considering, for ease of notation that $m_1 = m_2 = m$ and $\forall i, p_i = p$, the complexity of the m-product is $\mathcal{O}(m^2pn)$. This complexity scales only linearly with dimension n and is dominated by m^2 . The m-product of arbitrary tensors with the same dimensions would result in complexity $\mathcal{O}(m^2p^n)$.

4 M-Tensor least squares regression

In this section, we demonstrate the use of the m-tensor formalism to efficiently compute a multidimensional least squares regression. The aim is to make the tensor least squares (3.2) manageable in arbitrarily high dimensions. Provided the set of sampled parameters $\mathcal{P} = \{\mathbf{x}_k\}_{k=1}^m$ and corresponding evaluations of f in a vector \mathbf{y} , we define the basis functions tensor as

$$\varphi(\mathbf{x}) = \bigotimes_{i=1}^n \psi([\mathbf{x}]_i), \quad \text{with } \psi(x) = [\psi_1(x) \ \psi_2(x) \ \dots \ \psi_p(x)].$$

This structure amounts to combining one-dimensional bases $\psi_{1 \leq i \leq p}$ for all input parameters. Those can be chosen arbitrarily as we will see further in the examples. This gives us a large approximation capability, unreachable with a matrix-based regression. We remark that by stacking the rank-1 tensors φ , we actually build an m-tensor. To that end, we provide the

following definition for the i -th core, *i.e.* the matrix built of one-dimensional basis functions of the i -th components of each sample in \mathcal{P} .

$$\Psi_i = \begin{bmatrix} \psi_1([\mathbf{x}_1]_i) & \psi_2([\mathbf{x}_1]_i) & \cdots & \psi_p([\mathbf{x}_1]_i) \\ \psi_1([\mathbf{x}_2]_i) & \psi_2([\mathbf{x}_2]_i) & \cdots & \psi_p([\mathbf{x}_2]_i) \\ \vdots & \vdots & \ddots & \vdots \\ \psi_1([\mathbf{x}_m]_i) & \psi_2([\mathbf{x}_m]_i) & \cdots & \psi_p([\mathbf{x}_m]_i) \end{bmatrix}.$$

That definition allows us to define the m-tensor regression operator Φ as

$$\Phi = \bigotimes_{i=1}^n \Psi_i. \quad (4.1)$$

We can now reformulate the regression problem in m-tensorial setting. We seek the regression coefficients $\mathbf{C} \in \mathbb{R}^{p \times p \times \cdots \times p}$ such that

$$\Phi \times \mathbf{C} = \mathbf{y}. \quad (4.2)$$

A least squares solution $\hat{\mathbf{C}}$ to this problem can be built explicitly based on the general pseudoinverse formulation eq. (2.4). Provided that Φ has linearly independent rows, we have the solution

$$\hat{\mathbf{C}} = \Phi^\top (\Phi \times \Phi^\top)^{-1} \mathbf{y}. \quad (4.3)$$

We will systematically consider the linearly independent rows case because, in general, the exponential growth of the system with n will prevent it to have linearly independent columns as it would require m to grow exponentially with n . We generally cannot use coefficients formulated as eq. (4.3). The dimensionality of $\hat{\mathbf{C}}$ would require to store p^n elements, which rapidly become unmanageable with growing n . Instead, by leveraging properties of the m-tensor product, we give the approximation of f as

$$\hat{f}(\mathbf{x}) = \varphi(\mathbf{x}) \times \Phi^\top (\Phi \times \Phi^\top)^{-1} \mathbf{y}.$$

This expression purposefully explicits the use of the m-product to show how it results only in our efficient inner products. In fact, by denoting

$$P = \Phi \times \Phi^\top \in \mathbb{R}^{m \times m},$$

and defining

$$\mathbf{z} = P^{-1} \mathbf{y} \in \mathbb{R}^m$$

we give the online computation of \hat{f} requiring only the computation of the vector made from $\varphi(\mathbf{x}) \times \Phi^\top \in \mathbb{R}^m$ and its inner product with \mathbf{z} :

$$\hat{f}(\mathbf{x}) = \langle \varphi(\mathbf{x}) \times \Phi^\top, \mathbf{z} \rangle.$$

This computation results in complexity $\mathcal{O}(mnp)$. We provide a toy example of the m-tensor least squares in appendix section B for ease of comprehension.

remark The solution, although mathematically sound, proves numerically unstable because of the inversion of P . For that reason, we highlight the fact that by construction, P is symmetric positive definite and therefore admits a Cholesky factorization. We give that decomposition as

$$P = LL^\top,$$

with $L \in \mathbb{R}^{m \times m}$ a lower triangular matrix. Computation of \mathbf{z} can then be done through a two step procedure. We first solve for an intermediate result $\bar{\mathbf{z}}$ by upward backsubstitution

$$\bar{\mathbf{z}} = L^{-1}\mathbf{y}.$$

Then \mathbf{z} is solved for by downward backsubstitution with

$$\mathbf{z} = (L^\top)^{-1}\bar{\mathbf{z}}.$$

The use of Cholesky factorization results in the same complexity scaling with a better numerical stability. We take advantage of this in the algorithms we provide.

5 Regularization of m-tensor regression

Mimicking a matrix-based regression while allowing so many combinations of one-dimensional bases to express the problem, which is the greatest strength of our approach, plagues the m-tensor regression with ill-conditioning and a risk of overfitting. Therefore we address the crucial ability to regularize it. Generally the aim of regularization is to enhance the conditioning of a matrix. A matrix Φ with rank r has a condition number κ defined as

$$\kappa(\Phi) = \frac{\sigma_1}{\sigma_r}. \quad (5.1)$$

Here σ_1 denotes the largest singular value of Φ and σ_r its lowest nonzero singular value. A matrix is known to be *ill-conditioned* if κ is very large. Regularized regression will take the form of penalizations added to the least square minimization problem eq. (2.2). Let \mathcal{R} denote a penalization term, the new minimization problem in matrix form is formulated as

$$\hat{\mathbf{c}}_{\mathcal{R}} = \underset{\mathbf{c}}{\operatorname{argmin}} \left(\|\Phi\mathbf{c} - \mathbf{y}\|_2^2 + \mathcal{R}(\mathbf{c}, \boldsymbol{\lambda}) \right). \quad (5.2)$$

The penalization term \mathcal{R} is intended to privilege certain forms of solution in the minimization. It acts based on the terms of \mathbf{c} and a *regularization parameter* (or set of parameters) $\boldsymbol{\lambda}$. We suggest ways to act in a memory efficient manner to generalize the approach to tensor regression. The case of m-tensor regression yields a minimization problem similar to the classical one:

$$\hat{\mathbf{C}}_{\mathcal{R}} = \underset{\mathbf{C}}{\operatorname{argmin}} \left(\|\Phi \times \mathbf{C} - \mathbf{y}\|_2^2 + \mathcal{R}(\mathbf{C}, \boldsymbol{\lambda}) \right). \quad (5.3)$$

5.1 Tikhonov regularization

Matrix-based Tikhonov regularization solves for eq. (5.2) with

$$\mathcal{R}(\mathbf{c}, \lambda) = \lambda\|\mathbf{c}\|_2.$$

This leads to a minimization of the ℓ_2 norm of \mathbf{c} . This minimization problem admits an explicit solution given as

$$\hat{\mathbf{c}}_2 = \Phi^\top (\Phi\Phi^\top + \lambda^2 \mathbb{I})^{-1} \mathbf{y}. \quad (5.4)$$

This elevates the values of the lowest singular values, hence limiting the rise of κ in equation eq. (5.1). In fact the effect of Tikhonov regularization is to level up singular values of Φ below the set threshold.

We generalize this approach to m-tensor regression. Consider the m-tensor formulation (5.3), with regularization term

$$\mathcal{R}(\mathbf{C}, \lambda) = \lambda\|\mathbf{C}\|_2.$$

Analogously to the matrix case eq. (5.4), we can explicitly give the regularized solution as

$$\hat{\mathbf{C}}_2 = \mathbf{\Phi}^\top (\mathbf{\Phi} \times \mathbf{\Phi}^\top + \lambda^2 \mathbb{I}d)^{-1} \mathbf{y}.$$

The approximation of f with Tikhonov regularization is then given by evaluating the inner product

$$\hat{f}_2(\mathbf{x}) = \langle \boldsymbol{\varphi}(\mathbf{x}) \times \mathbf{\Phi}^\top, \mathbf{z}_2 \rangle. \quad (5.5)$$

In this expression \mathbf{z}_2 is defined as

$$\mathbf{z}_2 = (\mathbf{\Phi} \times \mathbf{\Phi}^\top + \lambda^2 \mathbb{I}d)^{-1} \mathbf{y}.$$

This approach's complexity is dominated by the construction of the m-product $\mathbf{\Phi} \times \mathbf{\Phi}^\top$ ($\mathcal{O}(m^2pn)$) and the inversion of $(\mathbf{\Phi} \times \mathbf{\Phi}^\top + \lambda^2 \mathbb{I}d)$. Inference cost, *i.e.* the online evaluation of the approximation given in (5.5), has a complexity $\mathcal{O}(mpn)$.

5.2 Spectral truncation

In the matrix-based regression context, spectral truncation aims at improving the condition number of the regression operator by simply removing (setting to zero) the singular values below a set threshold τ . This is equivalent to solving for eq. (5.2) with a penalization on the nuclear norm of Φ , *i.e.* the sum of its singular values. Let $\Phi = U\Sigma V^\top$ give the SVD of Φ , the least squares regression coefficients are defined as

$$\hat{\mathbf{c}} = V\Sigma^{-1}U^\top \mathbf{y}.$$

The spectral truncation regularization consists in replacing Σ by a matrix D_τ to get the regularized coefficients as

$$\hat{\mathbf{c}}_* = UD_\tau V^\top \mathbf{y} \quad \text{where} \quad [D_\tau]_{ii} = \begin{cases} \sigma_i & \text{if } \sigma_i \geq \tau \\ 0 & \text{else} \end{cases}.$$

This approach generalizes to the m-tensor regression by leveraging the SVD of $P = \mathbf{\Phi} \times \mathbf{\Phi}^\top$. We are able to truncate the rank so as to regularize the regression. We select the r first singular values, deemed sufficient to represent the system by comparing them to the input threshold τ . We denote U_r the matrix built from only the first r left singular vectors of P . The regularized coefficients are then given by

$$\hat{\mathbf{C}}_* = \mathbf{\Phi}^\top U_r U_r^\top P^{-1} \mathbf{y}.$$

We project the least squares regression onto the span only of the r first left singular vectors. This approach is chosen so as to keep using the inner products of rank-1 tensors. The resulting approximation is computed as the inner product

$$\hat{f}_*(\mathbf{x}) = \langle \boldsymbol{\varphi}(\mathbf{x}) \times \mathbf{\Phi}^\top, \mathbf{z}_* \rangle,$$

where \mathbf{z}_* is defined as

$$\mathbf{z}_* = U_r U_r^\top P^{-1} \mathbf{y}.$$

This approach results in a complexity dominated once again by the computation of $\mathbf{\Phi} \times \mathbf{\Phi}^\top$ and its SVD, leading to a higher computational overhead compared to Tikhonov regularization. Inference cost is the same as for the Tikhonov regularization. A randomized version of the SVD may be preferred to speed up its computation if the rank can be approximated as a given input.

5.3 ALI regularization

In this section we consider a regularization technique that aims at enhancing the conditioning of the system as well as improving the efficiency of the approximation. We consider a subset of $\tilde{m} < m$ rows of Φ , corresponding to a subsample $\tilde{\mathcal{P}} \subset \mathcal{P}$ that can approximate all rows of Φ up to a given tolerance ε . The tensor's rows are therefore considered Almost Linearly Dependent (ALD) [8] on this subset. We denote $\tilde{\Phi}$ the tensor made only of those \tilde{m} Almost Linearly Independent (ALI) rows of Φ . The ALD criterion is based on the following distance δ_k evaluation, by denoting $[\Phi]_k$ the k -th row of Φ

$$\forall k, 1 \leq k \leq m, \quad \delta_k = \min_{\mathbf{w} \in \mathbb{R}^{\tilde{m}}} \left\| [\Phi]_k - \mathbf{w}^\top \tilde{\Phi} \right\|_2^2 < \varepsilon.$$

δ_k evaluates the distance of each row $[\Phi]_k$ to the subspace spanned by the rows of $\tilde{\Phi}$. If the row subsample of size \tilde{m} is optimal with respect to the criterion ε , then there exists a weight matrix $W \in \mathbb{R}^{m \times \tilde{m}}$ such that

$$\left\| \Phi - W \tilde{\Phi} \right\|_2^2 < \varepsilon.$$

From that, we derive the ALI decomposition. We use this decomposition as it provides a bound with respect to the ℓ_2 optimal provided by the SVD while preserving the row structure (rank-1 tensors) of the retained *modes*. Error is controllable through the hyperparameter ε .

By reordering the rows of Φ we can restructure the linear system eq. (4.2) to group together rows of the subset $\tilde{\Phi}$ and the remainder that we denote $\bar{\Phi}$. We apply the same separation and reordering in the RHS \mathbf{y} to write

$$\begin{bmatrix} \tilde{\Phi} \\ \bar{\Phi} \end{bmatrix} \times \mathbf{C} = \begin{bmatrix} \tilde{\mathbf{y}} \\ \bar{\mathbf{y}} \end{bmatrix}.$$

We then only solve for the first subset, *i.e.* we seek $\tilde{\mathbf{C}}$ in the minimization problem

$$\tilde{\mathbf{C}} = \underset{\mathbf{C}}{\operatorname{argmin}} \left(\left\| \tilde{\Phi} \times \mathbf{C} - \tilde{\mathbf{y}} \right\|_2^2 \right).$$

We base the approach on the ALD criterion so the solution generalizes to the full system. With this subsystem approach, the ALI regularization is the only one that manages to reduce the size of the system to solve and provide a gain with respect to m along with the linearization of complexity with regard to n .

The subset selection can be done in one of two ways: the optimal or greedy one. The former is developed analogously to what is done by Gelß et al. [9] while the latter is an adaptation of the kernel recursive least squares developed by Engel et al. [8] for signal processing. The approach proposed by Engel et al. is enhanced by incrementally building the Cholesky factor instead of building and inverting the kernel matrix which proves numerically unstable. This is inspired by Baddoo et al. [3]. We demonstrate the efficiency of the greedy approach, which requires a single pass over the rows of Φ to build $\tilde{\Phi}$. We build the approximation the same way we did for previous regularization techniques, first we construct

$$\tilde{P} = \tilde{\Phi} \times \tilde{\Phi}^\top.$$

Then we compute the approximation

$$\tilde{f}(\mathbf{x}) = \langle \varphi(\mathbf{x}) \tilde{\Phi}^\top, \tilde{\mathbf{z}} \rangle,$$

where $\tilde{\mathbf{z}}$ is defined as

$$\tilde{\mathbf{z}} = \tilde{P}^{-1} \tilde{\mathbf{y}}.$$

The computational gain at the online stage can reveal massive compared to the least squares formulation if $\tilde{m} \ll m$. We evaluated the complexity of the least squares approximation inference at $\mathcal{O}(mpn)$. With the ALI regularization, we have a complexity $\mathcal{O}(\tilde{m}pn)$. Even small gains in this formulation outperforms all previous approaches at a small and controllable cost in terms of error.

Indeed Engel et al. gave a bound on the span of the selected rows of the ALI decomposition of Φ with respect to the optimal decomposition in the least squares sense, *i.e.* the SVD and on the mean squared error of the approximation. Without providing the full demonstration for which we redirect the reader to the work of Engel et al., we adapt here one of the bounds we inherit. By defining the projector onto the span of the rows of $\tilde{\Phi}$ as

$$\Pi_{\tilde{\Phi}} = \tilde{\Phi}^\top \left(\tilde{\Phi} \times \tilde{\Phi}^\top \right)^{-1} \tilde{\Phi},$$

we have

$$\frac{1}{m} \sum_{i=1}^m \|(\text{Id} - \Pi_{\tilde{\Phi}}) \times \varphi^\top(\mathbf{x}_i)\|_2^2 \leq \varepsilon.$$

This means that the mean squared error of the approximation of Φ by the ALI decomposition is directly bounded by ε . Full algorithms to obtain the optimal and greedy formulations of the ALI regression operator are given in appendix section C. In the end, the greedy construction results in a complexity bounded by $\mathcal{O}(m^2np)$ by considering $\forall i, p_i = p$ for ease of notations.

remark This decomposition relates to the s-PGD [14] formulation. Let us explicit the form of $\tilde{\mathbf{C}}$ by its elements definition. Let $\{\tilde{\Psi}_i\}_{1 \leq i \leq n}$ denote the cores of the ALI subset of rows of Φ , then

$$\begin{aligned} [\tilde{\mathbf{C}}]_{i_1 i_2 \dots i_n} &= \sum_{k=1}^{\tilde{m}} [\tilde{\mathbf{z}}]_k [\tilde{\Phi}]_{ki_1 i_2 \dots i_n}, \\ &= \sum_{k=1}^{\tilde{m}} [\tilde{\mathbf{z}}]_k \prod_{j=1}^n [\tilde{\Psi}]_{ki_j}. \end{aligned}$$

This formulation makes the *modes* and weights formulation appear just as in the case of the s-PGD. We share the non-orthogonality of the retained rank-1 tensors, in this case the rows of $\tilde{\Phi}$. We keep the same light formulation, especially well suited for inverse problems or surrogate evaluations.

6 Kernel regression interpretation

In this section we make a few remarks on the parallel between our approach and the kernel regression framework. In practice, we manage to always take advantage of the high efficiency of the dot product of rank-1 tensors. We do not need to select a type of kernel or combine from a limited set of kernels, instead we directly work on the mapping with a tensor product structure so as to build implicitly all combinations.

We leverage the tensor product's ability to augment dramatically the row dimensions. Therefore we can always consider the linearly independent rows pseudoinverse of the m-tensor. In that case, we make $P = \Phi \times \Phi^\top$ appear in the pseudoinverse formulation so that only dot products are needed. That way, an element of P is given by the inner product $\langle \varphi(\mathbf{x}_i), \varphi(\mathbf{x}_j) \rangle$, considering the mapping

$$\varphi : \begin{cases} \mathbb{R}^n & \longrightarrow \mathbb{R}^{p \times p \times \dots \times p} \\ \mathbf{x} & \longmapsto \bigotimes_{i=1}^n \psi([\mathbf{x}]_i) \end{cases},$$

we can derive a kernel function k defined as the inner product

$$k : \begin{cases} \mathbb{R}^n \times \mathbb{R}^n & \longrightarrow \mathbb{R} \\ (\mathbf{x}, \mathbf{x}') & \longmapsto \langle \boldsymbol{\varphi}(\mathbf{x}), \boldsymbol{\varphi}(\mathbf{x}') \rangle \end{cases} .$$

Therefore we have the insurance to build a kernel that verifies Mercer's theorem [22]. There is a direct equivalence between the matrix $P = \Phi \times \Phi^\top$ and the kernel matrix K defined with entries

$$[K]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j).$$

Usual kernel approaches avoid defining the mapping, whereas we manage to manipulate it in combination with the rank-1 tensor dot product to remain efficient with flexibility in the kernel definition. The efficiency of the rank-1 tensor dot product is a direct equivalence of the so-called *kernel trick* to our case.

We rely on the prior works of Schölkopf et al. [27, 26] to derive properties of kernel methods and apply them to this tensor-based approach. Noticeably we were able to obtain an m-tensor PCA from the well-known kPCA. From those observations we managed to derive powerful methods, making the use of m-tensors a salient tool to solve high dimensional regression problems.

7 Numerical examples

In this section we present several numerical example of use for the m-tensor regression framework. The aim is to demonstrate how easy m-tensors are to use and their efficiency to solve high dimensional problems. All those examples are available on GitHub² along with our m-tensor toolbox. For a toy example explaining all stages on a least squares problem, we develop a polynomial toy example in appendix section B. In what follows we consider the relative error, defined by denoting \mathbf{x} and $\hat{\mathbf{x}}$ the real system and its approximation respectively, as

$$E_2 = \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_2}{\|\mathbf{x}\|_2}.$$

7.1 Rosenbrock function approximation

We first demonstrate the scaling in complexity of the approach. To that end we build a regression of the Rosenbrock function generalized from 20 input parameters to 300: $n \in \{20, 50, 100, 200, 300\}$. We chose the Rosenbrock function for ease of implementation and its arbitrary definition with respect to the input's dimension. The Rosenbrock function is usually defined over the domain $\mathcal{D} = [-5, 10]^{\otimes n}$ in dimension n as

$$f(\mathbf{x}) = \sum_{i=1}^{n-1} 100 ([\mathbf{x}]_{i+1} - [\mathbf{x}]_i^2)^2 + ([\mathbf{x}]_i - 1)^2.$$

We explicit the tensor form and chose a polynomial mapping up to degree 4, defined as

$$\boldsymbol{\varphi}(\mathbf{x}) = \bigotimes_{i=1}^n [1 \quad [\mathbf{x}]_i \quad [\mathbf{x}]_i^2 \quad [\mathbf{x}]_i^3 \quad [\mathbf{x}]_i^4].$$

Which results in seeking an order n coefficient tensor $\mathbf{C} \in \mathbb{R}^{5 \times 5 \times \dots \times 5}$ so that the approximation is given as a n -way contracted product:

$$\hat{f}(\mathbf{x}) = \boldsymbol{\varphi}(\mathbf{x}) \times \mathbf{C}.$$

²GitHub repository with all examples: <https://github.com/RemiC-OV/MTensor>.

We consider that for a dimension n , we are given a Latin Hypercube Sampling (LHS) of m points over the definition domain \mathcal{D} . We will consider a linear scaling of the sampling with respect to the dimension to demonstrate the efficiency of the approach in massively underdetermined cases, *i.e.* $m = \alpha n$. We will evaluate the efficiency and accuracy of the approach for varying $\alpha \in \{5, 10, 20, 30, 40\}$. We denote $\mathcal{P} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$ the LHS and \mathbf{y} the evaluation of the Rosenbrock function on that sampling:

$$\forall k, 1 \leq k \leq m, \quad [\mathbf{y}]_k = f(\mathbf{x}_k).$$

We build the m-tensor from the definition of φ and the given samples \mathcal{P} . To do so, we simply evaluate the *core* matrices. In this case, those are Vandermonde matrices given for each input variable

$$\forall i, 1 \leq i \leq n, \quad \Psi_i = \begin{bmatrix} 1 & [\mathbf{x}_1]_i & [\mathbf{x}_1]_i^2 & [\mathbf{x}_1]_i^3 & [\mathbf{x}_1]_i^4 \\ 1 & [\mathbf{x}_2]_i & [\mathbf{x}_2]_i^2 & [\mathbf{x}_2]_i^3 & [\mathbf{x}_2]_i^4 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & [\mathbf{x}_m]_i & [\mathbf{x}_m]_i^2 & [\mathbf{x}_m]_i^3 & [\mathbf{x}_m]_i^4 \end{bmatrix}.$$

Given those matrices, the m-tensor regression operator is fully defined and we can readily build our approximation. We denote the m-tensor Φ defined as

$$\Phi = \bigotimes_{i=1}^n \Psi_i.$$

We plot surface results for models built for all combinations of dimension $n \in \{20, 50, 100, 200, 300\}$ and sample size per axis $\alpha \in \{5, 10, 20, 30, 40\}$. All models are built on a LHS, therefore we build them 3 times and plot an average to account for representativity issues of the sampling. There is no significant variability observed throughout samplings. The aim is to make complexity tendencies appear clearly. Figure 1 presents the *offline* construction cost of the models, it is evaluated in seconds without parallelization. Those costs are assumed only once. They demonstrate in practice the theoretical complexity scalings given for $m = \alpha n$ in table 1. The incremental size of the ALI regression system at each step of the algorithm makes it hard to exactly estimate the degree of the scaling on n but we can state that the upper bound provided in table 1 is far from being reached in the studied domain. The higher computational overhead observed for the spectral truncation is explained by the need for both SVD and inversion steps, both scaling as $\mathcal{O}(m^3)$.

Table 1: Theoretical complexity evaluations for construction and inference costs for different regressions.

	Least squares	Spect. Trunc.	ALI Reg.
Construction	$\mathcal{O}(\alpha^3 n^3)$	$\mathcal{O}(\alpha^3 n^3)$	$\mathcal{O}(\alpha^2 n^3)$ (up. bound)
Inference	$\mathcal{O}(\alpha n^2)$	$\mathcal{O}(\alpha n^2)$	$\mathcal{O}(\tilde{m}n)$

The observed quantity of figure 2 is the *unit inference time*. It corresponds to the average time to compute a solution using the approximation we built. Results are plotted for the least squares model and the ALI regression model. The inference time for the spectral truncation model is the same as for the least squares. Those plots demonstrate the ability of the ALI regression model to fully linearize the cost with respect to the dimension.

Figure 3 presents the relative ℓ_2 error of our models on a LHS of $3m$ unseen points. We notice that errors are worse for lower dimensions, where hyperparameters are harder to chose

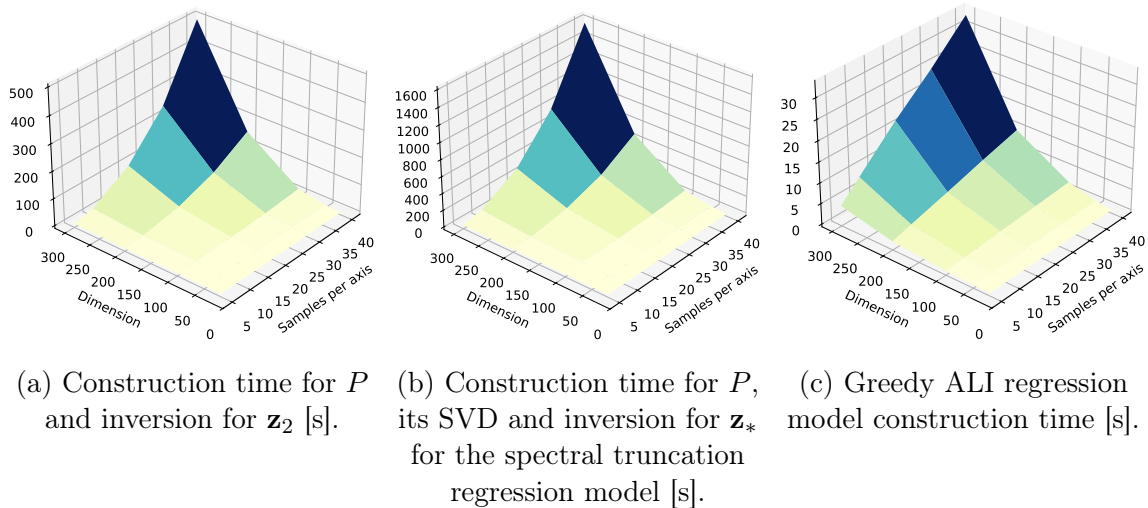


Figure 1: Evolution of the model construction times [s] for the least squares, spectral truncation and ALI regularization models in function of the number of samples per axis and the input dimension.

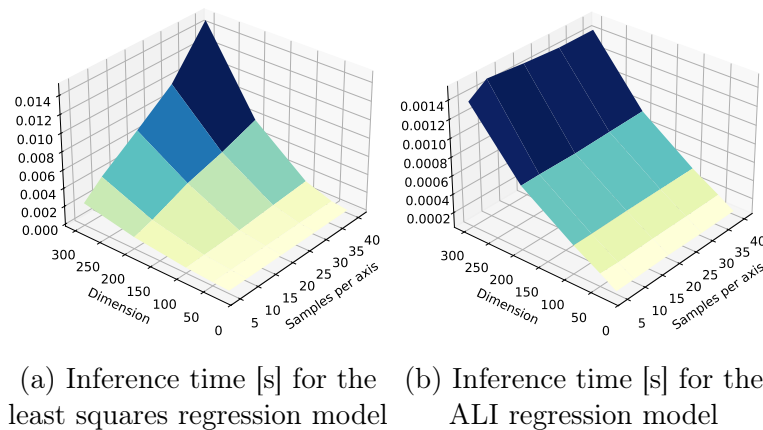


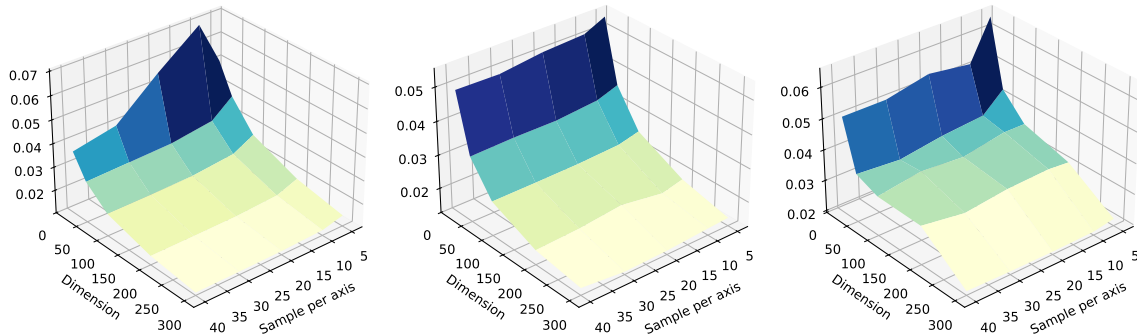
Figure 2: Evolution of the unit inference time for the least squares and ALI regularization models in function of the number of samples per axis and the input dimension.

and models appear more sensitive. In higher dimensions, all models perform at about 2% error. The low convergence rate of the error with the number of training samples is inherited by the linear scaling of the sampling whereas the *ideal sampling* (interpolative) would be scaling exponentially. Therefore we suspect that our plots must show an extremely localized point on a larger convergence curve that would require too many points to build but starts to shape out at dimension $n = 20$.

7.2 Prediction of dynamical systems in time

In the context of dynamical systems, regression is used to predict the evolution in time of the system. If the state is denoted $\mathbf{x}(t) \in \mathbb{R}^n$, then there exists a law f that describes the variation of the state in time. This results in the general form

$$\frac{\partial \mathbf{x}}{\partial t} = f(\mathbf{x}).$$



(a) Relative ℓ_2 error for the least squares regression. (b) Relative ℓ_2 error for the spectral truncation regression. (c) Relative ℓ_2 error for the ALI regression model.

Figure 3: Evolution of the relative ℓ_2 error for the least squares, spectral truncation and ALI regularization models on a LHS of unseen points in function of the number of samples per axis and the input dimension.

The function f might be nonlinear. The aim is to approximate at best this function in order to be able to predict future time increments. We propose to demonstrate the efficiency of the m-tensor regression to identify the dynamics. For small errors in the identified model, the approximation might diverge extremely fast from the real system, therefore this represents a good way to demonstrate the robustness of our approach.

7.2.1 3D Lorenz attractor

In a first example, we describe the approach in a didactic way and demonstrate the robustness of our method. Therefore, we build the regression of a 3D Lorenz attractor [20]. In this case the model is overdetermined, it is presented to show the implementation approach and its simplicity. The models provided by the different regularization proposed before are then used to integrate the system in time from random initial conditions.

The model is 3D, *i.e.* $\mathbf{x} \in \mathbb{R}^3$ and f is defined as

$$\begin{cases} \frac{\partial[\mathbf{x}]_1}{\partial t} &= \sigma([\mathbf{x}]_2 - [\mathbf{x}]_1) \\ \frac{\partial[\mathbf{x}]_2}{\partial t} &= \rho[\mathbf{x}]_1 - [\mathbf{x}]_2 - [\mathbf{x}]_1[\mathbf{x}]_3 \\ \frac{\partial[\mathbf{x}]_3}{\partial t} &= [\mathbf{x}]_1[\mathbf{x}]_2 - \beta[\mathbf{x}]_3 \end{cases} .$$

The parameters are chose as $\sigma = 10$, $\rho = 28$ and $\beta = 2.667$. The system is initialized as $\mathbf{x}(t=0) = [0 \ 1 \ 1.05]^\top$ and solved in time with the scheme

$$\mathbf{x}(t + \Delta t) = \mathbf{x}(t) + f(\mathbf{x}(t))\Delta t.$$

We chose a timestep $\Delta t = 0.001$ and evaluate the behaviour of the system over $m = 500$ steps. To apply the m-tensor methodology, we chose to define the mapping $\boldsymbol{\psi}(x) = [1 \ x]$ which allows us to build the m-tensor $\boldsymbol{\Phi}$ as

$$\boldsymbol{\Phi} = \bigotimes_{i=1}^3 \Psi_i$$

where we define row j of core i , denoted by $[\Psi_i]_j$, as

$$\forall i, 1 \leq i \leq 3, \forall j, 1 \leq j \leq m, [\Psi_i]_j = \boldsymbol{\psi}([\mathbf{x}]_i).$$

We then seek to predict the derivative $\frac{\partial \mathbf{x}}{\partial t}$ from the current state \mathbf{x} . For that reason, we build the right hand side of the system to solve as the matrix

$$\forall i, 1 \leq i \leq 3, \forall j, 1 \leq j \leq m, [Y]_{ji} = \frac{\partial [\mathbf{x}_j]_i}{\partial t}.$$

The resolution then follows the approaches presented before to give an approximation of f . This implicitly solves for

$$\Phi \times \mathbf{C} = Y$$

by building the approximation of f as

$$\hat{f}(\mathbf{x}) = \langle \varphi(\mathbf{x}) \times \Phi, Z \rangle,$$

where we define the matrix Z based on the regularization we want to apply and $\varphi(\mathbf{x}) = \bigotimes_{i=1}^3 \psi([\mathbf{x}]_i)$. We present our results for the least squares, spectral truncation (with rank 8) and ALI regressions (retaining only 8 rows of Φ , it means that for ALI regression, $Z \in \mathbb{R}^{8 \times 3}$).

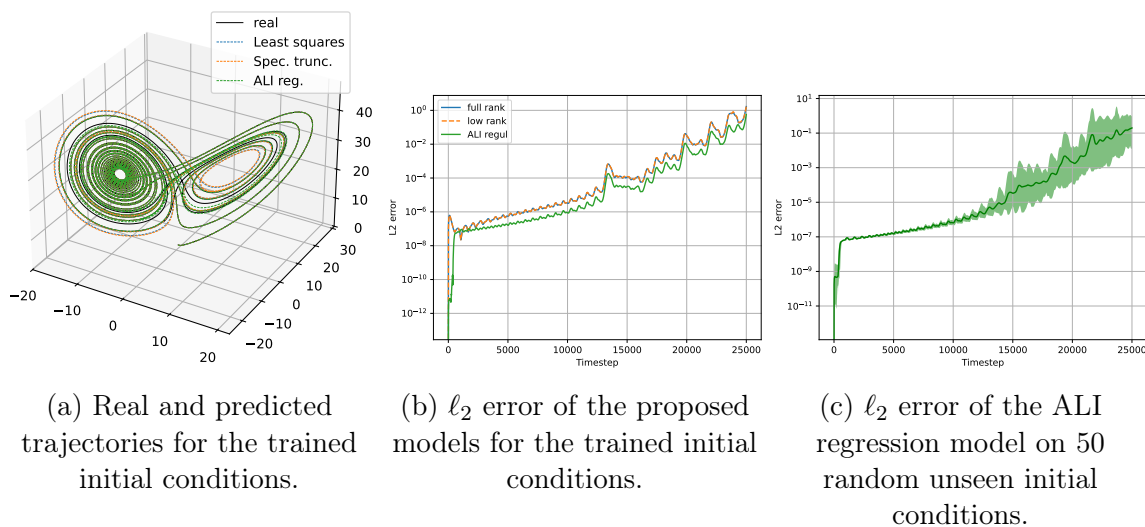


Figure 4: Evaluation of proposed regression models of a Lorenz attractor from a training on 500 timesteps on 25000 timesteps.

We show the results of our method in figure 4. Figure 4a shows the trajectories followed by all three models and the real one on the trained initial conditions for 25000 timesteps, *i.e.* 50 times the trained period. Figure 4b presents the ℓ_2 error over the same span, still on the trained initial conditions. We also test the ALI regression on 50 randomly selected initial conditions and plot the ℓ_2 error of the trajectories over a span of 25000 timesteps to demonstrate generalizability of the learnt model. This is the plot of figure 4c. The only way to make these models more robust would be to eliminate fully the coefficients close to zero that ultimately add up to make the model move away from the real trajectory. This translates into enforcing a ℓ_0 norm minimization which is an ongoing topic of research for this dimensionality of problems.

7.2.2 Kuramoto oscillators

The next example demonstrates the ability of the proposed tensor regression framework to scale well with dimension with only a few data points as input. We therefore take interest in the Kuramoto oscillator model [18]. The model is defined for an arbitrary number n of oscillators, all coupled through the equations

$$\forall i, 1 \leq i \leq n, \quad \frac{\partial [\mathbf{x}]_i}{\partial t} = \omega_i + \frac{K}{n} \sum_{j=1}^n \sin([\mathbf{x}]_j - [\mathbf{x}]_i).$$

The factor K is a coupling constant chosen here to be shared by all oscillators, each described by their phase $[\mathbf{x}]_{1 \leq i \leq n}$ (the state we estimate) and their natural frequency $\omega_{1 \leq i \leq n}$. We define the natural frequencies randomly in the interval $[-5, 5]$ and set the coupling factor $K = 2$.

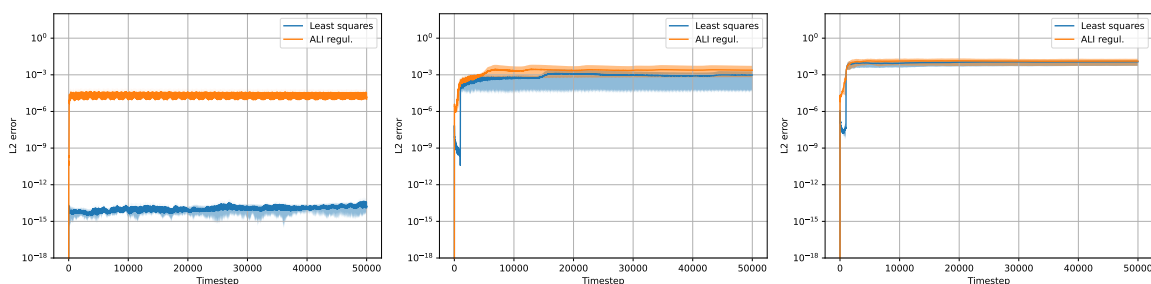
We build our model by including bases $\boldsymbol{\psi}_1$ and $\boldsymbol{\psi}_2$ defined for each variable as

$$\boldsymbol{\psi}_1(x) = [1 \quad \sin(x)] \quad \text{and} \quad \boldsymbol{\psi}_2(x) = [1 \quad \cos(x)].$$

The resulting m-tensor is therefore built as a larger m-tensor product of two bases by variable. By denoting the cores built on each input for both basis $\Psi_{1 \leq i \leq n}^1$ and $\Psi_{1 \leq i \leq n}^2$, the m-tensor of the regression problem is built as

$$\Phi = \left(\bigotimes_{i=1}^n \Psi_i^1 \right) \otimes \left(\bigotimes_{i=1}^n \Psi_i^2 \right).$$

The ability of our method to produce robust models is tested by building for dimensions $n \in \{3, 10, 100\}$ a model, 5 times each to evaluate average results. Each model is built with different random natural frequencies and initial conditions. We plot the ℓ_2 errors for predicted phases in time over 50000 timesteps in figure 5 for $n = 3, 10, 100$ respectively in figures 5a, 5b and 5c. Results are based on the same training period of 1000 timesteps and then run using the



(a) Relative error of the regression models of Kuramoto oscillators for $n = 3$.

(b) Relative error of the regression models of Kuramoto oscillators for $n = 10$.

(c) Relative error of the regression models of Kuramoto oscillators for $n = 100$.

Figure 5: Relative errors of the least squares and ALI regression models for Kuramoto oscillators for $n \in \{3, 10, 100\}$ over 50000 timesteps from 1000 seen timesteps.

same initial conditions as for training. Area between min and max ℓ_2 relative error are plotted around averaged results over the 5 run. We can see that ALI regression generally performs worse but systematically runs faster by retaining only 65, 360 and 370 (rounded averages) rows of Φ for $n = 3, 10, 100$ respectively. The growing error through dimensions is explained by the fact that the amount of given data does not evolve. Models trained over 1000 timesteps in the case $n = 100$ are massively underdetermined and still perform at about 2% error.

7.3 A practical note

A detail left until now in the implementation is that of numerical overflow. If we use the kernel trick to avoid computing all elements of the tensors at hand, it is strictly equivalent in practice. This means that working with a problem involving a dot product between two rank-1 tensors of order n and dimension (p, p, \dots, p) will effectively lead to a result that multiplies and sums p^n elements together. This leads to results that are unmanageable numerically. Therefore the one-dimensional mappings $\boldsymbol{\psi}$ are built with a hyperparameter that scales the input down in high-dimensional contexts.

This hyperparameter is inherent with the method and a right tuning provides massive gains in term of precision. As it results purely from numerical issues, *i.e.* the mathematical background is sound otherwise, there must exist a normalization dependent on the input to be systematically applied, however this normalization has not yet been determined. In our examples, bases in low dimension (<10) are not scaled down, but higher dimensions are scaled with a factor 10^{-7} .

The real significant choice that remains is that of one-dimensional bases ψ . This choice of basis function can be made different for each axis depending on the problem at hand but the first choice we have systematically made was to start with

$$\psi(x) = [1 \ x].$$

This choice, although simple builds all combinations of parameters at the lowest computational cost. If more precision is sought, the degree can be adapted incrementally, seeing how low the cost is to compute a solution. Cyclic behaviour can be represented using a sinusoidal polynomial basis or Fourier basis. A solution left unexplored in this paper is the choice of local bases along each axis. One may *discretize* each axis with local bases, even with compact support. The resulting tensor regression operator will capture localized effects in the n -dimensional mesh that is implicitly built. Moreover cores would be sparse, providing a potential further speedup to the method.

All examples presented here are publicly available from data generation to result plotting. The implementation uses a toolbox leveraging the previously presented operations based on m-tensors. Curious readers can access this code at: <https://github.com/RemiC-OV/MTensor>. The examples and tools presented were developed fully in a NumPy implementation, it is therefore simplistic and can be optimized. We leveraged algebraic structures to gain in memory use and computational efficiency but we did not implement any form of parallelism. It is one of the main development perspectives to demonstrate the gains possible. Most operations presented in this paper are parallelizable throughout dimensions and/or samples.

8 Discussion

We briefly discuss the advantages and drawbacks of the proposed approach. We chose to compare the performances of our formalism to that of kernel regression, TT-based approaches and s-PGD as we consider those to readily compare.

We consider our framework to fall into the category of kernel regression methods as we have seen in section 6. The way we enrich the area of kernel regression is by efficiently reaching all combinations of a choice of one-dimensional functions. We suffer the same limitations as kernel methods *i.e.* the necessity to select the bases to work with and poor conditioning, to which we answer by several regularization approaches.

Our formalism differs from that of tensor trains by specialising our tools to regression. Seeing the works of Gelß et al. with MANDy, one can see the direct equivalence of m-tensors to the diagonal slice structured cores. They gained efficiency in high dimensions by switching to kFSA which resorts to the kernel trick through inner products of tensor trains instead of the sequential computation of the TT-SVD and pseudoinverse. By scrapping the need for generalized algebra offered by tensor trains, we avoid storing and operating on a massive amount of zeros as is done in the TT equivalent to our method. It follows that the m-tensor formalism is not fit for as much generalization as the tensor trains. This might prove limiting in more complex applications of the m-tensor formalism.

We view our approach as a generalization of the PGD framework and the sparse-PGD for the case of scarce data. We avoid seeking greedily the best rank-1 tensor to add to the sought

coefficients and instead find without iterative scheme the least squares optimal coefficients. We avoid the ALS solution and gain greatly in efficiency at the solving stage. However the m-tensor regression results in coefficients that are not readily manageable as would be in the case of the PGD that may store them as a sum of few rank-one tensors. A key advantage for future works that is inherited from the structure of s-PGD is the ability to compute inverse problems [25]. The separated variables structure allows to evaluate derivatives with respect to each parameter by only computing only the derivatives of one-dimensional bases.

We did not explore comparisons with any Tucker-based formalism. Even though the Tucker format is a natural choice, it still inherently suffers the curse of dimensionality through the storage of its core. In fact few approaches are found that use the Tucker format above dimension 3. This is a huge limitation that requires tricks that make working with them in higher dimensions more conceptually complex. For instance one may decompose the Tucker core as a tensor train but we argue that it does not provide a straightforward understanding of the relationship to the choice of bases.

9 Conclusion and perspectives

In this paper we presented a tensor formalism for high dimensional regression of nonlinear models. The introduced m-tensor formalism is tailored for high dimensional regression in scarce data contexts. By leveraging tensor algebra, we provided properties from kernel methods with good scalability. The formalism smoothly generalizes well-known properties of matrix-based approaches to provide optimality to tensor structured models. We generalized different regularization techniques that seek to make models robust and generalizable from the low amount of data available. Those regularization techniques prove robust and yield low errors while limiting the scaling when faced with a high number of parameters. The proposed ALI regularization technique allows for further gains in computational and memory costs, achieving linear scaling at online stage on the provided examples.

The m-tensor formalism was shown to break the curse of dimensionality with, in the worst case, cubic scaling on the number of parameters. Linear scaling of memory use with the dimension of the problem and parallel capability are inherited from the separated variables approach.

The proposed approach was applied on a general benchmark to demonstrate scaling and low relative errors. Robustness was demonstrated by applying the different regularization techniques to identify the nonlinear dynamical behaviour of two systems. In the first case, we show good generalizability for a Lorenz chaotic system. In the second one we show again both generalizability and scalability up to 100 parameters with limited data on Kuramoto oscillators.

In the future, we will seek different developments, most probably stemming from randomized numerical linear algebra to further accelerate our tools and enhance their scaling. The same approach might lead us to ℓ_0 norm minimization. We will seek to expand the validity domain by exploiting rational forms of approximation. On the application side, the tool seems most adapted to model order reduction, data-driven operator inference, large-scale optimization or classification. Generalization of the m-tensor formalism might prove useful also for solving parametric PDEs, at this stage most likely through a collocation method.

Acknowledgements

We acknowledge fundings from the ENSAM RTE Chair, insightful discussions with Antonio Falcò and rich and fruitful exchanges with Jad Mounayer.

Bibliography

- [1] A. Ammar et al. “A new family of solvers for some classes of multidimensional partial differential equations encountered in kinetic theory modeling of complex fluids”. In: *Journal of Non-Newtonian Fluid Mechanics* 139.3 (2006), pp. 153–176. ISSN: 0377-0257. DOI: [10.1016/j.jnnfm.2006.07.007](https://doi.org/10.1016/j.jnnfm.2006.07.007).
- [2] Amine Ammar, Francisco Chinesta, and Antonio Falcò. “On the Convergence of a Greedy Rank-One Update Algorithm for a Class of Linear Systems”. In: *Arch Computat Methods Eng* 17 (2010), pp. 473–486. DOI: [10.1007/s11831-010-9048-z](https://doi.org/10.1007/s11831-010-9048-z).
- [3] Peter J. Baddoo et al. “Kernel learning for robust dynamic mode decomposition: linear and nonlinear disambiguation optimization”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 478.2260 (2022), p. 20210830. DOI: [10.1098/rspa.2021.0830](https://doi.org/10.1098/rspa.2021.0830).
- [4] R. Bellman, Rand Corporation, and Karreman Mathematics Research Collection. *Dynamic Programming*. Rand Corporation research study. Princeton University Press, 1957. ISBN: 9780691079516.
- [5] Gregory Beylkin and Martin J. Mohlenkamp. “Algorithms for Numerical Analysis in High Dimensions”. In: *SIAM Journal on Scientific Computing* 26.6 (2005), pp. 2133–2159. DOI: [10.1137/040604959](https://doi.org/10.1137/040604959).
- [6] Domenico Borzacchiello, José V. Aguado, and Francisco Chinesta. “Non-intrusive Sparse Subspace Learning for Parametrized Problems”. In: *Archives of Computational Methods in Engineering* 26 (2019), pp. 303–326. ISSN: 1886-1784. DOI: [10.1007/s11831-017-9241-4](https://doi.org/10.1007/s11831-017-9241-4).
- [7] Steven L. Brunton and J. Nathan Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. 2nd ed. Cambridge University Press, 2022. DOI: [10.1017/9781009089517](https://doi.org/10.1017/9781009089517).
- [8] Y. Engel, S. Mannor, and R. Meir. “The kernel recursive least-squares algorithm”. In: *IEEE Transactions on Signal Processing* 52.8 (2004), pp. 2275–2285. DOI: [10.1109/TSP.2004.830985](https://doi.org/10.1109/TSP.2004.830985).
- [9] Patrick Gelß et al. “Feature space approximation for kernel-based supervised learning”. In: *Knowledge-Based Systems* 221 (2021), p. 106935. ISSN: 0950-7051. DOI: [10.1016/j.knsys.2021.106935](https://doi.org/10.1016/j.knsys.2021.106935).
- [10] Patrick Gelß et al. “Multidimensional Approximation of Nonlinear Dynamical Systems”. In: *Journal of Computational and Nonlinear Dynamics* 14.6 (Apr. 2019), p. 061006. ISSN: 1555-1415. DOI: [10.1115/1.4043148](https://doi.org/10.1115/1.4043148).
- [11] S.A. Goreinov, E.E. Tyrtyshnikov, and N.L. Zamarashkin. “A theory of pseudoskeleton approximations”. In: *Linear Algebra and its Applications* 261.1 (1997), pp. 1–21. ISSN: 0024-3795. DOI: [10.1016/S0024-3795\(96\)00301-1](https://doi.org/10.1016/S0024-3795(96)00301-1).
- [12] Tamara G. Kolda Grey Ballard. *Tensor Decompositions for Data Science*. Cambridge University Press, 2025. ISBN: 9781009471671.
- [13] Wolfgang Hackbusch. *Tensor Spaces and Numerical Tensor Calculus*. 2nd ed. Vol. 42. Springer Cham, 2019. ISBN: 978-3-030-35553-1. DOI: [10.1007/978-3-030-35554-8](https://doi.org/10.1007/978-3-030-35554-8).
- [14] Rubén Ibáñez et al. “A Multidimensional Data-Driven Sparse Identification Technique: The Sparse Proper Generalized Decomposition”. In: *Complexity* 2018 (2018), pp. 1–11. DOI: [10.1155/2018/5608286](https://doi.org/10.1155/2018/5608286).

- [15] Kenneth L. Judd et al. “Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain”. In: *Journal of Economic Dynamics and Control* 44.C (2014), pp. 92–123. DOI: [10.1016/j.jedc.2014.03.00](https://doi.org/10.1016/j.jedc.2014.03.00).
- [16] Stefan Klus and Patrick Gelß. “Tensor-Based Algorithms for Image Classification”. In: *Algorithms* 12.11 (2019). ISSN: 1999-4893. DOI: [10.3390/a12110240](https://doi.org/10.3390/a12110240).
- [17] Tamara G. Kolda and Brett W. Bader. “Tensor Decompositions and Applications”. In: *SIAM Review* 51.3 (2009), pp. 455–500. DOI: [10.1137/07070111X](https://doi.org/10.1137/07070111X).
- [18] Yoshiki Kuramoto. “Self-entrainment of a population of coupled non-linear oscillators”. In: *International Symposium on Mathematical Problems in Theoretical Physics*. Ed. by Huzihiro Araki. Berlin, Heidelberg: Springer Berlin Heidelberg, 1975, pp. 420–422. ISBN: 978-3-540-37509-8.
- [19] Charles F. Van Loan. “The ubiquitous Kronecker product”. In: *Journal of Computational and Applied Mathematics* 123.1 (2000). Numerical Analysis 2000. Vol. III: Linear Algebra, pp. 85–100. ISSN: 0377-0427. DOI: [10.1016/S0377-0427\(00\)00393-9](https://doi.org/10.1016/S0377-0427(00)00393-9).
- [20] Edward N. Lorenz. “Deterministic Nonperiodic Flow”. In: *Journal of Atmospheric Sciences* 20.2 (1963), pp. 130–141. DOI: [10.1175/1520-0469\(1963\)020<0130:DNF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2).
- [21] Michael W. Mahoney and Petros Drineas. “CUR matrix decompositions for improved data analysis”. In: *Proceedings of the National Academy of Sciences* 106.3 (2009), pp. 697–702. DOI: [10.1073/pnas.0803205106](https://doi.org/10.1073/pnas.0803205106).
- [22] James Mercer. “XVI. Functions of positive and negative type, and their connection the theory of integral equations”. In: *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 209.441–458 (1909), pp. 415–446. DOI: [10.1098/rsta.1909.0016](https://doi.org/10.1098/rsta.1909.0016).
- [23] I. V. Oseledets. “Tensor-Train Decomposition”. In: *SIAM Journal on Scientific Computing* 33.5 (2011), pp. 2295–2317. DOI: [10.1137/090752286](https://doi.org/10.1137/090752286).
- [24] R. Penrose. “A generalized inverse for matrices”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 51.3 (1955), pp. 406–413. DOI: [10.1017/S0305004100030401](https://doi.org/10.1017/S0305004100030401).
- [25] Sebastian Rodriguez et al. “Hybrid twin of RTM process at the scarce data limit”. In: *International Journal of Material Forming* 16.4 (2023), p. 40. ISSN: 1960-6214. DOI: [10.1007/s12289-023-01747-2](https://doi.org/10.1007/s12289-023-01747-2).
- [26] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. “Nonlinear Component Analysis as a Kernel Eigenvalue Problem”. In: *Neural Computation* 10.5 (July 1998), pp. 1299–1319. ISSN: 0899-7667. DOI: [10.1162/089976698300017467](https://doi.org/10.1162/089976698300017467).
- [27] Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, Dec. 2001. ISBN: 9780262256933. DOI: [10.7551/mitpress/4175.001.0001](https://doi.org/10.7551/mitpress/4175.001.0001).
- [28] V.I. Slyusar. “A family of face products of matrices and its properties.” In: *Cybern Syst Anal* (1999), pp. 379–384. DOI: [10.1007/BF02733426](https://doi.org/10.1007/BF02733426).

A Tensor operations complexity analysis

In this section we demonstrate the low complexity of the tensor operations previously introduced. Those are what makes the proposed tensor formalism efficient and constitute the main tricks we leverage. In practice those mostly rely on the low computational complexity of rank-1 tensor operators.

In what follows, we consider that for a given matrix A or m-tensor \mathbf{T} , $[A]_k$ and $[\mathbf{T}]_k$ respectively denotes the k -th row (vector) of A and k -th row (rank-1 tensor) of \mathbf{T} . We defined in section 3 two specific tensor contractions, in this section we define a more general operator. Let $\mathbf{T} \in \mathbb{R}^{p_1 \times p_2 \times \dots \times p_n}$ denote an arbitrary order n tensor. Then the contraction operations are defined as summations over one or more dimensions of the tensor. We denote $\text{contract}_{k_1, k_2, \dots, k_m}(\cdot)$ the operator computing the contraction of a tensor over axes k_1, k_2, \dots, k_m . Elements of the resulting tensor are defined as

$$\begin{aligned} [\text{contract}_{k_1, k_2, \dots, k_m}(\mathbf{T})]_{i_1 i_2 \dots i_{k_1-1} i_{k_1+1} \dots i_{k_2-1} i_{k_2+1} \dots i_{k_m-1} i_{k_m+1} \dots i_n} = \\ \sum_{j_1=1}^{p_{k_1}} \sum_{j_2=1}^{p_{k_2}} \dots \sum_{j_m=1}^{p_{k_m}} [\mathbf{T}]_{i_1 i_2 \dots i_{k_1-1} j_1 i_{k_1+1} \dots i_{k_2-1} j_2 i_{k_2+1} \dots i_{k_m-1} j_m i_{k_m+1} \dots i_n} \\ \forall i_\kappa, 1 \leq i_\kappa \leq p_\kappa, \forall \kappa, 1 \leq \kappa \leq n, \kappa \notin \{k_1, k_2, \dots, k_m\}. \end{aligned}$$

And $\text{contract}_{k_1, k_2, \dots, k_m}(\mathbf{T}) \in \mathbb{R}^{p_1 \times p_2 \times \dots \times p_{k_1-1} \times 1 \times p_{k_1+1} \times \dots \times p_{k_2-1} \times 1 \times p_{k_2+1} \times \dots \times p_{k_m-1} \times 1 \times p_{k_m+1} \times \dots \times p_n}$, the contracted dimension are reduced to 1.

A.1 Rank-1 tensor contraction

We highlight the fact that the contraction of a rank-1 tensor consists in the product of the contractions over each axes on their own. Let $\mathbf{T} = \bigotimes_{i=1}^n \mathbf{u}_i$ with $\mathbf{u}_i \in \mathbb{R}^{p_i}$, then its contraction is computed as

$$\begin{aligned} \text{contract}_{1,2,\dots,n}(\mathbf{T}) &= \text{contract}_{1,2,\dots,n} \left(\bigotimes_{i=1}^n \mathbf{u}_i \right), \\ &= \sum_{j_1=1}^{p_1} \sum_{j_2=1}^{p_2} \dots \sum_{j_n=1}^{p_n} \left[\bigotimes_{i=1}^n \mathbf{u}_i \right]_{j_1 j_2 \dots j_n}, \\ &= \sum_{j_1=1}^{p_1} \sum_{j_2=1}^{p_2} \dots \sum_{j_n=1}^{p_n} \prod_{i=1}^n [\mathbf{u}_i]_{j_i}, \\ &= \sum_{j_1=1}^{p_1} [\mathbf{u}_1]_{j_1} \sum_{j_2=1}^{p_2} [\mathbf{u}_2]_{j_2} \dots \sum_{j_n=1}^{p_n} [\mathbf{u}_n]_{j_n}, \\ &= \prod_{i=1}^n \text{contract}_1(\mathbf{u}_i). \end{aligned}$$

If we consider $p_i = p_j = p$ for ease of notation, this results in a complexity of $\mathcal{O}(np)$.

A.2 Rank-1 tensors Hadamard product

We highlight a property of the Hadamard product of rank-1 tensors. Let $\forall i, 1 \leq i \leq M$, $\mathbf{T}_i = \bigotimes_{j=1}^n \mathbf{u}_j^i$ with $\mathbf{u}_j^i \in \mathbb{R}^{p_j}$ for any arbitrary number M of rank-1 tensors. We demonstrate that

the Hadamard product of M rank-1 tensors is the tensor product of Hadamard products:

$$\begin{aligned} \bigcirc_{i=1}^M \mathbf{T}_i &= \bigcirc_{i=1}^M \bigotimes_{j=1}^n \mathbf{u}_j^i, \\ &= \bigotimes_{j=1}^n \bigcirc_{i=1}^M \mathbf{u}_j^i. \end{aligned}$$

We can prove it by analyzing the behaviour of each component of the Hadamard product:

$$\left[\bigcirc_{i=1}^M \mathbf{T}_i \right]_{k_1 k_2 \dots k_n} = \prod_{i=1}^M \prod_{j=1}^n [\mathbf{u}_j^i]_{k_j}.$$

The product here being commutative, we can interchange Hadamard and tensor products.

A.3 M-product complexity

In this section we demonstrate the low complexity of the m-product inherited by the properties of the rank-1 tensor contraction and the commutativity of Hadamard and tensor products for rank-1 tensors. Let us consider the two m-tensors $\mathbf{T}_1 \in \mathbb{R}^{m_1 \times p_1 \times p_2 \times \dots \times p_n}$ and $\mathbf{T}_2 \in \mathbb{R}^{m_2 \times p_1 \times p_2 \times \dots \times p_n}$ be m-tensors respectively based on cores $\{A_i^1\}_{1 \leq i \leq n}$ and $\{A_i^2\}_{1 \leq i \leq n}$:

$$\mathbf{T}_1 = \bigotimes_{i=1}^n A_i^1, \quad \mathbf{T}_2 = \bigotimes_{i=1}^n A_i^2.$$

We can build an element of $\mathbf{T}_1 \times (\mathbf{T}_2)^\top \in \mathbb{R}^{m_1 \times m_2}$ as the inner product of order $n - 1$ rank-1 tensors. Elements of that product are defined as

$$\begin{aligned} [\mathbf{T}_1 \times \mathbf{T}_2]_{k_1 k_2} &= \langle [\mathbf{T}_1]_{k_1}, [\mathbf{T}_2]_{k_2} \rangle, \\ &= \sum_{i_1=1}^{p_1} \sum_{i_2=1}^{p_2} \dots \sum_{i_n=1}^{p_n} [\mathbf{T}_1]_{k_1 i_1 i_2 \dots i_n} [\mathbf{T}_2]_{k_2 i_1 i_2 \dots i_n}, \\ &= \text{contract}_r([\mathbf{T}_1]_{k_1} \odot [\mathbf{T}_2]_{k_2}), \\ &= \text{contract}_r \left(\left(\bigotimes_{i=1}^n [A_i^1]_{k_1} \right) \odot \left(\bigotimes_{i=1}^n [A_i^2]_{k_2} \right) \right), \\ &= \text{contract}_r \left(\bigotimes_{i=1}^n [A_i^1]_{k_1} \odot [A_i^2]_{k_2} \right), \\ &= \prod_{i=1}^n \text{contract}_1([A_i^1]_{k_1} \odot [A_i^2]_{k_2}), \\ &= \prod_{i=1}^n \langle [A_i^1]_{k_1}, [A_i^2]_{k_2} \rangle. \end{aligned}$$

The overall computation results in the Hadamard product of matrix product of corresponding cores of both m-tensors:

$$\mathbf{T}_1 \times \mathbf{T}_2^\top = \bigcirc_{i=1}^n A_i^1 (A_i^2)^\top.$$

Considering, for ease of notation that $m_1 = m_2 = m$ and $\forall i, 1 \leq i \leq n, p_i = p$ the complexity of the m-product is $\mathcal{O}(m^2 p n)$. This complexity scales only linearly with dimension n and is dominated by m^2 . The n -way contracted product of arbitrary tensors with the same dimensions would result in a complexity of $\mathcal{O}(m^2 p^n)$.

B Polynomial toy problem

In this section we give a very simplified vision of the approach on a polynomial toy problem. The model aims at approximating a two-dimensional polynomial function defined on $\mathcal{D} = [-1, 1]^{\otimes 2}$ as

$$f : \begin{cases} \mathcal{D} & \rightarrow \mathbb{R} \\ \mathbf{x} & \mapsto 5 - [\mathbf{x}]_1 + 3[\mathbf{x}]_1[\mathbf{x}]_2 - [\mathbf{x}]_1^2 - 15[\mathbf{x}]_1^2[\mathbf{x}]_2^2 - 3[\mathbf{x}]_1[\mathbf{x}]_2^2 - [\mathbf{x}]_1^2[\mathbf{x}]_2 \end{cases} .$$

This function is evaluated on a set of m points that we denote $\mathcal{P} \subset \mathcal{D}$. We define the RHS of the linear system to be solved as

$$[\mathbf{y}]_k = f(\mathbf{x}_k), \quad \forall k, 1 \leq k \leq m, \mathbf{x}_k \in \mathcal{P}.$$

We choose to represent one-dimensional problems as second order polynomials, so we define

$$\boldsymbol{\psi}(x) = [1 \quad x \quad x^2] .$$

This mapping is applied to each component of each point of \mathcal{P} to build the cores

$$[\Psi_{k \in \{1,2\}}]_{ij} = [\boldsymbol{\psi}([\mathbf{x}_i]_k)]_j .$$

Those cores define the m-tensor regression operator Φ as

$$\Phi = \Psi_1 \otimes \Psi_2 .$$

The matrix-based approach would have been built from their face-splitting product, *i.e.* the mode-1 unfolding of Φ . We consider for instance the three-points set

$$\mathcal{P} = \{(-1, -1), (0, 1), (1, 0)\} .$$

Cores are defined as

$$\Psi_1 = \begin{bmatrix} 1 & -1 & 1 \\ 1 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix} \quad \text{and} \quad \Psi_2 = \begin{bmatrix} 1 & -1 & 1 \\ 1 & 1 & 1 \\ 1 & 0 & 0 \end{bmatrix} .$$

Whereas the mode-1 unfolding of Φ (matrix-based equivalent) gives

$$\Phi_{(1)} = \begin{bmatrix} 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \end{bmatrix} .$$

Even that simple example shows the memory gain (1/3 in that case) of storing the factorization instead of building the full operator. The linear system to be solved writes for the m-tensor and matrix case respectively

$$\Phi \times C = \mathbf{y} \quad \text{and} \quad \Phi_{(1)} \mathbf{c} = \mathbf{y},$$

where

$$\mathbf{y} = [f(-1, -1) \quad f(0, 1) \quad f(1, 0)]^\top = [-2 \quad 5 \quad 3]^\top .$$

For the matrix case, the least squares solution $\hat{\mathbf{c}}$ is found by using the Moore-Penrose inverse of $\Phi_{(1)}$, denoted $\Phi_{(1)}^\dagger$. That pseudoinverse is computed from the SVD of the operator, if we consider $\Phi_{(1)} = U\Sigma V^\top$, then $\Phi_{(1)}^\dagger = V\Sigma^{-1}U^\top$. The least squares solution is then given as

$$\hat{\mathbf{c}} = \Phi_{(1)}^\dagger \mathbf{y} .$$

It yields numerically (with rounding at 10^{-3} for readability)

$$\hat{\mathbf{c}}^\top = [1.706 \quad 2.235 \quad 1.059 \quad 1.235 \quad -0.588 \quad 0.588 \quad 0.059 \quad 0.588 \quad -0.588].$$

In the case of the m-tensor least squares, we first build $P = \Phi \times \Phi$. We notice that it yields a 3×3 matrix. Then we solve for \mathbf{z} in

$$P\mathbf{z} = \mathbf{y}.$$

This is solved by inverting P , it yields (still rounded at 10^{-3})

$$\mathbf{z}^\top = [-0.588 \quad 1.647 \quad 0.647].$$

Then computing the coefficients \hat{C} requires using a tensor-dot product we did not present in this paper but comes naturally when familiar with tensor algebra. Formally it would be defined as

$$\begin{aligned} \hat{C} &= \text{contract}_c((\text{diag}(\mathbf{z})\Psi_1) \otimes \Psi_2), \\ &= \text{contract}_c(\Psi_1 \otimes (\text{diag}(\mathbf{z})\Psi_2)), \end{aligned}$$

where $\text{diag}(\mathbf{z})$ builds the 3×3 matrix whose diagonal is \mathbf{z} . In the end, elements of \hat{C} are defined as

$$[\hat{C}]_{ij} = \sum_{k=1}^3 [\Phi]_{kij} [\mathbf{z}]_k.$$

This definition results in $\hat{C} \in \mathbb{R}^{3 \times 3}$, explicitly

$$\hat{C} = \begin{bmatrix} 1.706 & 1.235 & 0.059 \\ 2.235 & -0.588 & 0.588 \\ 1.059 & 0.588 & -0.588 \end{bmatrix},$$

i.e. the same coefficients as in the matrix-based case with a preserved structure. We demonstrated for a toy problem that solving the m-tensor system is *strictly equivalent* to building the matrix system, even though that matrix rapidly becomes unmanageable with a growing number of input parameters. We do not further examine the quality of the solution for this low-dimensional toy problem.

The computational bottleneck in each case is in the inversion of the operator, $\Phi_{(1)}$ and P for the matrix and m-tensor approaches respectively. The first requires the computation of the pseudoinverse of a 3×9 matrix versus the inversion of a symmetric positive definite 3×3 matrix in the second case.

remark The computational cost of computing the pseudoinverse of a $m \times n$ matrix with $n > m$ is dominated by the computational complexity of the SVD, that is $\mathcal{O}(m^2n)$. The cost of inverting an unstructured $m \times m$ matrix is $\mathcal{O}(m^3)$. In our example ($m = 3$ and $n = 9$) this shows that the m-tensor least squares is clearly preferable to the matrix-based approach when variables are separable in the operator.

C Algorithms to compute ALI decomposition

In this section we provide algorithms to compute the ALI decomposition of m-tensors. As we stated in the document, we have two ways of constructing it. The first is the greedy way, inspired from online learning of streamed data. We qualify the second as the optimal decomposition. We first review the greedy way and then provide arguments for the more

costful optimal decomposition. We note that review of randomized numerical linear algebra and the work of Mahoney and Drineas on CUR matrix decomposition [21] lead us to seek a *near optimal* algorithm leveraging randomized approaches with a low computational cost.

As a reminder and to settle notations for this section, we consider an m-tensor \mathbf{T} with c-dim m and provide its ALI decomposition with a tolerance ε as

$$\|\mathbf{T} - W\tilde{\mathbf{T}}\|^2 \leq \varepsilon.$$

In the remainder of the section, we give a geometric interpretation to the incremental construction of the subset of rows from a given m-tensor. We consider each row to be a point in a high dimensional space. Providing the distance definition from the inner product of two of those points, *i.e.* the inner product of rank-1 tensors, we seek to maximize at each step the distance to the previous subspace.

C.1 Greedy decomposition

The greedy construction of the ALI decomposition of an m-tensor considers that we sweep once through the rows of an m-tensor \mathbf{T} or treat rows as they are streamed (or equivalently the rank-1 tensor mapped from streamed data) to select the best ones on-the-fly. This is considered a greedy approach because it builds the best subset from available data as it comes.

Element 1

The first provided row $[\mathbf{T}]_1$ is used by default to initialize $\tilde{\mathbf{T}}$. Weights W_1 are initialized as an array containing only 1. By convenience we define

$$L_1 = \|[\mathbf{T}]_1\|.$$

From that given point, we build a subspace $\text{span}(\{[\mathbf{T}]_1\})$. For the next point, we evaluate the distance to that subspace. For a new row $[\mathbf{T}]_2$, we compute δ_2 defined as

$$\delta_2 = \min_{w \in \mathbb{R}} \|[\mathbf{T}]_2 - w[\mathbf{T}]_1\|^2.$$

Given by the minimizing w defined as

$$w = \frac{\langle [\mathbf{T}]_1, [\mathbf{T}]_2 \rangle}{\|[\mathbf{T}]_1\|^2}.$$

This minimizer is obtained by expanding the norm in the definition of δ and deriving the minimum of a quadratic polynomial in w . Using that minimizer, δ will be given by

$$\delta_2 = \|[\mathbf{T}]_2\|^2 - \left(\frac{\langle [\mathbf{T}]_2, [\mathbf{T}]_1 \rangle}{\|[\mathbf{T}]_1\|^2} \right)^2.$$

By comparing δ_2 to the precision criterion ε , the point is kept or not. If it is kept, $[\mathbf{T}]_2$ is appended to $\tilde{\mathbf{T}}$ and L is updated as

$$L_2 = \begin{bmatrix} L_1 & 0 \\ s & \sqrt{\|[\mathbf{T}]_2\|^2 - s^2} \end{bmatrix},$$

where s is defined as

$$s = \frac{\langle [\mathbf{T}]_1, [\mathbf{T}]_2 \rangle}{L_1}.$$

The weight matrix W is updated as

$$W_2 = \begin{bmatrix} W_1 & 0 \\ 0 & 1 \end{bmatrix}.$$

Otherwise, if $\delta_2 < \varepsilon$, only the weight matrix is updated as

$$W_2 = \begin{bmatrix} W_1 \\ w \end{bmatrix}.$$

Element k

For any new row $[\mathbf{T}]_k$, we adopt the incremental construction for $\tilde{\mathbf{T}}$ as follows. Consider we retained \tilde{k} elements in $\tilde{\mathbf{T}}$ until that point. This means we need to evaluate the distance of $[\mathbf{T}]_k$ to a \tilde{k} -dimensional subspace. We define the vector \mathbf{b} defined from the already selected rows and the new one as

$$\mathbf{b} = \tilde{\mathbf{T}} \times [\mathbf{T}]_k^\top.$$

The distance δ_k between $[\mathbf{T}]_k$ and $\text{span}(\tilde{\mathbf{T}})$ is evaluated as

$$\delta_k = \min_{\mathbf{w}} \left\| [\mathbf{T}]_k - \mathbf{w}^\top \tilde{\mathbf{T}} \right\|^2 = \left\| [\mathbf{T}]_k \right\|^2 - \mathbf{b}^\top \tilde{P}^{-1} \mathbf{b},$$

where $\tilde{P} = \tilde{\mathbf{T}} \times \tilde{\mathbf{T}}^\top$ is symmetric positive definite and whose Cholesky factors are given by $L_{\tilde{k}}$, *i.e.*

$$\tilde{P} = L_{\tilde{k}} L_{\tilde{k}}^\top.$$

Therefore we define the vector \mathbf{s} from $L_{\tilde{k}}$ directly as

$$\mathbf{s} = L_{\tilde{k}}^{-1} \mathbf{b},$$

which can be computed by downward backsubstitution. Given the vector \mathbf{s} , the second term of the definition of δ can be efficiently computed as $\mathbf{b}^\top \tilde{P}^{-1} \mathbf{b} = \mathbf{s}^\top \mathbf{s}$. Then by confronting δ_k to the tolerance ε , $[\mathbf{T}]_k$ is kept or not. If $\delta_k > \varepsilon$, then $[\mathbf{T}]_k$ is appended to $\tilde{\mathbf{T}}$ and the Cholesky factor is updated as

$$L_{\tilde{k}+1} = \begin{bmatrix} L_{\tilde{k}} & 0 \\ \mathbf{s}^\top & \sqrt{\|[\mathbf{T}]_k\|^2 - \mathbf{s}^\top \mathbf{s}} \end{bmatrix}.$$

The weight matrix W is updated as

$$W_{k+1} = \begin{bmatrix} W_k & 0 \\ 0 & 1 \end{bmatrix}.$$

Otherwise, if $\delta_k < \varepsilon$, only the weight matrix is updated as

$$W_{k+1} = \begin{bmatrix} W_k \\ \mathbf{w}^\top \end{bmatrix},$$

with $\mathbf{w} = (L_{\tilde{k}}^\top)^{-1} \mathbf{s}$, which can be computed by upward backsubstitution. This construction is excessively efficient as it loops only once on the rows of \mathbf{T} , it does not suffer from numerical instability as we avoid computing the inverse \tilde{P}^{-1} . Instead we leverage its incrementally constructed Cholesky factors $L_{\tilde{k}}$. For that reason it is fit for streamed data if the hyperparameter ε can be chosen a priori. We describe this construction in the pseudoalgorithm 1.

In their work, Engel et al. [8] give a proof that the number of kept tensors \tilde{k} will always be finite under mild conditions. This means that the algorithm will eventually stop adding tensors to $\tilde{\mathbf{T}}$.

Algorithm 1 Greedy ALI decomposition

Input: $\varepsilon, [\mathbf{T}]_1$
Output: \mathbf{T}, W

Initialization

$\tilde{\mathbf{T}} \leftarrow [\mathbf{T}]_1$ ▷ Initialize reduced m-tensor
 $L_1 \leftarrow \|[[\mathbf{T}]_1]\|$ ▷ Initialize Cholesky factor
 $\tilde{k} \leftarrow 1$ ▷ Initialize row index in reduced m-tensor
 $W \leftarrow [1]$ ▷ Initialize weights

for any new $[\mathbf{T}]_k, k > 1$ **do** ▷ Loop on rows of initial \mathbf{T} or streamed data

Compute \mathbf{b}
 $\mathbf{b} \leftarrow \tilde{\mathbf{T}} \times [\mathbf{T}]_k^\top$
 Compute \mathbf{s} by backsubstitution
 $\mathbf{s} \leftarrow L_{\tilde{k}}^{-1} \mathbf{b}$
 Evaluate δ_k
 $\delta_k \leftarrow \|[[\mathbf{T}]_k]\|^2 - \mathbf{s}^\top \mathbf{s}$

if $\delta_k > \varepsilon$ **then** ▷ Compare distance to threshold

Update reduced m-tensor $\tilde{\mathbf{T}}$
 $\tilde{\mathbf{T}} \leftarrow \begin{bmatrix} \tilde{\mathbf{T}} \\ [\mathbf{T}]_k \end{bmatrix}$

Update number of kept tensors \tilde{k}
 $\tilde{k} \leftarrow \tilde{k} + 1$

Update Cholesky factor L
 $L_{\tilde{k}} \leftarrow \begin{bmatrix} L_{\tilde{k}-1} & 0 \\ \mathbf{s}^\top & \sqrt{\|[[\mathbf{T}]_k]\|^2 - \mathbf{s}^\top \mathbf{s}} \end{bmatrix}$

Update weights W
 $W \leftarrow \begin{bmatrix} W & 0 \\ 0 & 1 \end{bmatrix}$

else

Update weights W
 $W \leftarrow \begin{bmatrix} W \\ (L_{\tilde{k}}^\top)^{-1} \mathbf{s} \end{bmatrix}$

end if

end for
return $W, \tilde{\mathbf{T}}$

C.2 Optimal decomposition

Consider we are given a full m-tensor \mathbf{T} with c-dim m . We seek the m-tensor $\tilde{\mathbf{T}}$ of c-dim $\tilde{m} < m$ such that all rows of \mathbf{T} are *almost linearly dependent* on the rows of $\tilde{\mathbf{T}}$ with respect to a tolerance ε . To build such optimal decomposition, we propose an iterative approach during which we add the best choice among all rows of \mathbf{T} at each step.

Element 1

The first row of $\tilde{\mathbf{T}}$ is selected so that its span minimizes the sum of distances with the other rows. To evaluate the distance of all points to the subspaces spanned by each points, we build the matrix $\Delta_1 \in \mathbb{R}^{m \times m}$ as

$$[\Delta_1]_{ij} = \delta_{ij}^1, \quad \delta_{ij}^1 = \min_{w \in \mathbb{R}} \|[[\mathbf{T}]_i - w[\mathbf{T}]_j]\|^2.$$

Each element is defined by expanding the norm in the definition of distances provided here. This expansion provides a quadratic formulation in terms of w which admits a unique minimum that we inject to build Δ_1 with elements defined as

$$\delta_{ij}^1 = \|\mathbf{[T]}_i\|^2 - \left(\frac{\langle \mathbf{[T]}_i, \mathbf{[T]}_j \rangle}{\|\mathbf{[T]}_j\|} \right)^2.$$

We compute the contraction of each column vector of Δ_1 and select the first row of $\tilde{\mathbf{T}}$ to minimize that sum. If all distances in the column vector corresponding to the selected point were below the precision tolerance ε , then $\tilde{\mathbf{T}}$ would be sufficient and the algorithm would end there.

In practice, the elements of the matrix Δ_1 can be built from the elements of $P = \mathbf{T} \times \mathbf{T}^\top$. The squared norms are on its diagonal and the inner products in the upper (or lower equivalently) triangle. It is important to notice that Δ_1 is not symmetric.

Element k

The k -th row of $\tilde{\mathbf{T}}$, $k > 1$ is selected so that the subspace

$$\text{span} \left(\left\{ [\tilde{\mathbf{T}}]_i, 1 \geq i \geq k \right\} \right)$$

minimizes the distances to all other points. In order to evaluate these distances, we denote $\bar{\mathbf{T}}_k$ the m-tensor \mathbf{T} with rows previously selected in $\tilde{\mathbf{T}}$ removed. We define $\tilde{\mathbf{T}}_j$ as $\tilde{\mathbf{T}}$ appended with the j -th row of $\bar{\mathbf{T}}_k$

$$\tilde{\mathbf{T}}_j = \begin{bmatrix} \tilde{\mathbf{T}} \\ [\bar{\mathbf{T}}_k]_j \end{bmatrix}.$$

Then we build the matrix $\Delta_k \in \mathbb{R}^{(m-k) \times (m-k)}$ to evaluate the distances of each remaining point (not in $\tilde{\mathbf{T}}$) to each possible augmented subspaces.

$$[\Delta_k]_{ij} = \delta_{ij}^k, \quad \delta_{ij}^k = \min_{\mathbf{w} \in \mathbb{R}^k} \left\| [\bar{\mathbf{T}}_k]_i - \mathbf{w}^\top \tilde{\mathbf{T}}_j \right\|^2.$$

The minimizing weights \mathbf{w} can be computed by expanding the norm and solving for the unique minimum. This results in the minimized distances given by

$$\delta_{ij}^k = \|\mathbf{[\bar{T}}_k]_i\|^2 - (\mathbf{b}_k^{ij})^\top (\tilde{P}_k^j)^{-1} \mathbf{b}_k^{ij},$$

where we define $\mathbf{b}_k^{ij} = \tilde{\mathbf{T}}_j \times [\bar{\mathbf{T}}_k]_i^\top$, $\tilde{P}_k^j = \tilde{\mathbf{T}}_j \times \tilde{\mathbf{T}}_j^\top$. We notice that \tilde{P}_k^j is symmetric positive definite, and therefore admits a Cholesky factorization $\tilde{P}_k^j = L_k^j (L_k^j)^\top$. Seeing that each distance computation requires the inversion of \tilde{P}_k^j , we use that factor's definition and reformulate the second term of the definition of δ_{ij}^k as

$$(\mathbf{b}_k^{ij})^\top (\tilde{P}_k^j)^{-1} \mathbf{b}_k^{ij} = ((L_k^j)^{-1} \mathbf{b}_k^{ij})^\top ((L_k^j)^{-1} \mathbf{b}_k^{ij}) = (\mathbf{s}_k^{ij})^\top \mathbf{s}_k^{ij}.$$

With this construction, a single upward backsubstitution is needed to compute each \mathbf{s}_k^{ij} . We propose to alleviate the notations and use those formulations. Therefore we suggest the following matrix construction for Δ_k for more readability. We define at step k the matrix as

$$\Delta_k = \text{diag}(\bar{\mathbf{n}}) \mathbf{1}_{m-k} - S_k^\top S_k,$$

where $\mathbf{1}_{m-k}$ denotes the square matrix of size $m - k$ filled with ones. Then $\bar{\mathbf{n}}$ is defined as the vector containing the norms of each row of $\bar{\mathbf{T}}$, *i.e.*

$$\bar{\mathbf{n}}^\top = \left[\|\mathbf{[\bar{T}]}_1\| \quad \|\mathbf{[\bar{T}]}_2\| \quad \dots \quad \|\mathbf{[\bar{T}]}_{m-k}\| \right].$$

The matrix S_k is defined as

$$S_k = L_k^{-1} B_k,$$

where B_k is defined as

$$B_k = \tilde{\mathbf{T}} \times \bar{\mathbf{T}}^\top,$$

and L_k is the Cholesky factor of $\tilde{P}_k = \tilde{\mathbf{T}} \times \tilde{\mathbf{T}}^\top$. As for the first element selection, we sum the columns of Δ_k to obtain a vector $\boldsymbol{\delta}_k$. We select the index of the best choice in the rows of $\bar{\mathbf{T}}_k$ as the minimizer

$$\hat{i} = \operatorname{argmin}_i([\boldsymbol{\delta}_k]_i).$$

The corresponding row of $\bar{\mathbf{T}}_k$ is appended to $\tilde{\mathbf{T}}$. We check if all values of distances in the \hat{i} -th column of Δ_k are below ε , in which case the precision criterion is reached and the algorithm is stopped. Otherwise, operators are updated and the loop continues.

We note that L_k can be built incrementally, along with $\tilde{\mathbf{T}}$ with $L_1 = [||[\tilde{\mathbf{T}}]_1||]$ and for $k > 1$ with the formula

$$L_{k+1} = \begin{bmatrix} L_k & 0 \\ \mathbf{s}_k^\top & \sqrt{||[\tilde{\mathbf{T}}]_k||^2 - \mathbf{s}_k^\top \mathbf{s}_k} \end{bmatrix},$$

where \mathbf{s}_k is the \hat{i} -th column vector of S_k .

Once the optimal $\tilde{\mathbf{T}}$ of size \tilde{m} is computed, we evaluate the weight matrix using the projector onto the span of $\tilde{\mathbf{T}}$ denoted by $\Pi_{\tilde{\mathbf{T}}}$ and defined as

$$\Pi_{\tilde{\mathbf{T}}} = \tilde{\mathbf{T}}^\top \left(\tilde{\mathbf{T}} \times \tilde{\mathbf{T}}^\top \right)^{-1} \tilde{\mathbf{T}}.$$

That projector is used to circumvent the fact that our decomposition does not provide an orthonormal basis. Therefore we cannot simply build the weights in a unique way without defining a way to project elements on the ALI 'basis'. We identify the weights in the projection of \mathbf{T}

$$\mathbf{T} \times \Pi_{\tilde{\mathbf{T}}} = W \tilde{\mathbf{T}}.$$

Therefore the weight matrix W is given by

$$W = \mathbf{T} \times \tilde{\mathbf{T}}^\top \tilde{P}_m^{-1} = \mathbf{T} \times \tilde{\mathbf{T}}^\top (L_{\tilde{m}}^{-1})^\top L_{\tilde{m}}^{-1}.$$

We provide the pseudoalgorithm 2 to build the optimal m-tensor ALI decomposition from a full m-tensor. The algorithm uses the Cholesky factor and updates it in the loop. Therefore it is already an optimized and numerically stable option to implement.

remark We make a remark on the use in the context of regression. In that case, we consider we are given a set of sample \mathcal{P} of size m . We seek the subset $\tilde{\mathcal{P}}$ of size \tilde{m} such that the span of the elements of $\tilde{\mathcal{P}}$ through the mapping φ minimizes the ℓ_2 error of the projection of the full mapped set with a tolerance ε .

The first element \mathbf{x}_1 of $\tilde{\mathcal{P}}$ is selected so that the span of its mapping $\operatorname{span}(\{\varphi(\mathbf{x}_1)\})$ minimizes the sum of distances with the remainder of the mapped set. Then, the k -th element \mathbf{x}_k , $k > 1$ of the subset $\tilde{\mathcal{P}}$ is selected so that the subspace

$$\operatorname{span} \left(\left\{ \varphi(\mathbf{x}), \mathbf{x} \in \tilde{\mathcal{P}} \cup \{\mathbf{x}_k\} \right\} \right)$$

minimizes the distances to all other mapped points. The approach is exactly the same and leverages the fact that the mapping φ maps a point to a rank-1 tensor and a set to an m-tensor. That way a row $[\mathbf{T}]_i$ of \mathbf{T} is seen as the mapping of a point $\mathbf{x}_i \in \mathcal{P}$. The rows of $\tilde{\mathbf{T}}$ are the points of $\tilde{\mathcal{P}}$ mapped through φ .

Algorithm 2 Optimal ALI decomposition**Input:** \mathbf{T}, ε **Output:** $\tilde{\mathbf{T}}, W$ Build matrix Δ_1 Sum columns of Δ_1

$$[\boldsymbol{\delta}_1]_i \leftarrow \sum_{j=1}^m [\Delta_1]_{ji}$$

Select first element of $\tilde{\mathbf{T}}$

$$\tilde{\mathbf{T}} \leftarrow [\mathbf{T}]_{\hat{i}}, \hat{i} = \operatorname{argmin}_i([\boldsymbol{\delta}_1]_i)$$

Initialize Cholesky factor L_1

$$L_1 \leftarrow [||\tilde{\mathbf{T}}||]$$

Define reduced m-tensor $\bar{\mathbf{T}}_1$ by removing \hat{i} -th row of \mathbf{T} **for** $k, 1 < k \leq m$ **do**▷ Loop on rows of \mathbf{T} Build vector $\bar{\mathbf{n}}_k$ of norms of each row of $\bar{\mathbf{T}}_k$

$$\bar{\mathbf{n}}_k^\top \leftarrow [||[\bar{\mathbf{T}}_k]_1|| \quad ||[\bar{\mathbf{T}}_k]_2|| \quad \dots \quad ||[\bar{\mathbf{T}}_k]_{m-k}||]$$

Build matrix B_k

$$B_k \leftarrow \tilde{\mathbf{T}} \times \mathbf{T}_k^\top$$

Compute matrix S_k

$$S_k \leftarrow L_k^{-1} B_k$$

Compute matrix Δ_k

$$\Delta_k \leftarrow \operatorname{diag}(\bar{\mathbf{n}}_k) \mathbf{1} - S_k^\top S_k$$

Sum columns of Δ_k in vector $\boldsymbol{\delta}_k$

$$[\boldsymbol{\delta}_k]_i \leftarrow \sum_{j=1}^{m-k} [\Delta_k]_{ji}$$

Find optimal column index \hat{i}

$$\hat{i} = \operatorname{argmin}_i([\boldsymbol{\delta}_k]_i)$$

Append \hat{i} -th row of $\bar{\mathbf{T}}$ as k -th row of $\tilde{\mathbf{T}}$

$$\tilde{\mathbf{T}}^\top \leftarrow \begin{bmatrix} \tilde{\mathbf{T}}^\top \\ [\mathbf{T}]_{\hat{i}}^\top \end{bmatrix}$$

Select \hat{i} -th column vector of S_k

$$\mathbf{s}_k \leftarrow [S_k^\top]_{\hat{i}}$$

Update Cholesky factor L_{k+1}

$$L_{k+1} \leftarrow \begin{bmatrix} L_k & 0 \\ \mathbf{s}_k^\top & \sqrt{||[\mathbf{T}]_{\hat{i}}||^2 - \mathbf{s}_k^\top \mathbf{s}_k} \end{bmatrix}$$

Update $\bar{\mathbf{T}}_{k+1}$ by removing \hat{i} -th row of $\bar{\mathbf{T}}_k$ **if** $\forall j, 1 \leq j \leq m - k, [\Delta]_{j\hat{i}} \leq \varepsilon$ **then** ▷ If all distances to current subspace fall under ε

Break loop

end if**end for**Compute weights W with respect to final $\tilde{\mathbf{T}}$

$$W \leftarrow \mathbf{T} \times \tilde{\mathbf{T}}^\top L_k^{-\top} L_k^{-1}$$

return $\tilde{\mathbf{T}}, W$

remark As stated earlier, we realized the gain to be found from the field of randomized numerical linear algebra. By seeking a m-tensor CUR generalization, we found in the work of Mahoney and Drineas [21] a path for perspectives leveraging randomized approaches, instead of what we have done that relates more with the pseudoskeleton approach of Goreinov et al. [11].