

Entropic Descent Archetypal Analysis for Blind Hyperspectral Unmixing

Alexandre Zouaoui, Gedeon Muhawenayo, Behnood Rasti, Jocelyn Chanussot,

Julien Mairal

▶ To cite this version:

Alexandre Zouaoui, Gedeon Muhawenayo, Behnood Rasti, Jocelyn Chanussot, Julien Mairal. Entropic Descent Archetypal Analysis for Blind Hyperspectral Unmixing. 2022. hal-03788427v1

HAL Id: hal-03788427 https://inria.hal.science/hal-03788427v1

Preprint submitted on 26 Sep 2022 (v1), last revised 12 Jan 2024 (v2)

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.

Entropic Descent Archetypal Analysis for Blind Hyperspectral Unmixing

Alexandre Zouaoui, Gedeon Muhawenayo, Behnood Rasti, Senior Member, IEEE Jocelyn Chanussot, Fellow, IEEE, and Julien Mairal, Senior Member, IEEE

Abstract—In this paper, we introduce a new algorithm based on archetypal analysis for blind hyperspectral unmixing, assuming linear mixing of endmembers. Archetypal analysis is a natural formulation for this task. This method does not require the presence of pure pixels (i.e., pixels containing a single material) but instead represents endmembers as convex combinations of a few pixels present in the original hyperspectral image. Our approach leverages an entropic gradient descent strategy, which (i) provides better solutions for hyperspectral unmixing than traditional archetypal analysis algorithms, and (ii) leads to efficient GPU implementations. Since running a single instance of our algorithm is fast, we also propose an ensembling mechanism along with an appropriate model selection procedure that make our method robust to hyper-parameter choices while keeping the computational complexity reasonable. By using six standard real datasets, we show that our approach outperforms stateof-the-art matrix factorization and recent deep learning methods. We also provide an open-source PyTorch implementation: https://github.com/inria-thoth/EDAA.

Index Terms—Hyperspectral image, remote sensing, blind spectral unmixing, non-negative matrix factorization, archetypal analysis, entropic gradient descent

I. INTRODUCTION

Hyperspectral imaging (HSI) [1]–[5] consists of measuring the electromagnetic spectrum in a scene by using multiple narrow spectral bands. Thanks to its richer spectral information compared to traditional RGB images, HSI enables more accurate materials identification, leading to a broad range of applications including crop monitoring in agriculture [6], waste sorting [7], food safety inspection [8], or mineralogy [9].

Remote sensing [10], [11], such as airborne or satellite imagery, yields hyperspectral images whose pixels capture several objects or materials. As such, each pixel can include several pure spectral components (called *endmembers*), mixed in different proportions [12]. Any further analysis hence requires identifying and disentangling endmembers present in a scene before estimating their respective proportions, or fractional *abundances*, within each pixel of the HSI [13]. Since the endmembers spectrum signatures are not known beforehand and must be estimated from data, this operation is named *blind* hyperspectral unmixing (HU) [14], [15] owing to its link with blind source separation [16].

In this paper, we adopt a linear mixing model since it is often relevant in remote sensing scenes where mixtures occur between macroscopic materials. Therefore, we assume that each observed pixel can be represented as a linear combination of endmembers and some additive noise. In other words, we are interested in tackling unsupervised linear HU [13].

Further assumptions on the nature of endmembers are generally needed to estimate meaningful spectra. For instance, it can be assumed that there exists at least one pure pixel for each material present in the scene. The problem then requires finding these pure pixels within the original image. The pure pixel assumption is at the core of several geometrical endmember extraction methods including pixel purity index (PPI) [17], N-FINDR [18] and vertex component analysis (VCA) [19]. Once endmembers have been extracted, abundances can be estimated by minimizing the least squares errors between the original input spectra and the linearly reconstructed spectra as long as the abundances fractions satisfy the two physical constraints stating that they should be non-negative and sum to one for each pixel [20]. That being said, pure pixels are often missing in real scenarios. In the absence of pure pixels and in the case of linear models, endmembers and abundances can be simultaneously estimated by solving a constrained or penalized non-negative matrix factorization problem (NMF) [21]. For example, the authors of [22] have proposed a formulation that involves a data fidelity term and a minimum volume regularization term on endmembers, whose minimization consists in alternating between solving for endmembers and abundances.

In this work, we do not assume the existence of pure pixels as they are often missing in real data, since, for instance, the spectral signatures of endmembers in HSI can be significantly affected by various changes in atmospheric, illumination, and environmental conditions within the scene [23]. We mitigate the effect of spectral variability by (i) normalizing each pixel by the ℓ_2 -norm of its spectrum as a pre-processing step and (ii) modeling endmembers as convex combinations of pixels present in the scene. Not only HSI pixels are linear combinations of the estimated endmembers under the linear mixing model, but the estimated endmembers are also convex combinations of pixels. This corresponds to the archetypal analysis (AA) formulation introduced by Cutler and Breiman in [24]. Compared to plain NMF, AA provides greater interpretability given that each endmember can be traced back to pixels present in the HSI. In addition, since the estimated endmembers spectral signatures generally correspond to averaging the contributions of several pixels, the resulting estimation appears to be more robust to noise and spectral variability than pure pixel methods.

Alexandre Zouaoui, Jocelyn Chanussot, and Julien Mairal are with Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France. Gedeon Muhawenayo worked on this project while he was there.

Behnood Rasti is with Helmholtz-Zentrum Dresden-Rossendorf (HZDR)

A. Contributions and novelties

- We propose a new hyperspectral unmixing algorithm relying on entropic gradient descent for archetypal analysis. Our approach (i) provides better solutions for hyperspectral unmixing than traditional alternating optimization schemes based on projected gradient methods or active set algorithms, and (ii) allows more efficient GPU implementations.
- 2) The efficiency of our method enables us to make a key practical contribution, consisting of an ensembling mechanism along with an appropriate model selection procedure, which makes our method almost parameterfree and thus easy to use (the only sensitive parameter is the number of endmembers we want to estimate).
- Our approach, available in an open-source package¹, outperforms state-of-the-art matrix factorization and deep learning methods on six standard real datasets.

The remainder of this paper is organized as follows. Section II presents related works for unsupervised linear hyperspectral unmixing. Section III introduces our method. Section IV presents experimental results highlighting the performance of our proposed approach. Finally, we conclude the article and underline future research directions in Section V.

II. RELATED WORK ON UNSUPERVISED LINEAR HYPERSPECTRAL UNMIXING

In this section, we present the framework of non-negative matrix factorization (NMF) before introducing archetypal analysis (AA). Finally, we mention widely used deep learning architectures that tackle blind hyperspectral unmixing.

a) Non-negative matrix factorization (NMF): NMF [21], [25] is popular for hyperspectral unmixing. It consists of factorizing a matrix representing the HSI signal into the product of two matrices with non-negative entries, where one factor carries the endmembers whereas the other one represents the abundances for each pixel (typically with the constraint that abundances sum to one). Several variants of NMF have been proposed for this task. For instance, the method of minimum volume constrained non-negative matrix factorization (MVC-NMF) of [26] uses a minimum volume term for endmembers that effectively gets rid of the pure pixel assumption. In minimum dispersion constrained NMF (MiniDisCo) [27], the regularization function is called *dispersion*, which encourages endmembers with minimum variance to avoid degenerate solutions and improve the unmixing robustness to flat spectra. Other approaches design regularization functions for abundances such as [28], [29].

b) Archetypal analysis (AA): AA was first introduced in [24] and consists in modeling individual pixels present in the HSI as a linear mixture of archetypes (here the endmembers). Interestingly, archetypes are defined as convex combinations of the individual pixels present in the HSI. Motivated by the good interpretability of AA, Zhao *et al.* [30] have proposed a kernelized variant, which enables greater modeling flexibility at the cost of an extra hyper-parameter (the bandwidth of the RBF kernel functions). Notably, they use the relaxation form introduced in [31] to handle the cases where endmembers are located outside of the convex hull of the data.

Inspired by the robust AA formulation introduced in [32] that considers the Huber loss instead of the squared loss to reduce the impact of noise and outliers, Sun *et al.* [33] have introduced a robust kernel archetypal analysis (RKADA) method for blind hyperspectral unmixing. Their approach refines the standard AA formulation by adding a binary sparse constraint on the pixels contributions. Consequently, their method ensures that each endmember actually corresponds to real pixels rather than a sparse linear mixture of all pixels. The resulting optimization algorithm is a block-coordinate descent scheme that uses an active-set method [34] to solve the smooth least squares optimization subproblems with a simplicial constraint. As noted by [32], the active-set algorithm can be seen as an aggressive strategy that can leverage the underlying sparsity of the subproblems solutions.

Recently, Xu *et al.* [35] have proposed an ℓ_1 sparsityconstrained AA algorithm to increase the sparsity of the abundances. Rather than using an active-set algorithm, they adopt a projected gradient method [36] to solve the resulting optimization subproblems inside an alternating scheme. However their formulation leads to abundances that do not sum to one, which reduces the physical interpretability of the unmixing results. Moreover, their approach relies on other algorithms including PPI [17] for pre-processing and initialization which increases the complexity of their algorithm.

c) Autoencoders (AE): Interest in deep learning has been rapidly growing in many fields including remote sensing [37] thanks to the increasing amount of available data, rising computational power, and the development of suitable algorithms. Autoencoders for hyperspectral unmixing have been widely used since the pioneering work of Licciardi et al. [38]. In a nutshell, the encoder transforms an input into abundances maps that are then linearly decoded by the decoder. The encoder activations correspond to the fractional abundances while the parameters of the linear decoder layer correspond to the endmembers spectra. While the AE described above corresponds to a basic fully connected architecture, refinements have been introduced such as using a loss functions that involves the spectral angle distance and a dedicated sparsity penalty term as in [39], or convolutions to take advantage of the spatial structure in HSI [40]. See [41] for a comprehensive study on the various existing autoencoder architectures for blind hyperspectral unmixing. Finally, it is worth mentioning the recent effort by Rasti et al. [42] to incorporate a geometrically motivated penalty into deep learning-based approaches. Their method, called minimum simplex convolutional network (MiSiCNet), uses both geometrical and spatial information, by means of a convolutional encoder-decoder architecture, to tackle unsupervised hyperspectral unmixing.

Although the surge of deep learning methods has led to improving overall unmixing performances, training a single network per image remains costly due to the required extensive hyperparameters search. Moreover, training such networks requires GPU, and yet it can be considerably slower than traditional methods that run on CPU.

III. METHODOLOGY

In this section, we present our model formulation before describing its optimization. Next, we mention implementation details required to run our approach. Finally, we explain how to leverage our efficient GPU implementation and propose a procedure to make our model robust to hyper-parameter choices and thus easy to use in practice.

A. Model formulation

Let **X** in $\mathbb{R}^{L \times N}$ be a hyperspectral image (HSI) where *L* is the number of channels and *N* is the number of pixels. We assume a linear mixing model such that

$$\mathbf{X} = \mathbf{E}\mathbf{A} + \mathbf{N},\tag{1}$$

where **E** in $\mathbb{R}^{L \times p}$ is the *mixing matrix* composed of the discretized spectra of p endmembers over L channels, and **A** in $\mathbb{R}^{p \times N}$ is the *abundance matrix* that describes, for each pixel, the fraction relative to each endmember. Finally, **N** represents some classical noise occuring in hyperspectral imaging.

We are interested in tackling the *blind* unmixing setting where both the mixing and abundance matrices are unknown. However, we assume that the number p of endmembers is known. Since **E** represents the endmembers reflectance over Lchannels, it follows that its elements should be non-negative, *i.e.*, **E** ≥ 0 . This is also the case for **A**, whose columns represent the abundances for each pixel, which in addition, should sum to one. In other words, each column of **A** = $[\mathbf{a}_1, \ldots, \mathbf{a}_N]$ belongs to the simplex Δ_p defined as:

$$\Delta_p \triangleq \left\{ \mathbf{a} \in \mathbb{R}^p \text{ s.t. } \mathbf{a} \ge 0 \text{ and } \sum_{j=1}^p \mathbf{a}_j = 1 \right\}.$$
 (2)

The previous model and constraint yields the classical optimization problem

$$\begin{array}{ll} \underset{\mathbf{E}, \mathbf{A}}{\operatorname{arg min}} & \frac{1}{2} \| \mathbf{X} - \mathbf{E} \mathbf{A} \|_{F}^{2} \\ \text{s.t.} & \mathbf{E} \geq 0, \\ & \mathbf{a}_{i} \in \Delta_{p} \text{ for } 1 \leq i \leq N, \end{array}$$
(3)

which is a variant of non-negative matrix factorization.

The archetypal analysis formulation we consider adds a constraint and forces the endmembers to be convex combinations of the pixels present in **X**. Formally, it simply means that there exists a matrix **B** in $\mathbb{R}^{N \times p}$ such that $\mathbf{E} = \mathbf{XB}$ and the columns of **B** are in the simplex Δ_N (their entries are non-negative and sum to one). This formulation yields the archetypal analysis formulation [24]:

$$\begin{array}{ll} \arg\min_{\mathbf{A}, \mathbf{B}} & \frac{1}{2} \| \mathbf{X} - \mathbf{X} \mathbf{B} \mathbf{A} \|_{F}^{2}, \\ \text{s.t.} & \mathbf{a}_{i} \in \Delta_{p} \text{ for } 1 \leq i \leq N, \\ & \mathbf{b}_{j} \in \Delta_{N} \text{ for } 1 \leq j \leq p \end{array}$$

$$(4)$$

where $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_N]$ and $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_p]$.

B. Optimization

Minimizing (4) is difficult since the objective function is not jointly convex in (\mathbf{A}, \mathbf{B}) . However, it is convex with respect to (w.r.t.) one of the variables when the other one is fixed [31]. Consequently, it is natural to consider an alternating minimization scheme between \mathbf{A} and \mathbf{B} , which is guaranteed to asymptotically provide a stationary point of the problem [43]. Yet, because of the non-convexity of the objective, the choice of optimization algorithm is important as different stationary points may not have the same quality in terms of statistical estimation. In other words, different optimization procedures may have a different "implicit bias", a phenomenon that is currently the focus of a lot of attention in machine learning [44], especially for deep learning models, suggesting that it may also be key for blind HU.

In this paper, we adopt an optimization scheme called entropic gradient descent which has shown better theoretical properties in terms of convergence rates than gradient descent when optimizing over the simplex [45]. Our second motivation was the simplicity of the algorithm, which does not require performing orthogonal projections on the simplex, nor dealing with complicated active-set rules as in [32]. This enables us to take advantage of modern GPUs. We now present mathematical details before discussing implementation.

1) Entropic descent algorithm (EDA): As noted in [45], the entropic descent algorithm is simply a gradient descent method with a particular choice of a Bregman-like distance [46] generated by a specific function, here the negative entropy. As explained in [47], the choice of an appropriate distance-like function tailored to the geometry of the constraints, here the simplex, provide theoretical benefits in terms of convergence rates. The minimization over the simplex is precisely the reason why we adopt the negative entropy to derive the updates of the alternating minimization scheme. Formally, the negative entropy function h is defined as follows: for x in \mathbb{R}^d ,

$$h(x) = \sum_{i=1}^{d} x_i \ln(x_i) \text{ if } x \in \Delta_d, \quad +\infty \text{ otherwise,} \quad (5)$$

with the convention that $0 \ln 0 \equiv 0$.

h exhibits several desirable properties, including convexity on Δ_d . This enables us to consider D_h , the Bregman divergence [46] w.r.t. *h*, defined, for *x* and *y* in \mathbb{R}^d

$$D_h(x,y) = h(x) - h(y) - \nabla h(y)^{\top}(x-y),$$
 (6)

which is also called the Kullback-Leibler divergence. By convexity of h, we naturally have $D_h(x, y) \ge 0$. We start by considering the following generic optimization problem:

$$\min_{\mathbf{z} \in \Delta_{+}} f(\mathbf{z}), \tag{7}$$

where f is a convex Lipschitz continuous function with a gradient at $\mathbf{z} \in \Delta_d$ denoted by $\nabla f(\mathbf{z})$ and Δ_d corresponds to the *d*-dimensional simplex (2).

The algorithm to solve (7) we consider performs the following iterates, given $\mathbf{z}^k \in \Delta_d$,

$$\mathbf{z}^{k+1} \leftarrow \arg\min_{\mathbf{z}\in\Delta_d} \left\{ \nabla f(\mathbf{z}^k)^\top (\mathbf{z} - \mathbf{z}^k) + \frac{1}{\eta_k} D_h(\mathbf{z}, \mathbf{z}^k) \right\}.$$
(8)

If D_h was simply a squared Euclidean norm, we would recover a projected gradient descent algorithm. By using instead the Bregman distance function D_h derived from the negative entropy, we obtain the entropic descent method. Here D_h measures the distance between two vectors in Δ_d . As such, the next iterate \mathbf{z}^{k+1} should aim for the optimal balance between taking a gradient step and moving the least from the current iterate \mathbf{z}^k according to the geometry induced by h, with η_k controlling this trade-off.

We will now see how the negative entropy h yields explicit steps that effectively enforce the simplicial constraints. Indeed, after a short classical calculation (see, for instance [45]), it is possible to show that the update (8) is equivalent to the following one: for all j in $\{1, \ldots, d\}$,

$$\mathbf{z}_{j}^{k+1} = \frac{\mathbf{z}_{j}^{k} e^{-\eta_{k} \nabla f(\mathbf{z}^{k})_{j}}}{\sum_{l=1}^{d} \mathbf{z}_{l}^{k} e^{-\eta_{k} \nabla f(\mathbf{z}^{k})_{l}}},\tag{9}$$

where \mathbf{z}_{j}^{k} is the *j*-th entry of the vector \mathbf{z}^{k} and similarly, $\nabla f(\mathbf{z}^{k})_{j}$ is the *j*-th entry of $\nabla f(\mathbf{z}^{k})$.

It is thus easy to see that the iterates \mathbf{z}^k stay in the simplex Δ_d , and it is possible to show (see [45]) that the sequence $(\mathbf{z}^k)_k$ converges to the set of solutions of (7) with the appropriate step-sizes η_k . Interestingly, the update (9) can be implemented efficiently by using the softmax function, assuming the entries of \mathbf{z}^k are positive:

$$\mathbf{z}^{k+1} = \operatorname{softmax} \left(\log(\mathbf{z}^k) - \eta_k \nabla f(\mathbf{z}^k) \right).$$
 (10)

where $log(\mathbf{z}^k)$ is the vector carrying the logarithm of each entry of \mathbf{z}^k . This update immediately suggests a high compatibility with GPUs.

2) Alternating optimization scheme: We are now in shape to describe an alternating optimization scheme, consisting of performing, alternatively, K_1 updates of EDA for minimizing A when B is fixed, and vice versa using K_2 updates. This strategy is presented in Algorithm 1. Formally, by replacing the generic function f with the functions corresponding to the two optimization subproblems, we obtain the following updates:

$$\mathbf{A}^{k+1} = \operatorname{softmax} \left(\log(\mathbf{A}^k) + \eta_k \mathbf{B}^\top \mathbf{X}^\top (\mathbf{X} - \mathbf{X} \mathbf{B} \mathbf{A}^k) \right),$$
(11)

$$\mathbf{B}^{k+1} = \operatorname{softmax} \left(\log(\mathbf{B}^k) + \eta_k \mathbf{X}^\top (\mathbf{X} - \mathbf{X} \mathbf{B}^k \mathbf{A}) \mathbf{A}^\top \right),$$
(12)

where $log(\mathbf{A}^k)$ is the matrix carrying the logarithm of each entry of \mathbf{A}^k while softmax is applied in parallel on the columns of $\log(\mathbf{A}^k) + \eta_k \mathbf{B}^\top \mathbf{X}^\top (\mathbf{X} - \mathbf{X} \mathbf{B} \mathbf{A}^k)$. Note that when A and B are initialized with positive values, these iterates keep them positive.

C. Implementation details

a) Normalization: The input image $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]$ is ℓ_2 -normalized for each pixel: for all *i* in $\{1, \ldots, N\}$,

$$\mathbf{x}_i \leftarrow \frac{\mathbf{x}_i}{||\mathbf{x}_i||_2},\tag{13}$$

where \mathbf{x}_i denotes the *i*-th pixel. This step is important to gain invariance to illumination changes.

Algorithm 1 Entropic Descent Archetypal Analysis (EDAA)

- 1: Input: ℓ_2 -normalized data X in $\mathbb{R}^{L \times N}$; p (number of endmembers); T (number of outer iterations); K_1 (number of inner iterations for \mathbf{A}); K_2 (number of inner iterations for $\hat{\mathbf{B}}$).
- 2: Initialize $\hat{\mathbf{A}} \in \mathbb{R}^{p \times N}$ using (14).
- 3: Initialize $\hat{\mathbf{B}} \in \mathbb{R}^{N \times p}$ using (15).
- 4: Set η_1 according to (16).
- 5: Set η_2 according to (17).
- 6: for t = 1, ..., T do for $k = 1, ..., K_1$ do 7:
 - $\hat{\mathbf{A}} \leftarrow \operatorname{softmax} \left(\log(\hat{\mathbf{A}}) + \eta_1 \hat{\mathbf{B}}^\top \mathbf{X}^\top (\mathbf{X} \mathbf{X} \hat{\mathbf{B}} \hat{\mathbf{A}}) \right)$
- 8: ▷ log is applied element-wise; 9:
- 10: \triangleright softmax is applied along the first dimension. end for 11:
- for $k = 1, ..., K_2$ do 12: $\hat{\mathbf{B}} \leftarrow \operatorname{softmax} \left(\log(\hat{\mathbf{B}}) + \eta_2 \mathbf{X}^{\top} (\mathbf{X} - \mathbf{X} \hat{\mathbf{B}} \hat{\mathbf{A}}) \hat{\mathbf{A}}^{\top} \right)$ 13: 14: end for 15: end for
- 16: $\hat{\mathbf{E}} \leftarrow \mathbf{X}\hat{\mathbf{B}}$ 17: Return E, A ▷ Estimated endmembers, abundances.

b) Initialization: We initialize the abundance matrix A uniformly,

$$\mathbf{A}^0 = \frac{1}{p} \mathbb{1}_p \mathbb{1}_N^T, \tag{14}$$

where $\mathbb{1}_d$ denotes a *d*-dimensional vector of ones. This corresponds to the maximal entropy configuration for each pixel. The entropy for each pixel will naturally decrease as a result of the optimization, but the high entropy of the initialization will have a regularization effect.

The initialization of the pixel contribution matrix **B** is then also close to uniform. Nevertheless we, introduce a small random perturbation which is necessary to break the symmetry between the columns of B (otherwise, the updates of A and B will keep them invariant). Concretely, the entries of B are randomly sampled according to the uniform distribution on [0,1], $\mathcal{U}_{[0,1]}$. Next, they are rescaled by a factor 0.1. Finally, we apply the softmax function on each column so that the columns of **B** belong to the simplex Δ_N , for j in $\{1, \ldots, p\}$,

$$\mathbf{b}_i^0 = \operatorname{softmax}(0.1 \ \mathbf{u}), \tag{15}$$

where $\mathbf{u} \sim \mathcal{U}_{[0,1]^N}$. In practice, we observe that such an initialization leads to a matrix \mathbf{B}^0 that is very close to an uniform initialization $\frac{1}{N} \mathbb{1}_N \mathbb{1}_p^T$.

c) Step sizes: We use constant step sizes η_1 and η_2 , for A and B respectively.

$$\eta_1 = \frac{\gamma}{\sigma_{\max}^2},\tag{16}$$

where γ is a value in $S = \{0.125, 0.25, 0.5, 1, 2, 4, 8\}$ and $\sigma_{\rm max}$ is the largest singular value of the matrix XB⁰. We recover the classical convergence of gradient descent with fixed step size [48] up to the factor γ , since σ_{\max}^2 corresponds to the Lipschitz constant of the sub-problem related to (4) when minimizing w.r.t A, B being fixed. Having γ in S allows us to use slightly different step sizes and yields better performance in practice. Note that our model selection procedure, presented later, will automatically choose the right parameter γ , thus removing the burden for the user of having to deal with an extra hyper-parameter. Finally, η_2 is simply a rescaled version of η_1 to account for the matrices being of transposed dimensions:

$$\eta_2 = \sqrt{\frac{p}{N}} \eta_1. \tag{17}$$

d) Hyperparameters: For all experiments, we set T =100 and $K_1 = K_2 = 5$ as it provides a good trade-off between convergence speed and unmixing accuracy. Note also that these hyper-parameters are robust to different real datasets as detailed in section IV.

D. Model selection procedure

As stated above, the initialization of the matrix **B** is random, leading to different solutions for each run of the algorithm since the overall optimization problem is non convex. Besides, we allow for different step-sizes γ , which we draw randomly from the set S in practice, Since the convergence of the algorithm is very fast (see experimental section for concrete benchmarks), we are able to provide a large diversity of solutions given a dataset by running M times our method with different random seeds, while keeping the global computational complexity reasonable. A major question we address next is then how to select optimally the best solution in terms of unmixing accuracy.

For this, we take inspiration from classical model selection and sparse estimation theory [49]. First, we measure the fit of each solution in terms of residual error $\|\mathbf{X} - \mathbf{X}\mathbf{A}\mathbf{B}\|_{1}$, choosing the ℓ_1 -norm which is known to be more robust to outliers than the mean squared error. Second, we select the set of solutions that are in the same ball park as the best solution we have obtained in terms of ℓ_1 fit. This selection process is illustrated by the red dotted line in Figure 1, while the precise criterion is described in Algorithm 2.

From the subset of solutions with good fit, we then choose the one whose endmembers have the best incoherence, a desired property, which is classical in the theory of sparse estimation [50]–[52]. Indeed, dictionaries (here endmembers) with more incoherence will benefit from better theoretical guarantees in terms of estimation of abundances, making it a natural criterion for model selection in the context of unmixing. Formally, the coherence is simply defined as the maximal pairwise spectral correlation between the estimated endmembers. More precisely, for $\mathbf{E} = [\hat{\mathbf{e}}_1, \dots, \hat{\mathbf{e}}_n]$ the endmembers matrix, the coherence μ is defined as:

$$\mu = \max_{k \neq k'} \langle \hat{\mathbf{e}}_k, \hat{\mathbf{e}}_{k'} \rangle. \tag{18}$$

To the best of our knowledge, this is the first time the coherence μ is used as a model selection criterion for archetypal analysis. Our experiments, see next section, show that it is highly effective.

In summary, we automatically select the model whose endmembers have the lowest maximal pairwise spectral correlation among the ones that have a good ℓ_1 fit. This strategy is illustrated in figure 1 and described in Algorithm 2. In the experiments, the number of runs M was set to 50.

Algorithm 2 Model Selection Procedure

- 1: Input: M (number of runs); ℓ_2 -normalized data X in $\mathbb{R}^{L \times N}$; p (number of endmembers); T (number of outer iterations); K_1 (number of inner iterations for $\hat{\mathbf{A}}$); K_2 (number of inner iterations for \mathbf{B}). 2: for m = 1, ..., M do
- Set random seed s_m . 3:
- $\hat{\mathbf{E}}_m, \hat{\mathbf{A}}_m \leftarrow \text{EDAA}(\mathbf{X}, p, T, K_1, K_2, s_m)$ ⊳ See (1) 4: $\operatorname{fit}_m \leftarrow ||\mathbf{X} - \mathbf{E}_m \mathbf{A}_m||_1$
- 5:
- Compute coherence μ_m on $\hat{\mathbf{E}}_m$. 6: \triangleright See (18)

7: end for

- 8: $\operatorname{fit}_{\min} \leftarrow \min{\operatorname{fit}_1, \ldots, \operatorname{fit}_M}$
- 9: $\mathcal{I} \leftarrow \{m \mid \text{fit}_m \leq 1.05 \times \text{fit}_{\min}\}$ ▷ Subset of models.
- 10: best $\leftarrow \arg\min\{\mu_i, i \in \mathcal{I}\}$
- 11: Return: $\hat{\mathbf{E}}_{\text{best}}, \hat{\mathbf{A}}_{\text{best}}.$

IV. EXPERIMENTS

We have performed experiments on six standard real datasets whose descriptions are given below.

A. Hyperspectral data description

- 1) Samson: The Samson² hyperspectral image is a 95x95 pixels sub-region of a larger image captured using 156 bands spanning from 401 to 889 nm. Three main materials have been identified: Tree, Soil and Water. Note that we use a different ground truth from [42] that we selected for its sharper details on the abundances.
- 2) Jasper Ridge: The Jasper Ridge²hyperspectral image is a 100x100 pixels sub-region of a larger image initially captured using 224 bands spanning from 380 to 2500 nm. In total, 198 bands remain as 26 were removed as a pre-processing step due to dense water vapor and atmospheric effects. Four main materials have been identified: Tree, Dirt, Water and Road.
- 3) Urban4 and Urban6: The Urban²hyperspectral image is a 307x307 pixels image collected by the Hyperspectral Digital Image Collection Experiment (HYDICE) [53] sensor using 210 bands spanning from 400 to 2500 nm. In total, 162 bands remain as 48 were removed as a pre-processing step due to dense water vapor and atmospheric effects. There exists three versions of this dataset w.r.t. the number of endmembers. In this study, we focus on the two extremes: Urban4 contains 4 endmembers (Asphalt Road, Grass, Tree and Roof) and Urban6 contains two additional materials: Dirt and Metal, making it more challenging.
- 4) APEX: The APEX [54] hyperspectral image that we consider in this paper is a 111x122 pixels cropped region³ of a larger image captured over 285 bands

²downloaded at https://rslab.ut.ac.ir/data



Fig. 1: Illustration of the model selection procedure on three datasets using M = 50 runs. Runs are illustrated by blue dots and the selected one is in black. The selected run is the one with lowest coherence μ under the dashed red line representing the ℓ_1 fit threshold, see Alg. 2.

spanning from 413 to 2420nm. Four main materials were identified: Road, Tree, Roof and Water.

5) Washington DC Mall: The Washington DC Mall (WDC) hyperspectral dataset³consists in a 319x292 pixels image captured by the HYDICE [53] sensor over 191 bands spanning from 400 to 2400 nm. Six main materials were identified: Grass, Tree, Roof, Road, Water and Trail.

According to [55] (Samson, Jasper Ridge and Urban) and [42] (APEX and WDC), the endmembers spectra were manually selected from the images and the ground truth abundances were set by FCLSU. Illustrations of the datasets and their ground truth endmembers are available in the supplementary material.

B. Experimental setup

We compare our approach to five competitive methods from different unmixing categories:

- Geometrical unmixing baseline: FCLSU [20] using VCA [19] to extract endmembers. Our implementation of the FCLSU algorithm uses the *DecompSimplex* routine implemented in SPAMS⁴. This method relies on the active-set algorithm [34] that enables significantly faster convergence than generic quadratic programming solvers by leveraging the underlying sparsity of the abundances as noted by [32].
- Deep learning unmixing: Endnet⁵ [39] using VCA [19] to initialize the endmembers and MiSiCNet⁶ [42].
- NMF-based blind unmixing: non-negative matrix factorization quadratic minimum volume (NMF-QMV)⁷ [22] using the *boundary* term as the quadratic minimum volume penalty and AA [24] using the implementation from [32] developed in SPAMS.

To quantitatively evaluate the performance of the selected methods, we consider two metrics that are computed both globally and individually for each endmember. On one hand, we measure the quality of the generated abundances by means of the abundances root mean square error (RMSE) in percent between the ground truth and the estimated abundances:

$$\text{RMSE}(\mathbf{A}, \hat{\mathbf{A}}) = 100 \times \sqrt{\frac{1}{pN} \sum_{i=1}^{p} \sum_{j=1}^{N} \left(\mathbf{A}_{i,j} - \hat{\mathbf{A}}_{i,j}\right)^2}.$$
 (19)

On the other hand, we assess the quality of the estimated endmembers spectra by using the spectral angle distance (SAD) in degrees between the ground truth and the generated endmembers:

$$SAD(\mathbf{E}, \hat{\mathbf{E}}) = \frac{180}{\pi} \times \frac{1}{p} \sum_{i=1}^{p} \arccos\left(\frac{\langle \mathbf{e}_i, \hat{\mathbf{e}}_i \rangle}{||\mathbf{e}_i||_2 ||\hat{\mathbf{e}}_i||_2}\right), \quad (20)$$

where $\langle . \rangle$ denotes the inner product and \mathbf{e}_i denotes the *i*-th column of \mathbf{E} , *i.e.* the spectrum of the *i*-th endmember.

C. Unmixing experiments

Tables I and II report the unmixing accuracy in terms of abundances RMSE (19) and endmembers SAD (20) on six standard real datasets. The datasets are arbitrarily ranked based on their difficulty. For a fair comparison, all methods were evaluated on the ℓ_2 -normalized data (13) which induces slight changes compared to the results reported in [42].

Overall EDAA obtains the best abundances RMSE on all six datasets. As the unmixing difficulty increases, the gap between EDAA and the other methods grows larger. In particular, plain AA is not able to tackle more complicated mixing scenarios as in Urban6, APEX and WDC. We argue that our model selection technique is instrumental in avoiding collapsing runs in which endmembers spectra are highly correlated. This is underlined by the overall competitive SAD results obtained by EDAA across datasets, especially on Urban6, APEX and WDC where Endnet is the only contender.

The FCLSU baseline based on VCA obtains rather poor results except for WDC. This is likely due to the pure pixel

⁴http://thoth.inrialpes.fr/people/mairal/spams/

⁵implementation at https://github.com/burknipalsson/hu_autoencoders

⁶implementation at https://github.com/BehnoodRasti/MiSiCNet

⁷implementation at https://github.com/LinaZhuang/NMF-QMV_demo

		FCLSU [19], [20]	Endnet [39]	MiSiCNet [42]	NMF-QMV [22]	AA	EDAA
	Soil	11.28	11.61	6.47	13.67	6.16	5.74
Samson	Tree	9.13	7.72	5.38	8.4	4.00	3.77
	Water	5.05	6.86	3.47	11.61	2.30	<u>2.59</u>
	Overall	8.88	8.97	5.25	11.44	4.44	4.24
	Dirt	21.23	18.26	21.68	19.97	10.24	7.32
	Road	24.72	29.40	24.94	26.13	<u>9.79</u>	7.61
Jasper Ridge	Tree	11.20	<u>4.00</u>	3.41	14.55	6.32	6.63
	Water	13.61	22.38	7.07	19.81	<u>6.77</u>	5.69
	Overall	18.52	20.70	16.98	20.53	<u>8.46</u>	6.85
	Road	30.54	12.04	10.30	20.25	11.12	8.62
	Grass	32.99	17.94	12.35	20.22	<u>11.55</u>	9.28
Urban4	Roof	15.40	11.28	8.01	13.29	7.28	6.37
	Tree	20.02	11.69	8.78	21.56	8.62	6.27
	Overall	25.78	13.52	10.00	19.11	<u>9.80</u>	7.75
	Road	31.61	17.44	19.18	24.81	19.23	11.39
	Grass	23.62	31.47	18.84	27.97	9.73	18.61
	Roof	13.00	9.35	7.41	16.39	10.92	6.01
Urban6	Tree	16.14	15.50	<u>11.72</u>	19.95	14.45	9.85
	Dirt	25.06	30.67	<u>23.95</u>	24.31	27.76	15.94
	Metal	12.99	25.21	33.67	13.40	17.53	15.86
	Overall	21.54	23.09	<u>16.27</u>	21.74	17.66	13.63
	Road	33.31	29.32	32.66	31.43	34.12	16.54
APEX	Tree	20.97	18.18	19.97	24.62	22.49	14.48
	Roof	<u>14.15</u>	15.88	18.42	14.73	18.54	11.27
	Water	18.03	17.47	16.88	17.99	16.90	16.83
	Overall	22.77	<u>20.90</u>	22.86	23.10	23.98	14.94
WDC	Grass	30.90	27.35	31.61	34.69	31.92	32.53
	Tree	22.42	35.98	23.64	<u>19.90</u>	20.13	11.46
	Road	27.90	38.49	34.91	22.49	38.87	13.97
	Roof	8.71	27.04	<u>11.98</u>	19.81	20.89	29.31
	Water	17.76	12.94	14.72	20.34	14.06	9.63
	Trail	12.80	12.63	12.07	12.36	15.24	13.19
-	Overall	21.57	27.64	23.39	22.60	25.17	20.46

TABLE I: Abundances RMSE on six real datasets. The best results are shown in bold. The second best results are underlined.



Fig. 2: Estimated endmembers on the Jasper Ridge dataset. Ground truth endmembers are displayed in blue while their estimates are in orange.

assumption. Indeed, VCA selects a single pixel to represent the endmembers spectra, which is too stringent in real scenarios where spectral variability is ubiquitous.

Despite its quadratic minimum volume *boundary* term, NMF-QMV generally obtains worse results than the FCLSU baseline. Since it operates the unmixing in a subspace, NMF-QMV cannot prevent the endmembers spectra from having negative values, which breaks the physical interpretability of the estimates and subsequently harms the unmixing performance. This phenomenon can be observed in figures 2, 4, 6 and 8 for several endmembers. The associated abundances in figures 3, 5, 7 and 9 show that NMF-QMV produces maps that are too uniform and lack sparsity.

In contrast, Endnet achieves very good results in terms of SAD but tends to create overly sparse abundances which hinders its performance in terms of abundances RMSE. For instance, as can be seen in figures 8 and 9, the Road endmember is overlooked by Endnet even though EDAA recovers it neatly. Likewise, in figure 7, the Road endmember spreads too much compared to EDAA which appears closer to the ground truth.

MiSiCNet gives better unmixing results than Endnet in terms of abundances RMSE except for APEX although the

		FCLSU [19], [20]	Endnet [39]	MiSiCNet [42]	NMF-QMV [22]	AA	EDAA
	Soil	2.76	0.61	1.21	4.90	0.78	1.64
Samson	Tree	3.05	<u>1.93</u>	3.38	5.34	1.80	1.98
	Water	7.15	1.48	5.36	11.14	<u>1.38</u>	1.31
	Overall	4.32	<u>1.34</u>	3.32	7.13	1.32	1.64
	Dirt	13.03	1.63	4.26	12.40	2.44	2.74
	Road	40.39	32.85	20.04	45.66	8.00	3.10
Jasper Ridge	Tree	11.16	1.39	1.27	14.46	4.34	4.23
	Water	13.24	<u>3.21</u>	4.18	14.53	3.29	2.80
	Overall	19.46	9.77	7.44	21.77	4.52	3.22
	Road	15.40	6.40	5.73	14.51	3.73	6.01
	Grass	24.18	3.09	5.84	16.39	1.81	2.14
Urban4	Roof	47.56	3.76	16.10	36.31	15.59	10.49
	Tree	19.82	2.32	4.60	22.48	3.51	<u>2.81</u>
	Overall	26.74	3.89	8.07	22.42	6.16	<u>5.36</u>
	Road	13.43	3.26	7.47	26.60	6.74	4.85
	Grass	22.30	4.13	10.97	21.63	6.82	2.17
	Roof	15.65	17.76	13.97	15.55	18.07	13.70
Urban6	Tree	20.70	7.72	9.99	23.31	7.18	8.92
	Dirt	69.81	17.42	19.57	23.60	11.08	13.33
	Metal	39.35	<u>7.04</u>	9.75	68.73	40.97	4.52
	Overall	30.21	<u>9.56</u>	11.95	29.90	15.14	8.64
	Road	40.23	14.46	33.02	54.53	37.32	6.83
	Tree	14.13	7.53	<u>3.16</u>	16.06	2.39	7.68
APEX	Roof	8.25	4.36	11.31	7.98	10.20	<u>7.50</u>
	Water	7.15	2.83	6.02	9.71	2.95	2.21
	Overall	17.44	7.30	13.38	22.07	13.21	6.06
WDC	Grass	17.40	3.54	17.46	34.36	4.59	8.11
	Tree	23.73	12.85	12.36	17.70	<u>10.93</u>	1.81
	Road	32.56	26.76	33.20	17.28	46.73	7.67
	Roof	34.84	13.70	28.87	44.87	43.26	50.97
	Water	4.78	1.75	1.57	19.74	1.23	1.08
	Trail	9.94	1.49	3.24	9.60	5.32	4.75
	Overall	20.54	10.02	16.11	23.93	18.68	12.40

TABLE II: Endmembers SAD on six real datasets. The best results are shown in bold. The second best results are underlined.

SAD results falls in favor of Endnet except for Jasper Ridge. This is likely due to Endnet using the spectral angle distance on the input data in its loss which helps in achieving better SAD accuracy. However a good SAD is not sufficient to obtain good abundance maps, an area where MiSiCNet tends to shine as it incorporates spatial information by using convolutional filters and implicitly applying a regularizer on abundances.

Finally, AA is a very competitive method for the three simplest datasets yet its performance decreases drastically when dealing with more complicated mixtures. For example, in figures 2 and 3 only AA and EDAA are able to uncover the Road endmember in Jasper Ridge whereas all the other techniques fail. Yet, in figures 4 and 5, AA completely misses the Metal and Dirt endmembers in Urban6 while EDAA correctly identify all endmembers and produce reasonable abundance maps. Unlike EDAA, the random matrices initialization in AA cannot be leveraged to create an appropriate model selection procedure due to its slowness.

Additional qualitative results for the Samson and Urban4 datasets can be found in the supplementary material.

D. Computational cost

Table III reports the processing times for the different unmixing algorithms on the six real datasets. NMF-QMV was implemented in Matlab (2020b) while FCLSU, Endnet, MiSiCNet and the AA variants were implemented in Python (3.8). NMF-QMV, FCLSU and AA run on CPU whereas Endnet, MiSiCNet and EDAA run on GPU. The processing times were obtained using a computer with an Intel(R) Xeon(R) Silver 4110 processor (2.10 GHz), 32 cores, 64GB of memory, and a NVIDIA GeForce RTX (2080 Ti) graphical processing unit. The table shows that FCLSU is clearly faster than the other unmixing techniques, however it is a supervised method that relies on an endmembers extraction algorithm. In this case, VCA is used which is also fast. The deep learning methods are the slowest techniques despite running on GPU. Interestingly, EDAA requires a lower computational cost than NMF-QMV and AA although our approach consists in aggregating 50 runs obtained iteratively. For example, it takes on average 1.5 seconds for EDAA to perform a single unmixing task on the Urban6 dataset, which is three times faster than FCLSU. This demonstrates the efficiency of EDAA which allows us to use an adequate model selection procedure over several runs.

TABLE III: Processing times in seconds on six real datasets. The best results are in bold and the second best underlined. EDAA includes the model selection procedure over M = 50 runs.

	FCLSU	Endnet	MiSiCNet	NMF-QMV	AA	EDAA
Samson	0.3	≈ 560	≈ 80	20.1	17.6	11.2
JasperRidge	0.4	≈ 680	≈ 90	22.3	21.9	<u>9.6</u>
Urban4	3.6	≈ 720	≈ 411	112.5	237.6	66.3
Urban6	4.4	≈ 1000	≈ 417	158.4	343.0	75.2
APEX	0.6	≈ 720	≈ 92	27.2	44.8	15.4
WDC	4.0	≈ 1000	≈ 409	174.4	361.8	81.0

9

NMF-QMV AA EDAA GΤ FCLSU Endnet MiSiCNet Tree Water 1.00 0.75 0.50 0.25 0.00 Dirt Road

Fig. 3: Estimated abundances on the Jasper Ridge dataset. The ground truth is displayed on the right side.



Fig. 4: Estimated endmembers on the Urban6 dataset. Ground truth endmembers are displayed in blue while their estimates are in orange.

E. Ablation study

Finally, we study the sensitivity to hyper-parameters for Algorithm 1 and 2 in figure 10 where the Y-axis corresponds to the overall abundances RMSE. Given a fixed computational budget of 1000 updates, figure 10 (a) shows that the hyperparameters of EDAA are robust provided that the number of runs M in the model selection is large enough (here 100). Only the two extremes configurations ($K_1 = K_2 = 1$, T = 500 and $K_1 = K_2 = 50$, T = 10) are slightly worse, especially on Urban6. For the remaining experiments, we use $K_1 = K_2 = 5$. In figure 10 (b), we see that the number of outer iterations is quite stable except for WDC which requires more updates (1000, *i.e.* T = 100). Finally, we study the importance of the number of runs M from which to select the best candidate in figure 10 (c). We observe that the model selection procedure requires at least 50 runs to obtain very good performances, hence we use M = 50 in our unmixing experiments. On unknown datasets where real-time unmixing is not required, it is advised to use a large number of runs (at least 100) to ensure that the model selection procedure yields a good candidate. Detailed results for both abundances RMSE and SAD metrics are available in the supplementary material.

V. CONCLUSION

We have proposed a new algorithm based on archetypal analysis for blind hyperspectral unmixing. We have shown how to take advantage of its efficient GPU implementation in order to develop an adequate model selection procedure to obtain state-of-the-art performances. Remarkably, our simple and easy-to-use approach considerably improves the unmixing



Fig. 5: Estimated abundances on the Urban6 dataset. The ground truth is displayed on the right side.

results on a comprehensive collection of standard real datasets. Finally, we have made our results reproducible by releasing an open source codebase which also includes the plain archetypal analysis variant presented in this study.

Acknowledgments and Funding

This project was supported by the ERC grant number 714381 (SOLARIS project) and by ANR 3IA MIAI@Grenoble Alpes (ANR-19-P3IA-0003).

References

- D. Landgrebe, "Hyperspectral image data analysis," *IEEE Signal Processing Magazine*, vol. 19, no. 1, pp. 17–28, 2002.
- [2] A. Plaza, J. A. Benediktsson, J. W. Boardman, J. Brazile, L. Bruzzone, G. Camps-Valls, J. Chanussot, M. Fauvel, P. Gamba, A. Gualtieri, and others, "Recent advances in techniques for hyperspectral image processing," *Remote Sensing of Environment*, vol. 113, pp. 110–122, 2009. 1
- [3] M. E. Schaepman, S. L. Ustin, A. J. Plaza, T. H. Painter, J. Verrelst, and S. Liang, "Earth system science related imaging spectroscopy—An assessment," *Remote Sensing of Environment*, vol. 113, pp. 123–137, 2009. 1
- [4] A. F. Goetz, G. Vane, J. E. Solomon, and B. N. Rock, "Imaging spectrometry for earth remote sensing," *Science*, vol. 228, no. 4704, pp. 1147–1153, 1985. 1

- [5] R. O. Green, M. L. Eastwood, C. M. Sarture, T. G. Chrien, M. Aronsson, B. J. Chippendale, J. A. Faust, B. E. Pavri, C. J. Chovit, M. Solis, and others, "Imaging spectroscopy and the airborne visible/infrared imaging spectrometer (AVIRIS)," *Remote Sensing of Environment*, vol. 65, no. 3, pp. 227–248, 1998. 1
- [6] T. Adão, J. Hruška, L. Pádua, J. Bessa, E. Peres, R. Morais, and J. J. Sousa, "Hyperspectral imaging: A review on UAV-based sensors, data processing and applications for agriculture and forestry," *Remote sensing*, vol. 9, no. 11, p. 1110, 2017. 1
- [7] A. C. Karaca, A. Ertürk, M. K. Güllü, M. Elmas, and S. Ertürk, "Automatic waste sorting using shortwave infrared hyperspectral imaging system," in Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS), 2013, pp. 1–4.
- [8] A. A. Gowen, C. P. O'Donnell, P. J. Cullen, G. Downey, and J. M. Frias, "Hyperspectral imaging-an emerging process analytical tool for food quality and safety control," *Trends in food science & technology*, vol. 18, no. 12, pp. 590–598, 2007. 1
- [9] N. Fox, A. Parbhakar-Fox, J. Moltzen, S. Feig, K. Goemann, and J. Huntington, "Applications of hyperspectral mineralogy for geoenvironmental characterisation," *Minerals Engineering*, vol. 107, pp. 63–77, 2017. 1
- [10] R. N. Clark, G. A. Swayze, K. E. Livo, R. F. Kokaly, S. J. Sutley, J. B. Dalton, R. R. McDougal, and C. A. Gent, "Imaging spectroscopy: Earth and planetary remote sensing with the USGS Tetracorder and expert systems," *Journal of Geophysical Research: Planets*, vol. 108, no. E12, 2003. 1
- [11] J. M. Bioucas-Dias, A. Plaza, G. Camps-Valls, P. Scheunders, N. 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, 2013. 1



Fig. 6: Estimated endmembers on the APEX dataset. Ground truth endmembers are displayed in blue while their estimates are in orange.



Fig. 7: Estimated abundances on the APEX dataset. The ground truth is displayed on the right side.

- [12] P. Ghamisi, N. Yokoya, J. Li, W. Liao, S. Liu, J. Plaza, B. Rasti, and A. Plaza, "Advances in hyperspectral image and signal processing: A comprehensive overview of the state of the art," *IEEE Geoscience and Remote Sensing Magazine*, vol. 5, no. 4, pp. 37–78, 2017. 1
- [13] L. Parra, C. Spence, P. Sajda, A. Ziehe, and K.-R. Müller, "Unmixing hyperspectral data," Adv. in Neural Information Processing Systems (NIPS), vol. 12, 1999. 1
- [14] N. Keshava and J. F. Mustard, "Spectral unmixing," *IEEE Signal Processing Magazine*, vol. 19, no. 1, pp. 44–57, 2002.
- [15] 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 Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 5, no. 2, pp. 354–379, 2012. 1
- [16] P. Comon and C. Jutten, Handbook of Blind Source Separation: Independent component analysis and applications, 2010. 1
- [17] J. W. Boardman, F. A. Kruse, and R. O. Green, "Mapping target signatures via partial unmixing of AVIRIS data," 1995. 1, 2
- [18] M. E. Winter, "N-FINDR: An algorithm for fast autonomous spectral end-member determination in hyperspectral data," in *Imaging Spectrom*-

etry V, vol. 3753, 1999, pp. 266-275. 1

- [19] J. M. Nascimento and J. M. 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, 2005. 1, 6, 7, 8
- [20] D. C. Heinz and others, "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, 2001. 1, 6, 7, 8
- [21] D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," Adv. in Neural Information Processing Systems (NIPS), vol. 13, 2000. 1, 2
- [22] L. Zhuang, C.-H. Lin, M. A. Figueiredo, and J. M. Bioucas-Dias, "Regularization parameter selection in minimum volume hyperspectral unmixing," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 12, pp. 9858–9877, 2019. 1, 6, 7, 8
- [23] R. A. Borsoi, T. Imbiriba, J. C. M. Bermudez, C. Richard, J. Chanussot, L. Drumetz, J.-Y. Tourneret, A. Zare, and C. Jutten, "Spectral variability in hyperspectral data unmixing: A comprehensive review," *IEEE Geoscience and Remote Sensing Magazine*, vol. 9, no. 4, pp. 223–270, 2021.

12



Fig. 8: Estimated endmembers on the Washington DC dataset. Ground truth endmembers are displayed in blue while their estimates are in orange.

- [24] A. Cutler and L. Breiman, "Archetypal analysis," *Technometrics*, vol. 36, no. 4, pp. 338–347, 1994. 1, 2, 3, 6
- [25] P. Paatero and U. Tapper, "Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values," *Environmetrics*, vol. 5, no. 2, pp. 111–126, 1994. 2
- [26] L. Miao and H. Qi, "Endmember extraction from highly mixed data using minimum volume constrained nonnegative matrix factorization," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 3, pp. 765–777, 2007. 2
- [27] A. Huck, M. Guillaume, and J. Blanc-Talon, "Minimum dispersion constrained nonnegative matrix factorization to unmix hyperspectral data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 48, no. 6, pp. 2590–2602, 2010. 2
- [28] A. Zymnis, S.-J. Kim, J. Skaf, M. Parente, and S. Boyd, "Hyperspectral image unmixing via alternating projected subgradients," in *Asilomar Conference on Signals, Systems and Computers*, 2007, pp. 1164–1168.
- [29] Z. Yang, G. Zhou, S. Xie, S. Ding, J.-M. Yang, and J. Zhang, "Blind spectral unmixing based on sparse nonnegative matrix factorization," *IEEE Transactions on Image Processing*, vol. 20, no. 4, pp. 1112–1125, 2010. 2
- [30] G. Zhao, X. Jia, and C. Zhao, "Multiple endmembers based unmixing using archetypal analysis," in *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2015, pp. 5039–5042. 2
- [31] M. Mørup and L. K. Hansen, "Archetypal analysis for machine learning and data mining," *Neurocomputing*, vol. 80, pp. 54–63, 2012. 2, 3
- [32] Y. Chen, J. Mairal, and Z. Harchaoui, "Fast and robust archetypal analysis for representation learning," in *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), 2014. 2, 3, 6
- [33] W. Sun, G. Yang, K. Wu, W. Li, and D. Zhang, "Pure endmember extraction using robust kernel archetypoid analysis for hyperspectral imagery," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 131, pp. 147–159, 2017. 2
- [34] J. Nocedal and S. J. Wright, Numerical optimization, 1999. 2, 6
- [35] M. Xu, Z. Yang, G. Ren, H. Sheng, S. Liu, W. Liu, and C. Ye, "L1 Sparsity-Constrained Archetypal Analysis Algorithm for Hyperspectral Unmixing," *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2022. 2
- [36] C.-J. Lin, "Projected gradient methods for nonnegative matrix factorization," *Neural computation*, vol. 19, no. 10, pp. 2756–2779, 2007. 2
- [37] L. Zhang, L. Zhang, and B. Du, "Deep learning for remote sensing data: A technical tutorial on the state of the art," *IEEE Geoscience and Remote Sensing Magazine*, vol. 4, no. 2, pp. 22–40, 2016. 2
- [38] G. Licciardi, X. Ceamanos, S. Douté, and J. Chanussot, "Unsupervised nonlinear spectral unmixing by means of NLPCA applied to hyperspectral imagery," in *IEEE International Geoscience and Remote Sensing Symposium*, 2012, pp. 1369–1372. 2
- [39] S. Ozkan, B. Kaya, and G. B. Akar, "Endnet: Sparse autoencoder network for endmember extraction and hyperspectral unmixing," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 1, pp. 482–496, 2018. 2, 6, 7, 8
- [40] B. Palsson, M. O. Ulfarsson, and J. R. Sveinsson, "Convolutional autoencoder for spectral-spatial hyperspectral unmixing," *IEEE Transactions*

on Geoscience and Remote Sensing, vol. 59, no. 1, pp. 535–549, 2020.

- [41] B. Palsson, J. R. Sveinsson, and M. O. Ulfarsson, "Blind hyperspectral unmixing using autoencoders: A critical comparison," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 1340–1372, 2022. 2
- [42] B. Rasti, B. Koirala, P. Scheunders, and J. Chanussot, "Misicnet: Minimum simplex convolutional network for deep hyperspectral unmixing," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1– 15, 2022. 2, 5, 6, 7, 8
- [43] D. P. Bertsekas, "Nonlinear programming," Journal of the Operational Research Society, vol. 48, no. 3, pp. 334–334, 1997. 3
- [44] S. Pesme, L. Pillaud-Vivien, and N. Flammarion, "Implicit bias of sgd for diagonal linear networks: a provable benefit of stochasticity," Adv. in Neural Information Processing Systems (NeurIPS), 2021. 3
- [45] A. Beck and M. Teboulle, "Mirror descent and nonlinear projected subgradient methods for convex optimization," *Operations Research Letters*, vol. 31, no. 3, pp. 167–175, 2003. 3, 4
- [46] L. M. Bregman, "The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming," USSR computational mathematics and mathematical physics, vol. 7, no. 3, pp. 200–217, 1967. 3
- [47] M. Teboulle, "Entropic proximal mappings with applications to nonlinear programming," *Mathematics of Operations Research*, vol. 17, no. 3, pp. 670–690, 1992. 3
- [48] Y. Nesterov, Introductory lectures on convex optimization: A basic course, 2003, vol. 87. 4
- [49] T. Hastie, R. Tibshirani, and J. H. Friedman, *The elements of statistical learning: data mining, inference, and prediction*, 2009, vol. 2. 5
- [50] M. Elad and A. M. Bruckstein, "A generalized uncertainty principle and sparse representation in pairs of bases," *IEEE Transactions on Information Theory*, vol. 48, no. 9, pp. 2558–2567, 2002. 5
- [51] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Transactions on Information Theory*, vol. 49, no. 12, pp. 3320–3325, 2003. 5
- [52] J. Mairal, F. Bach, J. Ponce et al., "Sparse modeling for image and vision processing," Foundations and Trends® in Computer Graphics and Vision, vol. 8, no. 2-3, pp. 85–283, 2014. 5
- [53] L. J. Rickard, R. W. Basedow, E. F. Zalewski, P. R. Silverglate, and M. Landers, "HYDICE: An airborne system for hyperspectral imaging," in *Imaging Spectrometry of the Terrestrial Environment*, vol. 1937, 1993, pp. 173–179. 5, 6
- [54] M. E. Schaepman, M. Jehle, A. Hueni, P. D'Odorico, A. Damm, J. Weyermann, F. D. Schneider, V. Laurent, C. Popp, F. C. Seidel, and others, "Advanced radiometry measurements and Earth science applications with the Airborne Prism Experiment (APEX)," *Remote Sensing of Environment*, vol. 158, pp. 207–219, 2015. 5
- [55] F. Zhu, "Hyperspectral unmixing: ground truth labeling, datasets, benchmark performances and survey," arXiv preprint arXiv:1708.05125, 2017.

13



Fig. 9: Estimated abundances on the Washington DC dataset. The ground truth is displayed on the right side.



Fig. 10: Sensitivity analysis to the hyperparameters in Algorithms 1 and 2 measured in global abundances RMSE: (a) Varying inner and outer iterations K_1, K_2 and T for a constant number of updates (1000) and runs M = 100, (b) Varying outer iterations T using $K_1 = K_2 = 5$ and (c) Varying number of runs M using $K_1 = K_2 = 5$ and T = 100.

SUPPLEMENTARY MATERIAL A DETAILED DATASETS DESCRIPTION

In this section, we provide illustrations of the unmixing datasets used in the experiments. Each dataset is described with a false-color RGB image alongside the ℓ_2 -normalized ground truth endmembers.



Fig. 11: Samson dataset: (a) False colors RGB image (Red: 83rd band, Green: 43, Blue: 9) (b) ℓ_2 -normalized ground truth endmembers.



Fig. 12: JasperRidge dataset: (a) False colors RGB image (Red: 130th band, Green: 50, Blue: 5) (b) ℓ_2 -normalized ground truth endmembers.



Fig. 13: Urban dataset: (a) False colors RGB image (Red: 130th band, Green: 70, Blue: 30) (b) ℓ_2 -normalized ground truth endmembers for Urban6. Urban4 corresponds to the first 4 materials.

SUPPLEMENTARY MATERIAL B Additional results

a) Unmixing experiments: We provide qualitative results on the Samson and Urban4 datasets in figures 16, 17, 18 and 19.



Fig. 14: APEX dataset: (a) False colors RGB image (Red: 200th band, Green: 100, Blue: 10) (b) ℓ_2 -normalized ground truth endmembers (#0: Road, #1: Tree, #2: Roof, #3: Water).



Fig. 15: Washington DC Mall dataset: (a) False colors RGB image (Red: 150th band, Green: 75, Blue: 20) (b) ℓ_2 -normalized ground truth endmembers (#0: Grass, #1: Tree, #2: Road, #3: Roof, #4: Water, #5: Trail).

b) Ablation study: We report the detailed results obtained in the ablation study. For each dataset, the overall abundances RMSE and SAD are computed for all configurations. Table IV studies the sensitivity of the inner and outer iterations K_1 , K_2 and T in EDAA given a fixed computational budget. Table V studies the sensitivity of the outer iterations T in EDAA when we decrease the computational budget. Finally, table VI underlines the importance of the number of runs M in the model selection procedure.



Fig. 16: Estimated endmembers on the Samson dataset. Ground truth endmembers are displayed in blue while their estimates are in orange.



Fig. 17: Estimated abundances on the Samson dataset. The ground truth is displayed on the right side.



Fig. 18: Estimated endmembers on the Urban4 dataset. Ground truth endmembers are displayed in blue while their estimates are in orange.



Fig. 19: Estimated abundances on the Urban4 dataset. The ground truth is displayed on the right side.

TABLE IV: Sensitivity to hyperparameters of EDAA for a constant number of updates (1000). The abundances RMSE and SAD metrics are computed globally. The best results are in bold and the second best are underlined.

		$K_1 = K_2 = 1, T = 500$	$K_1 = K_2 = 5, T = 100$	$K_1 = K_2 = 10, T = 50$	$K_1 = K_2 = 20, T = 25$	$K_1 = K_2 = 50, T = 10$
Samoon	RMSE	4.42	4.24	4.15	4.48	4.34
Samson	SAD	1.65	1.64	1.61	1.78	1.68
Jaaman Didaa	RMSE	6.70	6.85	7.80	7.59	8.79
Jasper Kluge	SAD	3.48	3.22	4.28	3.12	4.79
	RMSE	7.40	7.51	7.43	7.90	9.49
UIDall4	SAD	5.72	5.87	5.53	6.05	5.46
Urban6	RMSE	17.92	13.52	12.35	12.71	17.92
UIDallo	SAD	11.71	7.95	8.74	8.52	8.79
ADEV	RMSE	14.25	14.66	14.20	14.15	16.46
AFEA	SAD	8.17	6.29	7.73	8.06	5.26
WDC	RMSE	22.91	20.47	22.16	20.92	21.20
WDC	SAD	<u>12.59</u>	12.33	16.05	13.46	13.51

		T = 100	T = 75	T = 50
Samoon	RMSE	4.24	4.34	3.97
Samson	SAD	<u>1.64</u>	1.69	1.46
Jacpar Didga	RMSE	6.85	5.90	6.44
Jasper Kluge	SAD	3.22	3.06	<u>3.19</u>
I Juhon 4	RMSE	7.51	7.46	7.72
Urban4	SAD	5.87	<u>5.01</u>	4.67
Urbanh	RMSE	13.52	13.54	13.99
Urbano	SAD	7.95	7.85	<u>7.93</u>
ADEV	RMSE	14.66	14.38	14.12
AFEA	SAD	6.29	<u>6.58</u>	6.99
WDC	RMSE	20.47	22.81	22.89
WDC	SAD	12.33	16.97	16.87

TABLE V: Sensitivity to the number of outer iterations T of Algorithm EDAA with $K_1 = K_2 = 5$. The abundances RMSE and SAD metrics are computed globally. The best results are in bold and the second best are underlined.

TABLE VI: Sensitivity to the number of runs M of the model selection procedure with $K_1 = K_2 = 5$ and T = 100. The abundances RMSE and SAD metrics are computed globally. The best results are in bold and the second best are underlined.

		M = 10	M = 20	M = 50	M = 100
Compon	RMSE	5.05	3.90	4.24	4.24
Samson	SAD	1.94	1.43	1.64	1.64
Jaaman Didaa	RMSE	10.27	10.27	6.85	6.85
Jasper Kluge	SAD	3.73	3.73	3.22	3.22
Urban4	RMSE	8.17	8.06	7.75	7.51
	SAD	5.43	<u>5.41</u>	5.36	5.87
Urban6	RMSE	16.50	16.50	13.63	13.52
	SAD	8.16	8.16	7.92	<u>7.95</u>
ADEV	RMSE	23.84	24.17	14.94	14.66
APEA	SAD	12.57	12.88	6.06	<u>6.29</u>
WDC	RMSE	24.39	24.39	20.46	20.47
wDC	SAD	11.52	11.52	12.40	12.33