

A scalable estimator of sets of integral operators

Valentin Debarnot
ITAV, CNRS, Toulouse, France
valentindebarnot@gmail.com

Paul Escande
Department of Applied Mathematics and Statistics
John Hopkins University, Baltimore, USA
paul.escande@gmail.com

Pierre Weiss
ITAV, CNRS, Toulouse, France
pierre.armand.weiss@gmail.com

July 5, 2019

Abstract

The main objective of this work is to estimate a low dimensional subspace of operators in order to improve the identifiability of blind inverse problems. We propose a scalable method to find a subspace $\hat{\mathcal{H}}$ of low-rank tensors that simultaneously approximates a set of integral operators. The method can be seen as a generalization of tensor decomposition models, which was never used in this context. In addition, we propose to construct a convex subset of $\hat{\mathcal{H}}$ in order to further reduce the search space. We provide theoretical guarantees on the estimators and a few numerical results.

1 Introduction

In many measurement devices, a signal v_0 living in some Hilbert space \mathcal{B}_n of dimension n is probed indirectly using an operator $H_0 : \mathcal{B}_n \rightarrow \mathcal{B}_m$, where \mathcal{B}_m is a Hilbert space of dimension m ¹. This yields a measurement vector $u_0 \in \mathcal{B}_m$ defined by

$$u_0 = f(H_0 v_0),$$

where f is some perturbation of the measurements (e.g. additive noise, modulus for phase retrieval, quantization,...). Solving an inverse problem consists in recovering an approximation \hat{v} of the signal v_0 using the measurements u_0 .

¹In all this paper, we assume that the operators are defined in finite dimensional spaces. An extension to infinite dimensional Hilbert spaces is feasible but requires additional discretization procedures. We decided to skip this aspect to clarify the exposition.

When the operator H_0 is known, many efficient solutions are now available. Unfortunately, in many cases, only a crude estimate of H_0 is available or it is even completely unknown. This is the field of bilinear or blind inverse problems. In that case, finding a reasonable approximation is far more involved. Significant theoretical progresses have been achieved in the last few years though, [18, 1, 23, 24, 22, 25].

One of the key ideas behind these methods is the principle of lifting. To apply it, it is common to assume that the operator H_0 and the signal v_0 live in known low dimensional vector spaces of operators $\mathcal{H} = \text{span}(P_1, \dots, P_{|S|})$ and signals $\mathcal{Q} = \text{span}(q_1, \dots, q_{|T|})$. Then, we can write that $H_0 = P\alpha_0$ and that $v_0 = Q\beta_0$ for some $\alpha_0 \in \mathbb{R}^{|S|}$ and some $\beta_0 \in \mathbb{R}^{|T|}$. Under those assumptions, the blind inverse problem is simplified to finding a pair of vectors $(\alpha, \beta) \in \mathbb{R}^{|S|} \times \mathbb{R}^{|T|}$ and the measurement associated to the pair can be written as

$$u_0 = (P\alpha)(Q\beta) = \sum_{s \in S, t \in T} \alpha_s \beta_t w_{s,t},$$

with $w_{s,t} = P_s q_t$, $S = \{1, \dots, |S|\}$ and $T = \{1, \dots, |T|\}$. This last expression only depends on the outer product $\alpha\beta^T$, allowing to lift the problem to the matrix space $\mathbb{R}^{|S| \times |T|}$. A typical way to attack the blind inverse problem is then to solve the following optimization problem:

$$\min_{M \in \mathbb{R}^{|S| \times |T|}, \text{rank}(M)=1} \frac{1}{2} \|\mathcal{W}M - y\|_2^2, \quad (1)$$

where $\mathcal{W} : M \mapsto \sum_{s \in S, t \in T} M_{s,t} w_{s,t}$. Various relaxations and algorithms can then be used to solve the lifted problem (1) and come with strong theoretical guarantees. We refer the interested reader to the above mentioned papers.

A critical issue to apply these techniques is the knowledge of the subspaces \mathcal{H} and \mathcal{Q} . In this paper, we will focus on the estimation of the subspace \mathcal{H} from a sampling set of operators $(H_l)_{l \in L}$ in $\mathcal{C} \subset \mathcal{H}$.

The interest is that determining a low dimensional set of operators with a small volume can significantly ease the problem of operator identification in blind inverse problems. While our primary motivation lies in the field of inverse problems, this problem can also be understood as a generic problem of approximation theory.

1.1 Application examples

Space varying blur An example that will be used in our numerical experiments is the case of space varying blurs in wide field microscopy. In this imaging modality, the blur varies spatially due to multiple effects such as scattering or defocus for instance. The possible family of blurs may vary depending on factors such as the focal screw, the temperature (which changes the refractive index of the immersion oil), small tilts with respect to the focal plane and many other parameters that are hard to model from a mathematical point of view. It is

possible to collect a family of operators $(H_l)_{l \in L}$ by observing fluorescent microbeads in a slide under various conditions and by using operator interpolation techniques such as [5].

Magnetic Resonance Imaging (MRI) In MRI, the traditional observation model simply states that the Fourier transform values of the image are observed. The reality is far more complex and complete image formation models comprise many unknowns such as inhomogeneities of the main magnetic field or spatial sensitivities [17]. To apply the proposed methodology to this device, the idea would be to first run many calibration scans to recover a list of operators $(H_l)_{l \in L}$ and then build a reduced model from this set.

Diffusion equations In many applications such as electrical impedance tomography [9], the operators H_l are given implicitly as solutions of partial differential equations (PDEs). For instance diffusion equations, which are widespread in applications, are of the form $\text{div}(c_l \nabla u) = v$, where c_l is a space varying diffusion coefficient that may change depending on external parameters. The application that maps v to u can be written as a linear integral operator H_l .

1.2 Contributions

The simplest approach to find a low dimensional vector space of operators \mathcal{H} is to apply a principal component analysis (PCA) on the set of vectorized operators $(H_l)_{l \in L}$. This approach is optimal in the sense of the Hilbert-Schmidt norm, but infeasible in practice. For instance, space varying blurring operators acting on small $2D$ images of size 1000×1000 can be encoded as matrices H_l of size $10^6 \times 10^6$, which can hardly be stored since each of them contains 9 Tera octets of data.

In this work, we therefore work under the assumption that the operators can be well approximated by low-rank tensors up to an invertible transformation. This hypothesis is reasonable for many applications of interest. For instance, it includes product-convolution expansions [16] and hierarchical matrices [19] as special cases. We then provide an estimator $\hat{\mathcal{H}}$ of the subspace of operators \mathcal{H} with an upper-bound of its rate of approximation. In addition, we propose to construct an estimator $\hat{\mathcal{C}}$ of \mathcal{C} , as the convex hull (in a matrix space) of the operators $(H_l)_{l \in L}$ projected onto $\hat{\mathcal{H}}$. To make further use of this convex hull, we propose a fast projection algorithm on $\hat{\mathcal{C}}$. We finally provide various numerical examples to highlight the strengths of the approach and its scalability.

1.3 Related works

To the best of our knowledge, the overall objective of this work is new, even though most of the individual tools that we combine together are well established. A related idea can be found in the literature of PDEs, where reduced

order bases [26, 31] or their variants [10] allow to solve families of PDEs efficiently. However, the objective there is to approximate the solutions of a PDE (usually linear) and not the associated operator. This is a significant difference, since approximating the operator (and its adjoint) allows to use the rich collection of nonlinear regularizers commonly used in the field of inverse problems to find regularized solutions.

2 Notation

In all the paper, I, J, K and L are the sets of integers ranging from 1 to $|I|, |J|, |K|$ and $|L|$. We assume that $u \in \mathcal{B}_m$ is defined over a set X of cardinality m . We let $u(x)$ denote the value of u at $x \in X$. Similarly, we assume that $Hu \in \mathcal{B}_n$ is defined over a set Y . The set of linear operators from \mathcal{B}_m to \mathcal{B}_n is denoted Ξ . An operator $H \in \Xi$ can either refer to an operator or its matrix representation in an arbitrary orthogonal basis. The entries in the matrix representation will be denoted $H(x, y)$. The Frobenius norm of H is defined by $\|H\|_F := \sqrt{\text{tr}(H^*H)}$. It is invariant by orthogonal transforms. The scalar products over all spaces will be denoted by $\langle \cdot, \cdot \rangle$.

The tensor product between two vectors $a \in \mathcal{B}_n$ and $b \in \mathcal{B}_m$ is defined by $(a \otimes b)(x, y) = a(x)b(y)$. The notation \odot stands for the element-wise (Hadamard) product and if X has a group structure and $a_1, a_2 \in \mathcal{B}_m$, $a_1 \star a_2$ denotes the convolution product between a_1 and a_2 .

Let $E = (e_i)_{i \in I}$ denote a family of elements in \mathcal{B}_m . The same notation will also apply to the matrix $E = [e_1, \dots, e_{|I|}]$ and to the subspace $E = \text{span}(e_i, i \in I)$. Let $W = (w_k)_{k \in K}$ denote a family of vectors with an SVD of the form $W = USV^T$ with $U = [u_1, \dots, u_n]$, then

$$\text{SVD}_{|I|}(w_k, k \in K) \stackrel{\text{def.}}{=} [u_1, \dots, u_{|I|}],$$

i.e. the $|I|$ left singular vectors associated to the largest singular values.

We let $\Delta_{N-1} = \{x \in \mathbb{R}^N, \sum_{i=1}^N x_i = 1\}$ denote the simplex of dimension N . We let \mathcal{K}_d denote the set of compact and convex sets of \mathbb{R}^d with non empty interior. The Hausdorff distance between C_1 and C_2 is defined by $\mathcal{D}(C_1, C_2) = \inf\{\epsilon \geq 0 : C_1 \subset C_2 + \epsilon B(0, 1), C_2 \subset C_1 + \epsilon B(0, 1)\}$, where $B(0, 1)$ is the unit Euclidean ball. Let $(X_n)_{n \in \mathbb{N}}$ be sequence of random variables and $(t_n)_{n \in \mathbb{N}}$ denote a sequence of real numbers, the notation $X_n = \mathcal{O}_{\mathbb{P}}(t_n)$ means that for any $\epsilon > 0$, there exists $M > 0$ and $N > 0$ such that $\mathbb{P}(|X_n/t_n| > M) < \epsilon$ for all $n > N$.

3 Operator representations

A critical requirement in this work is that the operators H_l can be approximated by (local) low-rank tensors, up to a linear transform. This need comes from the fact that arbitrary operators have no chance of being i) computed efficiently

in large scale applications and ii) approximated efficiently by low dimensional subspaces. We describe a few possible decompositions below.

3.1 Low-rank approximations

The simplest assumption is to state that every operator H_l is well approximated by a low-rank tensor of the form $H_l = \sum_{k \in K} \alpha_{k,l} \otimes \beta_{k,l}$, with $|K| \ll \min(m, n)$. Unfortunately, many observation operators met in practice are concentrated along their diagonal, making this assumption unrealistic.

3.2 Product-convolution expansions

Product-convolution expansions are a family of decompositions that were analyzed recently in [16]. They can be defined whenever $X = Y$ and X possesses a group structure. It amounts to assuming that

$$H_l(u) = \sum_{k \in K} \alpha_{k,l} \star (\beta_{k,l} \odot u). \quad (2)$$

This decomposition can be computed efficiently using fast Fourier transforms.

To understand its link with the low-rank assumption, it is handy to introduce the *space varying impulse response* (SVIR) of H_l defined by $S_l(x, y) = H_l(x + y, y)$. One can show that the SVIR of an operator S_l of the form (2) can be written as $S_l = \sum_{k \in K} \alpha_{k,l} \otimes \beta_{k,l}$. Hence, assuming that H_l can be approximated by a product-convolution expansion is equivalent to saying that its SVIR is nearly low-rank.

This assumption covers many practical applications. For instance, a sufficient condition for an operator H_l to be well approximated using this decomposition is that all the *impulse responses* $(S_l(\cdot, y))_{y \in Y}$ of the operators H_l can be simultaneously encoded in the basis $\text{span}(\alpha_{k,l}, k \in K)$.

3.3 Hierarchical matrices

Hierarchical matrix approximations [3, 19], are another popular method to approximate linear operators. It amounts to assuming that $H_l = \sum_{k \in K} \alpha_{k,l} \otimes \beta_{k,l}$, where $|K|$ is not necessarily small compared to m and n , but where most of the elements $\alpha_{k,l}$ and $\beta_{k,l}$ have a small support, allowing for fast matrix-vector products. It can be shown that many practical applications are well suited to those approximations. It is particularly popular in the fields of PDEs and some inverse problems. In addition, related approximations such as fast multipole methods [2] or wavelet expansions [4, 15] also fit this formalism.

3.4 A general setting

Overall, the most generic assumption on H_l can be formulated as follows.

Assumption 1. *There exists a left invertible linear mapping $\mathcal{R} : \Xi \rightarrow \Xi$ such that each sample $H_l \in \Xi$ satisfies:*

$$S_l = \mathcal{R}(H_l) = \sum_{k \in K} \alpha_{k,l} \otimes \beta_{k,l},$$

where for all $l \in L$, the sets $(\alpha_{k,l})_k \in \mathcal{A}$ and $(\beta_{k,l})_k \in \mathcal{B}$, where \mathcal{A} and \mathcal{B} are subspaces of $\mathcal{B}_m^{|K|}$ and $\mathcal{B}_n^{|K|}$ respectively.

Introducing the operator \mathcal{R} allows to encompass the usual low-rank assumption by taking $\mathcal{R} = \text{Id}$, but also the product-convolution expansions: going from the SVIR to the matrix representation can be expressed through an operator $\mathcal{R} : \Xi \rightarrow \Xi$ that shifts each column of H_l . The spaces \mathcal{A} and \mathcal{B} allow to incorporate support constraints, which are used for many decompositions such as the hierarchical matrices.

The final objective of this work is to estimate a subspace \mathcal{H} and a set \mathcal{C} . In fact, we will rather estimate $\mathcal{H}_{\mathcal{R}} = \mathcal{R}\mathcal{H}$ and $\mathcal{C}_{\mathcal{R}} = \mathcal{R}\mathcal{C}$, which is equivalent since \mathcal{R} is assumed to be left-invertible. In order to lighten the notation, we will skip the multiplication by \mathcal{R} in the rest of the paper.

4 Subspace estimation

In this section we provide an efficient and robust method to estimate the vector space of operators \mathcal{H} . We look for an estimator $\widehat{\mathcal{H}}$ of \mathcal{H} with a tensor product structure:

$$E \otimes F \stackrel{\text{def.}}{=} \text{span}(e_i \otimes f_j, (e_i)_{i \in I} \in \mathcal{E}_{|I|}, (f_j)_{j \in J} \in \mathcal{F}_{|J|}),$$

where the sets $\mathcal{E}_{|I|}$ and $\mathcal{F}_{|J|}$ can be thought of as:

- The set of orthogonal families of cardinality $|I|$ and $|J|$ defined by

$$\mathcal{E}_{|I|} = \{e_i \in \mathcal{B}_m, i \in I, \|e_i\|_2 = 1, \langle e_i, e_{i'} \rangle = \delta_{i,i'}\} \quad (3)$$

and

$$\mathcal{F}_{|J|} = \{f_j \in \mathcal{B}_n, j \in J, \|f_j\|_2 = 1, \langle f_j, f_{j'} \rangle = \delta_{j,j'}\}. \quad (4)$$

- The set of orthogonal families of cardinality $|I|$ and $|J|$ with support constraints.
- Additional knowledge on the operators, such as non-negativity, can possibly be added.

We impose a tensor product structure so that every operator living in $\widehat{\mathcal{H}}$ can be evaluated rapidly. The sets $\mathcal{E}_{|I|}$ and $\mathcal{F}_{|J|}$ do not necessarily coincide with the sets \mathcal{A} and \mathcal{B} , since it could be interesting to change the structure of the operators that are given as input to the algorithm.

The principle of our approach is to find a structured low-dimensional basis of operators that allows to approximate simultaneously all the sampled representations $(S_l)_{l \in L}$. This principle can be expressed through a variational problem, as follows:

$$(\widehat{E}, \widehat{F}) \stackrel{\text{def.}}{=} \underset{\substack{(e_i)_{i \in I} \in \mathcal{E}_{|I|} \\ (f_j)_{j \in J} \in \mathcal{F}_{|J|}}}{\text{argmin}} \phi(E, F), \quad (5)$$

with

$$\phi(E, F) \stackrel{\text{def.}}{=} \frac{1}{2} \sum_{l \in L} \|\Pi_{E \otimes F}(S_l) - S_l\|_F^2,$$

where $\Pi_{E \otimes F}$ is the projection onto the tensor product space $E \otimes F$.

4.1 The algorithm

Problem (5) is bi-convex: it is non-convex in the pair $(E, F) \in \mathcal{E}_{|I|} \times \mathcal{F}_{|J|}$, but it is convex when minimizing in $E \in \mathcal{E}_{|I|}$ for $F \in \mathcal{F}_{|J|}$ fixed and when minimizing in $F \in \mathcal{F}_{|J|}$ for $E \in \mathcal{E}_{|I|}$ fixed. This motivates the use of the alternating minimization procedure presented in Algorithm 1.

Algorithm 1 Alternating Least Squares (ALS)

Approximatively solve: Problem (5)

INPUT: $(S_l)_{l \in L}$, subspace constraints $\mathcal{E}_{|I|}$ and $\mathcal{F}_{|J|}$, initial guess (E_0, F_0) .

- 1: **procedure**
 - 2: Initialization: $t = 0$.
 - 3: **while** stopping criterion not satisfied **do**
 - 4: $E_{t+1} = \underset{E \in \mathcal{E}_{|I|}}{\text{argmin}} \phi(E, F_t)$.
 - 5: $F_{t+1} = \underset{F \in \mathcal{F}_{|J|}}{\text{argmin}} \phi(E_{t+1}, F)$.
 - 6: $t = t + 1$
 - 7: **end while**
 - 8: Return $\widehat{\mathcal{H}}_L = E_t \otimes F_t$.
 - 9: **end procedure**
-

This algorithm is tightly related to common methods found in the field of tensor decompositions. In the particular case where $\mathcal{E}_{|I|}$ and $\mathcal{F}_{|J|}$ are sets of orthogonal families of cardinality $|I|$ and $|J|$, Problem (5) coincides exactly with the Tucker2 model. This decomposition was first introduced by Tucker in [32]. It was then reinvented independently and given several names such as tensor PCA, 2DSVD, GLRAM, common component analysis, or tensor decompositions [32, 14, 34, 33]. We refer to the review papers [21, 12] for more insight on tensor decompositions. Computing this decomposition is a complex nonconvex problem, but the most standard approach to solve it takes the algorithmic form provided in Algorithm 1. It does not converge to the global minimizer in general and only provides approximate solutions. However, it is observed that it usually

yields estimates close to the global minimizer in practice with a properly chosen initialization.

4.1.1 Orthogonal constraints

In this section, we detail the algorithm, when the spaces $\mathcal{E}_{|I|}$ and $\mathcal{F}_{|J|}$ denote the set of orthogonal families of cardinality $|I|$ and $|J|$ respectively.

Initialization The initialization of Algorithm 1 is of major importance since Problem (5) is non convex. We suggest using the High Order Singular Value Decomposition (HOSVD) [13] in order to initialize the algorithm. This can be seen as a generalization of the SVD for tensors. As discussed in [21], this popular method provides a good starting point for an alternating algorithm.

From a variational point of view, the principle of the HOSVD consists in solving the following problems:

$$E_0 = \operatorname{argmin}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{l \in L} \|S_l - \sum_{k \in K} \Pi_E(\alpha_{k,l}) \otimes \beta_{k,l}\|_F^2, \quad (6)$$

$$F_0 = \operatorname{argmin}_{F \in \mathcal{F}_{|J|}} \frac{1}{2} \sum_{l \in L} \|S_l - \sum_{k \in K} \alpha_{k,l} \otimes \Pi_F(\beta_{k,l})\|_F^2, \quad (7)$$

i.e. to find the subspace E (resp. F) that captures most of the energy.

We will show below that we can leverage the specific low-rank structure of the operators S_l to evaluate the HOSVD rapidly. We let $A_l = [\alpha_{1,l}, \dots, \alpha_{|K|,l}]$ and $B_l = [\beta_{1,l}, \dots, \beta_{|K|,l}]$. We also diagonalize $A_l^T A_l \in \mathbb{R}^{|K| \times |K|}$ and $B_l^T B_l \in \mathbb{R}^{|K| \times |K|}$ as

$$A_l^T A_l = \Psi_{A_l} \Lambda_l \Psi_{A_l}^T \text{ and } B_l^T B_l = \Psi_{B_l} \Sigma_l \Psi_{B_l}^T$$

with $\Sigma_l = \operatorname{diag}(\sigma_{1,l}^2, \dots, \sigma_{|K|,l}^2)$ and $\Lambda_l = \operatorname{diag}(\lambda_{1,l}^2, \dots, \lambda_{|K|,l}^2)$.

Lemma 1 (Higher Order Singular Value Decomposition (HOSVD)). *Let $\tilde{A}_l = A_l \Psi_{B_l} = [\tilde{\alpha}_{1,l}, \dots, \tilde{\alpha}_{|K|,l}]$ and $\tilde{B}_l = B_l \Psi_{A_l} = [\tilde{\beta}_{1,l}, \dots, \tilde{\beta}_{|K|,l}]$. We have*

$$E_0 = \operatorname{SVD}_{|I|}(\sigma_{k,l} \tilde{\alpha}_{k,l}, k \in K, l \in L)$$

and

$$F_0 = \operatorname{SVD}_{|J|}(\lambda_{k,l} \tilde{\beta}_{k,l}, k \in K, l \in L)$$

Proof. We concentrate on E_0 only, since the proof for F_0 is similar. The first argument is to notice that Problem (6) is equivalent to

$$E_0 = \operatorname{argmax}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{l \in L} \left\| \sum_{k \in K} \Pi_E(\alpha_{k,l}) \otimes \beta_{k,l} \right\|_F^2,$$

since we are looking for the subspace that captures most of the energy. Expanding the squared Frobenius norm leads to:

$$\begin{aligned}
E_0 &= \operatorname{argmax}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{\substack{l \in L \\ k_1 \in K \\ k_2 \in K}} \langle \Pi_E(\alpha_{k_1, l}) \otimes \beta_{k_1, l}, \Pi_E(\alpha_{k_2, l}) \otimes \beta_{k_2, l} \rangle \\
&= \operatorname{argmax}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{\substack{l \in L \\ k_1 \in K \\ k_2 \in K}} \langle \Pi_E(\alpha_{k_1, l}), \Pi_E(\alpha_{k_2, l}) \rangle \langle \beta_{k_1, l}, \beta_{k_2, l} \rangle \\
&= \operatorname{argmax}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{l \in L} \langle \Pi_E(A_l)^T \Pi_E(A_l), B_l^T B_l \rangle.
\end{aligned}$$

Recalling that $\tilde{A}_l = A_l \Psi_{B_l} = [\tilde{\alpha}_{1, l}, \dots, \tilde{\alpha}_{|K|, l}]$ and $B_l^T B_l = \Psi_{B_l} \Sigma_l \Psi_{B_l}^T$, this leads to:

$$\begin{aligned}
E_0 &= \operatorname{argmax}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{l \in L} \langle \Psi_{B_l}^T \Pi_E(A_l)^T \Pi_E(A_l) \Psi_{B_l}, \Sigma_l \rangle \\
&= \operatorname{argmax}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{l \in L} \sum_{k \in K} \sigma_{k, l}^2 \|\Pi_E(\tilde{\alpha}_{k, l})\|_2^2 \\
&= \operatorname{argmax}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{l \in L} \sum_{k \in K} \|\Pi_E(\sigma_{k, l} \tilde{\alpha}_{k, l})\|_2^2 \\
&= \operatorname{SVD}_{|I|}(\sigma_{k, l} \tilde{\alpha}_{k, l}, k \in K, l \in L).
\end{aligned}$$

□

Lemma 1 shows that the computational cost of this initialization is dominated by the computation of two singular value decompositions: the first matrix is of size $m \times |L||K|$ and the second is of size $n \times |L||K|$. Depending on the cardinality $|L||K|$, this can be achieved either with standard linear algebra routines, or with randomized SVDs [20]. In the applications that we consider here, n and m would typically be very large, while the number of samples $|L|$ and the rank of the tensors $|K|$ are expected to be small. In that situation, the computation can be performed even for very large scale applications.

Apart from being computable, the HOSVD presents additional advantages: the cost function can be controlled by the tail of the square singular values and running the alternating least squares on top of this initialization procedure ensures that the cost function will not increase above this upper-bound [13]. In addition, the ranks $|I|$ and $|J|$ of the decomposition can be chosen automatically according to the decay of the singular values in the HOSVD.

The partial optimization problems The ALS algorithm requires to solve the two following partial optimization problem

$$\operatorname{argmin}_{(e_i)_{i \in I} \in \mathcal{E}_{|I|}} \phi(E, F_t), \tag{8}$$

and

$$\underset{(f_j)_{j \in J \in \mathcal{F}_{|J|}}}{\operatorname{argmin}} \phi(E_{t+1}, F), \quad (9)$$

where $E_t = [e_{t,1}, \dots, e_{t,|I|}]$ and $F_t = [f_{t,1}, \dots, f_{t,|J|}]$ are the output of Algorithm 1 at iteration $t \geq 0$. Solving the two subproblems requires the computation of two SVDs as in the previous section.

Lemma 2 (Partial optimization problem (8) and (9)). *Let $\tilde{A}_l = A_l(B_l^T F_t) = [\tilde{\alpha}_{1,l}, \dots, \tilde{\alpha}_{|J|,l}]$ and $\tilde{B}_l = B_l(A_l^T E_{t+1}) = [\tilde{\beta}_{1,l}, \dots, \tilde{\beta}_{|I|,l}]$. For all $t > 0$ we have*

$$E_{t+1} = \operatorname{SVD}_{|I|}(\tilde{\alpha}_{j,l}, j \in J, l \in L)$$

and

$$F_{t+1} = \operatorname{SVD}_{|J|}(\tilde{\beta}_{i,l}, i \in I, l \in L)$$

Proof. We concentrate on E_{t+1} only, since the proof for F_{t+1} is similar.

The projection $\Pi_{E \otimes F_t}(S_l)$ of the operator S_l onto the subspace $E \otimes F_t$ can be expressed as follows

$$\begin{aligned} \Pi_{E \otimes F_t}(S_l) &= \sum_{k \in K} \Pi_E(\alpha_{k,l}) \otimes \Pi_{F_t}(\beta_{k,l}) \\ &= \sum_{k \in K} \sum_{i \in I} \langle \alpha_{k,l}, e_i \rangle e_i \otimes \sum_{j \in J} \langle \beta_{k,l}, f_{t,j} \rangle f_{t,j} \\ &= \sum_{i \in I} \sum_{j \in J} \left\langle \sum_{k \in K} \langle \beta_{k,l}, f_{t,j} \rangle \alpha_{k,l}, e_i \right\rangle e_i \otimes f_{t,j} \\ &= \sum_{j \in J} \Pi_E(\tilde{\alpha}_{j,l}) \otimes f_{t,j}. \end{aligned}$$

Replacing this expression in (8), leads to solve the problem (6) again, with the difference that the second factors ($f_{t,j}$) form an orthogonal family. This allows to avoid the diagonalization step of Lemma 1:

$$\begin{aligned} E_{t+1} &= \operatorname{argmax}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{l \in L} \left\| \sum_{j \in J} \Pi_E(\tilde{\alpha}_{j,l}) \otimes f_{t,j} \right\|_F^2 \\ &= \operatorname{argmax}_{E \in \mathcal{E}_{|I|}} \frac{1}{2} \sum_{l \in L} \sum_{j \in J} \|\Pi_E(\tilde{\alpha}_{j,l})\|_F^2 \\ &= \operatorname{SVD}_{|I|}(\tilde{\alpha}_{j,l}, j \in J, l \in L). \end{aligned}$$

□

4.1.2 Hierarchical matrices

The results presented in the previous paragraph can readily be applied to the case of hierarchical decompositions. To this end, let $(\mathcal{T}_p)_{p \in P}$ denote a block-partition of $X \times Y$ [19]:

- each \mathcal{T}_p has a product structure: $\mathcal{T}_p = X_p \times Y_p$ for some $X_p \subset X$ and $Y_p \subset Y$.
- $\mathcal{T}_{p_1} \cap \mathcal{T}_{p_2} = \emptyset$ if $p_1 \neq p_2$.
- $X \times Y = \cup_{p \in P} \mathcal{T}_p$.

We assume that the subspaces \mathcal{A} and \mathcal{B} defining the operators S_l (see Assumption 1) encode support constraints. For each $l \in L$, the k -th tensor $\alpha_{k,l} \otimes \beta_{k,l}$ should satisfy:

$$\exists p \in P, \text{supp}(\alpha_{k,l} \otimes \beta_{k,l}) \subseteq \mathcal{T}_{p_k}$$

for some $p_k \in P$. In order to apply the proposed ideas, we can first define two vectors of ranks $(q_p)_{p \in P}$ and $(r_p)_{p \in P}$ and generate an estimate $\hat{\mathcal{H}}$ of \mathcal{H} of the form

$$\hat{\mathcal{H}} = \sum_{p \in P} E_p \otimes F_p,$$

with $\dim(E_p) = q_p$ and $\dim(F_p) = r_p$.

The estimation of the subspaces \hat{E}_p and \hat{F}_p can then be achieved with the same methodology as the one described for orthogonal matrices. In this setting, we can use HOSVD algorithm for each sub-blocks, this implies computing $|P|$ SVDs with lower dimensional matrices (depending of the size of support).

4.1.3 Non-negative decompositions

A common choice of family is the set of non-negative vectors, that is

$$\mathcal{E}_{|I|} = \{e \in \mathbb{R}_+^m, \|e\|_2 = 1\}^{|I|}$$

and

$$\mathcal{F}_{|J|} = \{f \in \mathbb{R}_+^n, \|f\|_2 = 1\}^{|J|},$$

where \mathbb{R}_+^m denotes the set of nonnegative vectors of \mathbb{R}^m . Problems of the form (5) can then be solved with approaches such as [7, 27, 11]. We do not explore this possibility further in this paper.

4.2 Theoretical guarantees

We are now ready to establish the theoretical guarantees of the estimator (\hat{E}, \hat{F}) under additional assumptions on the sampling model.

Assumption 2 (Sampling model). *The operators S_l are i.i.d. copies of a random operator S with $\|S\|_F \leq r$ almost surely. Let $\Phi(E, F) \stackrel{\text{def.}}{=} \frac{1}{2} \mathbb{E} \left(\|\Pi_{E \otimes F}(S) - S\|_F^2 \right)$.*

We assume that:

$$\inf_{(E \in \mathcal{E}_{|I|}, F \in \mathcal{F}_{|J|})} \Phi(E, F) = r^2 \kappa(I, J). \quad (10)$$

The scaling in r^2 in equation (10) is natural: if the random operator S is scaled by a constant factor, so will the approximation error. The bound (10) provides the best achievable estimate of subspace.

4.2.1 Arbitrary bounded errors

We let $(\widehat{E}_L, \widehat{F}_L)$ denote the solution of (5). In practice, we do not directly observe the operator S_l , but only an approximate version S_l^K of it. Hence we need to estimate the approximation error.

Assumption 3 (Approximation error). *The operators S_l^K satisfy the following inequality : $\|S_l^K - S_l\|_F \leq \kappa(K)\|S_l\|_F$ with $\kappa(K) \leq 1$.*

Theorem 1. *Assume that Assumptions 2 and 3 hold, then:*

$$\mathbb{P}\left(\Phi(\widehat{E}_L, \widehat{F}_L) \leq 6r^2 \max(\kappa(K), \kappa(I, J))\right) \geq 1 - 2 \exp\left(-8|L| \max(\kappa(K), \kappa(I, J))^2\right).$$

We first discuss the consequences of this Theorem 1 prior to detailing its proof. In case the relative approximation error $\kappa(K)$ is too large w.r.t. to $\kappa(I, J)$ there will be no guarantee to reach $\Phi(E^*, F^*)$ since the best achievable error will be of the order $r^2\kappa(K)$. This bound can be achieved with probability $1 - \delta$ by choosing $|L| = \frac{\log(2/\delta)}{8\kappa(K)^2}$. However, when the approximation gets finer i.e. $\kappa(K) < \kappa(I, J)$, the estimator $(\widehat{E}_L, \widehat{F}_L)$ becomes as good as possible up to a constant. This bound can be achieved with probability $1 - \delta$ by choosing $|L| = \frac{\log(2/\delta)}{8\kappa(I, J)^2}$.

Proof. We let

$$\Phi_L(E, F) \stackrel{\text{def.}}{=} \frac{1}{2|L|} \sum_{l \in L} \left(\|\Pi_{E \otimes F}(S_l) - S_l\|_F^2 \right)$$

and

$$\Phi_L^K(E, F) \stackrel{\text{def.}}{=} \frac{1}{2|L|} \sum_{l \in L} \left(\|\Pi_{E \otimes F}(S_l^K) - S_l^K\|_F^2 \right).$$

Step 1. We first control the bias term as follows

$$|\Phi_L^K(E, F) - \Phi_L(E, F)| \leq 3r^2\kappa(K)/2. \quad (11)$$

Let $G = E \otimes F$ and G^\perp denote its orthogonal complementary with respect to the Frobenius inner-product over the space of operators. We let $D_l = S_l^K - S_l$ and notice that $\|D_l\|_F \leq \kappa(K)\|S_l\|_F$ by Assumption 3. Now, we can decompose S_l as $S_l = S_l^G + S_l^{G^\perp}$ and S_l^K as $S_l^K = S_l^G + S_l^{G^\perp} + D_l^G + D_l^{G^\perp}$. This leads to

$$\|\Pi_G(S_l) - S_l\|_F^2 = \|S_l^{G^\perp}\|_F^2$$

and

$$\|\Pi_G(S_l^K) - S_l^K\|_F^2 = \|S_l^{G^\perp} + D_l^{G^\perp}\|_F^2.$$

So that

$$\begin{aligned} & \left| \|\Pi_G(S_l^K) - S_l^K\|_F^2 - \|\Pi_G(S_l) - S_l\|_F^2 \right| \\ &= \left| 2\langle S_l^{G^\perp}, D_l^{G^\perp} \rangle + \|D_l^{G^\perp}\|_F^2 \right| \\ &\leq 2r^2\kappa(K) + r^2\kappa(K)^2 \leq 3r^2\kappa(K). \end{aligned}$$

By summing this inequality over $l \in L$, we get the inequality (11).

Step 2. As in the previous step, we let $G = E \otimes F$. We show here that

$$\mathbb{P}(|\Phi_L - \Phi| \geq t) \leq 2 \exp\left(-\frac{8|L|t^2}{r^4}\right). \quad (12)$$

Let us introduce the random variable $Z_l = \|\Pi_G(S_l) - S_l\|_F^2 = \|S_l^{G^\perp}\|_F^2$. We have $\mathbb{E}(Z_l) = \Phi$ and by Assumption 2, we have $Z_l \in [0, r^2]$. Let $X \stackrel{\text{def}}{=} \sum_{l \in L} (Z_l - \mathbb{E}(Z_l))$. We have $X/(2|L|) = \Phi_L - \Phi$ and Hoeffding's inequality [6, Thm 2.8] ensures that for all $t > 0$ the random variable X satisfies

$$\mathbb{P}(|X| \geq t) \leq 2 \exp\left(-\frac{2t^2}{|L|r^4}\right).$$

Step 3. We are now ready to conclude the proof. We have

$$|\Phi_L^K - \Phi| \leq |\Phi_L^K - \Phi_L| + |\Phi_L - \Phi|.$$

The problem (10) has at least one solution denoted (E^*, F^*) . Indeed, the finite dimensional vector spaces E and F can be parameterized by $|I|$ and $|J|$ unit vectors. The tensor product of $|I||J|$ unit balls is a compact set and the function Φ is continuous, ensuring the existence of a minimizer. We get:

$$\begin{aligned} \Phi(\widehat{E}_L, \widehat{F}_L) &\leq \Phi_K^L(\widehat{E}_L, \widehat{F}_L) + |\Phi_L^K - \Phi_L| + |\Phi_L - \Phi| \\ &\leq \Phi_K^L(E^*, F^*) + 3/2r^2\kappa(K) + |\Phi_L - \Phi| \\ &\leq \Phi(E^*, F^*) + 3r^2\kappa(K) + 2|\Phi_L - \Phi|. \end{aligned}$$

Using the inequality (12), we get for all $t > 0$:

$$\mathbb{P}\left(\Phi(\widehat{E}_L, \widehat{F}_L) \leq r^2(\kappa(I, J) + 3\kappa(K)) + 2t\right) \geq 1 - 2 \exp\left(-\frac{8|L|t^2}{r^4}\right).$$

The first part of the theorem is obtained by selecting $t = r^2 \max(\kappa(K), \kappa(I, J))$. \square

4.2.2 Random errors

The bound in Theorem 1 may look a bit disappointing since it is impossible to reach the absolute best error $r^2\kappa(I, J)$. This is due to the fact that the approximation errors $D_l = S_l^K - S_l$ can be adversarial and create a bias in the estimation. If we add randomness assumptions on these errors, the situation can improve. We illustrate it below with random isotropic errors.

Theorem 2. *Suppose that assumptions 2 and 3 hold. Assume furthermore that the errors D_l have an isotropic distribution with $\mathbb{E}(\|D_l\|_F^2) = R^2$ and $\|D_l\|_F^2 \leq \kappa^2(K)r^2$ almost surely then:*

$$\mathbb{P}\left(\Phi(\widehat{E}_L, \widehat{F}_L) \leq r^2(\kappa(I, J) + \epsilon)\right) \geq 1 - 8 \exp\left(-\frac{2|L|\epsilon^2}{(6\kappa(K) + 1)^2}\right).$$

Theorem 2 shows that under isotropic random approximation errors, the estimator $(\widehat{E}_L, \widehat{F}_L)$ can become arbitrarily good. This bound can be achieved with probability $1 - \delta$ by choosing $|L| = \frac{(6\kappa(K)+1)^2}{2\epsilon^2} \log\left(\frac{8}{\delta}\right)$.

Proof. We let,

$$\Phi_L(E, F) \stackrel{\text{def.}}{=} \frac{1}{2|L|} \sum_{l \in L} \left(\|\Pi_{E \otimes F}(S_l) - S_l\|_F^2 \right)$$

and

$$\Phi_L^K(E, F) \stackrel{\text{def.}}{=} \frac{1}{2|L|} \sum_{l \in L} \left(\|\Pi_{E \otimes F}(S_l^K) - S_l^K\|_F^2 \right) - \frac{R^2}{mn} (mn - |I||J|)R^2,$$

where m and n are the dimension of the space \mathcal{B}_m and \mathcal{B}_n respectively.

The difference compared to the previous proof is that we can now bound $|\Phi_L^K - \Phi_L|$ by a quantity that vanishes with the number of observations L . For any pair of subspaces (E, F) , we have:

$$\mathbb{P} \left(|\Phi_L^K(E, F) - \Phi_L(E, F)| \geq t \right) \leq 2 \exp \left(-\frac{|L|t^2}{18r^4\kappa(K)^2} \right).$$

To prove this statement, let $G = E \otimes F$. We get:

$$|\Phi_L^K(E, F) - \Phi_L(E, F)| = \left| \frac{1}{2|L|} \sum_{l \in L} \left(2\langle S_l^{G^\perp}, D_l^{G^\perp} \rangle + \|D_l^{G^\perp}\|_F^2 \right) - \frac{R^2}{mn} (mn - |I||J|) \right|.$$

Letting $Z_l^{G^\perp} = 2\langle S_l^{G^\perp}, D_l^{G^\perp} \rangle + \|D_l^{G^\perp}\|_F^2$, we get $\mathbb{E}(Z_l^{G^\perp}) = (mn - |I||J|)R^2$ since D_l is isotropic. Indeed, $\mathbb{E} \left(2\langle S_l^{G^\perp}, D_l^{G^\perp} \rangle \right) = 0$ since $\mathbb{E}(D_l) = 0$, and by letting Π_{G^\perp} denote the projection onto G^\perp we get:

$$\begin{aligned} & \mathbb{E} \left(\|D_l^{G^\perp}\|_F^2 \right) \\ &= \mathbb{E} \left(\text{tr} \left(D_l^T \Pi_{G^\perp}^T \Pi_{G^\perp} D_l \right) \right) \\ &= \mathbb{E} \left(\text{tr} \left(\Pi_{G^\perp} D_l D_l^T \Pi_{G^\perp} \right) \right) = \text{tr} \left(\Pi_{G^\perp} \mathbb{E}(D_l D_l^T) \Pi_{G^\perp} \right) \\ &= \frac{R^2}{mn} \text{tr} \left(\Pi_{G^\perp}^T \Pi_{G^\perp} \right) = \frac{R^2}{mn} (mn - |I||J|) \end{aligned}$$

Noticing that $|Z_l^{G^\perp}| \leq r^2(2\kappa(K) + \kappa(K)^2) \leq 3r^2\kappa(K)$, we can use Hoeffding's inequality:

$$\mathbb{P} \left(\left| \frac{1}{2|L|} \sum_{l \in L} Z_l^{G^\perp} - \mathbb{E}(Z_l^{G^\perp}) \right| \geq t \right) \leq 2 \exp \left(-\frac{2|L|t^2}{9r^4\kappa(K)^2} \right).$$

We are now ready to conclude the proof. Similarly to the previous proof, we get:

$$\Phi(\widehat{E}_L, \widehat{F}_L) \leq \Phi_L^K(E^*, F^*) + |\Phi_L^K(\widehat{E}_L, \widehat{F}_L) - \Phi_L(\widehat{E}_L, \widehat{F}_L)| + |\Phi_L(\widehat{E}_L, \widehat{F}_L) - \Phi(\widehat{E}_L, \widehat{F}_L)|.$$

Using a union bound argument (given a set of events $(A_i)_{i \in \mathbb{N}}$, we have $\mathbb{P}(\cup_{i \in \mathbb{N}} A_i) \leq \sum_{i \in \mathbb{N}} \mathbb{P}(A_i)$), we get:

$$\mathbb{P}\left(\Phi(\widehat{E}_L, \widehat{F}_L) \leq \Phi_L^K(E^*, F^*) + t + t'\right) \geq 1 - 2 \exp\left(-\frac{2|L|t^2}{9r^4\kappa(K)^2}\right) - 2 \exp\left(-\frac{8|L|t'^2}{r^4}\right).$$

Using another union bound argument, we get:

$$\mathbb{P}\left(\Phi(\widehat{E}_L, \widehat{F}_L) \leq \Phi(E^*, F^*) + 2t + 2t'\right) \geq 1 - 4 \exp\left(-\frac{2|L|t^2}{9r^4\kappa(K)^2}\right) - 4 \exp\left(-\frac{8|L|t'^2}{r^4}\right).$$

By taking $t = 6\kappa(K)t'$ we get

$$\mathbb{P}\left(\Phi(\widehat{E}_L, \widehat{F}_L) \leq \Phi(E^*, F^*) + t'(12\kappa(K) + 2)\right) \geq 1 - 8 \exp\left(-\frac{8|L|t'^2}{r^4}\right).$$

Letting $\epsilon > 0$ and setting $t' = \frac{r^2\epsilon}{12\kappa(K)+2}$, we get:

$$\mathbb{P}\left(\Phi(\widehat{E}_L, \widehat{F}_L) \leq \Phi(E^*, F^*) + \epsilon r^2\right) \geq 1 - 8 \exp\left(-\frac{2|L|\epsilon^2}{(6\kappa(K) + 1)^2}\right).$$

Finally, given $\delta > 0$, we can select $|L| = \frac{(6\kappa(K)+1)^2}{2\epsilon^2} \log\left(\frac{8}{\delta}\right)$ so that the following holds true:

$$\mathbb{P}\left(\Phi(\widehat{E}_L, \widehat{F}_L) \leq \Phi(E^*, F^*) + r^2\epsilon\right) \geq 1 - \delta.$$

This concludes the proof. \square

5 Subset estimation and projection

5.1 Convex hull estimator

In this section, we assume that $(S_l)_{l \in L}$ are i.i.d. copies of the random operator S . We assume that the distribution of S is uniform over a convex, compact and non-empty set \mathcal{C} . Letting Π_L denote the projector onto $\widehat{\mathcal{H}}_L = \widehat{E}_L \otimes \widehat{F}_L$, we propose to construct an estimate $\widehat{\mathcal{C}}_L^{K, \Pi}$ of \mathcal{C} , by taking the convex hull of the projected and observed operators

$$\widehat{\mathcal{C}}_L^{K, \Pi} \stackrel{\text{def.}}{=} \text{conv}(\Pi_L(S_l^K), l \in L).$$

We can only expect $\widehat{\mathcal{C}}_L^{K, \Pi}$ to approximate $\Pi_L(\mathcal{C})$, and not \mathcal{C} directly, since some information is lost by the projection. The following proposition summarizes the rate of convergence of $\widehat{\mathcal{C}}_L^{K, \Pi}$.

Proposition 1. *Under the assumptions 2 and 3, we get the following result*

$$\mathcal{D}(\widehat{\mathcal{C}}_L^{K, \Pi}, \Pi_L(\mathcal{C})) \leq r\kappa(K) + \mathcal{O}_{\mathbb{P}}\left(\left(\frac{\ln|L|}{|L|}\right)^{-\frac{1}{\alpha}}\right),$$

where \mathcal{D} denotes the Hausdorff distance between sets and

- $\alpha = d$ if \mathcal{C} is a polytope,
- $\alpha = \frac{d+1}{2}$ if \mathcal{C} has C^3 boundary and positive curvature everywhere.

Proof. Step 1. The difficult part of this inequality is the rightmost term, which is due to [8] (Theorem 11). With our notation, his main result states that $\mathcal{D}(\widehat{\mathcal{C}}_L, \mathcal{C}) = \mathcal{O}_{\mathbb{P}}\left(\left(\frac{\ln|L|}{|L|}\right)^{-\frac{1}{\alpha}}\right)$, where $\widehat{\mathcal{C}}_L = \text{conv}(S_l, l \in L)$.

Step 2. In order to obtain our result, we first observe that since Π_L is a projection, it is also a contraction and $\mathcal{D}(\Pi_L(\widehat{\mathcal{C}}_L), \Pi_L(\mathcal{C})) \leq \mathcal{D}(\widehat{\mathcal{C}}_L, \mathcal{C})$.

Step 3. Now, let $\widehat{\mathcal{C}}_L^K \stackrel{\text{def.}}{=} \text{conv}(S_l^K, l \in L)$. We have

$$\mathcal{D}(\widehat{\mathcal{C}}_L, \widehat{\mathcal{C}}_L^K) \leq r\kappa(K) \quad (13)$$

Indeed, the distance function $d_{\widehat{\mathcal{C}}_L^K}(H) = \inf_{H' \in \widehat{\mathcal{C}}_L^K} \|H - H'\|_2$ is convex. Hence, the problem $\sup_{H \in \widehat{\mathcal{C}}_L} d_{\widehat{\mathcal{C}}_L^K}(H)$ appearing in the definition of the Hausdorff distance consists of finding the maximum of a convex function over a convex set. Hence the maximum is attained at an extremal point of $\widehat{\mathcal{C}}_L^K$, i.e. at a point S_l^K . All those points satisfy $\|S_l^K - S_l\|_F \leq r\kappa(K)$, hence $\sup_{H \in \widehat{\mathcal{C}}_L} d_{\widehat{\mathcal{C}}_L^K}(H) \leq r\kappa(K)$. A similar reasoning on the other part of the distance yields the inequality (13). Since Π_L is a contraction, we also get $\mathcal{D}(\Pi_L(\widehat{\mathcal{C}}_L), \Pi_L(\widehat{\mathcal{C}}_L^K)) \leq r\kappa(K)$.

Step 4. To conclude, we use the fact that the Hausdorff distance satisfies the triangle inequality. Hence:

$$\begin{aligned} \mathcal{D}(\widehat{\mathcal{C}}_L^{K, \Pi}, \Pi_L(\mathcal{C})) &\leq \mathcal{D}(\widehat{\mathcal{C}}_L^{K, \Pi}, \Pi_L(\widehat{\mathcal{C}}_L)) \\ &\quad + \mathcal{D}(\Pi_L(\widehat{\mathcal{C}}_L), \Pi_L(\mathcal{C})) \\ &\leq r\kappa(K) + \mathcal{O}_{\mathbb{P}}\left(\left(\frac{\ln|L|}{|L|}\right)^{\frac{1}{\alpha}}\right). \end{aligned}$$

□

Remark 1. *There are different ways to control the distance between sets. Another possibility is to use the Nikodym metric, i.e. the relative difference of volume between $\widehat{\mathcal{C}}_L^{K, \Pi}$ and $\Pi_L(\mathcal{C})$. For this metric, it can be shown that the convex hull estimator is a minimax operator (i.e. that it is optimal uniformly on the class of convex bodies) and we also obtain a convergence rate of the form $\mathcal{O}_{\mathbb{P}}(|L|^{-2/d+1})$ for a convex set \mathcal{C} with C^3 boundary and positive curvature everywhere.*

Remark 2. *Proposition 1 only characterizes the asymptotic behavior of this estimator. This result should be taken carefully since the constants in the $\mathcal{O}_{\mathbb{P}}$ depend on the geometry of the convex set \mathcal{C} . In particular, the sharper the corners of \mathcal{C} , the larger the constant.*

5.2 A projection algorithm

In what follows, we let $\widehat{\mathcal{C}} = \widehat{\mathcal{C}}_L^{K,\Pi}$ to simplify the notation. In the framework of blind inverse problems (see equation (1)), the knowledge of the convex set $\widehat{\mathcal{C}}$ may lead to the resolution of variational problems of the form

$$\min_{H \in \widehat{\mathcal{C}}, u \in \mathcal{W}} \frac{1}{2} \|Hu - y\|_2^2. \quad (14)$$

A critical tool to solve (14) is a projection operator $\Pi_{\widehat{\mathcal{C}}}$ onto the set $\widehat{\mathcal{C}}$. For instance, it would allow using a projected gradient descent. Let $H \in \Xi$ and $S = \mathcal{R}(H)$. The projection is defined as follows:

$$\Pi_{\widehat{\mathcal{C}}}(S) = \operatorname{argmin}_{\lambda \in \Delta_{|L|}} \frac{1}{2} \|M\lambda - S\|_F^2, \quad (15)$$

where $M : \lambda \rightarrow \sum_{l \in L} \lambda_l \Pi_L(S_l^K)$.

Depending on the number of samples $|L|$, different algorithms can be used to solve (15). For small $|L|$, interior point methods [29] are an excellent candidate, since they lead to high precision solutions in short computation times. For larger $|L|$, they become intractable and it is then possible to use lighter, but less precise first order solutions. We detail such an approach below.

First, we let $\tau = 1/\|M^*M\|_F$. This quantity can be computed using a power method for instance. We can then use the accelerated proximal gradient [28] descent described in Algorithm 2.

Algorithm 2 Projection onto convex hull of operators

INPUT: $\Pi_L(S_l^K)$, S , initial guess $\lambda_0 \in \Delta_{|L|}$.


OUTPUT: Projection of S onto $\widehat{\mathcal{C}}$.

- 1: **procedure**
 - 2: **for** $k = 1, 2, \dots, k_{end}$ **do**
 - 3: $\tilde{\lambda}_k = \Pi_{\Delta_{|L|}}(\lambda_k - \tau M^*(M\lambda_k - S))$
 - 4: $\lambda_{k+1} = \tilde{\lambda}_k + \frac{k-1}{k+2} (\tilde{\lambda}_k - \tilde{\lambda}_{k-1})$
 - 5: **end for**
 - 6: Return $M\lambda_{k_{end}}$
 - 7: **end procedure**
-

The projection on the $(|L|-1)$ -dimensional simplex can be computed in linear time and Algorithm 2 ensures that the cost function decays as $O(1/k^2)$. The matrix M^*M can be precomputed with a numerical complexity in $O(|L|^2(|I|^2n + |J|^2m))$. The product M^*S can also be computed efficiently, for operators S given in a tensor form. This is for instance the case if $S \in \widehat{\mathcal{H}}_L$.

6 Numerical experiments

In this section we illustrate the previous methods with a few numerical examples.



(a) Kernel operator 1. (b) SVIR operator 1. (c) Kernel operator 2. (d) SVIR operator 2.

Figure 1: Kernel and SVIR of two different inverse diffusion operators.

6.1 Approximation rate and computation times

6.1.1 The setting

We start with a one dimensional diffusion equation as introduced in Section 1.1. Our main aim here is to illustrate the computational complexity of the approach. We take $\mathcal{B}_n = \mathcal{B}_m = \mathbb{R}^n$ with $n = m$. We define the operator ∇ with forward finite differences and homogeneous Neumann boundary conditions. The divergence operator $\text{div} = -\nabla^*$, where ∇^* is the adjoint of ∇ .

We wish to find a family of estimators of the mapping $f \mapsto u$ for the following equation

$$\text{div}(c\nabla u) = f, \forall f \in \mathbb{R}^n,$$

and for diffusion coefficients $c \in \mathbb{R}^n$ living in a subset Ω of nonnegative vectors. We assume that we can access $|L|$ observations of c , denoted c_l for $l \in L$. We let

$$H_l : \mathbb{R}^n \mapsto \mathbb{R}^n \\ f \mapsto (\text{div}(c_l \nabla))^+ f$$

denote the operators of interest, where $+$ denotes the pseudo-inverse. In our simulations we consider diffusion coefficients c_l of the form:

$$c_l(x) = 3 + \sum_{p \in P} w_{l,1}(p) \cos(2\pi px) + w_{l,2}(p) \sin(2\pi px), \forall x \in \mathbb{R}^n,$$

with $w_{l,1}, w_{l,2}$ taken uniformly at random in the $|P| - 1$ -dimensional simplex $\Delta_{|P|-1}$. We assume that the operators H_l are given in a product-convolution form, or equivalently that their SVIR S_l can be written as $S_l = \sum_{k \in K} \alpha_{k,l} \otimes \beta_{k,l}$. In our numerical experiments, we compute the factors $\alpha_{k,l}$ and $\beta_{k,l}$ using a SVD of S_l . This is feasible since we work in 1D. The number of factors in the SVD is set to $|K| = 20$ which is enough to capture 97% percent of the energy on average. Two instances of operators H_l and their SVIR S_l are displayed in Figure 1.

6.1.2 Description of the approaches

Given $|I|$ and $|J|$, our aim is to find two families $(e_i)_{i \in I} \in \mathcal{E}_{|I|}$ and $(f_j)_{j \in J} \in \mathcal{F}_{|J|}$, with $\mathcal{E}_{|I|}$ and $\mathcal{F}_{|J|}$ defined as the sets of orthonormal families, see equations (3) and (4). We compare four approaches to estimate the subspace \mathcal{H} .

- SVD: We concatenate the vectorized representation of H_l in a matrix M . The family $(e_i)_{i \in I}$ is set to be the first $|I|$ left-eigenvectors, and the family $(f_j)_{j \in J}$ is set to be the first $|J|$ right-eigenvectors of M . This approach is optimal in terms of Frobenius norm but can only be applied because we work in a low dimensional 1D setting.
- DCT: We set e_i and f_j as the first elements of the discrete cosine transform, i.e. $e_i(x) = \cos(\pi/(n-1)ix)$ and $f_j(y) = \cos(\pi/(m-1)iy)$ with n corresponding to the number of elements in the discretization. The family $(e_i \otimes f_j)_{i \in I, j \in J}$ is in tensor product form and it is orthogonal, which allows making very fast computations.
- HOSVD: implements the decomposition in equations (6) and (7).
- ALS: use the Alternating Least Square Algorithm 1 with 15 iterations and the HOSVD as an initialization.

We first compare the four different methods in terms of their approximation quality for $|L| = 50$ observations. We evaluate the average relative projection error defined by $\mathbb{E} \left(\frac{\|H - \Pi_{\mathcal{H}}(H)\|_F}{\|H\|_F} \right)$. It can be evaluated through a Monte-Carlo simulation. Figure 2a shows the relative error for the different methods and various sizes $|I|$ with $|I| = |J|$.

The approximation rate given by the SVD is upper-bounded by the approximation properties of the considered family of operators. This is an illustration of Theorem 1 which describes the behavior of the approximation rate in terms of the constants $\kappa(I, J)$ and $\kappa(K)$. In this example, we distinguish two regimes: when $|I||J| < |L|$ the approximation rate is bounded by the constant $\kappa(I, J)$, and when $|I||J| \geq |L|$, the approximation rate is bounded by the constant $\kappa(K)$.


6.1.3 Computing times

We now examine the computational time for each method in Figure 2b.


The efficiency of the SVD has to be balanced by its important computational time. It becomes completely impractical on a usual workstation when the dimension n of the space \mathcal{B}_n is larger than 10^5 . We also observe that using the ALS algorithm instead of the HOSVD leads to negligible gains, despite a significantly higher computational burden. The runtime is basically proportional to the number of iterations.

6.2 Blind deblurring

In this section we apply the proposed method of subspace estimation to solve a blind deblurring problem. We use simulated operators and grayscale images.



(a) Relative approximation error versus the dimension $|I||J|$ of each basis.



(b) Computation time in seconds versus the dimension n of the problem (the operators $(H_l)_{l \in L}$ are of size $n \times n$).

Figure 2: Numerical behavior for 1D operators.

The setting We let $\mathcal{B}_n = \mathcal{B}_m = \mathbb{R}^{n \times n}$ with $n = 64$ and set $\mathcal{A} = \mathbb{R}^{n \times n}$ and $\mathcal{B} = \mathbb{R}^{n \times n}$. We generate random space varying impulse responses of the form

$$S_l = \sum_{k=1}^{|K|} \alpha_k \otimes \beta_{k,l}.$$

In the following, we set $|K| = 5$, let $\theta_k = \frac{\pi k}{6}$ and set for all $l \in L$ and all $(x_1, x_2) \in \{1, \dots, n\}^2$

$$\alpha_{k,l}(x_1, x_2) = \exp\left(-\frac{(\cos(\theta_k)x_1 - \sin(\theta_k)x_2)^2}{8} - \frac{(\sin(\theta_k)x_1 + \cos(\theta_k)x_2)^2}{2}\right).$$

This corresponds to anisotropic Gaussian functions rotated differently. We generate the maps $\beta_{k,l}$ as follows. For each $l \in L$:

1. We generate a matrix of $\mathbb{R}^{n \times n}$ where each element is a uniform random number in $[0, 1]$, independent of the others.
2. We compute a discrete convolution of this random matrix with an isotropic Gaussian blur. We then rescale it in $[0, 1]$, producing a discrete random field $f_l \in [0, 1]^{n \times n}$.




Figure 3: Examples of factors. Top: the full collection of $(\alpha_k)_{k \in K}$. Bottom: the factors $(\beta_{k,1})_{k \in K}$.




Figure 4: Learned family $(e_i)_{i \in I}$.

3. We then partition the domain Ω into $|K|$ sets $(\omega_{k,l})_{k \in K}$ defined as

$$\omega_{k,l} = f_l^{-1}([(k-1)/|K|, k/|K|]).$$

4. Finally, the factors $\beta_{k,l}$ are defined as the indicators of $\omega_{k,l}$ convolved with a Gaussian kernel.

We display the elements α_k and $\beta_{k,1}$ in Figure 3.

The output of our algorithm With those definitions, we get a list of random product-convolution operators H_l defined by

$$H_l u = \sum_{k \in K} \alpha_k \star (\beta_{k,l} \odot u), \forall u \in \mathbb{R}^{n \times n}.$$

From the collection of $(S_l)_{l \in L}$ we can use the initialization of Algorithm 1 to estimate a subspace $\widehat{\mathcal{H}}_{I,J}$. In this paragraph we index the estimator by I and J .

The families $(e_i)_{i \in I}$ and $(f_j)_{j \in J}$ produced by the initialization of Algorithm 1 are displayed in Figure 4 and 5. The family $(e_i)_{i \in I}$ is an orthogonalization of the family α_k . The family $(f_j)_{j \in J}$ is quite similar to the first elements of a Fourier basis. This is to be expected since the functions $\beta_{k,l}$ are smooth and Fourier bases optimally encode smooth function spaces, see e.g. [30].




Figure 5: Learned family $(f_j)_{j \in J}$.

A blind-deblurring experiment Using the notation of the introduction, we set $|S| = |I||J|$ and let $P_s = e_i \otimes f_j$ for $s = (i, j)$ denote the elementary operators constituting the subspace $\widehat{\mathcal{H}} = \text{span}(P_1, \dots, P_{|S|})$. We let $Q \in \mathbb{R}^{m \times |T|}$ denote a matrix with columns $(q_t)_{t \in T}$ corresponding to elements of the discrete Haar wavelet basis with $|T| = 274$. We let

$$\mathcal{Q} = \{Q\beta, \beta \in \mathbb{R}^{|T|}\}$$

denote the subspace containing the images of interest. We let $\beta_0 \in \mathbb{R}^{|T|}$ denote the coefficients of the true image in the subspace \mathcal{Q} , and $H_0 \in \mathcal{H}$ the true operator that we want to recover. Finally we let

$$u_0 = H_0 v_0 + \eta,$$

where η is an additive white Gaussian noise. We display the true image v_0 in Figure 6a and the blurry-noisy image u_0 in Figure 6b.

We wish to solve the following bilinear inverse problem

$$\min_{\substack{v \in \mathcal{Q}, \\ H \in \widehat{\mathcal{H}}}} \|Hv - u_0\|_2^2. \quad (16)$$

Using the lifting and convex relaxation techniques described in [1] leads to

$$\min_{M \in \mathbb{R}^{|T| \times |S|}} \|M\|_* + \frac{\lambda}{2} \|\mathcal{W}(M) - u_0\|_2^2 \quad (17)$$

where $\lambda > 0$ is a regularization parameter. For a matrix M , $\|M\|_*$ denotes the nuclear norm, i.e. the sum of the singular values of M . This convex program is solved using an accelerated proximal gradient method.

We let the algorithm run until the cost function stops decreasing. In Figure 6, we compare the reconstructed images with three different estimations $\widehat{\mathcal{H}}$ of \mathcal{H} :

- When $\widehat{\mathcal{H}} = \mathcal{H}$, we use the full subspace to solve (17), this yields the result in Figure 6c. It takes 390 seconds to solve the problem and we obtain a SNR of -3.0dB. The reason for this failure is that the dimension of the subspace is too large, making it impossible to identify the true image.




Figure 6: Blind deblurring experiment with different estimated subspaces. 6c: full dimensional search space. 6d: $|I| = 5$ and $|J| = 30$. 6e: $|I| = 5$ and $|J| = 8$. Reducing the search space makes the problem identifiable and reduces the computing times.

- When $\hat{\mathcal{H}} = \hat{\mathcal{H}}_{I,J}$ with $|I| = 5$ and $|J| = 30$, i.e we use subspace of dimension 150 to solve (17), we obtain the result in Figure 6d. The computing times are divided by three (120 seconds) compared to the full subspace $\hat{\mathcal{H}} = \mathcal{H}$. More importantly, the method succeeds to recover the sharp image and we obtain a SNR of 26.7dB.
- When $\hat{\mathcal{H}} = \hat{\mathcal{H}}_{I,J}$ with $|I| = 5$ and $|J| = 8$, we subspace of dimension 40 to solve (17), leading to the results in Figure 6e. This time, the computing times decay to 31 seconds, which is 12 times faster than the case $\hat{\mathcal{H}} = \mathcal{H}$. We also obtain a good result with a SNR of 26.2dB.

7 Conclusion

We presented a scalable approach to estimate a low dimensional subspace of linear operators from a sampling set of operators expressed as low rank tensors. This formalism covers several commonly used operator structures such as hierarchical matrices, wavelet expansions or product-convolutions. An important application lies in the field of blind inverse problems where the prior knowledge

of a low dimensional subspace can make it possible to identify both the signal and the operator.

The principle outlined in this paper changes the way a device is calibrated: instead of characterizing it through a single operator, we propose to describe all its potential states. For instance, in microscopy, variations of temperature change the refraction indexes and hence the associated measurement operators. With the proposed approach, we can capture these variations and hence use more precise models for image reconstruction. We hypothesize that the proposed formalism can improve reconstructions for many other practical problems.

Acknowledgments

This work was supported by Fondation pour la Recherche Médicale (FRM grant number ECO20170637521 to V.D), by the défi IMAG'IN and by the ANR project OMS. The authors thank Thomas Bonnafont for a preliminary study of the problem and Thomas Mangeat for providing inspiring data from microscopy, which motivated this work.

References

- [1] Ali Ahmed, Benjamin Recht, and Justin Romberg. Blind deconvolution using convex programming. *IEEE Transactions on Information Theory*, 60(3):1711–1732, 2014.
- [2] Rick Beatson and Leslie Greengard. A short course on fast multipole methods. *Wavelets, multilevel methods and elliptic PDEs*, 1:1–37, 1997.
- [3] Mario Bebendorf. *Hierarchical matrices*. Springer, 2008.
- [4] Gregory Beylkin, Ronald Coifman, and Vladimir Rokhlin. Fast wavelet transforms and numerical algorithms i. *Communications on pure and applied mathematics*, 44(2):141–183, 1991.
- [5] Jérémie Bigot, Paul Escande, and Pierre Weiss. Estimation of linear operators from scattered impulse responses. *Applied and Computational Harmonic Analysis*, 2017.
- [6] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [7] Rasmus Bro and Sijmen De Jong. A fast non-negativity-constrained least squares algorithm. *Journal of Chemometrics: A Journal of the Chemometrics Society*, 11(5):393–401, 1997.
- [8] Victor-Emmanuel Brunel. Methods for estimation of convex sets. <https://arxiv.org/abs/1709.03137>, 2017.

- [9] Margaret Cheney, David Isaacson, and Jonathan C Newell. Electrical impedance tomography. *SIAM review*, 41(1):85–101, 1999.
- [10] Abdellah Chkifa, Albert Cohen, and Christoph Schwab. High-dimensional adaptive sparse polynomial interpolation and applications to parametric pdes. *Foundations of Computational Mathematics*, 14(4):601–633, 2014.
- [11] Andrzej Cichocki, Rafal Zdunek, Seungjin Choi, Robert Plemmons, and Shun-Ichi Amari. Non-negative tensor factorization using alpha and beta divergences. In *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, volume 3, pages III–1393. IEEE, 2007.
- [12] Pierre Comon. Tensors: a brief introduction. *IEEE Signal Processing Magazine*, 31(3):44–53, 2014.
- [13] Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. A multilinear singular value decomposition. *SIAM journal on Matrix Analysis and Applications*, 21(4):1253–1278, 2000.
- [14] Chris Ding and Jieping Ye. 2-dimensional singular value decomposition for 2d maps and images. In *Proceedings of the 2005 SIAM International Conference on Data Mining*, pages 32–43. SIAM, 2005.
- [15] Paul Escande and Pierre Weiss. Sparse wavelet representations of spatially varying blurring operators. *SIAM Journal on Imaging Sciences*, 8(4):2976–3014, 2015.
- [16] Paul Escande and Pierre Weiss. Approximation of integral operators using product-convolution expansions. *Journal of Mathematical Imaging and Vision*, 58(3):333–348, 2017.
- [17] Jeffrey A Fessler. Model-based image reconstruction for mri. *IEEE Signal Processing Magazine*, 27(4):81–89, 2010.
- [18] Rémi Gribonval, Gilles Chardon, and Laurent Daudet. Blind calibration for compressed sensing by convex optimization. In *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pages 2713–2716. IEEE, 2012.
- [19] Wolfgang Hackbusch. *Hierarchical matrices: algorithms and analysis*, volume 49. Springer, 2015.
- [20] Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [21] Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.

- [22] Xiaodong Li, Shuyang Ling, Thomas Strohmer, and Ke Wei. Rapid, robust, and reliable blind deconvolution via nonconvex optimization. *Applied and Computational Harmonic Analysis*, 2018.
- [23] Yanjun Li, Kiryung Lee, and Yoram Bresler. Optimal sample complexity for blind gain and phase calibration. *IEEE Trans. Signal Processing*, 64(21):5549–5556, 2016.
- [24] Yanjun Li, Kiryung Lee, and Yoram Bresler. Identifiability in bilinear inverse problems with applications to subspace or sparsity-constrained blind gain and phase calibration. *IEEE Transactions on Information Theory*, 63(2):822–842, 2017.
- [25] Shuyang Ling and Thomas Strohmer. Self-calibration and bilinear inverse problems via linear least squares. *SIAM Journal on Imaging Sciences*, 11(1):252–292, 2018.
- [26] Yvon Maday and Anthony T Patera. Spectral element methods for the incompressible navier-stokes equations. In *IN: State-of-the-art surveys on computational mechanics (A90-47176 21-64)*. New York, American Society of Mechanical Engineers, 1989, p. 71-143. Research supported by DARPA., pages 71–143, 1989.
- [27] Morten Mørup, Lars Kai Hansen, and Sidse M Arnfred. Algorithms for sparse nonnegative tucker decompositions. *Neural computation*, 20(8):2112–2131, 2008.
- [28] Yu Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- [29] Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*, volume 13. Siam, 1994.
- [30] Allan Pinkus. *N-widths in Approximation Theory*, volume 7. Springer Science & Business Media, 1985.
- [31] Gianluigi Rozza, Dinh Bao Phuong Huynh, and Anthony T Patera. Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations. *Archives of Computational Methods in Engineering*, 15(3):1, 2007.
- [32] Ledyard R Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311, 1966.
- [33] Huahua Wang, Arindam Banerjee, and Daniel Boley. Common component analysis for multiple covariance matrices. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 956–964. ACM, 2011.
- [34] Jieping Ye. Generalized low rank approximations of matrices. *Machine Learning*, 61(1-3):167–191, 2005.