Does deblurring improve geometrical hyperspectral unmixing?

Simon Henrot, Charles Soussen, Manuel Dossot, David Brie

► To cite this version:

Simon Henrot, Charles Soussen, Manuel Dossot, David Brie. Does deblurring improve geometrical hyperspectral unmixing?. IEEE Transactions on Image Processing, 2014, 23 (3), pp.1169-1180. 10.1109/TIP.2014.2300822 . hal-00933013

HAL Id: hal-00933013

<https://hal.science/hal-00933013>

Submitted on 30 Jan 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Does deblurring improve geometrical hyperspectral unmixing?

Simon Henrot, Charles Soussen, *Member, IEEE*, Manuel Dossot, and David Brie, *Member, IEEE*

Abstract

In this paper, we consider hyperspectral unmixing problems where the observed images are blurred during the acquisition process, *e.g.* in microscopy and spectroscopy. We derive a joint observation and mixing model and show how it affects endmember identifiability within the geometrical unmixing framework. An analysis of the model reveals that nonnegative blurring results in a contraction of both the minimum-volume enclosing and maximum-volume enclosed simplex. We demonstrate this contraction property in the case of a spectrally-invariant point-spread function. The benefit of prior deconvolution on the accuracy of the restored sources and abundances is illustrated using simulated and real Raman spectroscopic data.

Index Terms

Hyperspectral unmixing, Minimum-volume simplex, Deconvolution

I. INTRODUCTION

Multispectral imaging refers to the acquisition of two-dimensional (2D) spatial images at different wavelengths, yielding one spectral vector per pixel of the image. Hyperspectral imaging takes this idea one step further by improving the spectral resolution to hundreds or thousands of overlapping spectral bands. Hyperspectral data may then be viewed as a collection of highly resolved spectra. In many applications, the image contains a small number of pure materials - termed *endmembers* - whose spectral

Simon Henrot, Charles Soussen and David Brie are with the CRAN, UMR 7039, Université de Lorraine, CNRS, Faculté des Sciences, B.P. 70239, F-54506 Vandoeuvre-lès-Nancy, France. Manuel Dossot is with the LCPME, UMR 7564, Université de Lorraine, CNRS, 405, rue de Vandoeuvre, F-54602 Villers-ls-Nancy, France. e-mail: firstname.lastname@univ-lorraine.fr.

signatures are mixed in each pixel because of limited spatial resolution. Spectral unmixing usually refers to the estimation of endmembers and their fractional contribution to each pixel, named *abundances*.

Many unmixing methods in the literature can be categorized as belonging to either a geometrical, a statistical (Bayesian) or a sparse regression framework [4]. In the former approach, each spectral vector belongs to a simplex whose vertices are the endmembers we seek [9]. Some algorithms assume that there exists at least one *pure pixel* (PP) per endmember and exploit this property by identifying the endmembers to the “purest” pixels in the hyperspectral scene. When the *pure pixel hypothesis* does not hold, the geometrical approach to unmixing then consists in fitting the *Minimum Volume Simplex* (MVS) to the data using one of many MV algorithms. However, highly mixed data are beyond the reach of geometrical algorithms because spectral signatures are located near the center of the true endmember simplex. One could then resort to statistical or sparse methods to unmix the data set. In this paper, we focus instead on analyzing and alleviating the effects of hyperspectral image blurring, which will be shown to possibly cause highly mixed data.

Hyperspectral images are usually degraded by the observation process, typically including blurring and corruption by noise. Data acquired by spectro-imagers dedicated to geoscience, remote sensing or planetology, result from a physics-based inversion pipeline ensuring that these images are free of blur. However, in other domains, the acquisition process does not include a deblurring step, resulting in blurry hyperspectral images. For instance, microscopy [25], Raman spectroscopy [10] or ground-based astrophysical [5], [22] images are blurred by diffraction through the objective aperture. We consider such hyperspectral images in this paper. Mathematically, each observed image plane is obtained as the two-dimensional convolution of the original image and the channel linear shift-invariant blur or *Point-Spread Function* (PSF). If the size of the PSF is less than that of an image pixel, the blur does not alter the spatial resolution of the data. However, if the PSF support extends beyond a pixel, neighboring pixels are combined in the observed data, resulting in additional mixing.

The latter case arises whenever one aims at obtaining a fine spatial resolution, e.g. in super-resolution reconstruction [1]. In this framework, some scenes that could *a priori* be geometrically unmixed may lose this property because of the blur. The operation of *deconvolution* which inverts the observation process may then restore the stack so that low computational complexity geometrical methods qualify to unmix it. In addition, deconvolution algorithms also provide the additional benefit of restoring images with higher signal-to-noise ratios (SNRs).

The ideal approach to account for these two physical phenomena would be to jointly deconvolve and unmix the data. Let us mention [26] where the authors perform joint deconvolution and abundance

estimation. They assume that endmembers belong to a large dictionary and the abundance estimation is formulated as a sparse regression problem. However, to the best of our knowledge, no algorithms have been designed to perform joint deconvolution and *blind* unmixing¹ of the data. Nevertheless, many off-the-shelf algorithms are available to perform deconvolution and unmixing separately. We propose to show in this paper that the sequential procedure of (deconvolution/unmixing) yield better results than a (unmixing/deconvolution) strategy.

To this end, we will study how blurring may affect the results of geometrical unmixing methods and assess whether deblurring hyperspectral images before unmixing them yields quantitative improvements. The remainder of the paper is organized as follows. Section II presents the linear mixing model and recalls its associated geometrical framework ; we also discuss identifiability issues by the geometrical unmixing approach. In section III, we factor in the general observation model and present experimental results on synthetic data to evaluate the gain of a deconvolution step before applying a spectral unmixing algorithm. Section IV focuses on the specific case where the PSF is invariant across the spectral channels. Finally, we present experimental results on real data in section V and we conclude in section VI.

II. HYPERSPECTRAL GEOMETRICAL UNMIXING FRAMEWORK

In this section, we introduce concepts and notations used throughout the rest of the paper. Most of the results given below are not new and can be found elsewhere in the literature: see [4] for a thorough review of hyperspectral unmixing and [6] for an introduction to convex geometry applied to endmember extraction.

A. Linear mixing model

Consider an hyperspectral image measuring radiance on L different spectral bands (*channels*) and N pixels. We gather the data in a $L \times N$ matrix $\mathbf{X}$ and use the following notations:

- 1) $\mathbf{x}^\ell$ is the ℓ -th row of $\mathbf{x}$, that is the 2D image at channel ℓ after lexicographical ordering into a row vector of length N ;
- 2) $\mathbf{x}_n$ is the n -th column of $\mathbf{x}$, i.e. the $L \times 1$ spectrum of the n -th pixel (also termed spectral vector or pixel vector).

Each spectral vector in the image is a linear combination of an known number P of endmembers $\{\mathbf{s}_1, \dots, \mathbf{s}_P\}$. When unknown, P can be obtained by some model order estimation method such as

¹Here, blind unmixing refers to the joint estimation of endmembers and abundances.

virtual dimensionality [7]. Ignoring noise for now, the linear mixing model (LMM) writes

$$\mathbf{X} = \mathbf{S}\mathbf{A}$$

where the p th column of $L \times P$ source matrix $\mathbf{S}$ indexes endmember $\mathbf{s}_p$ and the n th column $\mathbf{a}_n$ of $P \times N$ abundance matrix $\mathbf{A}$ contains the fractional abundance coefficients for $\mathbf{x}_n$:

$$\mathbf{x}_n = \mathbf{S}\mathbf{a}_n = \sum_{p=1}^P A_{p,n} \mathbf{s}_p. \quad (1)$$

The LMM is generally based on the following assumptions [4]:

- i) the number of endmembers P is much smaller than the number of bands L , that is $P \ll L$;
- ii) matrix $\mathbf{S}$ is of full column rank, *i.e.* endmembers $\{\mathbf{s}_1, \dots, \mathbf{s}_P\}$ are linearly independent;
- iii) *Abundance Nonnegativity Constraint* (ANC) : $A_{p,n} \geq 0$ for all p and n ;
- iv) *Abundance Sum Constraint* (ASC) : $\sum_{p=1}^P A_{p,n} = 1$ for all n .

Assumptions (i) and (ii) seem very reasonable in hyperspectral imaging since many bands are collected and the image is made up of a few distinct materials. Assumptions (iii) and (iv) come from the physical interpretation of abundance coefficient $A_{p,n}$ as the fractional spatial area occupied by the p th endmember in the n th pixel.

To model illumination variability due to surface topography [21], equation (1) sometimes writes $\mathbf{x}_n = \gamma_n \mathbf{S}\mathbf{a}_n$ where parameter γ_n is a scale factor for each pixel in the hyperspectral scene. This phenomenon can be accounted for in model (1) by simply rejecting hypothesis (iv). The geometrical consequences of this choice are detailed below.

B. Dimension reduction and affine projection

From equation (1) and assumptions (i) and (ii), it is immediate that pixel vectors lie in a P -dimensional vector subspace of $\mathbb{R}^L$. Suppose the subspace orthonormal basis $\mathbf{E} = [\mathbf{e}_1, \dots, \mathbf{e}_P]$ has been inferred from the data, e.g. by the HySime algorithm [3], then the n -th dimension-reduced pixel is given by

$$\hat{\mathbf{x}}_n = \mathbf{E}^T \mathbf{x}_n \in \mathbb{R}^P. \quad (2)$$

From equation (2) and assumption (iii), P -dimensional pixel vectors are enclosed in the convex cone spanned by dimension-reduced endmembers $\{\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_P\}$ [17] (see figure 1):

$$\{\hat{\mathbf{x}} | \hat{\mathbf{x}} = \sum_{p=1}^P \theta_p \hat{\mathbf{s}}_p, \theta_p \geq 0, p = 1 \dots P\}. \quad (3)$$

Fig. 1. Illustration in dimension $P = 3$ of the convex cone spanned by dimension-reduced endmembers $\{\hat{\mathbf{s}}_p, p = 1 \dots P\}$ enclosing dimension-reduced pixels $\hat{\mathbf{x}}_n$'s. $\hat{\mathbf{x}}_n$ is obtained by perspective projection of $\hat{\mathbf{x}}_n$ on the affine plane supported by $\{\hat{\mathbf{s}}_p, p = 1 \dots P\}$.

In order to satisfy hypothesis (iv), the P -dimensional points are then projected onto a $(P - 1)$ -dimensional affine plane (see figure 1). The choice of a particular affine plane is arbitrary and may be viewed as imposing the norm of the supporting endmembers.

Most algorithms use either an orthogonal projection [6] or the perspective projection [4] onto the $(P - 1)$ -dimensional affine plane : the choice of one method or another is related to the SNR, see [21]. For example, the perspective projection consists in scaling pixel vectors

$$\hat{\mathbf{x}}_n = \frac{\hat{\mathbf{x}}_n}{\hat{\mathbf{x}}_n^T \mathbf{u}} \quad (4)$$

to project them onto the hyperplane defined by $\hat{\mathbf{v}}^T \mathbf{u} = 1$ for every vector $\hat{\mathbf{v}} \in \mathbb{R}^P$, where $\mathbf{u} \in \mathbb{R}_+^P$ is chosen beforehand such that $\hat{\mathbf{x}}_n^T \mathbf{u} > 0$ for all n . From equations (1), (2) and (4), the n -th projected pixel vector rewrites

$$\hat{\mathbf{x}}_n = \sum_{p=1}^P \alpha_{p,n} \hat{\mathbf{s}}_p \quad (5)$$

where coefficients $\alpha_{p,n}$ are given by

$$\alpha_{p,n} = A_{p,n} \frac{\hat{\mathbf{s}}_p^T \mathbf{u}}{\hat{\mathbf{x}}_n^T \mathbf{u}} = \frac{A_{p,n} \hat{\mathbf{s}}_p^T \mathbf{u}}{\sum_q A_{q,n} \hat{\mathbf{s}}_q^T \mathbf{u}}. \quad (6)$$

It is easy to see that the $\alpha_{p,n}$'s satisfy the following properties:

- 1) $\alpha_{p,n} \geq 0$ for all p and n , under the ANC;

2) $\sum_p \alpha_{p,n} = 1$ for all n .

It follows that each projected pixel vector $\hat{\mathbf{x}}_n$ is a convex combination of the projected endmembers $\{\hat{\mathbf{s}}_p, p = 1 \dots P\}$. Hence, projected pixel vectors belong to the convex set $\mathcal{S}$:

$$\mathcal{S} = \{\hat{\mathbf{x}} | \hat{\mathbf{x}} = \sum_{p=1}^P \theta_p \hat{\mathbf{s}}_p | \theta_p \geq 0, p = 1 \dots P, \sum_{p=1}^P \theta_p = 1\}. \quad (7)$$

By definition, $\mathcal{S}$ is the simplex whose vertices are the projected endmembers $\{\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_P\}$.

C. Minimum Volume Simplex (MVS)

Most geometrical unmixing algorithms are based on the key idea introduced by Craig [9] that the **minimum-volume simplex** (MVS) enclosing all the pixel vectors identifies with $\mathcal{S}$ and thus the endmembers we seek can be extracted as the vertices of the MVS. The MV constraint allows to regularize the ill-posed nature of the nonnegative source separation problem [20]. Alternatively, one can look for the **maximum-volume simplex** inscribed in the dataset, provided the image contains pure pixels (PP). Recovering both simplices can be addressed as a combinatorial problem but is too complex to handle high dimensional data [12], [27]. Simple and computationally efficient approaches were proposed e.g. [24], [21] in the MVS case and [9], [8], [6] in the PP case, among many others.

Fig. 2. 2D examples of triangles enclosing the data for $P = 3$ endmembers. The true unknown simplex $\mathcal{S}$ is represented with a straight line while the other candidate triangle is plotted with a dotted line. Data points are represented as hollow circles and the vertices of the simplices are filled circles. (a): Statement (1) is true, i.e. $\mathcal{S}$ is the MVS. (b): both triangles are MVS's for this data configuration. (c): the dotted-lined triangle has a smaller volume than $\mathcal{S}$ and thus is the MVS.

We point out that unlike in the PP case, the MVS and the *true* simplex $\mathcal{S}$ whose vertices are the endmembers may be different simplices. Figure 2 illustrates the distinction between the MVS and $\mathcal{S}$ in

the simple case where $P = 3$, yielding a 2D affine plane. In what follows, we discuss the problem of endmember identifiability using the MVS approach. Craig's statement can be rewritten in the following way:

Statement 1. *the MVS identifies with $\mathcal{S}$.*

The first problem with the statement is that there might exist more than one MVS: see figure 2b. This case is not likely to arise in practice if enough data points have been collected so that there exists data points near the vertices of $\mathcal{S}$.

Assuming there exists only one MVS, the following result can be found in the literature [4]:

Sufficient condition 1. *If there is at least one data point on each vertex of $\mathcal{S}$, then statement 1 is true.*

This sufficient condition (SC) is known as the *pure pixel condition*. For $P = 3$, we make the following statement whose proof is given in appendix A (see figure 2c):

Sufficient condition 2. *If there are at least two data points per facet of $\mathcal{S}$ whose distance with the closest vertex is no greater than a third of the corresponding facet's length, then statement 1 is true.*

SC 2 is weaker than SC 1. We note that a recent paper establishes necessary and sufficient conditions for the identifiability of endmembers by Craig's criterion in the 3-endmember case, see [18]. Formulating such a statement for higher dimensions (that is, for more than three endmembers) is difficult and out of the scope of this paper. Rather, we stress that endmembers are more likely to be accurately identified by a MVS algorithm if there are pixel vectors near the vicinity of the vertices of $\mathcal{S}$ in the data. Obviously, if the data is highly mixed, the MVS will be enclosed in $\mathcal{S}$ and identifiability conditions will not be satisfied. Such a situation can arise because of two major reasons:

- 1) the abundance distributions themselves;
- 2) additional mixing due to the imaging system, as we will demonstrate in the next section.

In the latter case, this undesirable contribution can be removed by image restoration methods, thus improving endmember identifiability conditions.

III. EFFECT OF THE IMAGING PROCESS

A. Blurring of the linear mixture

In this section we now account for the fact that the image is degraded during the acquisition process. Under the common linear blur assumption, the 2D image $\mathbf{y}^\ell$ observed at a given channel ℓ is obtained

as the 2D convolution product of the true image and the channel point-spread function $\mathcal{H}^\ell$:

$$\mathbf{y}^\ell = \mathbf{x}^\ell \mathbf{H}^\ell \quad (8)$$

where the $N \times N$ observation matrix $\mathbf{H}^\ell$ is a convolution matrix corresponding to $\mathcal{H}^\ell$. For instance, when the blur is space-invariant, $(\mathbf{H}^\ell)^T$ is a block-Toeplitz matrix where each block is Toeplitz [14]². Each entry of the observed data matrix $\mathbf{Y}$ is given by

$$Y_{\ell,k} = \sum_{n=1}^N X_{\ell,n} H_{n,k}^\ell. \quad (9)$$

Using equations (1) and (9), the overall model combining noise, observation blurring and linear mixing of endmembers writes

$$Y_{\ell,k} = \sum_{n=1}^N \sum_{p=1}^P S_{\ell,p} A_{p,n} H_{n,k}^\ell + W_{\ell,k} \quad (10)$$

where $\mathbf{W}$ is the noise term and model (10) assumes that the SNR is high enough for the noise to be additive and i.i.d. Gaussian. We observe that the blurred data do not satisfy the LMM since the mixing coefficients $A_{p,n} H_{n,k}^\ell$ depend on the channel index ℓ . Thus, even in the absence of noise, **the $\mathbf{y}_k$'s do not necessarily belong to $\mathcal{S}$ anymore.**

B. Pixel trajectories

Since the abundances vary with the channel index, it is no longer possible to use barycentric coordinates to represent a pixel vector in the simplex formed by the endmembers [15]. To give some insight into the understanding of the effect of the channel dependent blur on the linear mixture, we propose to use the representation of figure 4. The endmembers correspond to the vertices of the simplex. Note that in this figure the simplices locations are arbitrarily fixed to $(0, 0)$, $(0, 1)$ and $(1/2, \sqrt{3}/2)$ resulting in an isocoles triangle. As the mixing coefficients are depending on the channel index, a mixed pixel can no longer be represented as a single point. For each channel, we have different mixing coefficients. We rather represent a channel dependent mixed pixel by a trajectory where each point of the trajectory corresponds to the barycentric coordinates at a particular channel.

Figure 3 displays a few pixel trajectories for a Gaussian PSF whose width linearly increases with the channel index³. Interestingly, this representation provides spatial information about the pixel at hand. For

²The transpose comes from the fact that $\mathbf{x}^\ell$ and $\mathbf{y}^\ell$ are row vectors.

³We restrict our analysis to the case of PSFs whose width increases with the channel index (i.e. wavelength) because this is the relevant case of conventional optical systems.

Fig. 3. Examples of trajectories of a few pixel vectors.

instance, a trajectory that goes from a vertex to another may indicate that the pixel is located in the boundary of two regions, each corresponding to a particular endmember; a trajectory pointing towards the center of the simplex corresponds to a highly mixed pixel.

As a consequence, even if applying a linear unmixing algorithm directly to the data is feasible, the resulting endmembers and abundances are literally unmeaningful. We rather propose to perform deblurring prior to the unmixing step and numerical experiments are conducted in the next section to assess the gain of doing so.

To conclude this section, we note that blurring caused by a nonnegative PSF results in a *contraction property* whose extent increases with the PSF width. We will come back to that point in section IV.

C. Experimental results

To quantitatively assess the effects of deblurring on hyperspectral unmixing, we use synthetic data so that the extracted endmembers may be compared to their true counterparts. First, we select $P = 3$ true independent endmembers from the United States Geological Survey (USGS) spectral library⁴. Each endmember signature is sampled on 224 spectral bands. We then synthesize a set of P random two-dimensional (2D) 30×30 abundance maps modeled as a mixture of 2D Gaussian functions with random locations and widths, under the constraint that the abundances for each pixel sum to one. The true data stack is then computed using equation (1).

⁴Available online from: <http://speclab.cr.usgs.gov/spectral-lib.html>

Fig. 4. Selected endmembers.

Fig. 5. Examples of abundance maps generated as spatial mixtures of 2D Gaussian distributions with random means and variances.

Image degradation is simulated by convolving each channel with a 2D Gaussian PSF whose width varies linearly with the channel index, from one pixel to the entire field of view. The approximation of physical PSFs by 2D Gaussian functions is actually reasonable for numerous optical systems, including fluorescence microscopy [25], Raman microscopy [10] and atmospheric turbulence [26]. We then add i.i.d. Gaussian noise at a specific SNR value. We wish to evaluate the effect of deconvolution for different noise levels. Hence, we run three separate experiments with respective SNRs of 5dB, 25dB and 50dB. The rationale for these values is that at 5dB, the role of the deconvolution algorithm is to both denoise and deblur the data, while at 50 dB its main role will consist in deblurring. A SNR of 25dB is viewed as a more standard situation.

Image deconvolution is achieved using the algorithm described in [14] which accounts for spectral and spatial smoothness and non-negativity of multispectral and hyperspectral data. The rationale for this choice is that the priors of the methods are well adapted to the data and its low complexity allows

Fig. 6. Original data simplex (with no blur and noise) obtained by Principal Component Analysis and orthogonal projection. Pixel vectors are denoted by the symbol '+', endmembers by circles.

a computationally efficient simulation. This algorithm requires tuning two regularization parameters, weighting respectively a spatial and spectral smoothness prior, denoted by μ and ν . For each situation, we run the algorithm by setting these parameters on a 10×10 logarithmic grid, with values ranging from 10^{-4} to 10^2 . Since we only evaluate hyperparameters on a finite grid, deconvolution of the data may not be optimal. However, the deconvolved image is always closer to the true image in the mean square error sense.

For each experiment, spectral unmixing is carried out on both the degraded data set and the deconvolved data set. Endmember extraction is performed using three widespread geometrical-based algorithms : *Vertex Component Analysis* (VCA) [21], *Minimum-Volume Enclosing Simplex* (MVES) [6] and *Simplex Identification via variable Splitting and Augmented Lagrangian* (SISAL) [2], using MATLAB code published by the authors. Abundance estimation is automatically carried out with MVES and we use Fully Constrained Least Squares inversion (FCLS) [13] to compute abundances for VCA and SISAL. While VCA is based on the pure pixel hypothesis, both MVES and SISAL are minimum-volume algorithms, hence the choice of these three methods allows the scrutiny of a wide range of situations. The goal here is to evaluate the increase of performance brought by a deconvolution step for all algorithms, rather than comparing the results of one algorithm to another, which would not be fair as they depend on the

SNR	Endmember extraction SAD						Abundance estimation MSE					
	VCA		MVES		SISAL		VCA		MVES		SISAL	
Data	Raw	Restored	Raw	Restored	Raw	Restored	Raw	Restored	Raw	Restored	Raw	Restored
5	0.8840	0.4344	2.2361	0.5206	0.8439	0.4936	0.0032	0.0029	0.0062	0.0031	0.0035	0.0027
25	0.8951	0.2780	0.6640	0.83	0.7856	0.4385	0.0046	0.0085	0.0037	0.0041	0.0035	0.0027
50	0.8926	0.2915	3.0823	0.4330	0.8527	0.4176	0.0046	0.0082	0.0062	0.0018	0.0035	0.0017

TABLE I
ENDMEMBER EXTRACTION AND ABUNDANCE ESTIMATION ERRORS FOR A VARIANT PSF.

situation at hand.

We choose the following performance metrics:

- 1) the *spectral angle distance* (SAD) between the true and extracted endmembers:

$$\text{SAD}_p = \cos^{-1} \left(\frac{\mathbf{s}_p^T \hat{\mathbf{s}}_p}{\|\mathbf{s}_p\|_2 \|\hat{\mathbf{s}}_p\|_2} \right), p = 1 \dots P. \quad (11)$$

The total SAD is given by $\text{SAD} = \sum_p \text{SAD}_p$:

- 2) The *Mean Square Error* (MSE) between the true and estimated abundances, given by

$$\text{MSE} = \|\hat{\mathbf{A}} - \mathbf{A}\|_F^2. \quad (12)$$

We then store the unmixing results on the raw data for each point of the hyperparameter grid and select the point which corresponds to the best extraction SAD. This choice is arbitrary : depending on the application, it may be better to select the deconvolution result yielding the smallest abundance estimation MSE. In either case, we point out that both these quantities could provide interesting criteria for hyperparameter selection.

Table I gives the various performances obtained for endmember extraction and abundance estimation by VCA, MVES and SISAL with and without a prior deconvolution step. Deconvolving the image prior to image unmixing improves the results for all algorithms.

IV. CASE OF PSF INVARIANCE

In this section, we focus on the specific case where the PSF is invariant across channels. This case is actually important in practice and arises in many imaging situations, such as Raman microscopy.

Fig. 7. Left : original projected pixel vectors, right : degraded projected pixel vectors. The data cloud is shrunk by the observation process and pure pixels become mixed in the observed data.

A. Bilinear model

In the specific case where the PSF is invariant across channels, the model reduces to the bilinear equation $\mathbf{Y} = \mathbf{SAH} + \mathbf{W}$ where each dimension-reduced pixel is given by

$$\hat{\mathbf{y}}_k = \sum_{n=1}^N \sum_{p=1}^P A_{p,n} H_{n,k} \hat{\mathbf{s}}_p + \hat{\mathbf{w}}_k. \quad (13)$$

The k -th projected pixel vector $\hat{\mathbf{y}}_k = \hat{\mathbf{y}}_k / \hat{\mathbf{y}}_k^T \mathbf{u}$ is then a noisy bilinear combination of the projected endmembers $\{\hat{\mathbf{s}}_p = \hat{\mathbf{s}}_p / \hat{\mathbf{s}}_p^T \mathbf{u}, p = 1 \dots P\}$:

$$\hat{\mathbf{y}}_k = \sum_{n=1}^N \sum_{p=1}^P \alpha'_{p,n,k} \hat{\mathbf{s}}_p + \frac{\hat{\mathbf{w}}_k^T \mathbf{u}}{\hat{\mathbf{y}}_k^T \mathbf{u}} \hat{\mathbf{w}}_k. \quad (14)$$

where coefficients $\alpha'_{p,n,k}$ are given by

$$\alpha'_{p,n,k} = A_{p,n} H_{n,k} \frac{\hat{\mathbf{s}}_p^T \mathbf{u}}{\hat{\mathbf{y}}_k^T \mathbf{u}} = \frac{A_{p,n} H_{n,k} \hat{\mathbf{s}}_p^T \mathbf{u}}{\sum_{q,m} A_{q,m} H_{m,k} \hat{\mathbf{s}}_q^T \mathbf{u}}. \quad (15)$$

From equation (15), we can see that coefficients $\alpha'_{p,n,k}$ satisfy the nonnegativity and sum-to-one properties:

- 1) $\alpha'_{p,n,k} \geq 0$ for all p, n and k under the ANC and knowing that PSFs entries are nonnegative;
- 2) $\sum_{p,n} \alpha'_{p,n,k} = 1$ for all $k = 1 \dots N$.

Thus, in the case where the PSF is invariant across channels and in the absence of noise, the observed pixel vectors $\{\hat{\mathbf{y}}_k, k = 1 \dots N\}$ remain in the simplex $\mathcal{S}$ (defined in equation (7)). However, the additional mixing introduced by the PSF may hinder spectral unmixing in a different way, as shown in section IV-B.

B. MVS contraction

How does the observation process affect the distribution of pixel vectors inside the simplex? The answer to the question obviously depends on the nature of the PSF. Since the entries of $\mathbf{H}$ are known to be nonnegative, the blurring process tends to average neighboring pixel intensities. Hence:

- 1) suppose the data include a pure pixel $\mathbf{x}_n$: then $\hat{\mathbf{x}}_n$ locates on a vertex of $\mathcal{S}$. If a mixed pixel $\mathbf{x}_m$ is in the spatial neighborhood of $\mathbf{x}_n$, the observed $\mathbf{y}_n$ will appear mixed in the observed image, which will relocate $\hat{\mathbf{y}}_n$ towards the inside of $\mathcal{S}$;
- 2) similarly, if a source $\mathbf{s}_p$ is not present in a given pixel $\mathbf{x}_n$ but appears in its neighborhood, the observed pixel vector $\mathbf{y}_n$ will include a contribution from $\mathbf{s}_p$ and $\hat{\mathbf{y}}_n$ will move towards the inside of $\mathcal{S}$.

This phenomenon inclines observed spectral vectors to cluster towards the center of $\mathcal{S}$. As a graphical illustration, consider for a simple multispectral image with $3 \times 3 = 9$ pixels and 3 spectral bands, each blurred with a 3×3 averaging kernel:

$$\mathcal{H} = \begin{pmatrix} 1/9 & 1/9 & 1/9 \\ 1/9 & 1/9 & 1/9 \\ 1/9 & 1/9 & 1/9 \end{pmatrix}. \quad (16)$$

Suppose our data set satisfies the pure pixel hypothesis, and also includes some mixed pixels, yielding the following 3×9 abundance matrix:

$$\mathbf{A} = \begin{pmatrix} 1 & 1/2 & 0 & 1/2 & 1/3 & 1/6 & 0 & 1/6 & 0 \\ 0 & 0 & 0 & 1/2 & 1/3 & 1/6 & 1 & 2/3 & 1/2 \\ 0 & 1/2 & 1 & 0 & 1/3 & 2/3 & 0 & 1/6 & 1/2 \end{pmatrix}. \quad (17)$$

The resulting simplices for the original and degraded data are given in figure 7. The data cloud clearly suffers a shrinkage inside of $\mathcal{S}$ from the observation process. Let us now formally state this contraction property, using the notion of *convex hull*. The convex hull of the data should not be confused with the MVS which necessarily has P vertices; in contrast, the convex hull has up to N vertices which all are pixels.

Proposition 1. *Let $\mathcal{S}_{\mathbf{x}}^*$ denote the MVS enclosing the true data $\mathbf{X}$, $\mathcal{S}_{\mathbf{y}}^*$ the one enclosing the blurred data $\mathbf{Y}$ and $\mathcal{V}$ denote the volume. If the PSF is nonnegative, then the convex hull of $\hat{\mathbf{Y}}$ is enclosed in the convex hull of $\hat{\mathbf{X}}$ and*

$$\mathcal{V}(\mathcal{S}_{\mathbf{y}}^*) \leq \mathcal{V}(\mathcal{S}_{\mathbf{x}}^*) \quad (18)$$

Proof: Assume the PSF coefficients are nonnegative. Then each observed pixel $\hat{\mathbf{y}}_k$ is a linear nonnegative combination of the original pixels $\{\hat{\mathbf{x}}_n, n = 1 \dots N\}$ (or possibly a subset of $\hat{\mathbf{X}}$ depending on the support of the PSF). Hence $\hat{\mathbf{y}}_k$ belongs to the conical hull of $\hat{\mathbf{X}}$ for each k . It follows that the convex hull of the projected observed pixel vectors $\{\hat{\mathbf{y}}_k, k = 1 \dots N\}$ is included in the convex hull of $\{\hat{\mathbf{x}}_n, n = 1 \dots N\}$.

Therefore $\mathcal{S}_{\mathbf{x}}^*$ encloses all pixels $\{\hat{\mathbf{y}}_k, k = 1 \dots N\}$. Since $\mathcal{S}_{\mathbf{y}}^*$ is the MVS enclosing $\{\hat{\mathbf{y}}_k, k = 1 \dots N\}$, we necessarily have

$$\mathcal{V}(\mathcal{S}_{\mathbf{y}}^*) \leq \mathcal{V}(\mathcal{S}_{\mathbf{x}}^*) \quad (19)$$

(Note however that while the volume of $\mathcal{S}_{\mathbf{y}}^*$ is smaller, $\mathcal{S}_{\mathbf{y}}^*$ is not necessarily included in $\mathcal{S}_{\mathbf{x}}^*$.) ■

In the case of Gaussian PSFs, Proposition 1 has an interesting consequence for the variant PSF case.

Proposition 2. *Consider two symmetric separable 2D Gaussian PSFs $\mathcal{H}_1$ and $\mathcal{H}_2$ of respective width satisfying $\sigma_1 > \sigma_2$. Let $\{\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2\}$ denote the respective observed projected pixels and $\{\mathcal{S}_1^*, \mathcal{S}_2^*\}$ the respective MVSs. Then the convex hull of $\hat{\mathbf{y}}_1$ is included in the convex hull of $\hat{\mathbf{y}}_2$ and $\mathcal{V}(\mathcal{S}_1^*) \leq \mathcal{V}(\mathcal{S}_2^*)$.*

Proof: The PSFs being separable, $\mathcal{H}_1$ may be written as the 2D convolution product of $\mathcal{H}_2$ and a Gaussian kernel of standard deviation $\sqrt{\sigma_1^2 - \sigma_2^2}$. Since all kernels are nonnegative, Proposition 2 is a straightforward application of Proposition 1. ■

Getting back to the idea of pixel trajectories outlined in section (III-B), when the PSF width increases with the channel, proposition (2) allows us to deduce that the convex hull of points at a given channel ℓ is included in the convex hull of points at channel $\ell - 1$. We now resume the discussion of the invariant PSF case.

Hence, we know from section II-C that directly applying a geometrical unmixing algorithm to the observed data may produce incorrect endmembers, and thus the subsequent estimation of abundances will also be biased. In the next section, we study the improvement of unmixing performance after applying a deconvolution algorithm to the data.

C. Experimental results

Here, we consider the effect of the standard deviation of the PSF on our performance metric, using the same synthetic data and experimental protocol as in section III-C. For a fixed SNR value, we consider each kernel size from one pixel ($\sigma = 0.1$) to the entire field of view of the image ($\sigma = 5$). For each value of σ , we run the deconvolution algorithm for spatial and spectral regularization parameters on the previously used grid. We display our experimental results on figure 9 and note a few key points:

Fig. 8. Illustration of the nesting property of convex hulls from channel to channel (variant PSF case).

- 1) As expected, the performance of both endmember extraction and abundance estimation are better after a deconvolution step, the margin between the two results tending to increase with σ except for a few localized values of σ . The reason why the unmixing algorithms perform slightly better on the raw data for these values certainly owes to the limited tuning of hyperparameters in the deconvolution step.
- 2) All unmixing results show a decrease in performance, both in terms of endmember extraction (increase of SAD) and abundance estimation (increase of MSE) when σ increases. Owing to the random initializations of the unmixing algorithms, the error curves are not purely monotonic. The decrease in performance seems very intuitive since the severity of the observed data degradation increases with σ .
- 3) The parameter search on the deconvolution step allows to stabilize the solution. Of course, the simulation makes use of ground truth while computing the best parameter values. In order to perform tuning in real imaging situations, one could use a library of spectral signatures and seek the parameters minimizing the spectral angle distance between the extracted endmembers and the corresponding ground truth endmembers that are a priori known to be present in the image.

Figure 10 is a geometrical view of the results of figure 9 illustrating the distance between the unknown and recovered sources, both on the raw and restored data for three values of the PSF spread σ . The results clearly show the MVS contraction caused by the blurring process and the reverse effect of the deconvolution step.

Let us sum up our experimental results in the invariant PSF case. A prior deconvolution step increases the performance of endmember extraction and abundance estimation for a vast majority of σ values. Of course, the wide variety of situations (abundance maps, endmembers, PSF spread, SNR, ...) make general statements difficult to quantify. However, in all our experiments, the unmixing results after deconvolution were always at least comparable and most of the time much better than the results obtained on the raw data.

V. EXPERIMENTAL RESULTS ON REAL DATA

In this section, we evaluate the gain of deconvolution on unmixing real Raman spectroscopy data. Our hyperspectral data set comprises images of size 98×131 pixels, each pixel being $100 \text{ nm} \times 100 \text{ nm}$, acquired on 337 bands ranging from 800 cm^{-1} to 1200 cm^{-1} . The scene of interest consists in a grain of *sodium acetate* (CH_3COONa) covered with *sodium carbonate* (Na_2CO_3) laying on a *silicon* layer (Si). Part of the sodium carbonate reacts with water vapor to yield hydrated sodium carbonate. These four chemical compounds are the endmembers we seek. A thorough inspection of the data reveals that the silicon compound contributes to all pixels of the image. The extraction of these endmembers is a challenging problem, since silicon is the only compound for which the pure pixel assumption is fulfilled. Because of the inherent high mixing of the data, even MV methods are not supposed to produce good results on this data set.

Given the limited spectral range, the PSF is considered to be invariant across channels. It is modeled as a 2D Gaussian function [10] with an experimentally measured full-width at half maximum of 300 nm. We apply our deconvolution algorithm to the data by setting regularization parameters through a trial-and-error process to $\mu = 20$ and $\nu = 5$.

VCA and FCLS are carried out on both the raw data matrix and its restored counterpart ; the results are displayed in figure 11. The first extracted endmember corresponds to the silicon layer, which presents a broad band at $910 - 960 \text{ cm}^{-1}$ due to the 2TO harmonic phonon mode of bulk silicon. The deconvolution step allows to denoise its abundance map and more importantly, uncovers structure that was distributed throughout other abundance maps. The second endmember is sodium acetate, with a peak at 930 cm^{-1} due to the intense C-C stretching mode of the acetate molecule [23]. It is well isolated by VCA (both for the raw and restored data), perhaps because the width of the grain make this endmember almost uniquely contribute to some pixels. However, the third and fourth endmembers present mixed spectral contributions from all compounds : both forms of sodium carbonate, hydrated and non hydrated (respective peaks at 1060 and 1080 cm^{-1}) as well as silicon and sodium acetate. The deconvolution step again reveals

structure in the fourth endmember that was not displayed in the case of raw data.

Further analysis of the spectra also show bands at 840 cm^{-1} and 1004 cm^{-1} , which reveals that the data have accidentally been polluted with sodium benzoate ($\text{NaC}_7\text{H}_5\text{O}_2$) [11]. Because this compound is associated with the sodium carbonate phase, experiments with 5 endmembers do not yield improvements, both compounds being mixed in a unique endmember.

Both MVES and SISAL extract endmembers with negative peaks from this data set, rendering them physically unmeaningful, so we do not include results from these algorithms. However, we employ another minimum-volume algorithm : the *Nonnegative Matrix Factorization* (NMF) with volume constraints (*NMF-vol*) on the inferred simplex [19]. The algorithm operates on both the raw data (where negative pixels have been clipped to zero since the algorithm imposes a nonnegative data matrix) and restored data. Due to computation times, it was not possible to include this method in our previous simulations on synthetic data; here we set 10000 as the maximum number of iterations. The resulting abundance maps and endmembers are also given in figure 11. The spectral shapes of endmembers extracted from the raw data appear noisier, a problem solved by the deconvolution step. The sodium acetate endmember displays more undesirable contribution from the silicon compound. However, the algorithm is able to separate the sodium carbonate compound (third endmember, characterized by the peak at 1080 cm^{-1}) from the hydrated sodium carbonate compound (fourth endmember, peak at 1060 cm^{-1}) ; both are mixed with the sodium benzoate, as expected. The main gain of deconvolution clearly appears on the third endmember, where the silicon contribution is completely suppressed. Another benefit of the restoration step is to reveal structure hidden in raw abundance maps (first and third endmember).

VI. CONCLUSION

In this paper, we have analyzed the effect of a linear observation system on geometrical hyperspectral unmixing. When the observation system varies with the spectral channels, the joint mixing and observation model is no longer bilinear. Blurring by a nonnegative PSF results in the contraction of the minimum volume simplex enclosing the data. We formally prove the contraction property in the case of a spectrally constant PSF and show that its extent increases with the PSF width. Hence, unmixing blurred hyperspectral data is expected to significantly benefit from a deconvolution procedure, which we demonstrate in terms of endmember extraction and abundance estimations using both synthetic and real Raman spectroscopic data. Comparing the benefits and computational costs of a sequential or joint deconvolution and unmixing approach remains an open problem, especially when one needs to resort to blind deconvolution methods. Perspectives also include identifying application domains where the issue might have been overlooked, e.g.

microscopy and ground-based astrophysical imaging, and addressing nonlinear unmixing schemes [16].

APPENDIX A

SC 2 PROOF

Suppose an hyperspectral image contains three endmembers. After dimension reduction and affine projection (see section (II)), pixel vectors are enclosed in a triangle $\mathcal{S}$ whose vertices correspond to the endmembers. Suppose further that there are exactly two data points per facet of $\mathcal{S}$ and denote these six points by $\{P_1 \dots P_6\}$. Each point P_i is assumed to be located at a distance $r * d$ from the closest vertex on its facet, where d is the facet length and $r \in [0, 1/2]$. As illustrated in section (II-C), there are two possible triangles fitting $\{P_1 \dots P_6\}$, $\mathcal{S}$ and $\mathcal{S}'$.

We first perform an affine transformation so that $\mathcal{S}$ is transformed into a right isosceles triangle. The respective areas $\mathcal{A}(\mathcal{S})$ and $\mathcal{A}(\mathcal{S}')$ are scaled by the same number (that is, the determinant of the transform matrix) and the distance ratio r is conserved since the transformation is affine. In the new domain, we have $\mathcal{A}(\mathcal{S}) = 1/2$. By first calculating the equations of the facets of $\mathcal{S}'$ and the coordinates of the vertices of $\mathcal{S}'$, it is easy to compute $\mathcal{A}(\mathcal{S}')$ as a function of r :

$$\mathcal{A}(\mathcal{S}') = \left| \frac{9}{2}r^2 - 6r + 2 \right| \quad (20)$$

Solving $\mathcal{A}(\mathcal{S}') > \mathcal{A}(\mathcal{S})$ for $r \in [0, 1/2]$ is trivial and yields $r < 1/3$.

ACKNOWLEDGMENT

This work was supported by the French Agence Nationale de la Recherche, with reference number ANR-09-BLAN-0336-04. The authors are very indebted to the anonymous reviewers for their insightful comments which significantly improved the contents of the paper.

REFERENCES

- [1] T. Akgun, Y. Altunbasak, and R.M. Mersereau. Super-resolution reconstruction of hyperspectral images. *IEEE Transactions on Image Processing*, 14(11):1860–1875, 2005.
- [2] J.M. Bioucas-Dias. A variable splitting augmented lagrangian approach to linear spectral unmixing. In *WHISPERS: First Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing*, pages 1–4, 2009.
- [3] J.M. Bioucas-Dias and J.M.P. Nascimento. Hyperspectral subspace identification. *IEEE Transactions on Geoscience and Remote Sensing*, 46(8):2435–2445, aug. 2008.
- [4] J.M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Qian Du, P. Gader, and J. Chanussot. Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(2):354–379, april 2012.

- [5] S. Bourguignon, D. Mary, and E. Slezak. Processing MUSE hyperspectral data: Denoising, deconvolution and detection of astrophysical sources. *Statistical Methodology*, 9:32–43, Jan.–Mar. 2012.
- [6] 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 Transactions on Signal Processing*, 57(11):4418–4432, nov. 2009.
- [7] C.-I. Chang and Q. Du. Estimation of number of spectrally distinct signal sources in hyperspectral imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 42(3):608–619, march 2004.
- [8] C.-I. Chang, C.-C. Wu, W.-M. Liu, and Y.-C. Ouyang. A new growing method for simplex-based endmember extraction algorithm. *IEEE Transactions on Geoscience and Remote Sensing*, 44(10):2804–2819, 2006.
- [9] M.D. Craig. Minimum-volume transforms for remotely sensed data. *IEEE Transactions on Geoscience and Remote Sensing*, 32(3):542–552, may 1994.
- [10] C.J. De Grauw, N.M. Sijtsema, C. Otto, and J. Greve. Axial resolution of confocal raman microscopes: Gaussian beam theory and practice. *Journal of Microscopy*, 188(3):273–279, 1997.
- [11] Marleen De Veij, Peter Vandenabeele, Thomas De Beer, Jean Paul Remon, and Luc Moens. Reference database of raman spectra of pharmaceutical excipients. *Journal of Raman Spectroscopy*, 40(3):297–307, 2009.
- [12] Peter Gritzmann and Victor Klee. On the complexity of some basic problems in computational convexity: I. containment problems. *Discrete Mathematics*, 136(1):129–174, 1994.
- [13] D.C. Heinz and C.-I. Chang. Fully constrained least squares linear spectral mixture analysis method for material quantification in hyperspectral imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 39(3):529–545, march 2001.
- [14] S. Henrot, C. Soussen, and D. Brie. Fast positive deconvolution of hyperspectral images. *IEEE Transactions on Image Processing*, 22(2):828–833, Feb. 2013.
- [15] P. Honeine and C. Richard. Geometric unmixing of large hyperspectral images: A barycentric coordinate approach. *IEEE Transactions on Geoscience and Remote Sensing*, 50(6):2185–2195, 2012.
- [16] P. Honeine, C. Richard, and N. H. Nguyen. Approches géométriques pour l'estimation des fractions d'abondance en traitement de données hyperspectrales extensions aux modèles de mélange non-linéaires. *Traitement du signal*, 30(1-2):61–86, 2013.
- [17] A. Ifarraguerri and C.-I. Chang. Multispectral and hyperspectral image analysis with convex cones. *IEEE Transactions on Geoscience and Remote Sensing*, 37(2):756–770, 1999.
- [18] C.-H. Lin, A. Ambikapathi, W.-C. Li, and C.-Y. Chi. On the endmember identifiability of craigs criterion for hyperspectral unmixing: A statistical analysis for three-source case. In *Proc. IEEE ICASSP, Vancouver, Canada*, 2013.
- [19] Lidian Miao and Hairong Qi. Endmember extraction from highly mixed data using minimum volume constrained nonnegative matrix factorization. *IEEE Transactions on Geoscience and Remote Sensing*, 45(3):765–777, march 2007.
- [20] S. Moussaoui, D. Brie, and J. Idier. Non-negative source separation: range of admissible solutions and conditions for the uniqueness of the solution. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) 2005.*, volume 5, pages v–289. IEEE, 2005.
- [21] J.M.P. Nascimento and J.M. Bioucas-Dias. Vertex component analysis: a fast algorithm to unmix hyperspectral data. *IEEE Transactions on Geoscience and Remote Sensing*, 43(4):898–910, april 2005.
- [22] E. Villeneuve, H. Carfantan, and D. Serre. PSF estimation of hyperspectral data acquisition system for ground-based astrophysical observations. In *Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS), 2011 3rd Workshop on*, pages 1–4. IEEE, 2011.

- [23] Liang-Yu Wang, Yun-Hong Zhang, and Li-Jun Zhao. Raman spectroscopic studies on single supersaturated droplets of sodium and magnesium acetate. *The Journal of Physical Chemistry A*, 109(4):609–614, 2005.
- [24] M. E Winter. N-findr: an algorithm for fast autonomous spectral end-member determination in hyperspectral data. In *SPIE's International Symposium on Optical Science, Engineering, and Instrumentation*, pages 266–275. International Society for Optics and Photonics, 1999.
- [25] B. Zhang, J. Zerubia, and J.-C. Olivo-Marin. Gaussian approximations of fluorescence microscope point-spread function models. *Applied Optics*, 46:1819–1829, Mar. 2007.
- [26] X.-L. Zhao, F. Wang, T.-Z. Huang, M. K. Ng, and R. J. Plemmons. Deblurring and sparse unmixing for hyperspectral images. *IEEE Transactions on Geoscience and Remote Sensing*, PP(99):1–14, 2013.
- [27] Yunhong Zhou and Subhash Suri. Algorithms for a minimum volume enclosing simplex in three dimensions. *SIAM Journal on Computing*, 31(5):1339–1357, 2002.

Fig. 9. Performance of unmixing methods as a function of the PSF spread σ . First column : VCA, second : MVES, third : SISAL. For each SNR section, the first row displays the spectral angle divergence between the observed / restored extracted endmembers and their true counterparts (respectively solid and dotted line). The second row displays the mean square error between the observed / restored estimated abundances and their true counterparts.

Fig. 10. Observed and SISAL-restored simplices for three values of σ and a SNR of 25dB. $\sigma = .1$ corresponds to no blurring (the PSF is only one pixel) while the PSF covers the entire field of view for $\sigma = 5$. Pixel vectors are denoted by the symbol '+' ; true endmembers by circles ; SISAL-estimated endmembers by triangles. The MVS contraction is increasingly severe when σ increases.

January 30, 2014

Fig. 11. Spectral unmixing results on the real Raman spectroscopy data set.

Fig. 12. $\mathcal{S}$ is the straight-lined triangle, $\mathcal{S}'$ is the dotted-line triangle. $r = .2$ in this illustration.