

SUnGP: A greedy sparse approximation algorithm for hyperspectral unmixing

Naveed Akhtar, Faisal Shafait and Ajmal Mian

School of Computer Science
and Software Engineering

The University of Western Australia
WA, 35 Stirling Highway, Crawley, 6009

Email: naveed.akhtar@research.uwa.edu.au, faisal.shafait@uwa.edu.au, ajmal.mian@uwa.edu.au

Abstract—Spectra measured at a pixel of a remote sensing hyperspectral sensor is usually a mixture of multiple spectra (endmembers) of different materials on the ground. Hyperspectral unmixing aims at identifying the endmembers and their proportions (fractional abundances) in the mixed pixels. Hyperspectral unmixing has recently been casted into a sparse approximation problem and greedy sparse approximation approaches are considered desirable for solving it. However, the high correlation among the spectra of different materials seriously affects the accuracy of the greedy algorithms. We propose a greedy sparse approximation algorithm, called SUnGP, for unmixing of hyperspectral data. SUnGP shows high robustness against the correlation of the spectra of materials. The algorithm employs a subspace pruning strategy for the identification of the endmembers. Experiments show that the proposed algorithm not only outperforms the state of the art greedy algorithms, its accuracy is comparable to the algorithms based on the convex relaxation of the problem, but with a considerable computational advantage.

I. INTRODUCTION

Modern spaceborne and airborne hyperspectral sensors measure the reflectance of the Earth's surface at hundreds of contiguous narrow bands [1], which results in hyperspectral image cubes with two spatial and one spectral dimension (see Fig. 1). Each pixel of a hyperspectral image is a vector that represents a spectral signature measured by the sensor. Due to low spatial resolution of the sensors and multiple scatterings, the spectra at a pixel is usually a mixture of multiple pure spectra (endmembers), corresponding to different materials on the ground. Hyperspectral unmixing aims at identifying the endmembers in a mixed pixel and computing their fractional abundances (i.e. their proportion in the pixel) [2]. Unmixing of the hyperspectral images is considered a major challenge in remote sensing data analysis [3].

Recently, Linear Mixing Model (LMM) has gained considerable attention for hyperspectral unmixing [2]. This model assumes that a mixed pixel is a linear combination of its constituent endmembers, weighted by their fractional abundances. Works that employ LMM, often use the geometric properties of hyperspectral data in unmixing. They exploit the fact that the convex hull of the pure endmembers in the data forms a simplex. Thus, finding the endmembers simplifies to finding the vertices of a simplex. Classical geometric approaches assume the presence of pure pixels for each endmember in the image. Vertex Component Analysis [4], N-FINDER [5], Pixel Purity Index [6] and Simplex Growing Algorithm [7] are some of the popular examples of such approaches.




Fig. 1: Illustration of a hyperspectral image cube: The XY-plane shows the spatial dimensions and the wavelength axis shows the spectral dimension. Pixels are recorded as vectors of reflectance values at different wavelengths. The cube shows reflectance patterns at ten wavelength bands. The data is collected by NASA's AVIRIS [12] over the Cuprite mines, NV.

In real world data, pure pixels are not usually present for each endmember in the image. Therefore, some approaches focus on extracting the pure spectral signatures from the image for hyperspectral unmixing. Minimum Volume Simplex Analysis [8], Iterative Constrained Endmembers (ICE) [9] and Sparsity Promoting ICE [10] are examples of such approaches. Extraction of pure spectra from the images fails in highly mixed scenarios. In this case, the above mentioned approaches generate artificial endmembers that cannot be associated with true materials [11]. For highly mixed scenarios, hyperspectral unmixing is often formulated as a statistical inferencing problem under the Bayesian paradigm [3]. However, the statistical inferencing process of the Bayesian approach is generally computationally expensive.

To overcome the above issues, hyperspectral unmixing has recently been formulated as a sparse approximation problem [13]. This approach assumes that a mixed pixel can be approximated by a sparse linear combination of pure spectra, already available in a dictionary. Iordache et al. [13] have analyzed different sparse approximation algorithms for hyperspectral unmixing, including the greedy algorithms (e.g., Orthogonal Matching Pursuit (OMP) [14]). Generally, the greedy algorithms are computationally more efficient than their convex relaxation based counterparts [16] and admit to simpler and faster implementations [17]. However, [13] shows that in comparison to the convex relaxation based algorithms (e.g., Basis Pursuit [15]), the accuracy of the greedy algorithms is severely affected by the high correlation of the spectra.

Shi et al. [18] argue strongly to exploit the greedy sparse approximation technique for hyperspectral unmixing because of its computational advantages. The authors have proposed a greedy algorithm for hyperspectral unmixing, called Simultaneous Matching Pursuit (SMP). SMP mitigates the problems caused by the high correlation of the spectra by processing the images in terms of spatial patches. The contextual information in a patch is exploited in identifying the endmembers. However, the image patch size becomes an important parameter for SMP and this parameter is image dependent. Furthermore, SMP assumes the existence of the contextual information in the complete image, which compromises its performance in the highly mixed scenarios, where this information is scarce.

In this work, we propose a greedy sparse approximation algorithm for hyperspectral unmixing. The algorithm performs Sparse Unmixing via the Greedy Pursuit strategy [14], hence it is named SUnGP. SUnGP approximates the mixed pixel by iteratively identifying its endmembers from a fixed dictionary. In each iteration, it selects a subspace that is further pruned to identify the endmember. The pruning strategy helps SUnGP in mitigating the adverse effects of the high correlation of the spectra. SUnGP is a pixel-based greedy algorithm, therefore its performance does not degrade in the highly mixed scenarios. Additionally, SUnGP can be modified to take advantage of the contextual information in the image, following the guidelines in [19]. Experiments with synthetic and real world hyperspectral data show that the proposed algorithm outperforms the existing state of the art greedy algorithms. Moreover, its results are comparable to the convex relaxation based sparse approximation algorithms, with a considerable computational advantage. In this work, we preprocess the hyperspectral data using spectral derivatives [29]. This results in better performance of greedy pursuit algorithms in hyperspectral unmixing.

II. HYPERSPECTRAL UNMIXING AS A SPARSE APPROXIMATION PROBLEM

A. Linear Mixing Model

Sparse approximation can exploit the Linear Mixing Model (LMM) of the pixels in hyperspectral unmixing. This model assumes that the reflectance y_i , measured at the i^{th} band of a mixed pixel, is a linear combination of the endmember reflectances at that band. Mathematically,

$$y_i = \sum_{j=1}^p l_{ij} \alpha_j + \epsilon_i \quad (1)$$

where, p is the total number of the endmembers in the pixel, l_{ij} is the reflectance of the j^{th} endmember at band i , α_j is the fractional abundance of the j^{th} endmember and ϵ_i represents the noise. If the image is acquired by a sensor with m spectral channels, we can write the LMM in the matrix form:

$$\mathbf{y} = \mathbf{L}\boldsymbol{\alpha} + \boldsymbol{\epsilon} \quad (2)$$

where, $\mathbf{y} \in \mathbb{R}^m$ represents the reflectance vector of the pixel, $\mathbf{L} \in \mathbb{R}^{m \times p}$ is a matrix with p endmembers, $\boldsymbol{\alpha} \in \mathbb{R}^p$ contains the corresponding fractional abundances of the endmembers and $\boldsymbol{\epsilon} \in \mathbb{R}^m$ represents the noise. The fractional abundances follow two constraints in LMM [3]. (1) **ANC: Abundance Non-negativity Constraint** ($\forall i, i \in \{1, \dots, p\}, \alpha_i > 0$) and (2) **ASC: Abundance Sum-to-one Constraint** ($\sum_{i=1}^p \alpha_i = 1$).

B. Sparse approximation of a mixed pixel

Let $\mathbf{D} \in \mathbb{R}^{m \times k}$ ($k > m$) be a matrix with each column $\mathbf{d}_i \in \mathbb{R}^m$ representing the spectral signature of a material. According to LMM, if \mathbf{D} contains a large collection of spectra, including the endmembers of the mixed pixel, then:

$$\mathbf{y} \approx \mathbf{D}\boldsymbol{\alpha} \quad (3)$$

where, $\boldsymbol{\alpha} \in \mathbb{R}^k$ has only p non-zero coefficients. Generally, in the remote sensing hyperspectral images, a pixel is a mixture of four to five spectra [13]. Therefore, it is safe to assume that $\boldsymbol{\alpha}$ is sparse ($p \ll k$). This fact allows us to formulate hyperspectral unmixing as the following sparse approximation problem:

$$(P_0^\eta) : \min_{\boldsymbol{\alpha}} \|\boldsymbol{\alpha}\|_0 \quad \text{s.t.} \quad \|\mathbf{D}\boldsymbol{\alpha} - \mathbf{y}\|_2 \leq \eta \quad (4)$$

where, $\|\cdot\|_0$ is the l_0 pseudo-norm that simply counts the number of non-zero elements in $\boldsymbol{\alpha}$, and η is the tolerance due to noise. In the sparse approximation literature, \mathbf{D} is known as the *dictionary* and its columns are called the *atoms*.

Minimization of l_0 pseudo-norm in (P_0^η) is, in general, an NP-hard problem [20]. In practice, its polynomial time approximation is achieved with the greedy algorithms. Relaxed convexification of the problem (P_0^η) is also possible. It is done by replacing l_0 pseudo-norm in (4) with the l_1 norm of $\boldsymbol{\alpha}$. Let us denote this version of the problem as (P_1^η) . Solving (P_1^η) is equivalent to solving the well known LASSO problem [21] with an appropriate Lagrangian multiplier λ [22]. The LASSO formulation of the problem is given below as (P_1^λ) :

$$(P_1^\lambda) : \min_{\boldsymbol{\alpha}} \frac{1}{2} \|\mathbf{y} - \mathbf{D}\boldsymbol{\alpha}\|_2 + \lambda \|\boldsymbol{\alpha}\|_1 \quad (5)$$

Previous works in sparse unmixing mainly focus on solving (P_1^η) and (P_1^λ) [18]. For instance, Sparse Unmixing by Splitting and Augmented Lagrangian (SUnSAL) [22] solves (P_1^λ) for sparse unmixing. The authors have also enhanced this algorithm to its constrained version, called CSUnSAL. CSUnSAL solves (P_1^η) with ASC as an additional constraint. Iordache et al. [13] have also used SUnSAL+ and CSUnSAL+ for sparse unmixing. These algorithms solve the corresponding problems by further imposing ANC on the solution space.

III. GREEDY SPARSE APPROXIMATION

Greedy algorithms approximate the signal \mathbf{y} in (P_0^η) by iteratively selecting the atoms of the dictionary. These atoms are selected such that the signal is approximated in minimum number of iterations. This greedy heuristic finally results in a sparse solution. We can unify the greedy sparse approximation algorithms under a base-line algorithm, which can be stated as the following sequential steps: 1) *Identification* of the atom(s) of \mathbf{D} , best correlated to the residual vector of the current approximation of \mathbf{y} . For initialization, \mathbf{y} itself becomes the residual vector. 2) *Augmentation* of a selected subspace with the identified atom(s). The selected subspace is empty in the first iteration. 3) *Residual update*, after approximating \mathbf{y} with the selected subspace. The above steps are repeated until some *stopping rule* is satisfied by the algorithm. Recently proposed greedy algorithms vary in different steps of the base-line algorithms.

OMP [14] follows the base-line algorithm very closely. After identifying a single atom and augmenting the selected subspace with it, OMP finds the least squares approximation of \mathbf{y} with the selected subspace. This approximation is subtracted from \mathbf{y} in the *residual update* step. This makes the residual vector *orthogonal* to the selected subspace. This notion is an enhancement over the Matching Pursuit (MP) algorithm [23], that updates the residual vector by deflating it with the last atom added to the selected subspace. A non-negative variant of OMP, henceforth denoted as OMP+, has also been proposed in [16]. OMP+ restricts its solution space only to the vectors with non-negative coefficients. Generalized-OMP (gOMP) [24] is a generalization of OMP. It identifies a fixed number of atoms in each iteration and augments the selected subspace with all of them. Then it updates the residue like OMP.

Subspace Pursuit (SP) [27], Compressive Sampling MP (CoSaMP) [25] and Regularized-OMP (ROMP) [26] are the algorithms that assume prior knowledge of the cardinality p of the signal \mathbf{y} . These algorithms identify a subspace of atoms in the step (1). In each iteration, SP identifies p atoms and augments the selected subspace with all of them. From this subspace it again selects p atoms that have maximal contribution to the approximation of \mathbf{y} in the least squares sense. SP updates the residue like OMP. CoSaMP identifies $2p$ atoms in each iteration and augments the selected subspace with them. Then, it selects p atoms like SP. However, it updates the residue by using the already computed coefficients of the selected p atoms. ROMP also identifies p atoms in each iteration. Then it drops off some of these atoms using a predefined regularization rule before the step (2). In ROMP, the residual vector is also updated following OMP.

IV. PROPOSED ALGORITHM

Hyperspectral data has its own characteristics, such as, the cardinality of a mixed pixel in an image is usually small (four to five) but unknown, the spectra of different materials (i.e. the atoms of the library) are highly correlated and the fractional abundances are non-negative quantities. The greedy sparse approximation algorithms reviewed in Section 3 were not originally proposed for hyperspectral unmixing, therefore they do not explicitly take care of the above mentioned characteristics of the hyperspectral data. In fact, to the best of our knowledge, no pixel-based greedy sparse approximation algorithm has ever been proposed specifically for the problem of sparse unmixing of hyperspectral data.

Here, we present Sparse Unmixing via Greedy Pursuit (SUnGP), a pixel-based greedy algorithm that has been designed particularly for the problem of hyperspectral unmixing. SUnGP is shown in Algorithm 1. Each iteration of the algorithm comprises the three steps of the base-line algorithm in Section 3. In the *identification* step, SUnGP first computes the correlations between the atoms of the dictionary and the residual vector of the current approximation of \mathbf{y} , where \mathbf{y} itself is considered as the residual vector at initialization. Then, SUnGP identifies the atoms of the dictionary corresponding to the L (an algorithm parameter) largest values of the computed correlations. These atoms are used for temporarily augmenting the selected subspace. Using this subspace, a non-negative least squares approximation of the mixed signal is computed (line ‘8’ in Algorithm 1). SUnGP then identifies the atom from the

forementioned L atoms that has the maximum contribution in this approximation. This atom is used for permanently augmenting the selected subspace (line ‘10’).

Notice that, SUnGP first identifies a subspace of L atoms (which are highly correlated) and later prunes it keeping in view the non-negativity of the solution space, to identify the single best atom. Once the best atom is identified, it is added to the selected subspace, never to be dropped off in the future iterations. This strategy stems directly from the aforementioned characteristics of the hyperspectral data. SUnGP follows OMP in the *residual update* step and uses a disjunction of three *stopping rules* (line ‘13’). The rules (a) and (b) are self explanatory. The rule (c) ensures that the algorithm stops if it is not able to reduce the l_2 norm of the residual vector at least by a fraction β in its last iteration. The rules (b) and (c) allow SUnGP to operate without prior information about the cardinality of the mixed pixel.

Algorithm 1 SUnGP

Initializaiton:

- 1: Iteration: $i = 0$
- 2: Initial solution: $\boldsymbol{\alpha}^0 = \mathbf{0}$
- 3: Initial residual: $\mathbf{r}^0 = \mathbf{y} - \mathbf{D}\boldsymbol{\alpha}^0 = \mathbf{y}$
- 4: Selected support: $\mathcal{S}^0 = \text{support}\{\boldsymbol{\alpha}^0\} = \emptyset$

Main Iteration: Update iteration: $i = i + 1$

Identification:

- 5: Compute $p_j = \frac{\mathbf{d}_j^T \mathbf{r}^{i-1}}{\|\mathbf{d}_j\|_2}, \forall j \in \{1, \dots, k\}$.
- 6: $\mathcal{N} = \{\text{indices of the atoms of } \mathbf{D} \text{ corresponding to the } L \text{ largest } p_j\}$
- 7: $\mathcal{S}_t = \mathcal{S}^{i-1} \cup \mathcal{N}$
- 8: $\boldsymbol{\alpha}_t = \min_{\boldsymbol{\alpha}} \|\mathbf{D}\boldsymbol{\alpha} - \mathbf{y}\|_2^2 \text{ s.t. } \text{support}\{\boldsymbol{\alpha}\} = \mathcal{S}_t, \boldsymbol{\alpha} \geq 0$
- 9: $j^* = \text{index } j \text{ of the largest coefficient of } \boldsymbol{\alpha}_t, \text{ s.t. } j \in \mathcal{N}$

Augmentation:

- 10: $\mathcal{S}^i = \mathcal{S}^{i-1} \cup \{j^*\}$

Residual update:

- 11: Compute $\boldsymbol{\alpha}^i = \min_{\boldsymbol{\alpha}} \|\mathbf{D}\boldsymbol{\alpha} - \mathbf{y}\|_2^2 \text{ s.t. } \text{support}\{\boldsymbol{\alpha}^i\} = \mathcal{S}^i$
- 12: $\mathbf{r}^i = \mathbf{y} - \mathbf{D}\boldsymbol{\alpha}^i$

Stopping rule:

- 13: If a) $i > \text{desired iterations}$, or b) $\|\mathbf{r}^i\|_2 < \epsilon_0$, or c) $\|\mathbf{r}^i\|_2 > \beta \|\mathbf{r}^{i-1}\|_2$ stop, otherwise iterate again.
-

V. PROCESSING THE DATA FOR GREEDY ALGORITHMS

In the sparse approximation literature, the correlation among the atoms of the dictionary is usually quantified by the mutual coherence $\mu \in [0, 1]$ of the dictionary.

$$\mu = \max_{i,j:j \neq i} \frac{\text{abs}(\mathbf{d}_i^T \mathbf{d}_j)}{\|\mathbf{d}_i\|_2 \|\mathbf{d}_j\|_2} \quad (6)$$

where \mathbf{d}_z is the z^{th} atom of the dictionary. In general, the greedy sparse approximation algorithms are able to identify the support of the solution more accurately, if μ is small for the dictionary. For instance, according to [17], if $\mu < 0.33$ OMP will always find the exact support for a signal that is a linear combination of two distinct atoms of the dictionary. However, in the sparse unmixing problem, usually $\mu \approx 1$ [13]. Researchers at the German Aerospace Center [29] have shown that for a dictionary of spectra sampled at a constant wavelength interval, μ can be reduced by taking the derivatives of

the spectra. The derivative of a spectra $\mathbf{d} \in \mathbb{R}^m$ is defined as

$$\Delta(\mathbf{d}) = \frac{d(b_i) - d(b_j)}{b_j - b_i}, \forall_i i \in \{1, \dots, m - c\} \quad (7)$$

where, b_z is the wavelength at the z^{th} band, $d(b_z)$ is the reflectance value at that wavelength and $j = i + c$, with i, j and c being positive integers.

Keeping in view its coherence reduction ability, we benefit from the derivative operation in hyperspectral unmixing with the greedy algorithms. We propose to use the following strategy for processing the hyperspectral data for the greedy sparse approximation algorithms:

- 1) Create the dictionary \mathbf{D}_Δ and the pixel \mathbf{y}_Δ by taking the derivative of each spectra in \mathbf{D} and the pixel \mathbf{y} , respectively.
- 2) Compute α with the greedy sparse approximation algorithm using \mathbf{D}_Δ and \mathbf{y}_Δ .
- 3) Estimate the fractional abundance vector $\tilde{\alpha}$, by the non-negative least squares approximation of \mathbf{y} using the atoms of \mathbf{D} corresponding to the support of α .

The unmixing performed on the differentiated data in step (2) is used to identify the correct support of the solution. Once the support is found, it is used for the actual estimation of the fractional abundances using the original data in the step (3). In the above strategy, we use the dictionaries after normalizing their atoms in l_1 norm. By doing so, the estimated fractional abundances automatically satisfy ASC. It is worth mentioning here that the derivative operation helps in coherence reduction, however the correlation among the spectra of the differentiated data still remains high enough to cause problems for the greedy sparse approximation algorithms. Therefore, the algorithms need to show robustness against the correlation of the spectra for effective hyperspectral unmixing.

VI. EXPERIMENTS

In this section, we present the results of the experiments performed with synthetic data and the real-world data. The experiments with the synthetic data are important because they provide quantitative evaluation of the approach, which is not possible with the real-world data. In all the experiments we used a fixed dictionary, which was created from the NASA Jet Propulsion Laboratory's Advanced Space-borne Thermal Emission and Reflectance Radiometer (ASTER) spectral library (<http://speclib.jpl.nasa.gov>). This library contains pure spectra of 2400 materials. To perform the experiments, we selected 425 of these spectra and resampled them in the wavelength range 0.4 to $2.5\mu\text{m}$, at a constant interval of 10nm . The resampling is performed to match the sampling strategy of the NASA's Airborne Visible and Infrared Imaging Spectrometer (AVIRIS) [12]. We dropped 24 bands of the spectra in the dictionary because of zero or very low reflectance values. This made \mathbf{D} a 200×425 matrix. The spectra were selected such that $\mu = 0.9986$ for the dictionary. We kept $\mu < 1$ in order to ensure that the spectra in the dictionary are unique.

A. Results on synthetic data

We simulated the synthetic hyperspectral data with 500 mixed pixels, where each pixel was a linear combination

of p randomly selected spectra from the dictionary. Following the experimental protocol in [13], we drew the fractional abundances of the endmembers in each pixel from a Dirichlet distribution. Therefore, the fractional abundances satisfy ASC. We added the Gaussian white noise to the data such that the pixels had $\text{SNR} = 50\text{dB}$. The algorithms were evaluated for the two goals of hyperspectral unmixing, namely, the endmembers identification and the estimation of the fractional abundances. For the former, we used the evaluation metric of *unmixing fidelity* $\Phi(\alpha) \rightarrow [0, 1]$. If $\mathcal{P} = \{x \mid x \text{ is the index of an endmember in } \mathbf{D}\}$ and $\mathcal{A} = \{a \mid a \text{ is a non-zero elements in } \alpha\}$, then:

$$\Phi(\alpha) = \frac{|\mathcal{P} \cap \mathcal{A}|}{|\mathcal{A}|} \quad (8)$$

where $|\cdot|$ denotes the cardinality of the set. Fig. 2a shows the results of the experiments performed to evaluate the endmember identification ability of the greedy algorithms. The values in these results, and the results to follow, are the mean values of the metrics, computed over the whole synthetic data. The results show that SUnGP performs better than the existing greedy sparse approximation algorithms. For SUnGP we used $L = 50$, which was optimized on a separate training data set. Different parameters of the other algorithms were also optimized on the same training data set. We have given the correct value of p to the greedy algorithms as their input parameters. gOMP is tuned to select 2 atoms in its each iteration. The figure shows some of the results in dotted plots. The algorithms corresponding to these plots assume a priori knowledge of p . Therefore, they are not of practical use in hyperspectral unmixing. However, we have included them in the analysis for a comprehensive comparison of SUnGP's performance with the state of the art greedy algorithms. Fig. 2b compares the algorithms for the different noise levels of data. SUnGP again shows good performance, especially at high SNR ¹.

In our experiments, we have processed the data according to the strategy discussed in Section 5. The unmixing fidelity of the solution can be computed at the step (2) of the strategy. Therefore, the results discussed above were computed at that step. Following [28], we chose $c = 2$ in Equation (7) for the derivative operation. To evaluate the estimation of the fractional abundances we also performed the step (3) of the strategy and compared the results of SUnGP with the results of the popular convex relaxation based algorithms that have been proposed specifically for hyperspectral unmixing. For this comparison, we did not assume a priori knowledge of p and used the stopping rule (c) (line '13', Algorithm 1) with $\beta = 0.9$ for SUnGP. The value of β was optimized on a separate training data set. For the convex relaxation based algorithms we chose $\lambda = 10^{-3}$. This value was also optimized on the training data. We used the Euclidean distance between the estimated fractional abundance vector $\tilde{\alpha}$ and the actual fractional abundance vector α_0 , as the evaluation metric. The comparison of the results is given in Fig. 2c, which shows that SUnGP's performance is comparable to the convex relaxation based algorithms. We do not present the results of fractional abundance estimation by the other greedy algorithms because

¹SNR as high as 400dB is a realistic value for hyperspectral remote sensing instruments [29]. Our analysis focuses only on the more challenging case of low SNR.

of space limitations. However, it is easy to see that better Φ will lead SUNGP to outperform the other greedy algorithms in fractional abundance estimation.

A major advantage of using the greedy algorithms in hyperspectral unmixing is the computational efficiency. Table 1 compares the computational timings of the algorithms. These timings have been computed for unmixing of a 500 pixel synthetic image, with $p = 5$ for each pixel and $\text{SNR} = 50\text{dB}$. Greedy algorithms, other than SUNGP, assume prior knowledge of p . The algorithms use the same parameter settings which are used in the results discussed above. Timings have been computed with authors' MATLAB implementations for each of the algorithms.

B. Results on the real-world data

We also performed sparse unmixing of the real-world hyperspectral data (http://aviris.jpl.nasa.gov/data/free_data.html), acquired by AVIRIS. From this data, we selected an image cube of dimensions $350 \times 350 \times 224$. The spatial dimensions (350×350) of this cube represent a region of Cuprite mines, Nevada. The Cuprite mining district has been studied well for its surface materials in the Geological Sciences literature. The USGS classification map in Fig. 3 shows the labels of the different materials in the region, as computed by the Tricorder software package (<http://speclab.cr.usgs.gov/PAPERS/tetracorder/>). For the region analyzed in this work, we separately show the *classification map* of Alunite (a mineral) computed by Tricorder. The figure also shows the *fractional abundance maps* of Alunite as computed by the different sparse approximation algorithms (only the best ones are shown because of the space limitations). From the figure, it is visible that SUNGP has estimated high fractional abundances for the pixels which have been classified as Alunite by Tricorder. The values are higher than those computed by any other algorithm.

Following [13] and [18], we have provided the results only for the visual comparison. The quantitative comparison of the results on the real-world data is not possible, as no ground truth values of the fractional abundances are available for the real-world data [18]. In the results shown, we have used 188 spectral bands (out of 224) of the image cube. The other bands were dropped because of zero or very low reflectance values. The corresponding bands were also dropped from the dictionary. We used the strategy discussed in Section 5 for the greedy algorithms. For CoSaMP we set $p = 5$. Rest of the algorithms use the same parameter settings as discussed in the previous section.

VII. CONCLUSION

We have proposed a pixel based greedy sparse approximation algorithm, called SUNGP, for hyperspectral unmixing. The proposed algorithm identifies different spectra in a mixed hyperspectral pixel by iteratively selecting a subspace of spectra from a fixed dictionary and pruning it. SUNGP has been tested on synthetic as well as the real-world remote sensing hyperspectral data. The algorithm has been shown to outperform the existing state of the art greedy sparse approximation algorithms. Furthermore, its results are comparable to the convex relaxation based sparse approximation algorithms, with a considerable computational advantage.

ACKNOWLEDGMENT

This research was supported by ARC Grant DP110102399

REFERENCES

- [1] K. Staenz, A. Mueller, A. Held, and U. Heiden, *Technical committees corner: International spaceborne imaging spectroscopy (ISIS) technical committee*, IEEE Geosci. Remote Sens. Lett., no. 165, pp.38-42, 2012.
- [2] J. B. Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, and P. Gader, *Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches*, IEEE J. Sel. Topics Appl. Earth Observ., 2012.
- [3] J. B. Dias, A. Plaza, G. Camps-Valls, P. Scheunders, N. Nasrabdi, and J. Chanussot, *Hyperspectral Remote Sensing Data Analysis and Future Challenges*, IEEE Geosci. Remote Sens. Mag., vol. 1, no. 2, 2013.
- [4] J. M.P. Nascimento and J. B. Dias. *Vertex component analysis: A fast algorithm to unmix hyperspectral data*. IEEE Trans. Geosci. Remote Sens., 2005.
- [5] M. E. Winter, *N-finder: An algorithm for fast autonomous spectral endmember determination in hyperspectral data*. in Proc. SPIE Imaging Spectrometry V, 2003, vol. 3753, pp. 266-275.
- [6] J. W. Boardman, F. A. Kruse, and R. O. Green. *Mapping target signatures via partial unmixing of aviris data*, In JPL Airborne Earth Sci. Workshop, pages 2326, 1995.
- [7] C. I. Chang, C. C. Wu, W. Liu, and Y. C. Ouyang, *A new growing method for simplex based endmember extraction algorithm*, IEEE Trans. Geosci. Remote Sens., vol. 44, no. 10, pp. 2804-2819, Oct 2006.
- [8] J. Li and J. B. Dias. *Minimum volume simplex analysis: A fast algorithm to unmix hyperspectral data*, In Geosci. Remote Sens. Symp., 2008.
- [9] M. Berman, H. Kiiveri, R. Lagerstrom, A. Ernt, R. Dunne, and J. F. Huntington, *ICE: A statistical approach to identify endmembers in hyperspectral images*, IEEE Trans. Geosci. Remote Sens., 2004.
- [10] A. Zare and P. Gader, *Sparsity promoting iterated constrained endmember detection for hyperspectral imagery*, IEEE Geosci. Remote Sens. Lett., vol. 4, no. 3, pp. 446-450, 2007.
- [11] M. Iordache, *A sparse regression approach to hyperspectral unmixing*, PhD thesis, Universidade Tecnica de Lisboa, Instituto Superior Tecnico, Nov 2011.
- [12] R. 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. Soils, M. R. Olah, and O. Williams, *Imaging spectroscopy and the airborne visible/infrared imaging spectrometer (aviris)*, Remote Sensing of Environment, vol. 65, no. 3, pp. 227-248, 1998.
- [13] M. D. Iordache, J. B. Dias, and A. Plaza. *Sparse unmixing of hyperspectral data*, IEEE Trans. Geosci. Remote Sens., vol. 49, no. 6, pp. 2014 - 2039, 2011.
- [14] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad. "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition." In Proc. IEEE Conf. Record of the Asilomar Conf. on Sig. Sys. Comp, 1993.
- [15] S. S. Chen, D. L. Donoho, Michael, and A. Saunders, *Atomic decomposition by basis pursuit*, SIAM Journal on Scientific Computing, vol. 20, pp. 33-61, 1998.
- [16] A. Bruckstein, M. Elad, and M. Zibulevsky, *On the uniqueness of non-negative sparse solutions to underdetermined systems of equation*, IEEE Trans. Inf. Theory, vol. 54, no. 11, pp. 4813-4820, Nov. 2008.
- [17] J. A. Tropp, *Greed is Good: Algorithmic Results for Sparse Approximation*, IEEE Trans. Inf. Theory, vol. 50, no. 10, pp. 2231-2242, 2004.
- [18] Z. Shi, W. Tang, Z. Duren, and Z. Jiang, *Subspace Matching Pursuit for Sparse Unmixing of Hyperspectral Data*, IEEE Trans. Geosci. Remote Sens., vol. PP, no. 99, pp. 1-19, 2013.
- [19] J. A. Tropp, A. C. Gilbert, and M. J. Strauss, *Algorithms for simultaneous sparse approximation: Part I: Greedy Pursuit*, Signal Process., vol. 86, no. 3, pp. 572-588, Mar. 2006.
- [20] B. Natarajan, *Sparse approximation solution to linear systems*, SIAM J. Comput., vol. 24, no. 2, pp. 227-234, 1995.
- [21] R. Tibshirani, *Regression shrinkage and selection via Lasso*, J. R. Stat. Soc., vol. 58, no.1, pp. 267- 288, 1996.

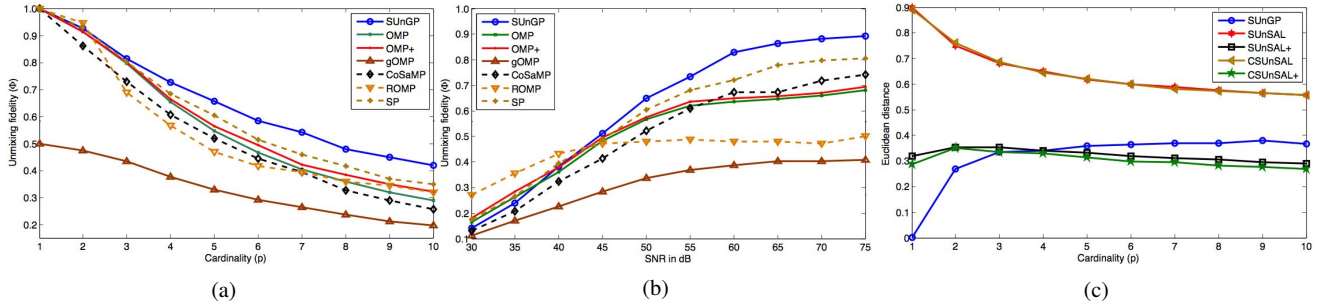


Fig. 2: Comparison of the results on synthetic data: a) Unmixing fidelity computed by different algorithms, as a function of the cardinality of the mixed pixels. The image contains the Gaussian white noise with $\text{SNR} = 50\text{dB}$ for each pixel. b) Unmixing fidelity as a function of the SNR of the pixels. The cardinality of each pixel is 5. c) Euclidean distance between α_0 and $\hat{\alpha}$, as a function of the cardinality of the mixed pixels. The comparison is between SUnGP and the convex relaxation based algorithms.

TABLE I: Processing time (in seconds) for 500 pixel synthetic image with 50 dB SNR. Each pixel is a mixture of 5 randomly selected spectra. The time is computed on a desktop PC equipped with an Intel Core i7-2600 CPU (at 3.4 GHz) and 8 GB RAM.

Algorithms	SP[27]	ROMP[26]	gOMP[24]	OMP[14]	OMP+[16]	CoSaMP[25]
Time (s)	0.83	0.85	0.89	1.76	2.27	2.82
Algorithms	SUnGP	SUnSAL+[22]	CSUnSAL+[22]	SUnSAL+[22]	CSUnSAL+[22]	
Time (s)	5.11	31.3	33.7	35.9	36.1	

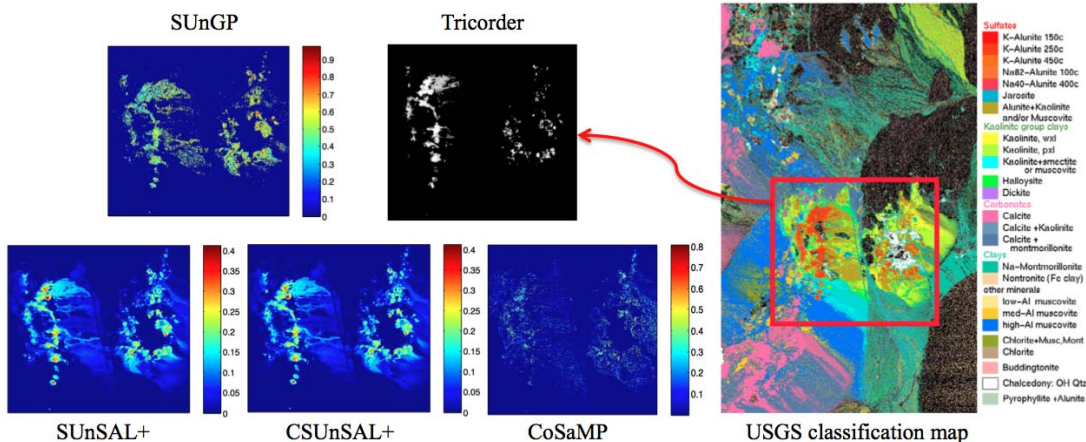


Fig. 3: Comparison of the results on the real-world data : The *USGS classification map* shows the labels of different materials in an image taken over Cuprite mines, NV. The labels are assigned with the Tricorder software. The 350×350 area inside the rectangle is analyzed in this work. The *Tricorder* map shows the classification map of Alunite (a mineral) in the region. Each of the other four maps show the fractional abundances of Alunite in the region, computed by the algorithms mentioned.

[22] J. M. Bi. Dias and M. A. T. Figueiredo, *Alternating direction algorithms for constrained sparse regression: Application to hyperspectral unmixing*, in Proc. 2nd WHISPERS, Jun. 2010, pp. 1-4.

[23] S. G. Mallat and Z. Zhang, *Matching Pursuits with Time-Frequency Dictionaries*, IEEE Transactions on Signal Processing, December 1993, pp. 3397-3415.

[24] J. Wang, S. Kwon, and B. Shim, *Generalized Orthogonal Matching Pursuit*, IEEE Transactions on Signal Processing vol. 60, no. 12, pp. 6202 - 6216, 2012.

[25] D. Needell, and J. A. Tropp, *CoSaMP: Iterative signal recovery from incomplete and inaccurate samples*, Appl. Comput. Harmon. Anal., vol. 26, no. 3, pp. 301 - 321, May 2009.

[26] D. Needell and R. Vershynin, *Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit*, IEEE J. Sel. Topics Signal Process., vol. 4, no. 2, pp. 310-316, Apr. 2010

[27] W. Dai and O. Milenkovic, *Subspace pursuit for compressive sensing signal reconstruction*, Information Theory, IEEE Transactions on, vol. 55, no. 5, pp. 2230-2249, May 2009.

[28] N. Akhtar, F. Shafait, and A. Mian, *Repeated constraint sparse coding with partial dictionaries for hyperspectral unmixing*, In Proceedings of IEEE WACV, 2014.

[29] J. Bieniarz, R. Mueller, X. Zhu, and P. Reinartz, *On The Use Of Overcomplete Dictionaries For Spectral Unmixing*, Proceedings of the 5th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing, 2013.