

On the identifiability and stable recovery of deep/multi-layer structured matrix factorization

François Malgouyres
 Institut de Mathématiques de Toulouse ; UMR5219
 Université de Toulouse ; CNRS
 UPS IMT, F-31062 Toulouse Cedex 9, France
<http://www.math.univ-toulouse.fr/~fmalgouy/>

Joseph Landsberg
 Department of Mathematics
 Mailstop 3368
 Texas A& M University
 College Station, TX 77843-3368
 Email: jml@math.tamu.edu

Abstract—We study a deep/multi-layer structured matrix factorization problem. It approximates a given matrix by the product of K matrices (called factors). Each factor is obtained by applying a fixed linear operator to a short vector of parameters (thus the name "structured"). We call the model deep or multi-layer because the number of factors is not limited. In the practical situations we have in mind, we typically have $K = 10$ or 20 .

We provide necessary and sufficient conditions for the identifiability of the factors (up to a scale rearrangement). We also provide a sufficient condition that guarantees that the recovery of the factors is stable. A practical example where the deep structured factorization is a convolutional tree is provided in an accompanying paper.

I. INTRODUCTION

A. Content of the paper

We consider the following matrix factorization problem: let $m_1 \dots m_{K+1} \in \mathbb{N}$, write $m_1 = m$, $m_{K+1} = n$, and let $X \in \mathbb{R}^{m \times n}$ be a matrix. We study the approximation of X by the product of factors $X_1 \dots X_K$ where $X_j \in \mathbb{R}^{m_j \times m_{j+1}}$. As usual, to avoid trivial solutions and construct useful factorizations, it is important that the factors are subject to strong restrictions (such as computational constraints or sparsity). In this paper, we impose the factors to be structured matrices defined by a small number S of unknown parameters. In practical application X contains some data and is only specified indirectly and/or approximately.

More precisely, for $k = 1 \dots K$, let

$$M_k : \mathbb{R}^S \longrightarrow \mathbb{R}^{m_k \times m_{k+1}}, \quad (1)$$

$$h \longmapsto M_k(h) \quad (2)$$

be a linear map.

We consider the structured matrix factorization problem involving K layers whose outputs are matrices $M_k(h_k^*)$ where the vectors $(h_k^*)_{k=1..K}$ are such that

$$M_1(h_1^*) \dots M_K(h_K^*) = X, \quad (3)$$

when such a factorization exists.

We establish simple conditions on the structure of the operators M_k and X such that, up to obvious scale rearrangement, the solution of (3) is unique.

When (3) has a solution, the set of all solutions coincides with the set of minimizers

$$\operatorname{argmin}_{(h_k)_{k=1..K} \in \mathbb{R}^{KS}} \|M_1(h_1) \dots M_K(h_K) - X\|^2. \quad (4)$$

However, when (3) does not have a solution but the operators M_k are such that the function

$$(h_k)_{k=1..K} \longmapsto \|M_1(h_1) \dots M_K(h_K) - X\|^2$$

is coercive (it clearly is continuous so (4) has solutions), a solution of (4) can be useful even when (3) does not have any solution. Therefore, the latter problem therefore appears to be slightly more general and it still makes sense to study whether (4) has a unique solution or not.

One motivation for this project was to understand whether the kernels of convolutional trees [1] can be recovered and to make progress concerning the recovery of deep sparse factorizations [2]. These topics will be discussed in an accompanying paper, together with complementary results and detailed proofs of the statements made in the publication.

The difficulty, when studying (3) and (4), comes from the fact that, because of the product of the factors, they are in general highly non-linear/non-convex. The non-uniqueness of the solution if therefore not easy to characterize and, to the best of our knowledge, there does not exist any general uniqueness guarantee for this problem.

In this article, we first provide known and preliminary results related to the considered structured deep factorization model. Then we establish simple geometric necessary and sufficient conditions guaranteeing the uniqueness of the solution to (4) (see Corollary 2). The condition is then relaxed to obtain a simpler sufficient condition (see Theorem 2)¹. Then, we generalize these properties and provide sufficient conditions guaranteeing the stable recovery of the factors (up to scaling ambiguity).

B. Examples of Multi-layer/deep structured matrix factorization

Matrix factorization problems are ubiquitous in statistics and data representation. In its simplest form, the model only

¹Notice that this condition is not truly sufficient, since it fails to provide the conclusion on a set of problems of measure 0. It, however, illustrates very well the situations in which identifiability should hold.

contains one layer (i.e. $K = 1$). This is for instance the case in linear approximation. In this case, X can be vectorized to form a column vector and the operator M_1 simply multiplies the column vector h_1 by a fixed (rectangular) matrix. Typically, in this setting, the latter matrix has more rows than columns and the uniqueness of the solution to (3) naturally relates to its column rank.

Notice that the above linear approximation is often improved using a "non-linear approximation" [3]. In this framework, the fixed matrix has more columns than rows and h_1 contains the inputs of a sparse vector whose support is also optimized to better approximate the data. This so called "non-linear approximation" is not under the scope of the present paper. The knowledge of the support is a good example of the kind of structure that we assume fixed in the present work. The main focus of the present work is *not* to study if the support is efficiently estimated (as is common in compressed sensing and in other work providing quantitative estimates of the reconstruction error).

The questions we are studying are mostly relevant when $K \geq 2$. In the case of such models, the non-linearity comes from the multiplicative nature of (3) and the identifiability is not easily guaranteed. Recently, in sparse coding and dictionary learning (see [4] for an overview), X contains the data and (most of the time) people consider two layers: $K = 2$. The layer $M_1(h_1)$ is an optimized dictionary of atoms and each column of $M_2(h_2)$ contains the code (or coordinates) of the corresponding column in X . In this case, the mapping M_2 maps a vector from a small vector space into a sparse matrix having a prescribed support. Again, we make the simplifying assumption that the supports are known (this is a strong restriction that rules the dictionary learning models out of the present study). Factorizations with $K = 2$ have also been considered for the purpose of blind deconvolution [5], [6] and blind deconvolution in random mask imaging [7].

Handcrafted deep matrix factorization of a few particular matrices are used in many fields of mathematics and engineering. In particular, most fast transforms are deep matrix factorization. This is for instance the case of the Cooley-Tukey Fast Fourier Transform, the Discrete Cosine Transform or the Wavelet transform.

The construction of optimized deep matrix factorization only started recently (see [8], [1] and references therein). In [8], [1], the authors consider compositions of sparse convolutions organized according to a tree, that is: a convolutional tree. In the simplified case studied in [8], X is a vector, the vectors h_k are the convolution kernels and each operator M_k maps h_k in a circulant (or block-circulant) matrix. The first layer corresponds to the coordinates/code of X in the frame obtained by computing the compositions of convolutions along the unique branch of the tree. Also, in [2], the authors consider a factorization involving several sparse layers. In that work, the authors simultaneously estimate the support and the coefficients of each sparse factor.

C. Identifiability of the matrix factors

When $K = 1$, the identifiability of the matrix factors is a simple, well investigated, issue.

When $K = 2$, the identifiability of the factors has been studied in many dictionary learning contexts and provides guarantees on the approximate recovery of both an incoherent dictionary and sparse coefficients when the number of samples is sufficiently large (i.e. n is large, in our setting). For instance, [9] provides such guarantees under general conditions which cover many practical settings. The problem significantly differs from the present setting, since the support of the sparse factor is assumed unknown while we assume it known.

When $K = 2$ some papers use the same lifting property we are using. They further propose to convexify the problem and provide sufficient conditions for obtaining the identifiability and stability in specific contexts: Blind deconvolution in random mask imaging [7] and Blind deconvolution [5]. A more general bilinear framework is considered in [6], where the analysis shares similarities with the results presented here but are restricted to the identifiability when $K = 2$. Our work extends these results by considering the identifiability and the stability of the recovery for any $K \geq 2$.

To the best of our knowledge, nothing is known concerning the identifiability and the stability of a matrix factorization when $K > 2$.

II. NOTATION

We continue to use the notation introduced in the introduction. For any integer $k \in \mathbb{N}$, the set $\mathbb{N}_k = \{1, \dots, k\}$.

We consider real valued tensors of order K whose axis are of size S , denoted by $T \in \mathbb{R}^{S \times \dots \times S}$. The space of tensors is abbreviated \mathbb{R}^{S^K} . The entries of T are denoted by T_{i_1, \dots, i_K} , where $(i_1, \dots, i_K) \in (\mathbb{N}_S)^K$. The index set is simply denoted \mathbb{N}_S^K . For $\mathbf{i} \in \mathbb{N}_S^K$, the entries of \mathbf{i} are $\mathbf{i} = (i_1, \dots, i_K)$ (for $\mathbf{j} \in \mathbb{N}_S^K$ we let $\mathbf{j} = (j_1, \dots, j_K)$ etc.). We either write $T_{\mathbf{i}}$ or T_{i_1, \dots, i_K} .

A collection of vectors is denoted $\mathbf{h} \in \mathbb{R}^{KS}$ (i.e. using bold fonts). Our collections are composed of K vectors of size S and the k^{th} vector is denoted $\mathbf{h}_k \in \mathbb{R}^S$. The i^{th} entry of the k^{th} vector is denoted $\mathbf{h}_{k,i} \in \mathbb{R}$. A vector not related to a collection of vectors is denoted by $h \in \mathbb{R}^S$ (i.e. using a light font).

All the vector spaces \mathbb{R}^{S^K} , \mathbb{R}^{KS} , \mathbb{R}^S etc. are equipped with the usual Euclidean norm. This norm is denoted $\|\cdot\|$ and the scalar product $\langle \cdot, \cdot \rangle$. In the particular case of matrices, $\|\cdot\|$ therefore corresponds to the Frobenius norm. We also use the usual infinity norm and denote it by $\|\cdot\|_\infty$.

Set

$$\mathbb{R}_*^{KS} = \{\mathbf{h} \in \mathbb{R}^{KS}, \forall k \in \mathbb{N}_K, \|\mathbf{h}_k\| \neq 0\}. \quad (5)$$

Define an equivalence relation in \mathbb{R}_*^{KS} : for any $\mathbf{h}, \mathbf{g} \in \mathbb{R}_*^{KS}$, $\mathbf{h} \sim \mathbf{g}$ if and only if there exists $(\lambda_k)_{k \in \mathbb{N}_K} \in \mathbb{R}^K$ such that

$$\prod_{k=1}^K \lambda_k = 1 \text{ and } \forall k \in \mathbb{N}_K, \mathbf{h}_k = \lambda_k \mathbf{g}_k.$$

Denote the equivalence class of $\mathbf{h} \in \mathbb{R}_*^{KS}$ by $[\mathbf{h}]$.

We say that a tensor $T \in \mathbb{R}^{S^K}$ is of rank 1 if and only if there exists a collection of vectors $\mathbf{h} \in \mathbb{R}^{KS}$ such that T is the outer product of the vectors \mathbf{h}_k , for $k \in \mathbb{N}_K$, that is, for any $\mathbf{i} \in \mathbb{N}_S^K$,

$$T_{\mathbf{i}} = \mathbf{h}_{1,i_1} \dots \mathbf{h}_{K,i_K}.$$

The set of all the tensors of rank 1 is denoted by Σ_1 .

The rank of any tensor $T \in \mathbb{R}^{S^K}$ is defined to be

$$\text{rk}(T) = \min\{r \in \mathbb{N} \mid \text{there exists } T_1, \dots, T_r \in \Sigma_1 \text{ such that } T = T_1 + \dots + T_r\}.$$

For $r \in \mathbb{N}$, let

$$\Sigma_r = \{T \in \mathbb{R}^{S^K}, \text{rk}(T) \leq r\}.$$

The * superscript refers to optimal solutions. A set with a * subscript means that 0 is ruled out of the set. In particular, $\Sigma_{1,*}$ denotes the non zero tensors of rank 1. Attention should be paid to \mathbb{R}_*^{KS} since its definition is not straightforward (see (5)).

III. FACTS ON THE SEGRE EMBEDDING AND TENSORS OF RANK 1 AND 2

Parametrize $\Sigma_1 \subset \mathbb{R}^{S^K}$ by the map

$$\begin{aligned} P : \mathbb{R}^{KS} &\longrightarrow \Sigma_1 \subset \mathbb{R}^{S^K} \\ \mathbf{h} &\longmapsto (\mathbf{h}_{1,i_1} \mathbf{h}_{2,i_2} \dots \mathbf{h}_{K,i_K})_{\mathbf{i} \in \mathbb{N}_S^K} \end{aligned} \quad (6)$$

The map \widehat{P} is called the Segre embedding and is often denoted by \widehat{Seg} in the algebraic geometry literature. We use a simpler notation in the present paper since it is the only non-linear mapping we are considering.

We will use the following standard facts:

- 1) **Identifiability of $[\mathbf{h}]$ from $P(\mathbf{h})$:** For \mathbf{h} and $\mathbf{g} \in \mathbb{R}_*^{KS}$, $P(\mathbf{h}) = P(\mathbf{g})$ if and only if $[\mathbf{h}] = [\mathbf{g}]$.
- 2) **Geometrical description of $\Sigma_{1,*}$:** $\Sigma_{1,*}$ is a smooth (i.e. C^∞) manifold of dimension $K(S-1) + 1$ (see, e.g., [10], chapter 4, pp. 103).
- 3) **Geometrical description of Σ_2 :** When $K = 2$, Σ_2 is a smooth manifold of dimension $4(S-1)$. When $K \geq 3$, there exists a closed set $C \subset \Sigma_2$, whose Hausdorff measure of dimension $2K(S-1) + 2$ is 0, such that $\Sigma_2 \setminus C$ is a smooth manifold of dimension $2K(S-1) + 2$. See, e.g., [10], chapter 5.

Define

$$\|[\mathbf{h}] - [\mathbf{g}]\|_\infty = \inf_{\mathbf{h}' \in [\mathbf{h}], \mathbf{g}' \in [\mathbf{g}]} \|\mathbf{h}' - \mathbf{g}'\|_\infty. \quad (7)$$

Theorem 1. Stability of $[\mathbf{h}]$ from $P(\mathbf{h})$

Let \mathbf{h} and $\mathbf{g} \in \mathbb{R}_*^{KS}$ be such that $\|P(\mathbf{h})\|_\infty \geq \|P(\mathbf{g})\|_\infty$ and

$$\|P(\mathbf{g}) - P(\mathbf{h})\|_\infty \leq \frac{1}{2} \|P(\mathbf{h})\|_\infty.$$

We have

$$\|[\mathbf{h}] - [\mathbf{g}]\|_\infty \leq 5 \|P(\mathbf{h})\|_\infty^{\frac{1}{K}-1} \|P(\mathbf{h}) - P(\mathbf{g})\|_\infty. \quad (8)$$

²These tensors are also called decomposable.

It is not difficult to build a simple example for which, modulo the 5 factor, the bound is reached. The upper bound in (8) is therefore tight, modulo a universal constant factor.

Notice that, if \mathbf{h} and $\mathbf{g} \in \mathbb{R}_*^{KS}$ are such that (for instance) $\mathbf{h}_1 = \mathbf{g}_1$, we have $\|[\mathbf{h}] - [\mathbf{g}]\|_\infty = 0$ even though we might have $\mathbf{h}_2 \neq \mathbf{g}_2$. In that case, typically, the infimum in (7) is not reached. This implies that our definition of $\|[\cdot] - [\cdot]\|_\infty$ is not a distance between equivalence classes. Using this property we can easily build two collections of vectors \mathbf{h} and $\mathbf{g} \in \mathbb{R}_*^{KS}$ such that $\|[\mathbf{h}] - [\mathbf{g}]\|_\infty = 0$ but $\|P(\mathbf{h}) - P(\mathbf{g})\|_\infty \neq 0$.

IV. THE LIFTING PRINCIPLE

Proposition 1. *The entries of the matrix*

$$M_1(\mathbf{h}_1)M_2(\mathbf{h}_2) \dots M_K(\mathbf{h}_K)$$

are multivariate polynomials whose variables are the entries of $\mathbf{h} \in \mathbb{R}^{KS}$. Moreover, every polynomial is the sum of monomials of degree K . Each monomial is a constant times $\mathbf{h}_{1,i_1} \dots \mathbf{h}_{K,i_K}$, for some $\mathbf{i} \in \mathbb{N}_S^K$.

Notice that any monomial $\mathbf{h}_{1,i_1} \dots \mathbf{h}_{K,i_K}$ is the entry $P(\mathbf{h})_{\mathbf{i}}$ in the tensor $P(\mathbf{h})$. Therefore every polynomial in the previous proposition take the form $\sum_{\mathbf{i} \in \mathbb{N}_S^K} c_{\mathbf{i}} P(\mathbf{h})_{\mathbf{i}}$ for some constants $(c_{\mathbf{i}})_{\mathbf{i} \in \mathbb{N}_S^K}$ independent of \mathbf{h} . In words, every entry of the matrix $M_1(\mathbf{h}_1)M_2(\mathbf{h}_2) \dots M_K(\mathbf{h}_K)$ is obtained by applying a linear form to $P(\mathbf{h})$. This leads to the following statement.

Corollary 1. *Let M_k , $k \in \mathbb{N}_K$ be as in (1). The map*

$$(\mathbf{h}_1 \dots \mathbf{h}_K) \mapsto M_1(\mathbf{h}_1)M_2(\mathbf{h}_2) \dots M_K(\mathbf{h}_K),$$

uniquely determines a linear map

$$\mathcal{A} : \mathbb{R}^{S^K} \longrightarrow \mathbb{R}^{m \times n},$$

such that for all $\mathbf{h} \in \mathbb{R}^{KS}$

$$M_1(\mathbf{h}_1)M_2(\mathbf{h}_2) \dots M_K(\mathbf{h}_K) = \mathcal{A}P(\mathbf{h}). \quad (9)$$

Notice that it is most of the time very difficult to provide a closed form expression for the operator \mathcal{A} . However, when the operators M_k are known, we can compute $\mathcal{A}P(\mathbf{h})$, using (9). Said differently, we can compute \mathcal{A} for any rank 1 entry. Therefore, since \mathcal{A} is linear, we can compute $\mathcal{A}T$ for any low rank tensor T . In practice, if the dimensionality of the problem permits it, it is possible to build a basis of \mathbb{R}^{S^K} made of rank 1 tensors. It is then possible to manipulate \mathcal{A} in this basis.

Using Corollary 1, we rewrite the problem (4) in the form

$$\mathbf{h}^* \in \operatorname{argmin}_{\mathbf{h} \in \mathbb{R}^{KS}} \|\mathcal{A}P(\mathbf{h}) - X\|^2. \quad (10)$$

We now decompose this problem into two sub-problems: A least-squares problem

$$T^* \in \operatorname{argmin}_{T \in \mathbb{R}^{S^K}} \|\mathcal{A}T - X\|^2 \quad (11)$$

and a non-convex problem

$$\mathbf{h}^{*'} \in \operatorname{argmin}_{\mathbf{h} \in \mathbb{R}^{KS}} \|\mathcal{A}(P(\mathbf{h}) - T^*)\|^2. \quad (12)$$

Proposition 2. *For any X , \mathcal{A} and any T^* solving (11):*

- 1) *Any solution $\mathbf{h}^{*'}$ of (12) also minimizes (10) and (4).*

2) Any solution \mathbf{h}^* of (4) and (10) also minimizes (12).

Proof. The proof relies on the fact that for any $T^* \in \operatorname{argmin}_{T \in \mathbb{R}^{S \times K}} \|\mathcal{A}T - X\|^2$, we have

$$\mathcal{A}^t(\mathcal{A}T^* - X) = 0,$$

where \mathcal{A}^t is the adjoint of \mathcal{A} , which implies that for any $T^* \in \operatorname{argmin}_{T \in \mathbb{R}^{S \times K}} \|\mathcal{A}T - X\|^2$ and any $\mathbf{h} \in \mathbb{R}^{KS}$

$$\begin{aligned} \|\mathcal{A}P(\mathbf{h}) - X\|^2 &= \|\mathcal{A}(P(\mathbf{h}) - T^*) + (\mathcal{A}T^* - X)\|^2, \\ &= \|\mathcal{A}(P(\mathbf{h}) - T^*)\|^2 + \|\mathcal{A}T^* - X\|^2 \\ &\quad + 2\langle \mathcal{A}(P(\mathbf{h}) - T^*), (\mathcal{A}T^* - X) \rangle, \\ &= \|\mathcal{A}(P(\mathbf{h}) - T^*)\|^2 + \|\mathcal{A}T^* - X\|^2. \end{aligned}$$

In words, $\|\mathcal{A}P(\mathbf{h}) - X\|^2$ and $\|\mathcal{A}(P(\mathbf{h}) - T^*)\|^2$ only differ by an additive constant. Moreover, since the value of the objective function $\|\mathcal{A}T^* - X\|^2$ is independent of the particular minimizer T^* we are considering, this additive constant is independent of T^* . As a consequence, a minimizer of $\|\mathcal{A}P(\mathbf{h}) - X\|^2$ also minimizes $\|\mathcal{A}(P(\mathbf{h}) - T^*)\|^2$ and vice versa. \square

From now on, because of the equivalence between solutions of (12) and (10), we stop using the notation \mathbf{h}'^* and write $\mathbf{h}^* \in \operatorname{argmin}_{\mathbf{h} \in \mathbb{R}^{KS}} \|\mathcal{A}(P(\mathbf{h}) - T^*)\|^2$.

V. IDENTIFIABILITY IN THE NOISE FREE CASE

Throughout this section, we assume that X is such that there exists $\bar{\mathbf{h}} \in \mathbb{R}^{KS}$ such that

$$X = M_1(\bar{\mathbf{h}}_1) \dots M_K(\bar{\mathbf{h}}_K). \quad (13)$$

Under this assumption, $X = \mathcal{A}P(\bar{\mathbf{h}})$, so

$$P(\bar{\mathbf{h}}) \in \operatorname{argmin}_{T \in \mathbb{R}^{S \times K}} \|\mathcal{A}T - X\|^2.$$

Moreover, we trivially have $P(\bar{\mathbf{h}}) \in \Sigma_1$ and therefore $\bar{\mathbf{h}}$ minimizes (12), (4) and (10).

We ask whether there exist guarantees that the resolution of (4) allows one to recover $\bar{\mathbf{h}}$ (up to the usual uncertainties).

In this regard, for any $\mathbf{h} \in [\bar{\mathbf{h}}]$, we have $P(\mathbf{h}) = P(\bar{\mathbf{h}})$ and therefore $\mathcal{A}P(\mathbf{h}) = \mathcal{A}P(\bar{\mathbf{h}}) = X$. Thus unless we make further assumptions on $\bar{\mathbf{h}}$, we cannot expect to distinguish any particular element of $[\bar{\mathbf{h}}]$ using only X . In other words, recovering an element of $[\bar{\mathbf{h}}]$ is the best we can hope for.

Definition 1. We say that $[\bar{\mathbf{h}}]$ is identifiable if the elements of $[\bar{\mathbf{h}}]$ are the only solutions of (4).

We say that \mathbb{R}^{KS} is identifiable if $[\bar{\mathbf{h}}]$ is identifiable for every $\bar{\mathbf{h}} \in \mathbb{R}^{KS}$.

Proposition 3. Characterization of the global minimizers

For any $\mathbf{h}^* \in \mathbb{R}^{KS}$, $\mathbf{h}^* \in \operatorname{argmin}_{\mathbf{h} \in \mathbb{R}^{KS}} \|\mathcal{A}P(\mathbf{h}) - X\|^2$ if and only if

$$P(\mathbf{h}^*) \in P(\bar{\mathbf{h}}) + \operatorname{Ker}(\mathcal{A}).$$

Corollary 2. Necessary and sufficient conditions of identifiability

1) For any $\bar{\mathbf{h}} \in \mathbb{R}^{KS}$: $[\bar{\mathbf{h}}]$ is identifiable if and only if

$$(P(\bar{\mathbf{h}}) + \operatorname{Ker}(\mathcal{A})) \cap \Sigma_1 = P(\bar{\mathbf{h}}).$$

2) \mathbb{R}^{KS} is identifiable if and only if

$$\operatorname{Ker}(\mathcal{A}) \cap \Sigma_2 = \{0\}. \quad (14)$$

The drawback of Corollary 2 is that, given a factorization model described by \mathcal{A} , it might be difficult to check whether the condition (14) holds or not. In reasonably small cases, one can use tools from numerical algebraic geometry such as those described in [11], [12].

We now establish a simpler condition on the rank of the operator \mathcal{A} related to the identifiability of \mathbb{R}^{KS} . More precisely, there is a set³ of matrices \mathcal{A} of measure 0 such that for any \mathcal{A} outside this set, the condition on the rank of \mathcal{A} is sufficient to guarantee the identifiability of \mathbb{R}^{KS} . The Theorem therefore provides a simple clue telling us if a factorization model can recover any $\bar{\mathbf{h}}$.

Theorem 2. Almost surely sufficient condition for Identifiability

When $K = 2$, for almost every \mathcal{A} such that $\operatorname{rk}(\mathcal{A}) \geq 2K(S - 1)$, \mathbb{R}^{KS} is identifiable.

When $K \geq 3$, for almost every \mathcal{A} such that $\operatorname{rk}(\mathcal{A}) \geq 2K(S - 1) + 2$, \mathbb{R}^{KS} is identifiable.

Proof. We only provide the proof when $K \geq 3$ since the proof when $K = 2$ is similar.

When $\operatorname{rk}(\mathcal{A}) \geq 2K(S - 1) + 2$, we have $\dim(\operatorname{Ker}(\mathcal{A})) = S^K - \operatorname{rk}(\mathcal{A}) \leq S^K - (2K(S - 2) + 2) = S^K - \dim(\Sigma_2)$. As a consequence,

$$\dim(\operatorname{Ker}(\mathcal{A})) + \dim(\Sigma_2) \leq S^K.$$

where we remind that S^K is the dimension of the ambient space. Therefore, by taking the Zariski closure of Σ_2 over the complex numbers, and since everything is invariant under rescaling, the non-intersection follows from a standard result of algebraic geometry (see [13], Thm 17.24). Since the manifolds do not intersect when defined using complex numbers, their restriction over \mathbb{R} will not intersect either. \square

VI. STABLE RECOVERY IN THE NOISY CASE

In this section, we assume that X is such that there exists $\bar{\mathbf{h}} \in \mathbb{R}^{KS}$ and some error e satisfying

$$X = M_1(\bar{\mathbf{h}}_1) \dots M_K(\bar{\mathbf{h}}_K) + e$$

with

$$\|e\| \leq \delta. \quad (15)$$

This can equivalently be written

$$X = \mathcal{A}P(\bar{\mathbf{h}}) + e.$$

The question we address is : *Can we establish guarantees that any solution of (4) is close to $\bar{\mathbf{h}}$ (up to the usual uncertainties)?*

First note that for any minimizer

$$\mathbf{h}^* \in \operatorname{argmin}_{\mathbf{h} \in \mathbb{R}^{KS}} \|\mathcal{A}P(\mathbf{h}) - X\|^2,$$

³It seems unfortunately difficult to know whether a particular matrix is in this set or not.

we have

$$\begin{aligned} \|\mathcal{A}(P(\mathbf{h}^*) - P(\bar{\mathbf{h}}))\| &\leq \|\mathcal{A}P(\mathbf{h}^*) - X\| + \|\mathcal{A}P(\bar{\mathbf{h}}) - X\| \\ &\leq 2\|\mathcal{A}P(\bar{\mathbf{h}}) - X\| \\ &\leq 2\delta \end{aligned}$$

Geometrically, this means that $P(\mathbf{h}^*)$ belongs to a cylinder centered at $P(\bar{\mathbf{h}})$ whose direction is $\text{Ker}(\mathcal{A})$ and whose section is defined using the operator \mathcal{A} .

If we further decompose (the decomposition is unique)

$$P(\mathbf{h}^*) - P(\bar{\mathbf{h}}) = T + T',$$

where $T \in \text{Ker}(\mathcal{A})$ and T' is orthogonal to $\text{Ker}(\mathcal{A})$, we have

$$\|\mathcal{A}(P(\mathbf{h}^*) - P(\bar{\mathbf{h}}))\| = \|\mathcal{A}T'\| \geq \sigma_{\min}\|T'\|, \quad (16)$$

where σ_{\min} is the smallest non-zero eigenvalue of \mathcal{A} . We finally obtain

$$\|P(\mathbf{h}^*) - P(\bar{\mathbf{h}}) - T\| \leq \frac{2\delta}{\sigma_{\min}}. \quad (17)$$

The term on the left-hand side corresponds to the distance between a point in Σ_2 (namely $P(\mathbf{h}^*) - P(\bar{\mathbf{h}})$) and a point in $\text{Ker}(\mathcal{A})$ (namely T). As already illustrated in the previous section, if the two manifolds intersect away from 0, $\|P(\mathbf{h}^*) - P(\bar{\mathbf{h}}) - T\|$ might be small while both $\|P(\mathbf{h}^*) - P(\bar{\mathbf{h}})\|$ and $\|T\|$ are large. This is neither compatible with the identifiability nor, of course, with the stable recovery of \bar{h} . In order to exclude this scenario, we define the following hypothesis on $\text{Ker}(\mathcal{A})$.

Definition 2. Let $\gamma > 0$, we say that $\text{Ker}(\mathcal{A})$ is γ -transverse to Σ_2 if there exists $\varepsilon > 0$ such that for any $T \in \Sigma_2$ and any $T' \in \text{Ker}(\mathcal{A})$ satisfying $\|T - T'\| \leq \varepsilon$, we have

$$\gamma(\|T\| + \|T'\|) \leq \|T - T'\|.$$

This means that, for $T \in \Sigma_2$ and $T' \in \text{Ker}(\mathcal{A})$, the only option in order to get $\|T - T'\| \leq \varepsilon$ is that both $\|T\|$ and $\|T'\|$ are small. Moreover, this implies that, in the vicinity of 0, $\text{Ker}(\mathcal{A})$ and Σ_2 are not tangential.

Theorem 3. Sufficient condition for stable recovery

Assume $\text{Ker}(\mathcal{A})$ is γ -transverse to Σ_2 , for $\gamma > 0$, for any

$$\mathbf{h}^* \in \operatorname{argmin}_{\mathbf{h} \in \mathbb{R}^{KS}} \|\mathcal{A}P(\mathbf{h}) - X\|^2,$$

and for δ sufficiently small,

$$\|P(\mathbf{h}^*) - P(\bar{\mathbf{h}})\| \leq \frac{2}{\gamma \sigma_{\min}} \delta,$$

where σ_{\min} is defined in (16) and δ is defined in (15). Moreover, if $\|\bar{\mathbf{h}}\| \neq 0$

$$\|\mathbf{h}^* - \bar{\mathbf{h}}\|_{\infty} \leq \frac{10}{\gamma \sigma_{\min}} \|P(\bar{\mathbf{h}})\|_{\infty}^{\frac{1}{K}-1} \delta, \quad (18)$$

Proof. Using (17) and the γ -transversality, we know that for $\frac{2\delta}{\sigma_{\min}} \leq \varepsilon$, we have

$$\frac{2\delta}{\sigma_{\min}} \geq \gamma \|P(\mathbf{h}^*) - P(\bar{\mathbf{h}})\|.$$

Using Theorem 1 and the fact that $\|P(\mathbf{h}^*) - P(\bar{\mathbf{h}})\|_{\infty} \leq \|P(\mathbf{h}^*) - P(\bar{\mathbf{h}})\|$, we finally get (18). \square

This proposition provides a sufficient condition to get stable recovery. The only significant hypothesis made on the factorization problem is that $\text{Ker}(\mathcal{A})$ is γ -transverse to Σ_2 . One might ask whether this hypothesis is sharp or not. In this regard, it seems clear that if it does not hold:

- either, away from 0, $\text{Ker}(\mathcal{A})$ and Σ_2 become arbitrary close: Again two situations might occur: either their closures intersect and we cannot expect stable recovery; or they become close, at infinity. The latter situation should not occur since the objective function is coercive.
- or near 0, $\text{Ker}(\mathcal{A})$ and Σ_2 are tangential. Again this is not compatible with stable recovery.

ACKNOWLEDGMENT

The authors started this collaboration during the thematic trimester on Machine Learning of the CIMI Excellence laboratory, funded by the ANR-11-LABX-0040-CIMI within the program ANR-11-IDEX-0002-02. Joseph Landsberg is supported by NSF DMS-1006353, DMS-1405348.

REFERENCES

- [1] H. W. J. T. O. Chabiron, F. Malgouyres, "Optimization of a fast transform structured as a convolutional tree," *preprint HAL*, no. hal-01258514, 2016.
- [2] L. L. Magoaru and R. Gribonval, "Flexible multi-layer sparse approximations of matrices and applications," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 4, pp. 688–700, 2016.
- [3] R. Devore, "Nonlinear approximation," *Acta Numerica*, vol. 7, pp. 51–150, 1998.
- [4] R. Rubinstein, A. M. Bruckstein, and M. Elad, "Dictionaries for sparse representation," *Proc. IEEE - Special issue on applications of sparse representation and compressive sensing*, vol. 98, no. 6, pp. 1045–1057, 2010.
- [5] A. Ahmed, B. Recht, and J. Romberg, "Blind deconvolution using convex programming," *Information Theory, IEEE Transactions on*, vol. 60, no. 3, pp. 1711–1732, 2014.
- [6] S. Choudhary and U. Mitra, "Identifiability scaling laws in bilinear inverse problems," *arXiv preprint arXiv:1402.2637*, 2014.
- [7] S. Bahmani and J. Romberg, "Lifting for blind deconvolution in random mask imaging: Identifiability and convex relaxation," *SIAM Journal on Imaging Sciences*, vol. 8, no. 4, pp. 2203–2238, 2015.
- [8] O. Chabiron, F. Malgouyres, J.-Y. Tourneret, and N. Dobigeon, "Toward fast transform learning," *International Journal of Computer Vision*, pp. 1–22, 2014.
- [9] R. Gribonval, R. Jenatton, F. Bach, M. Kleinstueber, and M. Seibert, "Sample complexity of dictionary learning and other matrix factorizations," *Information Theory, IEEE Transactions on*, vol. 61, no. 6, pp. 3469–3486, June 2015.
- [10] J. Landsberg, *Tensors: Geometry and Applications*, ser. Graduate studies in mathematics. American Mathematical Soc., 2012, vol. 128.
- [11] J. D. Hauenstein and A. J. Sommese, "Witness sets of projections," *Applied Mathematics and Computation*, vol. 217, no. 7, pp. 3349 – 3354, 2010.
- [12] —, "Membership tests for images of algebraic sets by linear projections," *Applied Mathematics and Computation*, vol. 219, no. 12, pp. 6809 – 6818, 2013.
- [13] J. Harris, *Algebraic geometry: a first course*. Springer Science & Business Media, 2013, vol. 133.