



Dictionary-Based Low-Rank Approximations and the Mixed Sparse Coding Problem

Jeremy E. Cohen*

Univ Lyon, INSA-Lyon, UCBL, UJM-Saint Etienne, CNRS, Insem, CREATIS UMR 5220, Villeurbanne, France

OPEN ACCESS

Edited by:

André Uschmajew,
Max Planck Institute for Mathematics
in the Sciences, Germany

Reviewed by:

Yunlong Feng,
University at Albany, United States
Jiajia Li,
College of William & Mary,
United States

*Correspondence:

Jeremy E. Cohen
jeremy.cohen@cnsr.fr

Specialty section:

This article was submitted to
Mathematics of Computation and
Data Science,
a section of the journal
Frontiers in Applied Mathematics and
Statistics

Received: 25 October 2021

Accepted: 19 January 2022

Published: 25 February 2022

Citation:

Cohen JE (2022) Dictionary-Based
Low-Rank Approximations and the
Mixed Sparse Coding Problem.
Front. Appl. Math. Stat. 8:801650.
doi: 10.3389/fams.2022.801650

Constrained tensor and matrix factorization models allow to extract interpretable patterns from multiway data. Therefore crafting efficient algorithms for constrained low-rank approximations is nowadays an important research topic. This work deals with columns of factor matrices of a low-rank approximation being sparse in a known and possibly overcomplete basis, a model coined as Dictionary-based Low-Rank Approximation (DLRA). While earlier contributions focused on finding factor columns inside a dictionary of candidate columns, i.e., one-sparse approximations, this work is the first to tackle DLRA with sparsity larger than one. I propose to focus on the sparse-coding subproblem coined Mixed Sparse-Coding (MSC) that emerges when solving DLRA with an alternating optimization strategy. Several algorithms based on sparse-coding heuristics (greedy methods, convex relaxations) are provided to solve MSC. The performance of these heuristics is evaluated on simulated data. Then, I show how to adapt an efficient MSC solver based on the LASSO to compute Dictionary-based Matrix Factorization and Canonical Polyadic Decomposition in the context of hyperspectral image processing and chemometrics. These experiments suggest that DLRA extends the modeling capabilities of low-rank approximations, helps reducing estimation variance and enhances the identifiability and interpretability of estimated factors.

Keywords: dictionary, tensors, sparse, low-rank, non-convex optimization, LASSO

1. INTRODUCTION

Low-Rank Approximations (LRA) are well-known dimensionality reduction techniques that allow to represent tensors or matrices as sums of a few separable terms. One of the main reasons why these methods are used extensively for pattern recognition is their ability to provide part-based representations of the input data. This is particularly true for Nonnegative Matrix Factorization (NMF) or Canonical Polyadic Decomposition (CPD), see section 2 for a quick introduction and the following surveys and book for more details [1–3]. In order to fix notations while retaining generality, let us make use of the following informal mathematical formulation of LRA.

DEFINITION 1 (Low-Rank Approximation). Given a data matrix $Y \in \mathbb{R}^{n \times m}$ and a small nonzero rank $r \in \mathbb{N}$, compute so-called factor matrices $A \in \mathbb{R}^{n \times r}$ and $B \in \mathbb{R}^{m \times r}$ in

$$\arg \min_{A \in \Omega_A, B \in \Omega_B} \|Y - AB^T\|_F^2. \quad (1)$$

The set Ω_A is an additional constraint set for A , such as nonnegativity elementwise. The set Ω_B is the structure required on B to obtain a specific LRA model; for instance in the unconstrained CPD of an order three tensor; Ω_B is the set of matrices written as the Khatri-Rao product of two factor matrices. A precise qualification of “small” rank depends on the sets Ω and is omitted for simplicity.

Among LRA models, identifiability¹ properties vary significantly. While CPD is usually considered to have mild identifiability conditions [4], NMF on the other hand is often not essentially unique [3]². Generally speaking, additional regularizations are often imposed in both matrix and tensor LRA models to help with the interpretability of the part-based representation. These regularizations may take the form of constraints on the parameters or penalizations, like the sparsity-inducing ℓ_1 norm [5, 6]. They can also take the form of a parameterization of the factors [7, 8] if such a parameterization is known.

This work focuses on imposing an implicit parameterization on A . More precisely, each column of factor A is assumed to be well represented in a known dictionary $D \in \mathbb{R}^{n \times d}$ using only a small number $k < n$ of coefficients. In other words, in this work the set Ω_A is the union of subspaces of dimension k spanned by columns of a given dictionary D , and there exist a columnwise k -sparse code matrix $X \in \mathbb{R}^{d \times r}$ such that $A = DX$. A low-rank approximation model which such a constraint on at least one mode is called Dictionary-based LRA (DLRA) thereafter.

DEFINITION 2 (Dictionary-based Low-Rank Approximation). Given a data matrix $Y \in \mathbb{R}^{n \times m}$, a small nonzero rank $r \in \mathbb{N}$, a sparsity level $k < n$ and a dictionary $D \in \mathbb{R}^{n \times d}$, compute columnwise k -sparse code matrix $X \in \mathbb{R}^{d \times r}$ and factor matrix $B \in \mathbb{R}^{m \times r}$ in

$$\arg \min_{X \in \Omega_X, \forall i \leq r, \|X_i\|_0 \leq k, B \in \Omega_B} \|Y - DXB^T\|_F^2, \quad (2)$$

where X_i is the i -th column of X . The set Ω_X is an additional constraint set for X , such as nonnegativity elementwise. The set Ω_B is the structure required on B to obtain a specific LRA model. Additionally, assume that B is full column-rank.

Abbreviations: LRA, low-rank approximation; CPD, canonical polyadic decomposition; DLRA, dictionary-based LRA; DCPD, dictionary-based CPD; DNMF, dictionary-based Nonnegative Matrix Factorization; NMF, nonnegative matrix factorization; DMF, dictionary-based matrix factorization; MSC, mixed sparse coding; AO, alternating optimization; HT, hard thresholding; IHT, iterative hard thresholding; OMP, orthogonal matching pursuit; FISTA, fast iterative soft thresholding algorithm; HOMP, hierarchical OMP; SNR, signal to noise ratio; SAM, spectral angular mapper.

¹Informally, the parameters of a model are identifiable if they can be uniquely recovered from the data.

²Essential uniqueness means uniqueness up to trivial scaling ambiguities and rank-one terms permutation.

1.1. Motivations

To better advocate for the usefulness of DLRA, below two particular DLRA models are introduced:

- Dictionary-based Nonnegative Matrix Factorization (DNMF) may be formulated as

$$\arg \min_{X \in \mathbb{R}_+^{d \times r}, \forall i \leq r, \|X_i\|_0 \leq k, B \in \mathbb{R}_+^{m \times r}} \|Y - DXB^T\|_F^2. \quad (3)$$

To ensure that $A = DX$ is nonnegative, dictionary D is supposed to be entry-wise nonnegative so that the constraint $X \geq 0$ is sufficient. NMF is known to have strict identifiability conditions [3, 9], and in general one should expect that NMF has infinitely many solutions. The dictionary constraint can enhance identifiability by restraining the set of solutions. For instance setting $D = Y$ and $k = 1$ yields the Separable NMF model, which is solvable in polynomial time and which factors are generically identifiable even in the presence of noise [10]. Moreover, DNMF is also a more flexible model than NMF. In section 5.2, it is shown that DNMF can be used to solve a matrix completion problem with missing rows in the data that NMF cannot solve.

- Dictionary-based Canonical Polyadic Decomposition (DCPD), for an input tensor $Y \in \mathbb{R}^{n \times m_1 \times m_2}$ may be formulated as

$$\arg \min_{X \in \mathbb{R}^{d \times r}, \forall i \leq r, \|X_i\|_0 \leq k, B \in \mathbb{R}^{m_1 \times r}, C \in \mathbb{R}^{m_2 \times r}} \|Y - DX(B \odot C)^T\|_F^2. \quad (4)$$

where \odot is the Khatri-Rao product [1]. Unlike NMF, CPD factor are identifiable under mild conditions [4]. But in practice, identifiability may still be elusive, and the approximation problem is in general ill-posed [11] and algorithmically challenging [12]. Therefore, the dictionary constraint can be used to reduce estimation variance given an adequate choice of dictionary D . It was shown when $k = 1$ in Equation (4) that it enhances identifiability of the CPD and makes the optimization problem well-posed [13].

Note that these models have known interesting properties when $k = 1$, but have not been studied in the general case. A motivation for this work is therefore to expand these previous works to the general case $n > k \geq 1$, focusing on the algorithmic aspects. Section 2 provides more details on the one-sparse case.

1.2. Contributions

The first contribution of this work is to propose the new DLRA framework. The proposed DLRA allows to constrain low-rank approximation problems so that some of the factor matrices are sparse columnwise in a known basis. This includes sparse coding the patterns estimated by constrained factorization models such as NMF, or tensor decomposition models such as CPD. The impact of the dictionary constraint on the LRA parameter estimation error is studied experimentally in section 5.2 dedicated to experiments with DLRA, where the flexibility of DLRA is furthermore illustrated on real applications. I show that DLRA allows to complete entirely missing rows in incomplete low-rank matrices. I also show the advantage of finding the

best atoms algorithmically when imposing smoothness in DCPD using B-splines.

A second contribution is to design an efficient algorithm to solve Problem (2). To that end, I shall focus on (approximate) Alternating Optimization (AO), understood as Block Coordinate Descent (BCD) where each block update consists in minimizing almost exactly the cost with respect to the updated block while other parameters are fixed. There are two reasons for this choice. First, AO algorithms are known to perform very well in many LRA problems. They are state-of-the-art for NMF [3] and standard for Dictionary Learning [14, 15] and tensor factorization problems [16, 17]. Nevertheless inexact BCD methods and all-at-once methods are also competitive [18–21] and an inertial Proximal Alternating Linearization Method (iPALM) for DLRA is quickly discussed in section 5.1.2. The proposed algorithm is coined AO-DLRA.

As a third contribution, I study the subproblem in Problem (2) with respect to X . Developing a subroutine to solve it for known B allows not only for designing AO-DLRA, but also for post-processing a readily available estimation of A . In fact a significant part of this manuscript is devoted to studying the subproblem of minimizing the cost in Problem (2) with respect to X , which is labeled Mixed Sparse Coding (MSC),

$$\arg \min_{\forall i \leq r, \|X_i\|_0 \leq k} \|Y - DXB^T\|_F^2. \quad (\text{MSC})$$

In this formulation of MSC, no further constraints are imposed on X and therefore Ω_X is $\mathbb{R}^{d \times r}$. While similar to a matrix sparse coding, which is obtained by setting B to the identity matrix, it will become clear in this manuscript that MSC should not be handled directly using sparse coding solvers in general. For instance, it is shown in section 3.1.4 that while having an orthogonal dictionary D yields a polynomial time algorithm to solve MSC when $r = 1$ using Hard Thresholding, this does not extend when $r > 1$. In this work, the MSC problem, which is NP-hard as a generalization of sparse coding, is studied in order to build several reasonable heuristics that may be plugged into an AO algorithm for DLRA.

1.3. Structure

The article is divided in three remaining sections. The first section provides the necessary background for this manuscript. The second one is devoted to studying MSC and heuristics to solve it. The third section shows how to compute DLRA using the heuristics developed in the first part. In section 3.1, we study formally the MSC problem and its relationship with sparse coding. In section 3.2, we study a simple heuristic based solely on sparse coding each column of YB^\dagger . Two ℓ_0 heuristics similar to Iterative Hard Thresholding (IHT) and Orthogonal Matching Pursuit (OMP) are also introduced, while two types of convex relaxations are studied in Section 3.3. Section 3.5 is devoted to compare the practical performance of the various algorithms proposed to solve MSC and shows that a columnwise ℓ_1 regularization coined Block LASSO is a reasonable heuristic for MSC. Section 5.1 shows how Block LASSO can be used to compute various DLRA models, while section 5.2 illustrates DLRA on synthetic and real-life source-separation problems.

All the proofs are deferred to the **Supplementary Material** attached to this article, as well as the pseudo-codes of some proposed heuristics and additional experiments.

1.4. Notations

Vectors are denoted by small letters x , matrices by capital letters X . The indexed quantity X_i refers to the i -th column of matrix X , while X_{ij} is the (i, j) -th entry in X . A subset S of columns of X is denoted X_S , while a submatrix with columns in S_i and rows in S_j is denoted X_{S_i, S_j} . The submatrix of X obtained by removing the i -th column is denoted by X_{-i} . The i -th row of matrix X is denoted by $X_{\cdot i}$ and X_I is the submatrix of X with rows in set I . The $\ell_0(x) = \|x\|_0$ map counts the number of nonzero elements in x . The product \otimes denotes the Kronecker product, a particular instance of tensor product [22]. The Khatri Rao product \odot is the columnwise Kronecker product. The support of a vector or matrix x , i.e., the location of its nonzero elements, is denoted by $\text{Supp}(x)$. If $S = \text{Supp}(x)$, the location of the zero elements is denoted \bar{S} . The list $[M, N, P]$ denotes the concatenation of columns of matrices M, N and P . A set I contains $|I|$ elements.

2. BACKGROUND

Let us review the foundations of the proposed work, matrix and tensor decompositions and sparse coding, as well as existing models closely related to the proposed DLRA.

2.1. Matrix and Tensor Decompositions

Matrix and tensor decompositions can be understood as pattern mining techniques, which extract meaningful information out of collections of input vectors in an unsupervised manner. Arguably one of the earliest form of interpretable matrix factorization is Principal Component Analysis [23, 24], which extract a few orthogonal significant patterns out of a given matrix while performing dimensionality reduction.

Other matrix factorization models exploit other constraints than orthogonality to mine interpretable patterns. Blind source separation models such as Independent Component Analysis historically exploited statistical independence [25], while NMF, which assumes all parameters are elementwise nonnegative, has received significant scrutiny over the last two decades following the seminal paper of Lee and Seung [26]. Sparse models such as Sparse Component Analysis exploit sparsity on the coefficients [27]. It can be noted that while all those factorization techniques aim at providing interpretable representations, they are typically identifiable under strict conditions not necessarily met in practice [9, 28, 29]. It should be noted that the most important underlying hypothesis in matrix factorization is the linear dependency of the input data with respect to the templates/principal components stored in matrix A following the notations of Equation (1).

In practice, to compute for instance NMF, one solves an optimization problem of the form

$$\arg \min_{A \geq 0, B \geq 0} \|Y - AB^T\|_F^2 \quad (5)$$

which is non-convex with respect to A, B jointly but convex for each block. A common family of methods therefore uses alternating optimization in the spirit of Alternating Least Squares [30]. Other loss functions can easily be used [26]. However, computing matrix factorization models is often a difficult task; in fact NMF and sparse component analysis are both NP-hard problems in general and existing polynomial time approximation algorithms should be considered heuristics [31, 32].

Tensor decompositions follow the same ideas of unsupervised pattern mining and linearity with respect to the representation basis, but extract information out of tensors rather than matrices. Tensors in this manuscript are considered simply as multiway arrays [1] as is usually done in data sciences. Tensors have become an important data structure as they appear naturally in a variety of applications such as chemometrics [33], neurosciences [34], remote sensing [35], or deep learning [36].

At least two families of tensor decomposition models can be considered, with quite different identifiability properties and applications. A first family contains interpretable models such as the Canonical Polyadic Decomposition, often called PARAFAC [37], or closely related models such as PARAFAC2 [38]. Contrarily to constrained matrix decomposition models, the addition of at least a dimension compared to matrix factorization fixes the rotation ambiguity inherent to matrix factorization models, and therefore the CPD model is identifiable under mild conditions [4, 39, 40]. Nevertheless, additional constraints are commonly imposed on the parameters of these models to refine the interpretability of the parameters, reduce estimation errors or improve the properties of the underlying optimization problem [41, 42]. A second family is composed of tensor formats, in particular the Tucker decomposition [43] and a wide class of tensor networks such as tensor trains [44]. These models are not used in general for solving inverse problems but rather for compression or dimensionality reduction. Nevertheless, they turn into interesting pattern mining tools given adequate constraints such as nonnegativity [45, 46].

2.2. Sparse Coding

Sparse Coding (SC) and other sparse approximation problems typically try to describe an input vector as a sparse linear combination of well-chosen basis vectors. A typical formulation for sparse coding an input vector $y \in \mathbb{R}^n$ in a code book or dictionary $D \in \mathbb{R}^{n \times d}$ is the following non-convex optimization problem

$$\arg, \min_{\|x\|_0 \leq k} \|y - Dx\|_2^2, \quad (\text{SC})$$

where $k \leq n$ is the largest number of nonzero entries allowed in vector x , i.e., the size of the support of x . As long as the dictionary D is not orthogonal, this problem is difficult to solve efficiently and is in fact NP-hard [31]. The body of literature of algorithms proposed to solve SC is very large, see for instance [47] for a comprehensive overview. Overall, there exist at least three kind of heuristics to provide candidate solutions to SC:

- Greedy methods, such as OMP [48, 49] which is described in more details in section 3.2.3. These methods select indices where x is nonzero greedily until the target sparsity level or a tolerance on the reconstruction error is reached. They benefit from optimality guarantees when the dictionary columns, called atoms, are “far” from each others. In that case the dictionary is said to be incoherent [50].
- Nonconvex first order methods, such as IHT [51], that are based on proximal gradient descent [52] knowing that the projection on the set of k -sparse vectors is obtained by simply clipping the $n - k$ smallest absolute values in x . Convergence guarantees for first order constrained optimization methods is a rapidly evolving topic out of the scope of this communication, but a discussion on optimality guarantees of IHT is available in section 3.2.2.
- Convex relaxation methods, such as Least Absolute Shrinkage and Selection Operation (LASSO). The non-convex ℓ_0 constraint in Equation (SC) is replaced by a surrogate convex constraint, typically using the ℓ_1 norm. This makes the problem convex and easier to solve, at the cost of potentially changing the support of the solution [53]. Solving the LASSO can be tackled with the Fast Iterative Soft Thresholding Algorithm (FISTA) [54, 55] which is essentially an accelerated proximal gradient method with convergence guarantees.

2.3. Models Closely Related to DLRA

While this work is interested in constraining the factor matrix A in Equation (1) so that its columns are sparse in a known dictionary (in other words, sparse coding the columns of A), a few previous works have been concerned with encoding each such column with only one atom. Most related to this work is the so-called Dictionary-based CPD [13], which I shall rename as one-sparse dictionary-based LRA. Within this framework, a model such as NMF becomes

$$\arg, \min_{\mathcal{K} \in \mathcal{P}_r([1, d]), B \geq 0} \|Y - D_{\mathcal{K}} B^T\|_F^2 \quad (6)$$

where $\mathcal{P}_r([1, d])$ is the set of all parts of $[1, d]$ with r elements, therefore \mathcal{K} is a set of r indices from 1 to d . This problem is a particular case of the proposed DLRA framework because $D_{\mathcal{K}}$ can be written as DX where the columns of X are one-sparse. It was shown that one-sparse DLRA makes low-rank matrix factorization identifiable under mild coherence conditions on the dictionary, and makes the computation of CPD a well-posed problem. Intuitively, the dictionary constraint may be used to enforce a set of known templates to be used as patterns in the pattern mining procedure, and therefore one-sparse DLRA may be seen as a glorified pattern matching technique.

Other works in the sparse coding literature are related to DLRA, in particular the so-called multiple measurement vectors or collaborative sparse coding [56–58], which extends sparse coding when several inputs collected in a matrix Y are coded with the same dictionary using the same support. Effectively this means solving a problem of the form

$$\arg, \min_{\|\sum_{i=1}^r Z_i^T\|_0 \leq kr} \|Y - DZ^T\|_F^2 \quad (7)$$

where the square is meant elementwise. DLRA can also be seen as a collaborative sparse coding by noticing that $Z := XB^T$ is at least kr row-sparse in Equation (2). However the low-rank hypothesis is lost, as well as the affectation of at most k atoms per column of matrix A .

3. MIXED SPARSE CODING HEURISTICS

3.1. Properties of Mixed Sparse Coding

In order to design efficient heuristics to solve MSC, let us first cover a few properties of MSC such as uniqueness of MSC solutions, relate MSC with other sparse coding problems, and check whether special simpler cases exists such as when support of the solution is known or when the dictionary is orthogonal. The following section covers this material.

3.1.1. Equivalent Formulations and Relation to Sparse Coding

Let us start our study by linking MSC to other sparse coding problems. For reference sake, in this work the standard matrix sparse coding problem is formulated as

$$\arg, \min_{\forall i \leq r, \|X_i\|_0 \leq k} \|Y - DX\|_F^2. \quad (8)$$

which is equivalent to solving r vector sparse coding problems as defined in Equation (SC).

Because the Frobenius norm is invariant to a reordering of the entries, it can be seen easily that Problem (MSC) is equivalent after vectorization³ to a structured vector sparse coding Problem (SC) with block-sparsity constraints:

$$\arg, \min_{\forall i \leq r, \|\text{vec}(X_i)\|_0 \leq k} \|\text{vec}(Y) - (D \otimes B) [\text{vec}(X_1), \dots, \text{vec}(X_r)]\|_2^2. \quad (9)$$

It appears that MSC is therefore a structured sparse coding problem since the dictionary is a Kronecker product of two matrices. Moreover, the sparsity constraint applies on blocks of coordinates in the vectorized input $\text{vec}(X)$, as studied in [60]. To the best of the author's knowledge, block sparse coding with Kronecker structured dictionary have not been specifically studied. In particular, the conjunction of Kronecker structured dictionary [61] and structured sparsity leads to specific heuristics detailed in the rest of this work. Nevertheless, in section 3.2, it is shown that MSC reduces to columnwise sparse coding under a small noise regime.

On a different note, consider the following problem

$$\min_{\|Y - DXB^T\|_F^2 \leq \epsilon_1} \max_i \|X_i\|_0. \quad (10)$$

for some positive constant ϵ_1 . Just like how Quadratically Constrained Basis Pursuit and LASSO are equivalent [47], one can show using similar arguments that Problems (MSC) and (10) have the same solution given a mild uniqueness assumption.

³Vectorization in this manuscript is row-first [59].

PROPOSITION 1. *Suppose that Problem (10) has a unique solution X^* with $\epsilon_1 \geq 0$. Then there exist a particular instance of Problem (MSC) such that X^* is the unique solution of (MSC). Conversely, a unique solution to Problem (MSC) is the unique solution to a particular instance of Problem (10).*

Formulation (10) of MSC could also be expressed with matrix induced norms. Indeed, defining $\ell_{0,0}(X) := \sup_{\|z\|_0=1} \|Xz\|_0 = \max_i \|X_i\|_0$ as an extension of matrix induced norms⁴ $\ell_{p,q}(X) := \sup_{\|z\|_p=1} \|Xz\|_q$ [62], Problem (10) is equivalent to

$$\min_{\|Y - DXB^T\|_F^2 \leq \epsilon_1} \ell_{0,0}(X), \quad (11)$$

a clear structured extension of quadratically constrained sparse coding [47]. Some authors preferably work with Problem (10) rather than (MSC) because in specific applications, choosing an error tolerance ϵ_1 is more natural than choosing the sparsity level k . While heuristics introduced further are geared toward solving Problem (MSC), they can be adapted to solve Problem (10).

The quadratically constrained formulation also sheds light on the fact that, if there exist some solution X to MSC such that the residual $\|Y - DXB^T\|_F^2$ is zero, then MSC is equivalent to matrix sparse coding Problem (8). Indeed, since it is assumed that B is full column rank, any such solution yields $YB(B^TB)^{-1} = DX$. Consequently, noiseless formulations of MSC will not be considered any further.

3.1.2. Generic Uniqueness of Solutions to MSC

In sparse coding, sparsity is introduced as a regularizer to enforce uniqueness of the regression solution. It is therefore natural to wonder if this property also holds for MSC. It is not difficult to observe that indeed generically the solution to MSC will be unique, under the usual spark condition on the dictionary D [63].

PROPOSITION 2. *Define $\text{spark}(D)$ as the smallest number of columns of D that are linearly dependent. Suppose that $\text{spark}(D) > 2k$ and suppose B is a full column rank matrix. Then the set of Y such that Problem (MSC) has strictly more than one solution has zero Lebesgue measure.*

This result basically states that in practice, most MSC instances will have unique solutions as long as the dictionary is not too coherent.

3.1.3. Solving MSC Exactly When the Support Is Known

Solving the NP-hard MSC problem exactly is difficult because the naive, brute force algorithm implies testing all combinations of supports for all columns, which means computing $\binom{k}{q}^r$ least squares and is significantly slower than brute force for Problem (8). However, in the spirit of sparse coding which boils down to finding the optimal support for a sparse solution, MSC also reduces to a least squares problem when the locations of the zeros in a solution X are fixed and known. Below is detailed

⁴Mind the possible confusion with the $\mathcal{L}_{p,q}$ convention $\ell_{0,\infty}$ commonly encountered.

how to process this least squares problem to avoid forming the Kronecker matrix $D \otimes B$ explicitly.

From the observation in Equation (9) that the vectorized problem is really a structured sparse coding problem, one can check that for a given support $S = \text{Supp}(\text{vec}(X))$, solving MSC amounts to solving a linear system

$$\arg, \min_{z \in \mathbb{R}^{kr}} \|y - (D \otimes B)_S z\|_2^2. \tag{12}$$

where $y := \text{vec}(Y)$.

It is not actually required to form the full Kronecker product $D \otimes B$ and then select a subset S of its columns, nor is it necessary to vectorize Y . More precisely, one needs to compute two quantities: $((D \otimes B)_S)^T (D \otimes B)_S$ and $((D \otimes B)_S)^T y$ to feed a linear system solver, and these products can be computed efficiently. Denote S_i the support of column X_i . Formally, one may notice that $(D \otimes B)_S$ is a block matrix $[D_{S_1} \otimes b_1, \dots, D_{S_r} \otimes b_r]$. Therefore, using the identity $(D \otimes b)^T (D \otimes b) = b^T b D^T D$, one may compute the cross product by first precomputing $U = D^T D^S$ and $V = B^T B$, then computing each block $(D_{S_i} \otimes b_i)^T (D_{S_j} \otimes b_j)$ as $v_{ij} U_{S_i, S_j}$.

For the data and mixing matrix product, notice that $(D_{S_i} \otimes b_i)^T y = D_{S_i} Y b_i^T$. Therefore, the model-data product can be obtained by first precomputing $N = Y B^T$, then computing each block $((D \otimes B)_S)^T y$ as $D_{S_i} N_i$.

Remark on regularization: It may happen that for a fixed support, the linear system (12) is ill-posed. This may be caused by a highly coherent dictionary D or the choice of a large rank r . To avoid this issue, when the system is ill-conditioned in later experiments, a small ridge penalization may be added.

3.1.4. MSC With Orthogonal Dictionary Is Easy for Rank One LRA

An important question is whether the problem generally becomes easier if the dictionary is left orthogonal. This is the case for Problem (8), where orthogonal D allows to solve the problem using only Hard Thresholding (HT) on $D^T Y$ columnwise. Below, it is shown that a similar HT procedure can be used when $r = 1$, but not in general when $r > 1$.

Supposing D is left orthogonal, the MSC problem becomes

$$\arg, \min_{\|X_i\|_0 \leq k \forall i \in [1, r]} \|D^T Y - X B^T\|_F^2. \tag{13}$$

When $r = 1$, matrix B^T is simply a row vector, which is a right orthogonal matrix after ℓ_2 normalization such that solving Problem (13) amounts to minimizing $\|\frac{1}{\|b\|_2} D^T Y b - x\|_2^2$. Then the solution is obtained using the thresholding operator $\text{HT}_k(\frac{1}{\|b\|_2} D^T Y b)$ which selects the k largest entries of its input.

Sadly when $r > 1$, matrix B^T is not right orthogonal in general and neither is $D \otimes B$. Therefore this thresholding strategy does not yield a MSC solution in general despite D being left orthogonal.

⁵If this is not possible because D is too large, one may instead compute the blocks U_{S_i, S_j} without pre-computing $D^T D$.

3.2. Non-convex Heuristics to Solve Mixed Sparse Coding

We now focus on heuristics to find candidate solutions to MSC. Similarly to sparse coding, MSC is an NP-hard problem for which obtaining the global solution typically requires costly algorithms [64, 65]. Therefore, in the following section, several heuristics are proposed that aim at finding good sparse approximations in reasonable time.

- A classical sparse coding algorithm [here OMP [49]] applied columnwise on $Y B (B^T B)^{-1}$. Indeed in a small noise regime, columnwise sparse coding on $Y B (B^T B)^{-1}$ is proven to be equivalent to solving MSC.
- A Block-coordinate descent algorithm that features OMP as a subroutine.
- A proximal gradient algorithm similar to Iterative Hard Thresholding.
- Two convex relaxations analog to LASSO [66], which are solved by accelerated proximal gradient in the spirit of FISTA [55]. Maximum regularization levels are computed, and a few properties concerning the sparsity and the uniqueness of the solutions are provided.

The performance of all these heuristics as MSC solvers is studied in section 3.5.

3.2.1. A Provable Reduction to Columnwise Sparse Coding for Small Noise Regimes

At first glance, one may think that projecting Y onto the row-space of B^T maps solutions of Problem (MSC) to solutions of Problem (8). Indeed, writing $Y = Y_B + Y_{-B}$ with $Y_B := Y B (B^T B)^{-1} B^T$ the orthogonal projection of Y on the row-space of B^T , it holds that Problem (MSC) has the same minimizers than

$$\min_{\|X_i\|_0 \leq k \forall i \in [1, r]} \|Y_B - D X B^T\|_F^2. \tag{14}$$

This problem is however not in general equivalent to matrix sparse coding because B^T distorts the error distribution.

Even though they are not equivalent, one could hope to solve MSC, in particular settings, by using a candidate solution to the matrix sparse coding Problem (8) with $Y B (B^T B)^{-1}$ as input. The particular case of $k = 1$ is striking: matrix sparse coding is solved in closed form whereas MSC has no general solution as far as I know. This kind of heuristic replacement of Problem (MSC) by Problem (8) has been used heuristically in [13] when $k = 1$, and can be interpreted as “First find Z that minimizes $\|Y - Z B^T\|_F^2$, then find a k -sparse matrix X that minimizes $\|Z - D X\|_F^2$ ”.

We show below that in a small perturbation regime, for a dictionary D and a mixing matrix B satisfying classical assumptions in compressive sensing, the matrix sparse coding solution has the same support than the MSC solution.

LEMMA 1. *Let X, X' columnwise k -sparse matrices and $\epsilon > 0, \delta > 0$ such that $\|Y - D X B\|_F^2 \leq \epsilon$ and $\|Y B (B^T B)^{-1} - D X'\|_F^2 \leq \delta$. Further suppose that $\text{spark}(D) > 2k$ and that B has*

full column-rank. Then

$$\|X - X'\|_F \leq \frac{1}{\sigma_{\min}^{(2k)}(D)} \sqrt{\delta + \frac{\epsilon}{\sigma_{\min}^2(B)}} \quad (15)$$

where $\sigma_{\min}(B)$ is the smallest nonzero singular value of B and $\sigma_{\min}^{(2k)}$ is the smallest nonzero singular value of all submatrices of D constructed with $2k$ columns.

Remarkably, if $k = 1$, then Problem (8) has a closed-form solution, and one may replace δ by the best residual to remove this unknown quantity in Lemma 1. On a different note, variants of this result can be derived under similar assumptions, for instance if D satisfies a $2k$ restricted isometry property.

We can now derive a support recovery equivalence between MSC and SC.

PROPOSITION 3. *Under the hypotheses of Lemma 1, supposing that X and X' are exactly columnwise k -sparse, if*

$$\min_{j \leq r} \sqrt{\min_{i \in \text{Supp}(X_j)} X_{ij}^2 + \min_{i' \in \text{Supp}(X'_{i'})} X'_{i'j}{}^2} > \frac{1}{\sigma_{\min}^{(2k)}(D)} \sqrt{\delta + \frac{\epsilon}{\sigma_{\min}^2(B)}}, \quad (16)$$

then $\text{Supp}(X) = \text{Supp}(X')$.

In practice, this bound can be used to grossly check whether support estimation in MSC may reduce to support estimation in matrix sparse coding or not. To do so, one may first tentatively solve Problem (8) with $YB(B^T B)^{-1}$ as input and obtain a candidate value X' as well as some residual δ . Furthermore, while this is costly, the values of $\sigma_{\min}(B)$ and $\sigma_{\min}^{(2k)}(D)$ may be computed. Then removing the term X_{ij} from Equation (16) yields a bound of the noise level ϵ :

$$\sigma_{\min}(B) \left(-\delta + (\sigma_{\min}^{(2k)}(D))^2 \min_{j, i \in \text{Supp}(X')} X_{ij}^2 \right) > \epsilon \quad (17)$$

under which the reduction is well-grounded.

These observations lead to the design of a first heuristic, which applies OMP to each column of $YB(B^T B)^{-1}$. The obtained support is then used to compute a solution X to MSC as described in section 3.1.3. This heuristic will be denoted TrickOMP in the following.

3.2.2. A First Order Strategy: Iterative Hard Thresholding

Maybe the most simple way to solve MSC is by proximal gradient descent. In sparse coding, this type of algorithm is often referred to as Iterative Hard Thresholding (IHT), and this is how I will denote this algorithm as well for MSC.

Computing the gradient of the differentiable convex term $f(X) = \|Y - DX\|_F^2$ with respect to X is easy and yields

$$\frac{\partial f}{\partial X} = -D^T YB + D^T DXBB^T. \quad (18)$$

Note that depending on the structure Ω_B , products $D^T YB$ and $B^T B$ may be computed efficiently.

Then it is required to compute a projection on the set of columnwise k -sparse matrices. This can easily be done columnwise by hard thresholding. However, for the algorithm to be well-defined and deterministic, we need to suppose that projecting on the set of columnwise k -sparse matrices is a closed-form single-valued operation. This holds if, when several solutions exist, i.e., when several entries are the k -th largest, one picks the right number of these entries in for instance the lexicographic order. The proposed IHT algorithm leverages the classic IHT algorithm and is summarized in the **Supplementary Material**. However, contrarily to the usual IHT, the proposed implementation makes use of the inertial acceleration proposed in FISTA [55]. Compared to using IHT for solving sparse coding, as hinted in Equation (9), here IHT is moreover applied to a structured sparse coding problem. This does not significantly modify the practical implementation of the algorithm.

IHT benefits from both convergence results as a (accelerated) proximal gradient algorithm with semi-algebraic regularization [51, 67] and support recovery guaranties for sparse coding [47, 68]. While the convergence results directly apply to MSC, extending support recovery guaranties to MSC is an interesting research avenue.

3.2.3. Hierarchical OMP

Before describing the proposed hierarchical algorithm, it might be helpful to understand how greedy heuristics for sparse coding, such as the OMP algorithm, are derived. Greedy heuristics select the best atom to reconstruct the data, then remove its contribution and repeats this process with the residuals. The key to understanding these techniques is that selecting only one atom is a problem with a closed-form solution. Indeed, for any $y \in \mathbb{R}^n$,

$$\arg \min_{\|x\|_0 \leq 1} \|y - Dx\|_2^2 = \arg \min_{z \in \mathbb{R}} \arg \min_{j \in [1, d]} \|y - zD_j\|_2^2 \quad (19)$$

and after some algebra, for a dictionary normalized columnwise,

$$\|y - zD_j\|_2^2 = \|y\|_2^2 + z^2 - 2zD_j^T y = \text{cst}(j) - 2zD_j^T y \quad (20)$$

which minimum with respect to j does not depend on the value z of the nonzero coefficient z in the vector x (except for the sign of z). Therefore the support of x is $\arg \max_j |D_j^T y|$.

Reasoning in the same manner for MSC does not yield a similar simple solution for finding the support. Indeed, even setting $k = 1$,

$$\arg \min_{\|X_i\|_0 \leq 1} \|Y - DXB^T\|_F^2 = \arg \min_{\forall i \leq r, z_i \in \mathbb{R}} \arg \min_{\forall i \leq r, j_i \in [1, d]} \|Y - \sum_{i=1}^r z_i D_{j_i} B_i^T\|_F^2 \quad (21)$$

and the interior minimization problem of the right-hand side is still difficult, referred to as dictionary-based low-rank matrix factorization in [13]. Indeed, the cost can be rewritten as $\|Y - D(\cdot, \mathcal{K}) \text{Diag}(z) B^T\|_F^2$ with \mathcal{K} the list of selected atoms in D . As discussed in section 3.2.1, the solutions to this problem in a noisy setting are in general not obtained by minimizing instead $\|YB(B^T B)^{-1} - D(\cdot, \mathcal{K}) \text{Diag}(z)\|_F^2$ with respect to \mathcal{K} which would be solved in closed form.

Therefore, even with $k = 1$, a solution to MSC is not straightforward. While a greedy selection heuristic may not be straightforward to design, one may notice that if the rank had been set to one, i.e., if $r = 1$, then for any k one actually ends up with the usual sparse coding problem.

Indeed, given a matrix $V \in \mathbb{R}^{n \times m}$, the rank-one case means we try to solve the problem

$$\arg \min_{x \in \mathbb{R}^d, \|x\|_0 \leq k} \|V - Dx\|_F^2 \quad (22)$$

which is equivalent to

$$\arg \min_{x \in \mathbb{R}^d, \|x\|_0 \leq k} \left\| \frac{1}{\|b\|_2} Vb - Dx \right\|_2^2. \quad (23)$$

This is nothing more than Problem (SC). Consequently, we are now ready to design a hierarchical, i.e. block-coordinate, greedy algorithm. The algorithm updates one column of X at a time, fixing all the others. Then, finding the optimal solution for that single column is exactly a sparse-coding problem.

This leads to an adaptation of OMP for MSC that I call Hierarchical OMP (HOMP), where OMP is used to solve each sparse coding subproblem, see **Algorithm 1**. The routine $\text{OMP}(x, D, k)$ applies Orthogonal Matching Pursuit to the input vector x with normalized dictionary D and sparsity level k and returns estimated codes and support. Note that after HOMP has stopped, it is useful to run a least square joint final update with fixed support as described in section 3.1.3 because the final HOMP estimates may not be optimal for the output support. HOMP does not easily inherit from OMP recovery conditions [50] because it employs OMP inside an alternating algorithm. On a practical side, any optimized implementation of OMP such as batch OMP [69] can be used to implement HOMP as a subroutine.

A note on restart: A restart condition is required, i.e., checking if the error increases after an inner iteration and rejecting the update in that case. Indeed OMP is simply a heuristic which, in general, is not guaranteed to find the best solution to the sparse coding subproblem. When restart occurs, simply compute the best update with respect to the previously known support and move to the next column update. If restart occurs on all modes, then the algorithm stops with a warning. Due to this restart condition, the cost always decreases after each iteration, therefore it is guaranteed that the HOMP algorithm either converges or stops with a warning.

3.3. Convex Heuristics to Solve MSC

The proposed greedy strategies TrickOMP and HOMP may not provide the best solutions to MSC or even converge to a critical point, and the practical performance of IHT is often not satisfactory (see section 3.5). Therefore, taking inspiration from existing works on sparse coding, one might as well tackle a convex problem which solutions are provably sparse. This means, first of all, finding convex relaxations to the $\ell_{0,0}$ regularizer.

In what follows, we study two convex relaxations and propose a FISTA-like algorithm for each. In both cases, the solution is

Algorithm 1 Hierarchical OMP.

Input: data Y , dictionary D , sparsity level k , initial value X .

Output: estimated codes X , support S

Precompute $D^T D$ and $D^T Y$ if memory allows.

while stopping criterion is not reached **do**

for p from 1 to r **do**

 Set $Vb = \frac{1}{\|B_p\|_2} (Y - DX_{-p} B_{-p}^T) B_p$.

 Compute $X_p, S_p = \text{OMP}(Vb, D, k)$

 If error increased, reject this update, and perform a least squares update with the previous support. End if rejection happened for each column.

end for

end while

Set X as the least squares solution following section 3.1.3 with support S .

provably sparse and there exist a regularization level such that the only solution is zero. Therefore, these regularizers force the presence of zeros in the columns of the solution.

3.3.1. A Columnwise Convex Relaxation: The Block LASSO Heuristic

A first convex relaxation of MSC is obtained by replacing each sparsity constraint $\|X_i\|_0 \leq k$ by an independent ℓ_1 regularization. This idea has already been theoretically explored in the literature, in particular in [60] where several support recovery results are established.

Practically, fixing a collection $\{\lambda_i\}_{i \leq r}$ of positive regularization parameters, the following convex relaxed problem, coined Block LASSO,

$$\arg \min_{X \in \mathbb{R}^{d \times r}} \frac{1}{2} \|Y - DXB^T\|_F^2 + \sum_{i=1}^r \lambda_i \|X_i\|_1 \quad (\text{CVX1})$$

provides candidates solutions to MSC. This is a convex problem since each term is convex. Moreover, the cost is coercive so that a solution always exists. However, it is not strictly convex in general, thus several solutions may co-exist.

Adapting the proof in [47], one can easily show that under uniqueness assumptions, solutions to Problem (CVX1) are indeed sparse.

PROPOSITION 4. *Let X^* a solution to Problem (CVX1), and suppose that X^* is unique. Then denoting S_i the support of each column X_i^* , it holds that D_{S_i} is full column-rank, and that $|S_i| \leq n$.*

This shows that the Block LASSO solutions have at most n non-zeros in each column. Therefore, the columnwise ℓ_1 regularization induces sparsity in all columns of X and solving Problem (CVX1) is relevant to perform MSC. But this does not show that the support recovered by solving Block LASSO is always the support of the solution of MSC, see [60] for such recovery results given assumptions on the regularizations parameters λ_i .

Conversely, it is of interest to know above which values of λ_i the solution X is null. Proposition 5 states that such a columnwise

maximum regularization can be defined. This can be used to set the λ_i regularization parameters individually given a percentage of $\lambda_{i,\max}$ and provide a better intuition for choosing each λ_i .

PROPOSITION 5. *A solution X^* of Problem (CVX1) satisfies $X^* = 0$ if and only if for all $i \leq r$, $\lambda_i \geq \lambda_{i,\max}$ where $\lambda_{i,\max} = \|D^T Y B_i\|_\infty$ (element-wise absolute value maximum). Moreover, if $X_i^* = 0$ is a column of a solution, then $\lambda_i \geq \lambda_{i,\max}$.*

Relation with other models: Vectorization easily transform the matrix Problem (CVX1) into a vector problem reminiscent of the LASSO. In fact, if the regularization parameters $\lambda_i = \lambda$ are set equal, then Problem (CVX1) is nothing else than the LASSO. A related but different model is SLOPE [70]. Indeed, in SLOPE, it is possible to have a particular λ_{j_i} for each **sorted** entry $X_{\sigma_i(j)_i}$, but in Problem (CVX1) the regularization parameters are fixed by blocks and the relative order of elements among the blocks changes in the admissible space. Thus the relaxed Problem (CVX1) cannot be expressed as a particular SLOPE problem.

Solving block LASSO with FISTA: A workhorse algorithm for solving the LASSO is FISTA [55]. Moreover, the regularization term is separable in X_i , and the proximal operator for each separable term is well-known to be the soft-thresholding operator

$$S_{\lambda_i}(x) = [|x| - \lambda]^+ \text{sign}(x), \quad S_{[\lambda_1, \dots, \lambda_r]}(X) = [S_{\lambda_1}(X_1), \dots, S_{\lambda_r}(X_r)] \quad (24)$$

understood element-wise. Therefore, the FISTA algorithm can be directly leveraged to solve Problem (CVX1), see **Algorithm 2** coined Block-FISTA hereafter. Convergence of the cost iterates of this proposed extrapolated proximal gradient method is ensured as soon as the gradient step is smaller than the inverse of the Lipschitz constant of the quadratic term, which is given by $\sigma(D)^2 \sigma(B)^2$ with $\sigma(M)$ the largest singular value of M . The resulting FISTA algorithm is denoted as Block FISTA. Note that after Block FISTA returns a candidate solution, this solution's support is extracted, truncated to be of size k columnwise using hard thresholding, and an unbiased MSC solution is computed with that fixed support.

3.3.2. Mixed ℓ_1 Norm for Tightest Convex Relaxation

The columnwise convex relaxation introduced in section 3.3.1 has the disadvantage of introducing a potentially large number of regularization parameters that must be controlled individually to obtain a target sparsity level columnwise. Moreover, this relaxation is not the tightest convex relaxation of the $\ell_{0,0}$ regularizer on $[-1, 1]$.

It turns out that using the tightest convex relaxation of the $\ell_{0,0}$ regularizer does solve the proliferation of regularization parameters problem, and in fact this yields a uniform regularization on the columns of X . This comes however at the cost of loosing some sparsity guaranties as detailed below.

PROPOSITION 6. *The tightest convex relaxation of $\ell_{0,0}$ on $[-1, 1]^{d \times r}$ is $\ell_{1,1} : X \mapsto \max_i \|X_i\|_1 = : \|X\|_{1,1}$.*

According to Proposition 6, Problem (MSC) may be relaxed into the following convex optimization problem coined Mixed

Algorithm 2 FISTA for Block LASSO (Block-FISTA)

Input: data Y , dictionary D , mixing matrix B , regularization ratio $\alpha \in [0, 1]^r$, sparsity level k , initial value X .

Output: estimated codes X , support S .

Precompute $D^T D$, $D^T Y B$ and $B^T B$ if memory allows.

Compute $\lambda_{i,\max}$ as in Proposition 5, and set $\lambda_i = \alpha_i \lambda_{i,\max}$, $\lambda = [\lambda_1, \dots, \lambda_r]$

Compute stepsize $\eta = \frac{1}{\sigma(D)^2 \sigma(B)^2}$

Initialize $Z = X$, $\beta = 1$.

while stopping criterion is not reached **do**

$X_{\text{old}} = X$

$X = S_{\eta\lambda} (Z - \eta(D^T D Z B^T B - D^T Y B))$

$\beta_{\text{old}} = \beta$

$\beta = \frac{1}{2}(1 + \sqrt{1 + 4\beta^2})$

$Z = X + \frac{\beta_{\text{old}} - 1}{\beta} (X - X_{\text{old}})$

end while

Estimate the support $S = S(X)$

Set X as the least squares solution following section 3.1.3 with support S .

LASSO hereafter:

$$\arg \min_{X \in \mathbb{R}^{d \times r}} \frac{1}{2} \|Y - D X B^T\|_F^2 + \lambda \|X\|_{1,1}. \quad (25)$$

To again leverage FISTA to solve Problem (25) and produce a support for a solution of MSC requires to compute the proximal operator of the regularization term. For the $\ell_{1,1}$ norm, the proximal operator has been shown to be computable exactly with little cost using a bisection search [71–73]. In this work I used the low-level implementation of [72]⁶. The resulting Mixed-FISTA algorithm is very similar to **Algorithm 2** but using the $\ell_{1,1}$ proximal operator instead of soft-thresholding, and its pseudocode is therefore differed to the **Supplementary Material**.

Properties of the Mixed LASSO: Differently from the Block LASSO problem, the Mixed LASSO may not have sparse solutions if the regularization is not set high enough.

PROPOSITION 7. *Suppose there exist a unique solution X^* to the Mixed LASSO problem. Let \mathcal{I} the set of indices such that for all i in \mathcal{I} , $\|X_i^*\|_1 = \|X^*\|_{1,1}$. Denote S the support of X^* . Then there exist at least one i in \mathcal{I} such that D_{S_i} is full column rank, and $\|X_i^*\|_0 \leq n$. Moreover, if D is overcomplete, $\mathcal{I} = \{1, \dots, r\}$.*

This result is actually quite unsatisfactory. The uniqueness condition, which is generally satisfied for the LASSO, seems a much stricter restriction in the Mixed LASSO problem. In particular all columns of the solution must have equal ℓ_1 norm in the overcomplete case. Moreover sparsity is only ensured for one column. However, from the simulations in the Experiment section, it seems that in general the Mixed LASSO problem does yield sparser solutions than the above theory predicts.

Maximum regularization: Intuitively, by setting the regularization parameter λ high enough, one expects the

⁶<https://github.com/bbejar/prox-l1oo>.

solution of the Mixed LASSO to be null. Below I show that this is indeed true.

PROPOSITION 8. $X^* = 0$ is the unique solution to the Mixed LASSO problem if and only if $\lambda \geq \lambda_{\max}$ where $\lambda_{\max} = \sum_{i=1}^r \|D^T Y B_i\|_{\infty}$.

Similarly to Block-FISTA, this result can be used to choose the regularization parameter as a percentage of λ_{\max} .

3.4. Nonnegative MSC

In this section I quickly describe how to adapt the Block LASSO strategy in the presence of nonnegativity constraints on X . This is a very useful special case of MSC since many LRA models use nonnegativity to enhance identifiability such as Nonnegative Matrix Factorization or Nonnegative Tucker Decomposition.

The nonnegative MSC problem can be written as follows:

$$\arg, \min_{X \geq 0, \|X_i\|_0 \leq k} \|Y - DXB^T\|_F^2. \quad (26)$$

Block LASSO minimizes a regularized cost which becomes quadratic and smooth due to the nonnegativity constraints:

$$\arg, \min_{X \geq 0} \|Y - DXB^T\|_F^2 + \sum_{i=1}^n \lambda_i 1^T X_i. \quad (27)$$

The Nonnegative least squares Problem (26) is easily solved by a modified nonnegative Block-FISTA, where the proximal operator is a projection on the nonnegative orthant.

While in the general case a least-squares update with fixed support is performed at the end of Block-FISTA to remove the bias induced by the convex penalty, in the non-negative case a nonnegative least squares solver is used on the estimated support with a small ridge regularization. One drawback of this approach is that the final estimate for X might have strictly smaller sparsity level than the target k in a few columns. In the **Supplementary Material**, a quick comparison between Block-FISTA and its nonnegative variant shows the positive impact of accounting for nonnegativity for support recovery.

3.5. Comparison of the Proposed Heuristics

Numerical experiments discussed hereafter provide a first analysis of the proposed heuristics to solve MSC. A few natural questions arise upon studying these methods:

- Are some of the proposed heuristics performing well or poorly in terms of support recovery in a variety of settings?
- Are some heuristics much faster than others in practice?

We shall provide tentative answers after conducting synthetic experiments. However, because of the variety of proposed methods and the large number of experimental parameters (dimensions, noise level, conditioning of B and coherence of D , distribution of the true X , regularization levels), it is virtually impossible to test out all possible combinations and the conclusions of this section can hardly be extrapolated outside our study cases. All the codes used in the experiments below are freely

available online⁷. In particular, all the proposed algorithms are implemented in Python, and all experiments and figures can be reproduced using the distributed code.

In the **Supplementary Material**, I cover additional questions of importance: the sensitivity of convex relaxation methods to the choice of the regularization parameter, the sensitivity of all methods to the conditioning of B and the sensitivity to random and zero initializations.

3.5.1. Synthetic Experiments

In general and unless specified otherwise we set $n = 50, m = 50, d = 100, k = 5, r = 6$. The variance of additive Gaussian white noise is tuned so that the empirical SNR is exactly 20dB. To generate D , its entries are drawn independently from the Uniform distribution on $[0, 1]$ and its columns are then normalized. The uniform distribution is meant to make atoms more correlated and therefore increase the dictionary coherence. The entries of B are drawn similarly. However, the singular value decomposition of B is then computed, its original singular values discarded and replaced with linearly spaced values from 1 to $\frac{1}{\text{cond}(B)}$ where $\text{cond}(B) = 200$. The values of the true X are generated by first selecting a support randomly (uniformly), then sampling nonzero entries from standard Gaussian i.i.d. distributions. The initial X is set to zero.

In most test settings, the metric used to assess performance is based on support recovery. To measure support recovery, the number of correctly found nonzero position is divided by the total number of non-zeros to be found, yielding a 100% recovery rate if the support is perfectly estimated and 0% if no element in the support of X is correctly estimated.

The input regularization in Mixed-FISTA and Block-FISTA is always scaled from 0 to 1 by computing the maximum regularization, see Propositions 5 and 8. To choose the regularization ratio α , before running each experiment, Mixed-FISTA and Block-FISTA are ran on three instances of separately generated problems with the same parameters as the current test using a grid $[10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}]$. Then the average best α for these three tests is used as the regularization level for the whole test. This procedure is meant to mimic how a user would tune regularization for a given problem, generating a few instances by simulation and picking a vaguely adequate amount of regularization.

The stopping criterion for all methods is the same: when the relative decrease in cost $\frac{\text{err}^{it+1} - \text{err}^{it}}{\text{err}^{it}}$ reaches 10^{-6} , the algorithm stops. The absolute value allows for increasing the cost. Note that the cost includes the penalty terms for convex methods. Additionally, the maximum number of iterations is set to 1,000.

3.5.2. Test 1: Support Recovery vs. Noise Level

For the first experiment, the noise level varies in power such that the SNR is exactly on a grid $[1, 100, 10, 50, 40, 30, 20, 15, 10, 5, 2, 0]$. A total of 50 realizations of triplets (Y, D, B) are used in this experiment,

⁷<https://github.com/cohenjer/mscode> and <https://github.com/cohenjer/dlra>.

as well as in Test 2 and 3. Results are shown in **Figure 1**.

It can be observed that IHT performs poorly for all noise levels. As expected, TrickOMP performs well only at low noise levels. HOMP and the FISTA methods have degraded performance when the SNR decreases, but provide with satisfactory results overall. Block-FISTA seems to perform the best overall.

3.5.3. Test 2: Support Recovery vs. Dimensions (k,d)

This time the sparsity level k is on a grid [1, 2, 5, 10, 20] while the number of atoms d is also on a grid [20, 50, 100, 200, 400]. **Figure 2** shows the heat map results.

The TrickOMP has strikingly worse performance than the other methods. Moreover, as the number of atoms d increases or when the number of admissible supports $\binom{d}{k}$, correct atom selection becomes more difficult for all methods. Again, Block-FISTA seems to perform better overall, in particular for large d .

3.5.4. Test 3: Runtime vs. Dimensions (n,m) and (k,d)

In this last test, all algorithms are run until convergence for various sizes $n = [10, 50, 1,000]$ and $m = [10, 50, 1,000]$, or for various sparsity parameters $k = [5, 10, 30]$ and $[50, 100, 1,000]$.

Table 1 provides runtime and number of iterations averaged for $N = 10$ runs.

From **Table 1**, it can be inferred that HOMP is often much slower than the other methods. TrickOMP is always very fast since it relies on OMP which runs in exactly k iterations. Block-FISTA generally runs faster than Mixed-FISTA. Moreover, it is not much slower than faster methods such as IHT and TrickOMP, in particular for larger sparsity values.

4. DISCUSSION

In all the experiments conducted above and in the **Supplementary Material**, the Block-FISTA algorithm provides the best trade-off between support recovery and computation

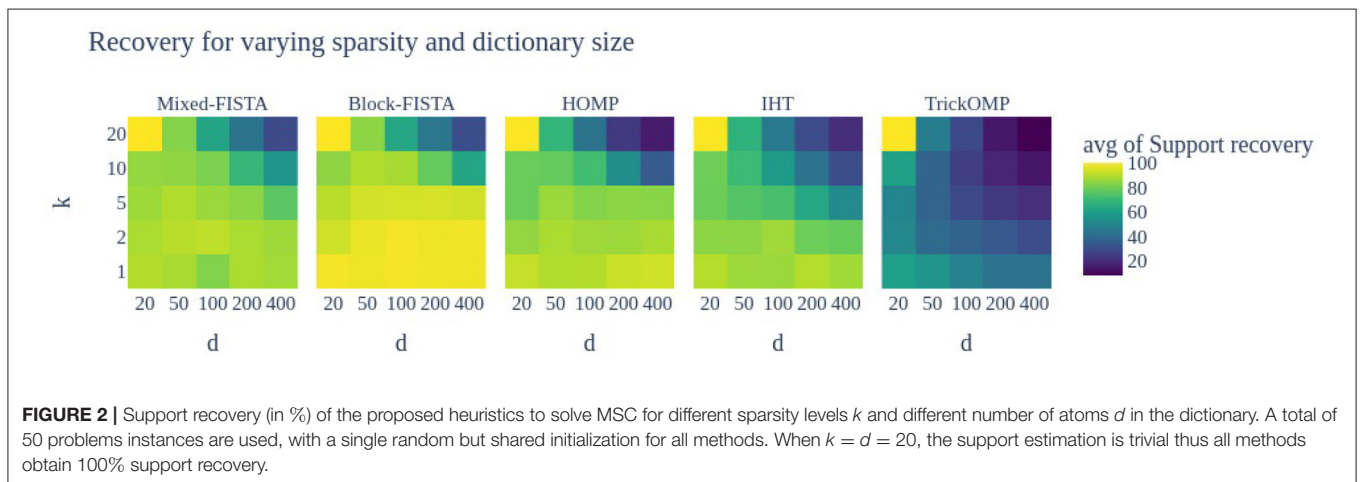
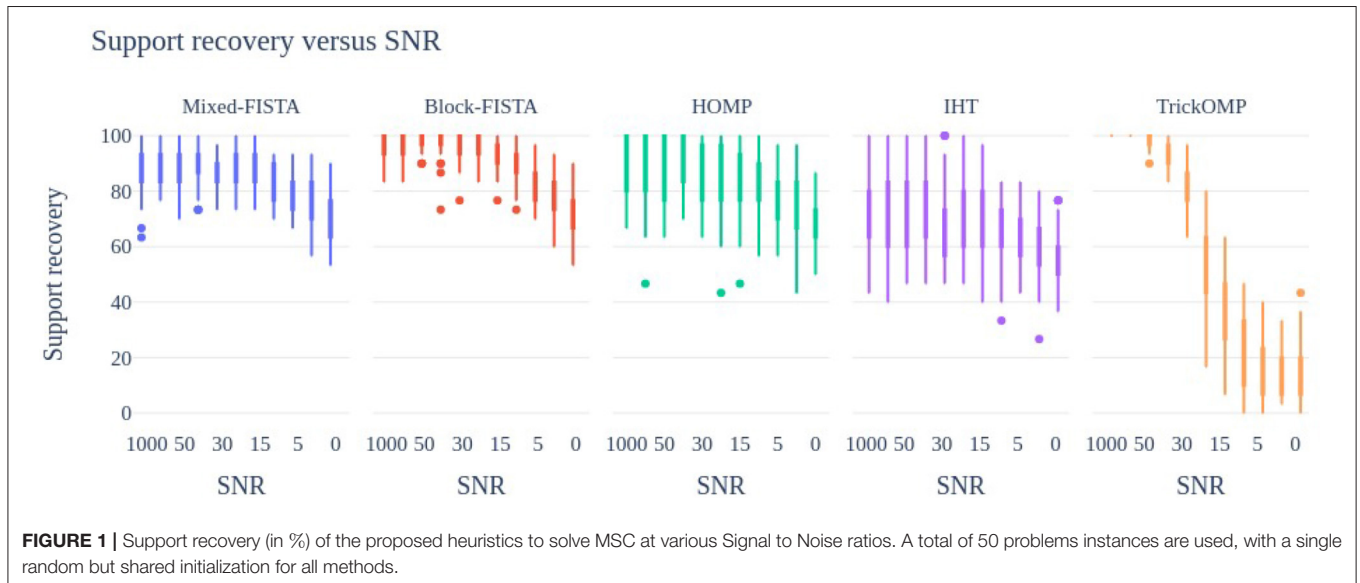


TABLE 1 | Average computation time in seconds and number of iterations with respect to (n, m) and (k, σ) .

| (n, m) | (10, 10) | (10, 50) | (10, 1,000) | (50, 10) | (50, 50) | (50, 1,000) | (1,000, 10) | (1,000, 50) | (1,000, 1,000) |
|----------|-----------|-----------|-------------|-----------|-----------|-------------|-------------|-------------|----------------|
| HOMP | 0.8, 148 | 0.9, 149 | 2, 88 | 2.6, 450 | 2, 330 | 10, 290 | 13, 280 | 13, 247 | 70, 376 |
| M.-FISTA | 0.4, 1000 | 0.4, 950 | 0.5, 875 | 0.4, 842 | 0.4, 799 | 1.6, 734 | 0.6, 393 | 1, 365 | 4.4, 388 |
| B.-FISTA | 0.3, 998 | 0.3, 898 | 0.3, 907 | 0.3, 843 | 0.3, 764 | 1, 724 | 0.5, 376 | 0.7, 369 | 4.2, 388 |
| IHT | 0.1, 356 | 0.2, 354 | 0.1, 254 | 0.07, 121 | 0.09, 105 | 0.2, 116 | 0.1, 92 | 0.2, 63 | 0.9, 80 |
| TrickOMP | 0.03, – | 0.03, – | 0.03, – | 0.03, – | 0.03, – | 0.04, – | 0.08, – | 0.07, – | 0.09, – |
| (d, k) | (50, 5) | (50, 10) | (50, 30) | (100, 5) | (100, 10) | (100, 30) | (1,000, 5) | (1,000, 10) | (1,000, 30) |
| HOMP | 0.59, 95 | 2.2, 185 | 64, 886 | 0.39, 50 | 4.3, 345 | 14, 176 | 1.3, 40 | 3.9, 53 | 53, 161 |
| M.-FISTA | 0.1, 356 | 0.3, 467 | 0.5, 456 | 0.37, 691 | 0.56, 814 | 0.78, 609 | 8.9, 1000 | 9.6, 1000 | 9.5, 1000 |
| B.-FISTA | 0.12, 492 | 0.23, 319 | 0.57, 373 | 0.23, 668 | 0.38, 679 | 0.63, 619 | 5.2, 1000 | 4.7, 369 | 5.4, 1000 |
| IHT | 0.05, 103 | 0.15, 92 | 0.69, 1000 | 0.06, 111 | 0.15, 100 | 0.52, 324 | 2.5, 215 | 3.3, 353 | 4.6, 528 |
| TrickOMP | 0.02, – | 0.08, – | 0.57, – | 0.43, – | 0.14, – | 0.55, – | 0.04, – | 0.13, – | 0.67, – |

The format is (time, it). Computations are single threaded, and use a Intel® Core™ i7-8650U CPU @ 1.90 GHz × 8 processor. TrickOMP exactly runs k steps of OMP equivalent to one HOMP inner iteration, therefore TrickOMP iterations are not reported in the table. Maximal number of iteration was set to 1,000.

time. Moreover, it is very easy to extend to nonnegative low-rank approximation models which are very common in practice. Therefore, to design an algorithm for DLRA, we shall make use of Block-FISTA (topped with a least-squares update with fixed support) as a solver for the MSC sub-problem. Note however than Block-FISTA required to select many regularization parameters, but the proposed heuristic using a fixed percentage of $\lambda_{i, \max}$ worked well in the above experiments.

5. DICTIONARY-BASED LOW RANK APPROXIMATIONS

5.1. A Generic AO Algorithm for DLRA

Now that a reasonably good heuristic for solving MSC has been found, let us introduce an AO method for DLRA based on Block-FISTA. The proposed algorithm is coined AO-DLRA and is summarized in **Algorithm 3**. It boils down to solving for X with Block-FISTA (**Algorithm 2**) and automatically computed regularization parameters, and then solving for the other blocks using any classical alternating method specific to the LRA at hand. Because solving exactly the MSC problem is difficult, even using Block-FISTA with well tuned regularization parameters, it is not guaranteed that the X update will decrease the global cost. In fact in practice the cost may go up, and storing the best update along the iterations is good practice.

5.1.1. Selecting Regularization Parameters

It had already been noted in section 3.5 that choosing the multiple regularization parameters λ_i of Block-FISTA can be challenging. In the context of Alternating Optimization, this is even more true. Indeed, the (structured) matrix B is updated at each outer iteration, therefore there is a scaling ambiguity between X and B that makes any arbitrary choice of regularization level λ_i meaningless. Moreover the values $\lambda_{i, \max}$ change at each outer iteration. Consequently, obtaining a sparsity regularization percentage α_i in each column of X at each iteration is challenging without some *ad-hoc* tuning in each outer iteration. To that end, the regularization percentages $\alpha = [\alpha_1, \dots, \alpha_r]$ are tuned inside each inner loop until the columnwise number of non-zeros reaches a target range $[k, k + \tau]$ where $\tau \geq 0$ is user-defined. This

Algorithm 3 An AO algorithm for DLRA (AO-DLRA)

Input: Initial guesses $X^{(0)}, B^{(0)}$, data Y , dictionary D , sparsity level $k \leq n$, initial regularization $\alpha \in [0, 1]^r$, iteration number l_{\max} , sparsity tolerance τ .

Output: Best estimated factors $X^{(l)}$ and $B^{(l)}$

Precompute $D^T D$, $D^T Y$ and $\sigma_D = \sigma(D^T D)$ if memory allows.

for $l = 1 \dots l_{\max}$ **do**

B update:

Compute $B^{(l)} \in \Omega_B$ that decreases the cost in Problem (2) with respect to B .

X update:

Stepsize evaluation: $\eta^{(l)} = \frac{1}{\sigma_D \sigma(B^{(l)T} B^{(l)})}$.

Compute efficiently data-factor product $(D^T Y) B^{(l)}$ and inner products $B^{(l)T} B^{(l)}$.

(*) Update $X^{(l)}$ using Block-FISTA (**Algorithm 2**) with regularization parameters α , stepsize $\eta^{(l)}$ and initial guess $X^{(l-1)}$

while $\exists i \leq r$, $\|X_i^{(l)}\|_0 \notin [k, k + \tau]$ **do**

for $i = 1 \dots r$ **do**

if $\|X_i^{(l)}\|_0 \leq k$ **then**

Decrease regularization $\alpha := \alpha/1.3$

else if $\|X_i^{(l)}\|_0 \geq k + \tau$ **then**

Increase regularization $\alpha := \min(1.01\alpha, 1)$

end if

end for

Go to (*) with $X^{(l-1)} := X^{(l)}$

end while

Unbiased estimation: update $X^{(l)}$ with fixed support $S_{X^{(l)}}$ as in section 3.1.3.

Store $X^{(l)}$ and $B^{(l)}$ if residuals $\|Y - DX^{(l)} B^{(l)}\|_F^2$ have improved with respect to previous best.

end for

range is deliberately shifted to the right so that the each column does not have size strictly less than k non-zeros. Indeed, in that situation, a few atoms would have to be chosen arbitrarily during the unbiased estimation. More precisely, when a column has too

many zeros, α_i is divided by 1.3, while it is multiplied by 1.01 if it has few non-zeros. Note that an interesting research avenue would be to use an adaption of homotopy methods [74] for Block LASSO instead of FISTA, which would remove the need for this heuristic tuning.

5.1.2. A Provably Convergent Algorithm: Inertial Proximal Alternating Linear Minimization (iPALM)

The AO-DLRA algorithm proposed above is a heuristic with several arbitrary choices and no convergence guarantees. If designing an efficient AO algorithm with convergence guarantees proved difficult, designing a convergent algorithm is in fact straightforward using standard block-coordinate non-convex methods. I focus in the following on the iPALM algorithm, which alternates between a proximal gradient step on X similar the one discussed in section 3.2.2 and a proximal gradient step on B . iPALM is guaranteed to converge to a stationary point of the DLRA cost, despite the irregularity of the semi-algebraic $\ell_{0,0}$ map. A pseudo-code for iPALM to initialize **Algorithm 3** is provided in the **Supplementary Material**.

5.1.3. Initialization Strategies

Because DLRA is a highly non-convex problem, one can only hope to reach some stationary point of the cost in Problem (2). Furthermore, the sparsity constraint on the columns of X implies that X_i must belong to a finite union of subspaces, making the problem combinatorial by nature. Using a local heuristic such as AO-DLRA or iPALM, it is expected to encounter many local minima—a fact also confirmed in practical experiments reported in section 5.2 and in previous works [13]. Therefore, providing an initial guess for X and B close to a good local minimum is important.

There are at least two reasonable strategies to initialize the DCPD model. First, for any low-rank approximation model which is mildly identifiable (such as NMF, CPD), the suggested method is to first compute the low-rank approximation with standard algorithms to estimate $A^{(0)}$ and $B^{(0)}$, and then perform sparse coding on the columns of $A^{(0)}$ to estimate $X^{(0)}$. The identifiability properties of these models should ensure that $A^{(0)}$ is well approximated by DX with sparse X . Second, several random initialization can be carried out, only to keep the best result. A third option for AO-DLRA would be to use a few iterations of the iPALM algorithm itself initialized randomly, since iPALM iterations are relatively cheap. However, it is shown in the experiments below that this method does not yield good results.

5.2. Experiments for DLRA

In the next section, two DLRA models are showcased on synthetic and real-life data. First, the Dictionary-based Matrix Factorization (DMF, see below) model is explored for the task of matrix completion in remote sensing. It is shown that DMF allows to complete entirely missing rows, something that low-rank completion cannot do. Second, nonnegative DCPD (nnDCPD) is used for denoising smooth images in the context of chemometrics, and better denoising performance are obtained with nnDCPD than with plain nonnegative CPD (nnCPD) or

when post-processing the results of nnCPD. Nevertheless, the goal of these experiments is not to establish a new state-of-the-art in these particular, well-studied applications, but rather to demonstrate the versatile problems that may be cast as DLRA and the efficiency of DLRA when compared to other low-rank strategies. The performance of AO-DLRA and iPALM in terms of support recovery and relative reconstruction error for DMF and DCPD is then further assessed on synthetic data.

5.2.1. Dictionary-Based Matrix Factorization With Application to Matrix Completion

Let us study the following Dictionary-based Matrix Factorization model:


$$\arg, \min_{X \in \mathbb{R}^{d \times r}, \forall i \leq r, \|X_i\|_0 \leq k, B \in \mathbb{R}^{m \times r}} \|Y - DXB^T\|_F^2. \quad (28)$$

Low-rank factorizations have been extensively used in machine learning for matrix completion [75], since the low-rank hypothesis serves as regularization for this otherwise ill-posed problem. A use case for DMF is the completion of a low-rank matrix which has missing rows. Using a simple low-rank factorization approach would fail since a missing row removes all information about the column-space on that row. Formally, if a data matrix $Y \approx AB$ has missing rows indexed by I , then the rows of matrix A in I cannot be estimated directly from Y . Dictionary-based low-rank matrix factorization circumvents this problem by expressing each column of A as a sparse combination of atoms in a dictionary D , such that $Y \approx DXB^T$ with X columnwise sparse. While fitting matrices X and B can be done on the known entries, the reconstruction DXB^T will provide an estimation of the whole data matrix Y , including the missing rows. Formally, first solve

$$\arg, \min_{X \in \mathbb{R}^{(n-|I|) \times r}, \forall i \leq r, \|X_i\|_0 \leq k, B \in \mathbb{R}^{m \times r}} \|Y_{:, \bar{I}} - D_{:, \bar{I}}XB^T\|_F^2 \quad (29)$$

using **Algorithm 3** and then build an estimation for the missing values in Y as $\hat{Y}_{:, I} = D_{:, I}XB^T$. Initialization is carried out using random factors sampled element-wise from a normal distribution.

In remotely acquired hyperspectral images, missing rows in the data matrix are common as they correspond to corrupted pixels. Moreover, it is well-known that many hyperspectral images are approximately low-rank. Therefore, in this toy experiment, a small portion of the Urban hyperspectral image is used to showcase the proposed inpainting strategy. Urban is often considered to be between rank 4, 5, or 6 [76]. I will use $r = 4$ hereafter. Urban is a collection of 307×307 images collected on 162 clean bands. For the sake of simplicity, only a 20×20 patch is considered, with 50 randomly-chosen pixels removed from this patch in all bands, and therefore after pixel vectorization the data matrix $Y \in \mathbb{R}^{400 \times 162}$ has 50 missing rows. Since columns of factor A in this factorization should stand for patches of abundance maps, they are reasonably sparse in a 2D-Discrete Cosine Transform dictionary D , here using $d = 400$ atoms. Note that similar strategies for HSI denoising have been studied in the literature albeit without columnwise sparsity imposed on X , see [77].



We measure the performance of two strategies: DMF computed on known pixels solving Problem (29), and OMP for each band individually. Both approaches use the same dictionary D . To evaluate performance, the test estimation error on missing pixels $\|Y_I - \hat{Y}_I\|_F / \|Y_I\|_F$ is computed alongside with the average Spectral Angular Mapper (SAM) $\frac{1}{|I|} \sum_{i \in I} \arccos(\frac{Y_i^T \hat{Y}_i}{\|Y_i\|_2 \|\hat{Y}_i\|_2})$ on all missing pixels. The sparsity level k is defined on the grid $[10, 30, 50, 70, 100, 120, 150^*, 200^*, 250^*]$, where k^* is not used for the OMP inpainting because of memory issues. Results are averaged on $N = 20$ different initializations. For AO-DLRA, the initial regularization α is set to 5×10^{-3} , and $\tau = 20$.

Figure 3 shows the reconstruction error and SAM obtained after each initialization for various sparsity levels. First, for both metrics, there exist a clear advantage of the DMF approach when compared to sparse coding band per band. In particular, the band-wise OMP reconstruction does not yield good spectral reconstruction. Second, DMF apparently works similarly to sparse coding approaches for inpainting: if the sparsity level is too low the reconstruction is not precise, but if the sparsity level is too large the reconstruction is biased. Therefore, DMF effectively allows to perform inpainting with sparse coding on the factors of a low-rank matrix factorization.

5.2.2. Dictionary-Based Smooth Canonical Polyadic Decomposition With Application to Data Denoising

The second study examines the DCPD discussed around Equation (4) to perform denoising using smoothness. In this context, the data tensor Y is noisy, meaning formally that

$$\tilde{Y} = Y + \epsilon, \text{ and } Y = A(B \odot C)^T \quad (30)$$

where ϵ has a large power compared to $A(B \odot C)^T$ (e.g. Signal to Noise Ratio at -8.7dB in the following). Furthermore, let us suppose that A has smooth columns. The dictionary constraint can enforce smoothness on A by choosing D as a large collection of smooth atoms, in this case B-splines⁸. Because the first mode is constrained such that $A = DX$, each column of A is a sparse combination of smooth functions and will therefore itself

be smooth. The sparsity constraint k prevents the use of too many splines and ensures that A is indeed smooth. Hereafter, the sparsity value is fixed to $k = 6$.

There has been significant previous works on smooth CPD, perhaps most related to the proposed approach is the work of [8]. Their method also consists in choosing a dictionary D of B-splines. However, this dictionary has very few atoms (the actual number is determined by cross-validation), and picking the knots for the splines requires either time-consuming hand crafting, or some cross validation set. The advantage of their approach is that no sparsity constraint is imposed since there are already so few atoms in D , and the problem becomes equivalent to CPD on the smoothed data.

The rationale however is that heavily crafting the splines is unnecessary. Using DCPD allows an automatic picking of good (if not best) B-splines. Furthermore, each component in the CPD may use different splines while the method of [8] uses the same splines for all the components. Hand-crafting is still required to build the dictionary, but one does not have to fear introducing an inadequate spline.

An advantage of B-splines is that they are nonnegative, therefore one can compute nonnegative DCPD by imposing nonnegativity on the sparse coefficients X as explained in section 3.4. Imposing nonnegativity in the method of [8] is not as straightforward albeit doable [35].

For this study the toy fluorescence spectroscopy dataset “amino-acids” available online⁹ is used. Its rank is known to be $r = 3$, and dimensions are $201 \times 61 \times 5$. Fluorescence spectroscopy tensors are nonnegative, low-rank and feature smooth factors. In fact factors are smooth on two modes related to excitation and emission spectra, therefore a double-constrained DCPD is also of interest. It boils down to solving

$$\underset{\|X_i^{(A)}\|_0 \leq k^{(A)}, \|X_i^{(B)}\|_0 \leq k^{(B)}, C \in \mathbb{R}^{m_2 \times r}}{\text{minimize}} \quad \|\tilde{Y} - D^{(A)}X^{(A)}(D^{(B)}X^{(B)} \odot C)^T\|_F^2. \quad (31)$$

where $D^{(A,B)}$ and $k^{(A,B)}$ are the respective dictionaries of sizes 201×180 and 61×81 and sparsity targets for modes one and

⁸The exact implementation of D is detailed in the code.

⁹http://www.models.life.ku.dk/Amino_Acid_fluo.

two. The third mode factor contains relative concentrations [78]. Additional Gaussian noise is added to the data so that the effective SNR used in the experiments is -8.7dB . Because CPD is identifiable, in particular for the amino-acids dataset, nnCPD of the noisy data is used for initialization, computed using Hierarchical Alternating Least Squares [45]. We compare sparse coding one or two modes of the output of HALS with AO-DCPD (i.e., AO-DLRA used to compute DCPD) with smoothness on one mode, two modes, and on two modes with nonnegativity (AO-nnDCPD). We set $\alpha = 10^{-3}$ and $\tau = 5$, and $k^{(A)} = k^{(B)} = 6$.

Figure 4 shows the reconstructed factors and the relative error $\|\hat{Y} - Y\|_F / \|Y\|_F$ with respect to the true data, and Figure 5 shows one slice (here the fourth one) of the reconstructed tensor. It can be observed both graphically and numerically that DCPD techniques are overall superior to post-processing the output of HALS. In particular, the factors recovered using nnDCPD with dictionary constraints on two modes are very close to the true factors (obtained from the nnCPD of the clean data) using only $k = 6$ splines at most.

5.2.3. Performance of AO-DLRA for DMF and DCPD on Synthetic Data

To assess the performance of AO-DLRA in computing DMF and DCPD, the support recovery and reconstruction error are monitored on synthetic data. For both model, a single initialization is performed for $N = 100$ problem instances with the same hyperparameters. Matrices D, X, B, C involved in both problems and the noise tensors are generated as in section 3.5. The rank is $r = 6$ for a sparsity level of $k = 8$ and the conditioning of B is set to 2×10^2 . In the DMF experiment, the sizes are $n = m = 50, d = 60$ and the SNR is 100dB. For DPCD, we set $n = 20, m_1 = 21, m_2 = 22, d = 30$ and the SNR is 30 dB. We set $\alpha = 10^{-2}, \tau = 20$ for AO-DMF and $\alpha = 10^{-4}, \tau = 20$ for AO-DCPD.

A few strategies are compared: AO-DLRA initialized randomly, iPALM initialized randomly, and AO-DLRA initialized with iPALM. The same random initialization is used for all methods in each problem instance. For DCPD, we also consider AO-DLRA and iPALM initialized with a CPD solver (here Alternating Least Squares), and sparse coding the output

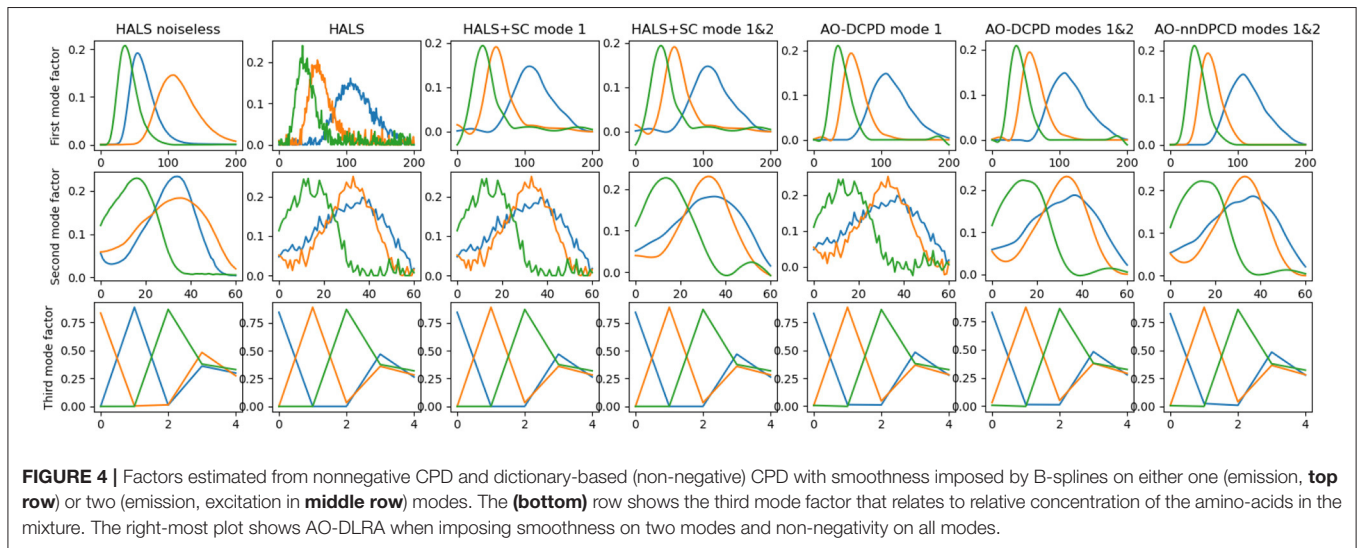


FIGURE 4 | Factors estimated from nonnegative CPD and dictionary-based (non-negative) CPD with smoothness imposed by B-splines on either one (emission, top row) or two (emission, excitation in middle row) modes. The (bottom) row shows the third mode factor that relates to relative concentration of the amino-acids in the mixture. The right-most plot shows AO-DLRA when imposing smoothness on two modes and non-negativity on all modes.

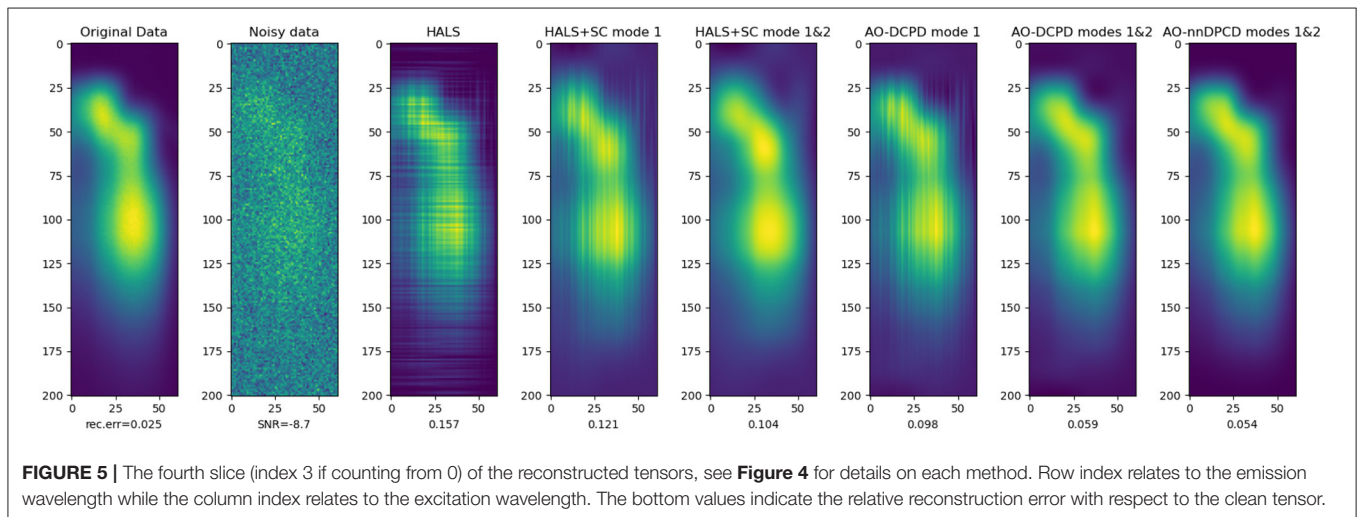


FIGURE 5 | The fourth slice (index 3 if counting from 0) of the reconstructed tensors, see Figure 4 for details on each method. Row index relates to the emission wavelength while the column index relates to the excitation wavelength. The bottom values indicate the relative reconstruction error with respect to the clean tensor.

of the ALS. AO-DLRA always stops after at most 100 iterations, while iPALM runs for at most 1,000 iterations or when relative error decrease is below 10^{-8} . The safeguard for stepsize in iPALM is set to $\mu = 0.5$ for DMF and $\mu = 1$ for DCPD.

Figure 6 shows the obtained results. It can be observed that both in DMF and DCPD, iPALM initialization is not significantly better than random initialization, and iPALM in fact provides quite poor results on its own. For DCPD, support recovery and reconstruction error is always better with AO-DLRA than when sparse coding the output of the ALS, even with random initialization. Overall, these experiments show that while the DLRA problem is challenging (the optimal support is almost always never found), reasonably good results are obtained using the proposed AO-DLRA algorithm. Furthermore, the support recovery scores of DMF using AO-DLRA are much higher than what could be obtained by chance, randomly picking elements in the support. This observation hints toward the identifiability of DMF. Remember indeed that without the dictionary constraint, matrix factorization is never unique when $r > 1$, thus any posterior support identification for a matrix A given a factorization $Y = AB$ should fail.

6. CONCLUSIONS AND OPEN QUESTIONS

In this manuscript, a Dictionary-based Low-Rank Approximation framework has been proposed. It allows to constrain any factor in a LRA to live in the union of k -dimensional subspaces generated by subsets of columns of a given dictionary. DLRA is shown to be useful for various signal processing tasks such as image completion or image denoising. A contribution of this work is an Alternating Optimization algorithm (AO-DLRA) to compute candidate solutions to DLRA. Additionally, the subproblem of estimating the sparse codes of a factor in a LRA, coined Mixed Sparse Coding, is extensively discussed. A heuristic convex relaxation adapted from LASSO is shown to perform very well for solving MSC when compared to other modified sparse coding strategies, and along the way, several theoretical results regarding MSC are provided.

There are several research directions that stem from this work. First, the identifiability properties of DLRA have not been addressed here. It was shown in previous works [13] that dictionary-based matrix factorization is identifiable when

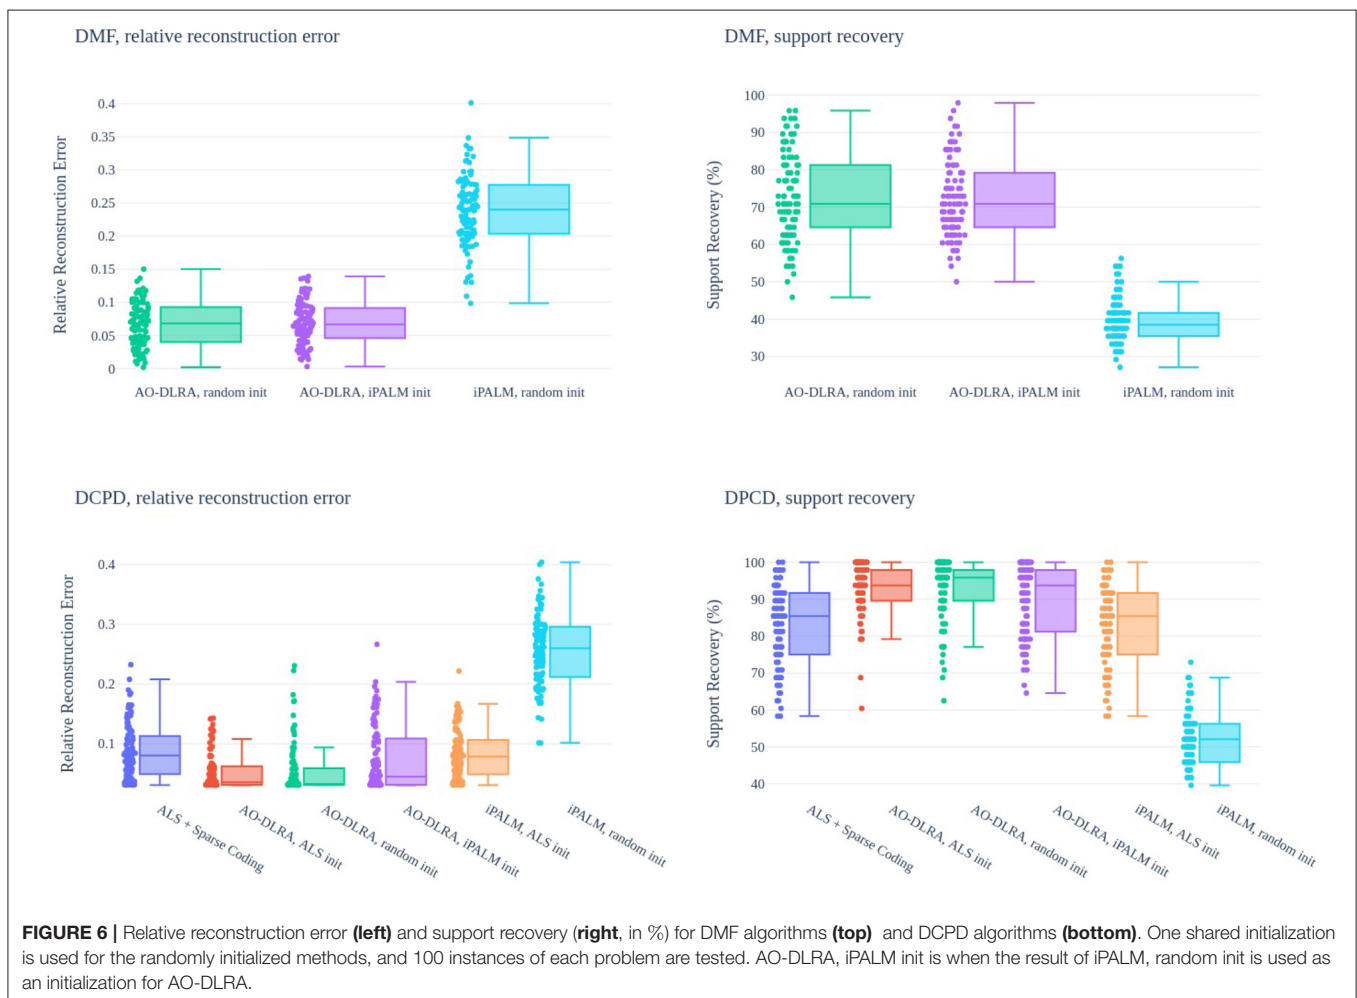


FIGURE 6 | Relative reconstruction error (left) and support recovery (right, in %) for DMF algorithms (top) and DCPD algorithms (bottom). One shared initialization is used for the randomly initialized methods, and 100 instances of each problem are tested. AO-DLRA, iPALM init is when the result of iPALM, random init is used as an initialization for AO-DLRA.

sparsity is exactly one, but the general case is harder to analyse. The identifiability analysis is furthermore model dependent while this work aims at tackling the computation of any DLRA. Second, while the proposed AO-DLRA algorithm proved efficient in practice, its convergence properties are lacking. It is reasonably easy to design an AO algorithm with guarantees for DLRA, but I could not obtain an algorithm with convergence guarantees which performance matched the proposed AO-DLRA. Finally, the proposed DLRA model could be extended to a supervised setting, where D is trained in a similar fashion to Dictionary Learning [15]. This would mean computing a DLRA for several tensors with the same dictionary, a problem closely related to coupled matrix and tensor factorization with linearly coupled factors [79].

DATA AVAILABILITY STATEMENT

The scripts used to generate the synthetic dataset for this study can be found alongside the code in the online repositories <https://github.com/cohenjer/mscode> and <https://github.com/cohenjer/dlra>. Further inquiries can be directed to the corresponding author/s.

REFERENCES

- Kolda TG, Bader BW. Tensor decompositions and applications. *SIAM Rev.* (2009) 51:455–500. doi: 10.1137/07070111X
- Sidiropoulos ND, De Lathauwer L, Fu X, Huang K, Papalexakis EE, Faloutsos C. Tensor decomposition for signal processing and machine learning. *IEEE Trans Signal Process.* (2017) 65:3551–82. doi: 10.1109/TSP.2017.2690524
- Gillis N. *Nonnegative Matrix Factorization*. Philadelphia: Society for Industrial and Applied Mathematics (2020). doi: 10.1137/1.9781611976410
- Kruskal JB. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear Algeb Appl.* (1977) 18:95–138. doi: 10.1016/0024-3795(77)90069-6
- Hoyer PO. Non-negative sparse coding. In: *IEEE Workshop on Neural Networks for Signal Processing*. Martigny (2002). p. 557–65.
- Mørup M, Hansen LK, Arnfred SM. Algorithms for sparse nonnegative Tucker decompositions. *Neural Comput.* (2008) 20:2112–31. doi: 10.1162/neco.2008.11-06-407
- Hennequin R, Badeau R, David B. Time-dependent parametric and harmonic templates in non-negative matrix factorization. In: *Proceedings of the 13th International Conference on Digital Audio Effects (DAFx)*. Graz (2010).
- Timmerman ME, Kiers HA. Three-way component analysis with smoothness constraints. *Comput Stat Data Anal.* (2002) 40:447–70. doi: 10.1016/S0167-9473(02)00059-2
- Fu X, Huang K, Sidiropoulos ND. On identifiability of nonnegative matrix factorization. *IEEE Signal Process Lett.* (2018) 25:328–32. doi: 10.1109/LSP.2018.2789405
- Gillis N, Luce R. Robust near-separable nonnegative matrix factorization using linear optimization. *J Mach Learn Res.* (2014) 15:1249–80.
- Silva VD, Lim LH. Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM J Matrix Anal Appl.* (2008) 30:1084–127. doi: 10.1137/06066518X
- Mohlenkamp MJ. The dynamics of swamps in the canonical tensor approximation problem. *SIAM J Appl Dyn Syst.* (2019) 18:1293–333. doi: 10.1137/18M1181389
- Cohen JE, Gillis N. Dictionary-based tensor canonical polyadic decomposition. *IEEE Trans Signal Process.* (2018) 66:1876–89. doi: 10.1109/TSP.2017.2777393

AUTHOR CONTRIBUTIONS

JC is the sole contributor to this work. He studied the theory around Mixed Sparse Coding and Dictionary-based low-rank approximations, implemented the methods, conducted the experiments, and wrote the manuscript.

FUNDING

This work was funded by ANR JCJC LoRAiA ANR-20-CE23-0010.

ACKNOWLEDGMENTS

The author thanks Rémi Gribonval for commenting an early version of this work.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fams.2022.801650/full#supplementary-material>

- Kjersti E, Sven OA, John HH. Method of optimal directions for frame design. *IEEE Int Conf Acoust Speech Signal Process Proc.* (1999) 5:2443–6.
- Mairal J, Bach F, Ponce J, Sapiro G. Online learning for matrix factorization and sparse coding. *J Mach Learn Res.* (2010) 11:19–60.
- Huang K, Sidiropoulos ND, Liavas AP. A flexible and efficient algorithmic framework for constrained matrix and tensor factorization. *IEEE Trans Signal Process.* (2016) 64:5052–65. doi: 10.1109/TSP.2016.2576427
- Fu X, Vervliet N, De Lathauwer L, Huang K, Gillis N. Computing large-scale matrix and tensor decomposition with structured factors: a unified nonconvex optimization perspective. *IEEE Signal Process Mag.* (2020) 37:78–94. doi: 10.1109/MSP.2020.3003544
- Aharon M, Elad M, Bruckstein AM. K-SVD and its non-negative variant for dictionary design. In: *Wavelets XI*. Vol. 5914. San Diego, CA: International Society for Optics and Photonics (2005). p. 591411. doi: 10.1117/12.613878
- Xu Y. Alternating proximal gradient method for sparse nonnegative Tucker decomposition. *Math Programm Comput.* (2015) 7:39–70. doi: 10.1007/s12532-014-0074-y
- Kolda TG, Hong D. Stochastic gradients for large-scale tensor decomposition. *SIAM J Math Data Sci.* (2020) 2:1066–95. doi: 10.1137/19M1266265
- Marmin A, Goulart JHM, Févotte C. Joint majorization-minimization for nonnegative matrix factorization with the β -divergence. *arXiv [preprint]. arXiv:210615214* (2021).
- Brewer J. Kronecker products and matrix calculus in system theory. *IEEE Trans Circuits Syst.* (1978) 25:772–81. doi: 10.1109/TCS.1978.1084534
- Pearson K. On lines and planes of closest fit to systems of points in space. *Lond Edinburgh Dublin Philos Mag J Sci.* (1901) 2:559–72. doi: 10.1080/14786440109462720
- Jolliffe I. *Principal Component Analysis*. Wiley Online Library (2002).
- Comon P. Independent component analysis, a new concept? *Signal Process.* (1994) 36:287–314. doi: 10.1016/0165-1684(94)90029-9
- Lee DD, Seung HS. Learning the parts of objects by non-negative matrix factorization. *Nature.* (1999) 401:788–91. doi: 10.1038/44565
- Gribonval R, Lesage S. A survey of sparse component analysis for blind source separation: principles, perspectives, and new challenges. In: *ESANN'06 Proceedings-14th European Symposium on Artificial Neural Networks*. Bruges (2006). p. 323–30.

28. Donoho DL, Stodden V. When does non-negative matrix factorization give a correct decomposition into parts? In: *Advances in Neural Information Processing 16*. Vancouver (2003).
29. Cohen JE, Gillis N. Identifiability of complete dictionary learning. *SIAM J Math Data Sci.* (2019) 1:518–36. doi: 10.1137/18M1233339
30. Gillis N, Glineur F. Accelerated multiplicative updates and hierarchical ALS algorithms for nonnegative matrix factorization. *Neural Comput.* (2012) 24:1085–105. doi: 10.1162/NECO_a_00256
31. Natarajan B. Sparse approximate solutions to linear systems. *SIAM J Comput.* (1995) 24:227–34. doi: 10.1137/S0097539792240406
32. Vavasis SA. On the complexity of nonnegative matrix factorization. *SIAM J Optim.* (2010) 20:1364–77. doi: 10.1137/070709967
33. Bro R. PARAFAC, tutorial and applications. *Chemom Intel Lab Syst.* (1997) 38:149–71. doi: 10.1016/S0169-7439(97)00032-4
34. Becker H, Albera L, Comon P, Gribonval R, Wendling F, Merlet I. Brain source imaging: from sparse to tensor models. *IEEE Signal Proc Mag.* (2015) 32:100–12. doi: 10.1109/MSP.2015.2413711
35. Cohen JE, Cabral-Farias R, Comon P. Fast decomposition of large nonnegative tensors. *IEEE Signal Process Lett.* (2015) 22:862–6. doi: 10.1109/LSP.2014.2374838
36. Kossaiji J, Lipton ZC, Kolbeinsson A, Khanna A, Furlanello T, Anandkumar A. Tensor regression networks. *J Mach Learn Res.* (2020) 21:1–21.
37. Hitchcock FL. The expression of a tensor or a polyadic as a sum of products. *J Math Phys.* (1927) 6:164–89. doi: 10.1002/sapm192761164
38. Harshman RA. PARAFAC2: mathematical and technical notes. *UCLA Work Pap Phonet.* (1972) 22:122215.
39. Sidiropoulos ND, Bro R. On the uniqueness of multilinear decomposition of N-way arrays. *J Chemometr.* (2000) 14:229–39. doi: 10.1002/1099-128X(200005/06)14:3<229::AID-CEM587>3.0.CO;2-N
40. Domanov I, De Lathauwer L. On the uniqueness of the canonical polyadic decomposition of third-order tensors - part i: basic results and uniqueness of one factor matrix. *SIAM J Matrix Anal Appl.* (2013) 34:855–75. doi: 10.1137/120877234
41. Lim LH, Comon P. Nonnegative approximations of nonnegative tensors. *J Chemometr.* (2009) 23:432–41. doi: 10.1002/cem.1244
42. Roald M, Schenker C, Cohen J, Acar E. PARAFAC2 AO-ADMM: constraints in all modes. In: *EUSIPCO*. Dublin (2021). doi: 10.23919/EUSIPCO54536.2021.9615927
43. Tucker LR. Some mathematical notes on three-mode factor analysis. *Psychometrika.* (1966) 31:279–311. doi: 10.1007/BF02289464
44. Oseledets IV. Tensor-train decomposition. *SIAM J Sci Comput.* (2011) 33:2295–317. doi: 10.1137/090752286
45. Phan AH, Cichocki A. Extended HALS algorithm for nonnegative Tucker decomposition and its applications for multiway analysis and classification. *Neurocomputing.* (2011) 74:1956–69. doi: 10.1016/j.neucom.2010.06.031
46. Marmoret A, Cohen JE, Bertin N, Bimbot F. Uncovering audio patterns in music with nonnegative tucker decomposition for structural segmentation. In: *ISMIR 2020-21st International Society for Music Information Retrieval*. Montreal (2020).
47. Foucart S, Rauhut H. In: Birkhausen, editor. *A Mathematical Introduction to Compressive Sensing*. Applied and Numeric Harmonic Analysis. Birkhauser; Springer (2013). doi: 10.1007/978-0-8176-4948-7
48. Mallat SG, Zhang Z. Matching pursuits with time-frequency dictionaries. *IEEE Trans Signal Process.* (1993) 41:3397–415. doi: 10.1109/78.258082
49. Pati YC, Rezaifar R, Krishnaprasad PS. Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In: *1993 Conference Record of The 27th Asilomar Conference on Signals, Systems and Computers*. Pacific Grove, CA (1993). p. 40–4.
50. Tropp JA. Greed is good: algorithmic results for sparse approximation. *IEEE Trans Inform Theory.* (2004) 50:2231–42. doi: 10.1109/TIT.2004.834793
51. Blumensath T, Davies ME. Iterative thresholding for sparse approximations. *J Four Anal Appl.* (2008) 14:629–54. doi: 10.1007/s00041-008-9035-z
52. Parikh N, Boyd SP. Proximal algorithms. *Found Trends Optim.* (2014) 1:127–239. doi: 10.1561/24000000003
53. Tropp JA. Just relax: Convex programming methods for identifying sparse signals in noise. *IEEE Trans Inform Theory.* (2006) 52:1030–51. doi: 10.1109/TIT.2005.864420
54. Nesterov Y. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. In: *Soviet Mathematics Doklady*. Vol. 27 (1983). p. 372–6.
55. Beck A, Teboulle M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J Imaging Sci.* (2009) 2:183–202. doi: 10.1137/080716542
56. Cotter SF, Rao BD, Engan K, Kreutz-Delgado K. Sparse solutions to linear inverse problems with multiple measurement vectors. *IEEE Trans Signal Process.* (2005) 53:2477–88. doi: 10.1109/TSP.2005.849172
57. Elhamifar E, Sapiro G, Vidal R. See all by looking at a few: sparse modeling for finding representative objects. In: *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Providence (2012). p. 1600–7. doi: 10.1109/CVPR.2012.6247852
58. Iordache MD, Bioucas-Dias JM, Plaza A. Collaborative sparse regression for hyperspectral unmixing. *IEEE Trans Geosci Remote Sensing.* (2014) 52:341–54. doi: 10.1109/TGRS.2013.2240001
59. Cohen JE. About notations in multiway array processing. *arXiv [preprint]. arXiv:151101306*. (2015).
60. Traonmilin Y, Gribonval R. Stable recovery of low-dimensional cones in Hilbert spaces: one RIP to rule them all. *Appl Comput Harm Anal.* (2018) 45:170–205. doi: 10.1016/j.acha.2016.08.004
61. Tsiiligkaridis T, Hero AO. Covariance estimation in high dimensions via Kronecker product expansions. *IEEE Trans Signal Process.* (2013) 61:5347–60. doi: 10.1109/TSP.2013.2279355
62. Golub GH, Loan CFV. *Matrix Computations*. Baltimore: The John Hopkins University Press (1989).
63. Donoho DL, Elad M. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. *Proc Natl Acad Sci USA.* (2003) 100:2197–202. doi: 10.1073/pnas.0437847100
64. Bourguignon S, Ninin J, Carfantan H, Mongeau M. Exact sparse approximation problems via mixed-integer programming: formulations and computational performance. *IEEE Trans Signal Process.* (2015) 64:1405–19. doi: 10.1109/TSP.2015.2496367
65. Nadisic N, Vandaele A, Gillis N, Cohen JE. Exact sparse nonnegative least squares. In: *ICASSP*. Barcelona (2020). p. 5395–9. doi: 10.1109/ICASSP40776.2020.9053295
66. Tibshirani R. Regression shrinkage and selection via the lasso. *J R Stat Soc Ser B.* (1996) 58:267–88. doi: 10.1111/j.2517-6161.1996.tb02080.x
67. Attouch H, Bolte J, Svaiter BF. Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized Gauss-Seidel methods. *Math Programm.* (2013) 137:91–129. doi: 10.1007/s10107-011-0484-9
68. Blumensath T, Davies ME. Iterative hard thresholding for compressed sensing. *Appl Comput Harm Anal.* (2009) 27:265–74. doi: 10.1016/j.acha.2009.04.002
69. Rubinstein R, Zibulevsky M, Elad M. Efficient implementation of the K-SVD algorithm using batch orthogonal matching pursuit. *CS Technion.* (2008) 40:1–15.
70. Bogdan M, van den Berg E, Sabatti C, Su W, Candès EJ. SLOPE-Adaptive variable selection via convex optimization. *Ann Appl Stat.* (2015) 9:1103–40. doi: 10.1214/15-AOAS842
71. Quattoni A, Carreras X, Collins M, Darrell T. An efficient projection for ℓ_1, ∞ regularization. In: *Proceedings of the 26th Annual International Conference on Machine Learning*. Montreal (2009). p. 857–64. doi: 10.1145/1553374.1553484
72. Bejar B, Dokmanic I, Vidal R. The fastest ℓ_1, ∞ prox in the west. *IEEE Trans Pattern Anal Mach Intell.* (2021). doi: 10.1109/TPAMI.2021.3059301
73. Cohen JE. Computing the proximal operator of the $\ell_1, 1$ induced matrix norm. *arXiv [preprint]. arXiv:200506804*. (2020).
74. Osborne MR, Presnell B, Turlach BA. A new approach to variable selection in least squares problems. *IMA J Numer Anal.* (2000) 20:389–403. doi: 10.1093/imanum/20.3.389
75. Candès EJ, Recht B. Exact matrix completion via convex optimization. *Found Comput Math.* (2009) 9:717–72. doi: 10.1007/s10208-009-9045-5
76. Zhu F. Hyperspectral unmixing: ground truth labeling, datasets, benchmark performances and survey. *arXiv [preprint]. arXiv:170805125* (2017).
77. Zhuang L, Bioucas-Dias JM. Fast hyperspectral image denoising and inpainting based on low-rank and sparse representations. *IEEE J*

- Select Top Appl Earth Observ Remote Sensing*. (2018) 11:730–42. doi: 10.1109/JSTARS.2018.2796570
78. Bro R. *Multi-way analysis in the food industry: models, algorithms, and applications* (Ph.D.). University of Amsterdam, Amsterdam, Netherlands (1998).
79. Schenker C, Cohen JE, Acar E. A Flexible optimization framework for regularized matrix-tensor factorizations with linear couplings. *IEEE J Select Top Signal Process*. (2020) 15:506–21. doi: 10.1109/JSTSP.2020.3045848

Conflict of Interest: The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Cohen. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.