



# A consistent and flexible framework for deep matrix factorizations

Pierre De Handschutter, Nicolas Gillis\*

Department of Mathematics and Operational Research, University of Mons, Belgium



## ARTICLE INFO

### Article history:

Received 19 June 2022

Revised 2 September 2022

Accepted 4 October 2022

### Keywords:

Deep matrix factorization

Loss functions

Constrained optimization

First-order methods

Hyperspectral unmixing

## ABSTRACT

Deep matrix factorizations (deep MFs) are recent unsupervised data mining techniques inspired by constrained low-rank approximations. They aim to extract complex hierarchies of features within high-dimensional datasets. Most of the loss functions proposed in the literature to evaluate the quality of deep MF models and the underlying optimization frameworks are not consistent because different losses are used at different layers. In this paper, we introduce two meaningful loss functions for deep MF and present a generic framework to solve the corresponding optimization problems. We illustrate the effectiveness of this approach through the integration of various constraints and regularizations, such as sparsity, nonnegativity and minimum-volume. The models are successfully applied on both synthetic and real data, namely for hyperspectral unmixing and extraction of facial features.

© 2022 Elsevier Ltd. All rights reserved.

## 1. Introduction

In the era of data science, the extraction of meaningful features in datasets is a crucial challenge. To do so, a fundamental class of unsupervised linear dimensionality reduction methods is low-rank matrix factorizations (LRMFs). Given a matrix  $X \in \mathbb{R}^{m \times n}$ , where each column is a data point in dimension  $m$ , and a factorization rank  $r$ , an LRMF approximates  $X$  as the product of two matrices of smaller inner dimension  $r$ ,  $W \in \mathbb{R}^{m \times r}$  and  $H \in \mathbb{R}^{r \times n}$ , such that  $X \approx WH$ . The columns of  $W$ , denoted  $W(:, k)$  for  $k = 1, \dots, r$ , in dimension  $m$ , are called the basis vectors, and the entries of the  $j$ th column of  $H$ , denoted  $H(:, j)$  for  $j = 1, \dots, n$ , indicate in which proportion each basis vector contributes to the corresponding data point, since  $X(:, j) \approx WH(:, j)$ .

Without any constraints on the factors,  $W$  and  $H$ , such a decomposition is highly non-unique. This has led researchers to use additional constraints, motivated by prior knowledge, to obtain uniqueness and interpretability of the factors. Two of the most widely used constraints are sparsity and nonnegativity. Such constrained LRMFs have had tremendous success in a wide variety of applications; see, e.g., [1–6], and the references therein.

Recently, surfing the wave of deep learning, LRMFs have been extended such that the original matrix  $X$  is decomposed as the product of more than two factors, which is referred to as deep MF. More precisely,  $L$  layers of decomposition are applied on the matrix  $X$  such that  $X \approx W_L H_L H_{L-1} \dots H_1$ . This is equivalent to considering

a hierarchical decomposition of  $X$  as follows:

$$\begin{aligned} X &\approx W_1 H_1, \\ W_1 &\approx W_2 H_2, \\ &\vdots \\ W_{L-1} &\approx W_L H_L. \end{aligned} \quad (1)$$

At each layer,  $W_l \in \mathbb{R}^{m \times r_l}$  and  $H_l \in \mathbb{R}^{r_l \times r_{l-1}}$  for  $l = 1, \dots, L$  with  $r_0 = n$ . The dimension  $r_l$  of each layer is the rank of the factorization at layer  $l$ . When the factorization is performed according to the scheme of (1), the ranks are generally chosen in decreasing order, that is  $r_1 \geq r_2 \geq \dots \geq r_L$ . This can be easily understood: the first levels of decomposition extract low-level features, closely related to data while the last ones consist in a few key high-level features. Also, increasing ranks would lead to trivial factorizations. In fact, if  $r_l \leq r_{l+1}$  for some  $l$ ,  $W_l = W_{l+1} H_{l+1} = (W_l 0) \begin{pmatrix} r_l \\ 0 \end{pmatrix}$  where  $0$  is the matrix of zeros of appropriate dimension, would be a feasible factorization that does not bring any informative feature.

Deep MFs have the power to decompose hierarchically an input dataset, with different levels of interpretation at each layer, and are currently used in many applications, such as hyperspectral unmixing [7], the extraction of facial features [8–10], recommender systems [11], and multi-view clustering [12–14] among others.

As for single-layer factorizations, additional constraints are typically assumed on the factors  $W_l$ 's and  $H_l$ 's. Indeed, without any constraint on the factors of deep MF, deep MFs degenerate into (overparametrized) classical MFs: the product of the factors  $H_l$ 's could be replaced by a single matrix whose rank is less than or equal to the minimum of the  $r_l$ 's. Hence, the LRMFs with constraints, such as sparsity and nonnegativity, have been extended to a multilayer framework; see [15,16] and the references therein

\* Corresponding author.

E-mail addresses: [pierre.dehandschutter@umons.ac.be](mailto:pierre.dehandschutter@umons.ac.be) (P. De Handschutter), [nicolas.gillis@umons.ac.be](mailto:nicolas.gillis@umons.ac.be) (N. Gillis).

for details. Similar to deep MF, deep concept factorization [17] is an extension of archetypal analysis to several layers.

A crucial question that has not been discussed thoroughly yet is the choice of the loss function used to assess the quality of the factorization. The loss functions proposed so far in the literature are not consistent because different losses are optimized at different layers; see the discussion in Section 2 for more details. Therefore, in this paper, we propose a general framework with appropriate loss functions to solve efficiently deep MFs with general constraints on the factors. Especially, in the numerical experiments, we illustrate the performance of our framework compared to the state of the art.

The remainder of the paper is organized as follows. In Section 2, we describe the motivations behind our work regarding the state of the art. Then, in Section 3, we propose two new loss functions and discuss why they are more meaningful and consistent compared to previous works. Based on these new loss functions, we present in Section 4 a generic optimization framework relying on an extrapolated projected gradient method in order to tackle deep MFs with various constraints and regularizations. In Section 5.1, we present results on synthetic data, and in Section 5.2, we compare the models on real hyperspectral and facial images, before concluding in Section 6.

## 2. State of the art and motivations of the work

The first factorization model with several layers was proposed by Cichocki et al. [8], and is usually dubbed “multilayer matrix factorization”. It aims at minimizing  $\|W_{l-1} - W_l H_l\|_F^2$ , with  $W_0 = X$  successively for each layer  $l = 1, \dots, L$ . In other words, each factorization of (1) is performed one at a time, from the first to the last level, and the model is merely a succession of standard one-layer MFs. This hierarchical decomposition is sketched in Algorithm 1, in which  $\mathcal{W}_l$  and  $\mathcal{H}_l$  denote the feasible sets respectively for the factors  $W_l$ 's and  $H_l$ 's,  $l = 1, \dots, L$ , regardless of the specific constraints that apply (e.g., nonnegativity or sparsity). For each layer, a classical two block coordinate descent (BCD) is performed. It updates  $W_l$  and  $H_l$  alternatively until some stopping criterion, such as a maximum number of iterations or an insufficient decrease of the loss function between two consecutive iterations, is reached. Then, the algorithm moves to the next layer, which is in turn factorized and so on until the last layer ( $l = L$ ). At lines and , *arg reduce* denotes any algorithm that leads to a decrease of the corresponding loss function with the constraints enforced on  $H_l$ 's and  $W_l$ 's.

---

### Algorithm 1 Multilayer matrix factorization [18].

---

**Input:** Nonnegative data matrix  $X$ , number of layers  $L$ , inner ranks  $r_l$ 's and feasible sets  $\mathcal{W}_l$  and  $\mathcal{H}_l$  for  $l = 1, \dots, L$ .  
**Output:** Matrices  $W_1, \dots, W_L$  and  $H_1, \dots, H_L$ .

- 1:  $W_0 = X$
- 2: **for**  $l = 1, \dots, L$  **do**
- 3:   Initialize  $W_l^{(0)}$  and  $H_l^{(0)}$
- 4:   **for**  $k = 1, \dots$  **do**
- 5:      $H_l^{(k)} = \underset{H \in \mathcal{H}_l}{\text{arg reduce}} \|W_{l-1} - W_l^{(k-1)} H\|_F^2$
- 6:      $W_l^{(k)} = \underset{W \in \mathcal{W}_l}{\text{arg reduce}} \|W_{l-1} - W H_l^{(k)}\|_F^2$
- 7:   **end for**
- 8: **end for**

---

A pitfall of the approach of Cichocki et al. is that the factors of the last layers do not have any influence on those of the first ones since the first layers are factorized before the last ones. To remedy this, an iterative procedure was suggested by Trigeorgis et al. [19], called “deep MF”. The initialization of the factors is performed through Algorithm 1 and is followed by iterative updates

until some stopping criterion is met, as described in Algorithm 2. More precisely, once all the factors  $W_l$ 's and  $H_l$ 's have been updated once through sequential factorizations as in multilayer MF, several rounds of updates (starting with the first layer) are performed, taking into account the previous updates of the factors of the other layers. For this purpose, a global loss function was proposed by Trigeorgis et al. [19, Equation (10)], namely

---

### Algorithm 2 Deep MF with general constraints [19].

---

**Input:** Data matrix  $X$ , number of layers  $L$ , inner ranks  $r_l$ 's, feasible sets  $\mathcal{W}_l$  and  $\mathcal{H}_l$  for  $l = 1, \dots, L$   
**Output:** Matrices  $W_1, \dots, W_L$  and  $H_1, \dots, H_L$

- 1: Compute initial matrices  $W_l^{(0)}$  and  $H_l^{(0)}$  for all  $l$  through a sequential decomposition of  $X$  (for example Algorithm 1)
- 2: **for**  $k = 1, \dots$  **do**
- 3:   **for**  $l = 1, \dots, L$  **do**
- 4:      $A_l^{(k)} = \begin{cases} W_l^{(k-1)} & \text{if } l = L \\ W_{l+1}^{(k-1)} H_{l+1}^{(k-1)} & \text{otherwise} \end{cases}$
- 5:      $B_l^{(k)} = H_{l-1}^{(k-1)} \dots H_1^{(k-1)}$
- 6:      $H_l^{(k)} = \underset{H \in \mathcal{H}_l}{\text{arg reduce}} \|X - A_l^{(k)} H B_l^{(k)}\|_F^2$
- 7:      $W_l^{(k)} = \underset{W \in \mathcal{W}_l}{\text{arg reduce}} \|X - W H_l^{(k)} B_l^{(k)}\|_F^2$
- 8:   **end for**
- 9: **end for**

---

$$\mathcal{L}_0(H_1, H_2, \dots, H_L; W_L) = \|X - W_L H_L \dots H_2 H_1\|_F^2. \quad (2)$$

It is the squared Frobenius norm of the difference between the original matrix  $X$  and the approximation obtained by unfolding the  $L$  layers of (1). This loss function (2) was reused by most of the papers in the deep MF literature.

To understand more clearly how Algorithm 2 works, let us consider the simple case where  $L = 3$ . In Table 1, we report the loss function minimized at each stage of Algorithm 2, regarding the updates of lines and with the updated factor in bold, the others being fixed.

Table 1 shows that Algorithm 2 minimizes three different loss functions depending on which factor matrix is updated. More precisely, only the updates of the last layer ( $l = 3$ ) and the one of  $H_2$  are performed according to the “global” loss function claimed in (2). For  $L > 3$ , the situation is even worse, hence Algorithm 2 does not appear to be coherent since different loss functions are minimized along the factors updates. Consequently, convergence guarantees on this global loss function cannot be derived from such an optimization framework; see Fig. 10 and 11 in Section 5.2.2 that shows that Algorithm 2 does not converge on a numerical example.

Note that one cannot simply minimize  $\mathcal{L}_0$  in (2) at every layer to compute  $(W_L, H_L, \dots, H_1)$ . In fact, assuming the ranks are decreasing, that is,  $r_{l+1} \leq r_l$  for all  $l$  (this is the most reasonable setting; see Section 1), any solution  $(H_1, \dots, H_L; W_L)$  can be transformed into a degenerate solution  $(H_1^{(*)}, \dots, H_L^{(*)}; W_L^{(*)})$  with the

**Table 1**  
Loss functions minimized at each step of Algorithm 2 for  $L = 3$ . In bold, we indicate the matrix being updated.

Layer $l$	Update of $H_l$	Update of $W_l$
1	$\ X - W_2 H_2 H_1\ _F^2$	$\ X - W_1 H_1\ _F^2$
2	$\ X - W_3 H_3 H_2 H_1\ _F^2$	$\ X - W_2 H_2 H_1\ _F^2$
3	$\ X - W_3 H_3 H_2 H_1\ _F^2$	$\ X - W_3 H_3 H_2 H_1\ _F^2$

same loss function, and with the following form

$$H_l^{(*)} = \begin{pmatrix} I_{r_l} & \mathbf{0}_{r_l \times (r_{l-1} - r_l)} \\ \mathbf{0}_{(r_1 - r_l) \times r_l} & \mathbf{0}_{(r_1 - r_l) \times (r_{l-1} - r_l)} \end{pmatrix} \text{ for } l = 2, \dots, L,$$

where  $I_{r_l}$  is the identity matrix of dimension  $r_l$ , and  $\mathbf{0}_{m \times n}$  is the  $m$ -by- $n$  zero matrix,  $H_1^{(*)} = H_L \dots H_1$  and  $W_L^{(*)} = W_L$ . The reason is that  $W_L H_L \dots H_1$  has rank at most  $\min_l r_l = r_L$ , and with the choice above, we have

$$W_1^{(*)} = W_L^{(*)} H_L^{(*)} \dots H_2^{(*)} = [W_L \mathbf{0}_{m \times (r_1 - r_L)}] \in \mathbb{R}^{m \times r_1}$$

such that  $\text{rank}(W_1^{(*)}) \leq r_L$ . This means that deep MF would simply reduce to an overparametrized LRMF [20]. Therefore, deep MFs models need to properly balance the importance of each layer to compute useful hierarchical decompositions.

Most of the recent works only rely on the loss given by (2) and follow Algorithm 2. However, for the reasons mentioned above, this loss function and the use of Algorithm 2 are not consistent across layers, but this has not been much discussed in the literature yet. Besides, we show in Section 5 that this framework is not able to retrieve the ground truth basis vectors in practice. In fact, the loss function (2) only minimizes the error at the last layer of factorization but does not control the accuracy of the factorizations of the previous layers, which yet may be crucial regarding the applications of deep MF.

To alleviate this, we propose in Section 3 two new loss functions that are consistent, that is, are guaranteed to diminish after each update and reflect the balance between the errors due to each layer of decomposition.

### 3. Consistent loss functions for deep MFs

As explained in the previous section, the framework proposed by Trigeorgis et al. [19] is inconsistent. In this section, we propose two global loss functions that can be used to optimize any of the factors in deep MF. Hence, it is straightforward to derive meaningful update rules that ensure the decrease of the global loss function after the update of any factor.

#### 3.1. Layer-centric loss function

The first loss function proposed consists of a weighted sum of the errors caused by each layer of decomposition, that is, by the layer-wise factorizations:

$$\mathcal{L}_1(H_1, H_2, \dots, H_L; W_1, W_2, \dots, W_L) = \frac{1}{2} \left( \|X - W_1 H_1\|_F^2 + \lambda_1 \|W_1 - W_2 H_2\|_F^2 + \dots + \lambda_{L-1} \|W_{L-1} - W_L H_L\|_F^2 \right). \quad (3)$$

This loss function is quite intuitive, with each term corresponding to a layer-wise error as the factorizations unfold. In fact, this can be seen as a globalization of the model of Cichocki et al. [18]: instead of trying to minimize the errors of each layer-wise factorization sequentially, all of them are aggregated within a global weighted loss function.

As  $l$  increases and the ranks decrease (recall  $r_l < r_{l-1}$  for all  $l$ ), the computational cost to evaluate each term in (3) decreases. More precisely, the  $l$ -th term requires  $m r_l r_{l-1}$  elementary operations to be computed hence the computational cost of evaluating (3) is  $\mathcal{O}(L m n r_1)$ . The loss of the first layer, which corresponds to a standard factorization of the input matrix, is the most expensive to compute.

#### 3.2. Data-centric loss function

The second loss function considers the errors between  $X$  and its successive approximations of ranks  $r_l$ 's:

$$\mathcal{L}_2(H_1, H_2, \dots, H_L; W_1, W_2, \dots, W_L) = \frac{1}{2} \left( \|X - W_1 H_1\|_F^2 + \mu_1 \|X - W_2 H_2 H_1\|_F^2 + \dots + \mu_{L-1} \|X - W_L H_L \dots H_2 H_1\|_F^2 \right). \quad (4)$$

While the loss function described in Section 3.1 focuses on layer-wise errors, this one is data-centric in the sense that it evaluates the errors between the data matrix  $X$  and its successive low-rank approximations. An advantage of this loss function is that the parameters  $\mu_l$ 's ( $l = 1, \dots, L-1$ ) are easier to tune since all the terms are likely to have a similar order of magnitude. However, since the successive approximations involve an increasing number of matrix multiplications, this loss function and the associated update rules are slightly more computationally expensive. Indeed, the most computationally costly term is the last one, which requires  $n(m r_L + r_L r_{L-1} + \dots + r_2 r_1)$  operations to be computed, which may become high if the number of layers is high and the ranks do not decrease rapidly.

### 4. General algorithmic framework

A general algorithm to minimize the two proposed loss functions (3) and (4) under general constraints on  $W_l$ 's and  $H_l$ 's is given in Algorithm 3. We denote  $W_{\rightarrow l}$  the set of matrices  $\{W_1, \dots, W_{l-1}\}$  for any  $l = 1, \dots, L$  and  $W_{l \rightarrow}$  the set of matrices  $\{W_{l+1}, \dots, W_L\}$ , and similarly for the  $H_l$ 's.

**Algorithm 3** Framework to solve deep MF with general constraints and consistent global loss function.

**Input:** Data matrix  $X$ , number of layers  $L$ , inner ranks  $r_l$ 's, feasible sets  $\mathcal{W}_l$  and  $\mathcal{H}_l$  for  $l = 1, \dots, L$ , a global loss function  $\mathcal{L}$  such as  $\mathcal{L}_1$  in (3) or  $\mathcal{L}_2$  in (4)  
**Output:** Matrices  $W_1, \dots, W_L$  and  $H_1, \dots, H_L$   
1: Compute initial matrices  $W_l^{(0)}$  and  $H_l^{(0)}$  for all  $l$  through a sequential decomposition of  $X$  (for example Algorithm 1)  
2: **for**  $k = 1, \dots, L$  **do**  
3:   **for**  $l = 1, \dots, L$  **do**  
4:      $H_l^{(k)} = \underset{H \in \mathcal{H}_l}{\text{arg reduce}} \mathcal{L}(W_{\rightarrow l}^{(k)}, W_l^{(k-1)}, W_{l \rightarrow}^{(k-1)}; H_{\rightarrow l}^{(k)}, \mathbf{H}, H_{l \rightarrow}^{(k-1)})$   
5:      $W_l^{(k)} = \underset{W \in \mathcal{W}_l}{\text{arg reduce}} \mathcal{L}(W_{\rightarrow l}^{(k)}, \mathbf{W}, W_{l \rightarrow}^{(k-1)}; H_{\rightarrow l}^{(k)}, H_l^{(k)}, H_{l \rightarrow}^{(k-1)})$   
6:   **end for**  
7: **end for**

Algorithm 3 consists in BCD over the factors of each layer. The subproblems in one factor matrix (in bold at lines 4 and 5) can be solved by various well-known techniques. In particular, when the feasible set is convex, these subproblems are convex. Moreover, many standard optimization schemes can be applied to such formulations and directly provide convergence guarantees to stationary points, e.g., the proximal alternating linearized minimization (PALM) [21] and the block successive minimization (BSUM) [22] optimization schemes. This general framework is also very flexible and we present some examples below, considering usual constraints on the factors of each layer, which will be illustrated in the experiments of Sections 5.1 and 5.2.

One possibility to implement the updates of the factor matrices in Algorithm 3, that is, to solve the *arg reduce* subproblems at lines 4 and 5, is a fast projected gradient method (FPGM), which is the one considered in the remainder of this paper. This method is easy to implement and scales relatively well (linearly with the size of the data). The FPGM is a well-known first-order optimization framework to update a general matrix  $M$  which can be any of the  $W_l$ 's or  $H_l$ 's,  $l = 1, \dots, L$ , see Algorithm. As each subproblem is convex, we choose  $\frac{1}{L}$  as the step size, with  $L$  the Lipschitz constant, except for the update of  $H_l$ 's for the second loss function

$\mathcal{L}_2$ . Indeed, in this case, the Lipschitz constant is quite costly to compute, as the derivatives are obtained by a sum over the layers (see (8)) hence, we simply compute the stepsize through a backtracking line search. The extrapolation step is based on Nesterov acceleration [23] and a restart guarantees the decrease of the loss function. In other words, if the error increases, the extrapolation step is not taken, as in [24]. Moreover, such an algorithm is guaranteed to converge to stationary points of the corresponding optimization problem, under suitable assumptions (in particular, the convexity of the subproblems in each factor matrix, which is our case here as long as  $\mathcal{L}$  is block-wise convex, and the feasible set is convex) [25,26].

---

**Algorithm 4** Restarted fast projected gradient method (FPGM).

---

**Input:** Initial matrix  $M^{(0)}$ , feasible set  $\mathcal{M}$ , loss function  $f(M)$ , parameter  $\alpha_1 \in (0, 1)$

**Output:** A matrix  $M$  that decreases  $f$ , that is,  $f(M) < f(M^{(0)})$

```

1: Compute Lipschitz constant  $L$  of  $f$ ;  $Y = M^{(0)}$ 
2: for  $k = 1, \dots$  do
3:    $M^{(k)} = \mathcal{P}_{\mathcal{M}}(Y - \frac{1}{L} \nabla f(Y))$ 
4:    $Y = M^{(k)} + \beta_k (M^{(k)} - M^{(k-1)})$  with  $\beta_k = \frac{\alpha_k(1-\alpha_k)}{\alpha_k^2 + \alpha_{k+1}}$  and  $\alpha_{k+1} = \frac{1}{2} (\sqrt{\alpha_k^4 + 4\alpha_k^2 - \alpha_k^2})$ 
5:   if  $f(M^{(k)}) > f(M^{(k-1)})$  then
6:      $Y = M^{(k-1)}$ ,  $\alpha_{k+1} = \alpha_1$  % Restart
7:   end if
8: end for

```

---

Let us compute the gradients with respect to  $W_l$ 's and  $H_l$ 's ( $l = 1, \dots, L$ ) of the loss functions (3) and (4), by introducing  $W_0 = X$ ,  $\lambda_0 = \mu_0 = 1$ . For  $\mathcal{L}_1$ , we have:

$$\frac{\partial \mathcal{L}_1}{\partial W_l} = \lambda_{l-1} (W_l H_l - W_{l-1}) H_l^T + \delta_l \lambda_l (W_l - W_{l+1} H_{l+1}) \quad (5)$$

where  $\delta_l = 0$  if  $l = L$  and 1 otherwise,

$$\frac{\partial \mathcal{L}_1}{\partial H_l} = \lambda_{l-1} W_l^T (W_l H_l - W_{l-1}). \quad (6)$$

For  $\mathcal{L}_2$ , let us call for all  $l$ ,  $D_l = H_{l-1} \dots H_1$ ,  $\tilde{H}_l = H_l \dots H_1 = H_l D_l$  and  $C_l^{(k)} = W_k H_k \dots H_{l+1}$  for all  $k \geq l$  (for  $k = l$ ,  $C_l^{(l)} = W_l$ ). Then, the gradients are given by:

$$\frac{\partial \mathcal{L}_2}{\partial W_l} = \mu_{l-1} (W_l \tilde{H}_l - X) \tilde{H}_l^T, \quad (7)$$

$$\frac{\partial \mathcal{L}_2}{\partial H_l} = \sum_{k=l}^L \mu_{k-1} C_l^{(k)T} (C_l^{(k)} H_l D_l - X) D_l^T. \quad (8)$$

We now implement this FPGM-based framework for various deep MF's models that include constraints on the factors, such as non-negativity in Section 4.1 and sparsity in Section 4.2, or add a regularization term to the loss function, such as a volume penalization in Section 4.3.

#### 4.1. Nonnegative deep MF

In deep nonnegative MF (deep NMF), the factors of the decomposition (1) are constrained to be nonnegative, that is,  $W_l \geq 0$ ,  $H_l \geq 0$  for all  $l = 1, \dots, L$ . The projection operator  $\mathcal{P}$  at line of Algorithm simply consists of the projection on the nonnegative orthant, that is,  $\mathcal{P}(A) = \max(A, 0)$ .

Other constraints could be easily incorporated, as long as the projection onto the feasible set can be computed efficiently. For example, a very common additional constraint consists in enforcing the entries of every column of the  $H_l$ 's to sum to 1, that is

$\|H_l(:, j)\|_1 = 1$  for all  $j = 1, \dots, r_{l-1}$ , for any  $l$ . This expresses that the coefficients of the linear combination of basis vectors corresponding to each data point sum to 1, hence can be interpreted as proportions; see e.g., [27].

#### 4.2. Sparse deep MF

Sparse matrix factorizations consist in enforcing some factors of the decomposition to be sparse to foster their interpretability. Numerous ways of tackling sparsity in MFs have been proposed in the literature, including targeting a row-wise (or column-wise)  $l_1$  norm for some factors [28], adding a  $l_1$  and/or  $l_2$  norm penalty [29] to the loss function, dictionary learning [30] and sparse component analysis [31], among others.

Recently, an efficient and fast method, referred to as grouped sparse projection (GSP) [32], was developed to avoid the drawback of the methods mentioned beforehand, that is, the tuning of many parameters. Especially in the context of deep MF, where the number of factors to update grows linearly with the number of layers, it may be convenient to limit the number of parameters. Hence, for a given factor, GSP aims to reach a target average sparsity of the whole matrix instead of each row/column separately, which confers much more flexibility to the sparsity pattern compared to standard approaches.

In Algorithm, it suffices to consider at line the grouped sparse projection given in Algorithm 1 of [32]. However, the feasible set of GSP is not convex, hence the convergence is not guaranteed when using the FPGM.

#### 4.3. Minimum-volume deep NMF

Minimum-volume one-layer NMF (minVolNMF) is a well-known NMF variant [33–35] that encourages the basis vectors, that is, the columns of  $W$ , to have a small volume. Intuitively, this boils down to trying to make them as close as possible to the data points, which enhances the interpretability of the decomposition. The minVolNMF model on which we will focus adds a penalty term for the volume to the reconstruction error in the loss function and is expressed as:

$$\min_{\substack{W \in \mathbb{R}^{m \times r}, W^T \mathbf{e} = \mathbf{e} \\ H \in \mathbb{R}_+^{r \times n}}} \frac{1}{2} (\|X - WH\|_F^2 + \kappa \log \det(W^T W + \delta I)) \quad (9)$$

where  $\kappa$  and  $\delta$  are parameters fixed by the user, and  $\mathbf{e}$  is the column vector of all ones. It is important to note that, while many models consider the intuitive constraint  $H^T \mathbf{e} = \mathbf{e}$ , more recent approaches showed that imposing the column-stochasticity of  $W$ , that is,  $W^T \mathbf{e} = \mathbf{e}$ , instead of  $H$  leads to better results in many applications. This is due to the better conditioning of  $W$  in this case, see the discussion in Section 4.3.3 of [6] for more details.

To the best of our knowledge, minVolNMF has not been extended to the deep context yet. We extend the approach of [35] by incorporating a volume contribution at every layer to (3) and (4). Hence, we add the following quantity to each term of both loss functions: for  $l = 1, 2, \dots, L$ ,

$$\kappa_l \log \det(W_l^T W_l + \delta I_{r_l}) \quad (10)$$

while imposing column-stochasticity on every  $W_l$ . To solve the minVolNMF problem, a majorization-minimization (MM) framework is usually considered. This consists in minimizing a surrogate function, namely a strongly convex upper approximation of the loss function, see [35] and [36] for the details. The FPGM of Algorithm can then be applied on this surrogate.

The volume contribution implies an additional term in the gradients of both  $\mathcal{L}_1$  and  $\mathcal{L}_2$  w.r.t.  $W_l$ 's. More precisely, the term  $\kappa_l W_l Z$ ,

multiplied either by  $\lambda_{l-1}$  or  $\mu_{l-1}$  is added to (5) and (7) respectively, with  $Z = (W_l^{(*)T} W_l^{(*)} + \delta I_r)^{-1}$  where  $W_l^{(*)}$  denotes the last iteration of  $W_l$  such that  $Z$  is constant during a given update of  $W_l$ .

A potential drawback of such an approach is the use of many regularization parameters, both for the weights of the linear combination of errors and the volume penalties at each layer. Indeed, in addition to the  $L - 1$  parameters  $\lambda_l$ 's of (3) or  $\mu_l$ 's of (4), the user has to fix the values of the  $L$  parameters  $\kappa_l$ 's involved in the volume regularization at each layer. In practice, the  $\kappa_l$  for a given layer is set as follows. Given the initial error  $err_l^{(0)} = \frac{1}{2} \|W_{l-1}^{(0)} - W_l^{(0)} H_l^{(0)}\|_F^2$  (for the layer-centric loss function (3)) or  $err_l^{(0)} = \frac{1}{2} \|X - W_l^{(0)} \tilde{H}_l^{(0)}\|_F^2$  (for the data-centric loss function (4)) and a first guess  $\tilde{\kappa}_l$ , the final value  $\kappa_l$  is given by  $\kappa_l = \tilde{\kappa}_l \frac{err_l^{(0)}}{|\log \det(W_l^{(0)T} W_l^{(0)} + \delta I_r)|}$ , such that the decomposition error and the volume term are of the same order of magnitude for a given layer.

### 5. Numerical experiments

In this section, we evaluate the models described in Section 4 on both synthetic (Section 5.1) and real (Section 5.2) data. One drawback of the experiments carried on deep MF models is the lack of ground truth in datasets. Indeed, despite the hierarchical structure of many datasets, few of them have available ground truth at each layer. Moreover, it is generally hard to guess in advance how features of a given layer can be interpreted, especially when various constraints are applied on the factors. Even with synthetic data, the setting should be chosen carefully to guarantee easy interpretability of the factors. For these reasons, we first present the results of our models on a simple yet meaningful toy example in dimension  $m = 3$ . The low dimensionality offers the advantage to control the ground truth basis vectors and easily interpret the features at each layer. We then study the performance of our models for hyperspectral unmixing and the extraction of facial features to show their efficiency on real-world challenges.

A Matlab implementation of the framework described above, with all the experiments, is available on <https://bit.ly/flexDeepMF>.

#### 5.1. Synthetic data

We consider for all the experiments a 2-layers network (that is,  $L = 2$ ) in dimension  $m = 3$ . The ranks  $r_l$ 's are set to  $r_1 = 6, r_2 = 3$ ,

and the target basis matrices  $W_1^*$  and  $W_2^*$  are given by

$$W_1^* = \begin{pmatrix} 0.1 & 0.1 & 0.4 & 0.4 & 0.5 & 0.5 \\ 0.4 & 0.5 & 0.1 & 0.5 & 0.1 & 0.4 \\ 0.5 & 0.4 & 0.5 & 0.1 & 0.4 & 0.1 \end{pmatrix} = W_2^* H_2^*,$$

where  $W_2^* = \begin{pmatrix} 1/2 & 0 & 1/2 \\ 0 & 1/2 & 1/2 \\ 1/2 & 1/2 & 0 \end{pmatrix}$  and  $H_2^* = \begin{pmatrix} 0.2 & 0 & 0.8 & 0 & 0.8 & 0.2 \\ 0.8 & 0.8 & 0.2 & 0.2 & 0 & 0 \\ 0 & 0.2 & 0 & 0.8 & 0.2 & 0.8 \end{pmatrix}$ .

Each column of the matrix  $H_2^*$  is generated according to a Dirichlet distribution of parameter  $\alpha = 0.05$ . The data matrix  $X$ , made of  $n = 1000$  points, is therefore generated as  $X = X^* + N$  where  $X^* = W_1^* H_1^*$  and  $N$  is additive Gaussian noise:  $N = \epsilon \|X^*\|_F \frac{Y}{\|Y\|_F}$  with  $Y \sim \mathcal{N}(0, 1)$ .

An example of a data set generated in that way in the noiseless case ( $\epsilon = 0$ ) is presented on Fig. 1, with the ground truth basis vectors at both layers. In the following, we consider 10 levels of noise:  $\epsilon = 10^{-2}, 2.51 \cdot 10^{-2}, 6.31 \cdot 10^{-2}, 9.49 \cdot 10^{-2}, 1.267 \cdot 10^{-1}, 1.585 \cdot 10^{-1}, 2.384 \cdot 10^{-1}, 3.182 \cdot 10^{-1}, 3.981 \cdot 10^{-1}, 1$ .

We compare five models, namely:

- Single-layer NMF of ranks  $r_l, l = 2, \dots, L$ ,
- The sequential multilayer MF of [18], as described in Algorithm 1, dubbed MMF. At the first layer, the solution of single-layer NMF corresponds to the one of MMF, by construction.
- The deep MF model from [19], see Algorithm 2, dubbed Tri-DMF. Although the updates of the original paper are performed with multiplicative updates (MU), we decided to solve the sub-problems at lines and with FPGM since MU are known to be slow.
- Deep MF with the layer-centric loss function (3) solved with Algorithms 3 and , dubbed LC-DMF,
- Deep MF with the data-centric loss function (4) solved with Algorithms 3 and , dubbed DC-DMF.

The global loss functions of these methods are not meaningfully comparable to each other due to, on the one hand, the absence of global loss function for multilayer MF and, on the other hand, the difference of magnitude between the terms of layer-centric and data-centric loss functions (see Section 3). Therefore, we report the average mean-removed spectral angle (MRSa) between the corresponding expected and computed basis vectors, that is, the columns of  $W_1^*$  and  $W_l$  respectively at each layer  $l$ . Given any

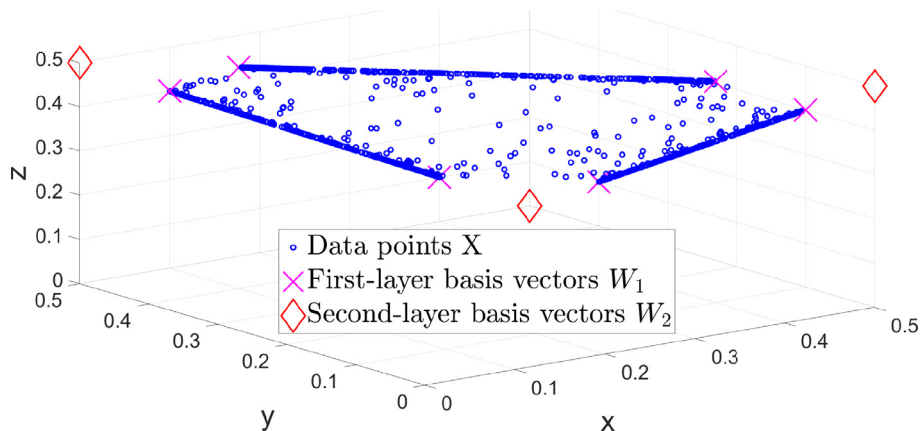


Fig. 1. Setting of the synthetic data considered in this section, in the noiseless case.




Fig. 2. Comparison of the MRS obtained at the first layer with minVol MMF, LC-DMF, DC-DMF and Tri-DMF on synthetic data in function of the noise level.




Fig. 3. Comparison of the MRS obtained at the second layer with minVol MMF, LC-DMF, DC-DMF, Tri-DMF and single-layer NMF on synthetic data in function of the noise level.

two vectors  $a$  and  $b$ , their MRS is defined as

$$MRS(a, b) = \frac{100}{\pi} \arccos\left(\frac{\langle a - \bar{a}, b - \bar{b} \rangle}{\|a - \bar{a}\|_2 \|b - \bar{b}\|_2}\right) \in [0, 100]$$

where  $\langle \cdot, \cdot \rangle$  indicates the scalar product of two vectors and  $\bar{\cdot}$  is the mean of a vector.

For all methods, the initial factors  $W_l^{(0)}$  and  $H_l^{(0)}$  are obtained by projecting onto the feasible sets the output of the successive nonnegative projection algorithm (SNPA) [37] applied on  $W_{l-1}^{(0)}$ , with  $W_0^{(0)} = W_0 = X$ . In a nutshell, SNPA is a column subset selection algorithm often used as an initialization technique for NMF.

**Remark 1** (Compared algorithms and goal of our experiments). Our main goal in this paper is to show that our new proposed loss functions, (3) and (4), are much more meaningful and lead to significantly better results in practice than the most widely used one, namely, that of Trigeorgis et al. [19], that is, Tri-DMF. In fact, to the best of our knowledge, most papers on deep MF rely on Tri-DMF. Their novelty is typically embedded into the regularizers and/or constraints added into their model to tackle specific applications: for example, a total variation regularizer for hyperspectral unmixing in [7], a regularizer to enhance good local characteristics for basis image extraction in [9], or a so-called community regularization for recommendation in social networks in [11]. We do not compare to such more recent deep MF models because, to have a meaningful and fair comparison, we would need to adapt our model as well, adding proper regularizers. Moreover,

we would need to provide background on each particular application, which is out of the scope of this paper. In summary, because our main goal is to provide theoretical and experimental evidence of the superiority of our proposed models, and the failure of the model of Trigeorgis et al. [19], we do not focus on particular models designed for specific applications and try to be as generic as possible. The regularizations and constraints we consider in this paper are very general and apply to a wide range of applications.

*MinVol deep NMF* We first consider the minVol deep NMF variant (minVol regularization together with non-negativity of the factors). The parameter  $\delta$  is fixed to  $\delta = 0.1$ . For the  $\tilde{\kappa}_l$ 's (see Section 4.3), we make a distinction between the first 4 levels of noise ( $\epsilon < 0.1$ ) and the last 6 ( $\epsilon > 0.1$ ), since minimizing the volume is more challenging when the noise increases. More precisely, we fix  $\tilde{\kappa}_1 = 10^{-3}$ ,  $\tilde{\kappa}_2 = 10^{-2}$  for the four first levels and  $\tilde{\kappa}_1 = 10^{-2}$ ,  $\tilde{\kappa}_2 = 10^{-1}$  for the six last levels of noise, for all the compared methods. To compute the parameters  $\lambda_l$ 's,  $l = 1, \dots, L - 1$  in (3), we proceed similarly to what is done for the minVol parameters  $\kappa_l$ 's, by always considering the first layer of decomposition as a baseline. More precisely, based on an initial guess  $\tilde{\lambda}_l$ , we set  $\lambda_l = \tilde{\lambda}_l \frac{err_k^{(0)}}{err_{l+1}^{(0)}}$ , where  $err_k^{(0)}$  denotes the  $k$ -th layer error  $\frac{1}{2} \|W_{k-1} - W_k H_k\|_F^2$  after the initialization. By doing so, the ratio between the  $(l + 1)$ -th and the first term of (3) is approximately equal to an arbitrary value  $\tilde{\lambda}_l$ , fixed by the user. In practice, we used  $\tilde{\lambda}_l = 10$




Fig. 4. Comparison of the MRSa obtained at the first layer with grouped sparse MMF, LC-DMF, DC-DMF and Tri-DMF on synthetic data in function of the noise level.

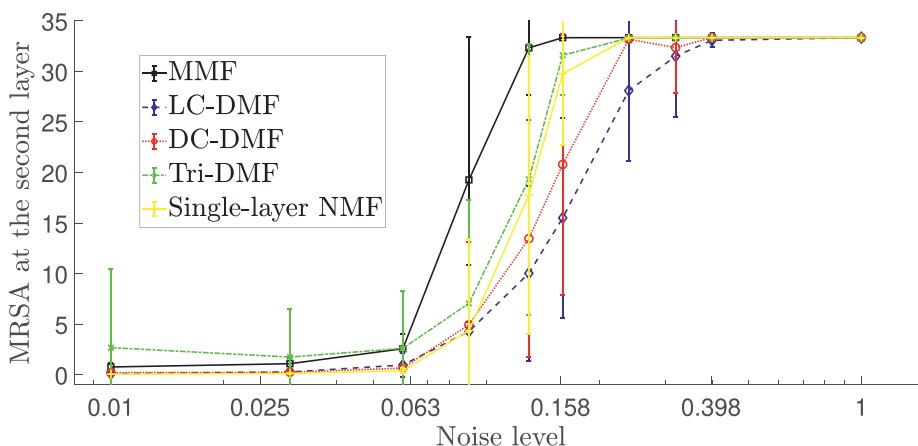


Fig. 5. Comparison of the MRSa obtained at the second layer with grouped sparse MMF, LC-DMF, DC-DMF, Tri-DMF and single-layer NMF on synthetic data in function of the noise level.

for all  $l = 1, \dots, L - 1$ . For the parameters  $\mu_l$ 's of the data-centric loss function, we fixed  $\mu_l = 1$  for all  $l = 1, \dots, L - 1$  since all the reconstruction error terms are expected to be of the same order of magnitude. We use the same values of these parameters for the experiments on real data in Section 5.2.

Fig. 2 and 3 display the MRSa (average and standard deviation over 25 runs) of MMF, LC-DMF, DC-DMF, Tri-DMF and single-layer NMF (for the second layer only) in function of the noise level for the first and second layers, respectively. For both layers, LC-DMF produces the lowest MRSa, except for very high noise levels. DC-DMF becomes more competitive as the noise increases. On the contrary, especially at the second layer, MMF produces a higher MRSa than DC-DMF and LC-DMF. This confirms that a weighted loss function together with an iterative framework such as Algorithm 3 is more efficient than the purely sequential approach of Cichocki et al. On the other side, the approach of Trigeorgis et al. completely fails to retrieve the correct basis vectors at the first layer, with one or two predicted basis vectors located inside the convex hull of the others. Even at the second layer, Tri-DMF is not the best method although the corresponding loss function is designed to minimize the reconstruction error at the last layer. Finally, single-layer NMF also performs worse than LC-DMF at the second layer and, of course, does not allow to automatically bind the features of consecutive layers, which is undoubtedly the main added value of deep approaches.

*Sparse deep MF* Let us now investigate sparse deep MF. Due to the structure of the ground truth factors, we only considered sparsity on the factors of the second layer, fixing the target grouped




Fig. 6. Urban hyperspectral image and its four main materials.

sparsity of both  $W_2$  and  $H_2$  to  $\frac{1}{3}$ . The values of the parameters  $\lambda_l$ 's and  $\mu_l$ 's are chosen in the same way as for the minVol variant.

Figs. 4 and 5 display the MRSa of MMF, LC-DMF, DC-DMF, Tri-DMF and single-layer NMF in function of the noise level for the first and second layers, respectively. The conclusions are similar to those of minVol deep MF: LC-DMF performs better in terms

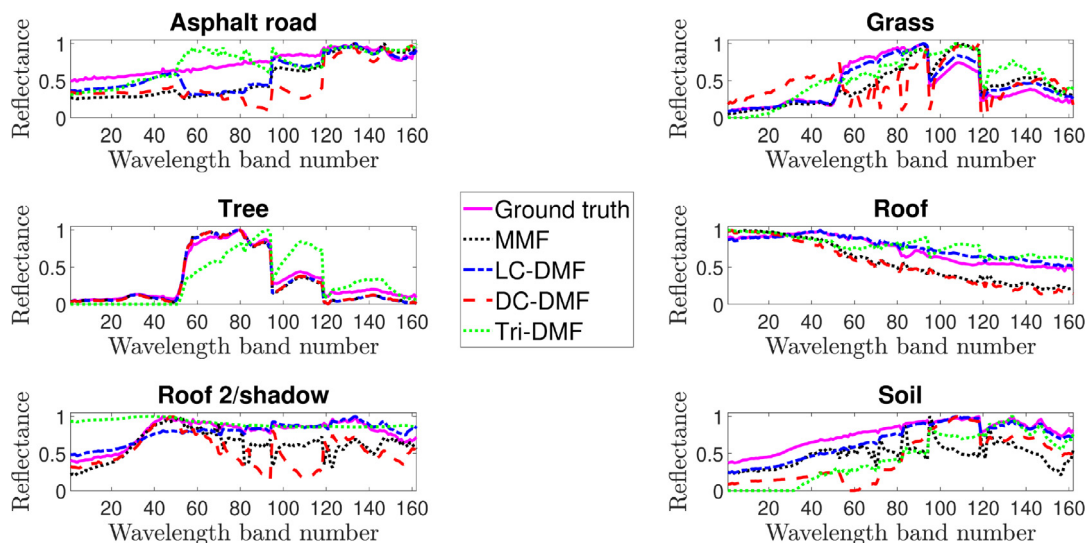


Fig. 7. Endmembers extracted by MMF, LC-DMF, DC-DMF and Tri-DMF in the Urban image at the first layer ( $r_1 = 6$ ), and the ground truth.




Fig. 8. Endmembers extracted by MMF, LC-DMF, DC-DMF, Tri-DMF and single-layer NMF in the Urban image at the second layer ( $r_2 = 4$ ), and the ground truth.

of MRSA than DC-DMF and MMF, and Tri-DMF completely fails to recover the basis vectors at the first layer. Moreover, at the second layer, single-layer NMF performs worse than weighted deep approaches such as LC-DMF and DC-DMF. This tends to confirm that LC-DMF is the most effective to recover the ground truth factors.

### 5.2. Real data

In this section, we present the performance of our framework on real data, namely hyperspectral images in Section 5.2.1 and faces in Section 5.2.2.

#### 5.2.1. Hyperspectral unmixing

A hyperspectral image (HI) is characterized by the reflectance values of  $n$  pixels in  $m$  wavelength spectral bands and is generally represented by a matrix  $X \in \mathbb{R}^{m \times n}$  where each column of  $X$  is the spectral signature of one pixel. Hyperspectral unmixing (HU) aims to identify the spectral signatures of  $r$  materials and under the linear mixing assumption, NMF has been widely used to solve HU [38]. When deep NMF is applied, the materials are extracted in a hierarchical manner [39,40].

We apply minVol deep NMF on the HYDICE Urban HI, which is made of  $n = 307 \times 307$  pixels in  $m = 162$  spectral bands [41]; see Fig. 6.

We consider a 3-layers network, with  $r_1 = 6$ ,  $r_2 = 4$  and  $r_3 = 2$ . To initialize the basis vectors of all layers, we use hierarchical clustering (HC) [42] instead of SNPA. Indeed, SNPA is designed to extract extreme points of the dataset which may not be the most


appropriate for the noisy, high-dimensional hyperspectral data. On the other hand, HC extracts clusters centroids. After several trials, it turns out that appropriate values for the minVol hyperparameters  $\tilde{\kappa}_l$ 's are  $10^{-2}$  for all  $l$ .

As the hyperspectral data are high-dimensional and tough to unmix, instead of simply initializing the factors of LC-DMF, DC-DMF and Tri-DMF with HC, we perform a few iterations of BCD after running HC at each layer, similarly to MMF (see Algorithm 1), before moving to the initialization of the next layer. After this "augmented" initialization, the framework of Algorithm 3 is applied. More precisely, we run the same number  $it = 500$  iterations of MMF, LC-DMF, DC-DMF and Tri-DMF, among which, for the last three methods,  $it_{in} = 50$  iterations are devoted to improve the initial factors of all layers and the remaining consist in the iterative updates of Algorithm 3. We also run single-layer NMF for  $r = 4$  and  $r = 2$  (at the first layer, it coincides with MMF) to evaluate the efficiency of deep approaches compared to a shallow one.


On Fig. 7 and 8, we plot the spectral signatures of the materials extracted by the pre-cited methods at the first and second layer respectively, that is, the columns of  $W_1$  and  $W_2$ , and the ground truth from [41], which is only available for  $r = 6$  and  $r = 4$  to the best of our knowledge. The MRSA's at both layers are presented in Table 2. At both layers, LC-DMF outperforms the other deep methods, including Tri-DMF and MMF. Moreover, at the second layer, LC-DMF achieves a MRSA very close to the one of single-layer NMF. Therefore, as for the synthetic data, LC-DMF seems to be the best loss function to minimize when tackling deep MF.

In Appendix A, we present the abundance maps indicating the proportion of every material in each pixel at each layer, ex-





**Fig. 9.** Features extracted by MMF on the CBCL face data set, with  $L = 3$ ,  $r_1 = 100$ ,  $r_2 = 49$ , and  $r_3 = 25$ . Each image contains the features extracted at a layer: (a) first layer  $W_1$ , (b) second layer  $W_2$ , and (c) third layer  $W_3$ .



**Fig. 10.** Layer-centric relative errors on the CBCL face data set, with  $L = 3$ ,  $r_1 = 100$ ,  $r_2 = 49$  and  $r_3 = 25$  at the (a) first, (b) second, and (c) third layer.

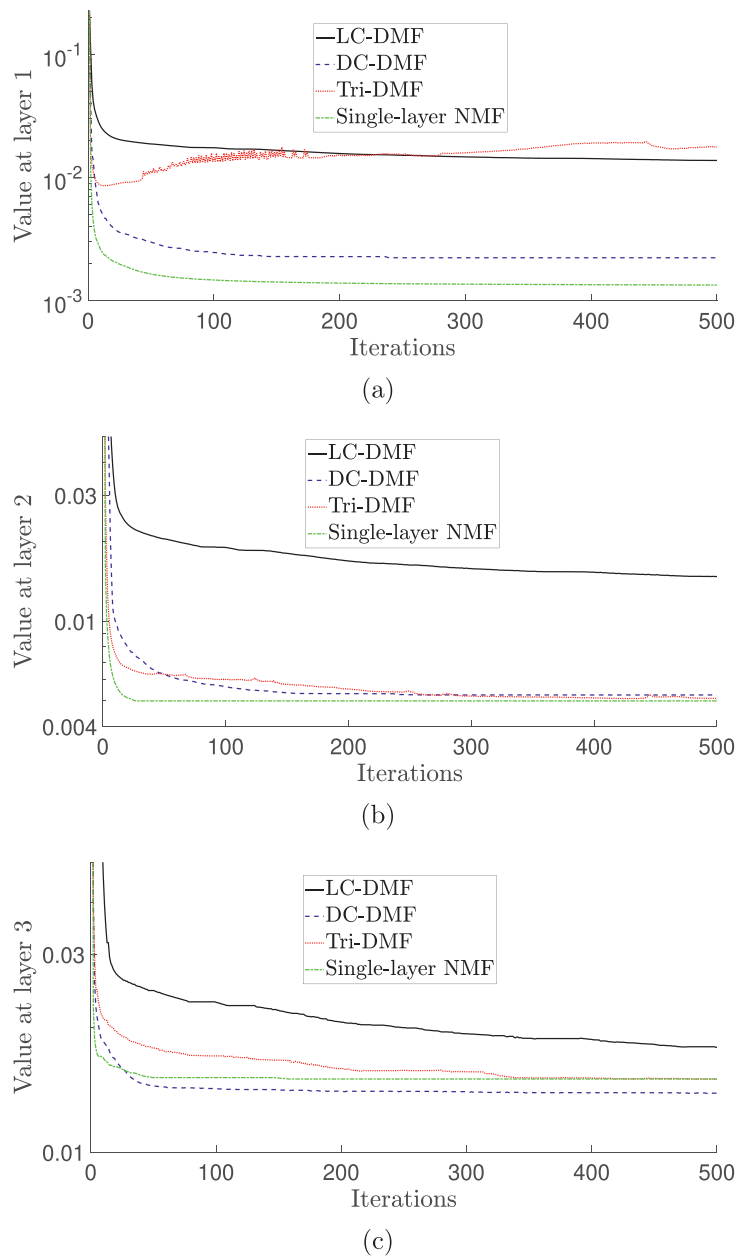


Fig. 11. Data-centric relative errors on the CBCL face data set, with  $L = 3$ ,  $r_1 = 100$ ,  $r_2 = 49$  and  $r_3 = 25$  at the (a) first, (b) second, and (c) third layer.

Table 2

MRSA of the compared methods at the first and second layer on the Urban hyperspectral image, with in bold the best value of each column.

Method	First layer ( $r_1 = 6$ )	Second layer ( $r_2 = 4$ )
MMF	16.98	12.35
LC-DMF	<b>9.48</b>	8.42
DC-DMF	22.95	14.15
Tri-DMF	26.07	20.07
Single-layer NMF	16.98	<b>7.74</b>

tracted by MMF, LC-DMF, DC-DMF, Tri-DMF and single-layer NMF on Figs. A1, A2, A3, A4, A5 and A6 respectively. In general, the six materials extracted at the first layer are asphalt road, grass, tree, two types of roof tops, and soil. Finally, at the last layer, two main categories of materials remain: on the one hand, grass and trees are mainly merged in a single "vegetal" cluster while on the other hand, the other materials are combined in a second "non-vegetal"

cluster. On the figures, the arrows between materials extracted at consecutive layers indicate which materials of a given layer contribute to those of the next layer (that is, the arrows represent the non-zero entries of the corresponding  $H_l$ ). When a material of the upper layer contributes to less than 10% to a material of the lower layer, the arrow is discarded. We observe that the hierarchy of materials extracted by LC-DMF is rather sparse, as each material is obtained by a combination of only a few materials of the previous layer. On the opposite, the decomposition of Tri-DMF is hardly interpretable since all materials contribute almost equally to those of the next layer.

### 5.2.2. Facial features extraction

Deep MF has also been shown efficient to extract facial features hierarchically [18,19]. The CBCL image<sup>1</sup> is made of 2429 grey-scale

<sup>1</sup> <http://www.ai.mit.edu/courses/6.899/lectures/faces.tar.gz>

images of  $19 \times 19$  pixels representing faces of different people exhibiting various expressions.

For this application, we consider grouped sparsity constraints on the factors. We set  $L = 3$ ,  $r_1 = 100$ ,  $r_2 = 49$  and  $r_3 = 25$ , and perform the initialisation of the factors with SNPA. We run MMF, LC-DMF, DC-DMF and Tri-DMF with a grouped sparsity level, that is, the average Hoyer sparsity of the columns, of 70% on  $W_1$ , 80% on  $W_2$  and 85% on  $W_3$ . Given any column vector  $x$ , its Hoyer sparsity [28] is defined as  $sp(x) = \frac{\sqrt{n} - \|x\|_1}{\sqrt{n-1} \|x\|_2}$ . In comparison, the average Hoyer sparsity of  $W_1$ ,  $W_2$  and  $W_3$  obtained with MMF without sparsity constraints are respectively 71.97%, 79.46% and 76.64%. The features extracted at all layers by MMF are displayed on Fig. 9. It shows the hierarchical decomposition of the data set by MMF: the larger facial features extracted at the first levels are made of smaller ones extracted at the deeper levels, such as eyes, mouths, and eyebrows. Sparsity allows the features, especially at the last layer, to contain only a few activated pixels. The features extracted by the other methods, that is, LC-DMF, DC-DMF, Tri-DMF and single-layer NMF are displayed in Appendix B on Figs. B1, B2, B3 and B4, respectively.

To compare quantitatively these methods, we plot the evolution of several loss functions along the 500 iterations. More precisely, Fig. 10 shows the relative layer-centric errors  $\frac{\|W_{l-1} - W_l H_l\|_F^2}{\|W_{l-1}^{(0)}\|_F^2}$  for  $l = 1, 2, 3$ , with  $W_0 = X$ , for all methods. Fig. 11 shows the relative data-centric errors  $\frac{\|X - W_l H_{1..l}\|_F^2}{\|X\|_F^2}$  for  $l = 1, 2, 3$ . Note that the first layer errors are both equal to  $\frac{\|X - W_1 H_1\|_F^2}{\|X\|_F^2}$  and appear in all global loss functions except Tri-DMF.

This experiment confirms the advantage of DC-DMF and LC-DMF over MMF and Tri-DMF:

- At the second and third layers, DC-DMF produces the lowest data-centric errors while LC-DMF produces the lowest layer-centric errors. Note however that DC-DMF has slightly larger errors than single-layer NMF at the first two layers, while it has a lower error at the third layer. This is expected since DC-DMF optimizes all layers simultaneously while single-layer NMF simply performs independent NMFs, that is, DC-DMF provides a hierarchy of intricate features.
- MMF has much higher relative errors than LC-DMF at the second and third layers (MMF is above  $10^{-1}$  while LC-DMF is below or about  $10^{-3}$ ). This comes from the sequential optimization procedure of MMF. More precisely, the factors of the first layer,  $W_1$  and  $H_1$ , are first extracted, then those of the second layer, and so on, without any possibility of “backpropagation”, see Section 2 for details.
- The first-layer error of Tri-DMF oscillates and does not converge (in fact, it appears to diverge), which is an expected consequence of the different loss functions minimized at each layer; see the discussion of Section 2 for more details.

## 6. Conclusion

In this paper, we have discussed the choice of the loss functions in deep MF models. Our motivation came from the fact that the mainstream framework proposed by Trigeorgis et al. [19], namely the loss function (2) and the corresponding Algorithm 2, is not consistent as it optimizes different losses at different layers. In

fact, we have shown that this approach leads to poor feature extraction, and divergence of the loss functions. We have therefore proposed two loss functions for deep MF that naturally weighs the different layers, namely a layer-centric loss and a data-centric loss, see (3) and (4), along with a simple block coordinate descent framework to compute the factors in the corresponding deep MF models. We focused on nonnegativity and sparsity constraints, and proposed a new deep MF model relying on minimum-volume regularization. We showed through experiments on synthetic and real data that our weighted loss functions allow us to outperform sequential MF, single-layer MF and the mainstream deep MF model from [19], while offering flexibility when various constraints or regularizations are used. In particular, we recommend to use the loss function defined in (3), that is, LC-DMF, to tackle deep MF as it performed best on average in our numerical experiments. Besides, it is computationally much cheaper than DC-DMF as the dimension reduces as the factorization unfolds, while DC-NMF uses the full data set at each layer (see Section 3 for the details). However, the best alternative will depend on the application and the data set at hand, and it is hard to know in advance which model will be the most appropriate. This also depends on the goal of the end user: for example, if the goal is to minimize the error w.r.t. the input data at all layers, then DC-DMF should be preferred.

An important direction of research is to find clever ways of choosing and tuning the regularization parameters in our proposed loss functions. Moreover, other loss functions could be investigated, as variants of the two proposed ones. An other perspective is to embed deep MF in a more powerful optimization framework such as TITAN [43] which has proven to be particularly efficient to tackle non-smooth non-convex problems, such as the grouped sparse variant of deep MF (see Section 4.2). Studying the identifiability, that is, the uniqueness of the factors retrieved by these models is also an important issue, which has not been investigated much for deep MF, except for some quite specific settings [44,45]. Finally, applying the different models on other applications, such as topic modeling, would also be insightful.

## Declaration of Competing Interest

None.

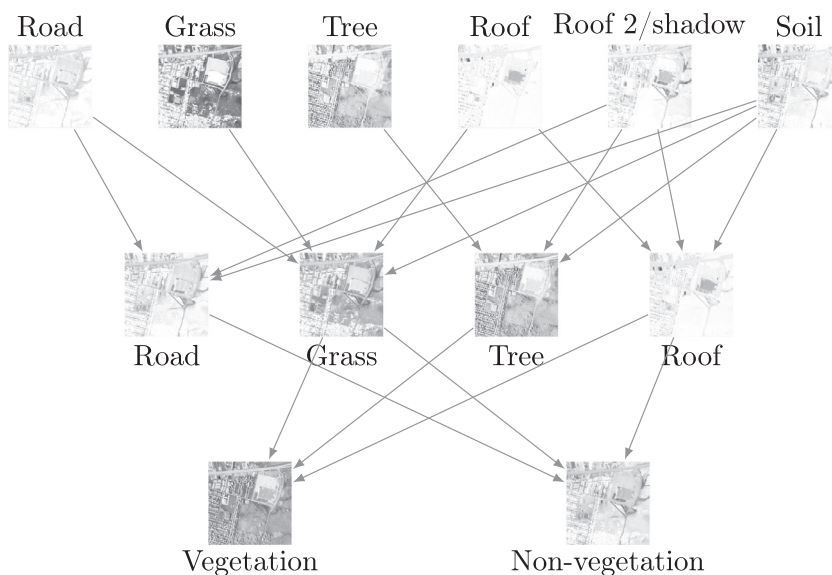
## Acknowledgement

We thank the reviewers for their insightful comments that helped us improve the paper.


This work was supported by the Fonds de la Recherche Scientifique - FNRS (F.R.S.-FNRS) and the Fonds Wetenschappelijk Onderzoek - Vlaanderen (FWO) under EOS Project no O005318F-RG47, and by the Francqui foundation. Pierre De Handschutter is a research fellow of the F.R.S.-FNRS. The authors claim no conflict of interest.

## Appendix A. Abundance maps of the Urban hyperspectral image

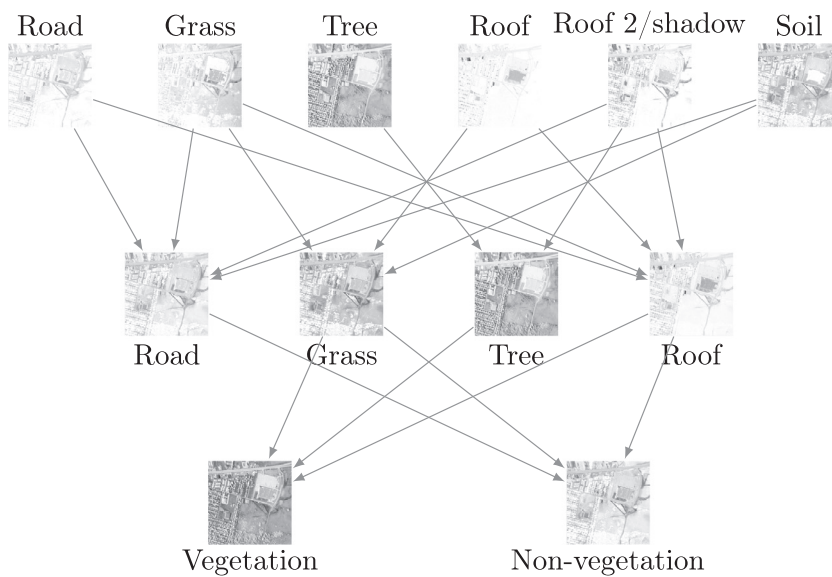
Figs. A1-A6 provide the abundance maps obtained with MMF, LC-DMF, DC-DMF, Tri-DMF and single-layer NMF applied on the Urban image for factorizations of depth  $L = 3$ , with  $r_1 = 6$ ,  $r_2 = 4$ ,  $r_3 = 2$  (see Section 5.2.1).



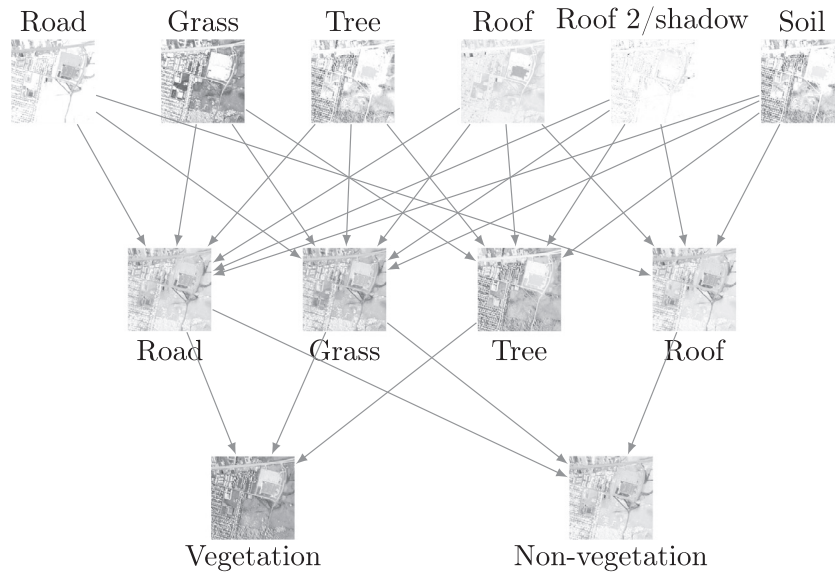
**Fig. A1.** Hierarchy of features extracted by MMF on the Urban hyperspectral image with  $L = 3$  layers,  $r_1 = 6$ ,  $r_2 = 4$ ,  $r_3 = 2$ .




**Fig. A2.** Hierarchy of features extracted by LC-DMF on the Urban hyperspectral image with  $L = 3$  layers,  $r_1 = 6$ ,  $r_2 = 4$ ,  $r_3 = 2$ .



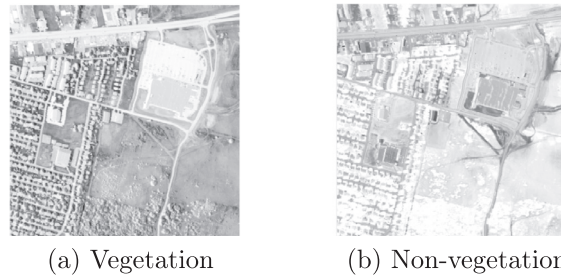
**Fig. A3.** Hierarchy of features extracted by DC-DMF on the Urban hyperspectral image with  $L = 3$  layers,  $r_1 = 6$ ,  $r_2 = 4$ ,  $r_3 = 2$ .



**Fig. A4.** Hierarchy of features extracted by Tri-DMF on the Urban hyperspectral image with  $L = 3$  layers,  $r_1 = 6$ ,  $r_2 = 4$ ,  $r_3 = 2$ .



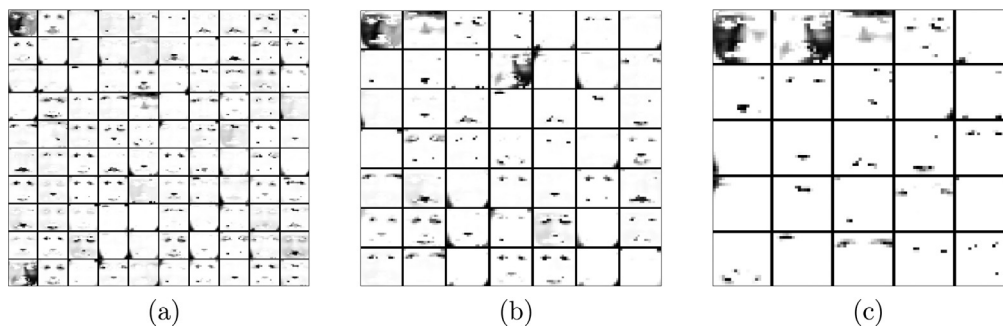
**Fig. A5.** Features extracted by single-layer NMF with  $r = 4$  on the Urban hyperspectral image.




**Fig. A6.** Features extracted by single-layer NMF with  $r = 2$  on the Urban hyperspectral image.

### Appendix B. Facial features of the CBCL dataset


In this section, we present the facial features extracted by LC-DMF, DC-DMF, Tri-DMF and single-layer NMF with grouped sparsity constraints on the CBCL dataset (see Section 5.2.2).




**Fig. B1.** Features extracted by LC-DMF on the CBCL face data set, with  $L = 3$ ,  $r_1 = 100$ ,  $r_2 = 49$ , and  $r_3 = 25$ . Each image contains the features extracted at a layer: (a) first layer  $W_1$ , (b) second layer  $W_2$ , and (c) third layer  $W_3$ .



**Fig. B2.** Features extracted by DC-DMF on the CBCL face data set, with  $L = 3$ ,  $r_1 = 100$ ,  $r_2 = 49$ , and  $r_3 = 25$ . Each image contains the features extracted at a layer: (a) first layer  $W_1$ , (b) second layer  $W_2$ , and (c) third layer  $W_3$ .



**Fig. B3.** Features extracted by Tri-DMF on the CBCL face data set, with  $L = 3$ ,  $r_1 = 100$ ,  $r_2 = 49$ , and  $r_3 = 25$ . Each image contains the features extracted at a layer: (a) first layer  $W_1$ , (b) second layer  $W_2$ , and (c) third layer  $W_3$ .



**Fig. B4.** Features extracted by single-layer NMF on the CBCL face data set, with (a)  $r = 49$  (b)  $r = 25$ .

## References

- [1] D.D. Lee, H.S. Seung, Learning the parts of objects by non-negative matrix factorization, *Nature* 401 (1999) 788–791.
- [2] A. d'Aspremont, L. Ghaoui, M. Jordan, G. Lanckriet, A direct formulation for sparse PCA using semidefinite programming, *Adv Neural Inf Process Syst* 17 (2004).
- [3] H. Zou, T. Hastie, R. Tibshirani, Sparse principal component analysis, *J. Comput. Graph. Stat.* 15 (2) (2006) 265–286.
- [4] R. Gribonval, K. Schnass, Dictionary identification sparse matrix-factorization via  $\ell_1$ -minimization, *IEEE Trans. Inf. Theory* 56 (7) (2010) 3523–3539.
- [5] X. Fu, K. Huang, N.D. Sidiropoulos, W.-K. Ma, Nonnegative matrix factorization for signal and data analytics: identifiability, algorithms, and applications, *IEEE Signal Process. Mag.* 36 (2) (2019) 59–80.
- [6] N. Gillis, *Nonnegative Matrix Factorization*, SIAM, Philadelphia, 2020.
- [7] X.-R. Feng, H.-C. Li, J. Li, Q. Du, A. Plaza, W.J. Emery, Hyperspectral unmixing using sparsity-constrained deep nonnegative matrix factorization with total variation, *IEEE Trans. Geosci. Remote Sens.* 56 (10) (2018) 6245–6257.
- [8] A. Cichocki, R. Zdunek, Multilayer nonnegative matrix factorization using projected gradient approaches, *Int J Neural Syst* 17 (06) (2007) 431–446.
- [9] Y. Zhao, H. Wang, J. Pei, Deep non-negative matrix factorization architecture based on underlying basis images learning, *IEEE Trans. Pattern Anal. Mach.* 43 (6) (2019) 1897–1913.
- [10] Y. Zhao, F. Deng, J. Pei, X. Yang, Progressive deep non-negative matrix factorization architecture with graph convolution-based basis image reorganization, *Pattern Recognit* (2022) 108984.
- [11] L. Wan, F. Xia, X. Kong, C.-H. Hsu, R. Huang, J. Ma, Deep matrix factorization for trust-aware recommendation in social networks, *IEEE Trans. Network Sci. Eng.* 8 (1) (2020) 511–528.
- [12] H. Zhao, Z. Ding, Y. Fu, Multi-view clustering via deep matrix factorization, in: *AAAI Conference on Artificial Intelligence*, 2017.
- [13] S. Huang, Z. Kang, Z. Xu, Auto-weighted multi-view clustering via deep matrix decomposition, *Pattern Recognit* 97 (2020) 107015.
- [14] K. Luong, R. Nayak, T. Balasubramaniam, M.A. Bashir, Multi-layer manifold learning for deep non-negative matrix factorization-based multi-view clustering, *Pattern Recognit* (2022) 108815.
- [15] P. De Handschutter, N. Gillis, X. Siebert, A survey on deep matrix factorizations, *Computer Science Review* 42 (2021) 100423.
- [16] W.-S. Chen, Q. Zeng, B. Pan, A survey of deep nonnegative matrix factorization, *Neurocomputing* (2022).
- [17] Y. Zhang, Z. Zhang, Z. Zhang, M. Zhao, L. Zhang, Z. Zha, M. Wang, Deep self-representative concept factorization network for representation learning, in: *SIAM International Conference on Data Mining*, 2020, pp. 361–369.
- [18] A. Cichocki, R. Zdunek, Multilayer nonnegative matrix factorisation, *Electron Lett* 42 (16) (2006) 947–948.
- [19] G. Trigeorgis, K. Bousmalis, S. Zafeiriou, B. Schuller, A deep matrix factorization method for learning attribute representations, *IEEE Trans. Pattern Anal. Mach.* 39 (3) (2016) 417–429.
- [20] S. Arora, N. Cohen, W. Hu, Y. Luo, Implicit regularization in deep matrix factorization, in: *Advances in Neural Information Processing Systems*, 2019, pp. 7411–7422.

- [21] J. Bolte, S. Sabach, M. Teboulle, Proximal alternating linearized minimization for nonconvex and nonsmooth problems, *Math Program* 146 (1) (2014) 459–494.
- [22] M. Razaviyayn, M. Hong, Z.-Q. Luo, A unified convergence analysis of block successive minimization methods for nonsmooth optimization, *SIAM J. Optim.* 23 (2) (2013) 1126–1153.
- [23] Y.E. Nesterov, A method of solving a convex programming problem with convergence rate  $O(1/k^2)$ , *Dokl. akad. nauk Sssr* 269 (1983) 543–547.
- [24] B. O’donoghue, E. Candès, Adaptive restart for accelerated gradient schemes, *Found. Comput. Math.* 15 (3) (2015) 715–732.
- [25] Y. Xu, W. Yin, A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion, *SIAM J. Imaging* 6 (3) (2013) 1758–1789.
- [26] Y. Xu, W. Yin, A globally convergent algorithm for nonconvex optimization based on block coordinate update, *J Sci Comput* 72 (2) (2017) 700–734.
- [27] M. Abdolali, N. Gillis, Simplex-structured matrix factorization: sparsity-based identifiability and provably correct algorithms, *SIAM Journal on Mathematics of Data Science* 3 (2) (2021) 593–623.
- [28] P.O. Hoyer, Non-negative matrix factorization with sparseness constraints, *Journal of machine learning research* 5 (9) (2004).
- [29] H. Kim, H. Park, Sparse non-negative matrix factorizations via alternating non-negativity-constrained least squares for microarray data analysis, *Bioinformatics* 23 (12) (2007) 1495–1502.
- [30] R. Gribonval, R. Jenatton, F. Bach, Sparse and spurious: dictionary learning with noise and outliers, *IEEE Trans. Inf. Theory* 61 (11) (2015) 6298–6319.
- [31] P. Georgiev, F. Theis, A. Cichocki, Sparse component analysis and blind source separation of underdetermined mixtures, *IEEE Trans. Neural Netw.* 16 (4) (2005) 992–996.
- [32] R. Ohib, N. Gillis, N. Dalmaso, S. Shah, V.K. Potluru, S. Plis, Explicit group sparse projection with applications to deep learning and NMF, [arXiv:1912.03896](https://arxiv.org/abs/1912.03896) (2019).
- [33] L. Miao, H. Qi, Endmember extraction from highly mixed data using minimum volume constrained nonnegative matrix factorization, *IEEE Trans. Geosci. Remote Sens.* 45 (3) (2007) 765–777.
- [34] T.-H. Chan, C.-Y. Chi, Y.-M. Huang, W.-K. Ma, A convex analysis-based minimum-volume enclosing simplex algorithm for hyperspectral unmixing, *IEEE Trans. Signal Process.* 57 (11) (2009) 4418–4432.
- [35] X. Fu, K. Huang, B. Yang, W.-K. Ma, N.D. Sidiropoulos, Robust volume minimization-based matrix factorization for remote sensing and document clustering, *IEEE Trans. Signal Process.* 64 (23) (2016) 6254–6268.
- [36] V. Leplat, A.M.S. Ang, N. Gillis, Minimum-volume rank-deficient nonnegative matrix factorizations, in: *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2019, pp. 3402–3406.
- [37] N. Gillis, Successive nonnegative projection algorithm for robust nonnegative blind source separation, *SIAM J. Imaging* 7 (2) (2014) 1420–1450.
- [38] J.M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader, J. Chanussot, Hyperspectral unmixing overview: geometrical, statistical, and sparse regression-based approaches, *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* 5 (2) (2012) 354–379.
- [39] L. Tong, J. Yu, C. Xiao, B. Qian, Hyperspectral unmixing via deep matrix factorization, *Int J Wavelets Multiresolut Inf Process* 15 (06) (2017) 1750058.
- [40] P. De Handschutter, N. Gillis, Deep orthogonal matrix factorization as a hierarchical clustering technique, in: *European Signal Processing Conference (EU-SIPCO)*, IEEE, 2021, pp. 1466–1470.
- [41] F. Zhu, Spectral unmixing datasets with ground truths, [arXiv:1708.05125](https://arxiv.org/abs/1708.05125) (2017).
- [42] N. Gillis, D. Kuang, H. Park, Hierarchical clustering of hyperspectral images using rank-two nonnegative matrix factorization, *IEEE Trans. Geosci. Remote Sens.* 53 (4) (2015) 2066–2078.
- [43] L.T.K. Hien, D.N. Phan, N. Gillis, An inertial block majorization minimization framework for nonsmooth nonconvex optimization, [arXiv:2010.12133](https://arxiv.org/abs/2010.12133) (2020).
- [44] F. Malgouyres, J. Landsberg, On the identifiability and stable recovery of deep/multi-layer structured matrix factorization, in: *2016 IEEE Information Theory Workshop (ITW)*, IEEE, 2016, pp. 315–319.
- [45] L. Zheng, E. Riccietti, R. Gribonval, Hierarchical identifiability in multi-layer sparse matrix factorization, [arXiv:2110.01230](https://arxiv.org/abs/2110.01230) (2021).

**Pierre De Handschutter** received the Master’s degree in computer science engineering from Université de Mons, Belgium, in 2019, where he is currently a F.R.S-FNRS research fellow with the Department of Mathematics and Operational Research. He started his Ph.D. in 2019 under the supervision of Prof. Nicolas Gillis.

**Nicolas Gillis** is Professor with the Department of Mathematics and Operational Research, University of Mons, Belgium. He is recipient of the Householder award and ERC starting grant. He is editor for *IEEE Transactions on Signal Processing*, and *SIAM Journals on Matrix Analysis and Applications* and on *Mathematics of Data Science*.