

GLUP: YET ANOTHER ALGORITHM FOR BLIND UNMIXING OF HYPERSPECTRAL DATA

Rita Ammanouil, André Ferrari, Cédric Richard, David Mary

Université de Nice Sophia-Antipolis, CNRS, Observatoire de la Côte d’Azur, France

ABSTRACT

This paper addresses the problem of blind fully-constrained linear unmixing of hyperspectral images. The endmembers and their cardinality are assumed unknown, but the endmember spectra are supposed to be present in the scene. Group Lasso regularization is used to extract the endmembers. The estimation problem is convex, and solved with the Alternating Direction Method of Multipliers. Comparisons with state-of-the-art methods on synthetic and real data sets show the efficiency of our approach.

Index Terms— Hyperspectral data, blind unmixing, linear mixing model, ADMM

1. INTRODUCTION

Hyperspectral imaging consists of capturing a scene over hundreds of narrow and adjacent spectral bands. The resulting images are characterized by both a spatial and a spectral dimension, which means that each pixel is a local spectrum. Unmixing is one of the most prominent techniques for analyzing hyperspectral data [1]. It consists of identifying the pure components in a scene, the so-called endmembers, and then estimating their spatial distribution, also known as abundance map. Most unmixing methods in the literature focus on the Linear Mixing (LM) [2], where each pixel is modeled by a linear combination of the endmember spectra weighted by their abundances. See [3–7] for extensions to nonlinear mixing models. With LM, the positivity and sum-to-one constraints are imposed on the abundances to ensure a physically meaningful model. Identifying the endmembers is a crucial task. It is usually the first unmixing output, and thus has serious consequences on the accuracy of abundance map estimation. Many algorithms have been developed to extract endmembers. Geometrical approaches, such as the well-known NFINDR [8], SGA [9] and those described in [10], search for the vertices of a simplex that encloses the data. Matching Pursuit methods [11] such as the recently published Self Dictionary Simultaneous Orthogonal Matching Pursuit (SDSOMP) [12] sequentially identify the indices of the pixels that have the largest contribution in the data. The algorithms mentioned previously assume that the number of endmembers is known, and rely on their presence in the scene. Some algorithms drop off the latter assumption and use predefined spectral libraries. With these algorithms, unmixing is formulated as a sparse regression problem [13]. The endmembers are sought in a library containing a huge number of endmember samples, and sparsity inducing regularizers are required to extract the desired endmembers. The main drawback of this trend is that the endmembers available in libraries were usually measured under different conditions.

This work was partly supported by the Agence Nationale pour la Recherche, France, (Hypanema project, ANR-12-BS03-003), and the Centre National de la Recherche Scientifique, France (Display project, Mastodons).

In this paper, we formulate the problem of endmember extraction as a sparse unmixing problem. We pay particular attention to the case where the endmembers are present among the available observations, using the data itself as the spectral library. We consider the Group Lasso $\ell_{2,1}$ -norm regularizer [14] to induce sparsity, and to impose the positivity and sum-to-one constraints on the abundances. A similar approach has been introduced in [15], however, the authors use the ℓ_∞ -norm and drop the sum-to-one constraint. Our optimization problem is solved with the Alternating Direction Method of Multipliers (ADMM) [16]. The rest of the paper is organized as follows. Section 2 describes the proposed model and the ADMM steps used to find the solution. Section 3 provides a comparison between the proposed approach and two algorithms, namely, NFINDR and SDSOMP. Finally, Section 4 concludes and provides some possible extensions for future work.

2. GROUP LASSO WITH UNIT SUM AND POSITIVITY CONSTRAINTS (GLUP)

Let us first introduce the linear mixing model and some useful notations. In matrix form, this model can be expressed as follows

$$\mathbf{S} = \mathbf{R}\mathbf{A} + \mathbf{E} \quad (1)$$

with $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_N)$, $\mathbf{R} = (\mathbf{r}_1, \dots, \mathbf{r}_M)$, $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_M)^\top$. Here, \mathbf{s}_j is the L -dimensional spectrum of the j -th pixel, L is the number of frequency bands, \mathbf{r}_i is L -dimensional spectrum of the i -th endmember, M denotes the number of endmembers, \mathbf{a}_i is the N -dimensional abundance map of the i -th endmember, N is the number of pixels in the image, and \mathbf{E} is the noise. All vectors are column vectors. Model (1) means that the (i, j) -th entry \mathbf{A}_{ij} of \mathbf{A} is the abundance of the endmember \mathbf{r}_i in pixel \mathbf{s}_j . Two constraints on the abundances are considered, the nonnegativity and sum-to-one constraints: $\mathbf{A}_{ij} \geq 0$ for all (i, j) , and $\sum_{i=1}^M \mathbf{A}_{ij} = 1$ for all j .

As mentioned in the introduction, we shall assume that the endmembers are unknown but present in the scene. Let ω be a subset of N' indexes in $\{1, \dots, N\}$ that contains at least the column index of each endmember. Under these assumptions, and without loss of generality, we observe that model (1) can be reformulated as

$$\mathbf{S} \approx \mathbf{S}_\omega \mathbf{X} + \mathbf{E} \quad (2)$$

with $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_{N'})^\top$ the abundance matrix, $N' \gg M$, and $\mathbf{S}_\omega = (\mathbf{s}_{\omega_1}, \dots, \mathbf{s}_{\omega_{N'}})$ the restriction of \mathbf{S} to its columns indexed by ω . Similarly as above, \mathbf{X}_{ij} is the abundance of \mathbf{s}_{ω_i} in \mathbf{s}_j and should thus obey the sum-to-one and positivity constraints. Equation (2) is an approximation rather than an exact equality due to the fact that \mathbf{S}_ω possibly contains noisy endmembers. However, we assume that the noise does not dramatically affect the factorization of the mixing process, which is valid for high SNR. In (2), note that if \mathbf{s}_{ω_i} is an endmember, \mathbf{x}_i has non-zero entries and represents the abundance map of \mathbf{s}_{ω_i} . However, if \mathbf{s}_{ω_i} is not an endmember, \mathbf{x}_i has

all its elements equal to zero. Consequently, \mathbf{X} admits $N' - M$ rows of zeros, the other rows being the rows of \mathbf{A} . This means that \mathbf{X} allows to identify the endmembers in \mathbf{S} through its non-zero rows.

Within this context, we shall look for a few columns of \mathbf{S}_ω that can effectively represent the whole scene. The unmixing problem requires that $\mathbf{S}_\omega \mathbf{X}$ matches \mathbf{S} , that \mathbf{X} only has a few rows different from zero, in addition to the non-negativity and sum-to-one constraints. This leads to the following convex optimization problem

$$\begin{aligned} \min_{\mathbf{X}} \quad & \frac{1}{2} \|\mathbf{S} - \mathbf{S}_\omega \mathbf{X}\|_F^2 + \mu \sum_{k=1}^N \|\mathbf{x}_k\|_2 \\ \text{subject to} \quad & \mathbf{X}_{ij} \geq 0 \quad \forall i, j \\ & \sum_{i=1}^N \mathbf{X}_{ij} = 1 \quad \forall j \end{aligned} \quad (3)$$

with $\mu \geq 0$ a regularization parameter. The algorithm has to force rows of \mathbf{X} to be zero vectors in order to identify the endmembers. To promote this effect, the Group Lasso $\ell_{2,1}$ -norm regularization is employed. It induces sparsity in the abundance matrix at the group level [14], by possibly driving all the entries in several rows \mathbf{x}_k of \mathbf{X} to zero. The solution of problem (3) can be obtained in a simple and flexible manner using the ADMM algorithm [16]. We consider the canonical form

$$\begin{aligned} \min_{\mathbf{X}, \mathbf{Z}} \quad & \frac{1}{2} \|\mathbf{S} - \mathbf{S}_\omega \mathbf{X}\|_F^2 + \mu \sum_{k=1}^N \|\mathbf{z}_k\|_2 + \mathcal{I}(\mathbf{Z}) \\ \text{subject to} \quad & \mathbf{A}\mathbf{X} + \mathbf{B}\mathbf{Z} = \mathbf{C} \end{aligned} \quad (4)$$

with

$$\mathbf{A} = \begin{pmatrix} \mathbf{I} \\ \mathbf{1}^\top \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} -\mathbf{I} \\ \mathbf{0}^\top \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} \mathbf{0} \\ \mathbf{1}^\top \end{pmatrix},$$

where \mathcal{I} is the indicator of the positive orthant guarantying the positivity constraint, that is, $\mathcal{I}(\mathbf{Z}) = 0$ if $\mathbf{Z} \succeq \mathbf{0}$ and $+\infty$ otherwise. The equality constraint imposes the consensus $\mathbf{X} = \mathbf{Z}$ and the sum-to-one constraint. In matrix form, the augmented Lagrangian for problem (4) is given by [17]

$$\begin{aligned} \mathcal{L}_\rho(\mathbf{X}, \mathbf{Z}, \mathbf{\Lambda}) = & \frac{1}{2} \|\mathbf{S} - \mathbf{S}_\omega \mathbf{X}\|_F^2 + \mu \sum_{k=1}^N \|\mathbf{z}_k\|_2 + \mathcal{I}(\mathbf{Z}) + \\ & \text{trace}(\mathbf{\Lambda}^\top (\mathbf{A}\mathbf{X} + \mathbf{B}\mathbf{Z} - \mathbf{C})) + \frac{\rho}{2} \|\mathbf{A}\mathbf{X} + \mathbf{B}\mathbf{Z} - \mathbf{C}\|_F^2 \end{aligned}$$

where $\mathbf{\Lambda}$ is the matrix of Lagrange multipliers, μ and ρ are positive parameters, respectively. ADMM flexibility lies in the fact that it splits the initial variable \mathbf{X} into two variables, \mathbf{X} and \mathbf{Z} , and equivalently the initial problem into two subproblems. At iteration $k + 1$, the ADMM algorithm consists of three sequential steps:

2.1. Minimization of $\mathcal{L}_\rho(\mathbf{X}, \mathbf{Z}^k, \mathbf{\Lambda}^k)$ with respect to \mathbf{X}

This step takes into account the previous estimates of \mathbf{Z} and $\mathbf{\Lambda}$. The augmented Lagrangian is quadratic in terms of \mathbf{X} . As a result, the solution has an analytical expression that is obtained by setting the gradient of $\mathcal{L}_\rho(\mathbf{X}, \mathbf{Z}^k, \mathbf{\Lambda}^k)$ to zero

$$\mathbf{X}^{k+1} = (\mathbf{S}_\omega^\top \mathbf{S}_\omega + \rho \mathbf{A}^\top \mathbf{A})^{-1} (\mathbf{S}_\omega^\top \mathbf{S} - \mathbf{A}^\top [\mathbf{\Lambda}^k + \rho (\mathbf{B}\mathbf{Z}^k - \mathbf{C})])$$

2.2. Minimization of $\mathcal{L}_\rho(\mathbf{X}^{k+1}, \mathbf{Z}, \mathbf{\Lambda}^k)$ with respect to \mathbf{Z}

After removing the terms that are independent of \mathbf{Z} , the minimization of $\mathcal{L}_\rho(\mathbf{X}^{k+1}, \mathbf{Z}, \mathbf{\Lambda}^k)$ with respect to \mathbf{Z} reduces to solving the following problem

$$\begin{aligned} \min_{\mathbf{Z}} \quad & \mu \sum_{k=1}^N \|\mathbf{z}_k\|_2 + \text{trace}(\mathbf{\Lambda}^\top \mathbf{B}\mathbf{Z}) + \frac{\rho}{2} \|\mathbf{A}\mathbf{X} + \mathbf{B}\mathbf{Z} - \mathbf{C}\|_F^2 \\ \text{subject to} \quad & \mathbf{Z} \succeq \mathbf{0} \end{aligned}$$

This minimization step can be split into N problems given the structure of matrices \mathbf{A} and \mathbf{B} , one for each row of \mathbf{Z} , that is,

$$\min_{\mathbf{z}} \quad \frac{1}{2} \|\mathbf{z} - \mathbf{v}\|_2^2 + \alpha \|\mathbf{z}\|_2 + \mathcal{I}(\mathbf{z}) \quad (5)$$

where $\mathbf{v} = \mathbf{x} + \rho^{-1} \boldsymbol{\lambda}$, $\alpha = \rho^{-1} \mu$, $\boldsymbol{\lambda}$, \mathbf{x} and \mathbf{z} correspond to a row in $\mathbf{\Lambda}$, \mathbf{X} and \mathbf{Z} respectively. Problem (5) admits a unique solution, given by the proximity operator [18] of $f(\mathbf{z}) = \alpha \|\mathbf{z}\|_2 + \mathcal{I}(\mathbf{z})$

$$\begin{cases} \mathbf{z}^* = \mathbf{0} & \text{if } \|(\mathbf{v})_+\|_2 < \alpha \\ \mathbf{z}^* = \left(1 - \frac{\alpha}{\|(\mathbf{v})_+\|_2}\right) (\mathbf{v})_+ & \text{otherwise} \end{cases} \quad (6)$$

where $(\cdot)_+ = \max(\mathbf{0}, \cdot)$.

2.3. Update of the Lagrange multipliers $\mathbf{\Lambda}$

Lagrange multiplier update is carried out at the end of each iteration. $\mathbf{\Lambda}^{k+1}$ represents the running sum of residuals. It gives an insight on the convergence of the algorithm. As k tends to infinity, the primal residual tends to zero and $\mathbf{\Lambda}^{k+1}$ converges to the dual optimal point.

$$\mathbf{\Lambda}^{k+1} = \mathbf{\Lambda}^k + \rho (\mathbf{A}\mathbf{X}^{k+1} + \mathbf{B}\mathbf{Z}^{k+1} - \mathbf{C}). \quad (7)$$

As suggested in [16], a reasonable stopping criteria is that the primal and dual residuals must be smaller than some tolerance thresholds, namely,

$$\|\mathbf{A}\mathbf{X}^{k+1} + \mathbf{B}\mathbf{Z}^{k+1} - \mathbf{C}\|_2 \leq \epsilon_{\text{pri}} \quad (8)$$

and

$$\|\rho \mathbf{A}^\top \mathbf{B} (\mathbf{Z}^{k+1} - \mathbf{Z}^k)\|_2 \leq \epsilon_{\text{dual}} \quad (9)$$

3. EXPERIMENTAL RESULTS

3.1. Synthetic data

The performance of the proposed approach was first evaluated using synthetic data. To this end, we extracted 8 endmembers with 420 spectral bands from the USGS library. Using these endmembers, we created a synthetic data set containing 200 pixels. The abundances for every simulated pixel were generated based on a Dirichlet distribution with unit parameter, as a consequence of which the resulting abundances obeyed the non-negativity and sum-to-one constraints. The data set was then corrupted with additive Gaussian noise.

We applied GLUP using $\rho = 1$, $\mu = 1$, the tolerances being set to 10^{-2} . In all the experiments we used $\mathbf{S}_\omega = \mathbf{S}$. Figure 1 shows the estimated abundance matrix for a SNR of 40 dB. White pixel indicates zero entries of \mathbf{X} . Eight rows were identified as abundance maps, and the corresponding columns in \mathbf{S} were identified as the endmembers. The processing time¹ of the algorithm was 1.367 seconds. The performance of GLUP obviously depends on parameter setting. For example, using smaller values of the tolerances results in increased accuracy at the expense of a longer processing time. The regularization parameter affects the sparsity of the solution. Smaller values of μ may result in a less sparse solution than the one in Figure 1. In other words, rows that do not correspond to endmember abundance maps might be different from zero. However, the entries in these rows are usually negligible. Table 1 compares the performance of GLUP, NFINDR and SDSOMP with two SNR levels, 40 and 20 dB respectively. We repeated the simulation 100 times using the same synthetic data set but with a different noise realization each time. For each realization, we examined the percentage of endmembers correctly identified. In the case of NFINDR and SDSOMP, this

¹Machine specifications: 2.2 GHz Intel Core i7 processor and 8 GB RAM

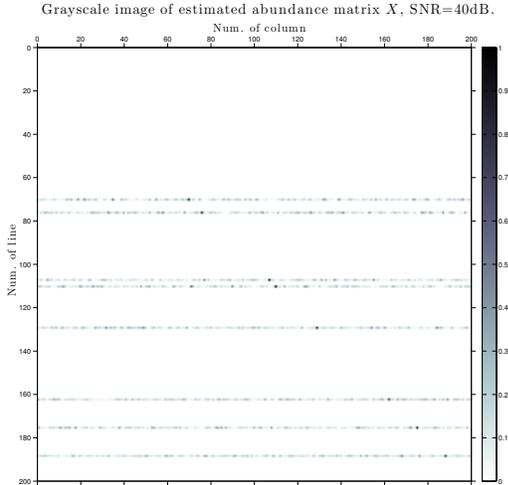


Fig. 1. Image of GLUP’s estimated abundance matrix.

Table 1. Percentage of correctly identified endmembers and mean processing time obtained with $M = 8$ and 100 realizations.

	SNR 40 dB	SNR 20 dB
SDSOMP	100 % (0.023 sec)	72.37 % (0.023 sec)
NFINDR	100 % (0.069 sec)	89.75% (0.068 sec)
GLUP	100 % (1.490 sec)	94.12% (3.737 sec)

task is straightforward. In the case of GLUP, the algorithm returns N candidate abundance maps that can be identified by searching for non-zero rows. As mentioned previously, some rows belonging to mixed pixels might be different from zero but negligible (especially with low SNR). For this reason, we kept the 8 rows with the largest mean value and then evaluated the percentage of correctly identified endmembers. For a SNR of 40 dB, the three algorithms were able to perfectly detect the endmembers in the scene. For a SNR of 20 dB, the performance of SDSOMP deteriorated because it determines the endmembers in a sequential way, one at each iteration. This property makes the algorithm sensitive to noise. Usually, an incorrect detection drives the subsequent iterations in a wrong direction. In addition, the algorithm does not take into account neither the positivity nor the sum-to-one constraints on the abundances. GLUP was able to outperform NFINDR and SDSOMP.

3.2. Real data

The performances of the three endmember extraction algorithms were compared using the Pavia Center data set collected by the ROSIS spectrometer. Pavia Center image has 1096×715 pixels, with 102 bands over the spectral range 430 – 860 nm. We used the left part of the image, composed of 1096×492 pixels.

We randomly selected 200 pixels from the image, and we applied GLUP to this subset. We set the regularization and penalty parameters to 10 and 1 respectively, the primal and dual tolerances to 10^{-2} . GLUP detected 6 endmembers. The same subset of pixels was used to extract the endmembers with NFINDR and SDSOMP. This time, the number of endmembers was preset to 6 for both algorithms. Using the endmembers extracted by the different methods, we applied the Fully Constrained Least Squares (FCLS) on the

Table 2. RMSE and average spectral angle obtained with Pavia.

	RMSE	Avg angle (rad)
SDSOMP	0.1103	0.392
NFINDR	0.0290	0.160
GLUP	0.0198	0.094

whole scene. Table 2 shows the Root Mean Square Error (RMSE) and the average spectral angle (in radian) between the reconstructed image and the original one. GLUP was able to provide lower reconstruction errors than NFINDR and SDSOMP in terms of RMSE and average spectral angle. SDSOMP had the largest values for both metrics, whereas GLUP scored slightly better values than NFINDR.

Figure 2 gives the abundance maps showing the spatial distribution of the estimated endmembers. The three algorithms describe the urban features of the scene as a combination of water, two types of roof tops, meadows, trees, and shadow. Due to the sequential nature of SDSOMP, the first endmember detected by this algorithm captured a dense combination of the materials in the scene rather than highlighting a specific one. NFINDR and GLUP did not encounter this problem. They were both able to improve the quality of the endmembers extraction in the sense that the endmembers can be easily discerned via their respective abundance maps.

4. CONCLUSION

In this paper, we provided a new approach for extracting the endmembers in an hyperspectral scene. We compared the proposed approach with a geometrical and Matching Pursuit approach, namely, NFINDR and SDSOMP. The proposed model uses the available observations, which are possibly noisy, in order to find the endmembers. In future work, we intend to account for the noise present in the observations while extracting the endmembers.

5. REFERENCES

- [1] N. Keshava, “A survey of spectral unmixing algorithms,” *Lincoln Laboratory Journal*, vol. 14, no. 1, pp. 55–78, 2003.
- [2] D. C. Heinz and C. I. Chang, “Fully constrained least squares linear spectral mixture analysis method for material quantification in hyperspectral imagery,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 3, pp. 529–545, 2001.
- [3] Y. Altmann, A. Halimi, N. Dobigeon, and J.-Y. Tourneret, “Supervised nonlinear spectral unmixing using a postnonlinear mixing model for hyperspectral imagery,” *IEEE Transactions on Image Processing*, vol. 21, no. 6, pp. 3017–3025, 2012.
- [4] J. Chen, C. Richard, and P. Honeine, “Nonlinear unmixing of hyperspectral data based on a linear-mixture/nonlinear-fluctuation model,” *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 480–492, 2013.
- [5] J. Chen, C. Richard, and P. Honeine, “Nonlinear estimation of material abundances in hyperspectral images with l_1 -norm spatial regularization,” *IEEE Transactions on Geoscience and Remote Sensing*, 2013.
- [6] J. Chen, C. Richard, and A. O. Hero, “Nonlinear unmixing of hyperspectral images using a semiparametric model and spatial regularization,” in *Proc. IEEE ICASSP*, 2014.

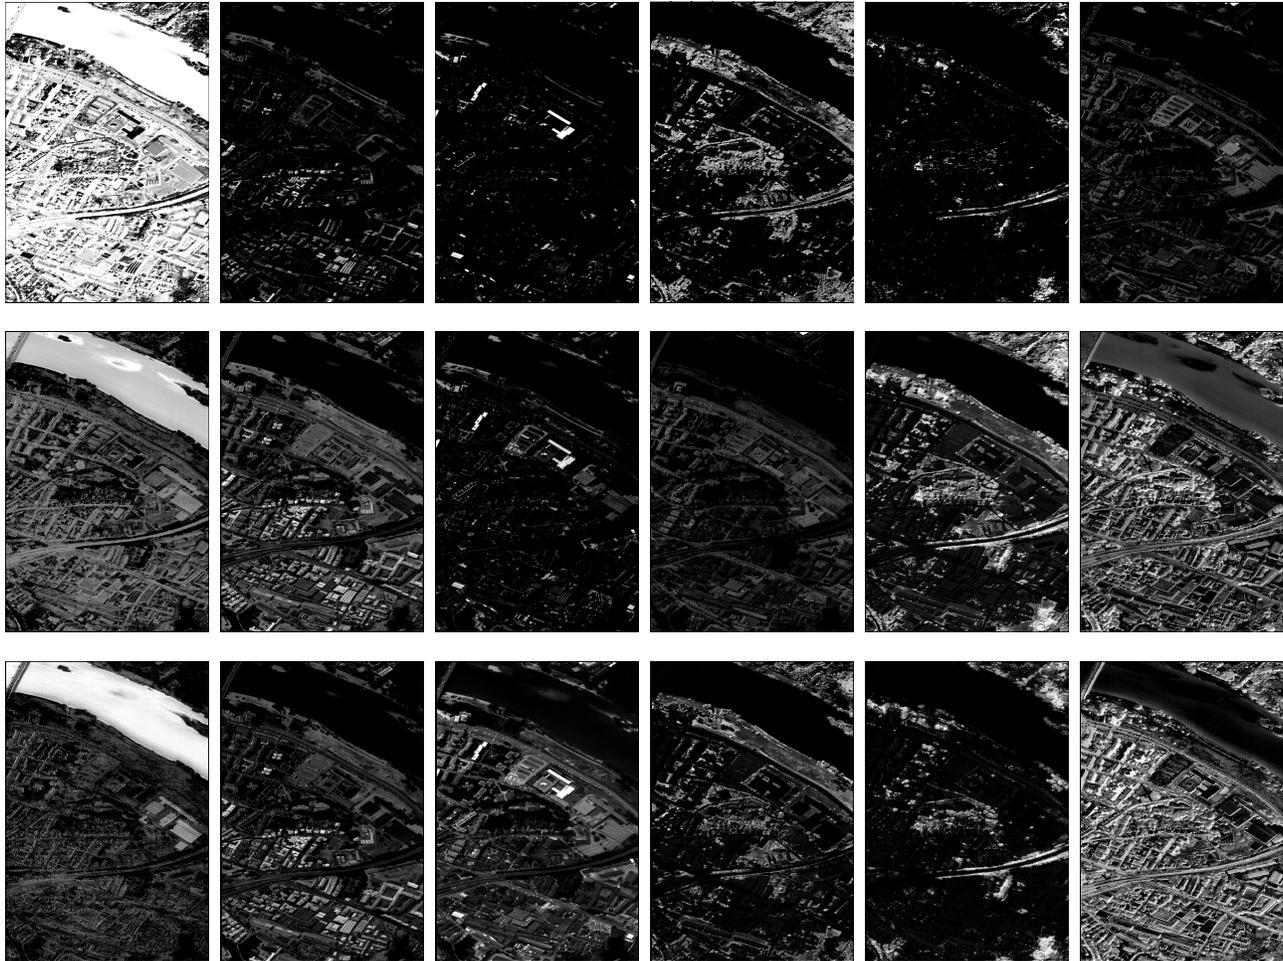


Fig. 2. Abundance maps for Pavia center image obtained with, from top to bottom: SDSOMP, NFINDR, GLUP.

- [7] N. Dobigeon, J.-Y. Tourneret, C. Richard, J.-C. M. Bermudez, S. McLaughlin, and A. O. Hero, "Nonlinear unmixing of hyperspectral images: Models and algorithms," *IEEE Signal Processing Magazine*, vol. 31, no. 1, pp. 82–94, 2014.
- [8] M. E. Winter, "N-FINDR: an algorithm for fast autonomous spectral endmember determination in hyperspectral data," in *Proc. SPIE Imaging Spectrometry*, October 1999.
- [9] C.-I. Chang, C.-C. Wu, W.-M. Liu, and Y.-C. Ouyang, "A new growing method for simplex-based endmember extraction algorithm," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 44, no. 10, pp. 2804–2819, 2006.
- [10] P. Honeine and C. Richard, "Geometric unmixing of large hyperspectral images: A barycentric coordinate approach," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 50, no. 6, pp. 2185–2195, 2012.
- [11] Z. Zhang S. G. Mallat, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [12] X. Fu, W. Ma, T. Chan, J. Bioucas-Dias, and M. Iordache, "Greedy algorithms for pure pixels identification in hyperspectral unmixing: A multiple-measurement vector viewpoint," in *Proc. EUSIPCO*, 2013.
- [13] D. Iordache, J. Bioucas-Dias, and A. Plaza, "Collaborative sparse regression for hyperspectral unmixing," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 52, no. 1, pp. 341–354, 2013.
- [14] M. Yuan and Y. Lin, "Model selection and estimation in regression with grouped variables," *Journal of the Royal Statistical Society: Series B (Statistical methodology)*, vol. 68, no. 1, pp. 49–67, February 2006.
- [15] E. Esser, M. Moller, S. Osher, G. Sapiro, and J. Xin, "A convex model for nonnegative matrix factorization and dimensionality reduction on physical space," *IEEE Transactions on Image Processing*, vol. 21, no. 7, pp. 3239–3252, 2012.
- [16] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [17] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, no. 1, pp. 293–318, 1992.
- [18] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *arXiv:0912.3522*, 2009.