



HAL
open science

New penalized criteria for smooth non-negative tensor factorization with missing entries

Amaury Durand, François Roueff, Jean-Marc Jicquel, Nicolas Paul

► **To cite this version:**

Amaury Durand, François Roueff, Jean-Marc Jicquel, Nicolas Paul. New penalized criteria for smooth non-negative tensor factorization with missing entries. *IEEE Transactions on Signal Processing*, 2024, 72, pp.2233-2243. 10.1109/TSP.2024.3392357 . hal-03615137

HAL Id: hal-03615137

<https://hal.science/hal-03615137>

Submitted on 21 Mar 2022

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.

New penalized criteria for smooth non-negative tensor factorization with missing entries

Amaury Durand^{*†}, François Roueff^{*}, Jean-Marc Jicquel[†], and Nicolas Paul[‡]

^{*} LTCI, Telecom Paris, Institut Polytechnique de Paris. 19 Place Marguerite Perey, 91120 Palaiseau, France.

[†] EDF R&D, TREE, E36, Lab Les Renardieres. Ecuelles, 77818 Moret sur Loing, France.

[‡]EDF R&D, PRISME. 6 quai Watier, 78400 Chatou, France.

Abstract—Tensor factorization models are widely used in many applied fields such as chemometrics, psychometrics, computer vision or communication networks. Real life data collection is often subject to errors, resulting in missing data. Here we focus in understanding how this issue should be dealt with for non-negative tensor factorization. We investigate several criteria used for non-negative tensor factorization in the case where some entries are missing. In particular we show how smoothness penalties can compensate the presence of missing values in order to ensure the existence of an optimum. This lead us to propose new criteria with efficient numerical optimization algorithms. Numerical experiments are conducted to support our claims.

Index Terms—Non-negative tensor decomposition, missing values, Tensor completion, smoothness, PARAFAC, CP decomposition.

I. INTRODUCTION

The ever growing literature on multi-way data has shown the effectiveness of tensor factorization models in many domains ranging from psychometrics and chemometrics to signal processing and machine learning (see [1]–[3] and the references therein). Key strengths of these models include 1) their flexibility in taking into account prior knowledge on the data such as non-negativity, sparsity or smoothness and 2) their ability to cope with missing data. As a result, they have proven very useful in real world applications, where factors are related to practical features. On the other hand, the theoretical study of tensor factorization is still a challenging research topic. In particular, finding the true factorization of a tensor is generally NP-Hard [4] and finding an approximated factorization requires solving a non-convex optimization problem often dealt with using an iterative algorithm based on alternating minimization or on gradient descent. Since the problem is non-convex, few guarantees exist on the convergence of the optimization method used except that the objective function decreases at each iteration and that the algorithm converges to a local minimum. A question which naturally arises in this context is the existence of a global minimum. Unfortunately, this existence is not always guaranteed and discrepancies have been explored in both theoretical and experimental works especially for the popular CANDECOMP/PARAFAC (CP) decomposition (see e.g. [5], [6] and the references therein).

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

However, it is common knowledge in the inverse problem and statistical learning communities, that ill-posed problems can be handled by controlling the complexity of the solution space. This can be done, for example, using regularization [7], [8]. In the context of tensor factorization, it has been shown in [6], [9] that adding non-negativity constraints ensures existence of a global solution for the CP decomposition. This decomposition is known as non-negative tensor factorization (NTF). In this paper, we provide similar guarantees for the NTF problem in the case where some, possibly many, entries are missing. This case is of particular interest because missing data are very common in practical settings where the data collection can be subject to errors. Another interest of dealing with missing entries is that it allows to use cross validation methods to select hyperparameters such as penalty parameters or the tensor rank.

A. Related work

The literature on tensor factorization from the past decades has given rise to a wide range of algorithms among which some are adapted to missing values and/or additional constraints. Methods handling missing values usually fall into one of the three following categories : imputation, weighted least squares and probabilistic models. In the first case, missing entries are estimated at each iteration resulting in an EM-like algorithm [10], [11]. In the second case the squared error is weighted with binary weights representing missing and observed entries (see e.g. [12], [13]). In the last case, prior distributions are proposed for the factors and their parameters are estimated from the observed data [14]–[16]. The problem of missing entries is also closely related to tensor completion where, in addition to the factors, one usually also tries to estimate a full tensor which coincides with the data tensor on observed entries (see e.g [17] for a recent survey). Because of the need for efficient algorithms to deal with large amount of data, the literature on tensor factorization is dominated by algorithmic considerations, especially when missing values are taken into account.

The effect of using smoothness constraints or penalties has also been thoroughly explored both for tensor factorization and tensor completion [18]–[25]. Such approaches are generally used for numerical purposes to regularize ill-posed optimization problems and for statistical purposes to compensate overfitting by incorporating prior knowledge or assumption

on the model. This can be beneficial for the interpretation of the factors [18], [19], [25], [26] but also for the accuracy of the factorization or completion. For example, total variation constraints are widely used to deal with natural images as they are able to capture their smoothness structure [22], [27], [28]. There are two main strategies to impose smoothness on the factors. The first strategy, which is studied in this paper, consists in adding a penalty term in the loss. Usual penalties for smoothness involve the total variation norm or the L^2 norm of the second derivative for spline smoothing, [19], [22], [25], [26], [28]. The second strategy consists in representing the factors within specific lower dimensional spaces, using splines, polynomials or kernels [18]–[21], [23], [24], [29]–[31].

B. Notation

The interval of integers between i and j is denoted by $[[i, j]] := \{i, \dots, j\}$. We use bold capital letters to denote tensors and matrices and bold lowercase letters for vectors. Standard font is used for the entries of the tensors, matrices and vectors. For example, $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_R] \in \mathbb{R}^{I \times R}$ means that $\mathbf{a}_r \in \mathbb{R}^I$ is the r -th column of \mathbf{A} and the (i, r) -th entry of \mathbf{A} is denoted by $A_{i,r}$ or $a_{i,r}$. A tensor $\mathbf{X} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ is indexed by a vector of integers $\mathbf{i} = i_{1:n} := (i_1, \dots, i_N) \in \prod_{n=1}^N [[1, I_n]]$ and we write $X_{\mathbf{i}} = X_{i_1, \dots, i_N}$. We denote the outer product between vectors by \circ and the Hadamard product between tensors by \otimes . We refer to [1] for the definitions of these usual tensor operations. We recall that the Frobenius scalar product of two tensors $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ is defined as $\langle \mathbf{X}, \mathbf{Y} \rangle_F = \sum_{\mathbf{i}} X_{\mathbf{i}} Y_{\mathbf{i}}$ and we denote by $\|\cdot\|_F$ its induced norm. The total variation p -norm of a vector $\mathbf{a} \in \mathbb{R}^I$ is $\|\mathbf{a}\|_{\text{TV}, p} = \left(\sum_{i=1}^{I-1} |a_i - a_{i+1}|^p \right)^{1/p}$ for $p \in [1, \infty)$ and $\|\mathbf{a}\|_{\text{TV}, \infty} = \max_{1 \leq i \leq I-1} |a_i - a_{i+1}|$. For a given norm or semi-norm ν on \mathbb{R}^I we respectively denote $\mathbb{S}_{\nu}^+ := \{\mathbf{a} \in \mathbb{R}_+^I : \nu(\mathbf{a}) = 1\}$ and $\mathbb{B}_{\nu}^+ := \{\mathbf{a} \in \mathbb{R}_+^I : \nu(\mathbf{a}) \leq 1\}$ the unit positive sphere and the unit positive ball. We also write for any spaces $\mathcal{X}_1, \dots, \mathcal{X}_N$, $x^{(1:N)} := (x^{(1)}, \dots, x^{(N)}) \in \prod_{n=1}^N \mathcal{X}_n$. Finally, throughout this paper we will denote the N dimensional grid of indices by $\mathcal{I} = \prod_{n=1}^N [[1, I_n]]$ and, for any tensor $\mathbf{X} \in \mathbb{R}^{I_1, \dots, I_N}$ and any set $A \subset \mathbb{R}$, we denote the set of indices at which \mathbf{X} 's entries fall into A by $\{\mathbf{X} \in A\}$. Specifically, we denote $\{\mathbf{X} = x\} = \{\mathbf{X} \in \{x\}\}$ the set of indices with entries equal to x or $\{\mathbf{X} > x\} = \{\mathbf{X} \in (x, +\infty)\}$.

The remaining of this paper is organized as follows. In Section II, we describe the tensor factorization problem and introduce two losses and their related optimization algorithms. The main theoretical contributions are gathered in Section III. Numerical experiments are conducted in Section IV for comparing the optimization algorithms of Section II. Proofs are postponed to Section VI.

II. LOSSES AND OPTIMIZATION ALGORITHMS

A tensor $\mathbf{Y} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ is said to admit a non-negative tensor factorization with rank R if there exists a sequence of N factor matrices $\mathbf{A}^{(1:N)} \in \prod_{n=1}^N \mathbb{R}_+^{I_n \times R}$ such that

$$\mathbf{Y} = \sum_{r=1}^R \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(N)}, \quad (1)$$

where we recall that $\mathbf{a}_r^{(n)}$ is the r -th column of $\mathbf{A}^{(n)}$. In order to avoid scaling indeterminacy in the tensor factorization, it is common to use an equivalent formulation of (1) using normalized factors based on N norms ν_1, \dots, ν_N on $\mathbb{R}^{I_1}, \dots, \mathbb{R}^{I_N}$. Let us set, for all $n \in [[1, N]]$ and $r \in [[1, R]]$, $\lambda_r = \prod_{n=1}^n \nu_n(\mathbf{a}_r^{(n)})$ and $\tilde{\mathbf{a}}_r^{(n)} = \frac{\mathbf{a}_r^{(n)}}{\nu_n(\mathbf{a}_r^{(n)})}$ if $\nu_n(\mathbf{a}_r^{(n)}) > 0$ and any arbitrary element of $\mathbb{S}_{\nu_n}^+$ otherwise. Then, the factorization (1) can be written as

$$\mathbf{Y} = \sum_{r=1}^R \lambda_r \tilde{\mathbf{a}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(N)}. \quad (2)$$

We study the problem of approximating a tensor $\mathbf{X} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ by a finite rank tensor using parameterization (1) or (2) in the case where some entries of the tensor are missing. Let $\mathbf{W} \in \mathbb{R}_+^{I_1, \dots, I_N}$ be a tensor of weights and define

$$L_{\mathbf{W}}(\mathbf{A}) := \left\| \mathbf{W} \otimes \left(\mathbf{X} - \sum_{r=1}^R \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(N)} \right) \right\|_F^2,$$

for all $\mathbf{A} = \mathbf{A}^{(1:N)} \in \prod_{n=1}^N \mathbb{R}_+^{I_n \times R}$. We also define the normalized equivalent,

$$\tilde{L}_{\mathbf{W}}(\boldsymbol{\lambda}, \tilde{\mathbf{A}}) := \left\| \mathbf{W} \otimes \left(\mathbf{X} - \sum_{r=1}^R \lambda_r \tilde{\mathbf{a}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(N)} \right) \right\|_F^2,$$

for all $\boldsymbol{\lambda} \in \mathbb{R}_+^R$ and $\tilde{\mathbf{A}} = \tilde{\mathbf{A}}^{(1:N)} \in \prod_{n=1}^N (\mathbb{S}_{\nu_n}^+)^R$, where we view $(\mathbb{S}_{\nu_n}^+)^R$ as the set of matrices in $\mathbb{R}_+^{I_n \times R}$ whose columns are valued in $\mathbb{S}_{\nu_n}^+$. In this framework, it is implicitly assumed that the zero entries of \mathbf{W} indicate missing entries and \mathbf{W} is usually taken as a binary tensor.

We also introduce two penalties based on N semi-norms μ_1, \dots, μ_N defined on $\mathbb{R}^{I_1}, \dots, \mathbb{R}^{I_N}$ respectively. Namely, given two integers d, p with $d \geq p \geq 1$ and $\boldsymbol{\alpha} \in \mathbb{R}_+^N$, we define, for all $\boldsymbol{\lambda} \in \mathbb{R}_+^R$ and $\tilde{\mathbf{A}} = \tilde{\mathbf{A}}^{(1:N)} \in \prod_{n=1}^N (\mathbb{S}_{\nu_n}^+)^R$,

$$\tilde{\mathcal{P}}_{\boldsymbol{\alpha}}(\boldsymbol{\lambda}, \tilde{\mathbf{A}}) := \sum_{n=1}^N \alpha_n \sum_{r=1}^R \lambda_r^d \mu_n^p(\tilde{\mathbf{a}}_r^{(n)}), \quad (3)$$

where we use the notation $\mu_n^p(\mathbf{a}) = (\mu_n(\mathbf{a}))^p$. This includes the penalty of [28] by taking $d = 2$ and $\mu_n = \|\cdot\|_{\text{TV}, p}$ for $p = 1, 2$. We propose a new penalty, defined as the unnormalized equivalent of (3), which reads as

$$\mathcal{P}_{\boldsymbol{\alpha}}(\mathbf{A}) := \sum_{n=1}^N \alpha_n \sum_{r=1}^R \nu_n^{d-p}(\mathbf{a}_r^{(n)}) \mu_n^p(\mathbf{a}_r^{(n)}) \prod_{m \neq n} \nu_m^d(\mathbf{a}_r^{(m)}), \quad (4)$$

for all $\mathbf{A} = \mathbf{A}^{(1:N)} \in \prod_{n=1}^N \mathbb{R}_+^{I_n \times R}$.

Then, let $f_{\mathbf{W}, \boldsymbol{\alpha}} := L_{\mathbf{W}} + \mathcal{P}_{\boldsymbol{\alpha}}$ and $\tilde{f}_{\mathbf{W}, \boldsymbol{\alpha}} := \tilde{L}_{\mathbf{W}} + \tilde{\mathcal{P}}_{\boldsymbol{\alpha}}$ and consider the two following equivalent optimization problems

$$\begin{aligned} \min_{\mathbf{A}^{(1:N)}} f_{\mathbf{W}, \boldsymbol{\alpha}}(\mathbf{A}^{(1:N)}) \\ \text{s.t. } \forall n \in [[1, N]], \mathbf{A}^{(n)} \in \mathbb{R}_+^{I_n \times R}, \end{aligned} \quad (5)$$

and

$$\begin{aligned} \min_{\boldsymbol{\lambda}, \tilde{\mathbf{A}}^{(1:N)}} \tilde{f}_{\mathbf{W}, \boldsymbol{\alpha}}(\boldsymbol{\lambda}, \tilde{\mathbf{A}}^{(1:N)}) \\ \text{s.t. } \forall r \in [[1, R]], \lambda_r \geq 0 \text{ and } \forall n \in [[1, N]], \tilde{\mathbf{a}}_r^{(n)} \in \mathbb{S}_{\nu_n}^+. \end{aligned} \quad (6)$$

The fact that the $\tilde{\mathbf{a}}_r^{(n)}$'s are constrained individually in Problem (6) naturally leads to the Hierarchical Alternating Least Squares (HALS) optimization method (see e.g. [2], [28]). This method consists in minimizing $\tilde{f}_{\mathbf{W},\alpha}$ alternatively in the $\tilde{\mathbf{a}}_r^{(n)}$'s and its updates are recalled in Algorithm 1 where we have defined $\mathbf{X}^{(r)} := \mathbf{X} - \sum_{s \neq r} \lambda_s \tilde{\mathbf{a}}_s^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_s^{(N)}$ and $\tilde{f}_{\mathbf{W},\alpha}^{(r,n)}(\tilde{\mathbf{a}}) := \left\| \mathbf{W} \otimes (\mathbf{X}^{(r)} - \lambda_r \tilde{\mathbf{a}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(n-1)} \circ \tilde{\mathbf{a}} \circ \tilde{\mathbf{a}}_r^{(n+1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(N)}) \right\|_F^2 + \alpha_n \lambda_r \mu_n^p(\tilde{\mathbf{a}})$.

Algorithm 1 HALS algorithm

Data: \mathbf{X} , \mathbf{W} and initial values for λ , $\tilde{\mathbf{A}}^{(1)}, \dots, \tilde{\mathbf{A}}^{(N)}$

$\mathbf{E} = \mathbf{X} - \sum_{r=1}^R \lambda_r \tilde{\mathbf{a}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(N)}$

repeat

for $r = 1, \dots, R$ **do**

$\mathbf{X}^{(r)} = \mathbf{E} + \lambda_r \tilde{\mathbf{a}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(N)}$

for $n = 1, \dots, N$ **do**

$\tilde{\mathbf{a}}_r^{(n)} = \operatorname{argmin}_{\tilde{\mathbf{a}} \in \mathbb{S}_{\nu_n}^+} \tilde{f}_{\mathbf{W},\alpha}^{(r,n)}(\tilde{\mathbf{a}})$

end

$\lambda_r = \operatorname{argmin}_{\lambda \geq 0} \left\| \mathbf{W} \otimes (\mathbf{X}^{(r)} - \lambda \tilde{\mathbf{a}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(N)}) \right\|_F^2$

$\mathbf{E} = \mathbf{X}^{(r)} - \lambda_r \tilde{\mathbf{a}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(N)}$

end

until Change of value of $\tilde{f}_{\mathbf{W},\alpha}$ is sufficiently small;

return $\lambda, \tilde{\mathbf{A}}^{(1)}, \dots, \tilde{\mathbf{A}}^{(N)}$

The update for $\tilde{\mathbf{a}}_r^{(n)}$ in Algorithm 1 requires minimizing the function $\tilde{f}_{\mathbf{W},\alpha}^{(r,n)}$ on the non-negative sphere $\mathbb{S}_{\nu_n}^+$. In the case where there is no non-negativity constraint, [28] proposes to use a projected gradient method. In our case, we propose to first solve $\min_{\tilde{\mathbf{a}} \in \mathbb{R}_+^{I_n}} \tilde{f}_{\mathbf{W},\alpha}^{(r,n)}(\tilde{\mathbf{a}})$ using a gradient-based method with bound constraints such as L-BFGS-B and then normalize the result by its ν_n -norm. The update in λ_r writes as $\lambda_r = \left[\frac{\langle \mathbf{W} \otimes \mathbf{X}^{(r)}, \mathbf{W} \otimes (\tilde{\mathbf{a}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(N)}) \rangle_F}{\left\| \mathbf{W} \otimes (\tilde{\mathbf{a}}_r^{(1)} \circ \dots \circ \tilde{\mathbf{a}}_r^{(N)}) \right\|_F^2} \right]_+$.

On the other hand, the constraints in Problem (5) reduce to a bound constraint on a vectorization of \mathbf{A} . Hence, in the case where the gradients of μ_n and ν_n are available, we can follow the approach of [13] and use a gradient-based method with bound constraints such as L-BFGS-B.

In the next section, we study the existence of a global optimum for Problem (6). The existence of a global optimum for Problem (5) follows immediately as they are two different parameterizations of the same criterion. However, using the normalized version allows for stronger results, such as coercivity, which are not possible in the unnormalized version as already noted in [6].

III. EXISTENCE OF A GLOBAL OPTIMUM FOR THE WEIGHTED NTF

In this section, we consider the normalized version of the weighed NTF problem, i.e. Problem (6), where, for all $n \in \llbracket 1, N \rrbracket$, ν_n is a norm on \mathbb{R}^{I_n} and μ_n is a semi-norm on \mathbb{R}^{I_n} . To simplify the presentation, we set $\Theta := \mathbb{R}_+^R \times \prod_{n=1}^N (\mathbb{S}_{\nu_n}^+)^R$

and $\theta = (\lambda, \mathbf{A}^{(1:N)}) \in \Theta$. We investigate the existence of a minimum with minimal assumptions on the semi-norms used in the penalty term. Namely, for $n = 1, \dots, N$, we define

$$\psi_n := \sup \{ \nu_n(\mathbf{a}) : \mathbf{a} \in \mathbb{B}_{\mu_n}^+ \}, \quad (7)$$

$$\psi'_n := \sup \{ \nu_n(\mathbf{a}) : \mathbf{a} \in \mathbb{B}_{\mu_n}^+, \exists i \in \llbracket 1, I_n \rrbracket, a_i = 0 \}, \quad (8)$$

and consider the following admissibility condition for the semi-norm μ_n .

(AC) We have $\psi'_n < \psi_n = \infty$,

Condition $\psi_n = \infty$ simply says that μ_n is not a norm on the positive cone (otherwise it would be equivalent to the norm ν_n). The condition $\psi'_n < \infty$ says that μ_n behaves as a norm on the positive sub-cone that have at least one zero entry. Condition (AC) holds for semi-norms typically used in smoothness penalties, as shown by the following lemma.

Lemma III.1. Condition (AC) holds for μ_n defined as one of the following semi-norms.

(i) For all $\mathbf{a} \in \mathbb{R}^{I_n}$, $\mu_n(\mathbf{a}) = \|\mathbf{a}\|_{\text{TV},p}$ for some $p \in [1, +\infty)$.

(ii) For all $\mathbf{a} \in \mathbb{R}^{I_n}$, $\mu_n(\mathbf{a}) = \left(\int_0^1 (a''(u))^2 du \right)^{1/2}$, where $a : [0, 1] \rightarrow \mathbb{R}$ is the natural cubic spline such that, for all $i \in \llbracket 1, I_n \rrbracket$, $a(u_i) = a_i$ for some $0 < u_1 \leq \dots \leq u_{I_n} < 1$.

In the next theorem, we provide a necessary and sufficient condition for $\tilde{f}_{\mathbf{W},\alpha}$ to be coercive on Θ , which means that $\tilde{f}_{\mathbf{W},\alpha}(\theta)$ tends to $+\infty$ as the norm of θ goes to $+\infty$.

The necessary and sufficient condition relies on the following definition of cylinders.

Definition III.1. For any non-empty subset $\mathcal{N} \subset \llbracket 1, N \rrbracket$, a \mathcal{N} -cylinder is a set defined as

$$\mathcal{C}((j_n)_{n \in \mathcal{N}}) := \{ \mathbf{i} \in \mathcal{I} : i_n = j_n, \forall n \in \mathcal{N} \}, \quad (9)$$

for some $(j_n)_{n \in \mathcal{N}} \in \prod_{n \in \mathcal{N}} \llbracket 1, I_n \rrbracket$. In particular, the whole set \mathcal{I} the unique \emptyset -cylinder.

Then the following result holds.

Theorem III.2. Let $\alpha \in \mathbb{R}_+^N$ and suppose that, for all $n \in \{ \alpha > 0 \}$, μ_n satisfies (AC). Then the two following assertions are equivalent.

(i) The set $\{ \mathbf{W} = 0 \}$ contains no $\{ \alpha = 0 \}$ -cylinder.

(ii) The function $\tilde{f}_{\mathbf{W},\alpha}$ is coercive on Θ .

In this case, both $\tilde{f}_{\mathbf{W},\alpha}$ and $f_{\mathbf{W},\alpha}$ admit global minima on Θ and on $\prod_{n=1}^N \mathbb{R}_+^{I_n \times R}$, respectively.

We assume that μ_n satisfy (AC) only for $n \in \{ \alpha > 0 \}$ because, for $n \in \{ \alpha = 0 \}$, μ_n vanishes in the penalty (3). Also note that Assertion (i) means that no $\{ \alpha = 0 \}$ -cylinder is missing and is not very restrictive. For example, if all modes are penalized, i.e. $\{ \alpha = 0 \} = \emptyset$, then Assertion (i) does not hold if and only if all entries are missing. In the experiments, we study the case of color image completion where $N = 3$ and the first two modes correspond to pixels and the third corresponds to the color channel. In this case, we penalize only the first two modes, i.e. $\{ \alpha = 0 \} = \{ 3 \}$. This means that Assertion (i) does not hold if and only if an entire color channel is missing.

IV. EXPERIMENTAL RESULTS

In this section, we compare the two optimization problems (5) and (6) with the penalties defined in (3) and (4) respectively where we take $p = d = 2$. The first optimization problem is solved using the `scipy` implementation of L-BFGS-B. The second optimization problem is solved using Algorithm 1. We propose two experiments. In the first one, we try to recover factors from an incomplete and noisy observation of a tensor of the type (2) with $N = 3$. In the second one, we apply the algorithms for color image completion. In both experiments and each of the algorithm, we take $\alpha = [\alpha, \alpha, 0]^T$ where $\alpha > 0$, and $\nu_n = \|\cdot\|_2$ for all n . The maximum number of iterations is set to 10^4 and the iterations are stopped if the relative improvement of the loss is lower than 10^{-6} . One iteration of the HALS algorithm consists of the steps in the **repeat-until** box of Algorithm 1 and one iteration of the L-BFGS-B method consists in one update of the gradient descent. In this case, during one iteration, all the factors are updated. The computational performances algorithms are compared using the average computing time per iteration denoted by TPI. All the experiments are run on a Linux Workstation with 40 Intel Xeon E5-2630 v4 2.20 GHz processors.

A. Factor estimation on toy data

In this experiment, we construct a tensor $\mathbf{X} = \mathbf{Y} + \sigma \mathbf{E} \in \mathbb{R}^{I \times I \times I}$ where \mathbf{Y} is as in (2) with $I = 50$, $N = 3$ and $R = 5$ and \mathbf{E} has a standard normal entries. The factors $\mathbf{a}_r^{(3)}$ are sampled uniformly on $[0, 1]$ and the factors $\mathbf{a}_r^{(1)}$ and $\mathbf{a}_r^{(2)}$ are constructed by taking non-negative random linear combination of 7 B-Spline functions of order 4. With this construction, the factors are non-negative and smooth for $n = 1, 2$ and we allow the $\mathbf{a}_r^{(2)}$'s to vanish on some intervals (see the first column of Figure 3). The standard deviation $\sigma > 0$ is computed as in [12], [13], i.e. $\sigma = (100/\nu - 1)^{-1/2} \frac{\|\mathbf{Y}\|_F}{\|\mathbf{E}\|_F}$, and we take $\nu = 10$. Since the true factors are generated from B-splines, we take μ_n as in (ii) of Lemma III.1. We generated 25%, 50% and 70% percent of missing data which are drawn randomly and uniformly on the grid $\llbracket 1, I \rrbracket^3$. For the unnormalized problem, we use an SVD-based initialization and use its normalized equivalent for the normalized problem. To evaluate the output $\hat{\mathbf{Y}} = \sum_{r=1}^R \hat{\lambda}_r \hat{\mathbf{a}}_r^{(1)} \circ \hat{\mathbf{a}}_r^{(2)} \circ \hat{\mathbf{a}}_r^{(3)}$, we use the normalized mean square error, $\text{NMSE} = \frac{\|\mathbf{Y} - \hat{\mathbf{Y}}\|_F^2}{\|\mathbf{Y}\|_F^2}$ and the similarity score, $\text{SIM} = \max_{\sigma \in \mathcal{S}_R} \frac{1}{R} \sum_{r=1}^R \prod_{n=1}^N \langle \tilde{\mathbf{a}}_r^{(n)}, \hat{\mathbf{a}}_{\sigma(r)}^{(n)} \rangle$, where \mathcal{S}_R denotes the set of permutations of $\llbracket 1, R \rrbracket$. Note that, since the $\tilde{\mathbf{a}}_r^{(n)}$'s and $\hat{\mathbf{a}}_r^{(n)}$'s have unit 2-norm, we have $\text{SIM} \in [0, 1]$. A lower NSME is interpreted as a better prediction of Y 's entries while a higher SIM is interpreted as a better estimation of its factors. An example of reconstruction is shown in Figure 3 where we observe that the gradient method is able to reconstruct the factors even when the proportion of missing data is high. In the remaining of this section, we discuss the choice of α and compare the methods with various values of I .

1) *Data-driven selection of α* : We propose to evaluate the selection of α using a 5-fold cross validation criterion (CV), which amounts to arbitrarily introduce additional missing

values in the objective function and evaluate their prediction errors. More precisely, we generate randomly 5 binary masks $(\mathbf{W}_k)_{k \in \llbracket 1, 5 \rrbracket}$ such that the sets $(\{\mathbf{W}_k > 0\})_{k \in \llbracket 1, 5 \rrbracket}$ create a partition of \mathcal{I} . For each fold $k \in \llbracket 1, 5 \rrbracket$ and each value of $\alpha \in \{0, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10\}$, we minimize $f_{\mathbf{W}_k \otimes \mathbf{W}_k, \alpha}$ (or $\tilde{f}_{\mathbf{W}_k \otimes \mathbf{W}_k, \alpha}$) and the cross validation score is computed using $f_{\mathbf{W}_k \otimes (\mathbf{1} - \mathbf{W}_k), \alpha}$ (or $\tilde{f}_{\mathbf{W}_k \otimes (\mathbf{1} - \mathbf{W}_k), \alpha}$), where $\mathbf{1}$ is the tensor with all entries equal to 1. We compare the cross validation scores with an oracle selection of α given by the NMSE and SIM scores which use the ground truth. The results are gathered in Figure 1 where we observe that the selection with the cross validation score seems consistent with the oracle based on NMSE, which indicates that CV is a suitable parameter selection method in our context. Note that the oracle selection using the SIM score can give different optimal values for α . This can be explained by the fact that the SIM score is more sensitive to small differences between the estimated factors and the true factors.




Fig. 1. CV selection of α (bottom), compared to the SIM (middle) and NMSE (top) criteria.

2) *Comparison with various dimensions*: We now compare the proposed methods for various values of I . We use the same ground truth represented by the first column of Figure 3 where we interpolated the factors to achieve higher values of I . The comparison is made in a best case scenario where, for each value of I , each proportion of missing data and each model, we use the oracle selection of $\alpha \in \{0, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 1, 10\}$ based on the NMSE score. The NMSE and SIM scores and the TPI are gathered in Figure 2. The two methods give very similar results for the NMSE and SIM whose values indicate almost perfect reconstruction in all cases. We observe that the reconstruction tends to be better for large values of I which is expected since the difficulty of the problem decreases as I increases (see

[13]). The advantage of the gradient-based method is, however, highlighted by the TPI, especially when the dimension increases. Note that we observed that the value of α does not affect much the TPI and therefore the comparison displayed in Figure 2 is representative of any value of α .




Fig. 2. Best case scenario for different dimensions.

B. Color image completion

In this experiment, we apply both methods to 3 color images (barbara, baboon and giant) of size $256 \times 256 \times 3$. Missing data are generated in two ways. In the first case, we remove all the color channels for 80% of the pixels selected randomly and uniformly. In the last case, we remove all the color channels for a mask obtained by scribbling the image. Note that, in all of these cases, Assertion (i) of Theorem III.2 holds. We use the quadratic variation penalty as in [28], i.e. $\mu_n = \|\cdot\|_{TV,2}$. The quality of completion, is evaluated by the peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) as in [28]. High PSNR and SSIM indicate a good completion. For both methods, we take a fix rank $R = 50$ and use a random initialization. For each image, each mask and each model, we use the oracle selection of α based on the SSIM. As observed in Figure 4, there is an optimal value of α which should neither be too small nor too large. This value is not necessarily the same for both algorithms. We also observe that, near the optimal value of α , the gradient-based method gives a better SSIM score than the HALS method. This is also highlighted by the completed images displayed in Figure 5. Note also that the gradient-based method has much lower TPI than the HALS one.

V. CONCLUSION

In this contribution, we extended previous results on the existence of a global minimum for the non-negative tensor factorization problem to the case where some entries are missing. We showed that, under non-restrictive assumptions on the observed entries, adding a penalty to the quadratic loss ensures the existence of a global minimum. We proposed two

formulations of the problem: a normalized one, which is solved using a HALS algorithm, and an unnormalized one, which is solved using a gradient-based method. The experimental study illustrates the advantages of the gradient-based method in computing time and in reconstruction error.

VI. PROOFS

A. Preliminary results

In this section, we provide preliminary results which are necessary to prove Theorem III.2. To simplify the notation, let us define, for all $\theta = (\lambda, \mathbf{A}^{(1:N)}) \in \Theta := \mathbb{R}_+^R \times \prod_{n=1}^N (\mathbb{S}_{\nu_n}^+)^R$,

$$g_{\mathbf{W}}(\theta) := \left\| \mathbf{W} \otimes \left(\sum_{r=1}^R \lambda_r \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(N)} \right) \right\|_F, \quad (10)$$

and, for all $\mathbf{a}^{(1:N)} \in \prod_{n=1}^N \mathbb{S}_{\nu_n}^+$,

$$h_{\mathbf{W}}(\mathbf{a}^{(1:N)}) := \left\| \mathbf{W} \otimes \left(\mathbf{a}^{(1)} \circ \dots \circ \mathbf{a}^{(N)} \right) \right\|_F^2. \quad (11)$$

Then the following proposition holds.

Proposition VI.1. *Let $\mathcal{A} = \prod_{n=1}^N \mathcal{A}_n^R$ be such that for all $n \in \llbracket 1, N \rrbracket$, $\mathcal{A}_n \subset \mathbb{S}_{\nu_n}^+$. Denote by $\overline{\mathcal{A}_n}$ the closure of \mathcal{A}_n in $\mathbb{S}_{\nu_n}^+$. Then the three following assertions are equivalent.*

- (i) *The function $\tilde{L}_{\mathbf{W}}$ is coercive on $\mathbb{R}_+^R \times \mathcal{A}$.*
- (ii) *The function $g_{\mathbf{W}}$ is coercive on $\mathbb{R}_+^R \times \mathcal{A}$.*
- (iii) *For all $\mathbf{a}^{(1:N)} \in \prod_{n=1}^N \overline{\mathcal{A}_n}$, there exists $\mathbf{i} \in \{\mathbf{W} > 0\}$ such that $\prod_{n=1}^N \mathbf{a}_{i_n}^{(n)} > 0$.*

Proof. The equivalence between (i) and (ii) is a direct consequence of the two triangular inequalities. Now, note that, by continuity of $h_{\mathbf{W}}$, Assertion (iii) is equivalent to

$$\inf \left\{ h_{\mathbf{W}}(\mathbf{a}^{(1:N)}) : \mathbf{a}^{(1:N)} \in \prod_{n=1}^N \mathcal{A}_n \right\} > 0. \quad (12)$$

Hence, it remains to prove that (ii) and (12) are equivalent.

Proof of (12) \Rightarrow (ii). Assume that Condition (12) holds and let $(\theta(m))_{m \in \mathbb{N}} \in (\mathbb{R}_+^R \times \mathcal{A})^{\mathbb{N}}$ with $\|\theta(m)\|_2 \xrightarrow{m \rightarrow +\infty} +\infty$ and, for all $m \in \mathbb{N}$, $\theta(m) = (\lambda(m), \mathbf{A}^{(1:N)}(m))$. Then, since \mathcal{A} is bounded, we must have $\|\lambda(m)\|_2 \xrightarrow{m \rightarrow +\infty} +\infty$ and, using the fact that the entries of $\theta(m)$ are all non-negative, we get

$$\begin{aligned} (g_{\mathbf{W}}(\theta(m)))^2 &\geq \sum_{r=1}^R (\lambda_r(m))^2 h_{\mathbf{W}}(\mathbf{a}_r^{(1:N)}(m)) \\ &\geq c_{\inf} \|\lambda(m)\|_2^2, \end{aligned}$$

where c_{\inf} is the inf in (12). Hence, under Condition (12), $g_{\mathbf{W}}(\theta(m))$ diverges to $+\infty$ as $m \rightarrow +\infty$ and Assertion (ii) follows.

Proof of (ii) \Rightarrow (12). Assume that Condition (12) does not hold and let us show that $\tilde{g}_{\mathbf{W}}$ is not coercive on $\mathbb{R}_+^R \times \mathcal{A}$ by constructing a sequence $(\theta(m))_{m \in \mathbb{N}} \in (\mathbb{R}_+^R \times \mathcal{A})^{\mathbb{N}}$ such that $\|\theta(m)\| \xrightarrow{m \rightarrow +\infty} +\infty$ and $g_{\mathbf{W}}(\theta(m)) \xrightarrow{m \rightarrow +\infty} 0$. Let us set, for all $m \in \mathbb{N}$, $\theta(m) := (\lambda(m), \mathbf{A}^{(1:N)}(m))$ with $\lambda(m) = [m, 0, \dots, 0]$ and, for all $n \in \mathbb{N}$, $\mathbf{A}^{(n)}(m) = [\mathbf{a}^{(n)}(m), \mathbf{a}_2^{(n)}, \dots, \mathbf{a}_R^{(n)}]$, where $\mathbf{a}_2^{(n)}, \dots, \mathbf{a}_R^{(n)}$ are any elements of \mathcal{A}_n and $\mathbf{a}^{(1:N)}(m) \in \prod_{n=1}^N \mathcal{A}_n$ is such that




Fig. 3. Estimated factors with gradient-based method and $\alpha = 0.1$. Rows represent the modes (i.e. n) and colors represent the components (i.e. r).

$h_{\mathbf{W}}(\mathbf{a}^{(1:N)}(m)) \leq 2^{-m}$. Note that such $\mathbf{a}^{(1:N)}(m)$ exists because Condition (12) does not hold. With this construction, we get $\|\boldsymbol{\theta}(m)\| \xrightarrow{m \rightarrow +\infty} +\infty$ and $(g_{\mathbf{W}}(\boldsymbol{\theta}(m)))^2 = m^2 h_{\mathbf{W}}(\mathbf{a}^{(1:N)}(m)) \leq m^2 2^{-m} \xrightarrow{m \rightarrow +\infty} 0$, which concludes the proof. \square

For the second result, we consider specific subsets

$$\mathcal{A}_n(\boldsymbol{\rho}, C) := \{\mathbf{a} \in \mathbb{S}_{\nu_n}^+ : \mu_n(\mathbf{a}) \leq \rho_n C\}, \quad (13)$$

$$\Theta(\boldsymbol{\rho}, C) := \mathbb{R}_+^R \times \prod_{n=1}^N (\mathcal{A}_n(\boldsymbol{\rho}, C))^R, \quad (14)$$

where $C \geq 0$ and $\boldsymbol{\rho} \in (0, +\infty]^N$. Note that, for all n such that $\rho_n = +\infty$, we have $\mathcal{A}_n(\boldsymbol{\rho}, C) = \mathbb{S}_{\nu_n}^+$. Let us define

$$C_{\mathbf{W}}(\boldsymbol{\rho}) := \inf \left\{ C \geq 0 : \tilde{L}_{\mathbf{W}} \text{ is not coercive on } \Theta(\boldsymbol{\rho}, C) \right\}. \quad (15)$$

The goal of the remaining of this section is to derive an explicit formulation of $C_{\mathbf{W}}(\boldsymbol{\rho})$. To this end, we need to introduce the following sets and constants. For any $n \in \llbracket 1, N \rrbracket$ and $\mathcal{J} \subset \llbracket 1, I_n \rrbracket$, define

$$\begin{aligned} \mathcal{A}_n(\mathcal{J}) &:= \{\mathbf{a} \in \mathbb{S}_{\nu_n}^+ : \forall j \in \mathcal{J}, a_j = 0\}, \\ m_n(\mathcal{J}) &:= \inf \{\mu_n(\mathbf{a}) : \mathbf{a} \in \mathcal{A}_n(\mathcal{J})\}, \end{aligned}$$

with the convention $\inf(\emptyset) = +\infty$. Moreover, for any $\mathcal{I}' \subset \mathcal{I}$ and $n \in \llbracket 1, N \rrbracket$, we define the projection of \mathcal{I}' onto the n -th coordinate as $\pi_n(\mathcal{I}') := \{i_n : \mathbf{i} \in \mathcal{I}'\}$ and denote by

$$\mathcal{J}(\mathcal{I}') := \left\{ \mathcal{J}_{1:N} : \begin{array}{l} \forall n \in \llbracket 1, N \rrbracket, \mathcal{J}_n \subseteq \pi_n(\mathcal{I}'), \\ \forall \mathbf{i} \in \mathcal{I}', \exists n \in \llbracket 1, N \rrbracket, i_n \in \mathcal{J}_n \end{array} \right\}$$

the class of sequences of sets $\mathcal{J}_{1:N} = (\mathcal{J}_1, \dots, \mathcal{J}_N)$, where each \mathcal{J}_n is a subset of all the n -th entries of the vector indices in \mathcal{I}' and such that each vector index \mathbf{i} in \mathcal{I}' has at least one entry, say the n -th, present in the corresponding \mathcal{J}_n .

Finally, we define, for all $\boldsymbol{\rho} \in (0, +\infty]^N$ and $\mathcal{I}' \subset \mathcal{I}$,

$$C(\boldsymbol{\rho}, \mathcal{I}') := \min \{c(\boldsymbol{\rho}, \mathcal{J}_{1:N}) : \mathcal{J}_{1:N} \in \mathcal{J}(\mathcal{I}')\}, \quad (16)$$

where

$$c(\boldsymbol{\rho}, \mathcal{J}_{1:N}) := \begin{cases} \infty & \text{if } \exists n, \mathcal{A}_n(\mathcal{J}_n) = \emptyset, \\ \max_{n \in \llbracket 1, N \rrbracket} \frac{m_n(\mathcal{J}_n)}{\rho_n} & \text{otherwise.} \end{cases}$$

The first case in the definition of $c(\boldsymbol{\rho}, \mathcal{J}_{1:N})$ amounts to use the convention $\inf(\emptyset)/\rho = +\infty$ for any $\rho \in (0, \infty]$. We have the following results.

Lemma VI.2. *Let $n \in \llbracket 1, N \rrbracket$, and $\mathcal{J} \subset \llbracket 1, I_n \rrbracket$, the following assertions hold.*

- (i) *We have $m_n(\mathcal{J}) = +\infty$ if and only if $\mathcal{J} = \llbracket 1, I_n \rrbracket$.*




Fig. 4. Evolution of the SSIM with α . Rows represent the three images and columns represent the types of mask.

(ii) If μ_n satisfies (AC), we have $m_n(\mathcal{J}) = 0$ if and only if $\mathcal{J} = \emptyset$.

Proof. Assertion (i) follows from the equivalence between $m_n(\mathcal{J}) = +\infty$ and $\mathcal{A}_n(\mathcal{J}) = \emptyset$, which is itself equivalent to $\mathcal{J} = \llbracket 1, I_n \rrbracket$. For Assertion (ii), recall that $\psi_n = \infty$ means that μ_n is a semi-norm which is not a norm on the positive cone, which implies $m_n(\emptyset) = 0$ thus showing the “if” implication. For the “only if” implication, note that $m_n(\mathcal{J}) > 0$ if and only if

$$\sup \{ \nu_n(\mathbf{a}) : \mathbf{a} \in \mathbb{B}_{\mu_n}^+(\mathbf{a}), \forall j \in \mathcal{J}, \mathbf{a}_j = 0 \} < \infty.$$

Hence, if $\psi'_n < \infty$, this happens whenever $\mathcal{J} \neq \emptyset$. \square

Lemma VI.3. Let $\mathcal{I}' \subset \mathcal{I}$ and $\rho \in (0, +\infty]^N$. Then the following assertions hold.

- (i) We have $C(\rho, \mathcal{I}') = +\infty$ if and only if $\mathcal{I}' = \mathcal{I}$.
- (ii) If μ_n satisfies (AC) for all $n \in \{\rho < +\infty\}$, then $C(\rho, \mathcal{I}') > 0$ if and only if $\mathcal{I} \setminus \mathcal{I}'$ contains no $\{\rho = +\infty\}$ -cylinder.

Proof. Using Assertion (i) of Lemma VI.2, we get that

$$\begin{aligned} C(\rho, \mathcal{I}') < +\infty \\ \Leftrightarrow \exists \mathcal{J}_{1:N} \in \mathcal{J}(\mathcal{I}'), \forall n \in \llbracket 1, N \rrbracket, m_n(\mathcal{J}_n) < +\infty \\ \Leftrightarrow \exists \mathcal{J}_{1:N} \in \mathcal{J}(\mathcal{I}'), \forall n \in \llbracket 1, N \rrbracket, \mathcal{J}_n \neq \llbracket 1, I_n \rrbracket \\ \Leftrightarrow \exists \mathcal{J}_{1:N} \in \mathcal{J}(\mathcal{I}'), \exists \mathbf{i} \in \mathcal{I}, \forall n \in \llbracket 1, N \rrbracket, i_n \notin \mathcal{J}_n \\ \Leftrightarrow \mathcal{J}(\mathcal{I}') \cap (\mathcal{J}(\mathcal{I}))^c \neq \emptyset. \end{aligned}$$

Hence, to conclude the proof of Assertion (i), we need to show that $\mathcal{J}(\mathcal{I}') \cap (\mathcal{J}(\mathcal{I}))^c \neq \emptyset$ if and only if $\mathcal{I}' \neq \mathcal{I}$. The “only if” implication is straightforward by contraposition. For the “if” implication, assume that $\mathcal{I}' \neq \mathcal{I}$ and take $\mathbf{i} \in \mathcal{I} \setminus \mathcal{I}'$.

Then it is easily seen that, taking $\mathcal{J}_n = \pi_n(\mathcal{I}') \setminus \{i_n\}$ for all $n \in \llbracket 1, N \rrbracket$, we get $\mathcal{J}_{1:N} \in \mathcal{J}(\mathcal{I}') \cap (\mathcal{J}(\mathcal{I}))^c$ which is therefore non-empty. This concludes the proof of Assertion (i).

For Assertion (ii), note that $C(\rho, \mathcal{I}') = 0$ is equivalent to

$$\exists \mathcal{J}_{1:N} \in \mathcal{J}(\mathcal{I}'), \forall n \in \llbracket 1, N \rrbracket, \rho_n^{-1} m_n(\mathcal{J}_n) = 0,$$

which, by Lemma VI.2 and (AC), is in turn equivalent to

$$\exists \mathcal{J}_{1:N} \in \mathcal{J}(\mathcal{I}'), \begin{cases} \forall n \in \{\rho = +\infty\}, \mathcal{J}_n \neq \llbracket 1, I_n \rrbracket, \\ \forall n \in \{\rho < +\infty\}, \mathcal{J}_n = \emptyset. \end{cases} \quad (17)$$

We now prove that (17) holds if and only if there exists $(j_n)_{n \in \{\rho = +\infty\}} \in \prod_{n \in \{\rho = +\infty\}} \llbracket 1, I_n \rrbracket$ such that $\mathcal{C}((j_n)_{n \in \{\rho = +\infty\}}) \subset \mathcal{I} \setminus \mathcal{I}'$. First, assume that (17) holds and take $j_n \in \llbracket 1, I_n \rrbracket \setminus \{\mathcal{J}_n\}$ for all $n \in \{\rho = +\infty\}$. Assume that $\mathcal{C}((j_n)_{n \in \{\rho = +\infty\}}) \not\subset \mathcal{I} \setminus \mathcal{I}'$. This means that there exists $\mathbf{i} \in \mathcal{C}((j_n)_{n \in \{\rho = +\infty\}}) \cap \mathcal{I}'$. Then, by definition of $\mathcal{J}(\mathcal{I}')$, there exists $n \in \{\rho = +\infty\}$ such that $j_n = i_n \in \mathcal{J}_n$ which contradicts the fact that $j_n \in \llbracket 1, I_n \rrbracket \setminus \{\mathcal{J}_n\}$. Hence $\mathcal{C}((j_n)_{n \in \{\rho = +\infty\}}) \subset \mathcal{I} \setminus \mathcal{I}'$ thus proving the “only if” implication. For the “if” implication, assume that there exists $(j_n)_{n \in \{\rho = +\infty\}} \in \prod_{n \in \{\rho = +\infty\}} \llbracket 1, I_n \rrbracket$ such that $\mathcal{C}((j_n)_{n \in \{\rho = +\infty\}}) \subset \mathcal{I} \setminus \mathcal{I}'$. Then, we get (17) by taking $\mathcal{J}_n = \emptyset$ for $n \in \{\rho < +\infty\}$ and $\mathcal{J}_n = \llbracket 1, I_n \rrbracket \setminus \{j_n\}$ for $n \in \{\rho = +\infty\}$. \square

Lemma VI.4. For all $\rho \in (0, +\infty]^N$, we have

$$C_{\mathbf{W}}(\rho) = C(\rho, \{\mathbf{W} > 0\}), \quad (18)$$


Proof. The proof relies on the identity

$$\begin{aligned} c(\rho, \mathcal{J}_{1:N}) = \\ \inf \left\{ \max_{n \in \llbracket 1, N \rrbracket} \rho_n^{-1} \mu_n(\mathbf{a}^{(n)}) : \forall n, \mathbf{a}^{(n)} \in \mathcal{A}_n(\mathcal{J}_n) \right\}. \quad (19) \end{aligned}$$

To prove (19), first note that, if there exists $n \in \llbracket 1, I_n \rrbracket$, such that $\mathcal{J}_n = \llbracket 1, I_n \rrbracket$, then the two terms of (19) are equal to $+\infty$. We now assume that, for all $n \in \llbracket 1, I_n \rrbracket$, $\mathcal{J}_n \neq \llbracket 1, I_n \rrbracket$. Then the inequality \leq of (19) is a straightforward consequence of the definition of $c(\rho, \mathcal{J}_{1:N})$. Let us now show that there exists $\mathbf{a}^{(1:N)} \in \prod_{n=1}^N \mathcal{A}_n(\mathcal{J}_n)$ such that $c(\rho, \mathcal{J}_{1:N}) = \max_{n \in \llbracket 1, N \rrbracket} \rho_n^{-1} \mu_n(\mathbf{a}^{(n)})$. It suffices to take, for $n \in \llbracket 1, N \rrbracket$, $\mathbf{a}^{(n)} \in \mathcal{A}_n(\mathcal{J}_n)$ such that $\mu_n(\mathbf{a}^{(n)}) = m_n(\mathcal{J}_n)$. Such $\mathbf{a}^{(n)}$ exists because $\mathcal{A}_n(\mathcal{J}_n)$ is compact and μ_n is continuous. This concludes the proof of (19).

Now, to prove (18), we show that, for all $C \geq 0$, $\tilde{L}_{\mathbf{W}}$ is not coercive on $\Theta(\rho, C)$ if and only if $C \geq C(\rho, \{\mathbf{W} > 0\})$. For the “only if” implication, assume that C is such that $\tilde{L}_{\mathbf{W}}$ is not coercive on $\Theta(\rho, C)$. Then, from Proposition VI.1 and closeness of the $\mathcal{A}_n(\rho, C)$ ’s defined in (13), we get that there exists $\mathbf{a}^{(1:N)} \in \prod_{n=1}^N \mathcal{A}_n(\rho, C)$ such that, for all $\mathbf{i} \in \{\mathbf{W} > 0\}$, there exists $n \in \llbracket 1, N \rrbracket$ such that $a_{i_n}^{(n)} = 0$. Using this observation, we construct $\mathcal{J}_{1:N} \in \mathcal{J}(\{\mathbf{W} > 0\})$ by the following procedure. Start with $\mathcal{J}_n = \emptyset$ for all $n \in \llbracket 1, N \rrbracket$ and then, for each $\mathbf{i} \in \{\mathbf{W} > 0\}$, select one of the n ’s such that $a_{i_n}^{(n)} = 0$ and put i_n in \mathcal{J}_n . With this construction, we have $\mathbf{a}^{(1:N)} \in \prod_{n=1}^N \mathcal{A}_n(\mathcal{J}_n) \cap \mathcal{A}_n(\rho, C)$ and therefore

$$C \geq \max_{n \in \llbracket 1, N \rrbracket} \rho_n^{-1} \mu_n(\mathbf{a}^{(n)}) \geq c(\rho, \mathcal{J}_{1:N}) \geq C(\rho, \{\mathbf{W} > 0\}),$$



(a) Uniformly missing pixels.

(b) Masked pixels.

Fig. 5. Image completion comparison with α giving the best SSIM.

where the first inequality comes from (13), the second from (19) and the last from (16).

For the “if” implication, let us take $C \geq C(\rho, \{\mathbf{W} > 0\})$ and show that Assertion (iii) of Proposition VI.1 does not hold. Note that we can take \mathbf{W} such that $\{\mathbf{W} > 0\} \neq \mathcal{I}$ because the other case is straightforward. Then, by definition of $C(\rho, \{\mathbf{W} > 0\})$, we know that there exists $\mathcal{J}_{1:N} \in \mathcal{J}(\{\mathbf{W} > 0\})$ such that $C(\rho, \{\mathbf{W} > 0\}) = c(\rho, \mathcal{J}_{1:N})$. Moreover, since $\{\mathbf{W} > 0\} \neq \mathcal{I}$, Assertion (i) of Lemma VI.2 gives that $c(\rho, \mathcal{J}_{1:N}) < +\infty$ and therefore we are in the case where the infimum in (19) is reached. Hence there exists $\mathbf{a}^{(1:N)} \in \prod_{n=1}^N \mathcal{A}_n(\mathcal{J}_n)$ such that $\max_{n \in \llbracket 1, N \rrbracket} \rho_n^{-1} \mu_n(\mathbf{a}^{(n)}) = C(\rho, \{\mathbf{W} > 0\})$. This gives that, for all $n \in \llbracket 1, N \rrbracket$, $\mu_n(\mathbf{a}^{(n)}) \leq \rho_n C(\rho, \{\mathbf{W} > 0\}) \leq \rho_n C$ and therefore $\mathbf{a}^{(n)} \in \mathcal{A}_n(\rho, C)$, as defined in (13). On the other hand, $\mathcal{J}_{1:N} \in \mathcal{J}(\{\mathbf{W} > 0\})$ means that, for all $\mathbf{i} \in \{\mathbf{W} > 0\}$, there exists $n \in \llbracket 1, N \rrbracket$ such that $i_n \in \mathcal{J}_n$ and we get that $\mathbf{a}_{i_n}^{(n)} = 0$ because $\mathbf{a}^{(n)} \in \mathcal{A}_n(\mathcal{J}_n)$. Hence Assertion (iii) of Proposition VI.1 does not hold and the proof is concluded. \square

B. Proofs of Theorem III.2 and Lemma III.1

Proof of Theorem III.2. In the proof, we use the functions $g_{\mathbf{W}}$ and $h_{\mathbf{W}}$ defined respectively in (10) and (11). Note that, as in Proposition VI.1, we have that $f_{\mathbf{W}, \alpha}$ is coercive on Θ if and only if $g_{\mathbf{W}}^2 + \tilde{\mathcal{P}}_{\alpha}$ is coercive on Θ .

Proof of (i) \Rightarrow (ii). Assume that (i) holds and let us show that $g_{\mathbf{W}}^2 + \tilde{\mathcal{P}}_{\alpha}$ is coercive on Θ . Take $\theta = (\lambda, \mathbf{A}^{(1:N)}) \in \Theta$ and let $s \in \llbracket 1, R \rrbracket$ be such that $\lambda_s = \max_{r \in \llbracket 1, R \rrbracket} \lambda_r = \|\lambda\|_{\infty}$.

Then, since all the entries of θ are non-negative, we have

$$\begin{aligned} & (g_{\mathbf{W}}(\theta))^2 + \tilde{\mathcal{P}}_{\alpha}(\theta) \\ & \geq (\lambda_s^2 \wedge \lambda_s^d) \left(h_{\mathbf{W}}(\mathbf{a}_s^{(1:N)}) + \sum_{n=1}^N \alpha_n \mu_n^p(\mathbf{a}_s^{(n)}) \right) \\ & \geq \left[\|\lambda\|_{\infty}^{d \wedge 2} - 1 \right]_{+} \eta, \end{aligned}$$

where

$$\eta := \inf_{\mathbf{a}^{(1:N)} \in \prod_{n=1}^N \mathcal{S}_{\nu_n}^+} \left(h_{\mathbf{W}}(\mathbf{a}^{(1:N)}) + \sum_{n=1}^N \alpha_n \mu_n^p(\mathbf{a}^{(n)}) \right).$$

Hence, to prove that $g_{\mathbf{W}}^2 + \tilde{\mathcal{P}}_{\alpha}$ is coercive on Θ , it suffices to prove that $\eta > 0$. Now let us set, for all $n \in \llbracket 1, N \rrbracket$, $\rho_n = \alpha_n^{-p}$ with the convention that $0^{-p} = +\infty$. Then, Lemma VI.4 and Assertion (ii) of Lemma VI.3 give that $C_{\mathbf{W}}(\rho) > 0$. Take $0 < C < C_{\mathbf{W}}(\rho)$ and set $\mathcal{A} := \prod_{n=1}^N \mathcal{A}_n(\rho, C)$ with $\mathcal{A}_n(\rho, C)$ defined as in (13). Then, by definition of $C_{\mathbf{W}}(\rho)$ in (15) and Relation (12), we have $\inf_{\mathcal{A}} h_{\mathbf{W}} > 0$. Moreover, for all $\mathbf{a}^{(1:N)} \in \mathcal{A}^c$, there exists $k \in \llbracket 1, N \rrbracket$ such that $\mathbf{a}^{(k)} \notin \mathcal{A}_k(\rho, C)$, i.e. $\alpha_k \mu_k^p(\mathbf{a}_r^{(k)}) > C^p$ and $\sum_{n=1}^N \alpha_n \mu_n^p(\mathbf{a}^{(n)}) \geq \alpha_k \mu_k^p(\mathbf{a}_r^{(k)}) > C^p$. Hence, we get $h_{\mathbf{W}} + \sum_{n=1}^N \alpha_n \mu_n^p \geq (\inf_{\mathcal{A}} h_{\mathbf{W}}) \mathbb{1}_{\mathcal{A}} + C^p \mathbb{1}_{\mathcal{A}^c}$ so that $\eta \geq (\inf_{\mathcal{A}} h_{\mathbf{W}}) \wedge C^p > 0$, thus concluding the proof.

Proof of (ii) \Rightarrow (i). Let us assume that (i) does not hold and show that $g_{\mathbf{W}}^2 + \tilde{\mathcal{P}}_{\alpha}$ is not coercive on Θ by constructing a sequence $(\theta(m))_{m \in \mathbb{N}} \in \Theta^{\mathbb{N}}$ such that $\theta(m) \xrightarrow{m \rightarrow +\infty} +\infty$ and $g_{\mathbf{W}}^2(\theta(m)) + \tilde{\mathcal{P}}_{\alpha}(\theta(m)) \xrightarrow{m \rightarrow +\infty} 0$. We set, for all $n \in \llbracket 1, N \rrbracket$, $\rho_n = \alpha_n^{-p}$ with the convention that $0^{-p} = +\infty$. Then, from Assertion (ii) of Lemma VI.3, we get that

$C_{\mathbf{W}}(\boldsymbol{\rho}) = 0$. By definition of $C_{\mathbf{W}}(\boldsymbol{\rho})$ and Relation (12), this gives that for all $m \in \mathbb{N}$, $\inf_{\mathcal{A}^{(m)}} h_{\mathbf{W}} = 0$ with $\mathcal{A}^{(m)} = \prod_{n=1}^N \mathcal{A}_n(\boldsymbol{\rho}, 2^{-m/p})$. In particular, we can find $\mathbf{a}^{(1:N)}(m) \in \mathcal{A}^{(m)}$ such that $h_{\mathbf{W}}(\mathbf{a}^{(1:N)}(m)) \leq 2^{-m}$. Now, take, for all $m \in \mathbb{N}$, $\boldsymbol{\theta}(m) = (\boldsymbol{\lambda}(m), \mathbf{A}^{(1:N)})$ with $\boldsymbol{\lambda}(m) = [m, 0, \dots, 0]$ and, for all $n \in \llbracket 1, N \rrbracket$, $\mathbf{A}^{(n)}(m) = [\mathbf{a}^{(n)}(m), \mathbf{a}_2^{(n)}, \dots, \mathbf{a}_R^{(n)}]$, where $\mathbf{a}_2^{(n)}, \dots, \mathbf{a}_R^{(n)}$ are arbitrary elements of $\mathbb{S}_{\nu_n}^+$. In this case, we have $\|\boldsymbol{\theta}(m)\|_2 \xrightarrow{m \rightarrow +\infty} +\infty$, but $(g_{\mathbf{W}}(\boldsymbol{\theta}(m)))^2 + \tilde{\mathcal{P}}_{\boldsymbol{\alpha}}(\boldsymbol{\theta}(m)) \leq (m^2 + Nm^d)2^{-m} \xrightarrow{m \rightarrow +\infty} 0$, which concludes the proof. \square

Proof of Lemma III.1. The fact that none of the μ_n 's in Points (i) and (ii) is a norm on the positive cone proves that $\psi_n = \infty$ in both cases. We now show that $\psi'_n < \infty$ holds in both cases. By equivalence of the norms, we can assume without loss of generality that $\nu_n = \|\cdot\|_{\infty}$. For μ_n as in Case (i), we can take $p = 1$ without loss of generality. Then, for all $i, j \in \llbracket 1, I_n \rrbracket$, we have $|a_j - a_i| \leq \sum_{k=j \wedge i}^{j \vee i - 1} |a_{k+1} - a_k| \leq \|\mathbf{a}\|_{\text{TV}, 1}$. This implies $\psi'_n < \infty$. Next, we take μ_n as in Case (ii). Let $\mathbf{a} \in \mathbb{B}_+^+$ such that $a_j = 0$ for some $j \in \llbracket 1, I_n \rrbracket$. Let $(\hat{a}_k)_{k \in \mathbb{Z}}$ be the Fourier coefficients of the corresponding Spline function a , $\hat{a}_k = \int_0^1 a(u) e^{-2i\pi k u} du$. Then, we have $\mu_n(\mathbf{a}) = \left(\sum_{k \in \mathbb{Z}} |\hat{a}_k|^2 (2\pi k)^2 \right)^{1/2} \leq 1$. Now, using the Cauchy-Schwarz inequality and the fact that $1/12 = \sum_{k \in \mathbb{Z}^*} (2\pi k)^{-2}$, where \mathbb{Z}^* is the set of non-zero relative integers, we get that

$$\sum_{k \in \mathbb{Z}^*} |\hat{a}_k| \leq \left(\sum_{k \in \mathbb{Z}^*} (2\pi k)^{-2} \right)^{1/2} \mu_n(\mathbf{a}) \leq \frac{1}{2\sqrt{3}}.$$

In particular, we get $\sum_{k \in \mathbb{Z}} |\hat{a}_k| < +\infty$ and therefore, for all $u \in [0, 1]$, we have $a(u) = \sum_{k \in \mathbb{Z}} \hat{a}_k e^{2i\pi k u}$. This implies that

$$\begin{aligned} \|\mathbf{a}\|_{\infty} &= \|\mathbf{a} - a_j\|_{\infty} \\ &\leq \sup_{u \in [0, 1]} |a(u) - a(u_j)| \\ &\leq \sum_{k \in \mathbb{Z}^*} |\hat{a}_k| \sup_{u \in [0, 1]} |e^{2i\pi k u} - e^{2i\pi k u_j}| \\ &\leq 2 \sum_{k \in \mathbb{Z}^*} |\hat{a}_k|. \end{aligned}$$

Hence, with the previous inequality, $\psi'_n \leq \frac{1}{\sqrt{3}}$ in this case, which concludes the proof. \square

REFERENCES

- [1] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, no. 3, pp. 455–500, September 2009.
- [2] A. Cichocki, R. Zdunek, A. H. Phan, and S. Amari, *Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-Way Data Analysis and Blind Source Separation*. Wiley Publishing, 2009.
- [3] N. D. Sidiropoulos, L. De Lathauwer, X. Fu, K. Huang, E. E. Papalexakis, and C. Faloutsos, "Tensor decomposition for signal processing and machine learning," *IEEE Transactions on Signal Processing*, vol. 65, no. 13, pp. 3551–3582, 2017.
- [4] C. J. Hillar and L.-H. Lim, "Most tensor problems are np-hard," *J. Acm*, vol. 60, no. 6, Nov. 2013. [Online]. Available: <https://doi.org/10.1145/2512329>
- [5] V. de Silva and L.-H. Lim, "Tensor rank and the ill-posedness of the best low-rank approximation problem," *SIAM J. Matrix Anal. Appl.*, vol. 30, no. 3, pp. 1084–1127, 2008. [Online]. Available: <https://doi.org/10.1137/06066518X>
- [6] L.-H. Lim and P. Comon, "Nonnegative approximations of nonnegative tensors," *Journal of Chemometrics*, vol. 23, no. 7-8, pp. 432–441, 2009. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cem.1244>
- [7] A. N. Tikhonov and V. Y. Arsenin, *Solutions of ill-posed problems*, ser. Scripta Series in Mathematics. V. H. Winston & Sons, Washington, D.C.: John Wiley & Sons, New York-Toronto, Ont.-London, 1977, translated from the Russian, Preface by translation editor Fritz John.
- [8] V. N. Vapnik, *Statistical learning theory*, ser. Adaptive and Learning Systems for Signal Processing, Communications, and Control. John Wiley & Sons, Inc., New York, 1998, a Wiley-Interscience Publication.
- [9] L.-H. Lim, "Optimal solutions to non-negative parafac/multilinear nmf always exist," in *Workshop on Tensor Decompositions and Applications, Centre International de rencontres Mathématiques, Luminy, France*, 2005.
- [10] R. Bro, "Parafac. tutorial and applications," *Chemometrics and Intelligent Laboratory Systems*, vol. 38, no. 2, pp. 149–171, 1997. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169743997000324>
- [11] C. A. Andersson and R. Bro, "The n-way toolbox for matlab," *Chemometrics and Intelligent Laboratory Systems*, vol. 52, no. 1, pp. 1–4, 2000. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016974390000071X>
- [12] G. Tomasi and R. Bro, "Parafac and missing values," *Chemometrics and Intelligent Laboratory Systems*, vol. 75, no. 2, pp. 163–180, 2005. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169743904001741>
- [13] E. Acar, D. M. Dunlavy, T. G. Kolda, and M. Mørup, "Scalable tensor factorizations for incomplete data," *Chemometrics and Intelligent Laboratory Systems*, vol. 106, no. 1, pp. 41–56, 2011, multiway and Multiset Data Analysis. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0169743910001437>
- [14] L. Xiong, X. Chen, T.-K. Huang, J. Schneider, and J. G. Carbonell, *Temporal Collaborative Filtering with Bayesian Probabilistic Tensor Factorization*, 2010, pp. 211–222. [Online]. Available: <https://pubs.siam.org/doi/abs/10.1137/1.9781611972801.19>
- [15] P. Rai, Y. Wang, S. Guo, G. Chen, D. Dunson, and L. Carin, "Scalable bayesian low-rank decomposition of incomplete multiway tensors," in *Proceedings of the 31st International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, E. P. Xing and T. Jebara, Eds., vol. 32, no. 2. Beijing, China: Pmlr, 22–24 Jun 2014, pp. 1800–1808. [Online]. Available: <https://proceedings.mlr.press/v32/rai14.html>
- [16] Q. Zhao, G. Zhou, L. Zhang, A. Cichocki, and S.-I. Amari, "Bayesian robust tensor factorization for incomplete multiway data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 27, no. 4, pp. 736–748, 2016. [Online]. Available: <https://doi.org/10.1109/TNNLS.2015.2423694>
- [17] Q. Song, H. Ge, J. Caverlee, and X. Hu, "Tensor completion algorithms in big data analytics," *ACM Trans. Knowl. Discov. Data*, vol. 13, no. 1, Jan. 2019. [Online]. Available: <https://doi.org/10.1145/3278607>
- [18] M. E. Timmerman and H. A. Kiers, "Three-way component analysis with smoothness constraints," *Computational Statistics & Data Analysis*, vol. 40, no. 3, pp. 447–470, 2002. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167947302000592>
- [19] M. M. Reis and M. M. C. Ferreira, "Parafac with splines: a case study," *Journal of Chemometrics*, vol. 16, no. 8-10, pp. 444–450, 2002. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cem.749>
- [20] T. Yokota, R. Zdunek, A. Cichocki, and Y. Yamashita, "Smooth nonnegative matrix and tensor factorizations for robust multi-way data analysis," *Signal Processing*, vol. 113, pp. 234–249, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0165168415000614>
- [21] T. Yokota and A. Cichocki, "Tensor completion via functional smooth component deflation," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 2514–2518.
- [22] X. Li, Y. Ye, and X. Xu, "Low-rank tensor completion with total variation for visual data inpainting," in *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, ser. Aaai'17. AAAI Press, 2017, p. 2210–2216.
- [23] M. Imaizumi and K. Hayashi, "Tensor decomposition with smoothness," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. International Convention Centre, Sydney, Australia: Pmlr, 06–11 Aug 2017, pp. 1597–1606. [Online]. Available: <http://proceedings.mlr.press/v70/imaizumi17a.html>
- [24] T. Sadowski and R. Zdunek, "Image completion with smooth non-negative matrix factorization," in *Artificial Intelligence and Soft Com-*

- puting, L. Rutkowski, R. Scherer, M. Korytkowski, W. Pedrycz, R. Tadeusiewicz, and J. M. Zurada, Eds. Cham: Springer International Publishing, 2018, pp. 62–72.
- [25] A. Durand, F. Roueff, J.-M. Jicquel, and N. Paul, “Smooth nonnegative tensor factorization for multi-sites electrical load monitoring,” in *Eusipco*, Dublin, Ireland, Aug. 2021. [Online]. Available: <https://hal.telecom-paris.fr/hal-03167498>
- [26] U. Henriët, S. and Şimşekli, S. Dos Santos, B. Fuentes, and G. Richard, “Independent-variation matrix factorization with application to energy disaggregation,” *IEEE Signal Processing Letters*, vol. 26, no. 11, pp. 1643–1647, 2019.
- [27] Y. Gousseau and J.-M. Morel, “Are natural images of bounded variation?” *SIAM Journal on Mathematical Analysis*, vol. 33, no. 3, pp. 634–648, 2001. [Online]. Available: <https://doi.org/10.1137/S0036141000371150>
- [28] T. Yokota, Q. Zhao, and A. Cichocki, “Smooth parafac decomposition for tensor completion,” *IEEE Transactions on Signal Processing*, vol. 64, no. 20, pp. 5423–5436, 2016.
- [29] R. Zdunek, A. Cichocki, and T. Yokota, “B-spline smoothing of feature vectors in nonnegative matrix factorization,” in *Artificial Intelligence and Soft Computing*, L. Rutkowski, M. Korytkowski, R. Scherer, R. Tadeusiewicz, L. A. Zadeh, and J. M. Zurada, Eds. Cham: Springer International Publishing, 2014, pp. 72–81.
- [30] A. A. Amini, E. Levina, and K. A. Shedden, “Structured regression models for high-dimensional spatial spectroscopy data,” *Electronic Journal of Statistics*, vol. 11, no. 2, pp. 4151 – 4178, 2017. [Online]. Available: <https://doi.org/10.1214/17-EJS1301>
- [31] C. Hautecoeur and F. Glineur, “Nonnegative matrix factorization over continuous signals using parametrizable functions,” *Neurocomputing*, vol. 416, pp. 256–265, 2020.