



HAL
open science

Continuous dictionaries meet low-rank tensor approximations

Clément Elvira, Jérémy E Cohen, Cédric Herzet, Rémi Gribonval

► **To cite this version:**

Clément Elvira, Jérémy E Cohen, Cédric Herzet, Rémi Gribonval. Continuous dictionaries meet low-rank tensor approximations. iTwist 2020 - International Traveling Workshop on Interactions between low-complexity data models and Sensing Techniques, Jun 2020, Nantes, France. pp.1-3. hal-02567115

HAL Id: hal-02567115

<https://hal.science/hal-02567115v1>

Submitted on 7 May 2020

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

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

Continuous dictionaries meet low-rank tensor approximations

Clément Elvira¹, Jeremy E. Cohen², Cédric Herzet³, Rémi Gribonval⁴.

¹ Univ. Bretagne Sud, ²Univ Rennes, Inria, CNRS, IRISA, ³ Inria, ⁴ Univ Lyon, Inria, CNRS, ENS de Lyon, UCB Lyon 1, LIP UMR 5668, F-69342, Lyon, France.

Abstract— In this short paper we bridge two seemingly unrelated sparse approximation topics: continuous sparse coding and low-rank approximations. We show that for a specific choice of continuous dictionary, linear systems with nuclear-norm regularization have the same solutions as a BLasso problem. Although this fact was already partially understood in the matrix case, we further show that for tensor data, using BLasso solvers for the low-rank approximation problem leads to a new branch of optimization methods yet vastly unexplored. In particular, the proposed Frank-Wolfe algorithm is showcased on an automatic tensor rank selection problem.

1 Sparse and low-rank reconstructions

Sparse decompositions. The “ k -sparse representation problem” consists in finding an accurate representation of some vector $\mathbf{y} \in \mathbb{R}^m$ as the linear combination of k elements from some set $\mathcal{D} \subset \mathbb{R}^m$, commonly called *dictionary*. Finding the best k -sparse representation of $\mathbf{y} \in \mathbb{R}^m$ in \mathcal{D} is known to be an NP-hard problem for most dictionaries, see [1, Section 2.3]. Hence, the design of tractable methods to find (provably) good sparse representations of certain families of signals \mathbf{y} has been a very active area of research during the last decades.

Although most contributions focus on the case where \mathcal{D} contains a finite number of elements, several authors have recently placed the sparse-representation problem in the context of “continuous” dictionaries. In this framework, \mathcal{D} is made up of an infinite (uncountable) number of elements continuously indexed by some parameter θ :

$$\mathcal{D} = \{\mathbf{d}(\theta) \in \mathbb{R}^m : \theta \in \Theta\} \quad (1)$$

where $\mathbf{d}: \Theta \rightarrow \mathbb{R}^m$ is a continuous function and $\Theta \subseteq \mathbb{R}^d$. The generalization of the well-known Lasso problem to context of continuous dictionaries (1) was considered in [2]. It involves the optimization with respect to a measure μ belonging to the space of Radon measures over Θ , say $\mathcal{M}(\Theta)$:

$$\mu^* \in \arg \min_{\mu \in \mathcal{M}(\Theta)} \frac{1}{2} \|\mathbf{y} - \mathcal{D}\mu\|_2^2 + \lambda \|\mu\|_{\text{TV}}, \quad (2)$$

where $\mathcal{D}\mu \triangleq \int \mathbf{d}(\theta) d\mu(\theta)$, and $\|\cdot\|_{\text{TV}}$ is the so-called “total variation” norm, see *e.g.*, [3, Eq (7)]. We note that, despite of its apparent complexity (optimization over a space of measures), efficient procedures based on the conditional gradient method [4, Section 2.2] have been proposed that operate on measures with finite support, see *e.g.*, [5, 6].

Low-rank approximation of matrices. Inspired by the success of the sparse model, several authors proposed to tackle inverse problems involving different types of low-dimensional structures. One striking example is the reconstruction of low-rank matrices considered by Candès *et al* in [7]. More specifically, the authors addressed the problem of reconstructing a

low-rank matrix $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ from partial observations $\mathbf{y} \in \mathbb{R}^m$ as the solution of the following optimization problem:

$$\mathbf{X}^* \in \arg \min_{\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}} \frac{1}{2} \|\mathbf{y} - \mathcal{A}\mathbf{X}\|_2^2 + \lambda \|\mathbf{X}\|_*, \quad (3)$$

Here $\mathcal{A}: \mathbb{R}^{n_1 \times n_2} \rightarrow \mathbb{R}^m$ plays the role of some linear observation operator and $\|\cdot\|_*$ is the so-called “nuclear” norm defined as:

$$\|\mathbf{X}\|_* = \sum_{\ell=1}^{\min(n_1, n_2)} \sigma_\ell(\mathbf{X}), \quad (4)$$

where $\sigma_\ell(\mathbf{X})$ denotes the ℓ th singular value of \mathbf{X} . It can be shown that the rank of solutions to (3) is related to λ , *i.e.* the nuclear norm induces low-rank solutions. The success of the method proposed by Candès *et al.* revolves around: (i) theoretical guarantees that the solution to (3) may be the solution of the NP-hard low-rank approximation problem; (ii) efficient solvers such as FISTA [8] which rely on the closed-form expression of the proximal operator of the matrix nuclear norm.

Low-rank approximation of tensors. Given the success of convex-relaxation methods for the reconstruction of low-rank matrices, it is tempting to extend the ingredients of problem (3) to the low-rank decomposition of t -dimensional tensors. To this end, following [9], the tensor nuclear norm is defined as the gauge function of the convex hull of normalized rank-one tensors:

$$\|\mathbf{X}\|_* = \inf \left\{ \sum_{\ell < \infty} |c_\ell| : \mathbf{X} = \sum_{\ell < \infty} c_\ell \mathbf{u}_\ell^{(1)} \otimes \dots \otimes \mathbf{u}_\ell^{(t)}, \mathbf{u}_\ell^{(i)} \in \mathbb{S}^{n_i} \right\},$$

where \otimes denotes the tensor product and \mathbb{S}^{n_i} is the n_i -dimensional unit-sphere. Indexation $\ell < \infty$ means the number of summands must be finite. It is well-known that this definition is equivalent to (4) when $t = 2$ [10].

With this generalized definition (and a proper definition of \mathcal{A}), problem (3) straightforwardly extends to the multi-dimensional tensorial case. Unfortunately, as shown in [10], the evaluation of the nuclear norm is NP-difficult as soon as $t > 2$. Although this seems to mean that (convex) nuclear norm regularized approximation (3) is not easier to solve than (non-convex) tensor low-rank approximation, in this contribution we propose a Frank-Wolfe algorithm working on a related BLasso problem that tries to solve (3) globally. We discuss research directions to hopefully provide appropriate assumptions for the global convergence of the algorithm, and show practical results for tensor rank selection.

2 From low rank to sparsity

The idea underlying the reformulation of the low-rank approximation problem as a sparse-representation problem is as follows: any rank- k tensor can be rewritten as the sum of k rank-1

tensors; hence, rank- k tensors have k -sparse representations in the (continuous) dictionary of rank-1 tensors. In the framework of sparse representations in continuous dictionaries, one may define the parameter set Θ as

$$\Theta = \left\{ (\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(t)}) : \mathbf{u}^{(i)} \in \mathbb{S}^{n_i} \right\} \quad (5)$$

and the ‘‘atom function’’ as

$$\mathbf{d}: \quad \Theta \quad \rightarrow \quad \mathbb{R}^m \\ (\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(t)}) \mapsto \mathcal{A}(\mathbf{u}^{(1)} \otimes \dots \otimes \mathbf{u}^{(t)}). \quad (6)$$

With these definitions, the continuous dictionary \mathcal{D} in (1) corresponds to the image of the set of unit-norm rank-1 t -dimensional tensors seen through operator \mathcal{A} . A sparse decomposition of \mathbf{y} in \mathcal{D} can then be searched via the resolution of problem (2). The pertinence of this approach with respect to our initial regularized tensor-approximation problem (3) is given in Proposition 1 below.

Before stating our result, we note that any solution μ^* of (2) defines a t -dimensional tensor as $\mathcal{D}\mu^*$, where \mathcal{D} is defined from (5)-(6). Conversely, any solution \mathbf{X}^* of (3) defines a discrete measure in $\mathcal{M}(\Theta)$ with mass $\sigma_\ell(\mathbf{X}^*)$ located at the corresponding singular vector, say $(\mathbf{u}_\ell^{(1)}, \dots, \mathbf{u}_\ell^{(t)})$. Keeping this remark in mind, our result reads:

Proposition 1. *Consider problem (2) with \mathcal{D} defined from (5)-(6). Any solution \mathbf{X}^* of (3) defines a discrete solution of (2). Conversely, any discrete measure μ^* which is solution of (2) defines a solution of (3).*

In other other words, Proposition 1 states that all the solutions of (3) can be found by identifying the discrete solutions of (2). Our proof revolves around the fact that the infimum in the nuclear norm is attained [10, Prop. 3.1]. In particular, it implies that there must exist discrete measures that are solutions to the BLasso problem. We can therefore exploit the algorithmic strategies available for solving the BLasso problem to search for the solutions of (3).

3 Algorithms

A popular algorithm to solve the BLasso problem is the so-called ‘‘Frank-Wolfe’’ algorithm [4, Section 2.2]. We consider hereafter a generalized version of this procedure called ‘‘Sliding Frank-Wolfe (SFW)’’, see [6, Chapter 4, Alg. 3]. It consists of three steps repeated until convergence. First, identify the best normalized atom by maximizing correlation with the residual. Second, compute the optimal coefficients with fixed selected atoms. Finally, optimize both coefficients and atom parameters for a fixed sparsity level s . In the tensor framework, this means alternately solving problems of the following form:

$$\arg \max_{\mathbf{u}^{(i)} \in \mathbb{S}^{n_i}} \langle \mathcal{A}^* \mathbf{r}, \mathbf{u}^{(1)} \otimes \dots \otimes \mathbf{u}^{(t)} \rangle \quad (7)$$

$$\arg \min_{\mathbf{c} \in \mathbb{R}^s} \frac{1}{2} \|\mathbf{y} - \mathbf{D}\mathbf{c}\|_2^2 + \lambda \|\mathbf{c}\|_1 \quad (8)$$

$$\arg \min_{\mathbf{c} \in \mathbb{R}^s, \mathbf{u}_\ell^{(i)} \in \mathbb{S}^{n_i}} \frac{1}{2} \|\mathbf{y} - \mathcal{A} \sum_{\ell=1}^s c_\ell \mathbf{u}_\ell^{(1)} \otimes \dots \otimes \mathbf{u}_\ell^{(t)}\|_2^2 + \lambda \|\mathbf{c}\|_1 \quad (9)$$

where $\mathbf{r} \in \mathbb{R}^m$ is the residual vector after a given number of iterations, matrix \mathbf{D} is computed straightforwardly from the set of identified atoms, and \mathcal{A}^* is the adjoint of \mathcal{A} , often easy to compute in practice.

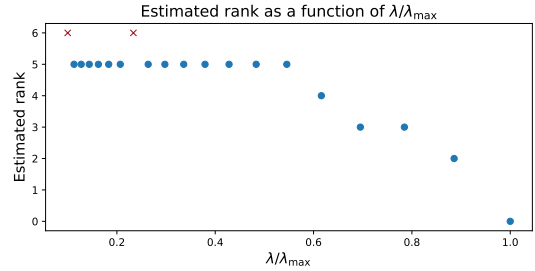


Figure 1: Tensor rank of the output of the proposed Sliding Frank-Wolfe algorithm, with λ varying from 0 to λ_{\max} as defined in [6]. Red crosses are bad estimates, that might be improved by a better implementation of steps (7) and (9) of the FW algorithm.

These three steps have very different levels of difficulty. Problem (8) is a simple Lasso problem, for which efficient solvers already exist. Problem (9) is in principle as difficult as the original tensor factorization problem with low-rank constraint. Nevertheless, it can be shown that convergence of SFW is still ensured as soon as the output of step (9) does not increase the BLasso cost function. Hence, problem (9) can be substituted by a few iterations of some local optimization procedures, see, e.g., [11].

Finally, solving (7) means finding the best rank-one approximation of a tensor $\mathcal{A}^* \mathbf{r}$, which is an NP-hard problem in general [12]. To the best of the author’s knowledge, no method is even proved to produce an estimate positively correlated with the optimal solution in general [13]. On the other hand, practical algorithms exist that empirically perform extremely well [10, 14]. Moreover the Higher-Order Singular Value Decomposition [15] provides an initialization which guarantees quasi-optimality of the fitting error in Frobenius norm. We are currently working on linking existing results on tensor best rank-one approximation with sufficient conditions for the FW algorithm to converge globally. Practically, we compute (7) using the Alternating Least Squares algorithm, initialized randomly ten times and picking the best result.

4 Experiments

Let us showcase the proposed SFW algorithm on a rank detection problem. Given a noised 3-dimensional data tensor $\mathbf{T} = \mathbf{X} + \sigma \mathbf{N}$ with \mathbf{X} of rank $k = 5$, we try to find the value of k by solving the tensor version of (3) for several values of λ . We set \mathcal{A} as the vectorization operator. The noise tensor has entries drawn from i.i.d. Gaussian distributions. Ground truth tensor \mathbf{X} is generated using factors drawn from i.i.d. Gaussian distributions with dimensions $20 \times 21 \times 22$, and ground-truth c_ℓ are drawn from the absolute value of i.i.d. Gaussian distributions with mean 0.5. The noise level σ is set to 0.005.

Figure 1 shows the results for various values of λ . We clearly observe proper rank detection, although there is a spurious error around $\lambda = 0.25\lambda_{\max}$. Moreover, we observed that the atoms selected with various λ values are very closely related, which would confirm some kind of global convergence behavior, despite our implementation behind largely improvable. Note that within the proposed framework, we may also, among others, tackle missing data imputation or directly low-rank approximation.

References

- [1] S. Foucart, H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, Springer New York, 2013. doi:10.1007/978-0-8176-4948-7.
- [2] K. Bredies, H. K. Pikkarainen, Inverse problems in spaces of measures, *ESAIM: Control, Optimisation and Calculus of Variations* 19 (1) (2013) 190–218. doi:DOI:10.1051/cocv/2011205.
- [3] V. Duval, G. Peyré, Exact Support Recovery for Sparse Spikes Deconvolution, *Foundations of Computational Mathematics* 15 (5) (2015) 1315–1355.
- [4] D. P. Bertsekas, *Nonlinear Programming*, Athena Scientific, USA, 2003.
- [5] N. Boyd, G. Schiebinger, B. Recht, The alternating descent conditional gradient method for sparse inverse problems, *SIAM Journal on Optimization* 27 (2) (2017) 616–639.
- [6] Q. Denoyelle, V. Duval, G. Peyré, E. Soubies, The sliding frank-wolfe algorithm and its application to super-resolution microscopy, *Inverse Problems* 36 (1) (2019) 014001. doi:10.1088/1361-6420/ab2a29.
- [7] E. J. Candès, B. Recht, Exact matrix completion via convex optimization, *Foundations of Computational Mathematics* 9 (6) (2009) 717–772. doi:10.1007/s10208-009-9045-5.
- [8] A. Beck, M. Teboulle, A fast iterative shrinkage thresholding algorithm for linear inverse problems, *SIAM Journal on Imaging Sciences* 2 (1) (2009) 183–202. doi:10.1137/080716542.
- [9] L.-H. Lim, P. Comon, Blind multilinear identification, *IEEE Trans. Inf. Theory* 60 (2) (2014) 1260–1280.
- [10] S. Friedland, L.-H. Lim, Nuclear norm of higher-order tensors, *Mathematics of Computation* 87 (311) (2017) 1255–1281. doi:10.1090/mcom/3239.
- [11] A. Uschmajew, Local convergence of the alternating least squares algorithm for canonical tensor approximation, *SIAM J. Matrix Analysis* 33 (2) (2012) 639–652.
- [12] S. Friedland, Best rank one approximation of real symmetric tensors can be chosen symmetric, *Frontiers of Mathematics in China* 8 (1) (2013) 19–40.
- [13] G. B. Arous, S. Mei, A. Montanari, M. Nica, The landscape of the spiked tensor model, *Communications on Pure and Applied Mathematics* 72 (11) (2019) 2282–2330.
- [14] A. P. da Silva, P. Comon, A. L. F. de Almeida, Rank-1 tensor approximation methods and application to deflation, *arXiv preprint arXiv:1508.05273* (2015).
- [15] L. D. Lathauwer, B. D. Moor, J. Vandewalle, A multilinear singular value decomposition, *SIAM J. Matrix Ana. Appl.* 21 (4) (2000) 1253–1278.