

Robust Volume Minimization-Based Matrix Factorization for Remote Sensing and Document Clustering

Xiao Fu, *Member, IEEE*, Kejun Huang, *Student Member, IEEE*, Bo Yang, *Student Member, IEEE*, Wing-Kin Ma, *Senior Member, IEEE*, and Nicholas D. Sidiropoulos, *Fellow, IEEE*

Abstract—This paper considers *volume minimization (VolMin)*-based structured matrix factorization. VolMin is a factorization criterion that decomposes a given data matrix into a basis matrix times a structured coefficient matrix via finding the minimum-volume simplex that encloses all the columns of the data matrix. Recent work showed that VolMin guarantees the identifiability of the factor matrices under mild conditions that are realistic in a wide variety of applications. This paper focuses on both theoretical and practical aspects of VolMin. On the theory side, exact equivalence of two independently developed sufficient conditions for VolMin identifiability is proven here, thereby providing a more comprehensive understanding of this aspect of VolMin. On the algorithm side, computational complexity and sensitivity to outliers are two key challenges associated with real-world applications of VolMin. These are addressed here via a new VolMin algorithm that handles volume regularization in a computationally simple way, and automatically detects and iteratively downweights outliers, simultaneously. Simulations and real-data experiments using a remotely sensed hyperspectral image and the Reuters document corpus are employed to showcase the effectiveness of the proposed algorithm.

Index Terms—Document clustering, hyperspectral unmixing, identifiability, matrix factorization, robustness against outliers, simplex-volume minimization (VolMin).

I. INTRODUCTION

STRUCTURED MATRIX FACTORIZATION (SMF) has been a popular tool in signal processing and machine learning. For decades, factorization models such as the singular value decomposition (SVD) and eigen-decomposition have been applied for dimensionality reduction (DR), subspace estimation, noise suppression, feature extraction, etc. Motivated by the influential paper of Lee and Seung [2], new SMF models such as *nonnegative matrix factorization* (NMF) have drawn much

attention, since they are capable of not only reducing dimensionality of the collected data, but also retrieving loading factors that have physically meaningful interpretations.

In addition to NMF, some related SMF models have attracted considerable interest in recent years. The remote sensing community has spent much effort on a class of factorizations where the columns of one factor matrix are constrained to lie in the unit simplex [3]. The same SMF model has also been utilized for document clustering [4], and, most recently, multi-sensor array processing and blind separation of power spectra for dynamic spectrum access [5], [6].

The first key question concerning SMF lies in identifiability—when does a factorization model or criterion admit unique solution in terms of its factors? Identifiability is important in applications such as parameter estimation, feature extraction, and signal separation. In recent years, identifiability conditions have been investigated for the NMF model [7]–[9]. An undesirable property of NMF highlighted in [9] is that identifiability hinges on both loading factors containing a certain number of zeros. In many applications, however, there is at least one factor that is dense. In hyperspectral unmixing (HU), for example, the basis factor (i.e., the spectral signature matrix) is always dense. On the other hand, very recent work [5], [10] showed that the SMF model with the coefficient matrix columns lying in the unit simplex admits much more relaxed identifiability conditions. Specifically, Fu *et al.* [5] and Lin *et al.* [10] proved that, under some realistic conditions, unique loading factors (up to column permutations) can be obtained by finding a minimum-volume enclosing simplex of the data vectors. Notably, these identifiability conditions of the so-called *volume minimization* (VolMin) criterion allow working with dense basis matrix factors; in fact, the model does not impose any constraints on the basis matrix except for having full-column rank. Since the NMF model can be recast as (viewed as a special case of) the above SMF model [11], such results suggest that VolMin is an attractive alternative to NMF for the wide range of applications of NMF and beyond.

Compared to NMF, VolMin-based matrix factorization is computationally more challenging. The notable prior works in [12] and [13] formulated VolMin as a constrained (log-)determinant minimization problem, and applied successive convex optimization and alternating optimization to deal with it, respectively. The major drawback of these pioneering works is that the algorithms were developed under a noiseless setting, and thus only work well for high signal-to-noise ratio (SNR) cases. Also, these algorithms work in the dimension-reduced domain, but the DR process may be sensitive to outliers and

Manuscript received March 3, 2016; revised June 8, 2016 and July 28, 2016; accepted August 6, 2016. Date of publication August 25, 2016; date of current version October 6, 2016. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Wee Peng Tay. This work was supported in part by the National Scientific Foundation (NSF) under Project NSF-ECCS 1608961 and in part by the Hong Kong Research Grants Council under Project CUHK 14205414. A part of this work was presented in the *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing 2016* [1].

X. Fu, K. Huang, B. Yang, and N. D. Sidiropoulos are with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, USA (e-mail: xfu@umn.edu; huang663@umn.edu; yang4173@umn.edu; nikos@umn.edu).

W.-K. Ma is with the Department of Electronic Engineering, The Chinese University of Hong Kong, Shatin, N.T., Hong Kong (e-mail: wkma@ee.cuhk.edu.hk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2016.2602800

modeling errors. The work [14] took noise into consideration, but the algorithm is computationally prohibitive and has no guarantee of convergence. Some other algorithms [15], [16] work in the original data domain, and deal with a volume-regularized data fitting problem. Such a formulation can tolerate noise to a certain level, but is harder to tackle than those in [12], [13], [16]—volume regularizers typically introduce extra difficulty to an already very hard bilinear fitting problem.

The second major challenge of implementing VolMin is that the VolMin criterion is very sensitive to outliers: it has been noted in the literature that even a single outlier can make the VolMin criterion fail [3]. However, in real-world applications, outlying measurements are commonly seen: in HU, pixels that do not always obey the nominal model are frequently spotted because of the complicated physical environment [17]; and in document clustering, articles that are difficult to be classified to any known category may also act like outliers. The algorithm in [18] is the state-of-the-art VolMin algorithm that takes outliers into consideration. It imposes a ‘soft penalty’ on outliers that lie outside the simplex that is sought, thereby allowing the existence of some outliers and achieving robustness. The algorithm works fairly well when the data are not severely corrupted, but it works in the reduced-dimension domain—and DR pre-processing can fail due to outliers.

Contributions: In this work, we explore both theoretical and practical aspects of VolMin. On the theory side, we show that two existing sufficient conditions for VolMin identifiability are in fact equivalent. The two identifiability results were developed in parallel, rely on different mathematical tools, and offer seemingly different characterizations of the sufficient conditions—so their equivalence is not obvious. Our proof ‘cross-validates’ the existing results, and thus leads to a deeper understanding of the VolMin problem.

On the algorithm side, we propose a new algorithmic framework for dealing with the VolMin criterion. The proposed framework takes outliers into consideration, without requiring DR pre-processing. Specifically, we impose an outlier-robust loss function onto the data fitting part, and propose a modified log-determinant loss function as the volume regularizer. By majorizing both functions, the fitting and the volume-regularization terms can be taken care of in a refreshingly easy way, and a simple *inexact* alternating optimization algorithm is derived. A Nesterov-type first-order optimization technique is further employed within this framework to accelerate convergence. The proposed algorithm is flexible—problem-specific prior information on the factors and different volume regularizers can be easily incorporated. Convergence of the proposed algorithm to a stationary point is also shown.

Besides a judiciously designed set of simulations, we also validate the proposed algorithm using real-life datasets. Specifically, we use remotely sensed hyperspectral image data and document data to showcase the effectiveness of the proposed algorithm in hyperspectral unmixing and document clustering applications, respectively. Notice that VolMin has never been used for document clustering before, to the best of our knowledge, and our work shows that VolMin is indeed very effective in this context, outperforming the state-of-art in terms of clustering accuracy.

A conference version of part of this work appears in [1]. Beyond [1], this journal version includes the equivalence of the identifiability conditions, first-order optimization-based updates, consideration of different types of regularization and constraints, proof of convergence, extensive simulations, and experiments using real data.

Notation: We largely follow common notational conventions in signal processing. $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{X} \in \mathbb{R}^{m \times n}$ denote a real-valued n -dimensional vector and a real-valued $m \times n$ matrix, respectively (resp.). $\mathbf{x} \geq \mathbf{0}$ (resp. $\mathbf{X} \geq \mathbf{0}$) means that \mathbf{x} (resp. \mathbf{X}) is element-wise non-negative. $\mathbf{x} \in \mathbb{R}_+^n$ (resp. $\mathbf{X} \in \mathbb{R}_+^{m \times n}$) also means that \mathbf{x} (resp. \mathbf{X}) is element-wise non-negative. $\mathbf{X} \succ \mathbf{0}$ and $\mathbf{X} \succeq \mathbf{0}$ mean that \mathbf{X} is positive definite and positive semidefinite, resp. The superscripts “ T ” and “ -1 ” stand for the transpose and inverse operations, resp. The ℓ_p norm of a vector $\mathbf{x} \in \mathbb{R}^n$, $p \geq 1$, is denoted by $\|\mathbf{x}\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$. The ℓ_p quasi-norm, $0 < p < 1$, is denoted by the same notation. The Frobenious norm and the matrix 2-norm are denoted by $\|\mathbf{X}\|_F$ and $\|\mathbf{X}\|_2$, respectively. The all-one vector is denoted by $\mathbf{1}$.

In this paper, we also make extensive use of convex analysis. Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_m]$. The convex cone of $\mathbf{x}_1, \dots, \mathbf{x}_m$ is denoted by $\text{cone}\{\mathbf{x}_1, \dots, \mathbf{x}_m\} = \text{cone}(\mathbf{X}) = \{\mathbf{y} | \mathbf{y} = \mathbf{X}\boldsymbol{\theta}, \boldsymbol{\theta} \geq \mathbf{0}\}$; the convex hull of $\mathbf{x}_1, \dots, \mathbf{x}_m$ is denoted by $\text{conv}\{\mathbf{x}_1, \dots, \mathbf{x}_m\} = \text{conv}(\mathbf{X}) = \{\mathbf{y} | \mathbf{y} = \mathbf{X}\boldsymbol{\theta}, \boldsymbol{\theta} \geq \mathbf{0}, \mathbf{1}^T \boldsymbol{\theta} = 1\}$; when $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ are linearly independent, $\text{conv}(\mathbf{X})$ is also called a simplex; the set of extreme rays of $\text{cone}(\mathbf{X})$ is denoted by $\text{ex}\{\text{cone}(\mathbf{X})\}$; and the dual cone of a convex \mathcal{X} is denoted by $\mathcal{X}^* = \{\mathbf{y} | \mathbf{y}^T \mathbf{x} \geq 0, \mathbf{x} \in \mathcal{X}\}$; $\text{bd}\mathcal{X}$ denotes the set of the boundary points of the *second order cone* \mathcal{X} . We point the readers to [5], [9], [19] for detailed illustration of the above concepts.

II. THE VOLMIN CRITERION AND IDENTIFIABILITY

In this section, we first give a brief introduction to the VolMin criterion for SMF and a concise review of the existing identifiability results. Then, we prove that the two independently developed identifiability results (using rather different mathematical tools) are equivalent.

A. Background

Consider the following signal model:

$$\mathbf{x}[\ell] = \mathbf{A}\mathbf{s}[\ell] + \mathbf{v}[\ell], \quad \ell = 1, \dots, L, \quad (1)$$

where $\mathbf{x}[\ell] \in \mathbb{R}^M$ is a measured data vector that is indexed by ℓ , $\mathbf{A} \in \mathbb{R}^{M \times K}$ is a basis which is assumed to have full column-rank, $\mathbf{s}[\ell] \in \mathbb{R}^K$ is the coefficient vector representing $\mathbf{x}[\ell]$ in the low dimensional subspace $\text{range}(\mathbf{A})$, and $\mathbf{v}[\ell] \in \mathbb{R}^M$ denotes noise. We assume that every $\mathbf{s}[\ell]$ satisfies

$$\mathbf{s}[\ell] \geq \mathbf{0} \text{ and } \mathbf{1}^T \mathbf{s}[\ell] = 1. \quad (2)$$

The model can be compactly written as $\mathbf{X} = \mathbf{A}\mathbf{S} + \mathbf{V}$, where $\mathbf{X} = [\mathbf{x}[1], \dots, \mathbf{x}[L]]$, $\mathbf{S} = [\mathbf{s}[1], \dots, \mathbf{s}[L]]$ and $\mathbf{V} = [\mathbf{v}[1], \dots, \mathbf{v}[L]]$.

The task of SMF is to factor \mathbf{X} into \mathbf{A} and \mathbf{S} . The simple model in (1) and (2) parsimoniously captures the essence of a large variety of applications. For *document clustering* or *topic mining* [4], estimating \mathbf{A} and \mathbf{S} can help recognize the most popular topics/opinions in textual data (e.g., documents, web




Fig. 1. Motivating examples: Hyperspectral unmixing and document clustering.

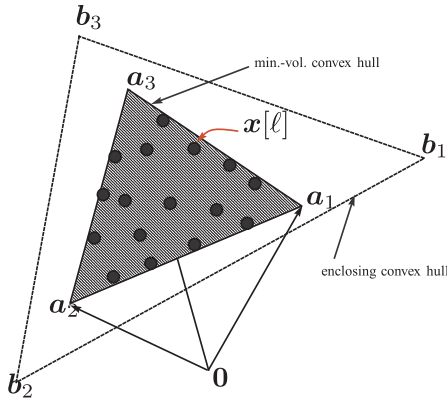


Fig. 2. The intuition of VolMin.

content, or social network posts), and cluster the data according to their weights on different topics/opinions. In *hyperspectral remote sensing* [3], [20], $x[\ell]$ represents a remotely sensed pixel using sensors of high spectral resolution, a_1, \dots, a_K denote K different spectral signatures of materials that comprise the pixel $x[\ell]$, and $s_k[\ell]$ denotes the proportion of material k contained in pixel $x[\ell]$. Estimating A enables recognition of the underlying materials in a hyperspectral image. See Fig. 1 for an illustration of these motivating examples. Very recently, the same model has been applied to *power spectra separation* [6] for dynamic spectrum access and *fast blind speech separation* [5]. In addition, many applications of NMF can also be considered under the model in (1) and (2), after suitable normalization [11].

Many algorithms have been developed for finding such a factorization, and we refer the readers to [3], [4] for a survey. Among these algorithms, we are particularly interested in the so-called VolMin criterion, which is identifiable under certain reasonable conditions. VolMin is motivated by the nice geometrical interpretation of the constraints in (2): Under these constraints, all the data points live in a *convex hull* spanned by a_1, \dots, a_K (or, a *simplex* spanned by a_1, \dots, a_K when the a_k 's are linearly independent); see Fig. 2. If the data points are sufficiently spread in $\text{conv}\{a_1, \dots, a_K\}$, then the minimum-volume enclosing convex hull coincides with $\text{conv}\{a_1, \dots, a_K\}$. Formally, the VolMin criterion can be formulated as

$$(\mathbf{A}, \{s[\ell]\}) = \arg \min_{\mathbf{B}, \{c[\ell]\}} \text{vol}(\mathbf{B}) \quad (3a)$$

$$\text{s.t. } \mathbf{x}[\ell] = \mathbf{B}\mathbf{c}[\ell], \quad (3b)$$

$$\mathbf{1}^T \mathbf{c}[\ell] = 1, \mathbf{c}[\ell] \geq \mathbf{0}, \forall \ell, \quad (3c)$$

where $\text{vol}(\mathbf{B})$ denotes a measure that is related or proportional to the volume of the simplex $\text{conv}\{b_1, \dots, b_K\}$, and (3b)–(3c) mean that every $x[\ell]$ is enclosed in $\text{conv}\{b_1, \dots, b_K\}$ (i.e., $x[\ell] \in \text{conv}\{b_1, \dots, b_K\}$). In the literature, various functions for $\text{vol}(\mathbf{B})$ have been used [5], [12]–[16], [18]. One representative choice of $\text{vol}(\mathbf{B})$ is

$$\text{vol}(\mathbf{B}) = \det(\bar{\mathbf{B}}^T \bar{\mathbf{B}}), \bar{\mathbf{B}} = [b_1 - b_K, \dots, b_{K-1} - b_K], \quad (4)$$

or its variants; see [13]–[15]. The reason of employing such a function is that $\sqrt{\det(\bar{\mathbf{B}}^T \bar{\mathbf{B}})}/((N-1)!)$ is the volume of the simplex $\text{conv}\{b_1, \dots, b_K\}$ by definition [21]. Another popular choice of $\text{vol}(\mathbf{B})$ is

$$\text{vol}(\mathbf{B}) = \det(\mathbf{B}^T \mathbf{B}); \quad (5)$$

see [5], [12], [16], [18]. Note that $\sqrt{\det(\mathbf{B}^T \mathbf{B})}/(N!)$ is the volume of the simplex $\text{conv}\{0, b_1, \dots, b_K\}$, which should scale similarly with the volume of $\text{conv}\{b_1, \dots, b_K\}$. The upshot of (5) is that (5) has a simpler structure than (4).

B. Identifiability of VolMin

The most appealing aspect of VolMin is its identifiability of A and S : Under mild and realistic conditions, the optimal solution to Problem (3) is essentially the true (A, S) . To be precise, let us make the following definition.

Definition 1: (VolMin Identifiability) Consider the matrix factorization model in (1), (2), and let (B^*, C^*) be any optimal solution to Problem (3). If every optimal (B^*, C^*) satisfies $B^* = \mathbf{A}\Pi$ and $C^* = \Pi^T S$, where Π denotes a permutation matrix, then we say that VolMin identifies the true matrix factors, or *VolMin identifiability holds*.

Fu *et al.* [5] have shown that.

Theorem 1: Let $\text{vol}(\mathbf{B})$ be the function in (5). Define a second order cone $\mathcal{C} = \{x \in \mathbb{R}^N \mid \mathbf{1}^T x \geq \frac{1}{\sqrt{N-1}} \|x\|_2\}$. Then VolMin identifiability holds if $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{S}) = K$ and

- i) $\mathcal{C} \subseteq \text{cone}(\mathbf{S})$; and
- ii) $\text{cone}(\mathbf{S}) \not\subseteq \text{cone}(\mathbf{Q})$ where \mathbf{Q} is any unitary matrix except the permutation matrices.

In plain words, a sufficient condition under which VolMin identifiability holds is when $\{s[\ell]\}_{\ell=1}^L$ are sufficiently scattered over the unit simplex, such that the second order cone \mathcal{C} is a subset of $\text{cone}(\mathbf{S})$. By comparing Theorem 1 to the identifiability conditions for NMF (see [7]–[9]), we see a remarkable advantage—VolMin does not have any restriction on A except being full-column rank. In fact, A can be dense, partially negative, or even complex-valued. This result allows us to apply VolMin to a wider variety of applications than NMF.

In [10], another sufficient condition for VolMin identifiability was proposed.

Theorem 2: Let $\text{vol}(\mathbf{B})$ be the function in (4). Assume $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{S}) = K$. Define $\mathcal{R}(r) = \{s \in \mathbb{R}^N \mid \|s\|_2 \leq r\} \cap \text{conv}\{e_1, \dots, e_N\}$ and $\gamma = \sup\{r \mid \mathcal{R}(r) \subseteq \text{conv}(\mathbf{S})\}$. Then VolMin identifiability holds if $\gamma > \frac{1}{\sqrt{N-1}}$.

Theorem 2 does not characterize its identifiability condition using convex cones like Theorem 1 did. Instead, it defines a




Fig. 3. Visualization of the sufficient conditions on the hyperplane $\mathbf{1}^T \mathbf{x} = 1$. The sufficient conditions in [5] and [10] both require that \mathcal{C} (the inner circle) is contained in $\text{cone}(\mathbf{S})$.

‘diameter’ r of the convex hull spanned by the columns of \mathbf{S} , and then develops an identifiability condition based on it.

The sufficient conditions presented in the two theorems seemingly have different flavors, but we notice that they are related in essence. To see the connections, we first note that Theorem 1 still holds after replacing $\text{cone}(\mathbf{S})$ and \mathcal{C} with convex hulls $\text{conv}\{\mathbf{s}[1], \dots, \mathbf{s}[L]\}$ and $\mathcal{C} \cap \text{conv}\{e_1, \dots, e_N\}$, respectively—since the $\mathbf{s}[\ell]$ ’s are all in $\text{conv}\{e_1, \dots, e_N\}$. In fact, $\mathcal{C} \cap \text{conv}\{e_1, \dots, e_N\}$ is exactly the set $\mathcal{R}(r)$ for $r \leq 1/\sqrt{N-1}$. Geometrically, we illustrate the conditions in Fig. 3 using $N = 3$ for visualization. We see that, if we look at the conditions in Theorem 1 at the 2-dimensional hyperplane that contains $\mathbf{1}^T \mathbf{x} = 1$, the two conditions both mean that the inner shaded region is contained in $\text{conv}\{\mathbf{s}[1], \dots, \mathbf{s}[L]\}$. Motivated by this observation, in this paper, we rigorously show that

Theorem 3: The sufficient conditions for VolMin identifiability in Theorem 1 and Theorem 2 are equivalent.

The proof of Theorem 3 can be found in Appendix A. Although the geometrical connection may seem clear on hindsight, rigorous proof is highly nontrivial. We first show that the condition in Theorem 1 is equivalent to another condition, and then establish equivalence between the ‘intermediate’ condition and the condition in Theorem 2.

Remark 1: The equivalence between the sufficient conditions in Theorem 1 and Theorem 2 is interesting and surprising—although the corresponding theoretical developments started from very different points of view, they converged to equivalent conditions. Their equivalence brings us deeper understanding of the VolMin criterion. The proof itself clarifies the role of regularity condition ii) in Theorem 1, which was originally difficult to describe geometrically—and now we understand that condition ii) is there to ensure $\gamma > \frac{1}{\sqrt{N-1}}$, i.e., the existence of a convex cone that is ‘sandwiched’ by $\text{cone}(\mathbf{S})$ and \mathcal{C} . In addition, the equivalence also suggests that the different cost functions in (4) and (5) ensure identifiability of \mathbf{A} and \mathbf{S} under the same sufficient conditions, and thus they are expected to perform similarly in practice. On the other hand, since the function in (5) is easier to handle, using it in practice is more appealing. As a by-product, since we have proved that condition ii) is equivalent to a condition that was used for NMF identifiability in [9] (cf.

Lemma 3), our result here also helps better understand the sufficient condition for NMF identifiability in [9] in a more intuitively pleasing way.

III. ROBUST VOLMIN VIA INEXACT BCD

In this section, we turn our attention to designing algorithms for dealing with the VolMin criterion. Optimizing the VolMin criterion is challenging. In early works such as [12], [18], linear DR with \mathbf{X} is assumed such that the basis after DR is a square matrix. This subsequently enables one to write the DR-domain VolMin problem as

$$\begin{aligned} \min_{\tilde{\mathbf{B}} \in \mathbb{R}^{K \times K}, \mathbf{C} \in \mathbb{R}^{K \times L}} \quad & \log |\det(\tilde{\mathbf{B}})| \\ \text{s.t.} \quad & \tilde{\mathbf{x}}[\ell] = \tilde{\mathbf{B}}\mathbf{c}[\ell] \\ & \mathbf{1}^T \mathbf{c}[\ell] = 1, \mathbf{c}[\ell] \geq \mathbf{0}, \end{aligned} \quad (6)$$

where $\tilde{\mathbf{x}}[\ell] \in \mathbb{R}^K$ is the dimension-reduced data vector corresponding to $\mathbf{x}[\ell]$, and $\tilde{\mathbf{B}} \in \mathbb{R}^{K \times K}$ is a dimension-reduced basis. Note that minimizing $\log |\det(\tilde{\mathbf{B}})|$ is the same as minimizing $\det(\tilde{\mathbf{B}}^T \tilde{\mathbf{B}})$. Problem (6) can be efficiently tackled via either alternating optimization [13] or successive convex optimization [12], [18]. The drawback with these existing algorithms is that noise was not taken into consideration. Also, these approaches require DR to make the effective \mathbf{A} square—but DR may not be reliable in the presence of outliers or modeling errors. Another major class of algorithms such as those in [15], [16] considers

$$\begin{aligned} \min_{\mathbf{B} \in \mathbb{R}^{M \times K}, \mathbf{C} \in \mathbb{R}^{K \times L}} \quad & \|\mathbf{X} - \mathbf{BC}\|_F^2 + \lambda \cdot \text{vol}(\mathbf{B}) \\ \text{s.t.} \quad & \mathbf{C} \geq \mathbf{0}, \mathbf{1}^T \mathbf{C} = \mathbf{1}^T, \end{aligned} \quad (7)$$

where $\lambda > 0$ is a parameter that balances data fidelity versus VolMin. The formulation in (7) avoids DR and takes noise into consideration. However, our experience is that volume regularizers, such as $\text{vol}(\mathbf{B}) = \det(\mathbf{B}^T \mathbf{B})$, are numerically harder to cope with, which will be explained in detail later. We should also compare Problem (7) with the VolMin formulation in (3). Problem (3) enforces a hard constraint $\mathbf{X} = \mathbf{BC}$, and thus ensures that every feasible \mathbf{B} and \mathbf{C} have full rank in the noiseless case. On the other hand, Problem (7) employs a fitting-based criterion, and an overly large λ could result in rank-deficient factors even in the noiseless case. Hence, λ should be chosen with caution.

Another notable difficulty is that outliers are very damaging to the VolMin criterion. In many cases, a single outlier can make the minimum-volume enclosing convex hull very different from the desired one; see Fig. 4 for an illustration. The state-of-the-art algorithm that considers outliers for the VolMin-based factorization is *simplex identification via split augmented Lagrangian* (SISAL) [18], but it takes care of outliers in the dimension-reduced domain. As already mentioned, the DR process itself may be impaired by outliers, and thus dealing with outliers in the original data domain is more appealing. Directly factoring \mathbf{X} in the original data domain also has the advantage of allowing us to incorporate any *a priori* information on \mathbf{A} and \mathbf{S} , such as nonnegativity, smoothness, and sparsity.




Fig. 4. The impact of outliers to VolMin. The dots are $\mathbf{x}[\ell]$'s; the shaded area is $\text{conv}\{\mathbf{a}_1, \dots, \mathbf{a}_N\}$, the triangles with dashed lines are data-enclosing convex hulls, and the one with solid lines is the minimum-volume enclosing convex hull. Left: the case where no outliers exist. Right: the case where a single outlier exists.

A. Proposed Robust VolMin Algorithm

We are interested in the VolMin-regularized matrix factorization, but we take the outlier problem into consideration. Specifically, we propose to employ the following optimization surrogate of the VolMin criterion:

$$\min_{\mathbf{B}, \mathbf{C}} \sum_{\ell=1}^L \frac{1}{2} \left(\|\mathbf{x}[\ell] - \mathbf{B}\mathbf{c}[\ell]\|_2^2 + \epsilon \right)^{\frac{p}{2}} + \frac{\lambda}{2} \log \det(\mathbf{B}^T \mathbf{B} + \tau \mathbf{I})$$

$$\text{s.t. } \mathbf{1}^T \mathbf{c}[\ell] = 1, \mathbf{c}[\ell] \geq \mathbf{0}, \forall \ell, \quad (8)$$

where $p \in (0, 2]$, $\lambda > 0$, $\epsilon > 0$, and $\tau > 0$. Here, $\epsilon > 0$ is a small regularization parameter, which keeps the first term inside its smooth region for computational convenience when $p < 1$; if $p \in (1, 2]$, we can simply let $\epsilon = 0$. The parameter $\tau > 0$ is also a small positive number, which is used to ensure that the cost function is bounded from below for any \mathbf{B} .

The motivation of using $\log \det(\mathbf{B}^T \mathbf{B} + \tau \mathbf{I})$ instead of the commonly used volume regularizers such as $\det(\mathbf{B}^T \mathbf{B})$ is computational simplicity: Although both functions are non-convex and conceptually equally hard to deal with, the former features a much simpler update rule because it admits a tight upper bound while the latter does not—this point will become clearer shortly. Interestingly, $\log \det(\mathbf{B}^T \mathbf{B} + \tau \mathbf{I})$ has been used in the context of low-rank matrix recovery [22], [23], but here we instead apply it for simplex-VolMin. The ℓ_2/ℓ_p - (quasi-) norm data fitting part is employed to downweight the impact of the outliers—when $0 < p < 2$, such a fitting criterion is less sensitive to large fitting errors and thus is robust against outliers. Other robust fitting criteria can also be considered—e.g., the ℓ_p norm-based criterion $\|\mathbf{X} - \mathbf{B}\mathbf{C}\|_p^p$ for $0 < p < 2$ where $\|\mathbf{Y}\|_p^p = \sum_{i=1}^m \sum_{j=1}^m |Y_{i,j}|^p$ is known to be robust to entry-level outliers [24]–[26]. Nevertheless, the type of outliers that matters in VolMin is column outliers (or gross outliers) which represents a point lying outside the ground-truth convex hull, and the proposed criterion is natural for fending against such outliers. In addition, computationally, the ℓ_2/ℓ_p mixed-norm criterion can be handled efficiently, as we will see.

Our primary objective is to handle Problem (8) efficiently. Nonetheless, we will also show that the proposed algorithmic framework can easily incorporate different volume-associated regularizers in the literature, such as the previously mentioned $\text{vol}(\mathbf{B}) = \det(\mathbf{B}^T \mathbf{B})$, and

$$\text{vol}(\mathbf{B}) = \sum_{i=1}^{K-1} \sum_{j=i+1}^K \|\mathbf{b}_i - \mathbf{b}_j\|_2^2; \quad (9)$$

see [27]. Notice that (9) is a coarse approximation of the volume of $\text{conv}\{\mathbf{b}_1, \dots, \mathbf{b}_K\}$, which measures the volume by simply adding up the squared distances between the vertices.

B. Update of \mathbf{C}

Our idea is to update \mathbf{B} and \mathbf{C} alternately, i.e., using block coordinate descent (BCD). Unlike classic BCD [28], we solve the partial optimization problems in an *inexact* fashion for efficiency. We first consider updating \mathbf{C} . The problem w.r.t. \mathbf{C} is separable w.r.t. ℓ and convex. Therefore, after t iterations with the current solution $(\mathbf{B}^t, \mathbf{C}^t)$, we consider:

$$\mathbf{c}^{t+1}[\ell] := \arg \min_{\mathbf{c}[\ell]} \frac{1}{2} \|\mathbf{x}[\ell] - \mathbf{B}^t \mathbf{c}[\ell]\|_2^2$$

$$\text{s.t. } \mathbf{1}^T \mathbf{c}[\ell] = 1, \mathbf{c}[\ell] \geq \mathbf{0}, \quad (10)$$

for $\ell = 1, \dots, L$. Since Problem (10) is convex, one can update \mathbf{C} by solving Problem (10) to optimality. An *alternating direction method of multipliers* (ADMM)-based algorithm was provided in the conference version of this work for this purpose; see the detailed implementation in [1]. Nevertheless, exactly solving Problem (10) at each iteration is computationally costly, especially when the problem size is large. Here, we propose to deal with Problem (10) using local approximation. Specifically, let

$$f(\mathbf{c}[\ell]; \mathbf{B}^t) = \frac{1}{2} \|\mathbf{x}[\ell] - \mathbf{B}^t \mathbf{c}[\ell]\|_2^2.$$

Then, $f(\mathbf{c}[\ell]; \mathbf{B}^t)$ can be locally approximated at $\mathbf{c}^t[\ell]$ by the following:

$$u(\mathbf{c}[\ell]; \mathbf{B}^t) = f(\mathbf{c}^t[\ell]; \mathbf{B}^t) + (\nabla f(\mathbf{c}^t[\ell]; \mathbf{B}^t))^T (\mathbf{c}[\ell] - \mathbf{c}^t[\ell])$$

$$+ \frac{L^t}{2} \|\mathbf{c}[\ell] - \mathbf{c}^t[\ell]\|_2^2,$$

where $L^t \geq 0$. On the right hand side (RHS) of the above, the first two terms constitute a first-order approximation of $f(\mathbf{c}[\ell]; \mathbf{B}^t)$ at $\mathbf{c}^t[\ell]$, and the second term restrains $\mathbf{c}^{t+1}[\ell]$ to be close to $\mathbf{c}^t[\ell]$ in terms of Euclidean distance. It is well-known that when $L^t \geq \|(\mathbf{B}^t)^T \mathbf{B}^t\|_2$,

$$u(\mathbf{c}[\ell]; \mathbf{B}^t) \geq f(\mathbf{c}[\ell]; \mathbf{B}^t), \forall \mathbf{c}[\ell] \in \mathbb{R}^K$$

holds for all $\mathbf{c}[\ell]$ and the equality holds if and only if $\mathbf{c}[\ell] = \mathbf{c}^t[\ell]$ [29]. In other words, when $L^t \geq \|(\mathbf{B}^t)^T \mathbf{B}^t\|_2$, $u(\mathbf{c}[\ell])$ is a ‘majorizing’ function of $f(\mathbf{c}[\ell]; \mathbf{B}^t)$. Given this majorizing function, we update $\mathbf{c}[\ell]$ by the following simple rule:

$$\mathbf{c}^{t+1}[\ell] = \arg \min_{\mathbf{1}^T \mathbf{c}[\ell]=1, \mathbf{c}[\ell] \geq \mathbf{0}} u(\mathbf{c}[\ell]; \mathbf{B}^t). \quad (11)$$

By re-arranging the terms and discarding constants, Problem (11) is equivalent to the following

$$\min_{\mathbf{1}^T \mathbf{c}[\ell]=1, \mathbf{c}[\ell] \geq \mathbf{0}} \left\| \mathbf{c}[\ell] - \left(\mathbf{c}^t[\ell] - \frac{1}{L^t} \nabla f(\mathbf{c}^t[\ell]; \mathbf{B}^t) \right) \right\|_2^2.$$

The RHS of the above can be considered as a gradient projection step with step size $1/L^t$. Letting $P_{L^t}(\mathbf{c}^t[\ell])$ denote the optimal solution of the above, we simplify the notation of updating $\mathbf{c}[\ell]$ as

$$\mathbf{c}^{t+1}[\ell] = P_{L^t}(\mathbf{c}^t[\ell]). \quad (12)$$

Problem (12) is a simple projection that can be solved with worst-case complexity of $\mathcal{O}(K \log K)$ flops; see [30] for a detailed implementation.

The described update of \mathbf{C} has light per-iteration complexity, but it could result in slow convergence of the overall alternating optimization algorithm; see Fig. 6 in the simulations. To improve the convergence speed in practice, and inspired by the success of Nesterov's optimal first-order algorithm and its related algorithms [31], [32], we propose the following update of \mathbf{C} :

$$\mathbf{c}^{t+1}[\ell] = P_{L^t}(\mathbf{y}^t[\ell]) \quad (13a)$$

$$q^{t+1} = \frac{1 + \sqrt{1 + 4(q^t)^2}}{2} \quad (13b)$$

$$\mathbf{y}^t[\ell] = \mathbf{c}^t[\ell] + \left(\frac{q^t - 1}{q^{t+1}} \right) (\mathbf{c}^t[\ell] - \mathbf{c}^{t-1}[\ell]), \quad (13c)$$

where $\{q^t\}_{t=1}^\infty$ is a sequence with $q^1 = 1$. Simply speaking, instead of locally approximating $f(\mathbf{c}[\ell]; \mathbf{B}^t)$ at $\mathbf{c}^t[\ell]$, we approximate it at an 'extrapolated point' $\mathbf{y}^t[\ell]$. Without the alternating optimization procedure, using extrapolation is probably much faster than using the plain gradient-based methods [31], [32]. Embedding extrapolation into alternating optimization was first considered in [33] in the context of tensor factorization, where acceleration of convergence was observed. In our case, the extrapolation procedure also substantially reduces the number of iterations for achieving convergence, as will be shown in the simulations.

C. Update of \mathbf{B}

The update of \mathbf{B} relies on the following two lemmas.

Lemma 1: [34] Assume $0 < p \leq 2$, $\epsilon > 0$, and let $\phi_p(w) := \frac{2-p}{2} \left(\frac{2}{p} w \right)^{\frac{p}{2-p}} + \epsilon w$. Then, we have $(x^2 + \epsilon)^{p/2} = \min_{w \geq 0} w x^2 + \phi_p(w)$. Also, the minimizer is unique and given by $w_{\text{opt}} = \frac{p}{2} (x^2 + \epsilon)^{\frac{p-2}{2}}$.

Lemma 2: [35] Let $\mathbf{E} \in \mathbb{R}^{K \times K}$ be any matrix such that $\mathbf{E} \succ \mathbf{0}$. Consider the function $f(\mathbf{F}) = \text{Tr}(\mathbf{F}\mathbf{E}) - \log \det \mathbf{F} - K$. Then, $\log \det \mathbf{E} = \min_{\mathbf{F} \succ \mathbf{0}} f(\mathbf{F})$, and the minimizer is uniquely given by $\mathbf{F}_{\text{opt}} = \mathbf{E}^{-1}$.

The lemmas provide two functions that majorize the data fitting part and the volume-regularization part in (8), respectively. Specifically, at iteration t and after updating \mathbf{C} , we have $(\hat{\mathbf{B}}^t, \{\mathbf{c}^{t+1}[\ell]\}_{\ell=1}^L)$. Then, the following holds:

$$\log \det(\mathbf{B}^T \mathbf{B} + \epsilon \mathbf{I}) \leq \text{Tr}(\mathbf{F}^t \mathbf{B}^T \mathbf{B}) - \log \det \mathbf{F}^t - K, \quad (14)$$

where $\mathbf{F}^t = ((\mathbf{B}^t)^T \mathbf{B}^t + \epsilon \mathbf{I})^{-1}$ and the equality holds when $\mathbf{B} = \mathbf{B}^t$. Similarly, we have

$$\begin{aligned} & \sum_{\ell=1}^L \frac{1}{2} \left(\|\mathbf{x}[\ell] - \mathbf{B} \mathbf{c}^{t+1}[\ell]\|_2^2 + \epsilon \right)^{\frac{p}{2}} \\ & \leq \sum_{\ell=1}^L \frac{w_\ell^t}{2} \|\mathbf{x}[\ell] - \mathbf{B} \mathbf{c}^{t+1}[\ell]\|_2^2 + \sum_{\ell=1}^L \phi_p(w_\ell^t), \end{aligned} \quad (15)$$

where $w_\ell^t = \frac{p}{2} (\|\mathbf{x} - \mathbf{B}^t \mathbf{c}^{t+1}[\ell]\|_2^2 + \epsilon)^{\frac{p-2}{2}}$ and the equality holds when $\mathbf{B} = \mathbf{B}^t$. Putting (14), (15) together and dropping

the irrelevant terms, we find \mathbf{B}^{t+1} by solving the following:

$$\begin{aligned} \mathbf{B}^{t+1} := & \arg \min_{\mathbf{B}} \sum_{\ell=1}^L \frac{w_\ell}{2} \|\mathbf{x}[\ell] - \mathbf{B} \mathbf{c}^{t+1}[\ell]\|_2^2 \\ & + \frac{\lambda}{2} \text{Tr}(\mathbf{F}^t (\mathbf{B}^T \mathbf{B})). \end{aligned} \quad (16)$$

Problem (16) is a convex quadratic program that admits the following closed-form solution:

$$\mathbf{B}^{t+1} := \mathbf{X} \mathbf{W}^t (\mathbf{C}^{t+1})^T (\mathbf{C}^{t+1} \mathbf{W} (\mathbf{C}^{t+1})^T + \lambda \mathbf{F}^t)^{-1}, \quad (17)$$

where $\mathbf{W}^t = \text{Diag}(w_1^t, \dots, w_L^t)$.

Remark 2: The expression in (16) reveals why the proposed criterion and algorithm can automatically downweight the effect brought by the outliers. Suppose that $(\mathbf{B}^t, \mathbf{C}^{t+1})$ is a "good enough" solution which is close to the ground truth. Then, w_ℓ^t is small when $\mathbf{x}[\ell]$ is an outlier since the fitting error term $\|\mathbf{x}[\ell] - \mathbf{B}^t \mathbf{c}^{t+1}[\ell]\|_2^2$ is large. Hence, for the next iteration, \mathbf{B}^{t+1} is estimated with the importance of the outlier $\mathbf{x}[\ell]$ downweighted.

Remark 3: In practice, adding constraints on \mathbf{B} by letting $\mathbf{B} \in \mathcal{B}$ is sometimes instrumental, since a lot of applications do have prior information that can be used to enhance performance. For example, in image processing, a nonnegative \mathbf{B} is often sought, and thus one can set $\mathcal{B} = \mathbb{R}_+^{M \times N}$. When \mathcal{B} is convex, the problem in (16) can usually be solved in an efficient manner; e.g., one can call general-purpose solvers such as interior-point methods. However, using general-purpose solvers here may lose efficiency since solving constrained least squares to a certain accuracy *per se* may require a lot of iterations. To simplify the update, we update \mathbf{B}^t following the same spirit of updating \mathbf{C} : Let $g(\mathbf{B}; \mathbf{C}^{t+1}) = \sum_{\ell=1}^L \frac{w_\ell}{2} \|\mathbf{x}[\ell] - \mathbf{B} \mathbf{c}^{t+1}[\ell]\|_2^2 + \frac{\lambda}{2} \text{Tr}(\mathbf{F}^t (\mathbf{B}^T \mathbf{B})) + \text{const}$, where $\text{const} = \sum_{\ell=1}^L \phi_p(w_\ell^t) - K$. We solve a local approximation of $g(\mathbf{B}; \mathbf{C}^{t+1})$:

$$\begin{aligned} \mathbf{B}^{t+1} := & \arg \min_{\mathbf{B} \in \mathcal{B}} g(\mathbf{B}^t; \mathbf{C}^{t+1}) + \nabla g(\mathbf{B}^t; \mathbf{C}^{t+1})^T (\mathbf{B} - \mathbf{B}^t) \\ & + \frac{\mu^t}{2} \|\mathbf{B} - \mathbf{B}^t\|_F^2 \\ := & \text{Proj}_{\mathcal{B}}(\mathbf{B}^t - \mu^t \nabla g(\mathbf{B}^t; \mathbf{C}^{t+1})), \end{aligned} \quad (18)$$

where $\mu^t \geq 0$ and

$$\begin{aligned} \nabla g(\mathbf{B}^t; \mathbf{C}^{t+1}) = & \mathbf{B}^t (\mathbf{C}^{t+1} \mathbf{W}^t (\mathbf{C}^{t+1})^T + \lambda \mathbf{F}^t) \\ & - \mathbf{X} \mathbf{W}^t (\mathbf{C}^{t+1})^T, \end{aligned}$$

is the partial derivative of the cost function in (16) w.r.t. \mathbf{B} at \mathbf{B}^t , and $\text{Proj}_{\mathcal{B}}(\mathbf{Z})$ denotes the Euclidean projection of \mathbf{Z} on \mathcal{B} . For some \mathcal{B} 's, the projection is easy to compute; e.g., when $\mathcal{B} = \mathbb{R}_+^M$, we have $\text{Proj}_{\mathcal{B}}(\mathbf{Z}) = \max\{\mathbf{Z}, \mathbf{0}\}$; see other easily implementable projections in [36]. Notice that the update in (18) can also easily incorporate extrapolation.

The robust volume minimization (RVolMin) algorithm is summarized in Algorithm 1. Its convergence properties are stated in Proposition 1, whose proof is relegated to Appendix B.

Proposition 1: Assume that L^t and μ^t are chosen such that $L^t \geq \|(\mathbf{B}^t)^T \mathbf{B}^t\|_2$ and $\mu^t \geq \|(\mathbf{F}^t)^T \mathbf{F}^t\|_2$, respectively. Also,

Algorithm 1: RVolMin.

input: \mathbf{X} ; $p \in (0, 2)$; K ; initial (\mathbf{B}, \mathbf{C}) ; ϵ ; τ .

- 1 $t = 0$;
- 2 $\mathbf{W}^t = \mathbf{I}$; $\mathbf{F}^t = \mathbf{I}$; $\mathbf{B}^t = \mathbf{B}$; $\mathbf{C}^t = \mathbf{C}$;
- 3 **repeat**
- 4 select L^t and μ^t ;
- 5 $\mathbf{C}^{t+1} \leftarrow (12), (13), \text{ or ADMM in [1]}$.
- 6 $\mathbf{B}^{t+1} \leftarrow (17) \text{ or } (18)$;
- 7 $t \leftarrow t + 1$;
- 8 $w_\ell^t \leftarrow \frac{p}{2} (\|\mathbf{x}[\ell] - \mathbf{B}^t \mathbf{c}^t[\ell]\|_2^2 + \epsilon)^{\frac{p-2}{2}}$ for $\ell = 1, \dots, L$;
- 9 $\mathbf{F}^t \leftarrow ((\mathbf{B}^t)^T \mathbf{B}^t + \tau \mathbf{I})^{-1}$.
- 10 **until** *Some stopping criterion is reached*;

output: \mathbf{C}^t ; \mathbf{B}^t .

assume that \mathcal{B} is a convex closed set. Then, if the initial objective value is finite, the whole solution sequence generated by Algorithm 1 converges to the set \mathcal{S} that consists of all the stationary points of Problem (8), i.e.,

$$\lim_{t \rightarrow \infty} d^{(t)}((\mathbf{B}^t, \mathbf{C}^t), \mathcal{S}) = 0,$$

where $d^{(t)}((\mathbf{B}^t, \mathbf{C}^t), \mathcal{S}) = \min_{\mathbf{Y} \in \mathcal{S}} \|\mathbf{Y} - (\mathbf{B}^t, \mathbf{C}^t)\|_F^2$.

Remark 4: As mentioned before, we may also use different volume regularizers. Let us consider the volume regularizer in (9) first. It was shown in [27] that this regularizer can also be expressed as $\text{vol}(\mathbf{B}) = \text{Tr}(\mathbf{G}\mathbf{B}^T\mathbf{B})$, where $\mathbf{G} = \mathbf{K}\mathbf{I} - \mathbf{1}\mathbf{1}^T$. Therefore, by letting $\mathbf{F}^t = \mathbf{G}$ in Algorithm 1, the updates can be directly applied to handle the regularizer in (9). Dealing with (5) is more difficult. One possible way is to make use of (18) since $\det(\mathbf{B}^T\mathbf{B})$ is differentiable. The difficulty is that a global upper bound of the subproblem w.r.t. \mathbf{B} may not exist. Under such circumstances, *sufficient decrease* at each iteration needs to be guaranteed for establishing convergence to a stationary point [37]. In practice, the *Armijo rule* is usually invoked to achieve this goal, which in general is computationally more costly compared to the cases where μ^t can be determined in closed form.

Remark 5: Problem (8) is a nonconvex optimization problem. Hence, a good starting point of RVolMin can help reach meaningful solutions quickly. In practice, different initializations can be considered:

- Existing VolMin algorithms. Many VolMin algorithms, such as the ones working in the reduced-dimension domain (e.g., the algorithms in [13], [18]), exhibit good efficiency. The difficulty is that these algorithms are usually sensitive to the DR process in the presence of outliers. Nevertheless, one can employ robust DR algorithms together with the algorithms in [13], [18] as an initialization approach. Nuclear norm-based algorithms [38] are viable options for robust DR, but are not suitable for large-scale problems because of the computational complexity. Under such circumstances, one may adopt simple alternatives such as that proposed in [39].

- Nonnegative matrix factorization. If \mathbf{A} is known to be nonnegative, any NMF algorithm can be employed as initialization. In practice, dealing with NMF is arguably simpler relative to VolMin, and many efficient solvers for NMF exist—see [40] for a survey. Although NMF usually does not provide a satisfactory result on its own in cases where it cannot guarantee the identifiability of its factors, using the NMF-estimated factors to

initialize the algorithms that provide identifiability guarantees can sometimes enhance the performance of the latter.

IV. SIMULATIONS

In this section, we provide simulations to showcase the effectiveness of the proposed algorithm. We generate the elements of $\mathbf{A} \in \mathbb{R}^{M \times K}$ from the uniform distribution between zero and one. We generate $\mathbf{s}[\ell]$ on the unit simplex and with $\max_i s_i[\ell] \leq \gamma$, where $\frac{1}{K} \leq \gamma \leq 1$ is given. We choose $\gamma = 0.85$, which results in a so-called ‘no-pure-pixel case’ in the context of remote sensing and is known to be challenging to handle; see [3], [4] for details. Zero-mean white Gaussian noise is added to the generated data. To model outliers, we define the outlier at data point ℓ as $\mathbf{o}[\ell]$ and let $\mathcal{A} \subseteq \{1, \dots, L\}$ be the index set of outliers. We assume that $\mathbf{o}[\ell] = \mathbf{0}$ if $\ell \notin \mathcal{A}$ and $\mathbf{x}[\ell] = \mathbf{o}[\ell]$ otherwise. We denote $N_o = |\mathcal{A}|$ as the total number of outliers. Those active outliers are generated following the uniform distribution between zero and one, and are scaled to satisfy problem specifications. For the proposed algorithm, we fix $p = 0.5$, $\epsilon = 10^{-12}$, and $\tau = 10^{-8}$ unless otherwise specified. We stop the proposed algorithm when the absolute change of the cost function is smaller than 10^{-5} or the number of iterations reaches 1000.

We define the signal-to-noise ratio (SNR) as $\text{SNR} = 10 \log_{10}(\frac{\mathbb{E}\{\|\mathbf{A}\mathbf{s}[\ell]\|_2^2\}}{\mathbb{E}\{\|\mathbf{v}[\ell]\|_2^2\}})$. Also, to quantify the corruption caused by the outliers, we define the signal-to-outlier ratio (SOR) as $\text{SOR} = 10 \log_{10}(\frac{\mathbb{E}\{\|\mathbf{A}\mathbf{s}[\ell]\|_2^2\}}{\mathbb{E}\{\|\mathbf{o}[\ell]\|_2^2\}})$. We use the mean-squared-error (MSE) of \mathbf{A} as a measure of factorization performance, defined as

$$\text{MSE} = \min_{\pi \in \Pi} \frac{1}{K} \sum_{k=1}^K \left\| \frac{\mathbf{a}_k}{\|\mathbf{a}_k\|_2} - \frac{\hat{\mathbf{a}}_{\pi_k}}{\|\hat{\mathbf{a}}_{\pi_k}\|_2} \right\|_2^2,$$

where Π is the set of all permutations of $\{1, 2, \dots, K\}$; and $\hat{\mathbf{a}}_k$ is the estimate of \mathbf{a}_k .

In this section, we use the SISAL algorithm proposed in [18] as a baseline. SISAL is a state-of-art robust VolMin algorithm that takes outliers into account by solving

$$\min_{\tilde{\mathbf{B}}, \mathbf{1}^T \mathbf{C} = \mathbf{1}^T, \{\tilde{\mathbf{x}}[\ell] = \tilde{\mathbf{B}}\mathbf{c}[\ell]\}} \log \det(\tilde{\mathbf{B}}) + \eta \|\mathbf{C}\|_h,$$

where $\|\cdot\|_h = \sum_{\ell=1}^L \sum_{k=1}^K \max(-c_k[\ell], 0)$ is an element-wise hinge function. The intuition behind SISAL is to penalize the outliers whose $\mathbf{c}[\ell]$ has negative elements, but still allowing them to exist, thereby having some robustness to outliers. The tuning parameter $\eta > 0$ in SISAL controls the amount of outliers that are ‘‘allowed,’’ and we test multiple η ’s for SISAL in the simulations. We run the original SISAL that uses SVD-based dimension reduction and the modified SISAL which uses the robust dimension reduction (RDR) algorithm in [39]. The latter is also used to initialize the proposed algorithm.

We first use an illustrative example to show the effectiveness of the proposed algorithm in the presence of outliers. In this example, we set $\text{SNR} = 18$ dB, $\text{SOR} = -10$ dB, $N_o = 20$, $(M, K) = (50, 3)$, and $L = 1000$. The results are projected onto the affine set that contains $\text{conv}\{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\}$, i.e., a two-dimensional hyperplane. In Fig. 5, we see that SISAL with different η ’s cannot yield reasonable estimates of \mathbf{A} since the

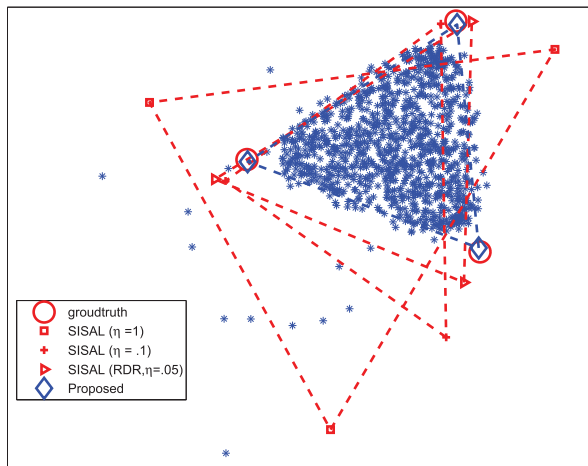


Fig. 5. The $\hat{\mathbf{A}}$'s estimated by various algorithms. Blue points are $\mathbf{x}[\ell]$'s.

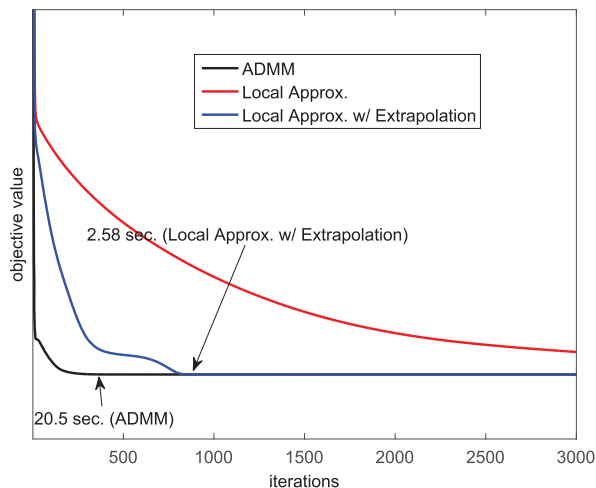


Fig. 6. Objective value vs. iterations, using different update strategies for \mathbf{C} .

DR stage threw the data to a badly estimated subspace. Using RDR, SISAL performs better, but is still not satisfactory. In this case, the proposed algorithm yields the most accurate estimate of \mathbf{A} .

In Fig. 6, we show the convergence curves of the algorithm under different update rules of \mathbf{C} , i.e., ADMM in [1], the proposed local approximation, and local approximation with extrapolation. We show the results averaged from 10 trials, where $\text{SNR} = 18$ dB and $\text{SOR} = -5$ dB. We see that using ADMM, the objective value converges within 400 iterations. Local approximation with extrapolation uses around 800 iterations to attain convergence of the objective value, but the objective value cannot converge within 3000 iterations without extrapolation. In terms of runtimes, the local approximation methods uses 0.003 second per iteration (a complete update of both \mathbf{C} and \mathbf{B}), while ADMM costs 0.05 second per iteration. Obviously, local approximation with extrapolation is the most favorable update scheme: its number of iterations for achieving convergence is around twice of that of ADMM, but it is 15 times faster relative to ADMM for completing an update of \mathbf{C} . Specifically, in the case under test, the average time for the algorithm using ADMM to update \mathbf{C} to achieve the pointed objective value in Fig. 6 is




Fig. 7. MSE of $\hat{\mathbf{A}}$ obtained by different algorithms vs. SNR. $(M, K) = (50, 5)$; $N_o = 20$; $\text{SOR} = -5$ dB.

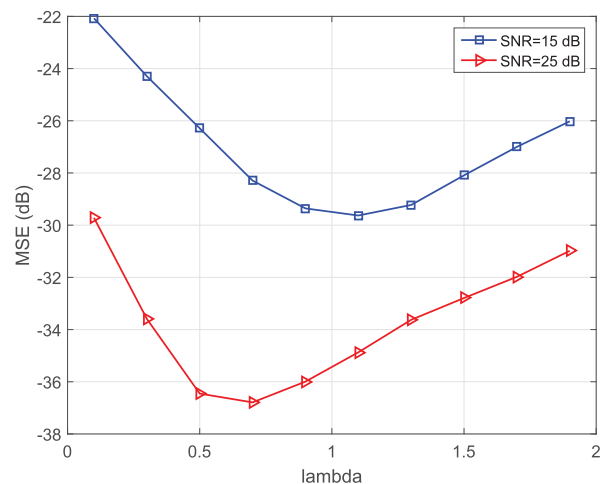


Fig. 8. MSE of proposed algorithm vs. λ . $(M, K) = (50, 5)$; $N_o = 20$; $L = 1000$; $\text{SOR} = -5$ dB.

20.5 seconds, while using local approximation with extrapolation costs 2.58 seconds to reach the same objective level. In the upcoming simulations, all the results of the proposed algorithms are obtained with the extrapolation strategy.

In Fig. 7, we show the MSE performance of the proposed algorithm versus SNR. We fix $\text{SOR} = -5$ dB, and let $N_o = 20$, $(M, K) = (50, 5)$, and $L = 1000$. In Fig. 7, we see that the original SISAL fails for different η 's. Using RDR, SISAL with $\eta = 0.1$ yields reasonable results for all the tested SNRs. The proposed algorithm with $\lambda = 1$ and $\lambda = 0.5$ gives the lowest MSEs. The MSEs given by RVolMin with $\lambda = 1$ are the lowest when $\text{SNR} \leq 20$ dB, and RVolMin with $\lambda = 0.5$ exhibits the best MSE performance when $\text{SNR} \geq 25$ dB. The results are consistent with the intuition behind selecting λ : when the SNR is low, a relatively large λ is needed to enhance the effect of the VolMin regularization.

To understand the effect of selecting λ , we plot the MSEs of the proposed algorithm versus λ in Fig. 8. We see that there exists an (SNR-dependent) optimal choice of λ for achieving the lowest MSE, but also note that any λ in the range considered yields satisfactory results in both cases.




Fig. 9. MSE of $\hat{\mathbf{A}}$ versus K . $M = 50$; $N_o = 20$; $L = 1000$; $\text{SOR} = -5$ dB.

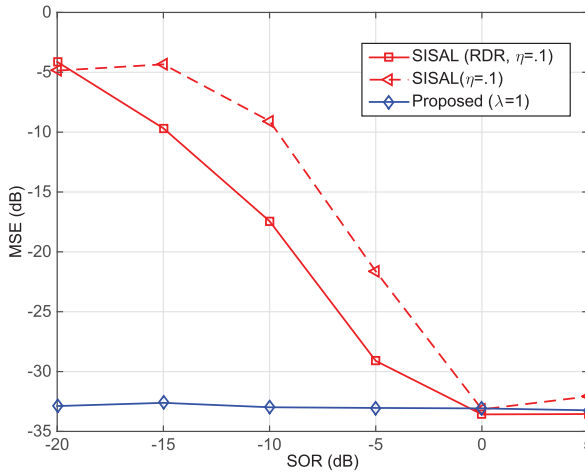


Fig. 10. MSE of $\hat{\mathbf{A}}$ versus SOR . $(M, K) = (50, 5)$; $N_o = 20$; $L = 1000$; $\text{SNR} = 20$ dB.

Fig. 9 shows the MSE performance of the algorithms versus K . We fix $\text{SNR} = 20$ dB and the other settings are the same as in the previous simulation. The results of SISAL and SISAL with RDR are also used as baselines. We run several η 's for SISAL and present the results of the one with the lowest MSEs. As expected, all the algorithms work better when the rank of the factorization model is lower—which is consistent with past experience on different matrix factorization algorithms, such as [40]. SISAL and SISAL with RDR work reasonably when $K = 3$, but deteriorate when $K \geq 6$. On the other hand, even when $K = 15$, the proposed algorithm still works well, giving the lowest MSE.

Fig. 10 shows the MSEs of the algorithms versus SORs. One can see that when some data are badly corrupted, i.e., when $\text{SOR} \leq -10$ dB, the proposed algorithm yields significantly lower MSEs than SISAL and SISAL with RDR. When $\text{SOR} \geq 0$ dB, all three algorithms provide comparable performance.

We also test the algorithms versus the number of outliers. In Fig. 11, one can see that the proposed algorithm is not very sensitive to the change of N_o : the MSE curve of the proposed algorithm is quite flat for different N_o 's in this simulation. SISAL

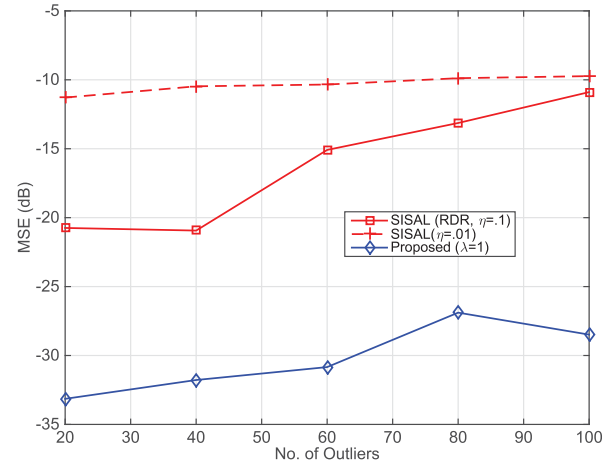


Fig. 11. MSE of $\hat{\mathbf{A}}$ versus N_o . $(M, K) = (50, 5)$; $L = 1000$; $\text{SOR} = -5$ dB; $\text{SNR} = 20$ dB.

TABLE I
THE MSEs OF THE ALGORITHMS UNDER ILL-CONDITIONED \mathbf{A} .
 $(M, K) = (50, 5)$; $L = 1000$; $N_o = 20$; $\text{SOR} = -5$ dB.

Algorithm	uniformly distributed \mathbf{A}		ill-conditioned \mathbf{A}	
	$\text{SNR} = -25$ dB	$\text{SNR} = -35$ dB	$\text{SNR} = -25$ dB	$\text{SNR} = -35$ dB
SISAL	-12.6327	-11.867	-11.8829	-11.8014
SISAL (RDR)	-22.3279	-24.7461	-13.2692	-13.3246
Proposed ($\lambda = 1$)	-35.5298	-39.7004	-24.6971	-25.435
Proposed ($\lambda = .5$)	-36.2388	-41.5057	-25.0232	-25.3715

TABLE II
THE MSEs OF THE PROPOSED ALGORITHM WITH DIFFERENT $\text{vol}(\mathbf{B})$'s.
 $(M, K) = (50, 5)$; $L = 1000$; $N_o = 20$; $\text{SNR} = 20$ dB.

$\text{vol}(\mathbf{B})$	measure	SOR (dB)			
		-10	-5	0	5
$\log \det(\mathbf{B}^T \mathbf{B} + \tau \mathbf{I})$	MSE	-32.3289	-33.1083	-33.0075	-32.9216
	TIME	6.262876	3.845999	3.328383	3.532759
$\det(\mathbf{B}^T \mathbf{B})$	MSE	-28.2461	-24.0797	-32.6881	-32.1538
	TIME	91.0616	35.11194	38.98876	39.80187
$\text{Tr}(\mathbf{G} \mathbf{B} \mathbf{B}^T)$	MSE	-28.4332	-28.5351	-28.413	-28.4776
	TIME	1.620391	1.540378	1.590791	1.517478

with RDR yields reasonable MSEs when $N_o \leq 40$, but its performance deteriorates when N_o is larger.

Table I presents the MSEs of the estimated $\hat{\mathbf{A}}$ under well- and ill-conditioned \mathbf{A} 's, respectively. To generate an ill-conditioned \mathbf{A} , we use a way that is similar to the method suggested in [11]: in each trial, we first generate $\tilde{\mathbf{A}}$ whose columns are uniformly distributed between zero and one, and such $\tilde{\mathbf{A}}$'s are relatively well-conditioned. Then, we apply singular value decomposition to obtain $\tilde{\mathbf{A}} = \mathbf{U} \tilde{\Sigma} \mathbf{V}^T$. Finally, we replace $\tilde{\Sigma}$ by $\hat{\Sigma} = \text{Diag}([1, 0.1, 0.01, 0.005, 0.001])$ and obtain $\mathbf{A} = \mathbf{U} \hat{\Sigma} \mathbf{V}^T$. This way, the condition number of the generated \mathbf{A} is 10^3 . The other settings are the same as those in Fig. 7. One can see from Table I that using such ill-conditioned \mathbf{A} , all the algorithms perform worse compared to the scenario where \mathbf{A} has uniformly distributed columns (cf. the first and second columns in Table I). Nevertheless, the proposed algorithm still gives the lowest MSEs.

In Table II, we present the MSE performance of the proposed algorithm using different volume regularizers. We see that using $\text{vol}(\mathbf{B}) = \text{Tr}(\mathbf{G} \mathbf{B} \mathbf{B}^T)$ has the shortest runtime since the

TABLE III
THE MSEs OF THE PROPOSED ALGORITHM WITH AND WITHOUT
NONNEGATIVITY CONSTRAINT ON \mathbf{B} UNDER VARIOUS SNRS.
(M, K) = (50, 5); $L = 1000$; $N_o = 20$; SOR = 5 dB.

Algorithm	measure	SNR (dB)			
		10	14	18	22
Proposed	MSE	-19.2268	-28.2876	-31.7755	-33.7787
	TIME	1.737251	2.033748	2.875241	4.152018
Proposed w/ nn	MSE	-19.6422	-28.6822	-31.9502	-33.8563
	TIME	5.09489	5.943189	7.863959	11.10508




Fig. 12. MSE of the proposed algorithm with different p 's under various SNRs. (M, K) = (50, 5); $L = 1000$; $N_o = 20$; SNR = 20 dB.

subproblem w.r.t. \mathbf{B} is convex and can be solved in closed form. When $\text{vol}(\mathbf{B}) = \det(\mathbf{B}^T \mathbf{B})$, the algorithm requires much more time compared to that of the other two regularizers. This is because the Armijo rule has to be implemented at each iteration. In terms of accuracy, using the $\log \det(\mathbf{B}^T \mathbf{B})$ regularizer gives the lowest MSEs when $\text{SOR} \leq -5$ dB. Using $\det(\mathbf{B}^T \mathbf{B})$ also exhibits good MSE performance when $\text{SOR} \geq 0$ dB. Using $\text{Tr}(\mathbf{G}\mathbf{B}\mathbf{B}^T)$ performs slightly worse in terms of MSE, since it is a coarse approximation to simplex volume. Interestingly, although our proposed log-determinant regularizer is not an exact measure of simplex volume as the determinant regularizer, it yields lower MSEs relative to the latter. Our understanding is that the performance gain results from the ease of computation.

Table III presents the MSE of the proposed algorithm with and without nonnegativity constraint on \mathbf{B} , respectively. We see that the MSEs are similar, with those of the nonnegativity-constrained algorithm being slightly lower. This result validates the soundness of our update rule for the constrained case, i.e., (18). In terms of speed, the unconstrained algorithm requires less time. We note that the nonnegativity constraint seems to only bring marginal performance gain in this simulation. This might be because the data are generated following the model in (1) and (2), and under this model VolMin identifiability does not depend on the nonnegativity of \mathbf{B} . However, when we are dealing with real data, adding nonnegativity constraints makes much sense, as will be shown in the next section.

Fig. 12 shows the effect of changing p . When $\text{SOR} = -10$ dB, we see that using $p \in [0.25, 0.75]$ gives relatively low MSEs. This is because using a small p is more effective in fending against outliers that largely deviate from the nominal model.




Fig. 13. The considered subimage of the Moffett data set.

It is interesting to note that using $p = 0.1$ gives slightly worse result compared to using $p \in [0.25, 0.75]$. Our understanding is that using a very small p may lead to numerical problems, since the weights $\{w_\ell\}_{\ell=1}^L$ can be scaled in a very unbalanced way in such cases, resulting in ill-conditioned optimization sub-problems. For the cases where $\text{SOR} = -5$ dB and 5 dB, a similar effect can be seen. In addition, a larger range of p , i.e., $p \in [0.25, 1.5]$, can result in good performance when $\text{SOR} = 5$ dB. The results suggest a strategy of choosing p : When the data is believed to be badly corrupted, using p around 0.5 is a good choice; and when the data is only moderately corrupted, using $p \in [1, 1.5]$ is preferable, since such a p gives good performance and can better avoid numerical problems.

V. REAL DATA VALIDATION

In this section, we validate the proposed algorithm using two real data sets, i.e., a hyperspectral image dataset with known outliers and a document dataset.

A. Hyperspectral Unmixing

Hyperspectral unmixing (HU) is the application where VolMin-based factorization is most frequently applied; see [3]. As introduced before, HU aims at estimating \mathbf{A} , i.e., the spectral signatures of the materials that are contained in a hyperspectral image, and also their proportions $s[\ell]$ in each pixel. It is well-known that there are outliers in hyperspectral images, due to the complicated reflection environment, spectral band contamination, and many other reasons [17]. In this experiment, we apply the proposed algorithm to a subimage of the real hyperspectral image that was captured over the Moffett Field in 1997 by the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS)¹; see Fig. 13. We remove the water absorption bands from the original 224 spectral bands, resulting in $M = 200$ bands for each pixel $\mathbf{x}[\ell]$. In this subimage with 50×50 pixels, there are three types of materials—water, soil, and vegetation. In the areas where different materials intersect, e.g., the lake shore, there are many outliers as identified by domain study. Our goal here is to test

¹Online available http://aviris.jpl.nasa.gov/data/image_cube.html.

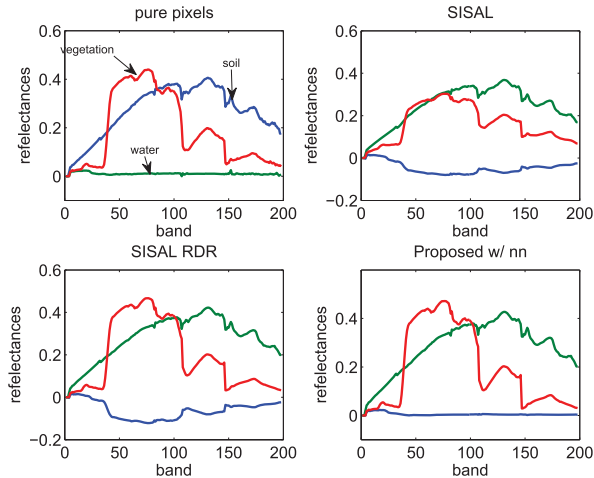


Fig. 14. The spectra of the manually selected pure pixels and the estimated spectra by the algorithms.

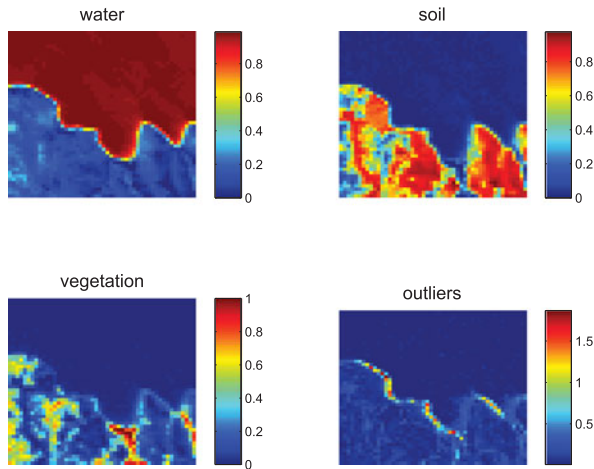


Fig. 15. The abundance maps of the materials and the distribution of the outliers obtained by the proposed algorithm.

whether our algorithm can identify the three materials and the outliers simultaneously.

We apply SISAL, SISAL with RDR, and the proposed algorithm to estimate \mathbf{A} . We set $\lambda = 1$ for our algorithm and tune η for SISAL carefully. Notice that we let $\mathcal{B} = \mathbb{R}_+^{M \times K}$ in this case, since the spectral signatures are known to be nonnegative. The estimated spectral signatures are shown in Fig. 14. As a benchmark, we also present the spectra of some manually selected pixels, which are considered purely contributed by only one material, and thus can approximately serve as the ground truth. We see that both SISAL and SISAL with RDR does not yield accurate estimates of \mathbf{A} . Particularly, the spectrum of water is badly estimated by both of these algorithms—one can see that there are many negative values of the spectra of water given by SISAL and SISAL with RDR. On the other hand, the proposed algorithm with nonnegativity constraint on \mathcal{B} gives spectra that are very similar to those of the pure pixels.

Fig. 15 shows the spatial distributions of the materials (i.e., the abundance maps $\{s_k[\ell]\}_{\ell=1}^L$ for $k = 1, \dots, K$) that are estimated by the proposed algorithm in the first three subimages. We see that the three materials are well separated, and their abundance maps are consistent with previous domain studies

TABLE IV
THE CLUSTERING ACCURACY ON REUTERS 21578 CORPUS

algorithm	number of clusters			
	3	4	5	6
LCCF	0.89042	0.8555	0.80367	0.73697
Proposed ($p = 2$)	0.88387	0.83681	0.77565	0.74542
Proposed ($p = 1.5$)	0.92221	0.87376	0.81392	0.78672
algorithm	number of clusters			
	7	8	9	10
LCCF	0.7131	0.65132	0.67441	0.6638
Proposed ($p = 2$)	0.72719	0.65828	0.65244	0.64662
Proposed ($p = 1.5$)	0.75217	0.67852	0.67149	0.67708

[41], [42]. In the last subimage of Fig. 15, we plot $1/w_\ell$ for $\ell = 1, \dots, L$. Notice that the weight w_ℓ corresponds to the ‘importance’ of the data $\mathbf{x}[\ell]$. The algorithm is designed to automatically give small w_ℓ to outlying pixels. We see that there are a lot of pixels on the lake shore that have very small weights, indicating that they are outliers. Physically, these outliers correspond to those areas where the solar light reflects several times between water and soil, resulting in nonlinearly mixed spectral signatures [41], [42]. The locations of the outlying pixels identified by our algorithm are also consistent with domain study [41], [42].

B. Document Clustering

We also present experimental results using the Reuters21578 document corpus². We use the subset of the full corpus provided by [43], which contains 8213 single-labeled documents from 41 clusters. In our experiment, we test our algorithm under different K (number of clusters), from 3 to 10. Following standard pre-processing, each document is represented as a term-frequency-inverse-document-frequency (tf-idf) vector, and *normalized cut weighting* is applied; see [43]–[45] for details. We apply our VolMin algorithm to factor the document data \mathbf{X} to ‘topics’ \mathbf{A} and a ‘weight matrix’ \mathbf{S} (cf. Fig. 1), and use \mathbf{S} to indicate the cluster labels of the documents. A regularized NMF-based approach, namely, *locally consistent concept factorization* (LCCF) [43] is employed as the baseline, which is considered a state-of-the-art algorithm for clustering the Reuters21578 corpus. For each K , we perform 100 Monte-Carlo trials by randomly selecting K clusters out of the total 41 clusters and 100 documents from each cluster. We report the performance by comparing the results with the ground truth. Performance is measured by a commonly used metric called *clustering accuracy*, whose detailed definition can be found in [43]—the clustering accuracy ranges from 0 to 1, and higher accuracies indicate better performances.

Table IV presents the results averaged from the 100 trials. For the proposed algorithm, we set $\lambda = 30$ and $p = 1.5$; we also present the result of $p = 2$ in this experiment, which we also use to initialize the $p < 2$ case. Note that here we use a larger λ relative to what was used in the simulations. The reason is that the document corpus contains considerable modeling errors and the fitting residue is relatively large. The rule of thumb for selecting λ is to set it at a similar level as the fitting error part; i.e., we let $\lambda = \mathcal{O}(\delta)$, where $\delta = \|\mathbf{X} - \hat{\mathbf{A}}\hat{\mathbf{S}}\|_F^2$ and can

²Online available: <http://www.daviddlewis.com/resources/testcollections/reuters21578/>

be coarsely estimated using plain NMF. This way, λ can balance the fitting part and the volume regularizer. From Table IV, we see that VolMin with $p = 2$ already yields comparable clustering accuracy with LCCF. This indicates that, even without outlier-robustness, modeling the document clustering problem using VolMin-based SMF is effective. Better accuracies can be seen by using $p = 1.5$, where we see that for most K , the proposed RVolMin algorithm gives the best accuracy. In particular, for $K = 6, 7$, more than 4% accuracy improvement can be seen, which is considered significant in the context of document clustering. Interestingly, further decreasing p does not yield better performance. This implies that the modeling error is not very severe, but outliers do exist, since using $p < 2$ gives better clustering result than using $p = 2$ which is not robust to modeling errors.

VI. CONCLUSION

In this work, we looked into theoretical and practical aspects of the VolMin criterion for matrix factorization. On the theory side, we showed that two independently developed sufficient conditions for VolMin identifiability are in fact equivalent. On the practical side, we proposed an outlier-robust optimization surrogate of the VolMin criterion, and devised an inexact BCD algorithm to deal with it. Extensive simulations showed that the proposed algorithm outperforms a state-of-the-art robust VolMin algorithm, i.e., SISAL. The proposed algorithm was also validated using real-world hyperspectral image data and document data, where interesting and favorable results were observed.

APPENDIX A

A. Proof of Theorem 3

To show Theorem 3, several properties of convex cones will be constantly used. They are:

Property 1: Let \mathcal{K}_1 and \mathcal{K}_2 be convex cones. Then, $\mathcal{K}_1 \subseteq \mathcal{K}_2 \Rightarrow \mathcal{K}_2^* \subseteq \mathcal{K}_1^*$.

Property 2: Let \mathcal{K}_1 and \mathcal{K}_2 be convex cones. Then, $(\mathcal{K}_1 \cap \mathcal{K}_2)^* = \text{conv}\{\mathcal{K}_1^* \cup \mathcal{K}_2^*\}$.

Property 3: If \mathbf{Q} is a unitary matrix. Then, $\text{cone}(\mathbf{Q})^* = \text{cone}(\mathbf{Q})$.

We show Theorem 3 step by step. First, we show the following lemma:

Lemma 3: Assume that $\mathcal{C} \subseteq \text{cone}(\mathbf{S})$ and \mathbf{Q} is any unitary matrix except the permutation matrices. Then, we have $\text{cone}(\mathbf{S})^* \cap \text{bd}\mathcal{C}^* = \{\lambda_i \mathbf{e}_i | i = 1, \dots, N\} \Leftrightarrow \text{cone}(\mathbf{S}) \not\subseteq \text{cone}(\mathbf{Q})$.

Proof: We first show the “ \Rightarrow ” part. Given a unitary \mathbf{Q} , suppose that

$$\text{cone}(\mathbf{S}) \subseteq \text{cone}(\mathbf{Q}).$$

By the basic properties of convex cones, we see that $\text{cone}(\mathbf{Q})^* \subseteq \text{cone}(\mathbf{S})^*$, and $\text{cone}(\mathbf{Q})^* = \text{cone}(\mathbf{Q})$. Combining, we see

$$\text{cone}(\mathbf{Q}) \subseteq \text{cone}(\mathbf{S})^*. \quad (19)$$

Also, we have

$$\mathcal{C} \subseteq \text{cone}(\mathbf{S}) \Rightarrow \text{cone}(\mathbf{S})^* \subseteq \mathcal{C}^* \Rightarrow \text{cone}(\mathbf{Q}) \subseteq \mathcal{C}^*. \quad (20)$$

Combining Eqs. (19) and (20), we have

$$\text{cone}(\mathbf{Q}) \subseteq \mathcal{C}^* \cap \text{cone}(\mathbf{S})^*. \quad (21)$$

We also know that the extreme rays of $\text{cone}(\mathbf{Q})$ lie in the boundary of \mathcal{C}^* , i.e., $\text{ex}\{\text{cone}(\mathbf{Q})\} \subseteq \text{bd}\mathcal{C}^*$ [9, Lemma 1]. Thus, we have

$$\text{ex}\{\text{cone}(\mathbf{Q})\} \subseteq \text{bd}\mathcal{C}^* \cap \text{cone}(\mathbf{S})^*. \quad (22)$$

Since we assumed $\text{cone}(\mathbf{S})^* \cap \text{bd}\mathcal{C}^* = \{\lambda_i \mathbf{e}_i | i = 1, \dots, N\}$, we have

$$\text{ex}\{\text{cone}(\mathbf{Q})\} \subseteq \{\mathbf{e}_1, \dots, \mathbf{e}_N\}. \quad (23)$$

Therefore, \mathbf{Q} can only be a permutation matrix.

We now show the “ \Leftarrow ” part. Following (22), and knowing that $\text{cone}(\mathbf{S})$ is a subset of the convex cone of some permutation matrix, we see that

$$\{\mathbf{e}_1, \dots, \mathbf{e}_N\} \subseteq \text{cone}(\mathbf{S})^* \cap \text{bd}\mathcal{C}^*. \quad (24)$$

Now, suppose that there are a set of vectors $\{\mathbf{r}_1, \dots, \mathbf{r}_p\}$ that does not include any unit vectors such that

$$\text{cone}(\mathbf{S})^* \cap \text{bd}\mathcal{C}^* = \{\mathbf{e}_1, \dots, \mathbf{e}_N, \mathbf{r}_1, \dots, \mathbf{r}_p\}.$$

Then, we see that we can represent $\text{cone}(\mathbf{S})^* = \text{conv}\{\mathbb{R}_+^N \cup \text{cone}\{\mathbf{r}_1, \dots, \mathbf{r}_p\}\}$. By Property 2, we see that

$$\begin{aligned} \text{cone}(\mathbf{S}) &= \text{conv}\{\mathbb{R}_+^N \cup \text{cone}\{\mathbf{r}_1, \dots, \mathbf{r}_p\}\}^* \\ &= \mathbb{R}_+^N \cap \text{cone}\{\mathbf{r}_1, \dots, \mathbf{r}_p\}^*. \end{aligned} \quad (25)$$

Since $\mathbb{R}_+^N \cap \text{cone}\{\mathbf{r}_1, \dots, \mathbf{r}_p\}^* = \text{cone}(\mathbf{S}) \subseteq \mathbb{R}_+^N$, (25) leads to

$$\text{cone}\{\mathbf{r}_1, \dots, \mathbf{r}_p\}^* \subseteq \mathbb{R}_+^N.$$

By Property 1, we see that

$$(\mathbb{R}_+^N)^* = \mathbb{R}_+^N \subseteq \text{cone}\{\mathbf{r}_1, \dots, \mathbf{r}_p\}.$$

This is a contradiction to the assumption that $\{\mathbf{r}_1, \dots, \mathbf{r}_p\}$ does not include the unit vectors. \blacksquare

Now we are ready to prove Theorem 3. We first notice that $\mathcal{C} \subseteq \text{cone}(\mathbf{S})$ is equivalent to $\gamma \geq \frac{1}{N-1}$. In fact, we see that $\mathcal{C} \cap \mathcal{S} = \mathcal{R}(\frac{1}{\sqrt{N-1}})$, where $\mathcal{S} = \{\mathbf{s} | \mathbf{1}^T \mathbf{s} = 1, \mathbf{s} \in \mathbb{R}^N\}$, and thus the claim holds. Thus, our remaining work is to show that condition (ii) in Theorem 1 is equivalent to restricting γ such that $\gamma > \frac{1}{\sqrt{N-1}}$.

Step 1: Let us consider a conic representation of Theorem 2. Specifically, the corresponding convex cone of $\mathcal{R}(r) = \{\mathbf{s} \in \mathbb{R}^N \mid \|\mathbf{s}\|_2 \leq r\} \cap \text{conv}\{\mathbf{e}_1, \dots, \mathbf{e}_N\}$ is $\tilde{\mathcal{R}}(r) = \mathcal{C}(r) \cap \mathbb{R}_+^N$, where

$$\mathcal{C}(r) = \{\mathbf{s} \in \mathbb{R}^N \mid \|\mathbf{s}\|_2 \leq r \mathbf{1}^T \mathbf{s}\},$$

and $\gamma = \sup\{r | \mathcal{C}(r) \subseteq \text{cone}(\mathbf{S})\}$ under this definition. It is also noticed that $\mathcal{C}(r)$ can be re-expressed as

$$\mathcal{C}(r) = \left\{ \mathbf{s} \mid \frac{\mathbf{1}^T \mathbf{s}}{\|\mathbf{1}\|_2 \|\mathbf{s}\|_2} \geq \frac{1}{r\sqrt{N}} \right\}.$$

In words, the vectors whose angles between $\mathbf{1}$ are less than or equal to $\arccos \frac{1}{r\sqrt{N}}$ comprises $\mathcal{C}(r)$. Therefore, by the definition of dual cone, i.e., $\mathcal{C}(r)^* = \{\mathbf{s} | \mathbf{s}^T \mathbf{y} \geq 0, \mathbf{y} \in \mathcal{C}\}$, $\mathcal{C}(r)^*$

contains all the vectors that have the angle with $\mathbf{1}$ less than or equal to $\arccos \frac{r}{\sqrt{r^2 N - 1}}$, which leads to

$$\mathcal{C}(r)^* = \mathcal{C} \left(\frac{r}{\sqrt{r^2 N - 1}} \right). \quad (26)$$

Step 2: Now, we consider the dual cone representation of Theorem 2. Let us define

$$\mathcal{T}(r) = \text{conv}(\mathcal{C}(r) \cup \mathbb{R}_+^N).$$

Then, following Property 2 and (26), we have

$$\tilde{\mathcal{R}}(r)^* = (\mathcal{C}(r) \cap \mathbb{R}_+^N)^* \quad (27a)$$

$$= \text{conv} \left(\mathcal{C} \left(\frac{r}{\sqrt{r^2 N - 1}} \right) \cup \mathbb{R}_+^N \right), \quad (27b)$$

$$= \mathcal{T} \left(\frac{r}{\sqrt{r^2 N - 1}} \right), \quad (27c)$$

where we have used $(\mathbb{R}_+^N)^* = \mathbb{R}_+^N$, i.e., Property 3. According to Property 1, we see that

$$\tilde{\mathcal{R}}(r) \subseteq \text{cone}(\mathcal{S}) \Leftrightarrow \text{cone}(\mathcal{S})^* \subseteq \mathcal{T} \left(\frac{r}{\sqrt{r^2 N - 1}} \right). \quad (28)$$

If we define $\kappa = \inf\{r | \text{cone}(\mathcal{S})^* \subseteq \mathcal{T}(r)\}$, we see from (28) and the definition of γ that

$$\gamma > \frac{1}{\sqrt{N-1}} \Leftrightarrow \kappa < 1. \quad (29)$$

Step 3: To show the equivalence between the sufficient conditions, let us begin from Theorem 1. The condition $\mathcal{C} \subseteq \text{cone}(\mathcal{S})$ means that $\mathcal{C}^* \subseteq \text{cone}(\mathcal{S})^*$ by Property (1), and it further implies that $\{e_1, \dots, e_N\} \subseteq \text{ex}\{\text{cone}(\mathcal{S})^*\}$ [9]. Suppose the rest of $\text{cone}(\mathcal{S})^*$'s extreme rays are $\mathbf{r}_1, \dots, \mathbf{r}_p$, and t is defined as $t = \inf\{r | \text{cone}\{\mathbf{r}_1, \dots, \mathbf{r}_p\} \subseteq \mathcal{C}(r)\}$. Then, we have

$$\begin{aligned} \text{cone}(\mathcal{S})^* &= \text{cone}\{\mathbf{r}_1, \dots, \mathbf{r}_p, e_1, \dots, e_N\} \\ &\subseteq \text{conv}\{\mathcal{C}(t) \cup \mathbb{R}_+^N\} = \mathcal{T}(t). \end{aligned}$$

Hence, by the definitions of t and κ , we have $t = \kappa$.

Given the above analysis, we first show that Theorem 1 implies Theorem 2. Now, assume that Condition (ii) in Theorem 1 is satisfied. We see that $\text{cone}(\mathcal{S})^* \cap \text{bd}\mathcal{C}^* = \{\lambda_i e_i | i = 1, \dots, N\}$ by Lemma 3. Then, $\mathbf{r}_1, \dots, \mathbf{r}_p$ are in the interior of $\mathcal{C}^* = \mathcal{C}(1)$. Therefore, we have $t < 1$, and subsequently $\kappa < 1$ and $\gamma > \frac{1}{\sqrt{N-1}}$ strictly.

Now we show the converse by contradiction. Assume that $\gamma > \frac{1}{\sqrt{N-1}}$ holds (the condition in Theorem 2 is satisfied). If at least one point in $\mathbf{r}_1, \dots, \mathbf{r}_p$ touches the boundary of \mathcal{C}^* (condition (ii) in Theorem 1 is not satisfied), then $t = 1$, and we cannot decrease it further while still contain $\text{cone}(\mathcal{S})^*$ in $\mathcal{T}(t)^*$. This means that $\kappa = 1$, or, equivalently, $\gamma = \frac{1}{\sqrt{N-1}}$, which contradicts our first assumption that $\gamma > \frac{1}{\sqrt{N-1}}$.

B. Proof of Proposition 1

In the following, we prove the proposition under the algorithmic structure without extrapolation. For the case where we update \mathbf{C} using (13), it is easy to see that $\mathbf{y}^t[\ell] \rightarrow \mathbf{c}^t[\ell]$ given

$t \rightarrow \infty$ by (13c). Hence, if the proposition holds for the algorithm without extrapolation, it also holds for the extrapolated version asymptotically.

First, let us cast the proposed algorithm into the framework of *block successive upper bound minimization* (BSUM) [29], [37]. Unlike the classic block coordinate descent algorithm that solves every block subproblem exactly [28], BSUM cyclically solves the upper-bound problems of every block subproblems. We consider the updates using (12) and (16) as an example. The proof of using other updates will follow. Our update rule in (12) and (16) can be equivalently written as

$$\mathbf{C}^{t+1} = \arg \min_{\mathbf{1}^T \mathbf{C} = \mathbf{1}^T, \mathbf{C} \geq \mathbf{0}} u_{\mathbf{C}}(\mathbf{C}; \mathbf{B}^t) \quad (30a)$$

$$\mathbf{B}^{t+1} = \arg \min_{\mathbf{B}} u_{\mathbf{B}}(\mathbf{B}; \mathbf{C}^{t+1}). \quad (30b)$$

where $u_{\mathbf{C}}(\mathbf{C}; \mathbf{B}^t) = \sum_{\ell=1}^L \frac{1}{2} (2u(\mathbf{c}[\ell]; \mathbf{B}^t) + \epsilon)^{p/2} + \frac{\lambda}{2} \log \det((\mathbf{B}^t)^T \mathbf{B}^t + \tau \mathbf{I})$, $u(\mathbf{c}[\ell]; \mathbf{B}^t)$ is defined as before, and

$$\begin{aligned} u_{\mathbf{B}}(\mathbf{B}; \mathbf{C}^{t+1}) &= \sum_{\ell=1}^L \frac{w_{\ell}}{2} \|\mathbf{x}[\ell] - \mathbf{B} \mathbf{c}^{t+1}[\ell]\|_2^2 \\ &\quad + \frac{\lambda}{2} \text{Tr}(\mathbf{F}^t (\mathbf{B}^T \mathbf{B})) + \text{const}, \end{aligned}$$

in which $\text{const} = \sum_{\ell=1}^L \phi_p(w_{\ell}^t) - K$. Note that solving (30a) is equivalent to solving (12) over different ℓ 's since the problems w.r.t. $\ell = 1, \dots, L$ are not coupled. Also denote $v(\mathbf{B}, \mathbf{C})$ as the objective value of Problem (8). When $L^t \geq \|(\mathbf{B}^t)^T \mathbf{B}^t\|_2$, we have

$$v(\mathbf{B}^t, \mathbf{C}) \leq u_{\mathbf{C}}(\mathbf{C}; \mathbf{B}^t), \quad \forall \mathbf{C} \quad (31a)$$

$$v(\mathbf{B}, \mathbf{C}^{t+1}) \leq u_{\mathbf{B}}(\mathbf{B}; \mathbf{C}^{t+1}), \quad \forall \mathbf{B}, \quad (31b)$$

where (31a) holds because under $L^t \geq \|(\mathbf{B}^t)^T \mathbf{B}^t\|_2$ we have

$$f(\mathbf{c}[\ell]; \mathbf{B}^t) = \frac{1}{2} \|\mathbf{x}[\ell] - \mathbf{B}^t \mathbf{c}[\ell]\|_2^2 \leq u(\mathbf{c}[\ell]; \mathbf{B}^t), \quad \forall \mathbf{c}[\ell],$$

and thus

$$\begin{aligned} v(\mathbf{B}^t, \mathbf{C}) &= \sum_{\ell=1}^L \frac{1}{2} (2f(\mathbf{c}[\ell]; \mathbf{B}^t) + \epsilon)^{\frac{p}{2}} \\ &\quad + \frac{\lambda}{2} \log \det((\mathbf{B}^t)^T \mathbf{B}^t + \tau \mathbf{I}) \\ &\leq \sum_{\ell=1}^L \frac{1}{2} (2u(\mathbf{c}[\ell]; \mathbf{B}^t) + \epsilon)^{\frac{p}{2}} \\ &\quad + \frac{\lambda}{2} \log \det((\mathbf{B}^t)^T \mathbf{B}^t + \tau \mathbf{I}) \\ &= u_{\mathbf{C}}(\mathbf{C}; \mathbf{B}^t); \end{aligned}$$

Eq. (31b) holds because of Lemmas 1 and 2; also note that the equalities hold when $\mathbf{C} = \mathbf{C}^t$ and $\mathbf{B} = \mathbf{B}^t$, respectively. Since all the functions above are continuously differentiable, we also have

$$\nabla_{\mathbf{C}} v(\mathbf{B}^t, \mathbf{C}^t) = \nabla_{\mathbf{C}} u_{\mathbf{C}}(\mathbf{C}^t; \mathbf{B}^t) \quad (32a)$$

$$\nabla_{\mathbf{B}} v(\mathbf{B}^t, \mathbf{C}^{t+1}) = \nabla_{\mathbf{B}} u_{\mathbf{B}}(\mathbf{B}^t; \mathbf{C}^{t+1}). \quad (32b)$$

Note that if we update \mathbf{C} using ADMM, we have $v(\mathbf{B}^t, \mathbf{C}) = u_{\mathbf{C}}(\mathbf{C}; \mathbf{B}^t)$ for all \mathbf{C} -Eqs. (31a) and (32a) still hold. Also, if

we update \mathbf{B} by (18), the conditions in (31b) and (32b) are also satisfied when $\mu^t \geq \|(\mathbf{F}^t)^T \mathbf{F}^t\|_2$ since we now we have

$$u_B(\mathbf{B}; \mathbf{C}^{t+1}) = g(\mathbf{B}^t; \mathbf{C}^{t+1}) + \nabla g(\mathbf{B}^t; \mathbf{C}^{t+1})^T (\mathbf{B} - \mathbf{B}^t) + \frac{\mu^t}{2} \|\mathbf{B} - \mathbf{B}^t\|_F^2,$$

and it can be shown that $v(\mathbf{B}, \mathbf{C}^{t+1}) \leq g(\mathbf{B}; \mathbf{C}^{t+1}) \leq u_B(\mathbf{B}; \mathbf{C}^{t+1})$ and

$$\nabla_{\mathbf{B}} v(\mathbf{B}^t, \mathbf{C}^{t+1}) = \nabla_{\mathbf{B}} g(\mathbf{B}^t; \mathbf{C}^{t+1}) = \nabla_{\mathbf{B}} u_B(\mathbf{B}^t; \mathbf{C}^{t+1}),$$

and the equalities hold simultaneously at $\mathbf{B} = \mathbf{B}^t$. Eqs. (31), (32) satisfy the sufficient conditions for a generic BSUM algorithm to converge (cf. Assumption 2 in [37]). In addition, by [37, Theorem 2 (b)], if we can show that $(\mathbf{B}^t, \mathbf{C}^t)$ lives in a compact set for all t , we can prove Proposition 1.

Next, we show that in every iteration, \mathbf{B}^t and \mathbf{C}^t are bounded. The boundness of \mathbf{C}^t is evident because we enforce feasibility at each iteration. To show that \mathbf{B}^t is bounded, we first note that

$$v(\mathbf{B}^t, \mathbf{C}^t) \geq v(\mathbf{B}^{t+1}, \mathbf{C}^{t+1}),$$

where the inequality holds since the non-increasing property of the BSUM framework [37]. By the assumption that $v(\mathbf{B}^0, \mathbf{C}^0)$ is bounded, i.e.,

$$v(\mathbf{B}^0, \mathbf{C}^0) \leq V,$$

where $V < \infty$, we have

$$\sum_{\ell=1}^L \frac{1}{2} \left(\|\mathbf{x}[\ell] - \mathbf{B}\mathbf{c}[\ell]\|_2^2 + \epsilon \right)^{\frac{\lambda}{2}} + \frac{\lambda}{2} \log \det(\mathbf{B}^T \mathbf{B} + \tau \mathbf{I}) \leq V$$

holds for every $(\mathbf{B}, \mathbf{C}) \in \{\mathbf{B}^t, \mathbf{C}^t\}_{t=1, \dots, \dots}$, that is generated by the algorithm. Since the first term on the left hand side of the above inequality is nonnegative, we have

$$\log \det(\mathbf{B}^T \mathbf{B} + \tau \mathbf{I}) \leq V \Leftrightarrow \log \left(\prod_{i=1}^N (\sigma_i^2 + \tau) \right) \leq V$$

$$\Rightarrow \log(\sigma_i^2 + \tau) \leq V - (N - 1) \log \tau, \forall i \quad (33a)$$

$$\Rightarrow \sigma_i^2 \leq \exp(V - (N - 1) \log \tau) - \tau, \forall i, \quad (33b)$$

where $\sigma_1, \dots, \sigma_N$ denote the singular values of \mathbf{B} , and (33a) holds since $\log(\sigma_i^2 + \tau) \geq \log \tau$ for all i . The right hand side of (33b) is bounded, which implies that every singular value of \mathbf{B}^t for $t = 1, 2, \dots$ is bounded. Since \mathcal{B} is a closed convex set, we conclude that the sequence $\{\mathbf{B}^t, \mathbf{C}^t\}_t$ lies in a compact set. Now, invoking [37, Theorem 2 (b)], the proof is completed.

ACKNOWLEDGMENT


The authors would like to thank Prof. N. Dobigeon for providing the subimage of the Moffet data.

REFERENCES

- [1] X. Fu, W.-K. Ma, K. Huang, and N. Sidiropoulos, "Robust volume minimization-based structured matrix factorization via alternating optimization," in *Proc. 2016 IEEE Int. Conf. Acoust., Speech, Signal Process.*, Mar. 2016, pp. 2534–2538.
- [2] D. Lee and H. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
- [3] W.-K. Ma *et al.*, "A signal processing perspective on hyperspectral unmixing," *IEEE Signal Process. Mag.*, vol. 31, no. 1, pp. 67–81, Jan. 2014.
- [4] N. Gillis, "The why and how of nonnegative matrix factorization," in *Regularization, Optimization, Kernels, and Support Vector Machines*, vol. 12. London, U.K.: Chapman & Hall, 2014, pp. 257–291.


- [5] X. Fu, W.-K. Ma, K. Huang, and N. D. Sidiropoulos, "Blind separation of quasi-stationary sources: Exploiting convex geometry in covariance domain," *IEEE Trans. Signal Process.*, vol. 63, no. 9, pp. 2306–2320, May 2015.
- [6] X. Fu, N. D. Sidiropoulos, and W.-K. Ma, "Tensor-based power spectra separation and emitter localization for cognitive radio," in *Proc. IEEE 8th Sensor Array Multichannel Signal Process. Workshop*, 2014, pp. 421–424.
- [7] D. Donoho and V. Stodden, "When does non-negative matrix factorization give a correct decomposition into parts?" in *Proc. 17th Annu. Conf. Neural Inf. Process. Syst.*, 2003, vol. 16, pp. 1141–1148.
- [8] H. Laurberg, M. G. Christensen, M. D. Plumbley, L. K. Hansen, and S. Jensen, "Theorems on positive data: On the uniqueness of NMF," *Comput. Intell. Neurosci.*, vol. 2008, 2008, Art. ID. 764206.
- [9] K. Huang, N. D. Sidiropoulos, and A. Swami, "Non-negative matrix factorization revisited: New uniqueness results and algorithms," *IEEE Trans. Signal Process.*, vol. 62, no. 1, pp. 211–224, Jan. 2014.
- [10] C.-H. Lin, W.-K. Ma, W.-C. Li, C.-Y. Chi, and A. Ambikapathi, "Identifiability of the simplex volume minimization criterion for blind hyperspectral unmixing: The no-pure-pixel case," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 10, pp. 5530–5546, Oct. 2015.
- [11] N. Gillis and S. Vavasis, "Fast and robust recursive algorithms for separable nonnegative matrix factorization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 4, pp. 698–714, Apr. 2014.
- [12] J. Li and J. M. Bioucas-Dias, "Minimum volume simplex analysis: A fast algorithm to unmix hyperspectral data," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2008, vol. 3, pp. 250–253.
- [13] T.-H. Chan, C.-Y. Chi, Y.-M. Huang, and W.-K. Ma, "A convex analysis-based minimum-volume enclosing simplex algorithm for hyperspectral unmixing," *IEEE Trans. Signal Process.*, vol. 57, no. 11, pp. 4418–4432, Nov. 2009.
- [14] A. Ambikapathi, T.-H. Chan, W.-K. Ma, and C.-Y. Chi, "Chance-constrained robust minimum-volume enclosing simplex algorithm for hyperspectral unmixing," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 11, pp. 4194–4209, Nov. 2011.
- [15] L. Miao and H. Qi, "Endmember extraction from highly mixed data using minimum volume constrained nonnegative matrix factorization," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 3, pp. 765–777, Mar. 2007.
- [16] G. Zhou, S. Xie, Z. Yang, J.-M. Yang, and Z. He, "Minimum-volume-constrained nonnegative matrix factorization: Enhanced ability of learning parts," *IEEE Trans. Neural Netw.*, vol. 22, no. 10, pp. 1626–1637, Oct. 2011.
- [17] N. Dobigeon, J.-Y. Tourneret, C. Richard, J. Bermudez, S. McLaughlin, and A. O. Hero, "Nonlinear unmixing of hyperspectral images: Models and algorithms," *IEEE Signal Process. Mag.*, vol. 31, no. 1, pp. 82–94, Dec. 2014.
- [18] J. M. Bioucas-Dias, "A variable splitting augmented lagrangian approach to linear spectral unmixing," in *Proc. 1st Workshop Hyperspectral Image Signal Process. Evol. Remote Sens.*, 2009, pp. 1–4.
- [19] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [20] X. Fu, W.-K. Ma, T.-H. Chan, and J. M. Bioucas-Dias, "Self-dictionary sparse regression for hyperspectral unmixing: Greedy pursuit and pure pixel search are related," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 6, pp. 1128–1141, Sep. 2015.
- [21] P. Grizzmann, V. Klee, and D. Larman, "Largest j -simplices in n -polytopes," *Discr. Comput. Geom.*, vol. 13, no. 1, pp. 477–515, 1995.
- [22] M. Fazel, H. Hindi, and S. P. Boyd, "Log-det heuristic for matrix rank minimization with applications to hankel and euclidean distance matrices," in *Proc. Amer. Control Conf. 2003*, 2003, vol. 3, pp. 2156–2162.
- [23] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, and Y. Ma, "Robust recovery of subspace structures by low-rank representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 171–184, Jan. 2013.
- [24] Y.-F. Liu, S. Ma, Y.-H. Dai, and S. Zhang, "A smoothing SQP framework for a class of composite l_q minimization over polyhedron," *Math. Program.*, pp. 1–34, 2015.
- [25] H. Eklblom, " l_p -methods for robust regression," *BIT Numer. Math.*, vol. 14, no. 1, pp. 22–32, 1974.
- [26] S. A. Vorobyov, Y. Rong, N. D. Sidiropoulos, and A. B. Gershman, "Robust iterative fitting of multilinear models," *IEEE Trans. Signal Process.*, vol. 53, no. 8, pp. 2678–2689, Aug. 2005.
- [27] M. Berman, H. Kiiveri, R. Lagerstrom, A. Ernst, R. Dunne, and J. Huntington, "ICE: A statistical approach to identifying endmembers in hyperspectral images," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 10, pp. 2085–2095, Oct. 2004.
- [28] D. P. Bertsekas, *Nonlinear Programming*. Belmont, MA, USA: Athena Scientific, 1999.
- [29] M. Hong, M. Razaviyayn, Z.-Q. Luo, and J.-S. Pang, "A unified algorithmic framework for block-structured optimization involving big data: With applications in machine learning and signal processing," *IEEE Signal Process. Mag.*, vol. 33, no. 1, pp. 57–77, Jan. 2016.

- [30] W. Wang and M. A. Carreira-Perpiñán, "Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application," 2013, arxiv preprint, arxiv: 1309.1541v1.
- [31] Y. Nesterov, *Introductory Lectures on Convex Optimization*, vol. 87. New York, NY, USA: Springer, 2004.
- [32] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [33] Y. Xu and W. Yin, "A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion," *SIAM J. Imag. Sci.*, vol. 6, no. 3, pp. 1758–1789, 2013.
- [34] X. Fu, K. Huang, W.-K. Ma, N. Sidiropoulos, and R. Bro, "Joint tensor factorization and outlying slab suppression with applications," *IEEE Trans. Signal Process.*, vol. 63, no. 23, pp. 6315–6328, Dec. 2015.
- [35] J. Jose, N. Prasad, M. Khojastepour, and S. Rangarajan, "On robust weighted-sum rate maximization in MIMO interference networks," in *Proc. IEEE Int. Conf. Commun.*, 2011, pp. 1–6.
- [36] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 123–231, 2013.
- [37] M. Razaviyayn, M. Hong, and Z.-Q. Luo, "A unified convergence analysis of block successive minimization methods for nonsmooth optimization," *SIAM J. Optim.*, vol. 23, no. 2, pp. 1126–1153, 2013.
- [38] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *J. ACM*, vol. 58, no. 3, p. 11, 2011.
- [39] F. Nie, J. Yuan, and H. Huang, "Optimal mean robust principal component analysis," in *Proc. 31st Int. Conf. Mach. Learn.*, 2014, pp. 1062–1070.
- [40] K. Huang and N. Sidiropoulos, "Putting nonnegative matrix factorization to the test: A tutorial derivation of pertinent Cramer–Rao bounds and performance benchmarking," *IEEE Signal Process. Mag.*, vol. 31, no. 3, pp. 76–86, May 2014.
- [41] C. Févotte and N. Dobigeon, "Nonlinear hyperspectral unmixing with robust nonnegative matrix factorization," *IEEE Trans. Image Process.*, vol. 24, no. 12, pp. 4810–4819, Aug. 2015.
- [42] A. Halimi, Y. Altmann, N. Dobigeon, and J.-Y. Tourneret, "Nonlinear unmixing of hyperspectral images using a generalized bilinear model," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 11, pp. 4153–4162, Nov. 2011.
- [43] D. Cai, X. He, and J. Han, "Locally consistent concept factorization for document clustering," *IEEE Trans. Knowl. Data Eng.*, vol. 23, no. 6, pp. 902–913, Jun. 2011.
- [44] W. Xu and Y. Gong, "Document clustering by concept factorization," in *Proc. 27th Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2004, pp. 202–209.
- [45] C. D. Manning, P. Raghavan, and H. Schütze, *Introduction to Information Retrieval*, vol. 1. Cambridge, U.K.: Cambridge Univ. Press, 2008.



Xiao Fu (S'12–M'15) received the B.Eng. and M.Eng. degrees in communication and information engineering from the University of Electronic Science and Technology of China, Chengdu, China, in 2005 and 2010, respectively and the Ph.D. degree in electronic engineering from the Chinese University of Hong Kong (CUHK), Shatin, N.T., Hong Kong, in 2014. From 2005 to 2006, he was an Assistant Engineer at China Telecom Co. Ltd., Shenzhen, China. He is currently a Postdoctoral Associate at the Department of Electrical and Computer Engineering,

University of Minnesota, Minneapolis, MN, USA. His research interests include signal processing and machine learning, with a recent emphasis on factor analysis and its applications. He received the Best Student Paper Award at IEEE International Conference on Acoustics, Speech and Signal Processing 2014, and was a finalist of the Best Student Paper Competition at IEEE Systems Analysis and Modeling Workshop 2014. He also co-authored a paper that received the Best Student Paper Award at IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing 2015.



Kejun Huang (S'13) received the bachelor's degree from Nanjing University of Information Science and Technology, Nanjing, China, in 2010, and the Ph.D. degree in electrical engineering from the University of Minnesota, Minneapolis, MN, USA, in 2016. His research interests include signal processing, machine learning, and optimization.



Bo Yang (S'15) received the B.Eng. and M.Eng. degrees in communication and information system from the Huazhong University of Science and Technology of China, Wuhan, China, in 2011 and 2014, respectively. Since 2014, he has been working toward the Ph.D. degree in the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN, USA. His research interests include data analytics and machine learning, with a recent emphasis on factor analysis. He received the Best Student Paper Award at International Workshop on Computational Advances in Multi-Sensor Adaptive Processing 2015.



Wing-Kin Ma (M'01–SM'11) received the B.Eng. degree in electrical and electronic engineering from the University of Portsmouth, Portsmouth, U.K., in 1995, and the M.Phil. and Ph.D. degrees, both in electronic engineering, from The Chinese University of Hong Kong (CUHK), Shatin, N.T., Hong Kong, in 1997 and 2001, respectively. He is currently an Associate Professor with the Department of Electronic Engineering, CUHK. From 2005 to 2007, he was also an Assistant Professor with the Institute of Communications Engineering, National Tsing Hua University, Taiwan, R.O.C. Prior to becoming a faculty member, he held various research positions with McMaster University, Canada; CUHK; and the University of Melbourne, Melbourne, Vic, Australia. His research interests include signal processing, communications and optimization, with a recent emphasis on MIMO transceiver designs and interference management, blind separation and structured matrix factorization, and hyperspectral unmixing in remote sensing.

He currently serves as a Senior Area Editor of IEEE TRANSACTIONS ON SIGNAL PROCESSING and Associate Editor of *Signal Processing*. Also, he has previously served as an Associate Editor and a Guest Editor of several journals, which include IEEE TRANSACTIONS ON SIGNAL PROCESSING, IEEE SIGNAL PROCESSING LETTERS, IEEE JOURNAL OF SELECTED AREAS IN COMMUNICATIONS, and IEEE SIGNAL PROCESSING MAGAZINE. He was a tutorial speaker in European Signal Processing Conference 2011 and International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2014. He is currently a Member of the Signal Processing Theory and Methods Technical Committee (SPTM-TC) and the Signal Processing for Communications and Networking Technical Committee (SPCOM-TC). He received several awards, such as ICASSP Best Student Paper Awards, in 2011 and 2014, respectively (by his students), WHISPERS 2011 Best Paper Award, Research Excellence Award 2013–2014 by CUHK, and 2015 IEEE SIGNAL PROCESSING MAGAZINE Best Paper Award.



Nicholas D. Sidiropoulos (F'09) received the Diploma degree in electrical engineering from the Aristotelian University of Thessaloniki, Thessaloniki, Greece, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland—College Park, College Park, MD, USA, in 1988, 1990, and 1992, respectively. He served as an Assistant Professor at the University of Virginia (1997–1999); an Associate Professor at the University of Minnesota, Minneapolis, MN, USA (2000–2002); a Professor at the Technical University of Crete, Greece (2002–2011); and a Professor at the University of Minnesota (2011–present). His current research interest includes signal and tensor analytics, with applications in cognitive radio, big data, and preference measurement. He received the NSF/CAREER award (1998), the IEEE Signal Processing Society (SPS) Best Paper Award (2001, 2007, 2011), and the IEEE SPS Meritorious Service Award (2010). He has served as an IEEE SPS Distinguished Lecturer (2008–2009), and a Chair of the IEEE Signal Processing for Communications and Networking Technical Committee (2007–2008). He received the Distinguished Alumni Award of the Department of Electrical and Computer Engineering, University of Maryland, College Park, in 2013, and was elected EURASIP Fellow in 2014.