



Contents lists available at ScienceDirect

## Linear Algebra and its Applications

journal homepage: [www.elsevier.com/locate/laa](http://www.elsevier.com/locate/laa)

## Smoothed separable nonnegative matrix factorization

Nicolas Nadisic<sup>a,\*</sup>, Nicolas Gillis<sup>a</sup>, Christophe Kervazo<sup>b</sup><sup>a</sup> Department of Mathematics and Operational Research, Faculté Polytechnique, Université de Mons, Rue de Houdain 9, 7000 Mons, Belgium<sup>b</sup> LTCI, Télécom Paris, Institut Polytechnique de Paris, 19 Place Marguerite Perey, 91120 Palaiseau, France

## ARTICLE INFO

*Article history:*

Received 4 July 2022

Received in revised form 19 April 2023

Accepted 12 July 2023

Available online 18 July 2023

Submitted by V. Mehrmann

*MSC:*

65F55

15A23

68U10

*Keywords:*

Nonnegative matrix factorization

Separability

Blind hyperspectral unmixing

Pure-pixel search algorithms

Latent simplex

Simplex-structured matrix

factorization

## ABSTRACT

Given a set of data points belonging to the convex hull of a set of vertices, a key problem in linear algebra, signal processing, data analysis and machine learning is to estimate these vertices in the presence of noise. Many algorithms have been developed under the assumption that there is at least one nearby data point to each vertex; two of the most widely used ones are vertex component analysis (VCA) and the successive projection algorithm (SPA). This assumption is known as the pure-pixel assumption in blind hyperspectral unmixing, and as the separability assumption in nonnegative matrix factorization. More recently, Bhattacharyya and Kannan (ACM-SIAM Symposium on Discrete Algorithms, 2020) proposed an algorithm for learning a latent simplex (ALLS) that relies on the assumption that there is more than one nearby data point to each vertex. In that scenario, ALLS is probabilistically more robust to noise than algorithms based on the separability assumption. In this paper, inspired by ALLS, we propose smoothed VCA (SVCA) and smoothed SPA (SSPA) that generalize VCA and SPA by assuming the presence of several nearby data points to each vertex. We illustrate the effectiveness of SVCA and SSPA over VCA, SPA and ALLS on synthetic data sets, on the unmixing of hyperspectral images, and on feature extraction on facial

\* Corresponding author.

E-mail addresses: [nicolas.nadisic@umons.ac.be](mailto:nicolas.nadisic@umons.ac.be) (N. Nadisic), [nicolas.gillis@umons.ac.be](mailto:nicolas.gillis@umons.ac.be) (N. Gillis), [christophe.kervazo@telecom-paris.fr](mailto:christophe.kervazo@telecom-paris.fr) (C. Kervazo).

images data sets. In addition, our study highlights new theoretical results for VCA.

© 2023 Elsevier Inc. All rights reserved.

## 1. Introduction

Low-rank matrix factorization is a widely used class of unsupervised data analysis models. Examples include principal component analysis [1], independent component analysis [2], sparse component analysis [3], and nonnegative matrix factorization [4], to cite a few. In this paper, we focus on the following matrix factorization model: given a set of data points within the convex hull of a set of vertices, estimate these vertices in the presence of noise. This problem can be formulated as follows.

**Problem 1.** Given  $X = WH + N \in \mathbb{R}^{m \times n}$  where  $H \in \mathbb{R}_+^{r \times n}$  is column stochastic and  $N$  is the noise, estimate  $W \in \mathbb{R}^{m \times r}$ .

Note that  $W$ ,  $H$  and  $N$  are unknown, only  $X$  is given. In this paper, we focus on the estimation of  $W$ . Once  $W$  is estimated, the matrix  $H$  can be estimated by solving a convex nonnegative least squares problem. Problem 1 is sometimes referred to as simplex-structured matrix factorization (SSMF), and generalizes nonnegative matrix factorization (NMF); see [5] and the references therein. It is also closely related to semi-NMF that only requires  $H$  to be nonnegative but not column stochastic [6]. Problem 1 is closely related to the column subset selection problem (see, e.g., [7,8] and the references therein), but the additional assumptions allow us to tackle it using convex geometry concepts.

In Problem 1, the columns of  $W$  are the vertices, while the columns of  $X$  are noisy data points within the convex hull of the columns of  $W$ ,

$$\text{conv}(W) = \{x \mid x = Wy, y \geq 0, e^\top y = 1\},$$

where  $e$  is the vector of all ones of appropriate dimension, and  $^\top$  is the transpose of a vector or a matrix. In fact, for all  $j$ ,

$$X(:, j) = WH(:, j) + N(:, j),$$

where  $X(:, j)$  denotes the  $j$ th column of  $X$ ,  $H(:, j) \geq 0$  and  $e^\top H(:, j) = 1$  (since  $H$  is column stochastic), which means that  $X(:, j) - N(:, j) \in \text{conv}(W)$  for all  $j$ . To be able to estimate  $W$  in Problem 1, appropriate assumptions on  $W$ ,  $H$  and  $N$  are required. In particular, the following assumptions are necessary [5]:

- No column of  $W$  is contained in the convex hull of the other columns of  $W$ , otherwise it is not possible to distinguish it from a data point.

- The data points must be sufficiently spread within  $\text{conv}(W)$ , having data points on each facet of  $\text{conv}(W)$ . This implies some degree of sparsity for  $H$ .
- The noise  $N$  must be bounded.

To obtain provable or practical algorithms, the above assumptions must be carefully and rigorously complemented. Different assumptions on  $W$ ,  $H$ , and  $N$  lead to different models for which different algorithms can be designed; see Section 2 for a literature review. We discuss below some applications of such algorithms.

*Applications in data analysis and machine learning.* The papers [9] and [10] describe in details various applications of solving Problem 1; in particular topic modeling via latent Dirichlet allocation, adversarial clustering and community detection via the mixed membership stochastic block model. Moreover, Problem 1 generalizes NMF and as such could be useful for all applications of NMF, such as feature extraction in sets of images [11], audio source separation [12], chemometrics [13], or blind hyperspectral unmixing (HU) [14,15]. Although the model considered in this work is very generic, we are motivated mainly by blind HU, a key problem in remote sensing. Let us briefly describe this problem; see the survey papers by [16,17] and the references therein for more details.

*Blind hyperspectral unmixing.* A hyperspectral image (HSI) is a picture of a scene acquired within a large number of spectral bands (usually between 100 and 200). Thus, for each pixel a precise electromagnetic spectrum is recorded, which gives an information concerning the materials present in the pixel; specifically, about their reflectances (fraction of incoming light they reflect) and/or their emissivity (which is due to the fact that the materials usually have non-zero temperatures). Unfortunately, despite their high spectral resolution, hyperspectral sensors generally have a low spatial resolution; as such, the spectrum recorded for each pixel might not correspond to the one of a single material, but rather to a mixture of the spectra of the different materials present within the pixel.

Given an HSI, blind HU thus aims to recover the set of materials present in the image, called endmembers, along with the abundances of each endmember in each pixel. The standard model used to solve blind HU is the linear mixing model. It assumes that the spectral signature of each pixel is a linear combination of the spectral signatures of the endmembers, where the weights of the linear combination are the abundances of the endmembers in the pixel. Typically, we represent a hyperspectral image as a matrix  $X$ , where the  $j$ th column,  $X(:, j)$ , corresponds to the spectral signature of the  $j$ th pixel in the scene. If the spectral signatures of the endmembers are also collected as the columns of a matrix  $W$ , then according to the linear mixing model, the  $j$ th pixel can be written as  $X(:, j) \approx \sum_{k=1}^r W(:, k)H(k, j) + N(:, j)$ , where  $H(k, j)$  is the abundance of the  $k$ th endmember in the  $j$ th pixel, and  $N(:, j)$  represents the noise and model misfit. This is exactly the setup of Problem 1.

*Pure-pixel assumption.* An important class of algorithms to solve blind HU are *pure-pixel search* algorithms. They rely on the assumption that for each endmember, there is at least one pixel in which this endmember appears almost alone, that is, purely, so that the endmember signature is close to the one of the corresponding pure pixel. Two of the most effective and widely used pure-pixel search algorithms are vertex component analysis (VCA) [18] and the successive projection algorithm (SPA) [19] which will be described in Section 2.

*Terminology.* In the NMF literature, the pure-pixel assumption is referred to as separability, and the corresponding algorithms are referred to as separable NMF algorithms, or near-separable NMF algorithms. Geometrically, the pure-pixel assumption requires that there is a data point (that is, a column of  $X$ ), close to each vertex (that is, to each column of  $W$ ). We will refer to such a data point as pure when the corresponding column of  $H$  is a unitary vector (that is, a column of the identity matrix).

*Contribution and outline.* When the separability assumption holds, there are typically more than one data point close to each vertex. In this paper, we leverage this observation by adapting VCA and SPA, providing two new algorithms, namely smoothed VCA (SVCA) and smoothed SPA (SSPA). The idea is to aggregate (for instance average) several data points around a vertex to obtain a better estimate of that vertex. In practice, it enables our two smoothed algorithms to achieve much better separation results than their non-smoothed counterparts, and to be more tolerant to noise.

SVCA can also be interpreted as an adaptation of the algorithm for learning a latent simplex (ALLS) proposed by [9], with two major modifications that we detail in Section 3. This observation allows us to provide theoretical guarantees for SVCA, and hence VCA which is a special case of SVCA. To the best of our knowledge, this is the first time a theoretical guarantee for VCA is provided in the presence of noise.

We believe that this work paves the way for a whole new branch of smoothed separable NMF algorithms. Although we focus in this work on VCA and SPA, two of the most well-known separable NMF algorithms, it is straightforward to extend the methodology to more recent algorithms such as SNPA [20].

The paper is organized as follows. In Section 2, we summarize the literature for solving Problem 1, with a focus on three algorithms, namely VCA, SPA and ALLS. In Section 3, we propose SVCA which is equivalent to applying VCA on a smoothed data set. SVCA is similar to ALLS, but two key differences make it empirically more efficient than ALLS. In Section 4, we propose SSPA which adapts SPA in the presence of multiple pure pixels. In Section 5, we show on synthetic and real-world hyperspectral data sets that SVCA and SSPA outperform VCA, SPA and ALLS in the presence of multiple pure pixels.

*Notation.* Given an  $m$ -by- $n$  real matrix  $X \in \mathbb{R}^{m \times n}$ ,  $X^\top$  is its transpose,  $X(:, j)$  is its  $j$ th column,  $X(i, :)$  is its  $i$ th row,  $X(i, j)$  is its entry at position  $(i, j)$ ,  $X(:, \mathcal{K})$  is the submatrix made of the columns of  $X$  in the index set  $\mathcal{K}$ , and similarly for  $X(\mathcal{K}, :)$  for

the rows. The  $i$ th singular value of  $X$  is denoted  $\sigma_i(X)$ , while  $K(W) = \max_j \|W(:, j)\|_2$ . If a matrix  $X$  is component-wise nonnegative, we write  $X \geq 0$ . The vector of all ones is denoted  $e$ . The  $m$ -by- $m$  identity matrix is denoted  $I_m$ .

## 2. Simplex-structured matrix factorization

In this section, we describe existing models and algorithms to solve Problem 1, that is, to solve SSMF, in order to identify the vertices of the convex hull of a set of data points,  $X$ . We focus on two key models, upon which our contribution is built:

1. Separable NMF: it assumes there is one data point close to each vertex of  $\text{conv}(W)$ , and that the noise added to each data point is bounded; see Section 2.1. Some authors refer to this model as near-separable NMF.
2. Learning a latent simplex: it is motivated by machine learning applications. It assumes that there is more than one data point close to each vertex of  $\text{conv}(W)$ , but it allows much larger noise levels as it only requires the  $\ell_2$  norm of  $N$  to be bounded, instead of each individual column; see Section 2.2.

Other models and algorithms exist to tackle Problem 1 relying on different assumptions. It is worth mentioning minimum-volume NMF [21], where  $W$  is regularized such that its convex hull  $\text{conv}(W)$  has the smallest possible volume, facet-based identification algorithms that identify the facets of  $\text{conv}(W)$  from which its vertices are recovered [22–24,5], and probabilistic simplex component analysis [25] that relies on a probabilistic model on the data (the columns of  $H$  are sampled using the Dirichlet distribution, and the entries of  $N$  using i.i.d. Gaussian noise). Discussing these approaches in detail is out of the scope of this article, whose focus is on the separable NMF problem.

### 2.1. Model 1: separable NMF

As already mentioned, an important class of NMF algorithms relies on the separability assumption, defined as follows.

**Assumption 1** (*Separability*). In Problem 1, there exists an index set  $\mathcal{K}$  of cardinality  $r$  such that  $H(\mathcal{K}, :) = I_r$  where  $I_r$  is the  $r$ -by- $r$  identity matrix.

Under this assumption, solving Problem 1 amounts to recover  $\mathcal{K}$  such that

$$X(:, \mathcal{K}) = W + N(:, \mathcal{K}) \approx W.$$

In blind HU, separability is known as the pure-pixel assumption.

The early algorithms building on this assumption emerged in the blind HU community. They include pixel purity index (PPI) [26], N-FINDR [27], and vertex component analysis

(VCA) [18]. Most of these algorithms were developed based on convex geometry concepts. These early works however did not analyze noise robustness, and in fact they are not guaranteed to recover the endmembers in the presence of noise.

In analytical chemistry, Problem 1 is closely related to the problem of self-modeling curve resolution [28]. As in blind HU, several algorithms were developed based on geometry concepts; in particular the successive projection algorithm (SPA) [19].

More recently, and motivated by applications in machine learning (in particular, topic modeling where pure data points are referred to as anchor words), [29] introduced the first provably robust separable NMF algorithms. Their robustness is deterministic: under some conditions, their algorithm is guaranteed to recover an approximation of the vertices. Arora et al. were not aware of the algorithms developed within the blind HU literature. Many provably robust algorithms have followed this seminal paper, including algorithms that use linear programming [30–32], a generalization of SPA [33], fast anchor words [34], and the successive nonnegative projection algorithm (SNPA) [20]. These deterministically robust algorithms guarantee that, in the presence of noise, the vertices are recovered, up to some error bounds that depend on the noise level and the conditioning of  $W$ ; see Section 2.1.2 for such a result for SPA. We refer the interested reader to [4, Chapter 7] for a detailed discussion and comparison of these algorithms.

In the following, we describe in more detail VCA and SPA that will be instrumental in proposing our new algorithms, smoothed VCA in Section 3 and smoothed SPA in Section 4.

### 2.1.1. Vertex component analysis

VCA [18] is a greedy separable NMF algorithm, that is, it identifies the indices of the subset  $\mathcal{K}$  sequentially. The index set is initialized with  $\mathcal{K} = \emptyset$ . At each of the  $r$  iterations of VCA, a random direction belonging to the subspace spanned by the  $r$  top left singular vectors of  $X$  is generated (this is equivalent to working with the best rank- $r$  approximation of  $X$ , and hence filters the noise). This direction is then projected onto the orthogonal complement of  $X(:, \mathcal{K})$ , and the index of the column of  $X$  that maximizes the absolute value of the inner product with that direction is added to  $\mathcal{K}$ . Algorithm 1 summarizes VCA.

The computational cost of VCA is  $\mathcal{O}(r \text{nnz}(X))$  operations, where  $\text{nnz}(X)$  is the number of non-zero entries of  $X$ . The main cost is to compute  $Y$  which can be done efficiently using the subspace power iteration, in  $\mathcal{O}(\text{nnz}(X)r)$  operations, and to compute the products  $(d_k^T P^\perp)X$  at each of the  $r$  iterations.

An important drawback of VCA is that it is not guaranteed to be deterministically robust to noise. In other words, for any noise  $N$  such that a data point goes outside  $\text{conv}(W)$ , there is a non-zero probability that VCA extracts this point. The reason is that VCA uses a linear function to identify the vertices of  $\text{conv}(W)$ ; see [4, Chapter 7.4] for a numerical example.

**Algorithm 1** Vertex Component Analysis (VCA) [18].

**Require:** The matrix  $X \in \mathbb{R}^{m \times n}$ , the number  $r$  of columns to extract.

**Ensure:** Index set  $\mathcal{K}$  of cardinality  $r$  such that  $X \approx X(:, \mathcal{K})H$  for some  $H \geq 0$ .

- 1: Let  $\mathcal{K} = \emptyset$ ,  $P^\perp = I_m$ ,  $V = []$ .
- 2: Let the column of  $Y \in \mathbb{R}^{m \times r}$  be the top  $r$  left singular vectors of  $X$ .
- 3: **for**  $k = 1 : r$  **do**
- 4:   Pick a random direction  $d_k \in \mathbb{R}^m$  in the subspace spanned by  $Y$ , e.g.,  $d_k \sim \mathcal{YN}(0, I_r)$  where  $\mathcal{N}(0, I_r)$  is the normal distribution of mean 0 and covariance matrix  $I_r$ .
- 5:   Compute  $u_k = (d_k^\top P^\perp)X \in \mathbb{R}^n$ , and let  $j_k = \operatorname{argmax}_{1 \leq j \leq n} |u_k(j)|$ .
- 6:   Let  $\mathcal{K} = \mathcal{K} \cup \{j_k\}$ .
- 7:   Update the projector  $P^\perp$  onto the orthogonal complement of  $W = X(:, \mathcal{K})$ :

$$v_k = \frac{P^\perp X(:, j_k)}{\|P^\perp X(:, j_k)\|_2},$$

$$V = [V \ v_k],$$

$$P^\perp \leftarrow (I_m - VV^\top).$$

8: **end for**

*2.1.2. Successive projection algorithm*

SPA [19] is very similar to VCA. The only difference is in the selection step, when adding an index to  $\mathcal{K}$ . SPA selects the column of  $P^\perp X$  with maximum  $\ell_2$  norm; see Algorithm 2.

**Algorithm 2** Successive Projection Algorithm (SPA) [19].

**Require:** The matrix  $X \in \mathbb{R}^{m \times n}$ , the number  $r$  of columns to extract.

**Ensure:** Index set  $\mathcal{K}$  of cardinality  $r$  such that  $X \approx X(:, \mathcal{K})H$  for some  $H \geq 0$ .

- 1: Let  $\mathcal{K} = \emptyset$ ,  $P^\perp = I_m$ ,  $V = []$ .
- 2: Let  $u_1(j) = \|X(:, j)\|_2^2$  for all  $j$ .
- 3: **for**  $k = 1 : r$  **do**
- 4:   Let  $j_k = \operatorname{argmax}_{1 \leq j \leq n} u_k(j)$ . (Break ties arbitrarily, if necessary.)
- 5:   Let  $\mathcal{K} = \mathcal{K} \cup \{j_k\}$ .
- 6:   Update the projector  $P^\perp$  (as in step 7 of Algorithm 1).
- 7:   Update the squared norms of the columns of  $P^\perp X$ : for all  $j$ ,

$$u_{k+1}(j) = u_k(j) - v_k^\top X(:, j) = \|P^\perp X(:, j)\|_2^2.$$

8: **end for**

It is interesting to note that VCA is equivalent to SPA if the direction  $u_k$  randomly chosen at each step is instead taken as the column of the residual  $P^\perp X$  with maximum  $\ell_2$  norm. We will use this observation for our proposed algorithm, smoothed SPA.

**Remark 1** (*On the projection  $P^\perp$* ). As opposed to VCA, SPA does not work on the subspace spanned by the first  $r$  singular vectors of  $X$ , and hence we stick to this variant

in this paper. However, in practice, projecting the data onto this subspace allows noise filtering and typically leads to better numerical performance.

*Robustness of SPA.* As opposed to VCA, SPA is deterministically robust to noise, provided the following assumption on top of separability (Assumption 1):

**Assumption 2** (*Column-wise bounded noise*). In Problem 1, the noise satisfies  $\|N(:, j)\|_2 \leq \epsilon$  for all  $j$  for some  $\epsilon > 0$  sufficiently small.

Let us state the robustness result for SPA.

**Theorem 1.** [33, Theorem 3] *Let  $X = WH + N$  as in Problem 1, and let Assumptions 1 and 2 be satisfied, that is,  $W = X(:, \mathcal{K}^*)$  for some index set  $\mathcal{K}^*$  of cardinality  $r$ , and  $\|N(:, j)\|_2 \leq \epsilon$  for all  $j$  where  $\epsilon \leq \mathcal{O}\left(\frac{\sigma_r^3(W)}{\sqrt{r}K(W)^2}\right)$  with  $K(W) = \max_j \|W(:, j)\|_2$ . Let also the  $r$ th singular value of  $W$  be positive, that is,  $\sigma_r(W) > 0$ , meaning that  $W$  has rank  $r$ . Let  $\mathcal{K}$  be the index set extracted by SPA. Then there exists a permutation  $\pi$  of  $\{1, 2, \dots, r\}$  such that for all  $k = 1, 2, \dots, r$ ,*

$$\|X(:, \mathcal{K}(k)) - W(:, \pi(k))\|_2 \leq \mathcal{O}\left(\frac{\epsilon K(W)^2}{\sigma_r^2(W)}\right),$$

where  $\mathcal{K}(k)$  denotes the  $k$ th index in  $\mathcal{K}$ .

Note that the bounds in Theorem 1 are relatively weak: the noise level has to be rather small to guarantee SPA to recover  $W$  approximately.

### 2.2. Model 2: learning a latent simplex

A drawback of separable NMF algorithms, such as VCA and SPA, is that they assume that there is only one data point close to each column of  $W$ . Therefore, to estimate  $W$ , the column-wise bounded noise assumption (Assumption 2) is necessary; see Theorem 1. This is a rather strong assumption, often not met in practical situations as typically many data points are affected by large amounts of noise.

The paper [9] rather proposes to leverage the fact that typically more than one data point are close to each column of  $W$ . This assumption, which is stronger than the pure-pixel one (requiring only a single pure-pixel), allows higher noise levels. It is called the *proximate latent points* assumption, and is defined as follows.

**Assumption 3** (*Proximate latent points*). In Problem 1, there exists  $r$  index sets,  $\mathcal{K}_k$  for  $k = 1, 2, \dots, r$ , of cardinality at least  $p = \delta n$  such that

$$\|WH(:, j) - W(:, k)\|_2 \leq \frac{4\sigma}{\delta} \text{ for all } j \in \mathcal{K}_k,$$

for some  $\delta \in [\frac{1}{n}, \frac{1}{r}]$  and  $\sigma > 0$ .

Under this assumption, instead of looking for one column of  $X$  to represent each vertex, like in VCA and SPA, algorithms should look for  $p$  of them and then estimate each vertex as the average of these  $p$  data points. We will refer to such algorithms as *smoothed separable NMF algorithms*. The main contribution of this paper is to propose two new such algorithms; in Sections 3 and 4.

This assumption is often met in the machine learning applications mentioned in Section 1; see the discussions in [9,10]. For high-resolution HSIs that satisfy the pure-pixel assumption, there are typically more than one pixel close to each endmember. This claim will be validated numerically in Section 5.2.

*2.2.1. Algorithm for learning a latent simplex*

To solve Problem 1 under Assumption 3, [9] proposed an algorithm similar to VCA, which we refer to as the algorithm for learning a latent simplex (ALLS). The main difference between ALLS and VCA is the selection step. Instead of picking a single column of  $X$ , ALLS averages over  $p$  columns for some  $p \in \{1, 2, \dots, \lfloor \frac{n}{r} \rfloor\}$ . More precisely, ALLS picks the  $p$  columns corresponding to the indices that maximize the absolute value of  $u_k$ ; see Algorithm 3.

The idea behind ALLS is to apply VCA on a smoothed data set. This smoothed data set is made of  $\binom{n}{p}$  data points which are the averages of all possible combinations of  $p$  data points, that is,  $p$  columns of  $X$ . Of course, constructing this smoothed data set explicitly is not practical, since  $\binom{n}{p}$  grows exponentially. However, by the linearity of the selection step in VCA, this is not necessary: the smoothed data point that maximizes a linear function is the average of the  $p$  data points that have the  $p$  largest values for that function. Algorithm 3 summarizes ALLS.

---

**Algorithm 3** Algorithm for Learning a Latent Simplex (ALLS) [9].

---

**Require:** The matrix  $X \in \mathbb{R}^{m \times n}$ , the number  $r$  of columns of  $W$ , the number  $p$  of columns of  $X$  to be averaged to obtain each column of  $W$ .

**Ensure:** A matrix  $W'$  such that  $X \approx W'H$  for some  $H \geq 0$ .

- 1: Let  $W' = []$ ,  $P^\perp = I_m$ ,  $V = []$ .
  - 2: Let the column of  $Y \in \mathbb{R}^{m \times r}$  be the top  $r$  left singular vectors of  $X$ .
  - 3: **for**  $k = 1 : r$  **do**
  - 4:   Pick a random direction  $d_k \in \mathbb{R}^m$  in the subspace spanned by  $Y$ , e.g.,  $d_k \sim Y\mathcal{N}(0, I_r)$ .
  - 5:   Compute  $u_k = (d_k^\top P^\perp) X \in \mathbb{R}^n$ .
  - 6:   Let  $\mathcal{S}_k$  be the set of  $p$  indices corresponding to the largest coordinates of  $u_k$  in absolute value.
  - 7:   Let  $W'(:, k)$  average of the columns of  $X(:, \mathcal{S}_k)$ .
  - 8:   Update the projector  $P^\perp$  (as in step 7 of Algorithm 1).
  - 9: **end for**
-

Note that ALLS with  $p = 1$  is equivalent to VCA.

*Computational cost.* The only additional cost of ALLS compared to VCA is to average  $p$  columns of  $X$ , which requires  $r$  times  $\mathcal{O}(pm)$  operations, which is negligible since  $p \ll n \leq \text{nnz}(X)$ .

2.2.2. Probabilistic robustness of ALLS

Let us describe the assumptions needed to prove the probabilistic robustness of ALLS. The condition on  $W$  is defined as follows:

**Assumption 4** (*Well-separatedness of  $W$* ). In Problem 1, the matrix  $W$  satisfies

$$\alpha(W) = \frac{\min_{k=1,2,\dots,r} \min_x \|W(:,k) - W(:,\bar{k})x\|_2}{K(W)} > 0, \tag{1}$$

where  $\bar{k} = \{1, 2, \dots, r\} \setminus \{k\}$ .

Assumption 4 holds if and only if  $\text{rank}(W) = r$ , in which case  $\text{conv}(W)$  is a simplex, that is, a polytope of dimension  $r - 1$  with  $r$  vertices (hence the name of the algorithm).

The condition on the noise is as follows.

**Assumption 5** (*Spectrally bounded perturbations*). In Problem 1,

$$\|N\|_2 = \sigma_1(N) \leq \sigma\sqrt{n},$$

where there exists some constant  $c$  such that

$$\sigma \leq \frac{\alpha^2 \sqrt{\delta}}{c r^9} \min_j \|W(:,j)\|_2, \tag{2}$$

where  $\alpha = \alpha(W)$  is defined in Assumption 4, and  $\delta$  and  $\sigma$  in Assumption 3 (recall,  $p = \delta n$  is the number of data points close to each column of  $W$ ).

It is key to note here that the noise allowed is not column wise as in Assumption 2, but on the spectral norm of  $N$ , which is rather different.

We can now state the robustness theorem for ALLS.

**Theorem 2.** [9] *Let us consider Problem 1 under Assumptions 3 (proximate latent points), 4 (well-separatedness of  $W$ ) and 5 (spectrally bounded perturbations). Then, with probability at least  $1 - c/r^{3/2}$ , ALLS computes a matrix  $W'$  such that upon permutation of its columns, for all  $k = 1, 2, \dots, r$ ,*

$$\|W(:,k) - W'(:,k)\|_2 \leq O\left(\frac{r^4 \sigma}{\alpha \sqrt{\delta}}\right). \tag{3}$$

Note that substituting (2) in (3) gives

$$\|W(:, k) - W'(:, k)\|_2 \leq O\left(\frac{\alpha}{c r^5}\right) \min_j \|W(:, j)\|_2.$$

*Implications for VCA.* Interestingly, since ALLS for  $p = 1$  coincides with VCA, Theorem 2 provides a probabilistic robustness result for VCA which is unknown in the blind HU literature.

### 2.2.3. Bounds of SPA versus ALLS

Theorem 2 might look somewhat weak because of the dependence in  $r^9$  in the bound (2) for  $\sigma$ . However, it is not known whether this bound is tight, although it is believed it could be improved [9,10]. A similar comment applies to SPA. Moreover, these bounds assume an adversarial setting, and noise robustness under particular generative models is also an interesting direction of research, as in [25].

In any case, Theorem 2 only requires a bound on  $\|N\|_2$  while SPA requires each column of the noise matrix  $N$  to be bounded, indicating that ALLS should perform better, in general, when  $p$  is sufficiently large. Since the theory is still not fully developed and the tightness of the theoretical bounds should be carefully studied, it is important to compare these algorithms empirically to shed light on their differences on practical problems; this will be done in Section 5.

## 3. Smoothed VCA

Inspired by VCA, and ALLS, we now propose smoothed VCA (SVCA); see Algorithm 4. SVCA has two key important differences compared to ALLS:

1. At step  $k$ , ALLS selects the  $p$  entries maximizing the absolute value of the vector of  $u_k$ , obtained as the inner product of  $X$  and a randomly generated direction  $d_k^\top P^\perp$ ; see steps 5-6 of Algorithm 3. This is not equivalent to maximizing (or minimizing) the linear function  $l(x) = d_k^\top P^\perp x$  over the smoothed polytope. In fact, by using the absolute value, this approach could select data points in opposite directions. For example, take the simple case with two vertices  $w_1 = (-1, 0)$  and  $w_2 = (1, 0)$ . For any direction  $d$ , we have  $|d^\top w_1| = |d^\top w_2|$  and hence it is very likely that data points close to both vertices will maximize  $|d^\top x|$ , and their average will be a poor approximation of both vertices.

Instead, to maximize (or minimize)  $l(x)$ , one should select the  $p$  indices maximizing  $u_k$  (or  $-u_k$ ). In SVCA, we therefore propose to select the  $p$  indices that maximize (resp. minimize)  $u_k$  if the median of the  $p$  largest values is larger (resp. smaller) than the absolute value of the median of the  $p$  smallest values of  $u_k$ . We have observed in practical experiments that this modification of ALLS is crucial to obtain competitive results in real-world hyperspectral images. In fact, we will see that SVCA outper-

forms ALLS, and the main reason is this modified selection step, see lines 6 to 10 of Algorithm 4.

2. Instead of averaging  $p$  columns of  $X$  at each step, we will also consider taking their median. This allows SVCA to be much more tolerant to non-Gaussian noise, such as gross corruptions and outliers, which are often present in HSI. While in the presence of Gaussian noise, using the average might in principle lead to the best unmixing results, the choice of a different aggregation function might still be appealing on real world data sets, in which the linear mixing model is not always perfectly accurate. In this context, other aggregation functions might be more robust to the model imperfections (this is for instance the case in HSI, where the data sets sometimes exhibit spectral variabilities – an aggregation function enabling to reduce the pure pixel variabilities impact would be appealing). Using the median also allows SVCA to be more tolerant to a misspecified value of  $p$ . For example, assume a scenario where there are exactly  $p'$  data points close to each vertex. For  $p < p'$ , one does not leverage optimally the presence of multiple pure data points. On the other side, as soon as  $p$  is larger than  $p'$ , ALLS will perform rather badly because it will average  $p'$  data points close to a vertex and  $p - p'$  data points potentially far away. If instead one takes the median, the algorithm remains able to extract the vertices accurately for any  $p < 2p'$ . In practice, as we will show in Section 5, using the median performs significantly better on real data sets.

Note that SVCA has the same computational cost as VCA, SPA and ALLS, namely  $\mathcal{O}(r\text{nnz}(X))$  operations.

---

**Algorithm 4** Smoothed Vertex Component Analysis (SVCA).

---

**Require:** The matrix  $X \in \mathbb{R}^{m \times n}$ , the number  $r$  of columns of  $W \in \mathbb{R}^{m \times r}$ , the number  $p$  of columns of  $X$  to be averaged to obtain each column of  $W$ , the aggregation method (median or mean).

**Ensure:** A matrix  $W$  such that  $X \approx WH$  for some  $H \geq 0$ .

- 1: Let  $W = []$ ,  $P^\perp = I_m$ ,  $V = []$ .
  - 2: Let  $Y \in \mathbb{R}^{m \times r}$  be the vector space spanned by the top  $r$  left singular vectors of  $X$ .
  - 3: **for**  $k = 1 : r$  **do**
  - 4:   Pick a random direction  $d_k \in \mathbb{R}^m$  in the subspace spanned by  $Y$ , e.g.,  $d_k \sim Y\mathcal{N}(0, I_r)$ .
  - 5:   Compute  $u_k = (d_k^\top P^\perp) X \in \mathbb{R}^n$ .
  - 6:   **if** the median of the  $p$  largest values of  $u_k$  is larger than the absolute value of the median of the  $p$  smallest values of  $u_k$  **then**
  - 7:     Let  $\mathcal{S}_k$  be the set of  $p$  indices maximizing  $u_k$ .
  - 8:   **else**
  - 9:     Let  $\mathcal{S}_k$  be the set of  $p$  indices minimizing  $u_k$ .
  - 10:   **end if**
  - 11:   Let  $W(:, k)$  be the median (or the mean) of the columns of  $X(:, \mathcal{S}_k)$ .
  - 12:   Update the projector  $P^\perp$  (as in step 7 of Algorithm 1).
  - 13: **end for**
-

*Recovery guarantees for SVCA.* SVCA is very similar to ALLS, and in fact the robustness analysis of ALLS applies to SVCA, that is, Theorem 2 applies to SVCA. The reason is that we guarantee SVCA to extract the data point in the smoothed data set that maximizes the absolute value  $l(x) = d_k^\top P^\perp x$ .

#### 4. Smoothed SPA

Since SVCA is equivalent to VCA for  $p = 1$ , it is not guaranteed to be deterministically robust to noise. This motivates us to propose smoothed SPA; see Algorithm 5. Unfortunately, it is not practical to apply SPA directly on the smoothed data set. Indeed, it would require to find the  $p$  columns of the smoothed data set with the largest  $\ell_2$  norm. The  $\ell_2$  norm being a nonlinear function, it would require to explicitly compute the  $\binom{n}{p}$  data points of the smoothed data set, which is computationally prohibitive.

Instead, we replace the random selection of  $u_k = Y\mathcal{N}(0, 1)$  in SVCA by the selection of the column of the residual  $P^\perp X$  with maximum  $\ell_2$  norm, that is,  $u_k = P^\perp X(:, j_k)$  for some  $j_k$  so that  $\|u_k\|_2 \geq \|P^\perp X(:, j)\|_2$  for all  $j$ . This allows us to combine the best of ‘both worlds’: deterministic robustness under separability when  $p = 1$ , and the use of the proximate latent point assumption when  $p > 1$  (Assumption 3).

---

#### Algorithm 5 Smoothed Successive Projection Algorithm (SSPA).

---

**Require:** The matrix  $X \in \mathbb{R}^{m \times n}$ , the number  $r$  of columns of  $W \in \mathbb{R}^{m \times r}$ , the number  $p$  of columns of  $X$  to be averaged to obtain each column of  $W$ , the aggregation method (median or mean).  
**Ensure:** A matrix  $W$  such that  $X \approx WH$  for some  $H \geq 0$ .

- 1: Let  $W = [], P^\perp = I_m, V = []$ .
- 2: Let  $u_1(j) = \|X(:, j)\|_2^2$  for all  $j$ .
- 3: **for**  $k = 1 : r$  **do**
- 4: Let  $j_k = \operatorname{argmax}_{1 \leq j \leq n} u_k(j)$ . (Break ties arbitrarily, if necessary.)
- 5: Let  $d_k = X(:, j_k)$ .
- 6: Compute  $u_k = (d_k^\top P^\perp)X \in \mathbb{R}^n$ .
- 7: **if**  $\max_i u_k(i) \geq -\min_i u_k(i)$  **then**
- 8: Let  $\mathcal{S}_k$  be the set of  $p$  indices maximizing  $u_k$ .
- 9: **else**
- 10: Let  $\mathcal{S}_k$  be the set of  $p$  indices minimizing  $u_k$ .
- 11: **end if**
- 12: Let  $W(:, k)$  be the median (or the mean) of the columns of  $X(:, \mathcal{S}_k)$ .
- 13: Update the projector  $P^\perp$  (as in step 7 of Algorithm 1).
- 14: Update the squared norms of the columns of  $P^\perp X$ : for all  $j$ ,

$$u_{k+1}(j) = u_k(j) - v_k^\top X(:, j) = \|P^\perp X(:, j)\|_2^2.$$

15: **end for**

---

*Recovery guarantees for SSPA.* For  $p = 1$ , SSPA coincides with SPA, and hence Theorem 1 applies to SSPA for  $p = 1$ , that is, it is deterministically robust to column-wise bounded noise. However, since the selection step of SSPA is deterministic, Theorem 2 does not apply to SSPA. A promising direction of further research would be to analyze noise robustness of SSPA for  $p > 1$ .

*Should you use SVCA or SSPA?* SVCA has the advantage to be a randomized algorithm, and hence can be run multiple times and the best solution, according to some criterion, can be kept. SSPA is deterministic and has the advantage to have stronger theoretical guarantees for  $p = 1$ . From a practical point of view, one could run SVCA several times, and SSPA once, and then keep the best solution.

*Comparing algorithms.* In this paper, when the ground truth  $W^*$  is not available, we will use the following criterion

$$Q_F(W) = \frac{\min_{H \geq 0} \|X - WH\|_F}{\|X\|_F} \in [0, 1],$$

to evaluate the quality of a solution  $W$ . Note that we do not use the sum-to-one constraint,  $H^\top e = e$ , because, in many practical situations, including hyperspectral imaging, this constraint is not satisfied for all columns of  $H$ , e.g., for pixels with low luminosity; see a discussion in [20].

In matrix factorization, it could be argued that in general  $Q_F(W)$  is not a good measure to assess the quality of a solution  $W$ , since any  $W$  such that  $\text{conv}(X) \subset \text{conv}(W)$  implies  $Q_F(W) = 0$ , even if  $W$  is very different from the ground truth  $W^*$ . In our case though,  $Q_F(W)$  is relevant, since all the considered algorithms generate solutions  $W$  which columns are close to  $\text{conv}(X)$ . Indeed, VCA, SPA and ALLS generate solutions within  $\text{conv}(X)$ . This is not always the case for SVCA and SSPA, because of the use of the median (a non-linear operator) for the aggregation of the columns of  $X$ . In practice though, they find solutions close to  $\text{conv}(X)$ .

## 5. Numerical experiments

In this section, we study and compare the performance of ALLS, SVCA, and SSPA. We first consider synthetic data sets. Then we tackle the unmixing of real-world hyperspectral images and the extraction of features in facial images data sets. The code and data are available online.<sup>1</sup> All algorithms are implemented in Matlab and run on a computer with an i5-8350U processor.

### 5.1. Synthetic data sets

In this section, we study the behavior of smooth separable NMF algorithms in several experimental setups. To build synthetic data sets, we first build  $W \in \mathbb{R}_+^{224 \times 10}$  by selecting 10 columns from the USGS hyperspectral library<sup>2</sup> using SPA. The condition number of the corresponding matrix is  $\kappa(W) = 33.88$ . Then, we generate a random  $H \in \mathbb{R}_+^{10 \times 1000}$  such that  $H = [I_{10}, H']$ , meaning there is at least one pure data point for every vertex.

<sup>1</sup> <https://gitlab.com/nnadisic/smoothed-separable-nmf>.

<sup>2</sup> <https://www.usgs.gov>.

**Table 1**

Generating the columns of  $H \in \mathbb{R}^{10 \times n}$  using the Dirichlet distribution of parameter  $\alpha e$ , this table reports the expected percentage,  $\delta$ , of pure data points close to each vertex. The  $j$ th data point is considered close to the  $i$ th vertex when  $H(i, j) > 0.95$ , hence this table reports the expected value of  $\frac{1}{n} |\{j \mid H(i, j) > 0.95\}|$  for all  $i$ . Since the Dirichlet distribution has the same parameters for all  $i$ , namely  $\alpha$ , this expected value is the same for all  $i$ .

$\alpha$	0.01	0.02	0.05	0.1	0.2	0.5
$\delta$	7.7%	5.9%	2.7%	0.75%	0.06%	0%

The coefficients of  $H'$  follow a Dirichlet distribution, which is usually a good model for the abundances in HSI [35], of parameters  $\alpha e$ , where  $\alpha$  controls the proportion of data points close to the vertices, see Table 1. The larger  $\alpha$ , the denser the columns of  $H$  and the less likely the ‘proximal latent points’ assumption is to be satisfied for large  $p$ .

Finally, we let  $X = WH + N$  where  $N$  is a normalized Gaussian noise: Given a noise level  $\epsilon$ , we first generate  $N(i, j) \sim \mathcal{N}(0, 1)$  for all  $(i, j)$ , then set

$$N \leftarrow \epsilon \frac{\|WH\|_F}{\|N\|_F} N,$$

so that  $\epsilon$  is the norm of the noise relative to  $WH$ :  $\|N\|_F = \epsilon \|WH\|_F$ .

Given the noisy data matrix  $X$  and a parameter  $p$ , we can compute  $W'$  with the algorithms ALLS, SVCA, and SSPA. We note ALLS( $p$ ), SVCA( $p$ ) and SSPA( $p$ ) these algorithms run with parameter  $p$ . Given the computed solution  $W'$ , we report the mean removed spectral angle (MRSa) to assess its quality. Given two spectral signatures  $x, y \in \mathbb{R}^m$ , the MRSa is defined as follows:

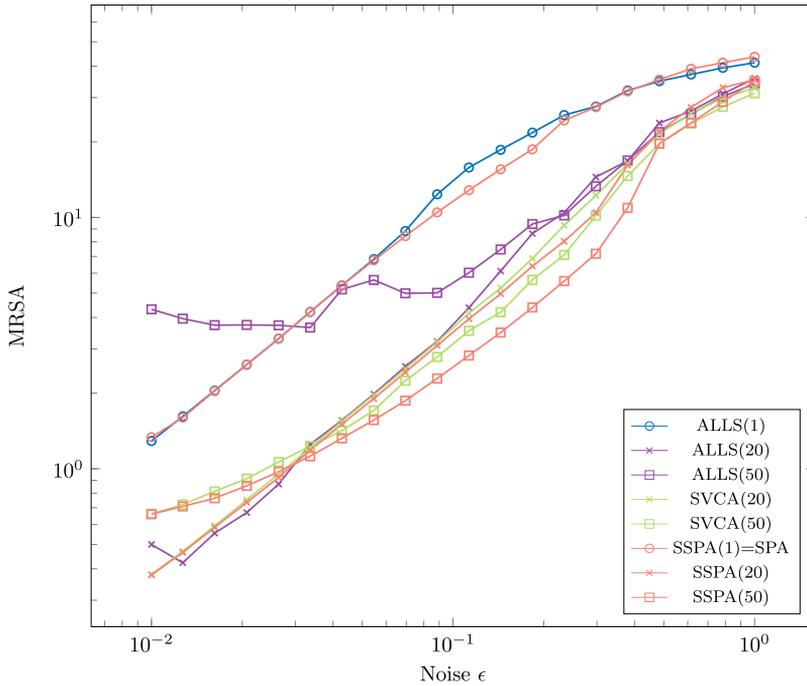
$$\phi(x, y) = \frac{1}{\pi} \arccos \left( \frac{(x - \bar{x})^\top (y - \bar{y})}{\|x - \bar{x}\|_2 \|y - \bar{y}\|_2} \right) \in [0, 1],$$

where for a vector  $z \in \mathbb{R}^m$ ,  $\bar{z} = (\sum_{i=1}^m z_i) e$  and  $e$  is the vector of all ones. Given two matrices, here the groundtruth  $W$  and the estimate  $W'$ , we define the MRSa as

$$\text{MRSa}(W, W') = \sum_{j=1}^r \phi(W(:, j), W'(:, j)),$$

after the columns of  $W'$  have been reordered so as to minimize the MRSa. The smaller the MRSa, the better the solution. For ALLS and SVCA, on a given data set, we run 30 trials and keep the median of the results. SSPA is deterministic so we only run it once. Unless stated otherwise, SVCA and SSPA are equipped with the median aggregation. In the following, we consider several experimental setups to highlight the property of these algorithms.

In Fig. 1, we test ALLS, SVCA, and SSPA with different values of the parameter  $p$ , when the noise  $\epsilon$  varies. Note that SVCA(1) and SSPA(1) are equivalent to

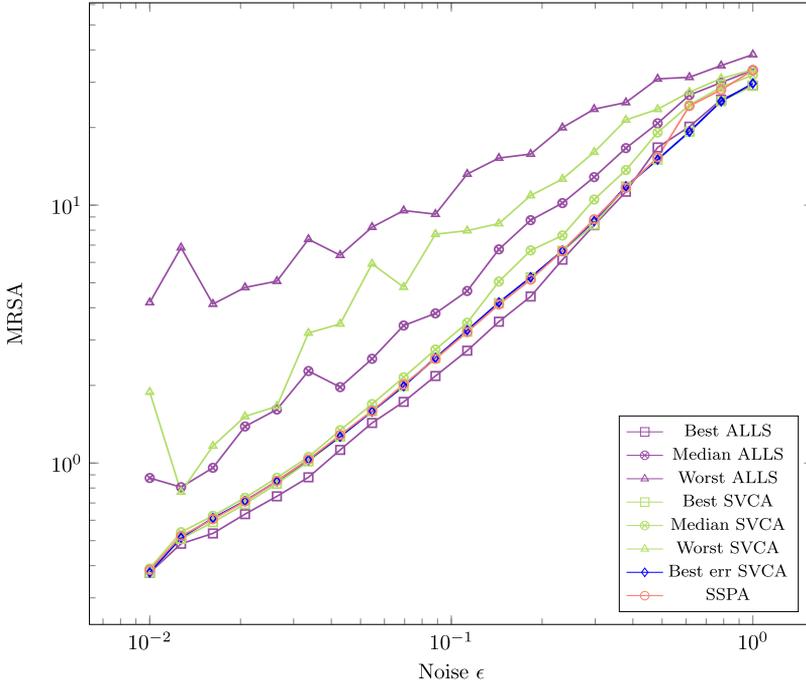


**Fig. 1.** Results for ALLS, SVCA, and SSPA for different values of  $p$ , when  $\epsilon$  varies, for fixed  $n = 1000$  and purity  $\alpha = 0.05$  ( $\delta = 2.7\%$ ). Values for ALLS and SVCA are the medians over 30 trials. Note that  $\text{ALLS}(1)=\text{SVCA}(1)=\text{VCA}$ . (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

their non-smoothed version VCA and SPA. Also note that  $\text{ALLS}(1)$  is equivalent to  $\text{SVCA}(1)$  and thus VCA. In this experiment, we observe that smoothing improves the algorithm performances. However, for ALLS, a parameter  $p$  set too large can in fact worsen the solution, especially when the noise level is small. Also, ALLS is outperformed by SVCA and SSPA; this will be confirmed in experiments on hyperspectral images in Section 5.2.

In Fig. 2, we compare the stability of ALLS and SVCA when  $\epsilon$  varies by showing the best, median, and worst result among 30 runs. We also compare them to the MRSA of the result of SVCA that has the smallest reconstruction error,  $Q_F(W')$ , and to SSPA which is deterministic. We see that the best result from ALLS is slightly better than other results. This is due to the use of the mean as an aggregation method, which works better with the centered Gaussian noise of the synthetic data,<sup>3</sup> see Fig. 4. However, the algorithm is less stable, as its median and worst results are worst than SVCA. With SVCA, the median results are close to the best. Also, the best results in terms of reconstruction error generally coincides with the best one in terms of MRSA, showing that the reconstruction

<sup>3</sup> Using the mean aggregation instead of the median aggregation in SVCA, its best MRSA result is always better than the best MRSA of ALLS in this experiment.



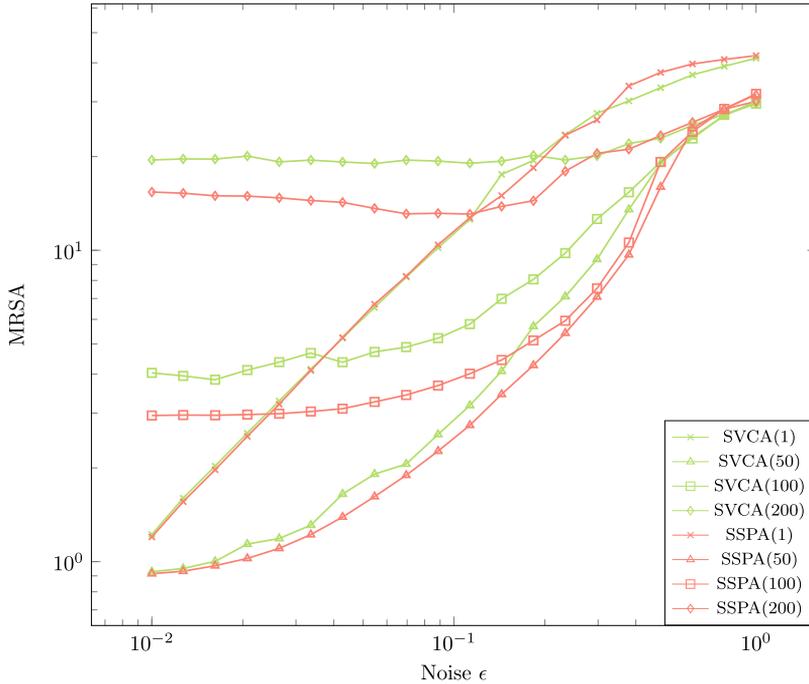
**Fig. 2.** Comparison of SSPA with best, median, and worst result for ALLS and SVCA among 30 runs, when  $\epsilon$  varies, for fixed  $n = 1000$ , purity  $\alpha = 0.05$  ( $\delta = 2.7\%$ ) and parameter  $p = 25$ . “Best err SVCA” represents the MRSA of the solution of SVCA that has the smallest relative reconstruction error.

error is a good proxy for the tested algorithms, as evoked in Section 4 (this will be useful when the groundtruth is unknown and the MRSA cannot be computed, for example in Section 5.2). The deterministic algorithm SSPA produces results that are slightly better than the median result of SVCA. However, the best result from SVCA is better than the result from SSPA for most noise levels.

In Fig. 3, we compare SVCA and SSPA with higher values of  $p$  when  $\epsilon$  varies. Again, we observe that the smoothing improves the algorithm performances, but an overestimated  $p$  worsens it. Interestingly, when the noise is very high (above 10%), smoothed algorithms outperform their non-smoothed counterpart even when  $p$  is overestimated. This is due to the fact that the value of  $p$  required to obtain the best estimation of  $W$  is not only determined by the purity but also by the noise level. For instance, consider a toy example with four data points and  $r = 2$ :

$$X = WH + N = W \begin{bmatrix} 1 & 0 & 0.99 & 0.01 \\ 0 & 1 & 0.01 & 0.99 \end{bmatrix} + N.$$

That is,  $x_3$  and  $x_4$  are not pure-pixel but are almost pure. Let us further assume  $N$  to follow a centered Gaussian law. If  $\epsilon = 0$  (noiseless mixing), the best estimation of  $W(:, 1)$  (resp.  $W(:, 2)$ ) from  $X$  is to extract  $X(:, 1)$  (resp.  $X(:, 2)$ ), yielding a perfect estimation. On the other hand, if  $\epsilon$  is large relatively to the distance of  $W(:, 1)$  and  $W(:, 2)$ , it is

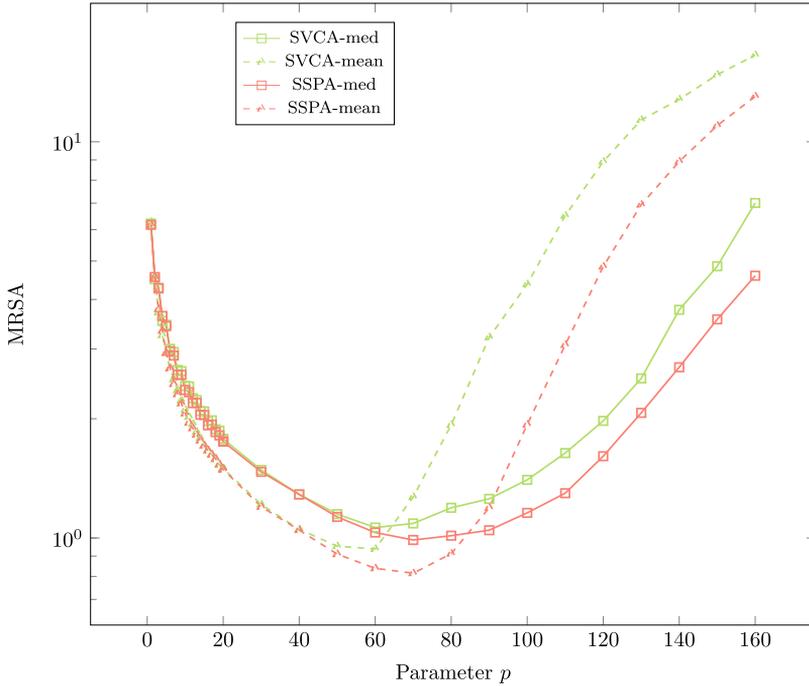


**Fig. 3.** Results for SVCA and SSPA for different values of  $p$ , when  $\epsilon$  varies, for fixed  $n = 1000$  and purity  $\alpha = 0.05$  ( $\delta = 2.7\%$ ). Values for SVCA are the medians over 30 trials.

better to choose as an estimate of  $W(:, 1)$  (resp.  $W(:, 2)$ ) the average – or median – of  $X(:, 1)$  and  $X(:, 3)$  (resp.  $X(:, 2)$  and  $X(:, 4)$ ), as the noise power is then divided by two.

In Fig. 4, we compare SVCA and SSPA equipped with either the median or the mean aggregation, for fixed data setup and when  $p$  varies. The reverse bell curve shows that the performance of the algorithms improves gradually as  $p$  grows, until it reaches an optimal value, after which the performance gradually worsens. We observe that the algorithms equipped with the median are more robust to an overestimation of  $p$ , but with the mean they are slightly better for smaller  $p$ . However the difference is small; this is expected as this synthetic data is generated with centered Gaussian noise, and as such the mean is expected to give the best estimation when  $p$  is well chosen. The difference is more obvious in hyperspectral images, for which the median aggregation always performs better than the mean aggregation; see Section 5.2 and Fig. 7.

In Fig. 5, we compare SVCA and SSPA equipped with either the median or the mean aggregation, for fixed data setup and when  $p$  varies, but where the synthetic data has been generated using the Poisson distribution instead of Gaussian noise. This data generation accounts for a counting process [36]. Given a noise level  $\epsilon$ , we first generate random matrices  $W$  and  $H$  as above. Then we compute  $c = \frac{1}{\text{mean}(WH)\epsilon^2}$ , where  $\text{mean}(WH)$  is the average of the entries of  $WH$ , and use the Matlab function `poissrnd` to generate a noisy data matrix  $X = \frac{1}{c}\text{poissrnd}(cWH)$ , that is, the entry  $X_{i,j}$  of  $X$  follows a Poisson distribution of parameter  $c(WH)_{i,j}$ , divided by  $c$ . This choice of  $c$  allows one to control

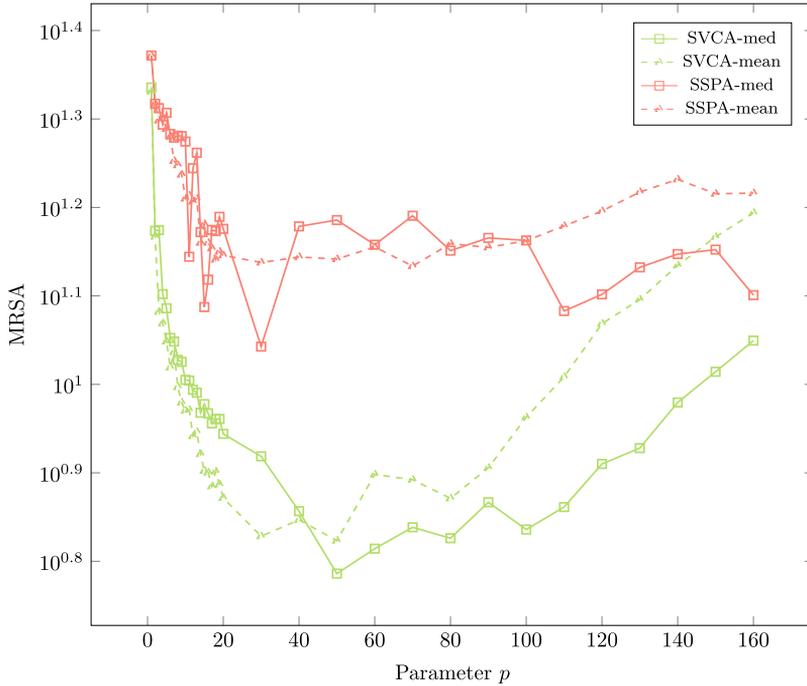


**Fig. 4.** Results for SVCA and SSPA using either the median or the mean to average points, when  $p$  varies, for fixed  $n = 1000$ , purity  $\alpha = 0.05$  ( $\delta = 2.7\%$ ), and noise  $\epsilon = 0.05$ . Values for SVCA are the medians over 30 trials.

the norm of the noise relative to the groundtruth  $WH$  (recall that the expectation of the Poisson distribution of parameter  $\lambda$  is  $\lambda$ , and the variance is  $\lambda$ ). The expected value of the entries of  $X$  remains unchanged (namely,  $\text{mean}(WH)$ ), while the standard deviation is equal to  $\epsilon \text{mean}(WH)$ . Hence, the entries of  $X$  are close to the interval  $[WH - \epsilon, WH + \epsilon]$ , and we have that  $\|X - WH\|_F / \|WH\|_F \approx \epsilon$ . We observe again that the smoothed algorithms outperform significantly VCA and SPA. A notable difference compared to experiments with Gaussian noise is that, with Poisson noise, the algorithms equipped with the median aggregation give, in general, better results than with the mean. The median aggregation is still more robust to an overestimation of  $p$ . As opposed to previous experiments, the median results of SVCA are better than the results of SSPA. This is expected because the selection step of SSPA uses the  $\ell_2$ -norm, which makes it more sensitive to non-Gaussian noise.

In Fig. 6, we compare SVCA and SSPA when  $p$  varies in setups with different values of purity  $\alpha$ . As expected, we observe that the shapes of the curves are similar, and that for a fixed noise level the value of the parameter  $p$  leading to the lowest MRSA decreases as the parameter  $\alpha$  increases.

To summarize, our synthetic experiments highlight that the two proposed algorithms consistently obtain better results than ALLS. They furthermore show that SVCA and SSPA outperform VCA and SPA when  $p$  is well chosen. Although the choice of  $p$  is



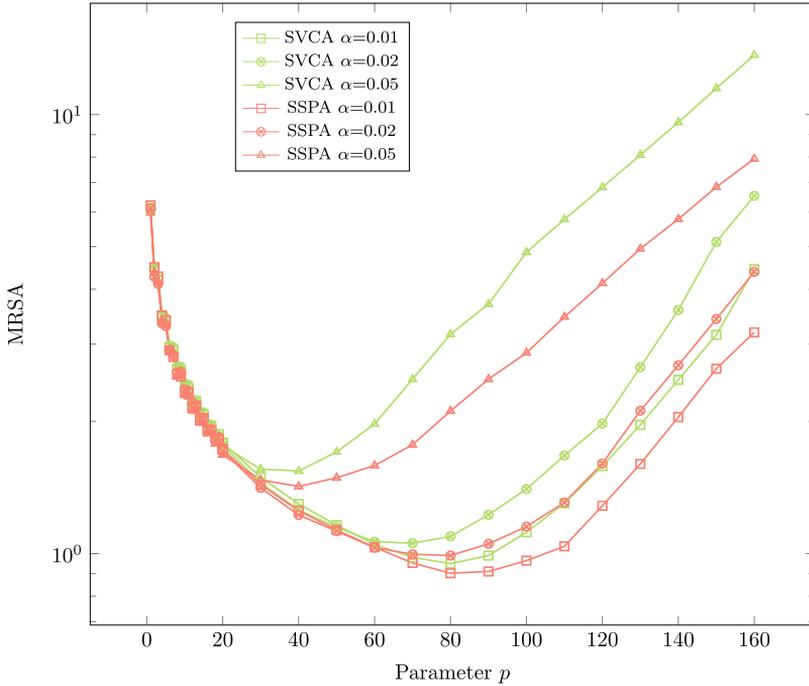
**Fig. 5.** Results for SVCA and SSPA using either the median or the mean to average points, when  $p$  varies, for Poisson noise of level  $\epsilon = 0.1$ , for fixed  $n = 1000$ , purity  $\alpha = 0.02$  ( $\delta = 5.9\%$ ). Values for SVCA are the medians over 30 trials.

important, as a bad value can worsen the results compared to the non-smoothed algorithms, the use of the median aggregation instead of the mean makes the algorithms less sensitive to this choice. The use of the median is also more adapted to sparse, Laplace or Poisson noises, whereas the mean gives better results in the case of centered Gaussian noise.

### 5.2. Hyperspectral images

In this section, we apply ALLS, SVCA, and SSPA to the unmixing of hyperspectral images, as described in Section 1. We consider three commonly used hyperspectral images,<sup>4</sup> San Diego, Urban, and Terrain. In hyperspectral data sets, extremely large values are commonly associated with sensor noise or interference. To avoid overfitting the factorizations to these interferences, the pixels corresponding to the 10 largest values of any wavelength range are zeroed out. Extreme pixels generally have extreme values in many wavelength ranges at once, so this preprocessing results in the removal of less than 0.1% pixels. The characteristics of these images are summarized in Table 2.

<sup>4</sup> Downloaded from <http://lesun.weebly.com/hyperspectral-data-set.html>.



**Fig. 6.** Results for SVCA and SSPA for different values of purity  $\alpha$ , when  $p$  varies, for fixed  $n = 1000$  and noise  $\epsilon = 0.05$ . Values for SVCA are the medians over 30 trials.

**Table 2**  
Summary of the hyperspectral images studied in this work.

Data set	$m$	$n$	$r$	Pixels zeroed out
San Diego	158	$400 \times 400 = 160000$	8	19
Urban	162	$307 \times 307 = 94249$	6	68
Terrain	188	$500 \times 307 = 153500$	6	107

Given a data matrix  $X \in \mathbb{R}^{m \times n}$ , we compute  $W \in \mathbb{R}^{m \times r}$  with the three algorithms. We then compute for each algorithm  $H \in \mathbb{R}^{r \times n}$  with a standard coordinate descent algorithm [37] and measure the relative reconstruction error  $\|X - WH\|_F / \|X\|_F$ . The smaller the error, the better the solution.

Some works such as [38] proposed groundtruths for these hyperspectral images, but they are computed using numerical methods and as such do not necessarily represent reality. Therefore, we lack a reference to assess the quality of the reconstruction, for example by measuring the MRSA. This is why we use the relative reconstruction error as the criterion for the quality of a solution. This is a satisfying criterion, as illustrated in Fig. 2.

In Table 3, we report results from the experiments. We observe that, when  $p > 1$ , the result is always improved. When  $p$  is too large, however, the solution can be worse than for  $p = 1$ , as expected. We observe that the best  $p$  varies between the algorithms.

**Table 3**

Relative reconstruction errors ( $\min_{H \geq 0} \|X - WH\|_F / \|X\|_F$ ) resulting from the unmixing of hyperspectral images with ALLS, SVCA, and SSPA, with different values of parameter  $p$ . SVCA(1) and SSPA(1) are equivalent to VCA and SPA. For non-deterministic algorithms ALLS and SVCA, we show the minimum, median, standard deviation, and maximum of the error over 30 trials.

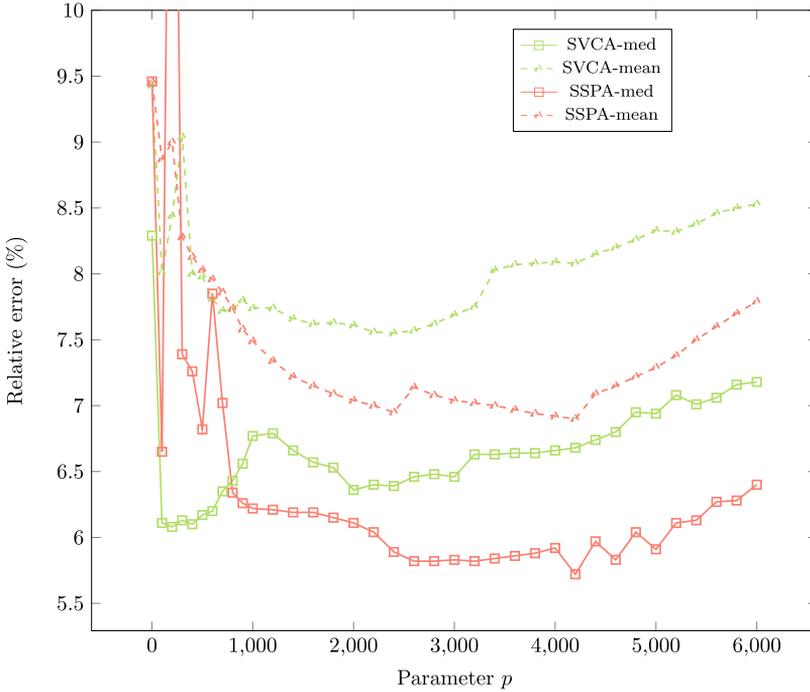
	p	San Diego			Urban			Terrain		
		Min	Med ± std	Max	Min	Med ± std	Max	Min	Med ± std	Max
ALLS	1	4.72	5.60 ± 0.67	8.25	5.39	9.14 ± 1.93	12.26	3.94	4.88 ± 0.73	7.08
	100	<b>4.27</b>	<b>5.35</b> ± 1.72	10.91	<b>6.37</b>	<b>9.28</b> ± 3.18	19.40	<b>3.84</b>	4.87 ± 0.87	6.85
	1000	4.64	6.14 ± 1.12	8.68	6.78	9.71 ± 2.16	14.20	<b>3.84</b>	4.71 ± 1.19	8.81
	2000	4.87	5.91 ± 1.62	11.79	6.96	9.93 ± 1.60	12.85	3.96	4.89 ± 0.88	7.63
	5000	5.51	7.42 ± 2.64	13.88	7.68	10.37 ± 1.89	14.98	4.28	5.26 ± 0.80	6.88
SVCA	1	3.95	5.42 ± 0.61	6.90	6.25	9.13 ± 1.78	12.23	4.03	5.11 ± 1.25	8.70
	100	<b>3.44</b>	4.92 ± 0.77	6.96	<b>5.08</b>	<b>6.10</b> ± 1.27	10.13	3.52	4.04 ± 0.67	6.52
	1000	3.82	4.95 ± 0.59	6.82	5.82	6.77 ± 1.23	10.84	<b>3.18</b>	<b>3.92</b> ± 0.38	4.70
	2000	3.73	<b>4.40</b> ± 0.51	5.81	5.66	6.36 ± 0.67	7.83	3.38	4.12 ± 0.45	4.95
	5000	4.01	4.66 ± 0.73	7.01	5.69	6.94 ± 1.21	11.74	3.70	4.19 ± 0.30	4.82
SSPA	1		5.90			9.46			5.01	
	100		9.29			6.65			<b>4.03</b>	
	1000		5.82			6.22			8.05	
	2000		<b>4.32</b>			6.11			7.86	
	5000		4.65			<b>5.91</b>			5.38	

For example, with Terrain, ALLS and SVCA perform best with  $p = 1000$  while SSPA performs best with  $p = 100$ . Larger values of  $p$  seem to give more stable results, as it produces solutions with smaller deviations. SVCA outperforms ALLS in all cases. SSPA performance is comparable to SVCA, and generally produces a better result than the median of SVCA, but never better than the best result obtained by SVCA.

In a few cases, SSPA produces solutions with a large error, for example in San Diego for  $p = 100$  and Terrain for  $p = 1000$ . We believe this behavior to originate from small groups of points with a very large norm, that could correspond to a rare material or to interference.

In Fig. 7, we compare SVCA and SSPA using either the median or the mean for the unmixing of the hyperspectral image Urban, with a varying  $p$ . We observe that  $p > 1$  always leads to better results for all algorithms and all data sets. Also, the median aggregation almost always gives better results than the mean. While the curves are not as regular as with synthetic data, we observe a similar tendency that the solution improves when  $p$  grows, until a certain point or zone after which it worsens again. However, SSPA-med has an irregular behavior for  $p = 200$ . This can be explained by the fact that SSPA is a greedy algorithm, so if it makes a bad choice in the first iterations, it will likely never compensate. Also, it is deterministic, so the error is not averaged over several runs.

In Fig. 8, we show the abundances maps corresponding to the unmixing of Urban; the corresponding spectral signatures can be found in the supplementary material. They indicate the proportion of every of the 6 extracted endmembers in the pixels of the image. We see that the smoothed algorithms obtain a better separation than the non-smoothed ones. For example, the fourth endmember extracted by SVCA and SSPA corresponds to grass, and it is well separated by these algorithms, while VCA and SPA mix it with



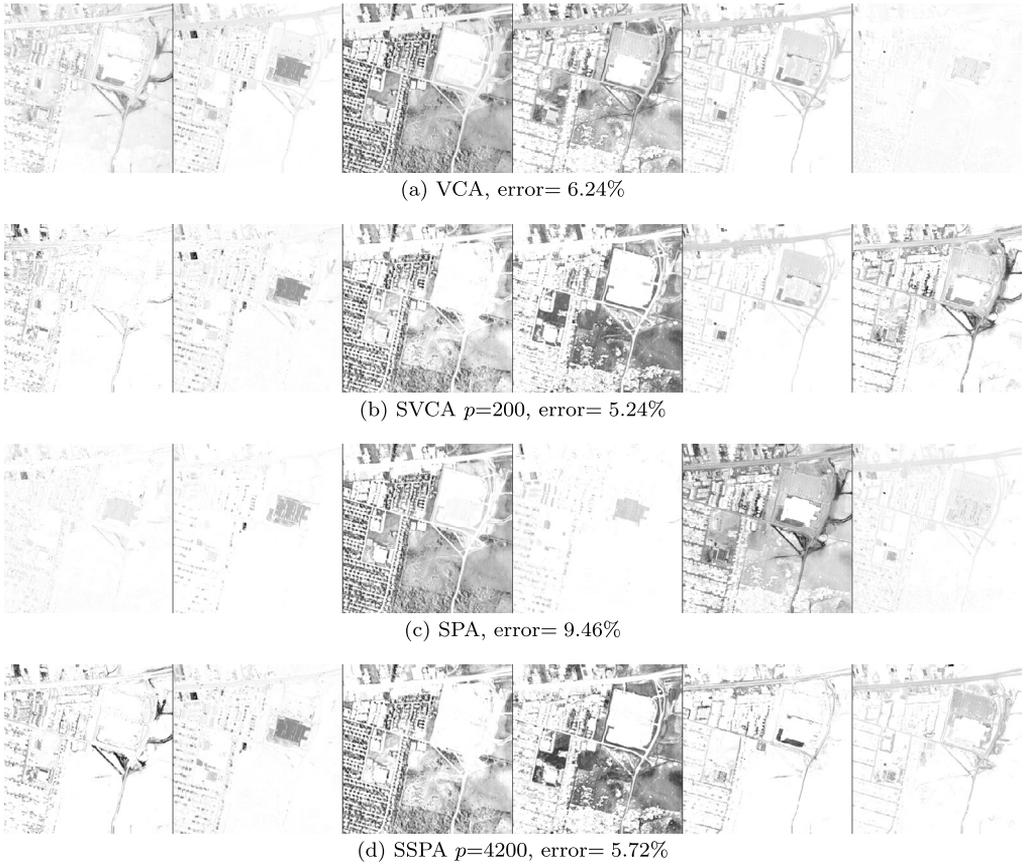
**Fig. 7.** Results of the unmixing of the hyperspectral image Urban. Values for SVCA are the medians over 30 trials. One point is out of the plot; for  $p = 200$ , SVCA-med has an error of 14.55%.

asphalt and dirt. The second endmember extracted by SVCA and SSPA corresponds to metallic rooftops, and it is well separated while VCA mixes it with other materials and SPA does not clearly identify it and produces a blurred picture.

In the supplementary material, we provide additional experiments, namely, the sensitivity to  $p$  of SVCA and SSPA, and the abundance maps and spectral signatures of the different algorithms for the hyperspectral images Terrain and San Diego.

### 5.3. Facial images

A popular application of NMF is the extraction of facial features, such as eyes, noses and lips, in a set of images, because it was first introduced in the seminal paper of Lee and Seung [11]. In this application, each row of  $X$  is a vectorized facial image, and an NMF,  $WH \approx X$ , extracts vectorized facial features as the rows of  $H$ . Separability makes sense in this context: it requires that, for each facial feature, there is at least one pixel that is only contained in that facial feature; see [4, p. 209] for a discussion and a numerical example. For the proximate latent points assumption to make sense, we need for each facial feature several pixels that appear mostly in that feature. Hence the more pixels in the images, the more likely this assumption will hold and the larger the parameter  $p$  can be chosen; this will be illustrated on the numerical examples below.



**Fig. 8.** Abundance maps of the unmixing of the Urban hyperspectral images (that is, reshaped rows of  $H$ ) with different algorithms. Endmembers have been reordered for easier comparison. Parameters  $p$  have been chosen as the best from Fig. 7. Error corresponds to  $\min_{H \geq 0} \|X - WH\|_F / \|X\|_F$ . For VCA and SVCA, we show the best solution over 30 trials.

Let us now apply ALLS, SVCA, and SSPA to extract facial features on four widely used data sets; see Table 4. Note that the first two data sets (CBCL, Frey) have significantly fewer pixels than the last two (ORL, UMIST) and hence we expect that  $p$  can be chosen larger for the last two. We choose  $r = 10$  for CBCL and Frey, and  $r = 20$  for ORL and UMIST.

Table 5 shows the results from the experiments of facial images data sets. The reconstruction errors are roughly between 15% and 25%, meaning that this type of data does not follow closely the NMF model. However, this model is still able to extract meaningful features as we show below, and the smoothed separable NMF algorithms do produce better results than the non-smoothed ones. We observe that the result is consistently improved for  $p > 1$  and  $p$  not too large. SVCA outperforms ALLS in all cases. The best solution from SVCA is always better than the solution from SSPA, and the median solution from SVCA is better than the solution from SSPA for all data sets except Umist. In this experiment the best  $p$  is the same for SVCA and SSPA in most cases. For the

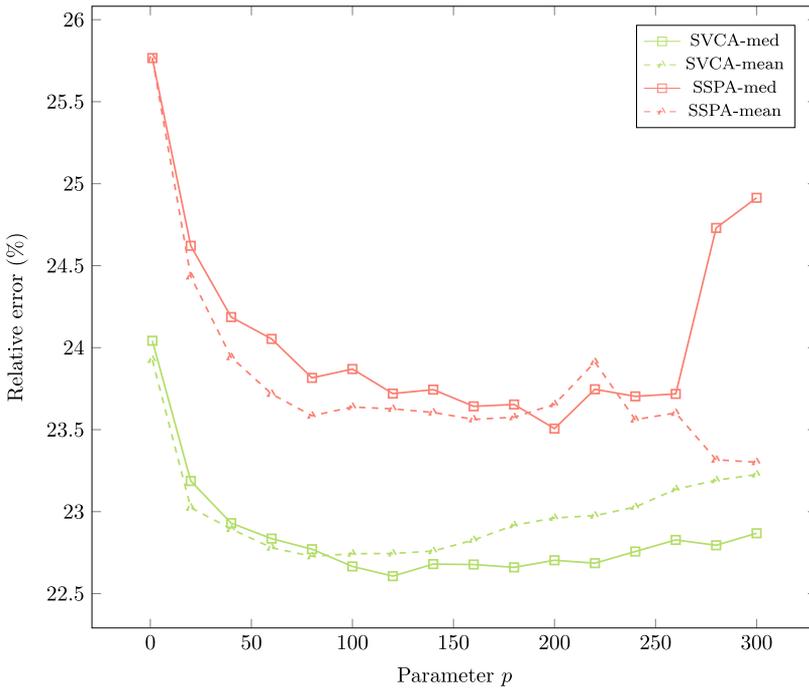
**Table 4**  
Summary of the facial images studied in this work.

Data set	$m$ (# faces)	$n$ (# pixels)	$r$
CBCL <sup>1</sup>	2429	$19 \times 19 = 361$	10
Frey <sup>2</sup>	1965	$28 \times 20 = 560$	10
ORL <sup>3</sup>	400	$112 \times 92 = 10304$	20
UMIST <sup>2</sup>	565	$112 \times 92 = 10304$	20

<sup>1</sup> <https://gitlab.com/ngillis/nmfbook/-/blob/master/data%20sets/CBCL.mat>.

<sup>2</sup> <http://www.cs.toronto.edu/~roweis/data.html>.

<sup>3</sup> <https://cam-orl.co.uk/facedatabase.html>.



**Fig. 9.** Results of the processing of the facial images data set ORL. Values for SVCA are the medians over 30 trials.

data sets with fewer pixels (CBCL, Frey), a value of  $p$  larger than 20 deteriorates the performance of SVCA and SSPA. This makes sense since there are fewer than 600 pixels, and hence the rule of thumb that  $p \ll \frac{n}{r}$  is not respected (recall that  $p = \frac{n}{r}$  would be the extreme case where there are exactly  $\frac{n}{r}$  data points close to each column of  $W$ ). For the data sets with more pixels (ORL, UMIST),  $p$  can be rather large and still leads to a decrease in the error for  $p$  up to 200.

Fig. 9 reports the relative errors for SVCA and SSPA on the data set ORL, using the mean and the median aggregation, for various values of  $p$ . Figures for other data sets are included in the supplementary material. We observe that, as for the synthetic data sets and hyperspectral images, the smoothed algorithms allow one to significantly

**Table 5**

Relative reconstruction errors ( $\min_{H \geq 0} \|X - WH\|_F / \|X\|_F$ ) in percent resulting from the processing of facial image datasets with ALLS, SVCA, and SSPA, with different values of parameter  $p$ . SVCA(1) and SSPA(1) are equivalent to VCA and SPA. For non-deterministic algorithms ALLS and SVCA, we show the minimum, median, standard deviation, and maximum of the error over 30 trials.

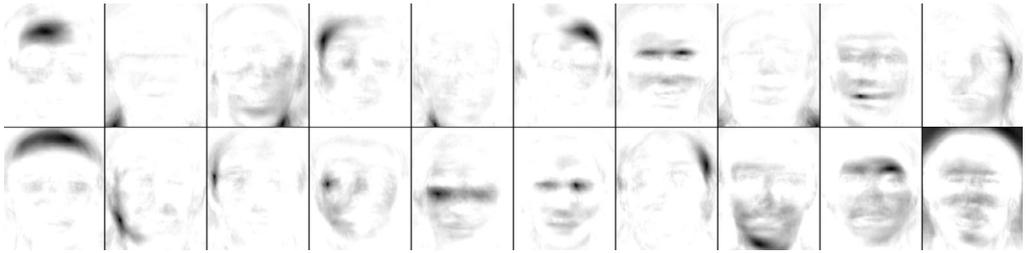
	p	CBCL			Frey		
		Min	Med ± std	Max	Min	Med ± std	Max
ALLS	1	19.36	20.08 ± 0.33	20.95	21.05	22.38 ± 0.58	23.83
	5	<b>18.53</b>	19.60 ± 0.51	20.41	<b>20.75</b>	<b>22.32</b> ± 0.73	23.96
	10	18.67	<b>19.44</b> ± 0.78	21.88	21.30	22.85 ± 0.85	24.57
	20	19.34	20.28 ± 1.17	23.83	22.30	24.05 ± 1.10	27.42
SVCA	1	19.25	19.81 ± 0.31	20.48	20.99	22.18 ± 0.64	24.02
	5	18.72	19.21 ± 0.34	20.14	20.69	21.57 ± 0.52	22.85
	10	18.48	<b>18.84</b> ± 0.35	19.76	<b>20.68</b>	<b>21.29</b> ± 0.43	22.57
	20	<b>18.44</b>	18.95 ± 0.25	19.55	20.79	21.70 ± 0.45	22.71
SSPA	1		20.76			22.57	
	5		20.05			21.65	
	10		<b>19.63</b>			<b>21.32</b>	
	20		20.83			21.58	
	p	ORL			Umist		
		Min	Med ± std	Max	Min	Med ± std	Max
ALLS	1	23.64	24.08 ± 0.18	24.35	15.38	16.00 ± 0.25	16.46
	50	<b>22.87</b>	<b>23.62</b> ± 0.42	24.92	<b>15.34</b>	<b>15.91</b> ± 0.32	16.48
	200	23.34	24.33 ± 0.60	25.93	16.27	16.96 ± 0.54	18.87
	400	24.50	25.31 ± 0.45	26.56	17.28	18.21 ± 0.49	19.07
SVCA	1	23.34	24.04 ± 0.32	24.70	15.40	15.95 ± 0.24	16.28
	50	22.31	22.89 ± 0.32	23.63	14.95	<b>15.37</b> ± 0.28	16.08
	200	<b>22.22</b>	<b>22.70</b> ± 0.29	23.53	<b>14.90</b>	<b>15.37</b> ± 0.25	15.91
	400	22.42	23.00 ± 0.34	23.86	15.29	15.85 ± 0.30	16.46
SSPA	1		25.76			16.56	
	50		24.14			15.72	
	200		<b>23.50</b>			<b>14.98</b>	
	400		24.50			15.40	

and consistently reduce the error, given that  $p$  is not too large. The difference between using the mean and the median is not significant, although the mean seems to perform slightly better for SVCA. The reason is that there is no outlier in these data sets, and the Gaussian assumption for the noise is reasonable. As opposed to the experiment on hyperspectral data, here SVCA outperforms SSPA on average.

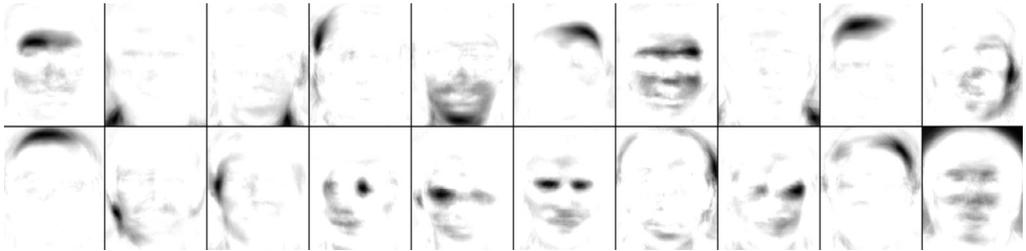
As an illustration, Fig. 10 shows the facial features extracted by VCA, SVCA, SPA, SSPA. They represent the proportion of every extracted facial feature in the pixels of the images. We observe that SVCA and SSPA, apart from reducing the error as discussed above, also produce sharper facial features. These features are more clearly delimited and separated, they grasp more finely the parts of the human face such as eyes, nose, and chin.

#### 5.4. Discussion

Let us make some additional points regarding SVCA and SSPA.



(a) VCA, error= 23.35%

(b) SVCA  $p=120$ , error= 22.18%

(c) SPA, error= 25.77%

(d) SSPA  $p=200$ , error= 23.51%

**Fig. 10.** Abundance maps of the unmixing of the facial images data set ORL (that is, reshaped rows of  $H$ ) with different algorithms. Features have been reordered for easier comparison. Parameters  $p$  have been chosen as the best from Fig. 9. Error corresponds to  $\min_{H \geq 0} \|X - WH\|_F / \|X\|_F$ . For VCA and SVCA, we show the best solution over 30 trials.

*Which algorithm should one use?* SVCA allows to generate different solutions, among which the best solution w.r.t. the reconstruction error can be found. Therefore, in practice and when time and resources allow, we recommend running SSPA once and SVCA several times, with different values of  $p$ , and keep the best solution.

*Which aggregation should one use?* If the noise present in the data is assumed to follow a Gaussian distribution, the mean aggregation is more adapted. In case of sparse noise or in the presence of outliers, then the median aggregation will perform better. Apart from the average and the median, other aggregation methods could perform better depending on the noise statistics and data set at hand. For instance, [39] uses an aggregation on manifold in the context of sparse matrix factorization to better take into account the structure of the columns of  $W$ . Exploring more advanced aggregation methods for smoothed separable NMF would be an interesting direction for future research.

*How to select  $p$ ?* Choosing a value for the parameter  $p$  is crucial and not trivial. Currently, one needs to have a prior knowledge on the number of data points close to the vertices, or use an empirical trial-and-error method. A strategy to determine the best  $p$  for a given setting is a particularly interesting direction of research. It would also be particularly meaningful to consider a different value of  $p$  for every vertex, as the number of proximal latent points typically varies for each vertex.

*Spectral variability in blind HU.* In blind HU, an issue with separable NMF algorithms is that they identify a single pixel to represent a material. It is however well-known that the spectral signature of an endmember may vary across the pixels of the image, for example because of differences in light intensity or orientation. This is known as *spectral variability*. By construction, most separable NMF algorithms, such as VCA and SPA, will identify pure pixels that do not represent well the average behavior of a material, but rather a pure pixel located at the boundary of the convex hull of the variations of the spectral signature of that endmember. Therefore, working on the smoothed data set, which averages every subset of  $p$  data points, allows to better represent this average behavior.

*Designing other smoothed separable NMF algorithms.* The proximal latent points assumption could be used to generalize other separable NMF algorithms. In this paper, we validated the idea on the two most widely used separable NMF algorithms, namely VCA and SPA, but the same idea can be applied to any separable NMF algorithms such as SNPA [20].

## 6. Conclusion

In this work, we investigated the smoothed separable NMF model, that strengthens the separability assumption by assuming the presence of several data points nearby each column of the basis matrix  $W$ . Inspired by the existing algorithm ALLS, we developed smoothed variants of two separable NMF algorithms, namely VCA and SPA. Empirically, we showed that our smoothed methods outperform both the non-smoothed ones and ALLS, for both synthetic data sets and for two real-world applications, namely the unmixing of hyperspectral images and feature extraction in facial images data sets. We

showed that the proximal latent points assumption is verified in hyperspectral images and in facial images data sets, and that smoothed separable NMF algorithms are a more effective tool for these applications.

Further works include the design of other smoothed separable NMF algorithms, and the use of SVCA and SSPA for other applications, in particular in machine learning [9,10]. Exploring new aggregation methods and smarter ways to choose the parameter  $p$  are also relevant directions for future research.

### Declaration of competing interest

There is no competing interest.

### Data availability

We have made available the code and data on a Git repository whose link is in the manuscript.

### Acknowledgements

We sincerely thank the reviewer for taking the time to carefully read the paper, and for the very detailed and insightful feedback that helped us improve our paper.

NN and NG acknowledge the support by the Fonds De La Recherche Scientifique - FNRS and the Fonds Wetenschappelijk Onderzoek - Vlaanderen (FWO) under EOS Project no. O005318F-RG47. NG also acknowledges the support of The Francqui Foundation.

### Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.laa.2023.07.013>.

### References

- [1] I.T. Jolliffe, *Principal Component Analysis*, Springer, 2011.
- [2] A. Tharwat, Independent component analysis: an introduction, *Appl. Comput. Inform.* 17 (2) (2021) 222–249.
- [3] R. Gribonval, M. Zibulevsky, Sparse component analysis, in: *Handbook of Blind Source Separation*, Elsevier, 2010, pp. 367–420.
- [4] N. Gillis, *Nonnegative Matrix Factorization*, SIAM, Philadelphia, 2020.
- [5] M. Abdolali, N. Gillis, Simplex-structured matrix factorization: sparsity-based identifiability and provably correct algorithms, *SIAM J. Math. Data Sci.* 3 (2) (2021) 593–623.
- [6] C.H. Ding, T. Li, M.I. Jordan, Convex and semi-nonnegative matrix factorizations, *IEEE Trans. Pattern Anal. Mach. Intell.* 32 (1) (2008) 45–55.
- [7] M. Dereziński, R. Khanna, M.W. Mahoney, Improved guarantees and a multiple-descent curve for column subset selection and the Nystrom method, *Adv. Neural Inf. Process. Syst.* 33 (2020) 4953–4964.

- [8] Y. Shitov, Column subset selection is np-complete, *Linear Algebra Appl.* 610 (2021) 52–58.
- [9] C. Bhattacharyya, R. Kannan, Finding a latent  $k$ -simplex in  $o^*(k \cdot \text{nnz}(\text{data}))$  time via subset smoothing, in: *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, 2020, pp. 122–140.
- [10] A. Bakshi, C. Bhattacharyya, R. Kannan, D.P. Woodruff, S. Zhou, Learning a latent simplex in input sparsity time, in: *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021, pp. 1–11.
- [11] D.D. Lee, H.S. Seung, Learning the parts of objects by non-negative matrix factorization, *Nature* 401 (6755) (1999) 788–791.
- [12] A. Ozerov, C. Févotte, Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation, *IEEE Trans. Audio Speech Lang. Process.* 18 (3) (2009) 550–563.
- [13] M. Thiel, N. Sauwen, T. Khamiakova, T. Maes, B. Govaerts, Comparison of chemometrics strategies for the spectroscopic monitoring of active pharmaceutical ingredients in chemical reactions, *Chemom. Intell. Lab. Syst.* 211 (2021) 104273.
- [14] V.P. Pauca, J. Piper, R.J. Plemmons, Nonnegative matrix factorization for spectral data analysis, *Linear Algebra Appl.* 416 (1) (2006) 29–47.
- [15] N. Gillis, R.J. Plemmons, Sparse nonnegative matrix underapproximation and its application to hyperspectral image analysis, *Linear Algebra Appl.* 438 (10) (2013) 3991–4007.
- [16] 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.
- [17] W.-K. Ma, J.M. Bioucas-Dias, T.-H. Chan, N. Gillis, P. Gader, A.J. Plaza, A. Ambikapathi, C.-Y. Chi, A signal processing perspective on hyperspectral unmixing: insights from remote sensing, *IEEE Signal Process. Mag.* 31 (1) (2014) 67–81.
- [18] J.M. Nascimento, J.M. Dias, Vertex component analysis: a fast algorithm to unmix hyperspectral data, *IEEE Trans. Geosci. Remote Sens.* 43 (4) (2005) 898–910.
- [19] U. Araújo, B. Saldanha, R. Galvão, T. Yoneyama, H. Chame, V. Visani, The successive projections algorithm for variable selection in spectroscopic multicomponent analysis, *Chemom. Intell. Lab. Syst.* 57 (2) (2001) 65–73.
- [20] N. Gillis, Successive nonnegative projection algorithm for robust nonnegative blind source separation, *SIAM J. Imaging Sci.* 7 (2) (2014) 1420–1450.
- [21] X. Fu, W.-K. Ma, K. Huang, N.D. Sidiropoulos, Blind separation of quasi-stationary sources: exploiting convex geometry in covariance domain, *IEEE Trans. Signal Process.* 63 (9) (2015) 2306–2320.
- [22] R. Ge, J. Zou, Intersecting faces: non-negative matrix factorization with new guarantees, in: *Proceedings of the 32nd International Conference on Machine Learning*, 2015, pp. 2295–2303.
- [23] C.-H. Lin, C.-Y. Chi, Y.-H. Wang, T.-H. Chan, A fast hyperplane-based minimum-volume enclosing simplex algorithm for blind hyperspectral unmixing, *IEEE Trans. Signal Process.* 64 (8) (2015) 1946–1961.
- [24] C.-H. Lin, R. Wu, W.-K. Ma, C.-Y. Chi, Y. Wang, Maximum volume inscribed ellipsoid: a new simplex-structured matrix factorization framework via facet enumeration and convex optimization, *SIAM J. Imaging Sci.* 11 (2) (2018) 1651–1679.
- [25] R. Wu, W.-K. Ma, Y. Li, A.M.-C. So, N.D. Sidiropoulos, Probabilistic simplex component analysis, preprint arXiv:2103.10027.
- [26] J.W. Boardman, F.A. Kruse, R.O. Green, Mapping target signatures via partial unmixing of AVIRIS data, in: *Proc. Summary JPL Airborne, Earth Science Workshop*, Pasadena, CA, 1995, pp. 23–26.
- [27] M.E. Winter, N-FINDR: an algorithm for fast autonomous spectral end-member determination in hyperspectral data, in: *Proc. SPIE Conf. Imaging Spectrometry V*, vol. 3753, International Society for Optics and Photonics, 1999, pp. 266–276.
- [28] J.-H. Jiang, Y. Liang, Y. Ozaki, Principles and methodologies in self-modeling curve resolution, *Chemom. Intell. Lab. Syst.* 71 (1) (2004) 1–12.
- [29] S. Arora, R. Ge, R. Kannan, A. Moitra, Computing a nonnegative matrix factorization—provably, in: *Proc. of the 44th Symp. on Theory of Computing (STOC '12)*, 2012, pp. 145–162.
- [30] B. Recht, C. Re, J. Tropp, V. Bittorf, Factoring nonnegative matrices with linear programs, in: *Advances in Neural Information Processing Systems (NIPS)*, 2012, pp. 1214–1222.
- [31] N. Gillis, Robustness analysis of Hottopixx, a linear programming model for factoring nonnegative matrices, *SIAM J. Matrix Anal. Appl.* 34 (3) (2013) 1189–1212.
- [32] N. Gillis, R. Luce, Robust near-separable nonnegative matrix factorization using linear optimization, *J. Mach. Learn. Res.* 15 (1) (2014) 1249–1280.

- [33] N. Gillis, S.A. Vavasis, Fast and robust recursive algorithms for separable nonnegative matrix factorization, *IEEE Trans. Pattern Anal. Mach. Intell.* 36 (4) (2013) 698–714.
- [34] S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu, M. Zhu, A practical algorithm for topic modeling with provable guarantees, in: *Proceedings of the 30th International Conference on Machine Learning*, 2013, pp. 280–288.
- [35] J.M. Nascimento, J.M. Bioucas-Dias, Hyperspectral unmixing based on mixtures of Dirichlet components, *IEEE Trans. Geosci. Remote Sens.* 50 (3) (2011) 863–878.
- [36] S.W. Hasinoff, Photon, Poisson noise, <http://people.csail.mit.edu/hasinoff/pubs/hasinoff-photon-2011-preprint.pdf>, 2014.
- [37] N. Gillis, F. Glineur, Accelerated multiplicative updates and hierarchical ALS algorithms for non-negative matrix factorization, *Neural Comput.* 24 (4) (2012) 1085–1105.
- [38] F. Zhu, Hyperspectral unmixing: ground truth labeling, datasets, benchmark performances and survey, preprint arXiv:1708.05125.
- [39] C. Kervazo, T. Liaudat, J. Bobin, Faster and better sparse blind source separation through mini-batch optimization, *Digit. Signal Process.* 106 (2020) 102827.