Endmembers as Directional Data for Robust Material Variability Retrieval in Hyperspectral Image Unmixing

Lucas Drumetz, Jocelyn Chanussot, A Iwasaki

ENDMEMBERS AS DIRECTIONAL DATA FOR ROBUST MATERIAL VARIABILITY RETRIEVAL IN HYPERSPECTRAL IMAGE UNMIXING

L. Drumetz[a], J. Chanussot[b], A. Iwasaki[c]

[a] IMT Atlantique, Lab-STICC, UBL, 29238 Brest, France
[b] Univ. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, Grenoble, France
[c] Research Center for Advanced Science and Technology, The University of Tokyo, Tokyo, Japan

ABSTRACT

Hyperspectral image unmixing is a source separation problem aiming at recovering the spectra of the pure materials of the observed scene (called endmembers), as well as their relative proportions in each pixel of the image (called abundances). The variability of the materials has recently received a lot of attention in the community. In particular, a consequent number of models and algorithms have been proposed to estimate pixel-wise endmembers to account for their variability. These algorithms often rely on classical endmember extraction algorithms to provide reference spectra. In difficult scenarios with shadows and significant variability these algorithms may fail. In this paper, we address this issue in the Extended Linear Mixing Model framework by considering that an endmember is a direction in the feature space, rather than a single point. Under this paradigm, we show that using k-means clustering with the cosine similarity outperforms geometric endmember extraction algorithms. We also design an algorithm to refine the estimation of the endmember directions, and to account for both illumination and intrinsic variability effects. We show the potential of the proposed algorithm on a synthetic dataset using real world spectra with variability, and a challenging real dataset of a natural scene.

Index Terms— Hyperspectral unmixing, endmember variability, Extended Linear Mixing Model, oblique manifold, convex optimization

1. INTRODUCTION

Hyperspectral remote sensing enables a fine automatic identification of the materials present in the observed scene, thanks to the fine spectral resolution of hyperspectral images [1]. Their spatial resolution is, however, limited, and the field of view a single pixel often comprises several materials of interest. The observed spectrum is then a mixture of the contributions of the different materials at this location. Spectral unmixing is a blind source separation problem whose goal is to retrieve the spectra of the pure materials (called endmembers), and to estimate their relative proportions (called abundances) in each pixel [2]. A Linear Mixing Model (LMM) [3] is usually adopted for this problem. It assumes that a pixel with index n among the N pixels of the image $\mathbf{x}_n \in \mathbb{R}^L$, where L is the number of spectral bands used, is decomposed into:

$$\mathbf{x}_n = \sum_{p=1}^{P} a_{pn}\mathbf{s}_p + \mathbf{e}_n \qquad (1)$$

where P is the number of materials considered, $\mathbf{s}_p \in \mathbb{R}^L$ is the signature of endmember p, supposed to perfectly represent the corresponding material, and a_{pn} is the abundance coefficient of this material in pixel n, and $\mathbf{e}_n$ is an additive noise. The abundance coeffcients, being proportions, are subject to the abundance nonnegativity constraint (ANC) and the abundance sum-to-one constraint (ASC), i.e. each pixel must be fully explained by positive contributions of the different materials. Geometrically, the data lie in a simplex spanned by the endmembers.

The LMM has been extensively used in the past two decades but over the years some of its limitations became apparent, namely nonlinearities in the mixing processes and material variability. The former occur in complex scenarios such as tree canopies or particulate materials (e.g. sand), and require more complex mixture models such as bilinear mixture models [4, 5]. Endmember variability, on the other hand, simply corresponds to the fact that each material cannot be completely represented by a single spectrum, and is always subject to some intra-class variations [6]. These can be due to multiple factors, the main ones being local changes in the illumination conditions (because of shadows or topography), and the intrinsic variability of the materials, corresponding to local physico-chemical changes in the composition of the materials. This variability can be considered in the spatial [7] or temporal [8] domains when dealing with sequences of images. Here we focus on variability within a single image. Two classes of methods to tackle it have been identified in [9]: bundle-based methods, where endmembers are represented by a collection of signatures, possibly extracted from the data [10], and statistical methods, where statistical distributions are assigned to the endmembers to allow them to be estimated in each pixel [11]. Recently, explicit variability models appeared to directly explain the possible variations of the endmembers [12, 13]. Typically, these are in the form of (1), except that the endmembers are now indexed by the pixels as well. One of such models is the Extended Linear Mixing Model (ELMM), proposed in [14]. It approximates the physical radiative transfer model of Hapke [15] into a tractable version. In [16], the variability of the materials due to illumination is indeed shown to be reasonably explained by pixel dependent positive scaling factors. In its simplest version, the model writes:

$$\mathbf{x}_n = \psi_n \sum_{p=1}^{P} a_{pn}\mathbf{s}_p + \mathbf{e}_n \qquad (2)$$

where ψ_n is a scaling factor accounting for brightness changes in the spectral signatures of the pixels. In this model, $\mathbf{s}_p$ is no longer the one signature of material p, but rather a reference endmember providing the *direction* of a straight line joining the origin and that point, on which each local endmember lies. This led to a fast algorithm to estimate the parameters of this model using nonnegative least squares, called Scaled (partially) Constrained Least Squares

Fig. 1. Geometric interpretation of the ELMM, as presented in [14]. Endmembers are further allowed to drift way from scaled versions of the references through a Gaussian prior.

Unmixing (SCLSU), see [13] for details. However, this algorithm cannot handle other types of variability than scaling factors. All intrinsic variability phenomena, which are much harder to model physically because of the diversity of their causes, are not taken into account.

A more complex version of the ELMM algorithm [17, 14] was proposed later. This version considers different scaling factors for each material, but also allows each local endmember to drift away from a scaled version of the reference, in terms of the Euclidean distance. From a Bayesian point of view, this amounts to putting a Gaussian prior on each local endmember, centered on the scaled version of the reference. This way, the recovered variability is more complex than just a scaling factor variation, and can explain intrinsic variability effects to some extent. This is done by solving the following optimization problem:

$$\underset{\mathbf{A}\in\Delta_P,\mathcal{S},\mathbf{\Psi}}{\arg\min} \frac{1}{2}\sum_{n=1}^{N}\left(\|\mathbf{x}_n - \mathbf{S}_n\mathbf{a}_n\|_2^2 + \lambda_S\|\mathbf{S}_n - \mathbf{S}_0\boldsymbol{\psi}_n\|_F^2\right) \quad (3)$$

where $\mathcal{S} = \{\mathbf{S}_n\}, n = 1, ..., N$, and $\mathbf{A} \in \Delta_P$ means that each abundance vector $\mathbf{a}_n \in \mathbb{R}^P$ in each pixel belongs to the unit simplex with P vertices, i.e. complies with the ASC and ANC. $\boldsymbol{\psi}_n \in \mathbb{R}^{P\times P}$ is a diagonal matrix comprising material dependent scaling factors on its diagonal. $\mathbf{S}_0$ is a matrix containing reference endmembers. $\|.\|_F$ denotes the Frobenius norm.

The geometric interpretation of the ELMM is presented in Fig. 1. The data lies in a convex cone, whose edges are the endmembers, and local instances of those (not represented here) are points close to specific locations on these lines (the scaled reference endmembers). Each pixel still belongs to a simplex. This formulation has proved useful, but relies critically on the reference endmembers $\mathbf{S}_0$. They have to be carefully chosen because they provide the edges of the convex cone, which condition the whole unmixing process. The Vertex Component Analysis (VCA) algorithm [18] is one of the best known endmember extraction algorithms, which rely on the convex geometry of the purely linear unmixing problem. However, in situations where shadows or other strong variability effects are present, these algorithms, which look for extreme points of the dataset, often recover signatures with very small magnitude. This can be explained in the framework of the conic model that is the ELMM. Small magnitude spectra are close to the origin and are due to shadows; if a simplex based model is assumed, then these points are indeed extreme in the dataset. They are poor representatives of the spectra of the materials and cause errors in the estimation of the unmixing parameters.

In this paper, we propose a new formulation of the ELMM and a new unmixing algorithm which tackles this issue. We show that using a simple k-means clustering algorithm with the cosine sim-

ilarity measure is able to provide reliable enough initial reference signatures, which we are able to refine in the unmixing algorithm, by modeling reference endmembers as directional data, i.e. data lying on the unit hypersphere [19]. A point on the sphere uniquely defines a direction in the feature space.

2. PROPOSED MODEL AND ALGORITHM

2.1. Presentation of the approach

According to the ELMM, a material should not be summarized by a single point anymore, but all the points on a line joining the origin and a reference represent the same material, with different brightness variations. Ideally, an endmember should then be seen as a direction (or a line passing through the origin) in the feature space. With this in mind, it makes more sense for an endmember extraction algorithm to look for directions rather than extreme points of the dataset. Finding reference directions can be simply achieved using the k-means algorithm as a pre-processing step, with the cosine similarity measure $d(\mathbf{x}_i, \mathbf{x}_j) = 1 - \frac{\mathbf{x}_i^\top \mathbf{x}_j}{\|\mathbf{x}_i\|_2\|\mathbf{x}_j\|_2}$. The rationale behind this is that this measure is insensitive to scaling variations of the spectra. Interestingly, it can be shown that k-means with the cosine similarity can be related to a mixture of Von Mises-Fisher distributions (an analog of the Gaussian distribution for directional data, i.e. data lying on the unit sphere [19]), in a similar way to the fact that the regular k-means is related to a mixture of Gaussian distributions [20], which supports the use of k-means and the cosine similarity in a conic model.

We can simply then use the centroids of the clusters as reference endmembers. However, these may not be perfect and too close to the center of the true convex cone, since the clustering makes hard assignments for each pixel. This means that the average of the pixels belonging to a cluster will tend to be biased toward the center of the cone due to the influence of mixed pixels.

It is then necessary to be able to correct these endmembers on the fly in the unmixing algorithm. A way to do this in the LMM literature is to use the volume of the simplex as a regularizer for the endmembers [21]. However, this regularizer is nonconvex and hard to handle, and a simpler convex relaxation is considered in [22] using sum of Euclidean distances between endmembers. It is not straightforward to adapt this to the conic model of the ELMM, because of the fact that each endmember can be represented by any point on the corresponding line. It makes no sense trying to compare the Euclidean distance between two points on two lines if they are allowed to have different magnitudes (as we will confirm in the experiments). Using the spectral angle would be a possibility, but this quantity is quite hard to handle in optimization problems.

To solve this issue, we propose to see the reference endmembers as *directional data*, that is as directions in the feature space. There are several ways to do this [23], the simplest being to constrain the reference endmembers to be normalized: $\|\mathbf{s}_p\|_2 = 1 \ \forall p$. As a matter of fact, each line passing through the origin is uniquely represented by a point on the unit hypersphere. Equivalently, the reference endmember matrix should have unit norm columns, i.e. $\mathbf{S}_0 \in \mathcal{OB}(L, P)$, the so-called oblique manifold [23]. This way, the distance between the references on the sphere directly impacts the position of the lines: the greater (resp. smaller) the associated regularization parameter, the closer (resp. further) to the line joining the origin and the centroid of the dataset the lines will be, with a worse (resp. better) data fitting. We will see that the position of the reference lines, (along with the looseness of the Gaussian prior on the local endmembers) will also impact the sparsity of the abundances.

Thus, we design a refined formulation of the ELMM which allows to iteratively adjust the position of the reference endmember

lines in the feature space. To do that, we propose to minimize the following cost function:

$$\underset{\mathbf{A},\mathcal{S},\boldsymbol{\Psi},\mathbf{S}_0}{\operatorname{argmin}} \quad \frac{1}{2}\sum_{n=1}^{N}\left(||\mathbf{x}_n - \mathbf{S}_n\mathbf{a}_n||_2^2 + \lambda_S\,||(\mathbf{S}_n - \mathbf{S}_0\boldsymbol{\psi}_n)||_F^2\right)$$

$$+ \frac{\lambda_{S_0}}{2}\,\operatorname{tr}\left(\mathbf{S_0VS_0^\top}\right)$$

$$\text{s.t.} \quad \mathbf{a}_n \in \Delta_P\;\forall n$$

$$\mathbf{S}_0 \in \mathcal{OB}(L, P) \tag{4}$$

where tr denotes the trace of a matrix, and $\mathbf{V} = P\mathbf{I}_P - \mathbb{1}_P\mathbb{1}_P^\top$ ($\mathbb{1}_P$, being a column vector of P ones), such that $\operatorname{tr}(\mathbf{S}_0\mathbf{V}\mathbf{S}_0^\top) = \sum_{i=1}^{P-1}\sum_{j=i+1}^{P}||\mathbf{s}_{0i} - \mathbf{s}_{0j}||_2^2$, i.e. the sum of pairwise Euclidean distances between reference endmembers [22]. λ_S and λ_{S_0} are regularization parameters. The term $\lambda_S\,||(\mathbf{S}_n - \mathbf{S}_0\boldsymbol{\psi}_n)||_F^2$, forces each local endmember to be close (but not equal) to scaled versions the (unit norm) representatives of the reference directions. The scaling factors capture illumination induced variability, while $\mathbf{S}_n$ can further account for intrinsic variability effects. λ_S is directly related to the variance of the Gaussian prior (the same for all endmembers out of simplicity). The fact that the reference endmembers are normalized also has the advantage of easily allowing to compare the magnitude of the scaling factors (and thus the impact of illumination induced variability) across different materials and images. Spatial regularizations could also be added if need be, as done in [14].

2.2. Optimization

Here, we propose an algorithm to obtain a stationary point of the cost function (4). This objective function is challenging for several reasons: it is nonconvex over all variables simulaneously, which usually calls for block coordinate descent methods to get a local minimum. In this case, this approach is made even more complex because the problem is not convex w.r.t. $\mathbf{S}_0$ either, because of the nonconvex unit norm constraints. However, we will see we can still obtain a local minimum for this variable by taking advantage of the Riemannian manifold structure of the constraint set. Before detailing the different steps of the optimization, we will briefly describe how we initialize the algorithm. We first run the k-means clustering algorithm (with the cosine similarity) to obtain centroids, which we normalize to initialize $\mathbf{S}_0$. We initialize $\mathbf{S}_n$ by assigning the appropriate column of this matrix to the current pixel $\mathbf{x}_n$, depending on its clustering label. The other columns are initialized using the remaining centroids. The abundance and scaling factor matrices are initialized using the SCLSU algorithm with the centroids as references, which is very fast. This way, we hope to obtain a good local minimum in spite of the complexity of the problem.

The optimization w.r.t. $\mathbf{A}$ is relatively easy, the objective function being smooth, convex and the constraint set (unit simplex) easy to project onto [24]. The global minimum of this subproblem can be then obtained pixel-by-pixel using (for instance) a projected gradient descent. The optimizations w.r.t to $\mathbf{S}_n$ and $\boldsymbol{\psi}_n$ are easy and enjoy closed form solutions (see [17] for details). Optimizing over $\mathbf{S}_0$ is harder because of the unit sphere constraints, depsite the smoothness of the objective. Using the fact that the constraint set has a Riemannian manifold structure for which a retraction mapping can be easily found, we perform a conjugate gradient descent on the oblique manifold [23] (we use the Manopt toolbox for MATLAB [25]). The convergence of each subproblem is guaranteed, but we cannot prove the convergence of the global algorithm to a stationary point of the objective, although in practice convergence is always observed.

We stop the algorithm whenever the relative variations between consecutive iterates of all blocks of variables go below $\epsilon = 10^{-3}$ (in norms). We note that the convergence is going to be slower than the original ELMM with fixed reference endmembers, because the latter are now iteratively updated and impact the whole geometry of the unmixing.

3. EXPERIMENTS

In this section, we present results obtained on a synthetic dataset whose materials incorporate realistic variability features, as well as a quite challenging real dataset with very correlated endembers and the presence of a significant proportion of shadowed areas.

3.1. Synthetic dataset

3.1.1. Data Generation

We generate a synthetic dataset in the following way. First, we use the ground truth of the well known Pavia University dataset[1] to provide us with labeled spectra (203 bands in the visible and near infrared domains) belonging to several classes of interest, incorporating their spectral variability. We consider three classes present in that image: vegetation, concrete and metallic roofs. Theses classes incorporate both illumination induced variability (the roofs and trees have different orientations with respect to the sun) and more intrinsic variability sources (especially in concrete and vegetation). In each pixel, we choose the local endmembers to be a random sample within each of these classes (after a normalization to make each representative lie on the unit sphere).

Scaling factors have been simulated using a mixture of 4 Gaussian distributions (fitted from the results of SCLSU on a subimage of the Pavia dataset), which reflects the fact that in real scenarios scaling factors often come from multimodal distributions (for example roofs with two different orientations, or areas with shadows).

The abundances have been designed to be relatively sparse, using a Dirichlet distribution such that the probability density is concentrated around the edges and vertices of the unit simplex (while still allowing a proportion of heavily mixed pixels).

The data was then generated using Eq. (2), adding Gaussian white noise such that the signal to noise ratio is 30dB. The generated image then benefits from realistic statistical properties.

3.1.2. Results

First, we run the VCA+SCLSU algorithm to quickly get unmixing results with variability using VCA derived endmembers. We show below that this approach fails. Then we focus on testing two algorithms with k-means derived references: SCLSU and the ELMM of [17]. Also, we denote by ELMM-SSD (Sum of Squared Distances) the ELMM augmented with the convex volume regularization, but without the oblique manifold constraint. Finally, we compare all those methods to the proposed one, denoted as RELMM (for Robust ELMM). Note that we do not compare the results to the classical Fully Constrained Least Squares Unmixing [26], because this algorithm assumes a simplex-based model and has been shown to fail in many endmember variability scenarios. For each algorithm, we empirically tune the regularization parameters to obtain the best possible performance (the chosen values are reported in Table 1). Quantitative results are presented using two metrics: the abundance Root Mean Squared Error (aRMSE) between the true abundances and the recovered ones: $\frac{1}{N\sqrt{P}}\sum_{n=1}^{N}||\hat{\mathbf{a}}_n - \mathbf{a}_n||_2$, and the mean (over all pixels and materials) Spectral Angle Mapper (SAM) $SAM = \frac{1}{NP}\sum_{n=1}^{N}\sum_{p=1}^{P}\operatorname{acos}\left(\frac{\hat{\mathbf{s}}_{pn}^\top\mathbf{s}_{pn}}{||\hat{\mathbf{s}}_{pn}||_2||\mathbf{s}_{pn}||_2}\right)$ between the true endmembers in each pixel and the recovered ones. These quantities are gathered in Table 1.

The VCA+SCLSU approach obtains very poor results both in abundance estimation and variability retrieval, because indeed two

---

[1]http://www.ehu.eus/ccwintco/index.php?title=Hyperspectral_Remote_Sensing_Scenes\sharpPavia_University_scene

	λ_S	λ_{S_0}	aRMSE	SAM (degrees)	Time (s)
VCA+SCLSU	×	×	0.2075	54.4	2
SCLSU	×	×	0.0654	6.32	2
ELMM	0.01	×	0.0642	5.62	57
ELMM+SSD	0.1	0.25	0.1718	10.41	1455
RELMM	0.1	0.5	**0.0560**	**3.48**	2257

Table 1. Quantitative results on the synthetic data. Except for VCA+SLCSU, all algorithms use k-means to obtain the initial reference endmember matrix. Regularization parameters values are reported when applicable.


(a) (b) (c)

Fig. 2. Scatterplots of the data (blue), the true endmembers (black) and the extracted ones (red) for (a) SCLSU (b) ELMM (c) RELMM.

of the extracted signatures are associated with pixels with small scaling factors, and have a very low magnitude. Using k-means instead, along with SCLSU leads to better results, but far from optimal because the variability is only explained by scaling factors. the ELMM does even better because of the additional Gaussian prior. We see all the importance of the oblique manifold constraint on the last two lines: ELMM+SSD fails because the regularization terms involves the comparison of references with possibly different scales, whereas introducing the constraint leads to the best results. We show in Fig. 2 qualitative results using scatterplots of the data (using the fisrt three principal components) along with the recovered and true endmembers for the three best algorithms (the other scatterplots show that the lines are much too far away from the true cone to be relevant). Similar conclusions can be drawn from this figure, showing that RELMM is able to find the best endmembers in each pixel.

3.2. Real dataset

The real dataset we use was acquired in 2009 by Japan Space Systems over the Tama Forest Science Garden in the western region of Tokyo, with the CASI-3 sensor (72 spectral bands in the visible and near-infrared domains) [27]. The spatial resolution is 1m. The image we use is a $207 \times 268 \times 72$ subset of the whole scene. An RGB representation is shown in Fig. 3 (a). This dataset has been used for supervised classification of tree species, using a ground truth and LiDAR data as an additional classification feature, since the different tree species are spectrally very close to one another. The image also comprises many shadowed areas because of the tree crowns, which were an important hurdle in previous studies [27]. Furthermore, other non vegetation endmembers are present, such as man made roofs, roads, and soil. We show here that using k-means instead of the VCA allows to distinguish between conifer and broadleaf trees in a completely unsupervised way. Some labeled conifer and broadleaf trees are shown in Fig. 3 (b). We show the scatterplot of the data and labeled pure pixels in Fig. 3(c).

We unmix the data using $P = 4$ materials using the above mentioned algorithms. We show in Fig. 4 the scatterplots of the data and recovered endmembers for VCA+SCLSU, SCLSU, ELMM, and RELMM. The abundance maps are shown in Fig. 5. For the ELMM, we set $\lambda_S = 0.01$, and for the RELMM, we set $\lambda_S = 0.5$ and $\lambda_{S_0} = 100$. As in the synthetic data case, the endmembers recovered by VCA are spurious because of shadow patches of the image, and the corresponding abundances are meaningless. Most of the data is projected on the closest line in the identified cone, which represents vegetation. Using k-means instead allows to distinguish be-


(a) (b) (c)

Fig. 3. (a) RGB representation of the data. (b) Ground truth for conifer (green) and broadleaf (red) trees. (c) Scatterplot of the data, as well as the ground truth (same color code).


(a) (b) (c) (d)

Fig. 4. Scatterplots of the data (blue) and the extracted endmembers (red) for (a) VCA+SCLSU (b) SCLSU (c) ELMM (d) RELMM.

tween conifer trees and broadleaf trees. Grass and shadows are also detected by large and low values of the scaling factors, respectively (not shown here because of space constraints). The abundances of SCLSU and the ELMM are rather similar, slightly sparser for the ELMM, because it is able to better capture variability effects than SCLU (as seen in Fig. 4 (c). The RELMM, thanks to being able to adjust the references, is able to obtain sparser abundance maps which match closely the ground truth of Fig. 3 (b). We see that the identified endmembers enclose the data very well and are the closest to the ground truth pixels of Fig. 3 (c).

4. CONCLUSION

In this paper, we have proposed a new algorithm to unmix hyperspectral data taking into account both illumination-induced and intrisic variability of the endmembers. This algorithm is able to obtain better estimates of the reference endmembers than VCA by taking advantage of the fact that endmembers are essentially directional data. The reference signatures are constrained to be on the unit sphere since each point on it completely represents a line in the feature space. We obtain robust estimates of the edges of the convex cone in the Extended Linear Mixing Model framework. Results on synthetic data show the soundess of the approach while very good performance is observed on a challenging real dataset. Future work will include a way to automatically estimate material dependent parameters (such as the variances of the Gaussian priors), and tests on the various ways to represent endmembers as directional data.


Conifer trees Man made Soil Broadleaf trees

Fig. 5. Abundances maps obtained by the tested algorithms.

5. REFERENCES

[1] J.M. Bioucas-Dias, A. Plaza, G. Camps-Valls, P. Scheunders, N.M. Nasrabadi, and J. Chanussot, "Hyperspectral remote sensing data analysis and future challenges," *IEEE Geoscience and Remote Sensing Magazine*, vol. 1, no. 2, pp. 6–36, June 2013.

[2] N. Keshava and J. F. Mustard, "Spectral unmixing," *IEEE Signal Processing Magazine*, vol. 19, no. 1, pp. 44–57, Jan 2002.

[3] 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*, vol. 5, no. 2, pp. 354–379, April 2012.

[4] R. Heylen, M. Parente, and P. Gader, "A review of nonlinear hyperspectral unmixing methods," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 7, no. 6, pp. 1844–1868, June 2014.

[5] N. Dobigeon, J. Y. Tourneret, C. Richard, J. C. M. Bermudez, S. McLaughlin, and A. O. Hero, "Nonlinear unmixing of hyperspectral images: Models and algorithms," *IEEE Signal Processing Magazine*, vol. 31, no. 1, pp. 82–94, Jan 2014.

[6] B. Somers, G.P. Asner, L. Tits, and P. Coppin, "Endmember variability in spectral mixture analysis: A review," *Remote Sensing of Environment*, vol. 115, no. 7, pp. 1603 – 1616, 2011.

[7] L. Drumetz, G. Tochon, M.A. Veganzones, J. Chanussot, and C. Jutten, "Improved local spectral unmixing of hyperspectral data using an algorithmic regularization path for collaborative sparse regression," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 6190–6194.

[8] S. Henrot, J. Chanussot, and C. Jutten, "Dynamical spectral unmixing of multitemporal hyperspectral images," *IEEE Transactions on Image Processing*, vol. 25, no. 7, pp. 3219–3232, July 2016.

[9] A. Zare and K.C. Ho, "Endmember variability in hyperspectral analysis: Addressing spectral variability during spectral unmixing," *IEEE Signal Processing Magazine*, vol. 31, no. 1, pp. 95–104, Jan 2014.

[10] B. Somers, M. Zortea, A. Plaza, and G. P. Asner, "Automated extraction of image-based endmember bundles for improved spectral unmixing," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 5, no. 2, pp. 396–408, April 2012.

[11] A. Halimi, N. Dobigeon, and J. Y. Tourneret, "Unsupervised unmixing of hyperspectral images accounting for endmember variability," *IEEE Transactions on Image Processing*, vol. 24, no. 12, pp. 4904–4917, Dec 2015.

[12] P.-A. Thouvenin, N. Dobigeon, and J.-Y. Tourneret, "Hyperspectral unmixing with spectral variability using a perturbed linear mixing model," *IEEE Transactions on Signal Processing*, vol. 64, no. 2, pp. 525–538, 2016.

[13] M.A. Veganzones, L. Drumetz, R. Marrero, G. Tochon, M. Dalla Mura, A. Plaza, J.M. Bioucas-Dias, and J. Chanussot, "A new extended linear mixing model to address spectral variability," in *Proc. IEEE Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2014.

[14] L. Drumetz, M. A. Veganzones, S. Henrot, R. Phlypo, J. Chanussot, and C. Jutten, "Blind hyperspectral unmixing using an extended linear mixing model to address spectral variability," *IEEE Transactions on Image Processing*, vol. 25, no. 8, pp. 3890–3905, Aug 2016.

[15] B. Hapke, *Theory of reflectance and emittance spectroscopy*, Cambridge University Press, 2012.

[16] L. Drumetz, *Endmember variability in hyperspectral image unmixing*, Ph.D. thesis, Université de Grenoble Alpes, 2016.

[17] L. Drumetz, S. Henrot, M. A. Veganzones, J. Chanussot, and C. Jutten, "Blind hyperspectral unmixing using an extended linear mixing model to address spectral variability," in *IEEE Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS 2015)*, 2015, pp. 1–4.

[18] 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*, vol. 43, no. 4, pp. 898–910, April 2005.

[19] K. V. Mardia and P. E. Jupp, *Directional statistics*, vol. 494, John Wiley & Sons, 2009.

[20] A. Banerjee, I. S. Dhillon, J. Ghosh, and S. Sra, "Clustering on the unit hypersphere using von Mises-Fisher distributions," *Journal of Machine Learning Research*, vol. 6, no. Sep, pp. 1345–1382, 2005.

[21] J. Li and J.M. Bioucas-Dias, "Minimum volume simplex analysis: A fast algorithm to unmix hyperspectral data," in *Geoscience and Remote Sensing Symposium, 2008. IGARSS 2008. IEEE International*, July 2008, vol. 3, pp. III – 250–III – 253.

[22] M. Berman, H. Kiiveri, R. Lagerstrom, A. Ernst, R. Dunne, and J. F. Huntington, "ICE: A statistical approach to identifying endmembers in hyperspectral images," *IEEE transactions on Geoscience and Remote Sensing*, vol. 42, no. 10, pp. 2085–2095, 2004.

[23] P-A Absil, Robert Mahony, and Rodolphe Sepulchre, *Optimization algorithms on matrix manifolds*, Princeton University Press, 2009.

[24] L. Condat, "Fast projection onto the simplex and the $\mathcal{L}_1$ ball," *Mathematical Programming*, pp. 1–11, 2014.

[25] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, "Manopt, a Matlab toolbox for optimization on manifolds," *Journal of Machine Learning Research*, vol. 15, pp. 1455–1459, 2014.

[26] D.C. Heinz and Chein-I Chang, "Fully constrained least squares linear spectral mixture analysis method for material quantification in hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 3, pp. 529–545, Mar 2001.

[27] T. Matsuki, N. Yokoya, and A. Iwasaki, "Hyperspectral tree species classification of japanese complex mixed forest with the aid of LiDAR data," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 5, pp. 2177–2187, 2015.