Universiteit Antwerpen

This item is the archived peer-reviewed author-version of:

Hyperspectral unmixing with endmember variability via alternating angle minimization

Hyperspectral unmixing with endmember variability via alternating angle minimization

Rob Heylen, Alina Zare, Paul Gader, Paul Scheunders

Abstract—In hyperspectral unmixing applications, one typically assumes that a single spectrum exists for every endmember. In many scenarios, this is not the case, and one requires a set or a distribution of spectra to represent an endmember or class. This inherent spectral variability can pose severe difficulties in classical unmixing approaches. In this paper, we present a new algorithm for dealing with endmember variability in spectral unmixing, based on the geometrical interpretation of the resulting unmixing problem, and an alternating optimization approach. This alternating angle minimization algorithm uses sets of spectra to represent the variability present in each class, and attempts to identify the subset of endmembers which produce the smallest reconstruction error. The algorithm is analogous to the popular multiple endmember spectral mixture analysis technique, but has a much more favorable computational complexity, while producing similar results. We illustrate the algorithm on several artificial and real data sets, and compare with several other recent techniques for dealing with endmember variability.

I. Introduction

HYPERSPECTRAL images are obtained from a relatively large distance in most remote sensing applications. The ground instantaneous field of view (GIVOF) of a pixel will become large, and will, for example, cover several meters squared. As a consequence, several different objects and/or materials can be expected in the GIVOF. All of the spectral signatures of these materials will combine to contribute to the measured spectrum.

Spectral mixture analysis aims to invert this process to learn the unique spectral signature of each material and determine their relative abundances in the GIVOF of each pixel. The unique spectral signatures of each material in a hyperspectral scene are termed *endmembers*. A large class of unmixing approaches assumes that a fixed and limited number of endmembers is present in a scene. The mixed spectra can then be decomposed into endmember contributions, using various techniques based on the inversion of linear or nonlinear mixing models, with many possible constraints on the model parameters, such as convexity or sparsity constraints. Several approaches which do not assume the presence of a mixing model exist as well, and we refer the interested reader to the recent review papers [1], [2], [3] for more details.

However, since the majority of these spectral unmixing methods assume that a fixed and limited number of endmembers (i.e., namely, a unique, fixed endmember spectral signature per material) are present in the scene, these approaches lack the ability to represent the spectral variability of the endmembers in a scene. The assumption that a single endmember spectrum exists for each class or material of interest is not valid in many unmixing applications. The spectral signature for a material varies within hyperspectral data collections due to reasons that range from environmental, atmospheric, scale, to temporal factors. For example, variation in illumination conditions [4], [5], [6], differences in the architecture of plant canopies, changes in the distribution of leaf orientation in vegetated regions (i.e., variation in spectral signature of the top to the underside of a leaf), varying building structure and layout in urban areas [7], composition and density of mineral grains in soil, and changing atmospheric conditions cause variation in spectral signature. Also scale differences will cause variability in many materials, as one is typically interested in the spectrum of an object at scales close to the image resolution. For example, a tree contains branches, which contain wood and bark, which contain mosses, dirt, and so on. At each scale, different spectra can be observed and defined, with different fractions, leading to an intrinsic variability in the spectrum of the entire tree due to differences in the makeup of each specimen.

Although spectral variability due to these sources is expected in many hyperspectral data sets, generally, methods in hyperspectral unmixing and endmember estimation do not account for spectral variability. As such, errors resulting from inaccurate or insufficient endmember representation will be propagated throughout analysis. The most prominent effects from inaccurate endmember representations are resulting errors in estimated proportion values, termed *proportion indeterminacy*, or the use of too many endmembers to represent a spectrum [5], [7]. In order to avoid these errors and to represent spectral variability during analysis, a number of spectral unmixing and endmember estimation algorithms that incorporate spectral variability have been developed in the literature. For recent overviews, see [8], [9]. A summary of the methods in the literature that account for spectral variability are shown in the diagram in Figure 1.

Most of these techniques fall under one of two classes: One either models the spectral variability within each class as a probability distribution, or one symbolizes each class by building a library of representative spectra. Examples of the former method can be found in [10] and [11], where the endmember variability is captured by modeling the endmember classes as normal or beta distributions respectively. The most well-known technique which represents endmember variability by employing libraries is multiple endmember spectra mixture analysis (MESMA) [12], on which we will also focus in this work. MESMA has been successfully used in a large number of fields, such as vegetation monitoring [13], [12], [14], urban remote sensing [15], [16], and landform [17] or fire mapping [18], [19], [20], [21].
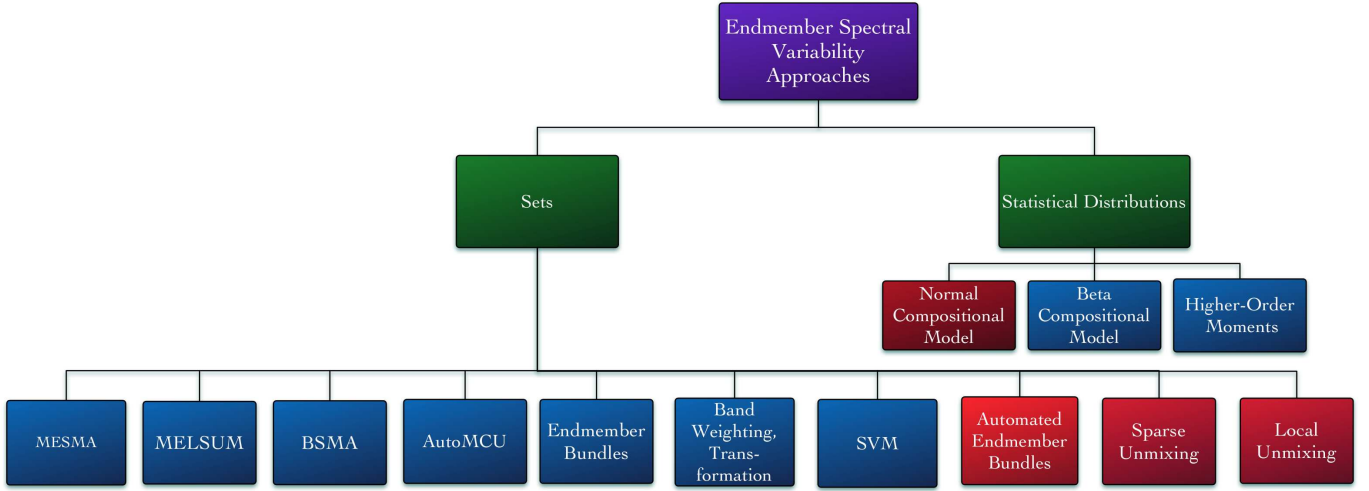
Fig. 1: Summary of spectral unmixing and endmember estimation algorithms that account for spectral variability. Methods shown in blue perform spectral unmixing. Methods shown in red perform spectral unmixing and endmember estimation. A more complete review of the methods listed in this diagram can be found in [8] and [9].

In the MESMA algorithm, one employs the linear mixing model (LMM), i.e., the assumption that every observed spectrum can be modeled as a convex linear combination of endmember spectra. Each endmember spectrum can be drawn from a library, which should contain enough spectra to correctly capture the variability which is present in the corresponding class. Once a single endmember from each class has been chosen, this set of endmembers, also called model, is used to unmix the target spectrum, and a reconstruction error (RE) is calculated. This is repeated for all possible models until a set of conditions has been fulfilled, such as constraints on the abundances, and REs below a predetermined threshold. Remark that most MESMA approaches assume the presence of a library containing only a photometric shade or zero vector endmember.

In the worst-case scenario, the target spectrum has to be unmixed with respect to all possible models, and the computational complexity of the algorithm scales as the product of the library sizes. If the number of libraries is not very small, or there are several large libraries present, the computational effort required to unmix with respect to all possible combinations becomes unrealistic to be carried out in practice. Therefore, several techniques have been investigated to speed up the MESMA algorithm, or to provide early stopping rules.

In the original version of MESMA presented in [12], one performs sum-to-one constrained unmixing with respect to all possible combinations of libraries and endmember candidates from these libraries. For each unmixing result, one checks whether the obtained abundances are within a given range (e.g., positive and smaller than one, with a given tolerance), whether the RE is below a given threshold, and if the number of iterations and the number of large band errors has not exceeded some predetermined values. With these additional demands, excessively long iteration times can usually be avoided, but the choice of thresholds and tolerances will have a large effect on both the runtime and the quality of the results

[22], [9].

Several authors [8] choose to perform an exhaustive search of all possible combinations, and simply retain the one with the lowest RE. Models which do not obey the abundance non-negativity constraint (ANC) or the abundance sum-to-one constraint (ASC) can be discarded, or avoided by using fully-constrained least-squares unmixing (FCLSU). This approach will invariably lead to computational difficulties for larger libraries, as no early stopping rules are enforced.

Approaches to reduce the sizes of the libraries have been proposed as well. In [21], the pixel purity index (PPI) algorithm is used to preselect library candidates from the image pixels. In [18], each endmember in a given library is written as an optimal linear combination of the other endmembers in that library by considering it as a separate unmixing problem which is solved with MESMA. The library members which occur the most in the unmixing results are then used as the members of a reduced library, as they are the most representative for the entire library. This way, severely reduced library sizes can be obtained: In [18], each library is reduced to contain only 3-5 members. A similar approach is used in [13], where the endmembers which produce the smallest error for modeling the other library endmembers are retained.

As the inclusion of more degrees of freedom always allows for a reduction in RE, the inclusion of more libraries in the unmixing of a given spectrum will result in lower REs, but not always in better results on the level of abundances. Therefore, several approaches demand a large relative decrease in RE before accepting unmixing results with a higher numbers of endmembers [23]. A different approach is that each (non-shade) endmember should contribute significantly to the total spectral signal, and should thus have an abundance larger than some predefined threshold.

Despite all these developments, the MESMA algorithm still contains a relatively simple but computationally demanding step in its implementation, where spectra have to be unmixed with respect to a large number of possible endmember combi-

nations. In this paper, we aim to alleviate this problem by introducing a new method for performing MESMA which scales linearly in the number of libraries and endmembers instead of as a product of the library sizes. This allows the application of MESMA to problems with a large number of large libraries, without compromising on the results with some early stopping rules.

The technique is based on an alternating optimization with respect to each library separately. The optimization itself is rewritten in terms of an angle minimization problem, which can be easily solved using standard algebraic relations. While alternating optimization in this sense does not guarantee convergence to the global optimum, we show that it is exactly the high dimensionality of the data which causes this technique to work in most practical cases. This way, alternating optimization techniques can be employed to estimate the model which will give the smallest RE, and unmixing with respect to all possible combinations, which is a combinatorial problem, can be avoided.

This paper is structured as follows: In section II we formally introduce the version of the MESMA algorithm employed in this manuscript, provide an efficient computational implementation, and discuss several of its properties, such as its computational complexity. In section III, we explain how optimization with respect to a single library is equivalent with a high-dimensional angle minimization problem, and employ this equivalence to construct an alternating optimization scheme which aims to solve the same problem as the MESMA algorithm. The resulting alternating angle minimization (AAM) algorithm is tested on several artificial and real-world data sets in section IV, followed by the conclusions in section V.

II. THE MESMA ALGORITHM

Because several versions of MESMA are available in the literature, with different preprocessing steps, early stopping rules, or other differences in their implementation details, we first provide a detailed explanation of the version of MESMA that we will consider in this manuscript.

In its simplest implementation, the MESMA algorithm will unmix a spectrum with respect to all possible models, where a model is a set of endmembers drawn from their respective endmember libraries. The model which yields the smallest reconstruction error and obeys the ANC and ASC is retained, and the unmixing results with respect to this model are returned.

Because the solution has to obey the ANC and ASC, one possible approach is to employ constrained optimization techniques, such as FCLSU or non-negatively constrained least squares (NCLS). Such an approach would have a very high computational complexity, and fortunately constrained optimization can be avoided if one considers all possible library combinations on top of all possible models drawn from a set of libraries.

Consider a single model, made up by the endmember set $\{e_1, \ldots, e_p\}$, and a single pixel x. According to the linear mixing model, the pixel x can be written as a convex combination of endmembers, and additional noise:

$$x = \sum_{i=1}^{p} a_i e_i + \eta, \qquad \begin{cases} \forall i : a_i \geq 0 & \text{(ANC)} \\ \sum_i a_i = 1 & \text{(ASC)} \end{cases} \tag{1}$$

If one performs FCLSU of x with respect to the endmember set $\{e_1, \ldots, e_p\}$, one solves the following optimization problem:

$$\underset{\{a_1, \ldots, a_p\}}{\mathrm{argmin}} \left\| x - \sum_{i=1}^{p} a_i e_i \right\|^2, \qquad \text{s.t.} \begin{cases} \forall i : a_i \geq 0 \\ \sum_i a_i = 1 \end{cases} \tag{2}$$

Because minimizing a least-squares error is equivalent with minimizing an Euclidean distance, this minimization problem can also be interpreted as a simplex projection problem, where one tries to find the point of minimal distance to x inside the simplex spanned by the endmembers [24], [25].

Many approaches for solving this constrained minimization problem have been proposed, such as active set approaches, quadratic programming, the alternating direction method of multipliers, extending NCLS techniques, non-negative matrix factorization, alternating projection methods, and geometrical approaches. Almost all of these approaches contain some iterative step which has to be repeated until some stopping criterion is obeyed, which makes them computationally heavy.

Because implementing the ASC is easy, a seemingly naive method for solving the FCLSU problem (2) can be derived by employing active sets: The active set is the set of endmembers with active constraints, hence with zero abundances. If one knows the active set beforehand, one can perform sum-to-one constrained least-squares unmixing (SCLSU) with respect to the endmembers not in the active set, and one will obtain the same unmixing result as FCLSU with respect to the entire endmember set. Because the active set is not known, one can simply evaluate every possible set, retain only those that obey the positivity constraints, and return the one corresponding with the smallest reconstruction error. In practice, one needs to unmix the target x with SCLSU with respect to every non-empty subset of endmembers drawn from $\{e_1, \ldots, e_p\}$, and retain the result that obeys the ANC and yields the smallest reconstruction error. As the number of non-empty subsets is $2^p - 1$, this approach can have a large computational complexity as well, especially for larger p.

However, in the case of MESMA, this approach will generally be computationally favorable over applying FCLSU to every model with p endmembers. Consider p endmember libraries $\{L_1, \ldots, L_p\}$, where library L_i has N_i elements. We can distinguish two cases:

- Apply FCLSU to every model of p endmembers. This requires executing the FCLSU algorithm $\prod_i N_i$ times.
- Apply SCLSU to every model with any number of endmembers, for every library combination. This requires executing the SCLSU algorithm $\prod_i (N_i + 1) - 1$ times.

Because FCLSU requires solving an optimization problem and SCLSU only requires some basic linear algebra operations, the latter approach can be orders of magnitude faster, especially for large data sets and libraries. Therefore, we present

---

Algorithm 1: MESMA algorithm

---

1 **for** *Every subset of libraries* **do**
2   **for** *Every model drawn from these libraries* **do**
3   Determine the abundances with SCLSU;
4   **if** *all abundances are non-negative* **then**
5   Determine the RE;
6   **else**
7   Set RE to ∞;

8 Return the solution with smallest reconstruction error;

Algorithm 1 as the MESMA algorithm we will use in this paper as a benchmark.

Line 1 can be easily implemented by counting from 1 to 2^p and using this counter as a bit mask to select the libraries, or by generating the banker's sequence [26]. Line 2 can be implemented by using a vector of overflowing counters. To determine the abundances with SCLSU, several approaches can be employed. First of all, one has to make the distinction between cases where a shadow endmember is present or not. Many MESMA approaches assume that a shadow or zero endmember always has to be present, which actually removes the sum-to-one constraint from the problem, as the abundance of the shadow endmember can be simply derived as one minus all other abundances. The problem then reduces to solving an unconstrained least-squares problem, for which many approaches can be employed.

Also when no shadow endmember is present, the same approach can be used. Because unmixing is a translation invariant operation, one can simply designate one of the endmembers as a shadow endmember, subtract it from all other endmembers and the target, and again perform the unconstrained least-squares unmixing. In practice, the Moore-Penrose pseudo inverse can be employed to solve the unconstrained least-squares problem. Let $E = (\boldsymbol{e}_1, \ldots, \boldsymbol{e}_p)$ be a matrix which lists the endmembers in the model as its columns, and $\boldsymbol{x}$ the target as a column vector. With $\boldsymbol{J} = (\boldsymbol{e}_2 - \boldsymbol{e}_1, \ldots, \boldsymbol{e}_p - \boldsymbol{e}_1)$, we can obtain the abundances by first calculating

$$\hat{\boldsymbol{a}} = (\boldsymbol{J}^T \boldsymbol{J})^{-1} \boldsymbol{J}^T (\boldsymbol{x} - \boldsymbol{e}_1) \tag{3}$$

The abundance vector will then be given as $\boldsymbol{a} = [1 - \sum_i \hat{a}_i; \hat{\boldsymbol{a}}]$, and the projected or reconstructed point as $E\boldsymbol{a}$. Remark that these equations can be simplified if $\boldsymbol{e}_1 = \overline{0}$, or simpler approaches can be used when $p = 1$ or $p = 2$.

Computationally, the presented version of the MESMA algorithm performs $\prod_i (N_i + 1)$ iterations, where the SCLSU algorithm is executed in each iteration. The SCLSU algorithm contains mainly matrix multiplications and pseudo-inversions, and hence has an order d^3 complexity, with d the dimensionality of the spectral vectors.

It must be noted that many versions of MESMA present in the literature possess one of several early stopping rules, or perform some selective pruning of the spectral libraries to reduce their sizes. This will generally reduce the computational requirements, at the cost of suboptimal solutions.

III. ALTERNATING ANGLE MINIMIZATION

The main computational bottleneck in the MESMA algorithm is the requirement to unmix with respect to a very large amount of different models. We propose an alternative approach, where an alternating optimization is carried out with respect to the libraries. The technique is based on a geometrical interpretation of the unmixing problem. This interpretation has been presented in many papers already [24], [25], [27], and we will merely repeat some of the main results here.

If we interpret each spectrum as a d-dimensional vector in a spectral space $[0, 1]^d$, and we assume that spectral noise is identically distributed in each spectral band, applying the maximum likelihood method to the linear unmixing problem will yield a least-squares problem. This least-squares problem is equivalent to finding a minimal Euclidean distance from the target data point $\boldsymbol{x}$ onto a given subset S of spectral space.

$$\boldsymbol{y} = P_S(\boldsymbol{x}) \iff \forall \boldsymbol{z} \in S : \|\boldsymbol{z} - \boldsymbol{x}\| \geq \|\boldsymbol{y} - \boldsymbol{x}\| \tag{4}$$

The subset S contains all points accessible under the employed mixing model (the LMM in this case), and depends on the employed constraints on the abundances:

- No constraints: the linear subspace spanned by the endmembers
- ASC: The affine space, or hyperplane, through the endmembers
- ANC: The convex cone of the endmember set
- ANC and ASC: The simplex spanned by the endmembers

The RE can then be defined as $\|\boldsymbol{x} - P_S(\boldsymbol{x})\|$, and the abundances are given by the barycentric coordinates of the projection $P_S(\boldsymbol{x})$ with respect to the endmembers [25], [27].

Let us first consider unmixing with only the ASC on the abundances. Solving the unmixing problem is then equivalent to projecting the target $\boldsymbol{x}$ orthogonally onto the hyperplane $H(E)$ through the endmembers in E. Remark that this hyperplane typically does not contain the origin, unless one of the endmembers is a zero or shadow endmember. Only in this case will the hyperplane be a linear subspace as well.

The RE is given by the orthogonal distance from $\boldsymbol{x}$ to $H(E)$, or as $\|\boldsymbol{x} - P_{H(E)}(\boldsymbol{x})\|$. If we define $F = E/\{\boldsymbol{e}_p\} = (\boldsymbol{e}_1, \ldots, \boldsymbol{e}_{p-1})$, using standard geometrical relations one can write this orthogonal distance as

$$\big\|\boldsymbol{x} - P_{H(E)}(\boldsymbol{x})\big\| = \big\|\boldsymbol{x} - P_{H(F)}(\boldsymbol{x})\big\| \sin(\theta) \tag{5}$$

with

$$\theta = \min(\alpha, \pi - \alpha) \tag{6}$$

$$\alpha = \cos^{-1}\left(\frac{(\boldsymbol{e}_p - P_{H(F)}(\boldsymbol{e}_p)) \cdot (\boldsymbol{x} - P_{H(F)}(\boldsymbol{x}))}{\|\boldsymbol{e}_p - P_{H(F)}(\boldsymbol{e}_p)\| \, \|\boldsymbol{x} - P_{H(F)}(\boldsymbol{x}))\|} \right) \tag{7}$$

$$= \sin^{-1}\left(\frac{\|\boldsymbol{x} - P_{H(E)}(\boldsymbol{x})\|}{\|\boldsymbol{x} - P_{H(F)}(\boldsymbol{x})\|} \right) \tag{8}$$

This is illustrated for a three-endmember scenario in Fig. 2. The orthogonal distance of $\boldsymbol{x}$ to the plane $H(E)$ through the endmembers in $E = (\boldsymbol{e}_1, \boldsymbol{e}_2, \boldsymbol{e}_3)$ can also be written as the orthogonal distance to the plane $H(F)$ through $F = (\boldsymbol{e}_1, \boldsymbol{e}_2)$, multiplied by the sine of the angle θ of this projection with
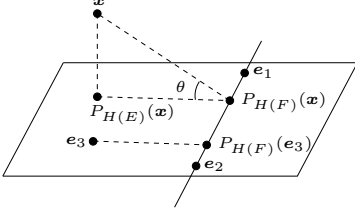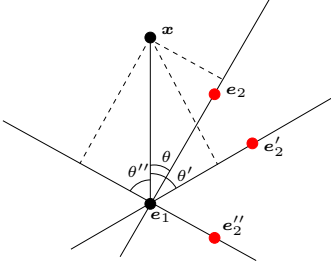
Fig. 2: Illustration of relation (5).

Fig. 3: Optimizing the RE corresponds with identifying the smallest angle θ

$H(E)$. The angle θ is also the angle between the vectors $\boldsymbol{x} - P_{H(F)}(\boldsymbol{x})$ and $\boldsymbol{e}_3 - P_{H(F)}(\boldsymbol{e}_3)$. This reasoning also works for higher dimensional spaces, and for higher numbers of endmembers in these spaces (with the constraint that $p \leq d$).

Remark that the roles of $\boldsymbol{x}$ and $\boldsymbol{e}_p$ can also be interchanged for calculating the angle θ. With $G = F \cup \boldsymbol{x} = (\boldsymbol{e}_1, \ldots, \boldsymbol{e}_{p-1}, \boldsymbol{x})$, we can also write

$$\alpha = \sin^{-1}\left(\frac{\|\boldsymbol{e}_p - P_{H(G)}(\boldsymbol{e}_p)\|}{\|\boldsymbol{e}_p - P_{H(F)}(\boldsymbol{e}_p)\|}\right) \qquad (9)$$

We can use this observation to easily find the optimal endmember in a single-library situation [28]. Suppose that the $p-1$ endmembers in $F = (\boldsymbol{e}_1, \ldots, \boldsymbol{e}_{p-1})$ are fixed, and the final endmember $\boldsymbol{e}_p$ can be chosen from a library L_p with N_p elements. Because $\|\boldsymbol{x} - P_{H(F)}(\boldsymbol{x})\|$ is a constant, (5) states that the endmember with the smallest angle θ will also be the one which will generate the lowest RE. Obtaining the optimal endmember can then be done by calculating the angles θ for every library element, and choosing the smallest angle.

An illustration in two dimensions, with a library of three endmembers, is given in Fig. 3. The point $\boldsymbol{x}$ is unmixed with respect to two endmembers, where one is a fixed point $\boldsymbol{e}_1$, and the second has to be chosen from a library containing three options $\boldsymbol{e}_2$, $\boldsymbol{e}_2'$ and $\boldsymbol{e}_2''$. It can be easily seen that the option with the lowest RE will also be the one with the smallest angle θ.

However, in a general scenario, each endmember can be drawn from its respective library. A strategy often used in optimization theory is to perform the optimization with respect to each library iteratively, until no further improvement is found or some maximal number of iterations has been reached. This is also the approach we use in the proposed alternating angle minimization (AAM) algorithm: Starting from a random initial model, we repeatedly optimize each endmember in turn using the minimal angle approach presented above, for a maximal number of iterations. As we only considered the ASC so far, we will employ a similar approach as in the MESMA algorithm in order to implement the ANC: We consider all possible subsets of libraries, perform the iterative optimization for each subset, and do fully-constrained unmixing with respect to the optimal model from each subset. The solution with the smallest RE is returned after considering all subsets. The AAM algorithm is given in Algorithm 2.

Algorithm 2: AAM algorithm

1 Select a number of iterations K.
2 **for** *Every non-empty subset $S \subset \{1, \ldots, p\}$* **do**
3   Let $q = |S|$ be the cardinality of S
4   Let $\boldsymbol{e}_j^i$ be shorthand for $L_{S_i}(j)$, endmember j from library L_{S_i}
5   Select a random index set $\{I_i\}_{i=1}^q, I_i \in 1, \ldots, N_{S(i)}$.
6   **for** *iteration $\in [1, \ldots, K]$* **do**
7   **for** *$i \in [1, \ldots, q]$* **do**
8          $F = \{\boldsymbol{e}_{I_1}^1, \ldots, \boldsymbol{e}_{I(i-1)}^{i-1}, \boldsymbol{e}_{I(i+1)}^{i+1}, \ldots, \boldsymbol{e}_{I_q}^q\}$
9          $G = F \cup \boldsymbol{x}$
10   **for** *$n \in [1, \ldots, N_{S(i)}]$* **do**
11             $p_n = \sin^{-1}\left(\frac{\|\boldsymbol{e}_n^i - P_{H(G)}(\boldsymbol{e}_n^i)\|}{\|\boldsymbol{e}_n^i - P_{H(F)}(\boldsymbol{e}_n^i)\|}\right)$
12   **if** $(\boldsymbol{x} - P_{H(F)}(\boldsymbol{x})) \cdot (\boldsymbol{e}_n^i - P_{H(F)}(\boldsymbol{e}_n^i)) < 0$ **then**
13                $p_n = \pi - p_n$
14          $I_i = \text{argmin}_n(\{p_n\}_{n=1}^{N_i})$
15    $E = \{\boldsymbol{e}_{I_1}^1, \ldots, \boldsymbol{e}_{I_q}^q\}$
16    $\boldsymbol{a}(S) = \text{FCLSU}(\boldsymbol{x}, E)$
17    $\text{RE}(S) = \|\boldsymbol{x} - E\boldsymbol{a}(S)\|$
18 Return results corresponding with $min(\text{RE})$

The hyperplane projections $P_{H(.)}$ can be performed with (3). In line 16, the abundances are determined by fully-constrained unmixing $\boldsymbol{x}$ with respect to the endmembers in E. Many algorithms exist to perform this step, and we used the FCLSU algorithm presented in [29]. In the angle minimization, ties are randomly broken, as these are not expected to occur much due to the high-dimensional vector inputs of real numbers.

For a given library subset, the RE will decrease after each update in the algorithm until a local minimum is reached or the maximal number of iterations has been performed. However, it cannot be guaranteed that the global minimum will be found, and one can easily construct situations in low dimensional spaces where the alternating optimization will not reach the global minimum. This is especially the case when the target spectrum lies inside the convex hull of the libraries. An example of this situation is given in Fig. 4 (a), where depending on the initialization, a different final result can be expected. When the point lies outside the convex hull determined by the libraries, as in Fig. 4 (b), the alternating optimization will reach the correct global minimum in this case.

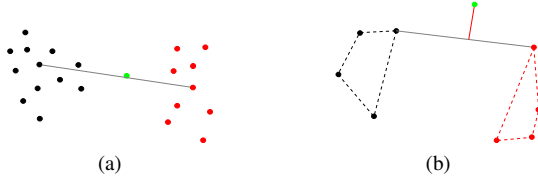Now, a feature of high-dimensional data sets is that gener-

Fig. 4: Two libraries (black and red points) in 2 dimensions, and a target to unmix (green point). The optimal solution is indicated with a gray line. The RE is the perpendicular distance to this line. Left: The target and some library spectra are inside the convex hull, and the problem is hard to solve. Right: No point is inside the convex hull of any set. The optimal model is easy to determine.
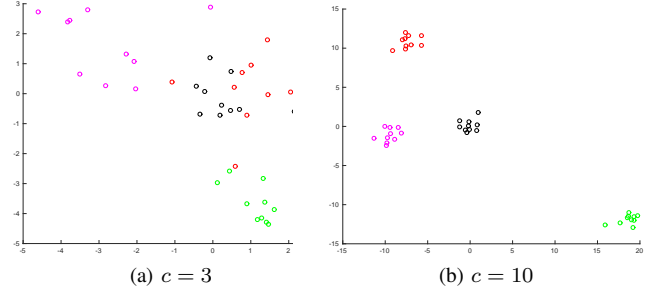

Fig. 5: Two artificial data sets in two dimensions. Three libraries are indicated with colored circles (red, green, cyan), while the data set is indicated with black circles. Two values of c are shown. For higher c, the libraries will show less correlation and overlap.

ally, every point will lie on the border of the convex hull, a consequence of the infamous "curse of dimensionality" [30]. For the AAM algorithm however, this will have a beneficial effect, as the situation depicted in Fig. 4 (a) will never occur in practical situations. While this is not a proof of why the AAM will always reach the global minimum (and counter-examples can be found in most data sets), this provides an intuitive explanation as to why the minimization functions better as the number of dimensions rises. We will show in the experimental section that the AAM algorithm indeed has a high dependence on the data dimensionality, and that the performance increases dramatically with increasing spectral dimension.

The computational complexity can be determined relatively easy. There are three nested for loops, where the first (line 2) will be executed $2^p - 1$ times, the second (line 6) K times, and the third (line 7) q times, with q the cardinality of the library subset. Because $\sum_{q=0}^{p} \binom{p}{q} q = p2^{p-1}$, the angle minimization (lines 8-14) will be executed $Kp2^{p-1}$ times in total, for each pixel.

The angle minimization involves the calculation of two projection matrices via pseudo-inversion of endmember matrices. Pseudo-inversion is typically solved via singular value decomposition, which has a worst-case behavior of d^3, with d the dimensionality of the vectors. The for-loop in line 10 can be avoided by employing matrix multiplication, which will yield an operation with complexity of order qdN_i. The maximal value for this term is $pdN_{\max}$, with $N_{\max} = \max_i\{N_i\}_i$. All other operations are of lower complexity. In total, this results in a highest-order and worst-case complexity of $NKp2^{p-1}(d^3 + pdN_{\max})$. Furthermore, the FCLSU algorithm will be executed $N(2^p - 1)$ times as well, each time to unmix a single pixel with respect to an endmember set with at most p endmembers.

This analysis shows that the proposed algorithm is linear in the number of pixels N and the library sizes N_i. The scaling with the number of libraries p is $p2^{p-1}$, which will increase quickly with increasing p. Remark that the MESMA algorithm scales as $\prod_i N_i$ with the library sizes, and also scales as 2^p with the number of libraries. A disadvantage of the AAM algorithm is that each pixel has to be processed concurrently, while the MESMA algorithm can treat an entire data set at once. Nevertheless, we will show in the experimental section that the AAM algorithm will be much faster than MESMA and other alternative algorithms in many practical scenarios, and

can treat situations which become computationally intractable with the alternatives.

IV. RESULTS

A. Artificial data sets

Several properties of the AAM algorithm can be tested and illustrated on artificial data sets, in order to be able to easily modify the data parameters. The data set consists of points in a $d = 200$ dimensional space. Each library is created as a d-dimensional random unit Gaussian distribution, centered around the point k_i. These centers $\{k_i\}_{i=1}^{p}$ are randomly selected from a Gaussian distribution with standard deviation c. The data set $\{x_i\}_{i=1}^{N}$ itself is selected from a unit Gaussian distribution around the origin. This way, the libraries and the data will show large overlaps (or correlation) for low values of c, and will be more independent for larger values of c. For simplicity, each library contained the same number of spectra. An illustration for 3 libraries of 10 endmembers in 2 dimensions is given in Fig. 5.

As a first experiment, we want to assess the evolution of the RE during the alternating angle minimization process. The RE is determined before each update in the AAM algorithm (i.e., between lines 7 and 8 in algorithm 2) for a given subset of libraries, and plotted in Fig. 6 as a function of the number of iterations. The number of libraries was $p = 10$, containing 10 endmember candidates each. As each iteration contains p updates, there are 10 data points in each interval.

This figure illustrates that in these artificial data sets, most of the decrease in RE happens during the first iteration, and that the RE reaches a stable minimum after 3 iterations for most points. Therefore, we set the number of iterations to $K = 3$ in all further experiments.

Next, we want to compare the results returned by the AAM algorithm with those obtained by MESMA on these artificial data sets. There are several ways to assess these differences between the unmixing results. First of all, both the AAM algorithm and MESMA return an endmember set for each pixel, which might show differences, and might even contain a different number of endmembers. Furthermore, after unmixing with these endmembers, the abundances with
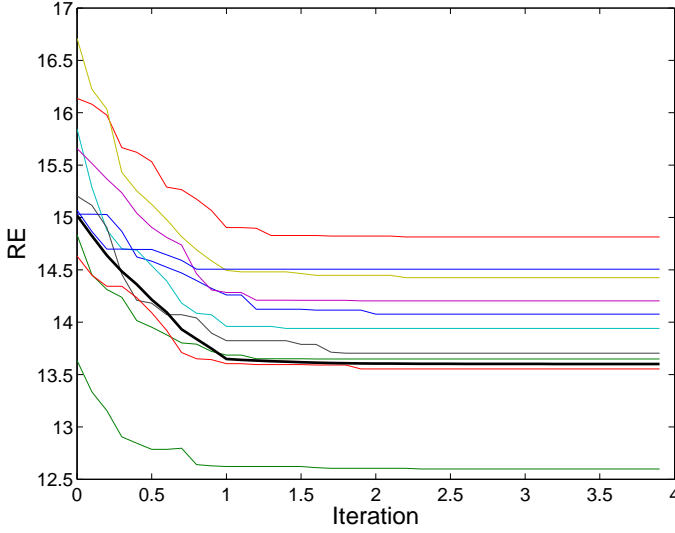
Fig. 6: The RE as a function of the number of iterations in the AAM algorithm, for 10 randomly selected data points (thin colored lines), and the average over 1000 random data points (black line).

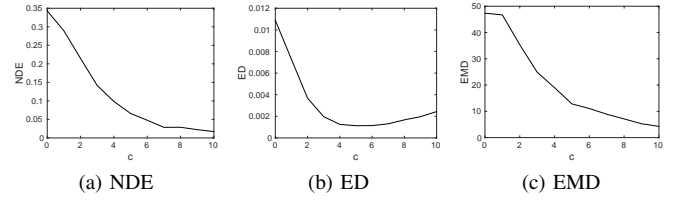(a) NDE   (b) ED   (c) EMD

Fig. 7: The NDE between the endmember sets, the ED between the abundance vectors, and the pixel-wise EMD as defined in [31], averaged over 100 random instances with 100 pixels each, with a given value of c.
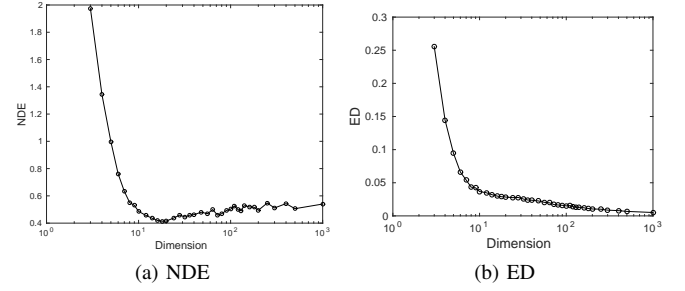

(a) NDE   (b) ED

Fig. 8: The NDE (a) and the ED (b) as a function of the data dimensionality d.

respect to each library are obtained. These can be compared as well, as different endmembers do not necessarily lead to large differences in the eventual unmixing results. Nevertheless, an approach which would be able to compare unmixing results simultaneously on both the endmember and the abundance level would be very advantageous.

Such a technique is proposed in [31], where the earth mover's distance (EMD) is used to assess differences between unmixing results. The main idea is to consider endmembers as bin centers in a high-dimensional space, and the corresponding abundances as bin values. With this interpretation, unmixing results can be represented as high-dimensional histograms, and histogram comparison measures such as the EMD can be employed to assess differences between unmixing results, even with different numbers of endmembers. Hence, to assess the per-pixel unmixing results, we employ three different measures:

- The number of different endmembers (NDE).
- The Euclidean distance (ED) between the abundance vectors

$$d(\boldsymbol{a}_1, \boldsymbol{a}_2) = \sqrt{(\boldsymbol{a}_1 - \boldsymbol{a}_2)^T(\boldsymbol{a}_1 - \boldsymbol{a}_2)} \qquad (10)$$

where we set the abundances of non-contributing libraries to zero.

- The EMD as proposed in [31], determined by the code distributed by the authors. The ground distance, the only parameter in the algorithm, is set to Euclidean.

These three measures are plotted as a function of the parameter c used to construct the data sets in Fig. 7. The displayed results are averaged over 100 random instances of 100 pixels each, with $p = 4$ libraries containing 10 members each. The dimensionality of the data was set to $d = 200$. These plots show that the NDE and the EMD decrease consistently as the endmember libraries become less related (i.e., have a

larger average distance from each other in spectral space). But even for $c = 0$, or a situation where all library members and data points are drawn from the same Gaussian distribution, the NDE graph shows that on average only 0.34 endmembers out of 4 differ between the AAM and the MESMA results. The ED between the abundance vectors is 0.011 in this case, which can be considered negligible for all practical scenarios. For larger values of c, this ED even decreases by an order of magnitude, and both the AAM and MESMA techniques will obtain the same unmixing results in virtually all cases. Note that the slight increase in ED for large values of c can be attributed to numerical reasons, as the magnitude of the endmembers scales with c.

As described in section III, the performance of the AAM algorithm depends strongly on the dimensionality d of the data set. To show this effect, we have repeated the previous experiment for several values of d, with $c = 0$ (i.e., all library and data points are drawn from the same standard Gaussian distribution). The ED between the abundance vectors and the NDE are displayed in Fig. 8, as a function of d, and averaged over 50 random data sets. We do not display the EMD as it depends strongly on d. These graphs show that the AAM algorithm functions very well for high d, but breaks down for lower dimensions.

In order to assess the computational runtimes, we have executed the AAM algorithm and MESMA on these artificial data sets, for varying numbers of libraries and library sizes. These runtimes are displayed in Fig. 9, and illustrate clearly that the computational cost rises quickly for MESMA as a function of the library sizes, but hardly increases for the AAM algorithm. For large libraries, the AAM algorithm will dramatically outperform MESMA, and this effect becomes more prominent as the number of libraries increases. For very
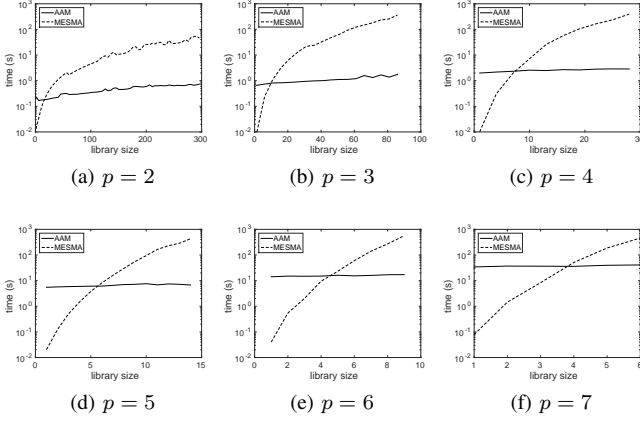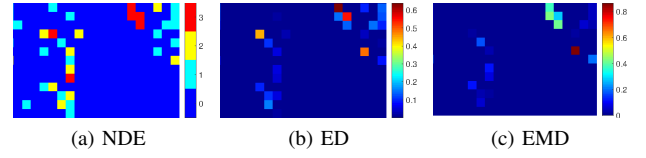
(a) NDE   (b) ED   (c) EMD

Fig. 11: Illustration of the differences between the MESMA and the AAM algorithm: The number of different endmembers (a), the ED between the abundance vectors (b) and the EMD between the unmixing results (c).

(a) $p = 2$ (b) $p = 3$ (c) $p = 4$

(d) $p = 5$ (e) $p = 6$ (f) $p = 7$

Fig. 9: The runtime in seconds of the AAM algorithm and MESMA as a function of library size, for different values of the number of libraries p. All libraries are equally sized.

small problems, MESMA would be the algorithm of choice.

B. Hyperspectral images

We have employed the proposed AAM algorithm to unmix several real-world data sets, and compare its performance with several alternative algorithms often used in the literature to deal with endmember variability. The algorithms used for comparison are the MESMA algorithm presented in section II, the normal compositional model (NCM) presented in [10], and the beta compositional model (BCM) presented in [11]. Both compositional models are solved with quadratic programming, and the spectral variant of the BCM is used. We refer to [32], [11] for more information on solving the NCM and BCM. Furthermore, we also provide the results that are obtained with FCLSU, where each library was replaced with its spectral average, and we provide known classification maps for illustrative purposes. Remark that there is no straightforward way of comparing abundance maps with classification maps, and classification ground truth can only serve as an illustrative example of where one might expect large abundances for a given class. This does not exclude other classes from contributing as well in the spectral mixture, often even with significant magnitude. Due to computational constraints, the compositional models are only executed on the data sets of limited size.

As the proposed AAM algorithm serves as an alternative for the MESMA algorithm, it should ideally return the same results. The alternatives based on the NCM and BCM allow comparisons with distribution-based techniques. We do not provide a detailed analysis of the relation between the abundance maps and possible classification ground truth, as the aim is to provide a fast alternative for MESMA, and not a new method for unmixing, nor a classification technique. We assess the differences between the MESMA and the AAM unmixing results by the NDE, the ED between the abundance vectors, and the EMD, represented as color maps.

The first hyperspectral image is also used in [11] to assess the BCM, and was collected in Long-beach, Missouri, USA.

This 13×19 pixel subset shows four different areas containing an asphalt road, a yellow curb next to the road, grass and oak trees. Ground truth spectra are available for each of these classes, obtained in the scene with a hand-held ASD spectrometer. The library sizes are 10 (asphalt), 10 (yellow curb), 50 (grass) and 10 (oak leaves). The number of spectral bands in both the image and the endmember libraries is 53.

The unmixing results obtained by the various methods are presented in Fig. 10, where each row corresponds with a class, and each column with an unmixing technique (or the ground truth). It can be observed that the AAM and the MESMA method both yield very similar results, and only differ in a handful of pixels. The NCM and the FCLSU method are very similar, while the BCM results show some differences with the NCM and FCLSU method. This similarity can be expected, as the spectral average employed in the FCLSU method will also serve as a highly probable candidate in the distribution-based approaches. One can conclude that the FCLSU and compositional models show similar results, while the library-based methods for endmember variability show a different behavior.

The differences between the results obtained by MESMA and the AAM algorithm are illustrated in Fig. 11, where the number of misidentified endmembers and the EDs between the abundance vectors are shown as color maps. It is clear that for the majority of pixels, identical endmembers and abundances are obtained. Possible deviations typically stay small.

Comparison of both maps indicates that one or two misidentified endmembers do not always automatically lead to large abundance errors. The endmembers obtained with the AAM algorithm can be very similar to those obtained with MESMA, or the corresponding abundance can be small. Both situations would lead to small abundance errors, even with misidentified endmembers.

The second data set is the well-known Pavia university data set, obtained over Pavia, Italy. The employed section of this data set has 510×200 pixels with 103 spectral bands, and has classification ground truth available. Four libraries were constructed, each containing 5 randomly selected spectra from the impervious classes, the meadows class, the trees class and the shadow class respectively. Using this library, the data set was unmixed with FCLSU, MESMA, and the AAM algorithm. We did not employ the compositional models as they required too much computational resources on data sets of this size.

The classification and abundance maps are displayed in Fig. 12. The maps obtained with MESMA and the AAM algorithm again show a very large similarity, while several notable differ-
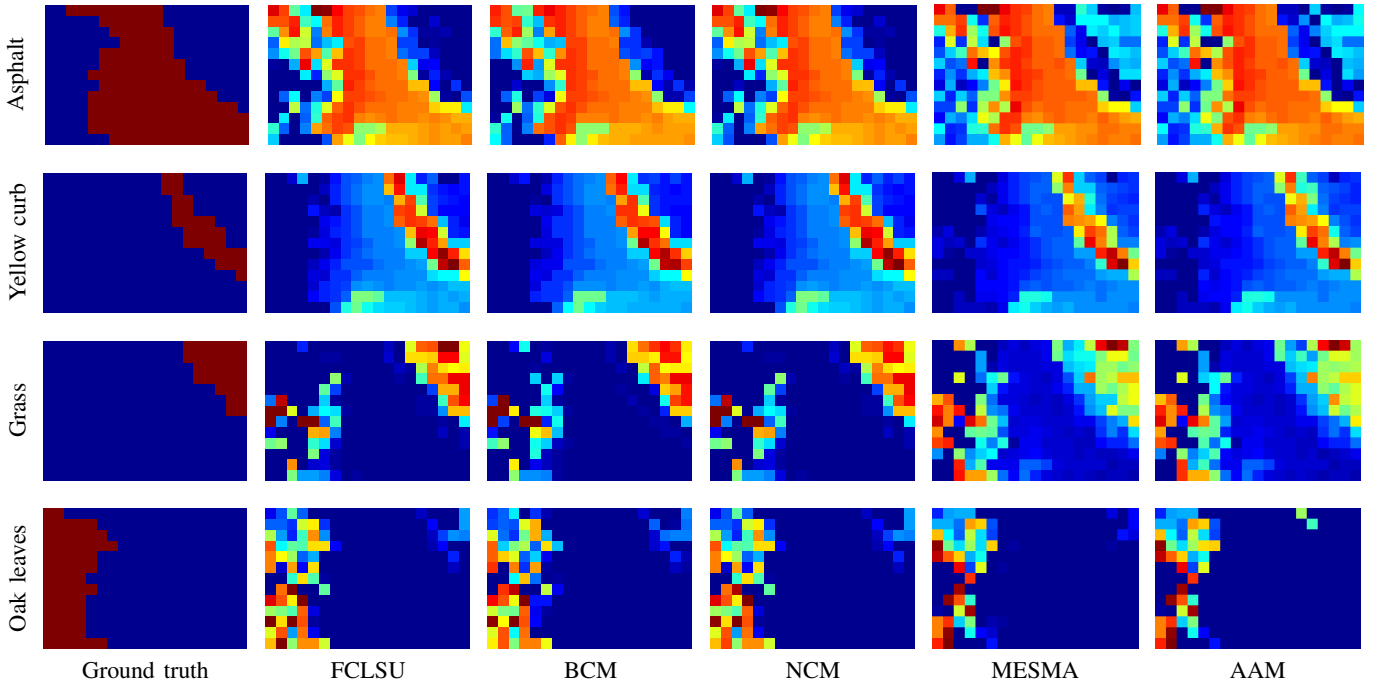
Fig. 10: Gulfport abundance maps

Algorithm	Long Beach	Pavia$_5$	Pavia$_{10}$	Pavia$_{15}$
FCLSU	0.17	86	88	85
BCM	91	N/A	N/A	N/A
NCM	12	N/A	N/A	N/A
MESMA	16	197	2207	9854
AAM	6.4	2080	2153	2232

TABLE I: The runtime in seconds of the algorithms for both data sets.

ences exist with the map obtained with FCLSU. The obtained unmixing results can again be assessed by considering the NDE, the ED and the EMD, which are displayed in Fig. 13 as colormaps. The NDE shows that for 69% of the pixels, the same endmembers are identified. The rest of the pixels have 1 (27%), 2 (3.5%) or 3 (0.4%) different endmembers. These are also the pixels where the ED and EMD will differ from zero.

The runtimes required for running the algorithms on these real hyperspectral data sets are shown in table I. For the Pavia data set, we included 3 versions, i.e., the version with 5 endmembers in each library used above, and one with 10 and 15 endmembers in each library, to illustrate the dependence on library size in real data sets. These are indicated as Pavia$_5$, Pavia$_{10}$ and Pavia$_{15}$ respectively. This table shows that FCLSU is the fastest algorithm for both data sets. On the Pavia data set, MESMA will be the fastest variant as long as the library sizes are less than 10, while AAM is faster for larger libraries. The runtime of MESMA also increases very fast for larger libraries, while AAM is much less sensitive to the library sizes. The compositional models could not be executed on the Pavia data set due to memory constraints.

V. CONCLUSIONS

We have introduced a new algorithm to deal with endmember variability in a constrained linear unmixing setting, based on alternating optimization with respect to each endmember library. The resulting AAM algorithm aims to solve the same optimization problem as the popular MESMA algorithm, but possesses a computational complexity which is linear in the size of the endmember libraries instead of multiplicative. This results in an algorithm which is capable of dealing with very large libraries as well without requiring large computational runtimes.

The AAM algorithm is illustrated on artificial and real data sets, and shown to yield very similar results as the MESMA algorithm in most cases, but with severely reduced runtimes. The measures employed for comparison of the unmixing results are the number of misidentified endmembers per pixel, the distances between the abundance vectors, and the EMD between the entire unmixing results, taking both endmember and abundance differences into account.

Future work concerns computational optimization of the current algorithm , and the inclusion of early stopping rules and sparsity, which will yield a further speedup, and might bring library-based methods for dealing with endmember variability into the reach of real-time or on-the-fly unmixing applications.

REFERENCES

[1] R. Heylen, M. Parente, and P. Gader, "A review of nonlinear hyperspectral unmixing methods," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 7, no. 6, pp. 1844–1868, June 2014.

[2] J.M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader, and J. Chanussot, "Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 5, no. 2, pp. 354–379, 2012.
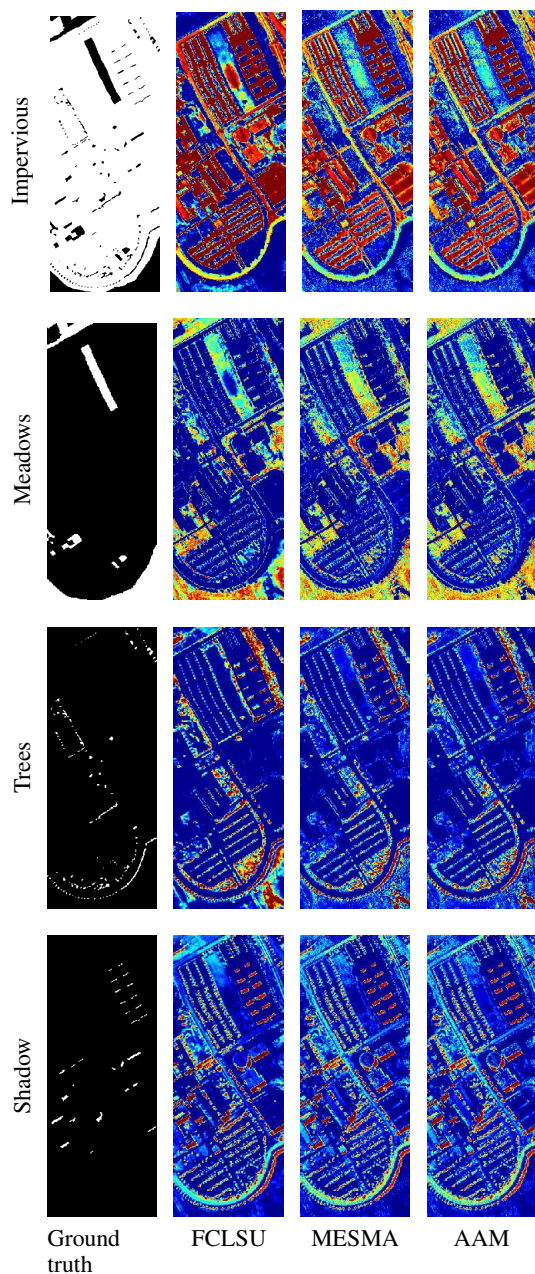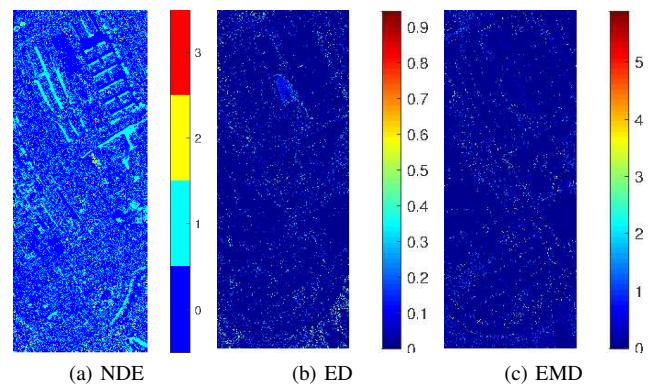
(a) NDE

(b) ED

(c) EMD

Fig. 13: Illustration of the differences between the MESMA and the AAM algorithm: The number of different endmembers (a), the ED between the abundance vectors (b) and the EMD between the unmixing results (c).

Ground truth — FCLSU — MESMA — AAM

Fig. 12: Pavia abundance maps

[3] N. Dobigeon, J.-Y. Tourneret, C. Richard, J.C.M. Bermudez, S. McLaughlin, and AO. Hero, "Nonlinear unmixing of hyperspectral images: Models and algorithms," *IEEE Sig. Process. Mag.*, vol. 31, no. 1, pp. 82–94, Jan 2014.

[4] John B. Adams, Donald E. Sabol, Valerie Kapos, Raimundo Almeida Filho, Dar A. Roberts, Milton O. Smith, and Alan R. Gillespie, "Classification of multispectral images based on fractions of endmembers: Application to land-cover change in the brazilian amazon," *Remote Sensing of Environment*, vol. 52, no. 2, pp. 137 – 154, 1995.

[5] M. Smith, D. Roberts, J. Hill, W. Mehl, B. Hosgood, J. Verdebout, G. Schmuck, C. Koechler, and J. Adams, "A new approach to quantifying abundances of materials in multispectral images," in *Proceedings of Geoscience and Remote Sensing Symposium IGARSS '94. Surface and Atmospheric Remote Sensing: Technologies, Data Analysis and Interpretation., International*, 1994, vol. 4, pp. 2372–2374.

[6] Philip E. Dennison, Kerry Q. Halligan, and Dar A. Roberts, "A comparison of error metrics and constraints for multiple endmember spectral mixture analysis and spectral angle mapper," *Remote Sensing of Environment*, vol. 93, pp. 359–367, 2004.

[7] F. J. Garcia-Haro, S. Sommer, and T. Kemper, "A new tool for variable multiple endmember spectral mixture analysis (vmesma)," *International Journal of Remote Sensing*, vol. 26, no. 10, pp. 2135–2162, May 2005.

[8] Ben Somers, Gregory P. Asner, Laurent Tits, and Pol Coppin, "Endmember variability in spectral mixture analysis: A review," *Remote Sensing Environ.*, vol. 115, no. 7, pp. 1603–1616, 2011.

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

[10] D. Stein, "Application of the normal compositional model to the analysis of hyperspectral imagery," *Advances in Techniques for Analysis of Remotely Sensed Data, 2003 IEEE Workshop on*, pp. 44–51, 2003.

[11] X. Du, A. Zare, P. Gader, and D. Dranishnikov, "Spatial and spectral unmixing using the beta compositional model," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 7, no. 6, pp. 1994–2003, June 2014.

[12] D.A. Roberts, M. Gardner, R. Church, S. Ustin, G. Scheer, and R.O. Green, "Mapping chaparral in the santa monica mountains using multiple endmember spectral mixture models," *Remote Sens. Environ.*, vol. 65, pp. 267–279, 1998.

[13] P. E. Dennison and D. A. Roberts, "Endmember selection for mapping chaparral species and fraction using multiple endmember spectral mixture analysis," *Remote Sensing Environ.*, vol. 41, pp. 123–135, 2003.

[14] K. N. Youngentob, D. A. Roberts, A. A. Held, P. E. Dennison, X. Jia, and D. B. Lindenmayer, "Mapping two eucalyptus subgenera using multiple endmember spectral mixture analysis and continuum-removed imaging spectrometry data," *Remote Sensing Environ.*, vol. 115, pp. 1115–1128, 2011.

[15] J. Franke, D. A. Roberts, K. Halligan, and G. Menz, "Hierarchical multiple endmember spectral mixture analysis (mesma) of hyperspectral imagery for urban environments," *Remote Sensing Environ.*, vol. 113, pp. 1712–1723, 2009.

[16] R. L. Powell, D. A. Roberts, P. E. Dennison, and L. L. Hess, "Sub-pixel mapping of urban land cover using multiple endmember spectral mixture analysis: Manaus, brazil," *Remote Sensing Environ.*, vol. 106, pp. 253–267, 2007.

[17] J. A. C. Ballantine, G. S. Okin, D. E. Prentiss, and D. A Roberts, "Mapping landforms in north africa using continental scale unmixing of modis imagery," *Remote Sensing Environ.*, vol. 97, pp. 470–483, 2005.

[18] D. A. Roberts, P. E. Dennison, M. Gardner, Y. Hetzel, S. L. Ustin, and C. Lee, "Evaluation of the potential of hyperion for fire danger assessment by comparison to the airborne visible/infrared imaging spectrometer," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, pp. 1297–1310, 2003.

[19] P. E. Dennison, K. Charoensiri, D. A. Roberts, S. H. Peterson, and R. O. Green, "Wildfire temperature and land cover modeling using hyperspectral data," *Remote Sensing Environ.*, vol. 100, pp. 212–222, 2006.

[20] T. C. Eckmann, D. A. Roberts, and C. J. Still, "Estimating subpixel fire sizes and temperatures from aster using multiple endmember spectral mixture analysis," *Int. J. Remote Sensing*, vol. 30, pp. 5851–5864, 2009.

[21] C. Quintano, A. Fernández-Manso, and D.A. Roberts, "Multiple

endmember spectral mixture analysis (mesma) to map burn severity levels from landsat images in mediterranean countries," *Remote Sensing Environ.*, vol. 136, pp. 76–88, 2013.

[22] P. Dennison, K. Halligan, and D. Roberts, "A comparison of error metrics and constraints for multiple endmember spectral mixture analysis and spectral angle mapper," *Remote Sensing Environ.*, vol. 93, pp. 359–367, Nov 2004.

[23] W. J. Okin, G. S. Okin, D. A. Roberts, and B. Murray, "Multiple endmember spectral mixture analysis: Endmember choice in an arid shrubland," *Proc. 5th AVIRIS Airborne Geosciences Workshop. JPL, CA*, 1998.

[24] P. Bajorski, "Simplex projection methods for selection of endmembers in hyperspectral imagery," *Proc. IEEE IGARSS*, vol. 5, pp. 3207–3210, 2004.

[25] R. Heylen and P. Scheunders, "Fully constrained least-squares spectral unmixing by simplex projection," *IEEE Trans. Geosci. Remote Sens.: Spec. iss. Remote Sensing*, vol. 49, no. 11, pp. 4112–4122, 2011.

[26] J. Loughry, J. I. Hemert, and L. Schoofs, "Efficiently enumerating the subsets of a set," *applied-math.org/subset.pdf*, Dec 2000.

[27] P. Honeine and C. Richard, "Geometric unmixing of large hyperspectral images: A barycentric coordinate approach," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 6, pp. 2185–2195, 2012.

[28] L. Tits, R. Heylen, B. Somers, P. Scheunders, and P. Coppin, "A geometric unmixing concept for the selection of optimal binary endmember combinations," *IEEE Geosci. Remote Sens. Lett.*, vol. 12, no. 1, pp. 82–86, Jan 2015.

[29] D.C. Heinz and C.-I Chang, "Fully constrained least squares linear spectral mixture analysis method for material quantification in hyperspectral imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 3, pp. 529–545, 2001.

[30] R. Heylen and P. Scheunders, "Multidimensional pixel purity index for convex hull estimation and endmember extraction," *IEEE Tran. Geosci. Remote Sens.*, vol. 51, no. 7, pp. 4059–4069, 2013.

[31] A. Zare and D. Anderson, "Earth movers distance-based simultaneous comparison of hyperspectral endmembers and proportions," *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 7, no. 6, pp. 1910–1921, June 2014.

[32] A. Zare, P. Gader, and G. Casella, "Sampling piecewise convex unmixing and endmember extraction," *IEEE Trans. Geosci. remote sens.*, vol. 51, no. 3, pp. 1655–1665, March 2013.