

Parcimonie, diversité morphologique et séparation robuste de sources

Thèse de doctorat de l'Université Paris-Saclay
préparée à l'Université Paris-Sud

École doctorale n°580 Sciences et technologies de l'information et de
la communication (STIC)
Spécialité de doctorat: Traitement du signal et des images

Thèse présentée et soutenue à Saclay, le 29 Septembre 2017, par

Cécile Chenot

Composition du Jury :

Pierre Chainais Professeur, Ecole Centrale de Lille	Rapporteur
Jean-Yves Tourneret Professeur, INP-ENSEEIH Toulouse	Rapporteur
Emilie Chouzenoux Maître de Conférences, Université Paris-Est Marne-La-Vallée	Examinatrice
Mike Davies Professeur, Université d'Edimbourg	Examineur
Thomas Rodet Professeur, ENS Paris Saclay	Président
Jérôme Bobin Chercheur, CEA Saclay	Directeur de thèse
Lenka Zdeborová CNRS, CEA Saclay	Invitée

Acknowledgements

I would like to express my appreciation to the members of my jury, E. Chouzenoux, M. Davies, T. Rodet and L. Zdeborová for their commitment. I am particularly grateful to P. Chainais and J.-Y. Tournet for reviewing thoroughly this manuscript and their valuable comments.

J'adresse mes plus sincères remerciements à mon directeur de thèse Jérôme Bobin pour ces trois années. Je lui suis reconnaissante de sa disponibilité, ses conseils avisés, sa confiance et son enthousiasme au cours de ces trois dernières années. Nul ne peut se prévaloir d'une parfaite compréhension des problèmes de factorisation de matrices... aussi je remercie Jérôme de m'avoir aidée à appréhender ces différentes problématiques avec ouverture d'esprit et intuition.

Je remercie également C.Gaonach, N.Aucuit et E.Cassan de l'IUT d'Orsay pour m'avoir encadrée pendant mon service d'enseignement et rendu cette expérience agréable; ainsi que N.Dobigeon et P.A.Thouvenin de l'IRIT pour nos discussions ces derniers mois.

Mes plus chaleureux remerciements sont adressés aux membres du laboratoire Cosmostat. En particulier, je remercie J.-L. Starck pour son accueil et son enthousiasme, et Florent Sureau pour son aide précieuse et nombreux conseils. Mes pensées vont aussi aux post-docs et doctorants du laboratoire et à nos activités extra-CEA - Austin, Kostas, Ming, Linc (~rou), Arnau, Mel, Sam, Christophe, Joana, plus récemment Jiaxin, Axel et Imane et mes contemporains de troisième année au SAp, Jérémy F. et Bilal. Je remercie Fred, François et Jérémy pour leurs nombreux conseils et réponses, ainsi que Morgan pour sa patience durant les monologues délibératifs dont il fût le témoin (le fameux effet coin-coin).

Je profite aussi de cette occasion pour remercier plus généralement mes proches et membres de la famille. Notamment, j'adresse mes humbles remerciements à Benjamin, Pierrick et Yoann, dont j'ai pu 'jouir du soutien indéfectible' [BM17] au quotidien. Je remercie mes parents de m'avoir encouragée à faire mes propres choix, et toujours soutenue. Enfin, mes plus tendres remerciements vont à Mathieu pour m'avoir secondée et soutenue ces derniers mois, et bien plus encore.

Contents

Résumé	III
Notations and Acronyms	V
1 Sparse Modeling and Blind Source Separation	1
1.1 Multivalued data analysis with BSS	2
1.2 Sparse modeling	7
1.3 Sparse BSS/GMCA	12
1.4 Organization of the manuscript	17
2 Robust Blind Source Separation	19
2.1 Outliers	19
2.2 Robust BSS in the literature	28
2.3 Summary	37
3 Algorithmic framework for sparse Matrix Factorization	39
3.1 Proximal algorithms	40
3.2 Block Coordinate algorithms for component separation problems	44
3.3 Heuristics for sparse BSS	56
3.4 Summary and extensions	73
4 Robust BSS in the presence of sparse outliers	75
4.1 Robust sparse BSS: recoverability	76
4.2 Robust GMCA	80
4.3 Robust AMCA	86
4.4 Numerical Experiments	96
4.5 Summary and future works	104
5 Robust BSS in transformed domains	107
5.1 Separation between the outliers and the sources	108
5.2 Robust GMCA in transformed domains	113

5.3	1D Monte-Carlo Simulations	120
5.4	Application to simulated astrophysical data	129
5.5	Summary and future works	136
6	Hyperspectral unmixing with spectral variabilities	137
6.1	Hyperspectral unmixing	138
6.2	Indeterminacy of the models	145
6.3	Algorithm	153
6.4	Numerical Experiments	157
6.5	Summary and perspectives	167
	Conclusion	169
	Appendix	173
	A Proximal Operators	173
	B Threshold value	174
	B.1 Noise removal/detection with sparse modeling	174
	B.2 Noise removal for the sources	175
	B.3 Noise removal for the outliers	177
	C Performance criteria	179
	D Publications	181
	Bibliography	183

Résumé

La Séparation Aveugle de Sources (SAS) est un outil privilégié pour l'analyse de données multivaluées. Elle cherche à estimer des signaux élémentaires dont différents mélanges sont observés. Cette approche non-supervisée permet d'extraire des composants interprétables des observations. La plupart des méthodes de SAS supposent que chacune des observations correspond à un mélange pondéré inconnu des sources recherchées. Toutefois, ce modèle est trop élémentaire pour rendre compte fidèlement de la complexité des données: la présence de phénomènes physiques inattendus, de bruit structuré ou de dysfonctionnements des capteurs, bien que fréquents en pratique, ne sont pas correctement modélisés. Cela fausse les méthodes classiques de SAS et nécessite le développement de stratégies robustes.

Nous nous intéressons dans cette thèse au développement de nouvelles méthodes capables d'estimer la matrice mélange, ainsi que de séparer précisément les données aberrantes des sources. Ces deux sous-problèmes ont été abordés dans la littérature de façon séparée, avec des hypothèses restrictives sur les composants qui ne sont pas toujours valables, notamment en imagerie multispectrale. Nous montrons que résoudre ces problèmes simultanément permet d'obtenir des méthodes plus robustes et précises, et notamment performantes pour un plus grand éventail de configurations incluant le cas déterminé. Pour cela, nous utilisons des représentations parcimonieuses des différents composants qui mettent en valeur les structures spectrales ainsi que spatiales des différents composants. C'est précisément en exploitant, de façon jointe, les différences spectrales et spatiales des composants à estimer qu'une séparation précise est possible. La première partie des travaux porte sur le démélange de sources ayant une morphologie similaire à celle des données aberrantes. Les déviations sont détectées grâce au degré de parcimonie des signaux estimés, et pénalisées lors du démélange. Cette approche est ainsi capable d'estimer de façon robuste la matrice de mélange, même si une séparation précise entre les sources et données aberrantes n'est pas possible sans contrainte restrictive.

Dans un deuxième temps, nous supposons que les sources et les déviations ont des morphologies différentes. La diversité morphologique est exploitée en utilisant des dictionnaires différents pour les représentations parcimonieuses des composants. Ainsi, en utilisant alternativement les diversités spectrales et spatiales entre les sources et les données aberrantes, l'algorithme proposé obtient une estimation précise de tous les composants et une plus grande robustesse pour les

applications complexes.

La dernière partie de cette thèse porte sur l'extension des travaux précédents pour l'estimation des variabilités spectrales en imagerie hyperspectrale terrestre. En effet, le modèle linéaire classique considère que les composants recherchés ont des signatures spectrales constantes sur l'ensemble de la région observée, alors que celles-ci varient potentiellement significativement. Dans l'étude préliminaire menée, nous proposons de contraindre les signatures spectrales afin qu'elles appartiennent à l'hypersphère. Cela constitue une hypothèse classique en SAS, mais qui n'avait pas été utilisée pour contraindre les modèles non-stationnaires. Le degré de parcimonie des signaux sur les données traitées ne permet pas leur estimation exacte, mais les résultats obtenus rivalisent avec les méthodes de l'état-de-l'art sur certains critères.

Les différents problèmes de minimisation associés aux approches proposées sont multi-convexes et paramétriques, ce qui les rend difficilement utilisables en pratique. Le développement d'algorithmes fiables et automatiquement paramétrés est ainsi nécessaire mais assez délicat. Les mécanismes de gestion de paramètres pour les régularisations avec contraintes de parcimonie sont étendus pour la SAS robuste. Cela confère aux algorithmes une certaine adaptabilité ainsi qu'une robustesse aux différentes corruptions et minima locaux. Cette estimation jointe des composants et des paramètres de régularisation, en particulier en présence de données aberrantes, nécessite plus largement le déploiement d'heuristiques couplées à des méthodes classiques pour garantir robustesse, fiabilité, précision et convergence des procédés de minimisation.

Des expérimentations montrent la fiabilité des méthodes et algorithmes dans de nombreuses configurations. Notamment, les approches proposées obtiennent de bons résultats dans le cas déterminé ainsi que dans le cadre d'applications difficiles, telle que la séparation d'émissions galactiques en présence de sources ponctuelles dans le contexte de la mission Planck.

Notations and Acronyms

Conventions

Matrix

- x : a scalar.
- \mathbf{x} : a vector.
- \mathbf{X} : a matrix.
- \mathbf{X}_i : the i th row of \mathbf{X} .
- \mathbf{X}^j : the j th column of \mathbf{X} .
- $\mathbf{X}_{i,j}$ the (i,j) th entry of \mathbf{X} .
- \mathbf{X}^T : transpose of \mathbf{X} .
- \mathbf{X}^\dagger : Moore-Penrose pseudo inverse of \mathbf{X} .
- $\mathbf{X}^{(k)}$: the value of \mathbf{X} in the k th loop of an algorithm.
- $\tilde{\mathbf{X}}$: estimated value of \mathbf{X} .

Operators

- $\|\mathbf{X}\|_p$, $p \in \mathbb{R}_+$: the (quasi)-norm of \mathbf{X} seen as a long vector. More precisely, for $p \in \mathbb{R}_+$ and $\mathbf{X} \in \mathbb{R}^{m \times t}$, $\|\mathbf{X}\|_p = \left(\sum_{i=1}^m \sum_{j=1}^t |\mathbf{X}_{i,j}|^p \right)^{\frac{1}{p}}$. In particular, $\|\mathbf{X}\|_2$ denotes the Frobenius norm of \mathbf{X} .
- $\|\mathbf{X}\|_0$: the number of non-zeros entries of \mathbf{X} .
- $\|\mathbf{X}\|_{s,2}$: spectral norm of \mathbf{X} (largest eigenvalue of \mathbf{X}).
- $\|\mathbf{X}\|_*$: nuclear norm of \mathbf{X} (sum of the eigenvalues of \mathbf{X}).
- $\|\mathbf{X}\|_{2,1} = \sum_{i=1}^t \|\mathbf{X}^i\|_2$, with $\mathbf{X} \in \mathbb{R}^{m \times t}$.
- \odot : Hadamard product.
- \otimes : tensor product.
- $*$: convolution product.
- $\mathbf{X} \geq 0$: $\mathbf{X}_{i,j} \geq 0$, $\forall i, j$
- $\langle \cdot, \cdot \rangle$: scalar product.
- $[\mathbf{X}]_+$: non-negative part of \mathbf{X} ($[[\mathbf{X}]_+]_{i,j} = 0$ if $[[\mathbf{X}]_+]_{i,j} \leq 0$ and $[\mathbf{X}]_{i,j}$ otherwise).

- prox_f : proximal operator of f .
- $\mathcal{S}_\Lambda(\mathbf{X})$: soft-thresholding operator $\mathcal{S}_\Lambda(\mathbf{X})_{i,j} = \max(0, |\mathbf{X}|_{i,j} - \Lambda_{i,j}) \times \text{sign}(\mathbf{X}_{i,j})$.
- $\chi_{\mathcal{C}}(\cdot)$: indicator function of the set \mathcal{C} . For every $\mathbf{M} \in \mathcal{C}$, $\chi_{\mathcal{C}}(\mathbf{M}) = 0$, and $+\infty$ otherwise.
- mad : median absolute deviation.
- Pr : probability.
- $\text{pct}(\mathbf{x}, k)$: k th percentile of the entries of \mathbf{x} .

Notations

- \mathbb{R} : the set of real numbers.
- \mathbb{R}_+ : the set of non-negative real numbers.
- $\bar{\mathbb{R}}$: extended real numbers.
- \mathbb{N}_+ : the set of non-negative natural numbers.
- $\mathbf{1}_{p \times k}$: the matrix of size $p \times k$ whose entries equal 1.
- $\mathbf{I}_{p \times p}$: the identity matrix of size $p \times p$.
- n : number of sources.
- m : number of observations.
- t : number of samples per observation/source.
- $\mathbf{X} \in \mathbb{R}^{m \times t}$: the observations.
- $\mathbf{A} \in \mathbb{R}^{m \times n}$: the mixing matrix.
- $\mathbf{S} \in \mathbb{R}^{n \times t}$: the sources.
- $\mathbf{O} \in \mathbb{R}^{m \times t}$: the outliers.
- $\Delta \mathbf{A}_i \in \mathbb{R}^{m \times t}$: the spectral variabilities associated with the i th component.
- $\Gamma(\cdot)$: function Gamma.
- $\mathcal{G}(\rho)$: generalized Gaussian with parameter shape ρ , centered and unit variance whose probability density function is given by $\frac{\rho}{2\Gamma(\frac{\rho}{2})} \exp^{-|x|^\rho}$, $\forall x \in \mathbb{R}$

Abbreviations

- BSS: Blind Source Separation
- ICA: Independent Component Analysis
- NMF: Nonnegative Matrix Factorization
- MCA: Morphological Component Analysis
- GMCA: Generalized Morphological Component Analysis
- rGMCA: robust GMCA
- tr-rGMCA: robust GMCA in transformed domains
- AMCA: Adaptive MCA

- rAMCA: robust AMCA

- PALS: projected alternated least squares
- PALM: proximal alternating linearized minimization
- BCD: block coordinate descent
- PBC: proximal block coordinate
- FB: forward-backward splitting algorithm
- GFBS: generalized forward-backward splitting algorithm

- rPCA: robust Principal Component Analysis
- PCP: Principal Component Pursuit
- β -div.: minimization of the β -divergence
- OP: Outliers Pursuit

- SV: spectral variabilities
- MDP: morphological diversity principle
- KL: Kullback-Leibler

- LC/MS: liquid chromatography mass spectrometry
- CMB: cosmic microwave background
- NMR: nuclear magnetic resonance

- SIR-SNR-SAR-SDR: Signal to Interference-Noise-Artifact-Distortion Ratio

Chapter 1

Sparse Modeling and Blind Source Separation

Summary

1.1 Multivalued data analysis with BSS	2
1.1.1 Illustrations	3
1.1.2 Linear Mixture Model	4
1.1.3 Blind Source Separation	6
1.2 Sparse modeling	7
1.2.1 Sparse modeling	8
1.2.2 Morphological diversity	10
1.3 Sparse BSS/GMCA	12
1.3.1 Rationale of sparse BSS	13
1.3.2 Generalized Morphological Component Analysis	15
1.4 Organization of the manuscript	17

We moved into a new era in which the amount and variety of data to be analyzed are tremendous: recommender systems, genomic sequencing, terrestrial remote sensing are but examples. The future missions in astrophysics including Euclid or S.K.A are revealing examples of the forthcoming challenges driven by the volume of the observations and the expected precision reached from their processing. The unsupervised extraction of meaningful, interpretable information from the produced data is one of them.

In this thesis, we will propose new robust methods for the analysis of multivalued data taking advantages of sparse modeling. This chapter aims at introducing these two founding topics.

1.1 Multivalued data analysis with BSS

Multivalued data are nowadays encountered in various domains of applications such as bioengineering [114], astrophysics [16], audio [130], terrestrial remote-sensing [10] and topic modeling [82] to only cite a few. More precisely, we designate by multivalued data, the observations of a same process, region etc. with different but coherent points of view. These data will be designated by $\mathbf{X} \in \mathbb{R}^{m \times t}$, where t denotes the number of samples of one observation, and m the number of observations.

For instance, a standard digital picture corresponds to multivalued data since a same scene is observed in three bands, red, green and blue - see for instance fig.1.1a. In that case, each \mathbf{X}_i corresponds to the scene acquired at one given frequency (blue, red or green), so that $m = 3$ (the total number of bands), and t is equal to the number of pixels of the image ($t = 512 \times 512$ in fig.1.1a).




Figure 1.1: From left to right: color image, acquisition in the blue band, red band and then green band.

We notice that the observations, along the different channels - $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3$ in fig.1.1a, are not independent: each sample \mathbf{X}^k , $k = 1, \dots, t$ corresponds to the sampled spectrum of the element at the k th pixel. Hence, supposing that a pepper has a constant spectrum, every pixel corresponding to this pepper has the same spectrum. As so, one can remark that these multivalued data are spectrally and spatially structured.

Blind Source Separation (BSS) aims to extract the meaningful information from these observations, by learning this underlying structure.

In this thesis, we will use 3 illustrative applications of multivalued data analysis, which are described below. Similarly to the digital picture, these data are spatially and spectrally (for the different samples and channels respectively in a more general framework) structured.




Figure 1.2: Illustration of the underlying data model with the Planck data.

1.1.1 Illustrations

Estimation of the CMB and foreground emissions from ESA-Planck data. According to the Big-Bang theory, the CMB (Cosmic Microwave Background) is the oldest observable of our Universe, the first light emitted 13 billions years ago. As so, it carries fundamental information on the early Universe and its evolution for cosmologists.

In 2009, the European Spatial Agency has launched the Planck satellite in order, among other things, to obtain a precise CMB map [37]. For this purpose, the satellite has observed the sky at $m = 9$ different frequencies (microwave range from 30GHz to 857GHz).

However, we do not have a direct access to the CMB map with the Planck data. We rather observe a combination of the different observable contributions fig.1.2: CMB and also galactic emissions such as the synchrotron, free-free, thermal and spinning dust emissions [86]. In order to retrieve the CMB map, it is necessary to separate these different contributions. This component separation problem is a Blind Source Separation problem [86], [16].

Hyperspectral unmixing. Thanks to the rapid development of hyperspectral sensors, the analysis of hyperspectral data has become a topical issue. In contrast with multispectral data (such as the ones of Planck), hyperspectral data consists of a hundred of spectral acquisitions ($m \approx 100$ versus $m \approx 10$ for multispectral). Consequently, it provides a very fine spectral description of the components and can detect slight spectral changes. This is of interest in various applicative domains such as medicine to detect skin cancer [96], in food engineering to automatically check the quality of the products [109] or in agriculture/ecology [75].

In the same spirit as for the analysis of the Planck data, it is possible to estimate the spectral signatures and the spatial distributions of the different elements composing the observed scene using BSS [80], [9], [98]. For instance in fig.1.3, hyperspectral unmixing can be used to detect

and estimate the presence of water at the surface of Mars [102].




Figure 1.3: Left: Mars Express data from the OMEGA spectrometer. Middle and right: top: estimated spectral signatures of the water and CO_2 , bottom: spatial distributions of these two components. Results from [102].

LC-MS data analysis. The types of the three previous examples (digital picture, Planck data, Omega observations) are similar: a same region/scene is observed at several frequencies. The following example has a different nature and is illustrative of the broaden variety of multivalued data.

The analysis of Liquid Chromatography - Mass Spectrometry (LC-MS) data permits a precise identification and quantification of the components of a liquid. This fluid is first driven through a chromatographic column. The compounds of the fluid leave the column at different instants, based on their properties. The output of a liquid chromatography, at different time samples, is analyzed with a mass spectrometer. We obtain a double separation fig.1.4a: in time with the chromatography and in mass to charge ratio with the mass spectroscopy (similarly to multi/hyperspectral images with the double separation: frequency-space domains). By estimating the elution times and mass to charge ratio of the different components with BSS, fig.1.4b,1.4c, one can recover the composition of the analyzed fluid.

1.1.2 Linear Mixture Model

In BSS, it is generally assumed that the observations follow the so called Linear Mixture Model. It states that the m observations \mathbf{X} (e.g. the 9 observations at different wavelengths for the Planck data) correspond to the sum of n elementary contributions: the one of the CMB, of the synchrotron and so on, fig.1.2. Besides, each of these elementary contributions is rank-1: in the




Figure 1.4: Top: acquisition process and example of LC-MS data \mathbf{X} for a given liquid. Bottom, from left to right: the elution profile and mass to charge ratio of a component. Figures from [115].

example fig.1.2, the i th contribution can be factorized as the product between its spectral signature $\mathbf{A}^i \in \mathbb{R}^{m \times 1}$ and its spatial distribution-intensity map, called source, $\mathbf{S}_i \in \mathbb{R}^{1 \times t}$.

In other words, each observation \mathbf{X}_j corresponds to a weighted linear combination of the n different sources $\{\mathbf{S}_i\}_{i=1..n}$, where n designates the number of components: $\mathbf{X}_j = \sum_{i=1}^n \mathbf{A}_j^i \mathbf{S}_i$. The LMM model is illustrated in fig.1.5 for the general framework, which is similar to the figure 1.2 representing the underlying structure of the Planck data.

This model is generally recast in the following matrix form:

$$\mathbf{X} = \mathbf{AS},$$

where

- $\mathbf{X} \in \mathbb{R}^{m \times t}$ stands for the observations. In multi/hyperspectral imaging, each row of \mathbf{X} corresponds to the observed scene at a given frequency, fig.1.3,1.2.
- $\mathbf{A} \in \mathbb{R}^{m \times n}$ contains the mixture weights, and is named mixing matrix. Each column of \mathbf{A} represents the variations of a component across the different channels. For instance in multi/hyperspectral imaging, fig.1.3, 1.2, each column of \mathbf{A} represents the spectrum/endmember of one component (e.g. ice water spectra in fig.1.3, or the spectral signature of the CMB in fig.1.2). When analyzing LC-MS data, the columns of \mathbf{A} equal the elution profiles of the components fig.1.4.




Figure 1.5: Illustration of the underlying structure of multivalued data with the LMM model.

- $\mathbf{S} \in \mathbb{R}^{n \times t}$, the sources, corresponds to the variability of the components from one sample to another. A so-called source, also named abundance in hyperspectral unmixing, represents the spatial distribution of a component over the observed region in multi/hyperspectral imaging (e.g. the CMB map in fig.1.2 or the ice water spatial distribution in fig.1.3).

An accurate estimation of the mixing matrix \mathbf{A} and the sources \mathbf{S} leads to a precise analysis of the multivalued data: estimating the natures/spectra of the elements at the surface of the observed scene amounts to recover \mathbf{A} , and evaluating their spatial distributions to determine \mathbf{S} . *Estimating both \mathbf{A} and \mathbf{S} only from the observations \mathbf{X} is a BSS problem.*

1.1.3 Blind Source Separation

Blind Source Separation methods aim to find a tuple (\mathbf{A}, \mathbf{S}) such that $\mathbf{X} = \mathbf{A}\mathbf{S}$. Unfortunately, this is an ill posed problem since there is an infinite number of tuples (\mathbf{A}, \mathbf{S}) fulfilling this condition. For instance, for any invertible matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$, the tuple $(\mathbf{A}\mathbf{M}, \mathbf{M}^{-1}\mathbf{S})$ is also a solution.

Nonetheless, we are not interested in recovering any tuple (\mathbf{A}, \mathbf{S}) , but the one with having a physical meaning, e.g. the spectra and spatial distributions of the components in multi/hyperspectral unmixing. In order to retrieve this meaningful tuple, additional information should be added to discriminate between the solutions. Depending on the priors, the BSS strategies can be divided into three classes: Independent Component Analysis (ICA), Non-negative Matrix Factorization (NMF), and sparse Blind Source Separation.

Independent Component Analysis. In contrast with sparse BSS and NMF, ICA is a statistical approach, similarly to PCA. While PCA only returns decorrelated sources based on second order statistics, the exact separation between the sources can be achieved with ICA [39], [27], [78], [40].

It is assumed that the sources are mutually independent and that at most one source follows a Gaussian distribution, with $m = n$. In that case, it has been proven that there exists a unique unmixing matrix \mathbf{B} so that $\mathbf{S} = \mathbf{B}\mathbf{X}$, [39] - up to scaling and permutation indeterminacies.

In practice, the observations are first whitened with PCA, (yielding $m = n$), the intermediate sources have a unit variance and are uncorrelated. Then, the ICA methods look for the unmixing matrix - a rotation - whose corresponding sources are mutually independent. Several implementations have been proposed: minimization of the mutual information (using the Kullback-Leibler divergence for instance [28]), or with other contrast functions providing approximations for the mutual information - see [40] for a review.

These methods do not explicitly estimate the sources but only the unmixing matrix \mathbf{B} . As pointed out for instance in [43], they do not necessarily yield a good estimate of the sources in the presence of noise. More critically, the mutual independence of the sources is not a valid hypothesis in several applications, including for the Planck data (the foreground emissions are in fact quite correlated [16]), or for terrestrial hyperspectral unmixing [103] to only cite two examples.

Nonnegative Matrix Factorization. Sparse BSS and NMF based methods are special cases of the matrix factorization problem. In NMF, the sources and mixing matrix are both non-negative, [87], [106]. The non-negativity of the components appears naturally in many applications, such as in terrestrial remote sensing (the sources correspond to the spatial distributions of the components and \mathbf{A} to the spectral signatures), or with the LC/MS data previously presented in this chapter.

Nonetheless, non-negativity is not a discriminative criterion for the separation, and further assumptions are needed to recover the initial \mathbf{A}, \mathbf{S} . In particular, it has been proven in [46] that the model is identifiable if every source is active alone (presence of pure pixels). This condition is however restrictive and not valid in many applications. To overcome this recoverability issue, the non-negativity assumption can be coupled with another prior such as sparsity [113].

Sparse Blind Source Separation. This last strategy assumes that the sources are sparsely represented in a given dictionary. Sparse BSS has first been presented in the Bayesian framework in [89] and [144] with a Laplacian prior on the source coefficients. It is nonetheless not restricted to this framework [66], [14]. Besides, in contrast with ICA, it is better suited to deal with the under-determined case [68], or Gaussian noise [14]. Further details on sparse modeling and sparse BSS are given in the remaining of this chapter.

1.2 Sparse modeling

Sparse modeling has a central role in many breakthrough works in signal processing of the last twenty years. We will focus in this chapter on the necessary background for the remaining of this thesis, and refer the reader to [124], [99], [52] and references therein for further details and examples of applications building upon sparse modeling.

1.2.1 Sparse modeling

Sarsity. Let $\mathbf{x} \in \mathbb{R}^{1 \times t}$ a real signal composed of t samples. We will say that \mathbf{x} is exactly sparse if only few entries of \mathbf{x} are non-zero. The number of non-zero entries of \mathbf{x} is given by its ℓ_0 ‘norm’. For instance in fig.1.6a , \mathbf{x} is exactly sparse, and $\|\mathbf{x}\|_0 = 18 \ll t$.

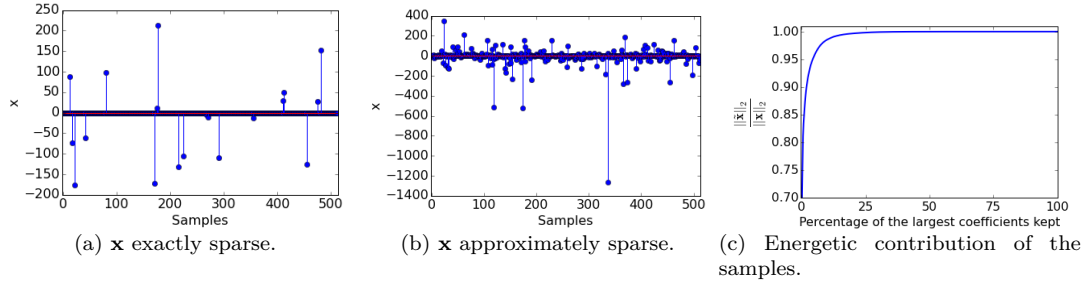


Figure 1.6: Example of exactly sparse (left) and approximately sparse (middle) signals. Right: energetic contribution of the samples for the approximately sparse signal; we display $\frac{\|\tilde{\mathbf{x}}\|_2}{\|\mathbf{x}\|_2}$, where $\tilde{\mathbf{x}}$ is composed of only a given percentage (x-axis) of the largest entries of \mathbf{x} .

If the signal \mathbf{x} is composed of only few significant samples, then \mathbf{x} is said to be approximately sparse, such as in fig.1.6b. In that case, $\|\mathbf{x}\|_0 = t$, but only few coefficients are needed to provide a good approximation of the signal (similarly for the exactly sparse signal, whose energy is contained in only 18 samples in fig.1.6a). For example, the approximately sparse signal of fig.1.6b can be approximated with an exactly sparse signal with less than 10% of non-zero entries, with an energetic error of less than 5% fig.1.6c. In other words, all the information content of sparse or approximately sparse signals is contained in only few entries.

Sparse signals are naturally met in spectrometry (c.f. the previous LC-MS data) for example. And more interestingly, many signals can be sparsely represented using an adequate transformation (broadly, a change of basis).

Sarsity with transformed domain. Most of the signals of interest are not sparse in the domain of observations (for instance the galactic emissions for the Planck data previously seen). Fortunately, these signals can be sparsely represented in a transformed domain.

For illustrative purpose, let us consider an easy example. Let \mathbf{x} be a 1D oscillating signal fig.1.7a. This signal can be sparsely represented using an orthonormal transformation, the DCT:

$$\mathbf{x} = \sum_{i=1}^t \alpha_i \Phi_{\mathbf{S}}^i = \alpha \Phi_{\mathbf{S}},$$

where $\alpha \in \mathbb{R}^{1 \times t}$ corresponds to the expansion coefficients of \mathbf{x} in the DCT dictionary $\Phi_{\mathbf{S}} \in \mathbb{R}^{t \times t}$, and α is composed of only few significant samples fig.1.7b.




Figure 1.7: Top: Left oscillating 1D signal \mathbf{x} , non-sparse in the domain of observations. Right: approximately expansion coefficients of \mathbf{x} in DCT, α . Bottom, from left to right: synchrotron emission map, histogram of the pixels of the emission in the direct domain, and then histogram of the expansion coefficients of the emission in wavelets.

A more realistic example is also provided in fig.1.7. A spatial distribution map of synchrotron emission is displayed in fig.1.7c. It does not have a sparse representation in the direct domain: we can see that numerous entries are significant fig.1.7d. However, this image has a sparse representation in the wavelet domain (undecimated Daubechies-4, with 2 scales) fig.1.7e: only few entries of the corresponding α are significant.

The sparsity level of the expansion coefficients is highly dependent on the chosen dictionary, transformation for a given signal. Besides, we point out that these dictionaries may not be orthonormal and using over-complete/redundant dictionaries can be of interest to sparsify the representations [122], [36]. Some widely employed transformations include: DCT for oscillating texture, wavelets for smooth piece-wise functions [42], [122] such as many natural images (e.g. galactic emissions for the Planck data), curvelets and ridgelets for lines, curves and cartoons [120], [22] to only cite a few.

1.2.2 Morphological diversity

The sparsity prior on the signal to be retrieved in a given dictionary has been intensively used in many ill-posed inverse problems: denoising, super-resolution, inpainting or compressive sensing for instance [52], [124]. Sparse modeling is also a powerful concept for Blind Source Separation and component separation in general since it highlights the difference between the sources. This principle has been designated as morphological diversity [121], [54], [12] and is described below.

Transformed domains. The morphological diversity principle has been first stated and employed in the framework of monochannel component separation, for components having different geometrical features [123], [15], [54].

For illustrative purpose, we consider two images having different geometrical features: a map of synchrotron fluctuations \mathbf{x}_1 and point source emissions \mathbf{x}_2 , fig.1.8. The point source emissions correspond to Dirac emissions convolved with a Gaussian kernel with known width. These two images are summed in the observation $\mathbf{x} = \mathbf{x}_1 + \mathbf{x}_2$.




Figure 1.8: From left to right: monochannel observation made of the sum of a natural image (synchrotron fluctuation map $\tilde{\mathbf{x}}_1$) and spurious point sources (Gaussian emissions $\tilde{\mathbf{x}}_2$). The components $\tilde{\mathbf{x}}_1$ and $\tilde{\mathbf{x}}_2$ have been retrieved with the MCA algorithm.

The two images have different morphologies. The synchrotron emission \mathbf{x}_1 is quite broadly distributed in the observed region, and is sparsely distributed in wavelet domain (undecimated Daubechies-20 for the experiment), noted Φ_S . On the other hand, the point source emissions \mathbf{x}_2 are already sparse in the direct domain, and even sparser in the convolutive dictionary Φ_O (we can build implicitly the dictionary Φ_O by deconvolving the emissions with the known Gaussian kernel).

The key point is that, if the components to be separated have a different morphology and if their corresponding dictionaries in which they are sparsely represented are wisely chosen, then

each morphological component has its sparsest decomposition in its own dictionary. Namely, the image \mathbf{x}_1 is more sparsely represented in $\Phi_{\mathbf{S}}$ than in $\Phi_{\mathbf{O}}$, and \mathbf{x}_2 is sparsest in $\Phi_{\mathbf{O}}$ than in $\Phi_{\mathbf{S}}$. The morphological diversity has been more formally stated as [14]:

$$\|\mathbf{x}_1\Phi_{\mathbf{S}}^T\|_0 < \|\mathbf{x}_1\Phi_{\mathbf{O}}^T\|_0 \quad \text{and} \quad \|\mathbf{x}_2\Phi_{\mathbf{O}}^T\|_0 < \|\mathbf{x}_2\Phi_{\mathbf{S}}^T\|_0. \quad (1.1)$$

Morphological Component Analysis. It is then possible to separate the different morphological components \mathbf{x}_1 and \mathbf{x}_2 given the observation \mathbf{x} by looking for the sparsest decomposition of the observation in the two dictionaries $\Phi_{\mathbf{O}}$ and $\Phi_{\mathbf{S}}$. The *Morphological Component Analysis* has been designed to perform this separation problem. It assumes that the observation is a linear combination of morphological components $\mathbf{x} = \sum_{i=1}^n \mathbf{x}_i$. The latter have different morphologies, and are sparsely represented in a given dictionary Φ_i . It is further assumed that the morphological diversity principle (MDP) is valid, meaning that the morphological components have a sparser representation in their respective dictionary than in the others eq.1.1. The MCA algorithm is looking for the sparsest decomposition of the image into n elements by minimizing:

$$\underset{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n}{\text{minimize}} \left\| \mathbf{x} - \sum_{i=1}^n \mathbf{x}_i \right\|_2^2 + \sum_{i=1}^n \lambda_i \|\mathbf{x}_i\Phi_i^T\|_0, \quad (1.2)$$

where the first term is the data-fidelity term, well suited to deal with Gaussian noise, and the second terms enforce the sparsity of the morphological components in their respective dictionary. Further details on the setting of the regularization parameters $\{\lambda_i\}$ and implementation can be found in [15] for instance. The output of the MCA algorithm with the synchrotron fluctuation map and the point source emissions are displayed in fig.1.8.

The morphological diversity principle for the separation of morphological components has been also employed for multichannel observations with the multichannel Morphological Component Analysis - BSS problem in which the sources have different morphologies - and to the Generalized Morphological Component Analysis - BSS problem in which the sources can share a same morphology [14]. For this last case, the morphological diversity principle needs to be extended to also consider the case of signals with similar geometrical features.

In a same transformed domain. The morphological diversity principle has also been extended for BSS in the presence of sources sharing a same morphology and sparsely represented in a same dictionary $\Phi_{\mathbf{S}}$. In that case, the sources share similar geometrical features, and the arguments previously employed to separate morphological components are not valid. That is the case for instance for the two natural images presented in fig.1.9.

Sparse representations can still be employed to highlight the differences between the sources.




Figure 1.9: From left to right: two natural images with similar geometrical contents (e.g. contours, texture), and then some of their wavelet coefficients.

Indeed, the sources are represented by only few significant entries in the dictionary Φ_S . Since these significant entries encode most of the information content of the sources, they are very likely to be active at different positions. This is illustrated in fig.1.9: the two natural images S_1 and S_2 are both sparsely represented in the wavelet domain, but one can notice that indeed, the supports of the largest entries of their sparse representations are disjoint.

We point out that this notion of morphological diversity encompasses the previous case with transformed domains. Indeed, if we concatenate the different dictionaries $\Phi_S = [\Phi_1 \ \Phi_2 \ \dots \Phi_n]$, then we end up with the framework of this more general aspect of the morphological diversity (initial framework of GMCA [14]).

We also underline that if the source entries are identically and independently distributed from a same narrow distribution, e.g. Generalized Gaussian, Laplacian or Bernoulli-Gaussian, then the sources are very likely to respect the MDP. In other words, the unmixing problems handled by ICA based methods (with narrow distribution) can also be performed by sparse BSS strategies.

1.3 Sparse BSS/GMCA

In the remaining of this thesis, we will assume that:

- The number of sources n is known, and that $t \geq m \geq n$.
- The matrices A and S are full-rank.
- The columns of A are normalized for the ℓ_2 norm. This assumption is commonly used to avoid the scaling indeterminacy between A and S ($\forall \alpha > 0$, $(A^i \alpha)(S_i \frac{1}{\alpha}) = A^i S_i$).
- The sources S share the same morphology and are sparsely represented in the given dictionary Φ_S (known a priori). Besides, the sources respect the morphological diversity principle in Φ_S .

These basic assumptions are common in sparse BSS, see for instance [14] and [124]. We

propose in this section to illustrate why the morphological diversity principle is so powerful for BSS [14].

1.3.1 Rationale of sparse BSS

We have previously observed that the BSS problem is ill-posed and that it is necessary to add prior information on the sources to be retrieved. Assuming that the sources are sparsely represented in the dictionary $\Phi_{\mathbf{S}}$, and that the expansion coefficients respect the morphological diversity principle in $\Phi_{\mathbf{S}}$ are efficient hypotheses for solving BSS [14].

In that case, in order to retrieve the sought-after sources, one can seek for the tuple $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$ so that $\tilde{\mathbf{S}}$ has the jointly sparsest representation in $\Phi_{\mathbf{S}}$. Indeed, while the sources are not correctly unmixed, their representations in $\Phi_{\mathbf{S}}$ are less sparse than the ones we aim to recover.

1D toy example. For illustrative purposes, let us consider that 2 sources are sparse in the direct domain, $\Phi_{\mathbf{S}} = \mathbf{I}$, and are mixed with \mathbf{A} into 3 noiseless observations \mathbf{X} , such as fig.1.10. The supports of the two sources are disjoint, so that on the first half of the grid, the entries of \mathbf{S}_1 are the only ones that can be active, while only those of \mathbf{S}_2 can be non-zero on the second half of the grid. We will note supp_{S_1} and supp_{S_2} the supports of \mathbf{S}_1 and \mathbf{S}_2 respectively (the set of non-zero entries of the signals). These two sources respect the MDP.

The ‘estimated’ mixing matrix $\tilde{\mathbf{A}}$ is chosen so that $\tilde{\mathbf{A}}\tilde{\mathbf{S}} = \mathbf{X}$ and $\tilde{\mathbf{S}} = \tilde{\mathbf{A}}^\dagger \mathbf{X}$ (noiseless observations). Hence, the spans of the ground truth \mathbf{A} and $\tilde{\mathbf{A}}$ are the same, and $\mathbf{A}\mathbf{S} = \tilde{\mathbf{A}}\tilde{\mathbf{S}}$. However, one can notice that the two tuples are different, fig. 1.10.

In particular, the estimated sources are clearly less sparse than the sought-after ones: $\|\mathbf{S}_1\|_0 = 10$ and $\|\mathbf{S}_2\|_0 = 12$ whereas $\|\tilde{\mathbf{S}}_1\|_0 = \|\tilde{\mathbf{S}}_2\|_0 = 22$. Indeed, since the observations are not correctly unmixed, the estimated sources correspond to a mixture of the two initial sources fig. 1.10. That is why, the supports of the two estimated sources are given by the union of supp_{S_1} and supp_{S_2} . In other words, the interferences created by an incorrect unmixing lead to less sparse solutions than \mathbf{S} . The only way to reach the sparsity level of \mathbf{S} is to return a mixing matrix $\tilde{\mathbf{A}}$ (such that $\tilde{\mathbf{A}}\tilde{\mathbf{S}} = \mathbf{X}$) avoiding any interference and thus being equal to \mathbf{A} (up to the permutation indeterminacy).

2D toy example. This second example is more realistic. Two images \mathbf{S}_1 and \mathbf{S}_2 are mixed into 2 noiseless observations fig.1.11. Similarly to the previous example, let us choose a mixing matrix $\tilde{\mathbf{A}}$, such that $\mathbf{X} = \tilde{\mathbf{A}}\tilde{\mathbf{S}}$, and $\tilde{\mathbf{S}} = \tilde{\mathbf{A}}^\dagger \mathbf{X}$ (noiseless observations). The corresponding sources $\tilde{\mathbf{S}}$ are displayed in fig.1.11. The source images \mathbf{S} are sparsely represented in the wavelet domain $\Phi_{\mathbf{S}}$. In contrast with the first toy example, the sources are compressible in $\Phi_{\mathbf{S}}$ but not exactly sparse. As we will see, we can relax the use of the ℓ_0 norm, and measure the compressibility level of the sources with the ℓ_1 norm. We display in fig.1.12a, the scatter plot of $\mathbf{S}_1\Phi_{\mathbf{S}}^T$ and $\mathbf{S}_2\Phi_{\mathbf{S}}^T$.

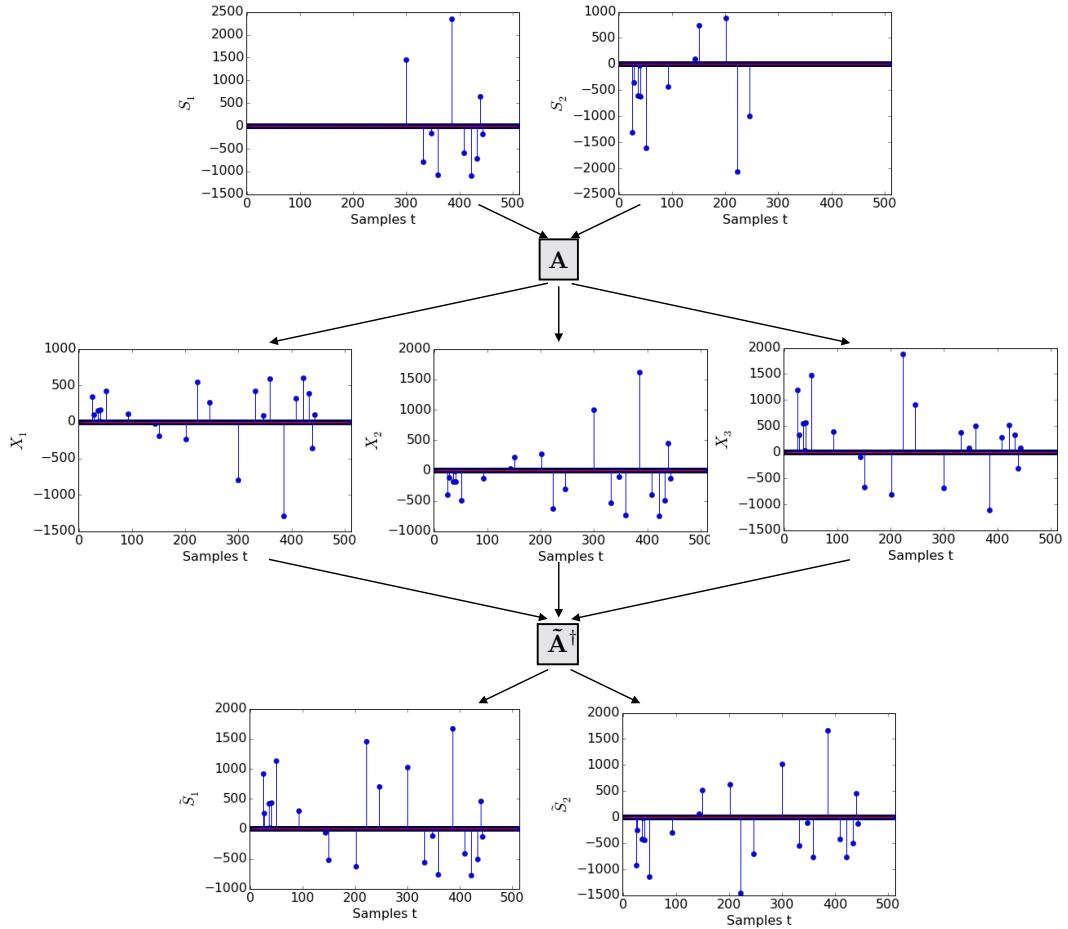


Figure 1.10: Two exactly sparse sources \mathbf{S}_1 and \mathbf{S}_2 are mixed with \mathbf{A} into 3 noiseless observations \mathbf{X} . Two estimated sources $\hat{\mathbf{S}}$, given by $\hat{\mathbf{A}}^\dagger \mathbf{X}$ (noiseless case) are represented. In this case, the tuple $(\hat{\mathbf{A}}, \hat{\mathbf{S}})$ fits perfectly the observations $\mathbf{A}\mathbf{S} = \hat{\mathbf{A}}\hat{\mathbf{S}}$, but the sources $\hat{\mathbf{S}}$ are not the ones that we are looking after.

One can observe that the expansion coefficients $\mathbf{S}\Phi_{\mathbf{S}}^T$ respect the MDP: when one coefficient is significant, it is mainly representative of one source. More generally, when the sources respect the MDP in $\Phi_{\mathbf{S}}$, the scatter plot of the expansion coefficients has a distinctive shape of cross (exactly sparse and disjoint support) or diamond, such as in fig.1.12a.

In contrast, one can notice that the mixed sources $\tilde{\mathbf{S}} = \tilde{\mathbf{A}}^\dagger \mathbf{X} = \tilde{\mathbf{A}}^\dagger \mathbf{A}\mathbf{S}$ do not respect the MDP. Indeed, the largest samples are significant for the 2 sources simultaneously due to the interferences between the 2 sources, fig.1.12b. In practice, we can measure the loss of “compressibility” of the mixed sources with the ℓ_1 norm: the wavelet coefficients $\mathbf{S}\Phi_{\mathbf{S}}^T$ are enclosed in an ℓ_1 ball with a smaller radius than the one enclosing the wavelet coefficients $\tilde{\mathbf{S}}\Phi_{\mathbf{S}}^T$ as shown in fig.1.12. The radius of the ℓ_1 ball of minimal radius enclosing $\mathbf{S}\Phi_{\mathbf{S}}^T$ is broadly given by the maximal entry of $\mathbf{S}\Phi_{\mathbf{S}}^T$ since their largest entries are nearly disjoint thanks to the MDP. On the other hand, the wavelet



Figure 1.11: From the left to the right: the reference 2 sources, the 2 observations, and possible solutions $\tilde{\mathbf{S}}$.

coefficients $\tilde{\mathbf{S}}\Phi_{\mathbf{S}}^T$ are not jointly sparse, and consequently very likely to be enclosed in an ℓ_1 ball with a larger radius (corresponding to the ℓ_1 norm of the entry k having simultaneously large coefficients $(\tilde{\mathbf{S}}_1\Phi_{\mathbf{S}}^T)^k$ and $(\tilde{\mathbf{S}}_2\Phi_{\mathbf{S}}^T)^k$), fig.1.12. Hence, by recovering the sources whose expansion coefficients are enclosed in the ℓ_1 ball of minimal radius, and consequently respecting the MDP, one should avoid the interferences between the sources, and thus recover the mixing matrix of interest.

1.3.2 Generalized Morphological Component Analysis

The GMCA algorithm, proposed in [14], has been designed to perform BSS for sources sparsely represented in an (overcomplete) dictionary $\Phi_{\mathbf{S}}$ (the redundant dictionary can be a union of bases for instance, if the components have different morphologies). The GMCA algorithm performs by minimizing a cost function whose elements are based on the following assumptions:

- Data-fidelity term: The observations are corrupted by a Gaussian noise. The data-fidelity term is consequently chosen as the standard square Frobenius norm: $\frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_2^2$.




Figure 1.12: Left: scatter plot of $\mathbf{S}_1 \Phi_S^T$ versus $\mathbf{S}_2 \Phi_S^T$, the wavelet coefficients of the reference images. Right: scatter plot of $\tilde{\mathbf{S}}_1 \Phi_S^T$ versus $\tilde{\mathbf{S}}_2 \Phi_S^T$, the wavelet coefficients of two possible sources. In green, the ℓ_1 ball with minimal radius enclosing the wavelet coefficients of the reference sources, and in red, the one enclosing the wavelet coefficients of $\tilde{\mathbf{S}}$.

- Compressibility of the sources in Φ_S . The sources have an approximately sparse representation in Φ_S . Different regularization functions can be employed to enforce the sparsity of the sources in Φ_S , namely $\|\Lambda \odot \mathbf{S} \Phi_S^T\|_p$, with $p \in [0, 1]$, and where $\Lambda \in \mathbb{R}^{n \times t}$ denotes the regularization parameters, and \odot the Hadamard product. GMCA has been initially proposed with $p = 0$ and $p = 1$. In this thesis, we will prefer the ℓ_1 norm to the ℓ_0 pseudo-norm for two reasons. First, the ℓ_1 norm is convex (while still enforcing sparsity), and is therefore easier to minimize [48]. Second, it has been observed in [112], that using the ℓ_1 norm favors a solution $\tilde{\mathbf{A}}$ so that the corresponding expansion coefficients $\tilde{\mathbf{S}} \Phi_S^T$ are clustered along such axes fig.1.12.
- Scaling indeterminacy: any BSS problem suffers from two indeterminacies: permutation (there is no order for the sources) and scaling. We will assume that the columns of \mathbf{A} are normalized, which is a standard assumption in BSS [14]. This is enforced using the indicator function $\chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A})^1$.

Therefore, the GMCA algorithm aims at minimizing the following cost function [14]:

$$\underset{\mathbf{A}, \mathbf{S}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_2^2 + \|\Lambda \odot \mathbf{S}\|_1 + \chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}). \quad (1.3)$$

¹The assumption can be relaxed in practice: one can simply enforce the norm of the columns to be smaller than 1 (indicator function of a convex set). Let us consider a fixed product $\mathbf{A}\mathbf{S}$, where \mathbf{A} has normalized columns and let α be a positive scalar. We note $\tilde{\mathbf{A}} = \alpha\mathbf{A}$, and $\tilde{\mathbf{S}} = \frac{1}{\alpha}\mathbf{S}$. For every scalar $\alpha \in (0, 1]$, the columns of $\tilde{\mathbf{A}}$ belong to the unit ball. However, the ℓ_1 norm of the associated sources $\tilde{\mathbf{S}}$ increases when α decreases: for fixed product $\mathbf{A}\mathbf{S}$, the minimum of the cost function is reached when the columns have a unit norm.

In practice, the GMCA algorithm looks for the ℓ_1 ball with minimal radius enclosing the source samples fig.1.12: as we have observed in the previous section, if the source samples are not correctly unmixed, they are not jointly sparse and are consequently enclosed in a ℓ_1 ball with a larger radius than the one needed to enclose the sought-after sources.

Nonetheless, we point out that minimizing the previous cost function is not simple:

- Regularization parameters: The values of the regularization parameters are crucial in sparse BSS (setting Λ to zero results to a matrix factorization problem without any constraint, and Λ to infinity implies that $\mathbf{S} = 0$). Given that cross-validation is not always feasible, the regularization parameters should be set automatically, according to the variable values, which is done by GMCA.
- Non-convex problem: This problem is multi-convex: namely, it is convex when one of the two variables is fixed, but not for the tuple (\mathbf{A}, \mathbf{S}) . Minimizing such cost function is thus challenging since one should avoid local minima.

These two issues require the development of efficient heuristics, such as the ones presented in [14]. They will be introduced and discussed in Chapter 3.

1.4 Organization of the manuscript

This first chapter presents a brief overview of the analysis of multivalued-data by means of sparse BSS. The sparse BSS method GMCA yields accurate results if the standard linear mixture model provides a precise description of the observations. Unfortunately, deviations from the linear model are frequently encountered in practice, and hamper standard BSS methods. Some examples of deviations are presented in Chapter 2, as well as the current state-of-the-art methods for robust matrix factorization/anomaly detection. Chapter 3 presents an introduction to the optimization for matrix factorization problems, as well as the heuristics developed for the choice of the regularization parameters. Then, Chapter 4 focuses on robust BSS in the presence of deviations sharing the same morphology as the sought-after sources. In Chapter 5, we propose another robust BSS method aiming to deal with the presence of deviations having a different morphology from the one of the components of interest. Last, in Chapter 6, the robust BSS methods are extended in order to recover spectral variabilities in terrestrial hyperspectral images.

Chapter 2

Robust Blind Source Separation

Summary

2.1	Outliers	19
2.1.1	Observations deviating from the standard linear model	20
2.1.2	Towards a more realistic data model	23
2.1.3	Influence of the deviations	25
2.2	Robust BSS in the literature	28
2.2.1	Removal of the outliers in the low-rank regime	29
2.2.2	Robust metrics	33
2.2.3	Component Separation	36
2.3	Summary	37
2.3.1	Summary of the different approaches in the literature	37
2.3.2	Identifiability of the outliers	38
2.3.3	Proposed approach: a component separation strategy	38

While the presence of large deviations from the standard linear model is ubiquitous in many real-world applications, most standard BSS methods are sensitive to gross errors. This mandates the development of robust BSS methods, able to handle outliers.

In this chapter, we motivate the developments on robust BSS by presenting some real-world applications for which the observations are not correctly described by the standard linear mixture model. Then, the main approaches proposed in the literature to handle the presence of outliers are reviewed.

2.1 Outliers

The presence of Gaussian noise is generally taken into account in the data model of standard BSS methods, and fairly handled [41], [14], [40]. However, it can only represent dense and relatively

small noise. In numerous applications, large deviations from the linear model $\mathbf{AS} + \mathbf{N}$ are encountered and encompass the presence of unexpected physical events [126], [129], instrumental artifacts [91] or non-linearity of the underlying physical process [58]. More precisely, the so-called deviations act as a significant additional component, structurally different from the source contribution, and spatially structured (in contrast to the Gaussian noise \mathbf{N}). More details will be provided after the 3 following examples.

2.1.1 Observations deviating from the standard linear model

By ways of illustrations, we will focus on the three applications that have been presented in the introduction 1.1.1: LC-MS data analysis in bioengineering [114], estimation of the foreground emissions from the ESA-Planck data [86], and unmixing of hyperspectral images for terrestrial remote sensing [10]. The presence of deviations from the standard linear mixture model is common in these three different applications.

Composition of a liquid from LC-MS data in bioengineering. In [114], the authors analyze LC-MS data with BSS so as to estimate the composition of a liquid.

These data are noisy, but the underlying noise model is not completely understood. A part




Figure 2.1: Zoom on some non-linearities present in LC-MS data, intensity in log-scale. Figure from [115].

of the deviations is dense, with moderate amplitude, and is generally modeled with an additive Gaussian noise. However, it would be better described by a multiplicative term as pointed out in [114]. Large deviations are also encountered and certainly correspond to electromagnetic noise, fig.2.1.

In the case of fig.2.1, the deviations are easily detectable and can be manually discarded. However, the hand-operated removing of the deviations has several drawbacks: it is time consuming,

and may be not precise (deviations not detectable without further analysis, or hidden contribution of the signals of interest).

CMB and foreground emission separation from ESA-Planck data in cosmology. The BSS methods aiming to estimate the CMB from the ESA-Planck data assume that the observations are well described by the standard linear model.

In practice, the observations can significantly deviate from this simple model:




Figure 2.2: Left: Sky map of some compact source emissions visible in the microwave range. Right: combined mask for the compact sources and the spectra variabilities.

- Some extra-components are not modeled, such as point source emissions fig.2.2¹. Each of these compact emissions has a unique emission law, and as so, should be modeled as an individual source. This is in practice impossible since they are too numerous, and their number is not known. Nonetheless, the position of the brightest ones is known, and the corresponding observations are masked prior to the unmixing.
- The underlying physical process is too complex to be modeled with the linear model. Indeed, the emission laws of the foreground components depend on physical parameters such as the temperature or gas densities, which are strongly varying across the sky. Consequently, we face consequential spectral variabilities of the galactic components, which are unfortunately dominating the CMB contribution. As well in practice, a part of the galactic center is masked prior to the unmixing fig.2.2.²

One may seek for accurate *full sky* estimates of the foregrounds and CMB, and not only for the statistical properties of the CMB. In that case, we cannot longer mask the galactic center. The deviations from the linear model (spectral variability and point sources) should be carefully taken into account.

¹Map from <http://sci.esa.int/jump.cfm?oid=56131>

²Map from https://wiki.cosmos.esa.int/planckpla/index.php/CMB_and_astrophysical_component_maps

Unmixing from hyperspectral images in terrestrial remote sensing.

As stated in the introduction 1.1.1, the observations are commonly modeled with the standard linear model $\mathbf{X} = \mathbf{AS} + \mathbf{N}$. However, this linear approximation can be too simplistic to represent accurately some terrestrial hyperspectral data [138], [45], [119], [2]:

- Variations of illumination: This is especially problematic if the abundances are assumed to 'sum-to-one' ($\sum_{i=1}^n \mathbf{S}_i^k = 1, \forall k = 1, \dots, t$), [80]. For instance, let us consider two pure pixels of an element, with two different illuminations. The energies of the pixels are different but the amplitudes of the two samples should both set to 1 (pure pixel): the corresponding endmembers should thus be different so as to take into account the variations of energy.
- Intrinsic variations: It is likely that a material varies slightly in the observed area. For illustrative purpose, some endmembers corresponding to pure pixels of vegetation from a MOFFETT scene are displayed in fig.2.3: one cannot expect different plant species to have exactly a same spectral signature.
- Intimate mixing: At a microscopic level, the components interact together. For instance, if observing a sea-coast, assuming that the material (wet sand for example) can be described only by the spectra of the (dry) sand and water is maybe too simple.
- Multiple paths reflections: The observed reflection may not come straightforwardly from a component, but rather reflects off several materials.




Figure 2.3: Examples of spectral variabilities of a component.

The presence of strong deviations is detrimental for some state-of-the-art unmixing methods [45], or for high precision classification. Besides in some applications, the spectral variations are the point of interest (ex: for remote monitoring in agriculture [90]), and it would thus be of

interest to estimate them precisely.

2.1.2 Towards a more realistic data model

As illustrated in the previous examples, the standard linear model $\mathbf{X} = \mathbf{AS} + \mathbf{N}$ can be too simplistic to fairly describe the observed physical process.

Explicit modeling of the outliers. In the following, the deviations will be explicitly taken into account in the data model by adding an extra linear term, $\mathbf{O} \in \mathbb{R}^{m \times t}$, *the outliers*, so that:

$$\mathbf{X} = \mathbf{AS} + \mathbf{O} + \mathbf{N}, \quad (2.1)$$

where \mathbf{AS} corresponds to the sought-after rank- n source contribution, \mathbf{N} to a dense and non-structured noise (assumed to be Gaussian i.i.d), and \mathbf{O} , the deviations/outliers.

Defining the different contributions. In principle, this data model could represent every kind of observations (for instance, by taking $\mathbf{O} = \mathbf{X}$). That is why, defining precisely what *are* the outliers is necessary to identify every contribution.

Describing the outliers on their own is difficult, and generally data dependent. For instance in 2.1.1, the deviations in the LC-MS data correspond to the diagonally-shaped (non-linear) contribution, with variable masses. If the deviations are no longer drifted, further analysis is needed to detect the deviations: in this case, the outliers can only be defined and detected by contrast with the source contribution.

- The source contribution \mathbf{AS} is linearly structured and of rank- n . In particular, the expansion coefficients of \mathbf{S} in $\Phi_{\mathbf{S}}$ (dictionary in which the sources are sparsely represented) are assumed to respect the MDP presented in 1.2.2 and thus to be jointly sparse fig.2.5b. Consequently, the source contribution $\mathbf{AS}\Phi_{\mathbf{S}}^T$ is clustered along the directions given by the columns of \mathbf{A} fig.2.5a.
- The outliers correspond to a structured corruption in contrast with the Gaussian noise. In particular, we assume that the outliers are spatially coherent: their morphology is unchanged from one observation to another. Besides, we consider that the morphology of the outliers is known a priori: there is a dictionary $\Phi_{\mathbf{O}}$ (possibly $\Phi_{\mathbf{O}} = \Phi_{\mathbf{S}}$) in which, the expansion coefficients $\mathbf{O}\Phi_{\mathbf{O}}^T$ are sparse. For instance in 2.1.1, the deviating masses in the LC-MS data are sparsely represented in the direct domain, similarly to the spurious point emissions in the Planck data. Contrary to the expansion coefficients of the sources in $\Phi_{\mathbf{S}}$, the samples $\mathbf{O}\Phi_{\mathbf{O}}^T$ do not need to be jointly sparse.

Besides, in order to separate the outliers from the source contribution, we will assume that at least, one of the following properties holds true:

- The outliers and the sources do not share the same morphology. In the spirit of MCA (see 1.8), the outliers and the sources have clearly different geometrical contents. This difference of spatial behaviors will be exploited in Chapter 5.
- The outlier contribution does not cluster in Φ_S .

In contrast with the source contribution which is clustered along the directions given by \mathbf{A} in Φ_S , we will assume that the outliers are well spread out in all the directions in Φ_S . This ensures that if Φ_O and Φ_S are too similar, we can still distinguish \mathbf{AS} from \mathbf{O} using the fact that the outliers are “aberrant”: fewer entries share a same behavior/spectrum. Ideally, the non-zero samples $\{(\mathbf{O}\Phi_S^T)^k\}_{k:(\mathbf{O}\Phi_S^T)^k \neq \mathbf{0}_{\mathbb{R}^m}}$ projected onto the unit sphere of \mathbb{R}^m are uniformly distributed (which can be achieved if the entries of $\{(\mathbf{O}\Phi_S^T)^k\}_{k:(\mathbf{O}\Phi_S^T)^k \neq \mathbf{0}_{\mathbb{R}^m}}$ are generated with a same normal distribution), such as in fig.2.4a.

This condition may not be met in practice: for instance, one point source emission in the Planck data will certainly active several atoms in Φ_S , and these active samples would share a same spectral signature. Nevertheless, these active samples sharing a same spectrum are less numerous than the active samples corresponding to a sought-after source, such as the synchrotron emission for instance fig.1.8³.

This difference of spectral structure will be exploited in the Chapter 4.

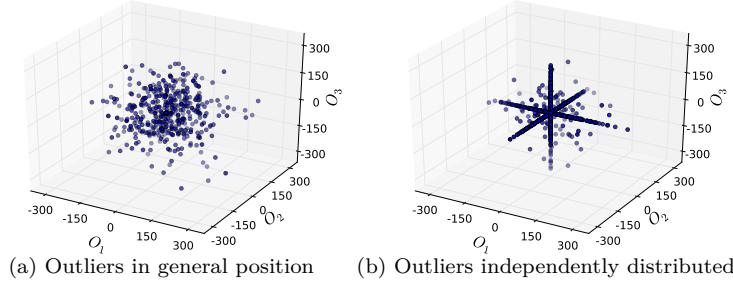


Figure 2.4: In (a), the outliers are column sparse whereas they are row and column sparse in (b), the amplitudes of the active entries are identical for (a) and (b).

One can notice that with our proposed characterization of the outliers, the presence of numerous corrupted entries, column and row sparse in Φ_S (outliers uniformly drawn at random in Φ_S) may not be defined as outliers. Indeed, such outliers are clustered along the canonical axes, as in fig.2.4b. We will rather consider that, ideally, the outliers, if sharing the same morphology as the sources, are column sparse in Φ_S , with i.i.d. entries. More details will be provided in Chapter 4.

³The following criterion measures an ‘admissible’ level of clustering of the outliers in Φ_S . Let d_s be the number of atoms of Φ_S , such that $\mathbf{O}\Phi_S^T \in \mathbb{R}^{m \times d_s}$. Let $\mathbf{B} \in \mathbb{R}^{m \times d_o}$ be a spectral dictionary whose atoms are normalized for the ℓ_2 norm, such that $\mathbf{O}\Phi_S^T = \mathbf{B}\mathbf{T}$, the rows of $\mathbf{T} \in \mathbb{R}^{d_o \times d_s}$ respect the MDP, and d_o be the minimal dimension permitting such decomposition. We will assume that $\max_i \frac{\|\mathbf{T}_i\|_1}{\|\mathbf{T}_i\|_2} < \min_j \frac{\|(\mathbf{S}\Phi_S^T)_j\|_1}{\|(\mathbf{S}\Phi_S^T)_j\|_2}$.

2.1.3 Influence of the deviations

There are applications where the deviations from the linear model are the objects of interest e.g. target and anomaly detection [70]. In this case, using a method able to estimate the deviations precisely, by taken them into account explicitly, makes sense. However, it may be not clear yet, why robust methods should also be employed if we are only seeking for the mixing matrix or the sources: *the presence of important deviations deteriorates the estimation of \mathbf{A} and \mathbf{S} , whatever the prior on the sources is used [62], [79], [33]*. Given that in this thesis we focus on sparse BSS, we propose to study the influence of the outliers in this particular context with toy examples.

For illustrative purpose, let us consider that 3 sources are mixed into 3 observations. For the sake of clarity, the sources are exactly sparse in the domain of observations. That is, they are jointly sparse (sparse and respect the MDP 1.2.2): one entry is significant for at most one source such as in fig.2.5b. Besides, the data samples are clustered along the directions given by the columns of \mathbf{A} fig.2.5a. The algorithm GMCA, presented in 1.3.2 and introduced in [14] would be particularly well suited to retrieve the sources and mixing matrix⁴. We recall that it performs the separation by minimizing the following cost function:

$$\underset{\mathbf{A}, \mathbf{S}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_2^2 + \|\Lambda \odot \mathbf{S}\|_1,$$

where Λ is a matrix of size $n \times t$ containing the regularization parameters.

Among the possible tuples $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$, GMCA looks for one satisfying $\mathbf{X} \approx \tilde{\mathbf{A}}\tilde{\mathbf{S}}$, with the sparsest sources (enclosed in the ℓ_1 ball with the smallest radius), Section 1.3.1, [14]. In this ideal setting, the results obtained by GMCA are very convincing (cannot be visually distinguished from the ground truth).




Figure 2.5: Three exactly sparse sources in the direct domain are mixed into 3 noiseless observations. Left: scatter plot of \mathbf{X}_1 versus \mathbf{X}_2 . Right: scatter plot of \mathbf{S}_1 versus \mathbf{S}_2 . The red arrows symbolize the directions of \mathbf{A} .

We consider then that the observations are corrupted by outliers. We can broadly distinguish

⁴The associated algorithm is presented in Section 3.3.5

two kinds of outliers (the energy of \mathbf{O} is the same in the two examples):

- i) The outliers are also sparse in the same domain as the sources (the direct domain here). More precisely, only few samples of \mathbf{O} are active: in fig. 2.6b, 10% of the entries are active. In contrast with the source contribution, the outliers are in general position and do not cluster in specific directions. This set-up corresponds to the variable masses in the LC-MS data presented in 2.1.1, and will be studied in the Chapter 4.
- ii) The outliers are dense/non-sparse in the domain in which the sources are sparsely represented: all the entries are corrupted, by an additive noise (not Gaussian i.i.d. such as \mathbf{N}), fig.2.6a. Similarly, the outliers are in general position, they do not cluster in any specific direction. This setting could describe the contribution of spectral variations encountered in the Planck and hyperspectral data in Section 2.1.1 and will be studied in the chapters 5 and 6.

In the presence of moderate outliers, it is indeed possible to retrieve the mixing matrix with GMCA, [67]. However, whenever the outliers are significant, they hamper the separation. The unmixing is hindered differently depending on the type of the outliers: sparse or dense.




Figure 2.6: Left: scatter plot of \mathbf{X}_1 versus \mathbf{X}_2 when the observations are corrupted by dense outliers. Right: scatter plot of \mathbf{X}_1 versus \mathbf{X}_2 when the observations are corrupted by sparse outliers, with a same energy. The red arrows symbolize the directions of \mathbf{A} .

Dense outliers. When the outliers are dense, they act as a large Gaussian noise. Indeed, the projected sources $\mathbf{A}^\dagger \mathbf{X} = \mathbf{S} + \mathbf{A}^\dagger \mathbf{O}$ equal the sources plus an additional dense and relatively small term $\mathbf{A}^\dagger \mathbf{O}$, which affects the sources similarly to a Gaussian noise 2.7a. The largest samples still belong to the source contribution (plus the noise) fig.2.6a, and the projected sources are still jointly sparse 2.7a, even if they are less precisely clustered along the directions given by the columns of \mathbf{A} 2.6a, 2.7a.

Consequently, it is still possible to recover a broad estimate of \mathbf{A} with GMCA: we can see that the directions are fairly but not exactly recovered by GMCA fig.2.7b. Besides, if we try to further denoise the sources, then, only few coefficients are kept fig.2.7c since the noise level is

quite similar to the amplitude of the sources. One can notice that the ℓ_1 norm has introduced a very significant bias in the recovered samples fig.2.7c. This effect can be reduced by using a reweighting ℓ_1 [23] or the ℓ_0 pseudo-norm: similarly, few entries would be kept, but they would not have a biased amplitude.




Figure 2.7: Left: scatter plot of $(\mathbf{A}^\dagger \mathbf{X})_1$ versus $(\mathbf{A}^\dagger \mathbf{X})_2$ with dense outliers, where the red arrows symbolize the directions of the true \mathbf{A} . Middle: scatter plot of $(\tilde{\mathbf{A}}^\dagger \mathbf{X})_1$ versus $(\tilde{\mathbf{A}}^\dagger \mathbf{X})_2$ with dense outliers. Right: some of the samples of the recovered source $\tilde{\mathbf{S}}_1$. The tuple $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$ is estimated with GMCA.

In the presence of dense outliers, a fair estimation of the variables can be obtained with standard sparse BSS algorithms. However, if the energy of the outliers is very significant, such as in fig.2.7, the separation between outliers and source contribution cannot be very precise.

Sparse outliers. Large and sparse outliers are more damaging for sparse BSS methods. The basic assumptions needed for GMCA are no longer valid in the presence of sparse and large outliers (the consequences of the MDP are no longer valid, Section 1.3). First, the largest samples can be related to the outliers and not to the sources fig.2.6b. Besides, the projected sources $\mathbf{A}^\dagger \mathbf{X} = \mathbf{S} + \mathbf{A}^\dagger \mathbf{O}$ are not jointly sparse fig.2.8a.

This last point is particularly damaging for GMCA. Indeed, this method looks for the tuple $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$ so that $\tilde{\mathbf{S}}$ is enclosed in the ℓ_1 ball with the smallest radius. In the presence of outliers, there may be another tuple $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$ so that the resulting sources $\tilde{\mathbf{S}}$ are enclosed in a ℓ_1 ball with a smaller radius. That is the case in fig.2.8: the estimated sources $\tilde{\mathbf{S}} \approx \tilde{\mathbf{A}}^\dagger \mathbf{A} \mathbf{S} + \tilde{\mathbf{A}}^\dagger \mathbf{O}$ fig.2.8b are sparser (ℓ_1 ball in red) than the projected sources $\mathbf{S} + \mathbf{A}^\dagger \mathbf{O}$ (ℓ_1 ball in green) fig.2.8a. That is why the mixing matrix is not recovered. In addition to the inaccurate unmixing, the sources (whose corresponding columns are recovered) are corrupted by the projected outliers $\tilde{\mathbf{A}}^\dagger \mathbf{O}$. Given that this contribution is also sparse, it cannot be easily distinguished from the sources entries and a fortiori discarded with the thresholding, fig.2.8c.




Figure 2.8: Left: scatter plot of $(\mathbf{A}^\dagger \mathbf{X})_1$ versus $(\mathbf{A}^\dagger \mathbf{X})_2$ with sparse outliers. Middle: scatter plot of $(\tilde{\mathbf{A}}^\dagger \mathbf{X})_1$ versus $(\tilde{\mathbf{A}}^\dagger \mathbf{X})_2$ with sparse outliers. The green dashed lines symbolize the ℓ_1 ball in which is enclosed $(\mathbf{A}^\dagger \mathbf{X})$ and the red dashed lines the ℓ_1 ball enclosing $(\tilde{\mathbf{A}}^\dagger \mathbf{X})$. Right: some of the samples of the recovered source $\tilde{\mathbf{S}}_1$. The tuple $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$ is estimated with GMCA.

Standard sparse BSS are clearly hampered by sparse and large outliers. They are misled by large corrupted entries and fail to retrieve the mixing matrix. Besides, given the outlier contribution is also sparse, an exact separation between \mathbf{O} and \mathbf{AS} cannot be achieved using only standard methods such as GMCA.

2.2 Robust BSS in the literature

The presence of dense outliers can be fairly handled by the methods which are robust to the presence of Gaussian noise. However, the standard strategies lead to erroneous results in the presence of large and sparse outliers. This mandates the development of *robust* BSS methods, which should tackle the following tasks:

- i) Separation of the source contribution \mathbf{AS} from the outliers \mathbf{O} so as to return denoised estimated sources.
- ii) Unmixing of the sources, *i.e.* estimating precisely the mixing matrix \mathbf{A} .

Only few robust methods have been proposed in the literature for dealing with *large and sparse outliers*. They can be classified into three different groups according to their strategies: replacement of the sensitive metrics in the cost-functions of optimization-based methods (*i.e.* only task ii)), removal of the outliers prior to the unmixing (*i.e.* task i) followed by task ii)), and joint estimation of \mathbf{O} , \mathbf{S} and \mathbf{A} (*i.e.* tasks i) and ii) simultaneously).

2.2.1 Removal of the outliers in the low-rank regime

We first focus on the outlier removal strategies permitting a precise separation between $\mathbf{L} = \mathbf{A}\mathbf{S}$ and \mathbf{O} such that:

$$\mathbf{X} = \mathbf{L} + \mathbf{O}.$$

Separating \mathbf{L} from \mathbf{O} is a challenging problem which requires additional priors, such as assuming that the source contribution \mathbf{L} has low-rank $m \gg n$, and that the outliers are sparse.

Outlier detection. These two properties are very effective to detect the outliers Fig.2.9: the corrupted sample \mathbf{X}^k cannot be detected by visual inspection in the determined case ($m = n$) fig.2.9a, while the column-span of \mathbf{L} and the corrupted sample \mathbf{X}^k are easily identified in the over-determined setting ($m > n$)fig.2.9a. Only the samples whose outlier contribution does not lie in the column-span of \mathbf{A} can be detected: this is likely if \mathbf{L} has low-rank and the outliers are sparse.

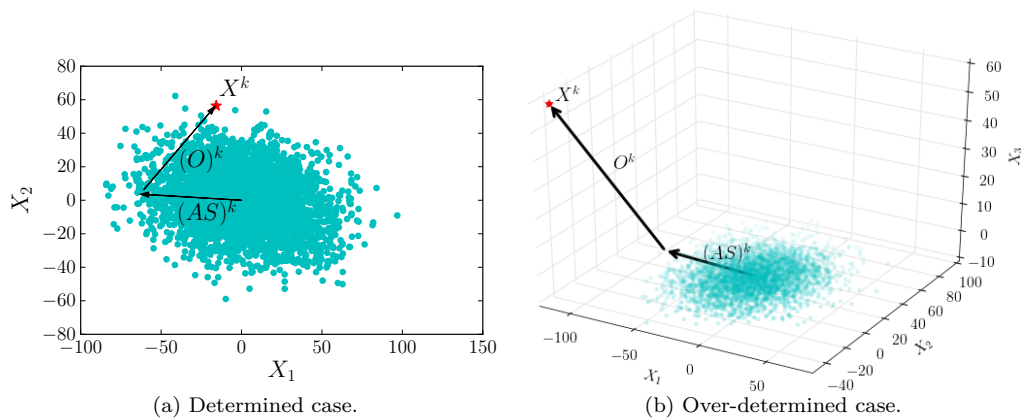


Figure 2.9: On the left, 2 sources are mixed into 2 corrupted observations (blue points). On the right, 2 sources are mixed into 3 observations. For both, the red star symbolizes the corrupted sample, at the k th column and the arrows symbolize the two contributions to this sample \mathbf{O}^k and $(\mathbf{AS})^k$.

Outlier identification. Even if the corrupted samples are detected, identifying clearly the two contributions \mathbf{L} and \mathbf{O} requires further assumptions since for every corrupted column k , there is an infinite number of tuples $(\mathbf{L}^k, \mathbf{O}^k)$ such that $\mathbf{X}^k = \mathbf{L}^k + \mathbf{O}^k$ and \mathbf{L}^k lies in the estimated column-span of \mathbf{L} .

Exact separation between \mathbf{L} and \mathbf{O} . An analogous framework, leading to the exact identifiability of the two components, has been proposed in the two breakthrough papers [23]

and [29]. In order to be fully identifiable, the components can meet the following properties (sufficient but not necessary conditions):

- \mathbf{L} should not be sparse. This condition has been characterized differently in [23] and [29]: by ensuring that the tangent space of \mathbf{L} with respect to the matrices of rank equal or less than n have well spread out entries (and so is \mathbf{L}) in [29] or that the right and left singular vectors of \mathbf{L} are well spread out for [23].

This last characterization is designated as *incoherence condition with parameter μ* [21] and states that for $\mathbf{L} = \mathbf{U}\Sigma\mathbf{V}^* = \sum_{i=1}^n \sigma_i \mathbf{U}^i \mathbf{V}^{i*}$, where $\sigma_1, \dots, \sigma_n$ denote the positive singular values of \mathbf{L} :

Definition 2.2.1 (Incoherence condition with parameter μ)

$$\max_i \|\mathbf{U}^* e_i\|_2^2 \leq \frac{\mu n}{m}, \quad \max_i \|\mathbf{V}^* e_i\|_2^2 \leq \frac{\mu n}{t}, \quad \|\mathbf{U}\mathbf{V}^*\|_\infty \leq \sqrt{\frac{\mu n}{mt}}$$

- The outliers should not be concentrated in a row or a column, to ensure that the rank and column-span of \mathbf{L} is still recoverable (the outliers plus a column or row of \mathbf{L} should not be zero). For this purpose, the authors of [23] assume that outliers are column and row sparse, and that the support of outliers is drawn uniformly at random. In [29], the authors consider that the matrices belonging to the tangent space of \mathbf{O} , with respect to matrices having at most the same number of active entries as \mathbf{O} , have a spectrum (set of its eigenvalues) which is well spread out (*i.e.* the energy of the outliers is well spread out in all the directions).

If these requirements are satisfied, it is then possible to recover \mathbf{L} and \mathbf{O} by minimizing the following problem [29], [23], named *Principal Component Pursuit* (PCP) in [23]:

$$\underset{\mathbf{L}, \mathbf{O}: \mathbf{X}=\mathbf{L}+\mathbf{O}}{\text{minimize}} \|\mathbf{L}\|_* + \lambda \|\mathbf{O}\|_1, \quad (2.2)$$

where $\|\mathbf{L}\|_*$ denotes the nuclear norm, the convex relaxation of $\text{rank}(\mathbf{L})$, which enforces the low-rankness of \mathbf{L} , and $\|\mathbf{O}\|_1$ promotes sparse outliers.

In brief, the algorithm exploits the incoherence of structure between the two terms to distinguish them: the spectrum of \mathbf{L} is sparse (low-rankness) whereas the one of \mathbf{O} is dense; and the entries of \mathbf{O} are sparse while the ones of \mathbf{L} are dense.

In [29, Theorem 2], the authors propose a deterministic recovery guarantee, depending on how broadly distributed are the entries of \mathbf{L} and how sparse is \mathbf{O} . The recovery guarantee in [23, Theorem 1.1] is probabilistic (due to the assumption on the outliers) and assume that m and t are large enough (with respect to n and the number of active entries of \mathbf{O}):

Theorem 2.2.1 (Theorem 1.1, [23]) *Suppose \mathbf{L} obeys the condition 2.2.1 with parameter μ , and that the support of \mathbf{O} is uniformly distributed among all sets of cardinality $\|\mathbf{O}\|_0$. Then there*

is a numerical constant c such that with probability at least $1 - ct^{-10}$, the PCP decomposition 2.2 with $\lambda = \frac{1}{\sqrt{t}}$ is exact, provided that:

$$n \leq \frac{\rho_r m}{\mu(\log(t))^2} \text{ and } \|\mathbf{O}\|_0 \leq \rho_o t m,$$

where ρ_o and ρ_r are two positive numerical constants.

In particular, it appears from this previous theorem that there is no parameter to tune: taking $\lambda = \frac{1}{\sqrt{t}}$ provides a correct separation if the previous requirements are met.

Numerous works have followed these two papers such as [143] (presence of Gaussian noise), [142] (faster implementation and presence of Gaussian noise), or [105] (utilization of a truncated nuclear norm) - a review of the different methods can be found in [18]. Beyond the methodological developments, PCP based techniques have been very popular in applicative domains including background estimation of video [18] and hyperspectral image denoising [139], [91], [92]. Given that hyperspectral data cubes are made of hundreds of observations of few components, the component of interest \mathbf{L} is indeed low-rank ($m \gg n$). In [139], the authors use a PCP-based technique to remove the outliers (impulse noise, stripping noise) which are assumed to be row and column sparse (*i.e.* only few pixels are corrupted over few bands). In contrast, the noise is assumed to be column-sparse in [92] and the ℓ_1 norm in PCP is replaced by the $\ell_{2,1}$ to promote column-sparse outliers. Nonetheless, dealing with column-sparse outliers is much more challenging than with row and column sparse outliers.

Relaxing the sparsity pattern of outliers: column-sparse outliers. In many applications, the outliers are more likely to be column-sparse than uniformly drawn at random (column and row sparse). For example in hyper/multichannel data, unexpected physical events, such as the presence of point source emissions in cosmology [126], Section 2.1.1, are probably visible at all the frequencies, and so are column-sparse.

Identifying precisely the two contributions in this case is not possible. Indeed, for every corrupted sample k , one should recover the initial decomposition $\mathbf{O}^k + \mathbf{L}^k$, while having no prior on the vectors, only global priors on \mathbf{L} (low-rank), and \mathbf{O} (column-sparse). Hence, even if restricting the vector \mathbf{L}^k to lie in the column-span of \mathbf{L} , the identification problem is still ill-posed.

In contrast, if the outliers are assumed to be sparse, the corresponding column \mathbf{O}^k is likely to be 1 sparse: \mathbf{O}^k belongs to one of the canonical axes, which greatly reduces the possibility for \mathbf{O}^k fig.2.10a. On the other hand, if the outliers are column-sparse, the problem cannot be solved exactly. In the fig.2.10b for instance, since the corrupted sample \mathbf{X}^k belongs to the z -axis, the PCP algorithm would return $\tilde{\mathbf{O}}^k = \mathbf{X}^k$ and $\mathbf{L}^k = 0$ (which minimizes jointly the nuclear and ℓ_1 norms) but not the solution of interest.

In [135], the authors study the recovery of \mathbf{L} and \mathbf{O} with column sparse outliers with the




Figure 2.10: 2 sources are mixed into 3 corrupted observations (blue points). On the left, the outlier sample \mathbf{O}^k is 1 sparse, and it is column-sparse (all the entries of \mathbf{O}^k are active) on the right. The blue arrows symbolize the directions of \mathbf{A} , the red arrow the contribution of \mathbf{O}^k to \mathbf{X}^k , and in white, the one of \mathbf{L}^k .

Outlier Pursuit (OP) algorithm:

$$\underset{\mathbf{L}, \mathbf{O}: \mathbf{X} = \mathbf{L} + \mathbf{O}}{\text{minimize}} \|\mathbf{L}\|_* + \lambda \|\mathbf{O}\|_{2,1},$$

where the $\ell_{2,1}$ is used to favor column-sparse outliers.

In that case, one can only recover exactly the column-span of \mathbf{L} and the support of \mathbf{O} given that: \mathbf{L} is not also column sparse (but it can be row-sparse, in contrast with [23]), and the number of corrupted samples and the rank n of \mathbf{L} are small enough with respect to the number of samples t [135].

Unmixing from denoised observations. Even though the low-rank regime is favorable to the detection of corrupted samples and the estimation column-span of \mathbf{A} , separating the outliers from \mathbf{L} is ill-posed and cannot be guaranteed with mild hypotheses (only if L has low-rank, and the outliers are column and row sparse).

If the separation between \mathbf{O} and \mathbf{L} is not perfect, the following unmixing on $\mathbf{X} - \tilde{\mathbf{O}}$ would be inaccurate if not erroneous. This applies more generally to the 'two-steps' methods, combining first a denoising method and then an unmixing strategy.

The first obvious reason is that an outlier residual would be still present and would hinder the following unmixing and the sources estimation. The second point is that, the denoising method may have not preserve the key properties of \mathbf{L} , for example:

- If a monochannel denoising method is used (the outliers are removed channel by channel, independently), the structural property of \mathbf{AS} (linearity from one observation to another) may not be preserved - as illustrated in the numerical experiments in 5.3.2.

- The parameter λ in PCP-based methods (PCP or Outlier-Pursuit) is not easy to tune. In [23], [29], [135], the authors propose a value for λ . However, even if this value achieves perfect asymptotic results, it may be not the best choice in practice: another value of λ can lead to better unmixing results.

Removing the outliers prior to the unmixing is not an easy task. The exact separation between \mathbf{L} and \mathbf{O} requires sparse outliers and low-rankness, two conditions which are not valid in some applications, for example for the removing of the point source emissions from the ESA-Planck data. In the following section, we will present a group of methods which do not rely on low-rankness.

2.2.2 Robust metrics

Another strategy for withstanding large deviations consists in replacing the most sensitive metrics involved in the cost-functions of the deterministic methods such as the Frobenius norm or the Kullback-Leiber (KL) divergence. Namely, since the outliers are assumed to be large, and so correspond to large values of the residue $\mathbf{X} - \mathbf{A}\mathbf{S}$, the robust metrics should not penalize much large entries. They should rather penalize small deviations, which are more likely to correspond to a mis-estimation of the variables, for the inliers points. This approach has been proposed for robust ICA and NMF.

Robust ICA. In the outlier-free setting, based on the assumptions that the sources are independent and that $m = n$, the ICA based methods look for an unmixing matrix \mathbf{B} such that the estimated sources $\tilde{\mathbf{S}} = \mathbf{B}\mathbf{X}$ are statistically independent [40], Section 1.1.3. One way to measure the independence of the estimated sources is to use the mutual information of the sources, defined as the Kullback-Leibler (KL) divergence between the product of their marginal distributions $\prod_{i=1}^n p_{\mathbf{S}}(\tilde{\mathbf{S}}_i)$ and their joint distribution $p_{\mathbf{S}}(\tilde{\mathbf{S}})$ [39]:

$$\mathbb{D}_{KL}(p_{\mathbf{S}}(\tilde{\mathbf{S}}) \parallel \prod_{i=1}^n p_{\mathbf{S}}(\tilde{\mathbf{S}}_i)) = \int p_{\mathbf{S}}(\tilde{\mathbf{S}}) \log \left(\frac{p_{\mathbf{S}}(\tilde{\mathbf{S}})}{\prod_{i=1}^n p_{\mathbf{S}}(\tilde{\mathbf{S}}_i)} \right) d\tilde{\mathbf{S}}.$$

The value of the KL-divergence is non-negative and equal to zero if and only if the estimated sources are independent. Unfortunately, it is not robust to the presence of outliers [62]. To overcome this problem, the authors of [100] propose to replace the KL-divergence by the β -divergence, for which the influence of the outliers is bounded [100]. Similarly, the β -divergence between the product of the marginal densities of the estimated sources and their joint density is non-negative and equal to zero if and only if the estimated sources are independent (for $1 > \beta > 0$,

and for $\beta = 0$ we refer to the KL divergence):

$$\mathbb{D}_\beta(p_{\mathbf{S}}(\tilde{\mathbf{S}}) \parallel \prod_{i=1}^n p_{\mathbf{S}}(\tilde{\mathbf{S}}_i)) = \frac{1}{\beta} \int p_{\mathbf{S}}(\tilde{\mathbf{S}}) \left(p_{\mathbf{S}}^\beta(\tilde{\mathbf{S}}) - \prod_{i=1}^n p_{\mathbf{S}}^\beta(\tilde{\mathbf{S}}_i) \right) d\tilde{\mathbf{S}} - \frac{1}{\beta+1} \int \left(p_{\mathbf{S}}^{\beta+1}(\tilde{\mathbf{S}}) - \prod_{i=1}^n p_{\mathbf{S}}^{\beta+1}(\tilde{\mathbf{S}}_i) \right) d\tilde{\mathbf{S}}.$$

As the value of β increases, the divergence becomes less sensitive to outliers, but also less effective since it is on the overall less influenced by the deviations to the observations (the best unmixing performances, without outliers, are obtained with very small values of β , for a cost function close to the K-L divergence).

The strategy proposed in [100] amounts to maximize the quasi log-likelihood of the β -divergence between the empirical marginal densities of the estimated sources and their empirical joint density. It does not need any whitening of the data, but only that $m = n$. Nevertheless, the efficiency of this approach highly depends on the choice of the parameter β , which is generally problematic in practice.

Divergences have also been employed in the framework of NMF.

Robust NMF. This strategy has been quite popular for NMF. Indeed, it is possible to change the Frobenius norm, which is commonly used as a data fidelity term, for a more robust metric while keeping the fast multiplicative updates.

Divergences The utilization of divergences in the NMF framework has a long history [87]. The data fidelity term is of the form $\mathcal{D}(\mathbf{X} \parallel \mathbf{AS}) = \sum_{j=1..m, i=1..n} d(\mathbf{X}_{i,j} \parallel (\mathbf{AS})_{i,j})$, where d denotes the entry-wise divergence. The most common divergences can be seen as a β -divergence, which is defined as the following for $\beta \in \mathbb{R}$ and $x, y \in \mathbb{R}_+$, [51], [59]:

$$d(x \parallel y) = \begin{cases} \frac{1}{\beta(\beta-1)} (x^\beta + (\beta-1)y^\beta - \beta xy^{\beta-1}), \forall \beta \in \mathbb{R} \setminus \{0, 1\} \\ x \log\left(\frac{x}{y}\right) + (y-x), \beta = 1 \text{ KL} \\ \frac{x}{y} - \log\left(\frac{x}{y}\right) - 1, \beta = 0 \text{ Itakuro-Saito} \end{cases} \quad (2.3)$$

For $\beta = 2$, the β divergence corresponds to the Euclidean distance (well suited for Gaussian noise) [88], the value $\beta = 1$ to the KL divergence (maximum likelihood for Poisson noise) [87] and $\beta = 0$ to the Itakuro-Saito divergence (maximum likelihood for multiplicative gamma noise) [56]. For $0 \leq \beta < 2$, these divergences are non-symmetric ($d(x \parallel x + \Delta) < d(x \parallel x - \Delta)$ for $\Delta > 0$), and the influence of the large entries are lessen as β tends to 0: for $\Delta > 0$: $d(\Delta \times x \parallel \Delta \times y) = \Delta^\beta d(x \parallel y)$ [56]. Hence, it is possible to lessen the influence of large errors, and so to be robust to outliers. In [35], the authors further generalized the framework to the $\alpha\beta$ divergences and studied their robustness with respect to outliers and parameters variations.

Whereas the theoretical properties of these divergences and their robustness motivate their utilization, setting correctly the parameters of the divergence is not straightforward, especially in the presence of an additional Gaussian noise (setting β to a small value would make the method robust to large outliers but more sensitive to the Gaussian noise).

Numerous works, using robust distances instead of parametric divergences, have also been proposed for NMF as presented below.

Robust distances In [81], the authors used the $\ell_{2,1}$ norm for the data fidelity term with the non-negativity of \mathbf{A} and \mathbf{S} in order to cope with column-Laplacian outliers. They propose an implementation procedure based on multiplicative updates, as well as the extension for the ℓ_1 norm which will be more well-suited to deal with an entry-wise Laplacian noise (entries are Laplacian i.i.d).

The use of the ℓ_1 norm was already proposed by Ke and Kanade in [79] (the components are not assumed to be non-negative). However, they did not present an efficient implementation for this problem (they proposed to smooth the ℓ_1 with the Huber norm [76]). Besides, they proposed a weighting procedure for the ℓ_1 , which can be useful in the presence of missing data (the weighting is a binary mask) or to further penalize the entries with a large residue (*i.e.* which are likely to be corrupted).

The authors of [50] also exploit a weighting procedure in a more generic framework. Using half-quadratic minimization, they show that many robust cost-functions ($\ell_{2,1}$, ℓ_1 , but also correntropy, Huber loss and so on) can be easily minimized with weighted NMF.

Loss functions such as the Huber loss [76], the correntropy and Welsch estimator [95] or the hypersurface [30], [74] are worth pointing out. Indeed, using such functions may be more efficient than the ℓ_1 norm because they do not penalize similarly small and large entries. Indeed, as stated above, the small entries of the residue are likely to correspond to errors created by a mis-estimation of the variables, but not to outliers, which are assumed to be very significant. The aforementioned loss-functions penalize proportionally more the small errors than the large ones: quadratic behavior for the small values and then linear for the largest ones for the Huber and hypersurface losses, and quadratic, then linear and then constant behavior for the Welsch loss, see fig.2.11.




Figure 2.11: Examples of error functions. The width of the quadratic regime can be adapted to the data.

We further underline that the quadratic behavior of these error functions is also interesting to deal with the remaining Gaussian noise \mathbf{N} present in the data. More precisely, in [117], the authors have shown that the value of the Huber loss of a residue \mathbf{R} , $\|\mathbf{R}\|_{\text{Huber},\beta}$ with cut-off parameter β corresponds to the minimizer:

$$\|\mathbf{R}\|_{\text{Huber},\beta} = \min_{\mathbf{R}=\mathbf{R}_Q+\mathbf{R}_S} \frac{1}{2} \|\mathbf{R}_Q\|_2^2 + \beta \|\mathbf{R}_S\|_1.$$

Namely, the Gaussian residue \mathbf{R}_Q is penalized by the quadratic function, while the sparse outliers \mathbf{R}_S by the ℓ_1 norm with regularization parameter β . That is, the Huber loss permits a more precise discrimination of the two residual contributions and the results should be the same as if explicitly estimating the outliers with an ℓ_1 norm. However, the explicit estimation of the outliers is worthwhile if additional information on the outliers (such as non-negativity, [58]) is known and can be added to further constrain the problem.

2.2.3 Component Separation

Last, the third strategy consists in estimating jointly \mathbf{A} , \mathbf{S} and \mathbf{O} . This framework is more flexible since several priors can be considered for every component, what leads to a better identifiability of the variables. This has been proposed in a Bayesian framework [4] but more frequently with deterministic strategies. In that case, the tuple $(\mathbf{A}, \mathbf{S}, \mathbf{O})$ is estimated by minimizing a cost function of the form:

$$\underset{\mathbf{A}, \mathbf{S}, \mathbf{O}}{\text{minimize}} \mathcal{D}(\mathbf{X} \|\mathbf{AS} + \mathbf{O}) + \mathcal{J}_1(\mathbf{A}) + \mathcal{J}_2(\mathbf{S}) + \mathcal{J}_3(\mathbf{O}),$$

where \mathcal{D} is the divergence used as data fidelity term, and \mathcal{J}_i , $i = 1..3$ correspond to the penalty functions depending on the prior information on the variables.

This kind of approach has been used in the NMF framework with sparse outliers in [57], [140], and [118]. The authors of [118] use the standard Frobenius norm for the data fidelity term, non-negative priors on \mathbf{A} and \mathbf{S} and the $\ell_{1,2}$ norm for \mathbf{O} which are assumed to be sparse whereas the authors of [140] use the ℓ_1 norm to promote sparsity.

In [57], the authors assume that all the variables are non-negative (including \mathbf{O}), that the entries of each source sample sum to one (see Chapter 6), that the outliers are column-sparse, which is promoted with the $\ell_{2,1}$ norm, and use a β divergence for the data-fidelity term:

$$\text{rNMF: } \underset{\mathbf{A} \geq 0, \mathbf{S} \geq 0, \mathbf{1}_{1 \times n} \mathbf{S} = \mathbf{1}_{1 \times t}, \mathbf{O} \geq 0}{\text{minimize}} \mathcal{D}(\mathbf{X} \|\mathbf{AS}) + \|\mathbf{O}\|_{2,1}, \quad (2.4)$$

where the divergence is defined as in eq.2.3.

It seems in [57] that the use of the β divergence does not bring much in comparison to the Frobenius norm for the studied terrestrial hyperspectral images.

On the overall, these works highly rely on the low-rankness of \mathbf{AS} and the non-negativity of the components to identify the outliers. Indeed, the components \mathbf{A} and \mathbf{S} are only enforced to be non-negative (with indicator functions). On the other hand, the outliers are penalized with a norm (ℓ_1 , $\ell_{2,1}$ or $\ell_{1,2}$) so as to enforce their spatial rareness. This creates an imbalance between the two terms \mathbf{AS} and \mathbf{O} since only the energy of \mathbf{O} is penalized. Hence, the part $[\mathbf{AO}]_+$ will necessarily leak toward the estimated sources since it will be costless, what can be significant if $m \approx n$.

2.3 Summary

2.3.1 Summary of the different approaches in the literature

The different approaches are summarized in the following table Tab.2.1.

Strategies	Framework	Assumption	Advantages	Weaknesses
Separation $\mathbf{AS/O}$	PCP-based methods [23]	Low-rankness and \mathbf{O} sparse	Exact identification of \mathbf{O}	Restrictive hypotheses
	OP-based methods [135]	Low-rankness and column-sparse \mathbf{O}	Estimation of the support of \mathbf{O} and column-span of \mathbf{A}	No exact estimation of \mathbf{O} . Outliers residual hinders the following unmixing.
Robust Metrics	ICA with β divergence [100]	$m = n$	Theoretically robust to outliers. Can handle the determined case.	No estimation of \mathbf{O} . Determined case only. Choice of β crucial but challenging.
	NMF [35], [81]	Non-negativity of \mathbf{A} and \mathbf{S}	Some strategies are easy to implement	No estimation of \mathbf{O} and non-negativity.
Component Separation	NMF [57]	Low-rankness and non-negativity	Well suited for hyperspectral imaging	No identification of \mathbf{O} . Restrictive hypotheses.

Table 2.1: Strategies for robust BSS and outliers estimation.

2.3.2 Identifiability of the outliers

Few works have studied the identifiability of the outliers: namely only in the framework of PCP [23], [29], [136].

In the low-rank regime ($m \gg n$), it has been proven that the exact separation between \mathbf{AS} and \mathbf{O} is possible if the outliers are sparse and their sparsity pattern is uniformly drawn at random. These hypotheses are restrictive: the low-rank regime is not always valid, and the outliers are frequently column-sparse in multi/hyperspectral imaging (an anomaly is visible at several wavelengths).

Even in the low-rank regime, if the outliers are column-sparse, the exact separation between \mathbf{AS} and \mathbf{O} is not possible. However, it is possible to recover the column-span of \mathbf{A} and the support of \mathbf{O} , as long as the active samples of \mathbf{O} do not lie in the span of \mathbf{A} .

We further underline that non-negativity, combined with low-rankness, is not sufficient for the identification of \mathbf{O} . Unless the projection of the outliers onto the column-span of \mathbf{A} is non-positive, what is quite unlikely, the part $[\mathbf{A}^\dagger \mathbf{O}]_+$ would leak towards the estimated sources. This leakage can be controlled, or reduced if the sources are more penalized, such as with the sum-to-one assumption [57] (not only with the indicator function enforcing the non-negativity, which is insensitive to the energy of the estimated sources).

On the overall, the identifiability of the outliers cannot be guaranteed in many applications. However, it is still possible to determine and characterize the potential limits of a method, and the leakages that will be encountered in practice. In the remaining of this thesis, efforts will be undertaken to clarify these points.

2.3.3 Proposed approach: a component separation strategy

There is currently no method able to tackle the BSS problem in the presence of outliers in a wide range of settings, including the determined case.

We propose a new approach, based on the component separation strategy estimating jointly \mathbf{A} , \mathbf{S} and \mathbf{O} . This approach is efficient and flexible since several priors can be added on the components to constrain them. More precisely, we will build upon the sparse modeling of the components to separate the different variables by exploiting efficiently the difference of morphology between the sources and the outliers (Chapter 5), or their spectral (column space) structure (Chapter 4).

Chapter 3

Algorithmic framework for sparse Matrix Factorization

Summary

3.1 Proximal algorithms	40
3.1.1 Proximal operator	40
3.1.2 Proximal algorithm: Forward-Backward Splitting	42
3.1.3 Generalized Forward Backward Splitting	43
3.2 Block Coordinate algorithms for component separation problems	44
3.2.1 Framework	45
3.2.2 Projected least-squares	47
3.2.3 Block Coordinate Descent	49
3.2.4 Prox-linear updates	51
3.2.5 Proximal updates	52
3.2.6 Remarks and summary	54
3.3 Heuristics for sparse BSS	56
3.3.1 Influence of the thresholds: illustration with toy-examples	57
3.3.2 Influence of the parameters Λ : summary	65
3.3.3 Adaptive threshold strategy	66
3.3.4 Benefits and drawbacks of the adaptive threshold strategy	69
3.3.5 Enhancing the results of GMCA	70
3.4 Summary and extensions	73

In many applications, BSS methods aim to retrieve \mathbf{A} , \mathbf{S} , meaningful components from a physical point of view, and not only at estimating a denoised subspace \mathbf{AS} . To this end, efficient BSS algorithms must meet these conflicting features:

- Robustness. The unmixing should not be influenced by the noise.

- Accuracy. The sources and mixing matrix should be precisely estimated.
- Reliability. Matrix factorization methods based on optimization suffer from the non-convexity of the involved cost-functions. The robustness against local minima is crucial for many real-world applications.
- Convergence. The proposed methods should be convergent or empirically stable.
- Easy-to-use. The regularization parameters involved in sparse factorization problems should be automatically set to guarantee an easy handling of the methods, and to avoid time-consuming cross-validations.

After presenting a short background on proximal calculus, we will review some standard optimization strategies for sparse BSS that will be employed in this thesis. Then, we will discuss the influence of the regularization parameters in sparse BSS and an automatic setting strategy, which will also be of paramount importance for the proposed robust BSS methods. The final optimization scheme together with the thresholding strategy will meet the aforementioned desired properties.

3.1 Proximal algorithms

Suppose that we aim to solve an inverse problem with the associated cost function:

$$\operatorname{argmin}_{\mathbf{Y}} f(\mathbf{Y}) + g(\mathbf{Y}),$$

where $f : \mathbb{R}^{d_1 \times d_2} \rightarrow \bar{\mathbb{R}}$, with $d_1, d_2 \in \mathbb{N}_+$, is a convex, proper and differentiable function, and $g : \mathbb{R}^{d_1 \times d_2} \rightarrow \bar{\mathbb{R}}$ be a convex, proper, lower semi-continuous function, not necessarily differentiable. This formulation encompasses many inverse problems, for which f is a differentiable and convex data fidelity term, and the function g a regularization for \mathbf{Y} .

Given that the function g can be not differentiable, the standard optimization methods such as the gradient descent [19] cannot be employed. By contrast, proximal calculus provides an efficient and unified framework to tackle non-smooth convex minimization problems [107]. This is particularly interesting for sparse BSS since sparsity enforcing regularizations such as the ℓ_1 norm are generally not differentiable.

Consequently, we start this chapter by presenting a necessary background on proximal calculus, a tool which will be widely employed in this thesis.

3.1.1 Proximal operator

Definition. Let $g : \mathbb{R}^{d_1 \times d_2} \rightarrow \bar{\mathbb{R}}$ be a convex, proper and lower-semi continuous function. In the following, the functions satisfying these criteria will be called *proximable*. Many functions

commonly employed in inverse problems are proximal e.g., any norm or indicator function of a closed, convex and non-empty set is proximal.

The *proximal operator* of the function g is given by:

$$\text{prox}_g : \mathbb{R}^{d_1 \times d_2} \longrightarrow \mathbb{R}^{d_1 \times d_2} \quad (3.1)$$

$$\mathbf{Y} \mapsto \underset{\mathbf{Z}}{\text{argmin}} \frac{1}{2} \|\mathbf{Z} - \mathbf{Y}\|_2^2 + g(\mathbf{Z}) \quad (3.2)$$

The proximal operator at \mathbf{Y} can be seen as a local minimizer of g , close to \mathbf{Y} . That is, if \mathbf{Y} is a local minimum of g , then the proximal operator of g at \mathbf{Y} is \mathbf{Y} (the fixed point property is key in the analysis of proximal algorithms [107]).

Examples. The proximal operators of the regularizations that will be used in this thesis can be found in App.A. By way of illustrations, we will present some of them:

Proximal operators of indicator functions: we consider that g is the indicator function of the non-negative orthant. Its proximal operator is given by:

$$\text{prox}_g(\mathbf{Y}) = \underset{\mathbf{Z}}{\text{argmin}} \frac{1}{2} \|\mathbf{Z} - \mathbf{Y}\|_2^2 + \chi_{\mathbf{U}:\geq 0}(\mathbf{Y}) \quad (3.3)$$

$$= [\mathbf{Y}]_+ \quad (3.4)$$

The proximal operator is nothing but the projection of \mathbf{Y} onto the non-negative orthant, denoted as $[\mathbf{Y}]_+$. This is also the case if g corresponds to the indicator function of a non-empty convex and closed set. More generally, many proximal operators can be linked to projections.

Proximal operators of the ℓ_1 and other norms: The proximal operator of the $\lambda \ell_1$ norm is the soft-thresholding operator $\mathcal{S}_\lambda(\cdot)$:

$$\mathcal{S}_\lambda(\mathbf{Y}) = \text{sign}(\mathbf{Y}) \times \max(0, |\mathbf{Y}| - \lambda).$$

This proximal operator, as well as the proximal operator of any norm $\|\cdot\|_p$, $p \geq 1$, can also be related to projections [107]. Indeed, the Moreau decomposition gives us [107]:

$$\text{prox}_{\lambda \|\cdot\|_p}(\mathbf{Y}) = \mathbf{Y} - \Pi_{\mathbf{U}:\|\mathbf{U}\|_{\frac{p}{p-1}} \leq \lambda}(\mathbf{Y}), \quad (3.5)$$

where $\Pi_{\mathbf{U}:\|\mathbf{U}\|_{\frac{p}{p-1}} \leq \lambda}$ denotes the projection on the convex set $\{\mathbf{U} : \|\mathbf{U}\|_{\frac{p}{p-1}} \leq \lambda\}$. For the $\lambda \|\cdot\|_1$ norm, it leads to: $\text{prox}_{\lambda \|\cdot\|_1}(\mathbf{Y}) = \mathbf{Y} - \Pi_{\mathbf{U}:\|\mathbf{U}\|_\infty \leq \lambda}(\mathbf{Y})$.

3.1.2 Proximal algorithm: Forward-Backward Splitting

FB Algorithm. The Forward Backward Splitting algorithm (FB - [38]) aims to solve our former non-smooth convex problem:

$$\underset{\mathbf{Y}}{\operatorname{argmin}} f(\mathbf{Y}) + g(\mathbf{Y}),$$

where f is a proper, convex and differentiable function whose gradient is L -Lipschitz, and g is a proximable function (the decomposition of the problem is not unique).

As stated above, these criteria are valid for many convex inverse problems: the use of the squared Frobenius norm for the data fidelity term associated with a norm or an indicator function for the regularization is common and meet the conditions e.g., the LASSO or group-LASSO [61].

The FB algorithm is given in Alg.1. It converges towards the global minimum of the cost function as long as the step-size $\gamma \in (0, \frac{2}{L})$, [38].

Algorithm 1 FB algorithm

```

1: procedure FB( $\tilde{\mathbf{Y}}^{(0)}$ )
2:   while do not converge do
3:      $\tilde{\mathbf{Y}}^{(k+1)} \leftarrow \operatorname{prox}_{\gamma g}(\tilde{\mathbf{Y}}^{(k)} - \gamma \nabla f(\tilde{\mathbf{Y}}^{(k)}))$  ▷ Proximal gradient step
   return  $\tilde{\mathbf{Y}}^{(k)}$ .

```

Remarks. We only make few remarks concerning the interpretation of this algorithm, more details can be found in [107, Section 4.2].

First, if there is no regularization g , the FB boils down to the standard gradient descent, for minimizing f . On the other hand, if f is null, the FB amounts to minimize g using the proximal operator of g [107].

Besides, if g corresponds to the indicator function of a convex set (its proximal operator is the projection onto this set), then the FB algorithm corresponds to the *projected gradient algorithm* [94], [20].

The FB algorithm can also be seen as a *majorization-minimization algorithm* whenever $\gamma \leq \frac{1}{L}$. Hence, the $k + 1$ th update of the FB algorithm can be reformulated as:

$$\tilde{\mathbf{Y}}^{(k+1)} = \underset{\mathbf{Y}}{\operatorname{argmin}} f(\tilde{\mathbf{Y}}^{(k)}) + \nabla f(\tilde{\mathbf{Y}}^{(k)})^T (\mathbf{Y} - \tilde{\mathbf{Y}}^{(k)}) + \frac{1}{2\gamma} \left\| \mathbf{Y} - \tilde{\mathbf{Y}}^{(k)} \right\|_2^2 + g(\mathbf{Y}), \quad (3.6)$$

where the function to be minimized, generally denoted as *prox-linearization* [136], is a tight upper bound of the initial cost function at the point $\tilde{\mathbf{Y}}^{(k)}$.

The second order approximation of f at $\tilde{\mathbf{Y}}^{(k)}$, assuming that f is twice differentiable at $\tilde{\mathbf{Y}}^{(k)}$

with an Hessian $\mathbf{H}^{(k)}$ is given by:

$$f(\tilde{\mathbf{Y}}^{(k)}) + \nabla f(\tilde{\mathbf{Y}}^{(k)})^T (\mathbf{Y} - \tilde{\mathbf{Y}}^{(k)}) + \frac{1}{2} (\mathbf{Y} - \tilde{\mathbf{Y}}^{(k)})^T \mathbf{H}^{(k)} (\mathbf{Y} - \tilde{\mathbf{Y}}^{(k)}).$$

In the prox-linear update, the second order term is simply upper-bounded by $\frac{1}{2\gamma} \|\mathbf{Y} - \tilde{\mathbf{Y}}^{(k)}\|_2^2$. The prox-linearization is also used by the PALM algorithm [17] for Block Coordinate minimization as we will see later 3.2.4.

This algorithm is efficient if the proximal operator of the function g is explicit. Otherwise, if the regularization g corresponds to a combination of proximable terms, whose proximal operators are explicit, the Generalized Forward Backward Splitting algorithm can be employed.

Example: inverse problem with sparsity enforcing regularization. The FB algorithm can be used conveniently to solve the LASSO problem [61], corresponding to the minimization of the cost-function used for sparse BSS (3.11) for a fixed mixing matrix \mathbf{A} :

$$\operatorname{argmin}_{\mathbf{S}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_2^2 + \|\Lambda \odot \mathbf{S}\|_1$$

The data-fidelity term is convex, proper and differentiable with a L -Lipschitz gradient. The Lipschitz constant L is equal to the largest eigenvalue of $\mathbf{A}^T \mathbf{A}$, *i.e.* the spectral norm $\|\mathbf{A}^T \mathbf{A}\|_{2,S}$. The ℓ_1 norm is proximable. Hence, it is possible to minimize this cost-function with the FB algorithm. This popular algorithm is also known as ISTA (Iterative Soft-Thresholding Algorithm - [107]), and is presented in Alg.2.

Algorithm 2 ISTA

- 1: **procedure** ISTA(\mathbf{X}, \mathbf{A})
 - 2: Choose a step-size $L_S \in (0, \frac{2}{\|\mathbf{A}^T \mathbf{A}\|_{2,S}})$
 - 3: **while** not converge **do**
 - 4: $\tilde{\mathbf{S}}^{(j)} \leftarrow \mathcal{S}_{\Lambda \times L_S}(\tilde{\mathbf{S}}^{(j-1)} + L_S \times \tilde{\mathbf{A}}^T (\mathbf{X} - \tilde{\mathbf{A}} \tilde{\mathbf{S}}^{(j-1)}))$ \triangleright Proximal gradient update
 - 5: $j \leftarrow j + 1$
 - return** $\tilde{\mathbf{S}}^{(j)}$.
-

3.1.3 Generalized Forward Backward Splitting

GFBS Algorithm. Let us assume that the proximable function g can be split into several proximable functions, whose proximal operators are explicit: $g = \sum_{j=1}^k g_j$. The corresponding problem $\operatorname{argmin}_{\mathbf{Y}} f(\mathbf{Y}) + \sum_{j=1}^k g_j(\mathbf{Y})$ can be quite conveniently solved using the Generalized Forward Backward Splitting algorithm (GFBS- [110]) presented in Alg.3.

Algorithm 3 GFBS algorithm

```

1: procedure GFB( $\tilde{\mathbf{Y}}^{(0)}$ )
2:   Set  $\gamma \in (0, \frac{2}{L})$ ,  $\mu \in (0, \min(\frac{3}{2}, \frac{1+\frac{2}{L\gamma}}{2}))$ ,  $w_j \in (0, 1)$  such that  $\sum_1^k w_j = 1$ ,  $\mathbf{Z}_j \leftarrow \mathbf{Y}^{(0)}, \forall j = 1, \dots, n$ .
3:   while do not converge do
4:     for  $j = 1, \dots, n$  do
5:        $\mathbf{Z}_j \leftarrow \mathbf{Z}_j + \mu \text{prox}_{\frac{\gamma}{w_j} g_j}(2\mathbf{Y}^{(k)} - \mathbf{Z}_j - \gamma \nabla f(\mathbf{Y}^{(k)})) - \mu \mathbf{Y}^{(k)}$ 
6:        $\mathbf{Y}^{(k+1)} \leftarrow \sum_{j=1}^k w_j \mathbf{Z}_j$ 
7:        $k \leftarrow k + 1$ 
   return  $\tilde{\mathbf{Y}}^{(k)}$ .
```

```

1: procedure PALM-RGMCA( $\mathbf{X}, \mathbf{A}^{(0)}, \mathbf{S}^{(0)}, \mathbf{O}^{(0)}$ )
2:   while does not converge do
3:     Update  $L_A$ 
4:      $\mathbf{A}^{(k+1)} \leftarrow \text{prox}_{\frac{1}{L_A} g_A}(\mathbf{A}^{(k)} + \frac{1}{L_A}(\mathbf{X} - \mathbf{A}^{(k)}\mathbf{S}^{(k)} - \mathbf{O}^{(k)})\mathbf{S}^{(k)T})$ 
5:     Update  $L_S$ 
6:      $\mathbf{S}^{(k+1)} \leftarrow \text{prox}_{\frac{1}{L_S} g_S}(\mathbf{S}^{(k)} + \frac{1}{L_S}\mathbf{A}^{(k+1)T}(\mathbf{X} - \mathbf{A}^{(k+1)}\mathbf{S}^{(k)} - \mathbf{O}^{(k)}))$ 
7:      $\mathbf{O}^{(k+1)} \leftarrow \text{prox}_{g_O}(\mathbf{X} - \mathbf{A}^{(k+1)}\mathbf{S}^{(k)})$ 
8:      $k \leftarrow k + 1$ 
   return  $\mathbf{A}^{(k)}, \mathbf{S}^{(k)}, \mathbf{O}^{(k)}$ .
```

Example of application. The GFBS algorithm can be of interest for image denoising. Given a noisy version \mathbf{X} of a natural image \mathbf{Y} to be retrieved, sparsely represented in a transformed domain $\Phi_{\mathbf{S}}$ and non-negative, one can minimize the following cost-function:

$$\underset{\mathbf{Y}}{\text{argmin}} \frac{1}{2} \|\mathbf{X} - \mathbf{S}\|_2^2 + \|\Lambda \odot \mathbf{Y} \Phi_{\mathbf{S}}^T\|_1 + \chi_{\mathbf{U}: \mathbf{U} \geq 0}(\mathbf{Y}),$$

where the first term is the data-fidelity term, well suited to deal with Gaussian noise, the second term enforces the sparsity of the image in the given dictionary $\Phi_{\mathbf{S}}$, and the last term ensures that the recovered image is non-negative.

The regularization $\|\Lambda \odot \mathbf{Y} \Phi_{\mathbf{S}}^T\|_1 + \chi_{\mathbf{U}: \mathbf{U} \geq 0}(\mathbf{Y})$ is proximable but does not admit a closed form expression. However, $\|\Lambda \odot \mathbf{Y} \Phi_{\mathbf{S}}^T\|_1$ and $\chi_{\mathbf{U}: \mathbf{U} \geq 0}(\mathbf{Y})$ are proximable and their proximal operators have a closed form expression App.A: the GFBS algorithm can be interestingly employed.

3.2 Block Coordinate algorithms for component separation problems

We will now move to the review of some optimization strategies which can conveniently tackle sparse matrix-factorization problems.

The optimization-based approach has taken the lion's share for matrix-factorization and component separation problems. Thanks to the recent advances in optimization, it provides a flexible

framework in which our prior knowledge on the components can be easily promoted.

The special structure of BSS and component estimation problems makes the *Block Coordinate methods* (BC) particularly well suited to deal with the associated optimization problems. We will present in this section some of the standard BC methods¹ and their application to sparse BSS.

3.2.1 Framework

Considered cost-function. Let us assume that we aim to estimate n real components $\{\mathbf{X}_i\}_{i=1,\dots,n}$, with $\mathbf{X}_i \in \mathbb{R}^{m_i \times t_i}$, by minimizing a cost function h with the following expression:

$$h(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n) = f(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n) + \sum_{i=1}^n g_i(\mathbf{X}_i), \quad (3.7)$$

where the functions f and $\{g_i\}_{i=1,\dots,n}$ are such that:

- $f : \mathbb{R}^{m_1 \times t_1} \times \mathbb{R}^{m_2 \times t_2} \dots \times \mathbb{R}^{m_n \times t_n} \rightarrow \overline{\mathbb{R}}$ represents the data fidelity term. We assume that the observations are corrupted by an additive Gaussian noise. Consequently, *the function f is chosen as the quadratic loss* (squared Frobenius norm).

For every component $\mathbf{X}_i^{(k)}$ at the k th iteration of an algorithm, we define $f_i^k(\mathbf{X}) : \mathbb{R}^{m_i \times t_i} \rightarrow \overline{\mathbb{R}}, \mathbf{X} \mapsto f(\mathbf{X}_1^{(k)}, \mathbf{X}_2^{(k)}, \mathbf{X}_{i-1}^{(k)}, \mathbf{X}, \mathbf{X}_{i+1}^{(k-1)}, \dots)$ or more generally f_i , and ∇f_i^k the gradient of this function. The gradients ∇f_i^k are supposed to be L_i^k -Lipschitz.

The functions f_i of $\mathbf{X}_i, \forall i = 1, \dots, n$ are convex, but f is not necessarily convex for the tuple $(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n)$.

- Each real valued function g_i designates the *convex* penalization associated with the variable \mathbf{X}_i , such that $g_i : \mathbb{R}^{m_i \times t_i} \rightarrow \overline{\mathbb{R}}$. In this chapter and the remaining of this thesis, we will assume that these functions correspond to a combination of norms, mixed norms [83], and indicator functions of closed convex subsets, not empty. As so, they are real-valued, proper, lower semi-continuous, convex and continuous on their effective domain.

Examples.

Standard sparse BSS [14]. In GMCA, Section 1.3.2, the sources are assumed to be sparsely represented in a dictionary $\Phi_{\mathbf{S}}$. The sparsity of the sources in $\Phi_{\mathbf{S}}$ can be promoted by the penalization $\|\Lambda \odot \mathbf{S}\Phi_{\mathbf{S}}^T\|_1$, where Λ designates the regularization parameters. In order to avoid the scaling indeterminacy between \mathbf{A} and \mathbf{S} , the columns of \mathbf{A} are assumed to be normalized for the ℓ_2 norm. This feature is enforced by using the indicator function $\chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A})$:

$$\underset{\mathbf{A}, \mathbf{S}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_2^2 + \chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}) + \|\Lambda \odot \mathbf{S}\Phi_{\mathbf{S}}^T\|_1. \quad (3.8)$$

¹The algorithms are presented in a restrictive setting, corresponding to the framework in which they will be employed in the remaining of this thesis.

This problem and the associated cost-function are the basics of the remaining work in this thesis. As so, several implementations of this problem will be proposed in this chapter.

Non-negative BSS and sparsity in a transformed domain [113]. In many applications, the sources and the mixing matrix are non-negative. This can be the case in multispectral imaging for instance: the sources correspond to the spatial distributions of the components, and the columns of the mixing matrix to their spectra. The non-negativity is promoted by adding two characteristics functions $\chi_{\mathbf{Y}:\geq 0}$:

$$\underset{\mathbf{A}, \mathbf{S}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_2^2 + \|\Lambda \odot \mathbf{S}\Phi_{\mathbf{S}}^T\|_1 + \chi_{\mathbf{Y}:\geq 0}(\mathbf{S}) + \chi_{\mathbf{Y}:\geq 0}(\mathbf{A}) + \chi_{\mathbf{Y}:\|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}).$$

Unmixing of hyperspectral data [102]. In hyperspectral imaging, the spectra of the components are sampled at hundreds of contiguous frequencies. As so, they are reasonably smooth, and sparsely represented in the wavelet domain associated with the matrix Φ_A . It is thus possible, with the sparse representations of the sources and the spectra, to highlight the morphological and the spectral diversities between the sources with a cost function of the form:

$$\underset{\mathbf{A}, \mathbf{S}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_2^2 + \|\Lambda \odot \mathbf{S}\Phi_{\mathbf{S}}^T\|_1 + \chi_{\mathbf{Y}:\geq 0}(\mathbf{S}) + \|\Upsilon \odot \mathbf{A}\Phi_A^T\|_1 + \chi_{\mathbf{Y}:\geq 0}(\mathbf{A}),$$

where Υ denotes the regularization parameters.

We remark that in contrast with the previous problems, the roles of \mathbf{A} and \mathbf{S} are symmetric.

Sparse BSS in the presence of sparse outliers. Sparse deviations from the model can be encountered for instance in multispectral analysis: some elements not explicitly taken into account can contribute significantly to the total energy at a given wavelength (*e.g.* the CO at 100GHz in the Planck-ESA data). This contribution is row and column sparse in the same dictionary as the one of the sources $\Phi_{\mathbf{S}}$, what can be enforced using the ℓ_1 norm:

$$\underset{\mathbf{A}, \mathbf{S}, \mathbf{O}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S} - \mathbf{O}\|_2^2 + \chi_{\mathbf{Y}:\|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}) + \|\Lambda \odot \mathbf{S}\Phi_{\mathbf{S}}^T\|_1 + \|\Upsilon \odot \mathbf{O}\Phi_{\mathbf{S}}^T\|_1, \quad (3.9)$$

where Υ corresponds to the regularization parameters.

Overview of the minimization strategies. The resulting cost function h is a *multi-convex optimization problem*: it is globally non-convex but the subproblems \mathbf{P}_i with all but one variables fixed are convex:

$$\mathbf{P}_i : \underset{\mathbf{X}_i}{\text{argmin}} f_i(\mathbf{X}_i) + g_i(\mathbf{X}_i). \quad (3.10)$$

Hence, it is customary to optimize this type of cost function by iteratively and alternately minimize it with respect of each variable \mathbf{X}_i . This is commonly designated as *Block Coordinate method*.

We can distinguish four main minimization strategies:

- Projected least-squares: each subproblem is solved without taking into account the regularization, and then, the solution is projected on the constraints. It provides only an approximate solution, and may not converge.
- Block Coordinate Descent: each subproblem is exactly minimized.
- Prox-linear updates: the problem is solved by minimizing a prox-linear approximation of each subproblem.
- Proximal updates: each subproblem is approximated with a proximal regularization.

Details on these four strategies are provided below. For illustrative purpose, we will also present the algorithms associated with these 4 optimization methods for solving the sparse BSS problem 3.8.

3.2.2 Projected least-squares

PALS Algorithm. The projected least-squares² has a long history, and was first introduced in the NMF framework with the Alternating Least Squares [106].

The update of a component is two-steps: first, the least-squares estimate is computed (minimization of the subproblems without taking into the regularization), and then this solution is projected on the constraints with the proximal operator of the regularization. A general formulation of the projected alternating least squares (PALS) can be found in Alg.4.

Algorithm 4 Projected Alternating Least Squares Algorithm

```

1: procedure PALS( $\mathbf{X}_1^{(0)}, \mathbf{X}_2^{(0)}, \dots, \mathbf{X}_n^{(0)}$ )
2:   for  $k=1, \dots, K$  do                                     ▷ Outer loop -  $k$ th update of the components
3:     for  $i=1, \dots, n$  do                                   ▷ Inner loop - update of the  $i$ th component
4:        $\mathbf{X}_i^{(k)} \leftarrow \operatorname{argmin}_{\mathbf{X}_i} f_i^{(k)}(\mathbf{X}_i)$                                      ▷ Least-squares estimate
5:        $\mathbf{X}_i^{(k)} \leftarrow \operatorname{prox}_{g_i}(\mathbf{X}_i^{(k)})$                                        ▷ Projection/Proximal operator
   return  $\mathbf{X}_1^{(k)}, \mathbf{X}_2^{(k)}, \dots, \mathbf{X}_n^{(k)}$ .

```

Remarks. PALS suffers from two drawbacks:

- First, it is only interesting if the proximal operators of the regularizations g_i are explicit, so that the 'proximal projections' are one step. In some applications, this is not the case and PALS loses *its main interests: simplicity and rapidity*.
- Second, given that the cost function and the sub-problems \mathbf{P}_i are *not properly minimized* (in general), the implementation Alg.4 may be not precise.

If the problem is simple (e.g. standard NMF or sparse BSS), then it generally performs fairly.

²The so-called 'Projected Least-Squares' method encompasses indeed the standard projected least-squares approach, but is presented in a more general framework. The term f is not restricted to the quadratic loss, and the regularization to the indicator function of a closed convex set.

However, if the problem is slightly more complex (e.g. sparsity and non-negativity constraints in the direct domain), then a significant performance discrepancy between PALS and more rigorous minimization schemes are generally encountered [112].

For the same reason, *this algorithm cannot be proved to converge*, or even to be stable. In fact, the implementation Alg.4 can even diverge. We will see in the following example that some modifications of PALS can be added to improve the stability of the algorithm.

Projected least squares are convenient because they are easy to implement and fast. However, while the sparsity constraint can be correctly handled using this strategy, it becomes necessary to take into account the different constraints simultaneously and more rigorously in complex settings, for instance with the two next strategies.

Example. The standard sparse BSS problem, with sources sparsely represented in the direct domain is given by:

$$\underset{\mathbf{A}, \mathbf{S}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS}\|_2^2 + \chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}) + \|\Lambda \odot \mathbf{S}\|_1. \quad (3.11)$$

The Block Coordinate methods including PALS aim to recover \mathbf{A} and \mathbf{S} by alternatively updating the components, whose associated subproblems are given by:

$$\mathbf{P}_S : \underset{\mathbf{S}}{\text{argmin}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS}\|_2^2 + \|\Lambda \odot \mathbf{S}\|_1 \quad (3.12)$$

$$\mathbf{P}_A : \underset{\mathbf{A}}{\text{argmin}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS}\|_2^2 + \chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}) \quad (3.13)$$

For PALS, the updates of the components are given by:

- Source Update: The least-squares estimate of \mathbf{S} is \mathbf{XA}^\dagger . The proximal operator of the ℓ_1 is the soft-thresholding operator. Hence, the update of \mathbf{S} at the k th iteration is given by:

$$\tilde{\mathbf{S}}^{(k)} = \mathcal{S}_\Lambda(\mathbf{XA}^{(k)\dagger}).$$

- Mixing matrix Update: Similarly, the least-squares estimate of \mathbf{A} is $\mathbf{S}^\dagger \mathbf{X}$. The proximal operator of the indicator function is approximated by the projection on the ℓ_2 ball. At the k th iteration of the algorithm, the update of \mathbf{A} is thus:

$$\tilde{\mathbf{A}}^{(k)i} \leftarrow \frac{(\mathbf{XS}^{(k)\dagger})^i}{\|(\mathbf{XS}^{(k)\dagger})^i\|_2}, \forall i = 1..n$$

The sensitivity to the condition number of the variables is the main default of this approach, based on pseudo-inversion.

To overcome this weakness, we propose to replace the pseudo-inversion with an 'interrupted'

Algorithm 5 Modified PALS for sparse BSS

```

1: procedure MPALS( $\mathbf{X}, \tilde{\mathbf{A}}^{(0)}$ )
2:   for  $k=1, \dots, K$  do
3:     if  $\text{cond}(\tilde{\mathbf{A}}^{(k-1)}) < 10000$  then                                 $\triangleright$  Pseudo-inverse if small condition number
4:        $\tilde{\mathbf{S}}^{(k)} \leftarrow \tilde{\mathbf{A}}^{(k-1)\dagger} \mathbf{X}$ 
5:     else:                                                                 $\triangleright$  Incomplete Gradient Descent
6:        $L = \|\tilde{\mathbf{A}}^{(k-1)T} \tilde{\mathbf{A}}^{(k-1)}\|_{2,s}$ 
7:       for  $j=1, \dots, 100$  do:
8:          $\mathbf{S}^{(k)} \leftarrow \mathbf{S}^{(k)} + \frac{1}{L} \tilde{\mathbf{A}}^{(k-1)T} (\mathbf{X} - \tilde{\mathbf{A}}^{(k-1)} \mathbf{S}^{(k)})$ 
9:        $\tilde{\mathbf{S}}^{(k)} \leftarrow \mathcal{S}_\Lambda(\tilde{\mathbf{S}}^{(k)})$                                            $\triangleright$  Soft-thresholding
10:      if  $\text{cond}(\tilde{\mathbf{S}}^{(k)}) < 100$  then                                        $\triangleright$  Pseudo-inverse if small condition number
11:         $\tilde{\mathbf{A}}^{(k)} \leftarrow \mathbf{X} \tilde{\mathbf{S}}^{(k)\dagger}$ 
12:      else:                                                                 $\triangleright$  Incomplete Gradient Descent
13:         $L = \|\tilde{\mathbf{S}}^{(k-1)} \tilde{\mathbf{S}}^{(k-1)T}\|_{2,s}$ 
14:        for  $j=1, \dots, 100$  do:
15:           $\tilde{\mathbf{A}}^{(k)} \leftarrow \tilde{\mathbf{A}}^{(k)} + \frac{1}{L} (\mathbf{X} - \tilde{\mathbf{A}}^{(k)} \mathbf{S}^{(k)}) \tilde{\mathbf{S}}^{(k)T}$ 
16:           $\tilde{\mathbf{A}}^{(k)i} \leftarrow \frac{(\tilde{\mathbf{A}}^{(k)i})^i}{\|\tilde{\mathbf{A}}^{(k)i}\|_2}, \forall i = 1, \dots, n$                  $\triangleright$  Projection
return  $\tilde{\mathbf{A}}^{(k)}, \tilde{\mathbf{S}}^{(k)}$ .

```

gradient descent for the least-squares estimate if necessary, Alg.5. With the incomplete inversion, the components are still updated in the directions of the gradient but there is not any strong amplification of some components due to a large condition number. This greatly enhances the stability of PALS but does not dramatically increase the computational time.

3.2.3 Block Coordinate Descent

BCD Algorithm. The Block Coordinate Descent (BCD, [128]) consists in estimating alternatively and exactly the subproblems \mathbf{P}_i such as presented in Alg. 6.

Remarks. In contrast with PALS, BCD can be proved to converge toward a stationary point of the cost function h under mild conditions [128], [136], which are fulfilled by the problems in this thesis.

Unfortunately, minimizing the subproblems may be not easy since the regularization functions may be not differentiable. Different minimizers based on proximal calculus [38] can be deployed for that purpose: the conditions that we have imposed on f and g_i allow us to solve the subproblems with the FB algorithm (or GFBS if necessary) - see Section 3.1. We will illustrate how BCD can be employed together with proximal calculus techniques to solve the sparse BSS problem (3.11).

Algorithm 6 Block Coordinate Descent

```

1: procedure BCD( $\mathbf{X}_1^{(0)}, \mathbf{X}_2^{(0)}, \dots, \mathbf{X}_n^{(0)}$ )
2:   while do not converge do
3:     for  $i=1, \dots, n$  do
4:        $\mathbf{X}_i^{(k)} \leftarrow \operatorname{argmin}_{\mathbf{X}_i} f_i^k(\mathbf{X}_i) + g_i(\mathbf{X}_i)$ 
5:      $k \leftarrow k + 1$ 
   return  $\mathbf{X}_1^{(k)}, \mathbf{X}_2^{(k)}, \dots, \mathbf{X}_n^{(k)}$ .

```

Example. The two subproblems associated with the updates of \mathbf{A} and \mathbf{S} are given in eq.3.13 and 3.12. The two regularization functions are not smooth, but the FB algorithm can be used to solve the subproblems efficiently.

Indeed, the proximal operators of $\chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A})$ and $\|\Lambda \odot \mathbf{S}\|_1$ are explicit (have a closed form expression). The quadratic term is differentiable with L_A and L_S Lipschitz gradients for \mathbf{A} and \mathbf{S} respectively³. We point out that updating \mathbf{S} with the FB strategy corresponds to the ISTA algorithm Alg.2.

The BCD implementation of standard sparse BSS eq.3.8 with the FB implementation for solving the subproblems is given in Alg.7.

Algorithm 7 Block Coordinate Descent for sparse BSS with FB implementation of the subproblems

```

1: procedure BCD SPARSE BSS( $\mathbf{X}, \tilde{\mathbf{A}}^{(0)}$ )
2:   while not converge do ▷ BCD
3:      $\tilde{\mathbf{S}}^{(0)} \leftarrow \tilde{\mathbf{S}}^{(k)}$ 
4:     Compute  $L_S = \|\tilde{\mathbf{A}}^{(k)T} \tilde{\mathbf{A}}^{(k)}\|_{s,2}$ 
5:     while not converge do ▷ FB for PS-ISTA
6:        $\tilde{\mathbf{S}}^{(j)} \leftarrow \mathcal{S}_{\frac{\Lambda}{L_S}} \left( \tilde{\mathbf{S}}^{(j-1)} + \frac{1}{L_S} \tilde{\mathbf{A}}^{(k)T} (\mathbf{X} - \tilde{\mathbf{A}}^{(k)} \tilde{\mathbf{S}}^{(j-1)}) \right)$ 
7:        $j \leftarrow j + 1$ 
8:      $\tilde{\mathbf{S}}^{(k+1)} \leftarrow \tilde{\mathbf{S}}^{(j)}$ 
9:     Compute  $L_A = \|\tilde{\mathbf{S}}^{(k)} \tilde{\mathbf{S}}^{(k)T}\|_{s,2}$ 
10:     $\tilde{\mathbf{A}}^{(0)} \leftarrow \tilde{\mathbf{A}}^{(k)}$ 
11:    while not converge do ▷ FB for PA
12:       $\tilde{\mathbf{A}}^{(j)} \leftarrow \tilde{\mathbf{A}}^{(j-1)} + \frac{1}{L_A} (\mathbf{X} - \tilde{\mathbf{A}}^{(j-1)} \tilde{\mathbf{S}}^{(k+1)}) \tilde{\mathbf{S}}^{(k+1)T}$  ▷ Gradient descent
13:       $(\tilde{\mathbf{A}}^{(j)})^i \leftarrow \frac{(\tilde{\mathbf{A}}^{(j)})^i}{\max(1, \|(\tilde{\mathbf{A}}^{(j)})^i\|_2)} \quad \forall i = 1, \dots, n$  ▷ Proximal Operator
14:       $j \leftarrow j + 1$ 
15:     $\tilde{\mathbf{A}}^{(k+1)} \leftarrow \tilde{\mathbf{A}}^{(j)}$ 
16:     $k \leftarrow k + 1$ 
  return  $\tilde{\mathbf{A}}^{(k)}, \tilde{\mathbf{S}}^{(k)}$ .

```

The BCD strategy converges towards a stationary point with this setting [128, Lemma 3.1 a)

³Since the variables evolve during the iterative process, the Lipschitz constants are not constant and need to be reevaluated at every loop.

and Theorem 4.1 b):

- The level sets on $\mathcal{X} = \{(\tilde{\mathbf{A}}, \tilde{\mathbf{S}}) : h(\tilde{\mathbf{A}}, \tilde{\mathbf{S}}) \leq h(\tilde{\mathbf{A}}^{(0)}, \tilde{\mathbf{S}}^{(0)})\}$ are compact. The cost function h is continuous on \mathcal{X} , if the starting tuple indeed satisfies the conditions enforced by the indicator function.
- The cost function is multi-convex.
- The quadratic term is differentiable and its effective domain is open.

3.2.4 Prox-linear updates

PALM Algorithm. This strategy, coined Proximal Alternating Linearized Minimization (PALM - [136], [17]) can be seen as a generalization of the FB algorithm to the framework of Block Coordinate minimization. It updates alternatively the components by minimizing a prox-linearization of each subproblem 3.6. At the k th iteration of the algorithm, the i th component is updated with

$$\mathbf{X}_i^{(k+1)} = \underset{\mathbf{X}_i}{\operatorname{argmin}} f_i^k(\mathbf{X}_i^{(k)}) + \nabla f_i^k(\mathbf{X}_i^{(k)})^T (\mathbf{X}_i - \mathbf{X}_i^{(k)}) + \frac{L_i^k}{2} \|\mathbf{X}_i - \mathbf{X}_i^{(k)}\|_2^2 + g_i(\mathbf{X}_i),$$

where L_i^k denotes a numerical constant, for which we will take the Lipschitz constant of ∇f_i^k in the following. This can be reformulated using the proximal operator of the regularization g_i :

$$\mathbf{X}_i^{(k+1)} = \operatorname{prox}_{\frac{1}{L_i^k} g_i}(\mathbf{X}_i^{(k)} - \frac{1}{L_i^k} \nabla f_i^k(\mathbf{X}_i^{(k)})).$$

The PALM algorithm for the general framework is presented in Alg.8.

Algorithm 8 PALM algorithm

```

1: procedure PALM( $\tilde{\mathbf{X}}_1^{(0)}, \tilde{\mathbf{X}}_2^{(0)}, \dots, \tilde{\mathbf{X}}_n^{(0)}$ )
2:   while do not converge do
3:     for  $i=1, \dots, n$  do
4:       Compute the Lipschitz constant  $L_i^k$  of  $\nabla f_i^k$ 
5:        $\tilde{\mathbf{X}}_i^{(k+1)} \leftarrow \operatorname{prox}_{\frac{1}{L_i^k} g_i} \left( \tilde{\mathbf{X}}_i^{(k)} - \frac{1}{L_i^k} \nabla f_i^k \left( \tilde{\mathbf{X}}_i^{(k)} \right) \right)$  ▷ Proximal gradient step
6:      $k \leftarrow k + 1$ 
   return  $\tilde{\mathbf{X}}_1^{(k)}, \tilde{\mathbf{X}}_2^{(k)}, \dots, \tilde{\mathbf{X}}_n^{(k)}$ .

```

Remarks. The PALM algorithm converges towards a stationary point in our framework [17, Theorem 3.1]:

- The functions g_i are proximable, and satisfy the Kurdyka-Lojasiewicz property (see the examples in [17] and [136]).
- f (quadratic loss) is C^2 and also satisfies the Kurdyka-Lojasiewicz property.

- The function f is differentiable for the n variables, and the gradient of the i th component at the k th iteration is L_i^k - Lipschitz, L_i^k is bounded.

In contrast with BCD, PALM does not exactly minimize the subproblems, the components are only updated with one proximal gradient step Alg.9 and Alg.7. From this point of view, PALM differs from the BCD and the projected gradient as well: the updates of PALM are functions of all the variables, whereas for BCD and PALS, the update of the component \mathbf{X}_i does not depend on the previous estimate of \mathbf{X}_i .

Example. The specific implementation for solving sparse BSS (3.11) can be found in Alg.9.

Algorithm 9 PALM sparse BSS

```

1: procedure PALM SPARSE BSS( $\tilde{\mathbf{A}}^{(0)}, \tilde{\mathbf{S}}^{(0)}$ )
2:   while do not converge do
3:     Compute  $L_S = \|\tilde{\mathbf{A}}^{(k-1)T} \tilde{\mathbf{A}}^{(k-1)}\|_{s,2}$ 
4:      $\tilde{\mathbf{S}}^{(k)} \leftarrow \mathcal{S}_{\frac{\mathbf{A}}{L_S}} \left( \tilde{\mathbf{S}}^{(k-1)} + \frac{1}{L_S} \tilde{\mathbf{A}}^{(k-1)T} (\mathbf{X} - \tilde{\mathbf{A}}^{(k-1)} \tilde{\mathbf{S}}^{(k-1)}) \right)$   $\triangleright$  Proximal gradient step for  $\mathbf{S}$ 
5:     Compute  $L_A = \|\tilde{\mathbf{S}}^{(k)} \tilde{\mathbf{S}}^{(k)T}\|_{s,2}$ 
6:      $\tilde{\mathbf{A}}^{(k)} \leftarrow \tilde{\mathbf{A}}^{(k-1)} + \frac{1}{L_A} (\mathbf{X} - \tilde{\mathbf{A}}^{(k-1)} \tilde{\mathbf{S}}^{(k)}) \tilde{\mathbf{S}}^{(k)T}$   $\triangleright$  Gradient descent for  $\mathbf{A}$ 
7:      $(\tilde{\mathbf{A}}^{(k)})^i \leftarrow \frac{(\tilde{\mathbf{A}}^{(k)})^i}{\max(1, \|(\tilde{\mathbf{A}}^{(0)})^i\|_2)} \quad \forall i = 1, \dots, n$   $\triangleright$  Proximal Operator for  $\mathbf{A}$ 
8:      $k \leftarrow k + 1$ 
   return  $\tilde{\mathbf{A}}^{(k)}, \tilde{\mathbf{S}}^{(k)}$ .

```

The difference between PALM and BCD is clearly appearing with this problem: only one proximal gradient step updates the components with PALM Alg.9, whereas each component is updated until convergence with a proximal gradient step with BCD Alg.7

3.2.5 Proximal updates

PBC Algorithm. This last algorithm, which has been studied in [136], [6], minimizes alternatively a proximal approximation of each subproblem. It has not been intensively employed in the literature, neither during this thesis. However, it has interesting properties which deserve to be mentioned.

For the i th component \mathbf{X}_i , the modified subproblem to be minimized at the k th iteration is given by:

$$\mathbf{X}_i^{(k+1)} \leftarrow \underset{\mathbf{X}_i}{\operatorname{argmin}} f_i^k(\mathbf{X}_i) + g_i(\mathbf{X}_i) + \frac{1}{2L_i^k} \|\mathbf{X}_i - \mathbf{X}_i^{(k)}\|_2^2, \quad (3.14)$$

or equivalently

$$\mathbf{X}_i^{(k+1)} \leftarrow \underset{L_i^k(f_i^k + g_i)}{\operatorname{prox}} (\mathbf{X}_i^{(k)}),$$

where L_i^k is a numerical constant, strictly lower bounded by 0 and upper bounded (cannot be ∞). Hence, this version amounts to minimize locally the cost function, as a proximal operator does.

Algorithm 10 Proximal Block Coordinate algorithm

```

1: procedure PBC( $\tilde{\mathbf{X}}_1^{(0)}, \tilde{\mathbf{X}}_2^{(0)}, \dots, \tilde{\mathbf{X}}_n^{(0)}$ )
2:   while do not converge do
3:     for i=1..n do
4:       Choose the constant  $L_i^k$ ,
5:        $\tilde{\mathbf{X}}_i^{(k+1)} \leftarrow \operatorname{argmin}_{\tilde{\mathbf{X}}_i} f_i^k(\tilde{\mathbf{X}}_i) + g_i(\tilde{\mathbf{X}}_i) + \frac{1}{2L_i^k} \left\| \tilde{\mathbf{X}}_i - \tilde{\mathbf{X}}_i^{(k)} \right\|_2^2$ 
6:      $k \leftarrow k + 1$ 
   return  $\tilde{\mathbf{X}}_1^{(k)}, \tilde{\mathbf{X}}_2^{(k)}, \dots, \tilde{\mathbf{X}}_n^{(k)}$ .

```

Remarks. Similarly to PALM, PCB does not exactly minimize each subproblem, and the update of $\mathbf{X}_i^{(k)}$ depends on $\mathbf{X}_i^{(k-1)}$. However, the approximation is different, as well as its leading effect on the algorithm: in the PALM algorithm, the function f_i^k is replaced by its local second order approximation (whose curvature is upper bounded thanks to the Lipschitz-constant of the gradient of f_i^k), whereas for the proximal update, only a proximity term is added (the curvature of the cost function is modified and 'increased').

Besides, if the proximal operators of g_i are explicit, then every step of PALM has a closed form expression. In contrast, minimizing eq.3.6 may necessitate the use of the FB algorithm. In that case, the Proximal Block Coordinate (PBC) implementation with FB sub-routines is given in Alg.11.

Algorithm 11 Proximal Block Coordinate algorithm

```

1: procedure PBC-FB( $\tilde{\mathbf{X}}_1^{(0)}, \tilde{\mathbf{X}}_2^{(0)}, \dots, \tilde{\mathbf{X}}_n^{(0)}$ )
2:   while do not converge do
3:     for i=1..n do
4:       Choose the constant  $L_i^k$ , and the Lipschitz constant  $L_i^k$  of  $(\nabla f_i^k + \frac{1}{L_i^k} \mathbf{I})$ 
5:        $\tilde{\mathbf{X}}_i^{(0)} \leftarrow \tilde{\mathbf{X}}_i^{(k-1)}$ 
6:       while do not converge do
7:          $\tilde{\mathbf{X}}_i^{(j+1)} \leftarrow \operatorname{prox}_{\frac{1}{L_i^k} g_i} \left( \tilde{\mathbf{X}}_i^{(j)} - \frac{1}{L_i^k} \left( \nabla f_i^k \left( \tilde{\mathbf{X}}_i^{(j)} \right) + \frac{1}{L_i^k} \left( \tilde{\mathbf{X}}_i^{(j)} - \tilde{\mathbf{X}}_i^{(k-1)} \right) \right) \right)$ 
8:          $j \leftarrow j + 1$ 
9:        $\tilde{\mathbf{X}}_i^{(k)} \leftarrow \tilde{\mathbf{X}}_i^{(j)}$ 
10:     $k \leftarrow k + 1$ 
   return  $\tilde{\mathbf{X}}_1^{(k)}, \tilde{\mathbf{X}}_2^{(k)}, \dots, \tilde{\mathbf{X}}_n^{(k)}$ .

```

This algorithm is quite similar to the BCD implementation, using the FB algorithm for solving

the subproblems. However, the proximal gradient steps are different since a proximity term has been added in PBC. This proximal gradient step, for PBC, can be also written as:

$$\tilde{\mathbf{X}}_i^{(j+1)} \leftarrow \text{prox}_{\frac{1}{L_i^k} g_i} \left(\tilde{\mathbf{X}}_i^{(j)} + \underbrace{\frac{1}{L_i^k L_i^k} (\tilde{\mathbf{X}}_i^{(k-1)} - \tilde{\mathbf{X}}_i^{(j)})}_{\text{'extrapolation'}} - \frac{1}{L_i^k} (\nabla f_i^k (\tilde{\mathbf{X}}_i^{(j)})) \right).$$

Extrapolation is commonly used to accelerate FB implementation (see FISTA for instance [8]). In our case, this is not, strictly speaking an extrapolation, since the weighting depends on the previous estimate of the outer-loop and not of the FB implementation. However, one can still expect PBC to require less outer loops than BCD, thanks to the proximity term: there are less strong changes from one update to another, what can smooth the regularization path, and thus leads to a fastest convergence, especially if not initialized far from the optimum.

This proximal gradient step can also be rewritten as:

$$\tilde{\mathbf{X}}_i^{(j+1)} \leftarrow \text{prox}_{\frac{1}{L_i^k} g_i} \left(\tilde{\mathbf{X}}_i^{(j)} + \frac{1}{L_i^k L_i^k} \tilde{\mathbf{X}}_i^{(k-1)} - \frac{1}{L_i^k} \underbrace{\left(\nabla f_i^k + \frac{1}{L_i^k} \mathbf{I} \right)}_{\text{'preconditioning'}} (\tilde{\mathbf{X}}_i^{(j)}) \right).$$

The presence of the proximity operator can also be seen as a sort of Tikhonov regularization. If f_i^k has a large condition number (the Lipschitz constant of its gradient is likely to be large), this solution, with an appropriate L_i^k can greatly fasten the process, and may enhance the solutions.

Therefore, PBC may be a good alternative to BCD. This strategy is shown to converge for our setting in [136], and [6]. Nonetheless, no systematic and efficient strategy for the choice of the L_i^k has been proposed in the literature, yet.

Example. The PBC implementation for solving sparse BSS eq.3.8 in the direct domain is presented in Alg.12.

3.2.6 Remarks and summary

Summary. In this section, we have focused on the most common optimization strategies for sparse matrix factorization. A summary can be found in Tab.3.1.

A more general framework. Last, we recall that we have restricted the framework of these optimization strategies on several points, which are not exhaustive:

Algorithm 12 Proximal Block Coordinate for sparse BSS with FB implementation of the sub-problems

```

1: procedure PBC SPARSE BSS( $\mathbf{X}, \tilde{\mathbf{A}}^{(0)}$ )
2:   while not converge do ▷ PBC
3:      $\tilde{\mathbf{S}}^{(0)} \leftarrow \tilde{\mathbf{S}}^{(k)}$ 
4:     Choose the term  $L$  and compute  $L_S = \|\tilde{\mathbf{A}}^{(k)T} \tilde{\mathbf{A}}^{(k)} + \frac{1}{L} \mathbf{I}\|_{s,2}$ 
5:     while not converge do ▷ FB
6:        $\tilde{\mathbf{S}}^{(j)} \leftarrow \mathcal{S}_{\frac{\Lambda}{L_S}} \left( \tilde{\mathbf{S}}^{(j-1)} + \frac{1}{L_S} (\tilde{\mathbf{A}}^{(k)T} (\mathbf{X} - \tilde{\mathbf{A}}^{(k)} \tilde{\mathbf{S}}^{(j-1)}) + \frac{1}{L} (\tilde{\mathbf{S}}^{(j-1)} - \tilde{\mathbf{S}}^{(k-1)})) \right)$ 
7:        $j \leftarrow j + 1$ 
8:      $\tilde{\mathbf{S}}^{(k+1)} \leftarrow \tilde{\mathbf{S}}^{(j)}$ 
9:     Choose  $L$  and compute  $L_A = \|\tilde{\mathbf{S}}^{(k)} \tilde{\mathbf{S}}^{(k)T} + \frac{1}{L} \mathbf{I}\|_{s,2}$ 
10:     $\tilde{\mathbf{A}}^{(0)} \leftarrow \tilde{\mathbf{A}}^{(k)}$ 
11:    while not converge do ▷ FB
12:       $\tilde{\mathbf{A}}^{(j)} \leftarrow \tilde{\mathbf{A}}^{(j-1)} + \frac{1}{L_A} ((\mathbf{X} - \tilde{\mathbf{A}}^{(j-1)} \tilde{\mathbf{S}}^{(k+1)}) \tilde{\mathbf{S}}^{(k+1)T} + \frac{1}{L} (\tilde{\mathbf{A}}^{(j-1)} - \tilde{\mathbf{A}}^{(k-1)}))$  ▷
13:      Gradient descent
14:       $(\tilde{\mathbf{A}}^{(j)})^i \leftarrow \frac{(\tilde{\mathbf{A}}^{(j)})^i}{\max(1, \|(\tilde{\mathbf{A}}^{(j)})^i\|_2)} \quad \forall i = 1, \dots, n$  ▷ Proximal Operator
15:       $j \leftarrow j + 1$ 
16:       $\tilde{\mathbf{A}}^{(k+1)} \leftarrow \tilde{\mathbf{A}}^{(j)}$ 
17:       $k \leftarrow k + 1$ 
18:    return  $\tilde{\mathbf{A}}^{(k)}, \tilde{\mathbf{S}}^{(k)}$ .
```

- Convexity: The regularizations are assumed to be convex, whereas it is not necessary: the ℓ_0 pseudo-norm can also be used for example [17].
- Mixed strategy: The three strategies have been presented separately, but it is possible to have a 'mixed approach' by minimizing completely one subproblem, and using a prox-linearization for the other variables [136].
- Cyclic updates: The updates with the PALM algorithm do not need to be cyclic (i.e. the i th component is updated at the k th iteration if $k \equiv 0 \pmod{i}$, k is a multiple of i): several gradient steps can be done for one of the variable and the components can be updated with an non-cyclic order [34].
- Acceleration and preconditioning: Accelerated versions (for FB [107], [8] and PALM [136]) have not been presented. As well, adaptive preconditioned versions (GFBS [111], or PALM [34]) are of interest to fasten the process.

Sparse BSS. Last, we point out that in the framework of the sparse BSS, if the sources and mixing matrix are orthonormal, PALS, BCD and PALM lead to the same algorithm. Otherwise, they are all different and may converge (PALS cannot be guaranteed to converge) to different stationary points.

Strategies	Advantages	Weaknesses	Comments
Projected Gradient / PALS [14], [106]	Fast. Easy to implement.	May not converge. Pointless if the proximal operators of the regularizations are not explicit.	Lack of stability and precision.
BCD [128], [136]	Provably convergent.	Can be time consuming.	Exact minimization of each subproblem.
Prox-linear - PALM [17], [136]	Provably convergent. Can be fast. General framework.	Does not minimize the subproblems exactly. Computationally efficient only if the proximal operators are explicit.	The updates of the components depend on the current estimates.
Proximal - PBC [6] [136]	Provably convergent. Can be faster than BCD. General framework.	Does not minimize the subproblems exactly. Computationally less efficient than PALS and PALM.	The updates of the components depend on the current estimates. Choice of the numerical constants for the proximity terms.

Table 3.1: Outline of some common optimization strategies for sparse matrix factorization

3.3 Heuristics for sparse BSS

The recent advances in optimization, in particular in proximal calculus, have broadened the variety of cost-functions that can be proposed and minimized efficiently. However, proposing a data model and the associated cost function is just the tip of the iceberg.

Two main problems arise in practice: setting the regularization parameters and proposing an efficient minimization scheme for the considered problem.

Indeed, setting the hyperparameters involved in cost functions has a crucial influence on the returned solutions. For instance in eq.3.8, the parameter Λ determines the trade-off between the reconstruction error and the sparsity level of the sources. Unfortunately, they are challenging to determine, since only few information is known a priori.

Second, with the general framework of proximal calculus, numerous optimization schemes are now available to minimize various cost-functions. However, they do not exhibit similar performances, in term of computational time, precision or reliability. This last feature is of particular importance in some real-world applications for which the multi-starts approach is unreasonable.

In this section, we introduce the strategies implemented during this thesis for solving sparse BSS: the automatic setting of the thresholds and a two-steps minimization scheme (robust ini-

tialization and precise refinement). We present and extend the rationales of the strategies first introduced in [14], [112].

3.3.1 Influence of the thresholds: illustration with toy-examples

We start by discussing the influence of the threshold values in the framework of standard sparse BSS 3.11. For this purpose, we will compare the performances of the different algorithms⁴ solving 3.8 when varying the values of Λ . The experiments will be carried on 3 illustrative settings.

3.3.1.1 Data setting. For illustration purposes, we consider 3 settings. In these 3 experiments, 10 sources are mixed into 10 observations, with $t = 4096$ samples. They are mixed with the same mixing matrix \mathbf{A} , whose entries are drawn from a centered Gaussian law and with its columns normalized. The experiments then differ on the following points:

- Case A - Ideal setting - fig.3.1a: the sources are exactly sparse, generated from a Bernoulli-Gaussian law, with an activation parameter of 10%, and a standard deviation of 1000 (centered Gaussian law). This case is almost noiseless (SNR of 120 dB).
- Case B - Noisy setting - fig.3.1b: the same sources as for case A are generated. In contrast with case A, the observations are corrupted with a Gaussian noise \mathbf{N} . The SNR is set to 30 dB.
- Case C - Approximately sparse sources - fig.3.1c: the sources are generated with a Generalized Gaussian law with a probability density function given by $\frac{\rho}{2\Gamma(\frac{1}{\rho})} \exp^{-|x|^\rho}$, $\forall x \in \mathbb{R}$ (unit variance and centered generalized Gaussian law with shape parameter ρ , noted $\mathcal{G}(\rho)$) with shape parameter $\rho = 0.5$. The amplitude of the sources is set so that the energy of these sources is equal to the energy of the sources in the cases A and B (with an amplitude of the significant spikes around 1000). The SNR is set to 120dB. Expansion coefficients of natural images using wavelets would have similar distributions. Thus, this setting is more realistic.




Figure 3.1: Scatter plot of two projected sources $(\mathbf{A}^\dagger \mathbf{X})_1$ and $(\mathbf{A}^\dagger \mathbf{X})_2$ for the 3 cases.

The initialization matrix $\tilde{\mathbf{A}}^{(0)}$ is generated so that $\tilde{\mathbf{A}}^{(0)} = \mathbf{A} + d\mathbf{A}$. The entries of $d\mathbf{A}$ are

⁴The PBC implementation has not been intensively exploited during this thesis. It will not be used in this second section.

drawn from a Gaussian law, centered. The columns of $d\mathbf{A}$ have a norm equal to $2 \times \sin(\frac{\theta}{2})$, where θ is a scalar that will vary in the following experiments. The columns of the mixing matrix are then normalized. This assures that the spectral angles made between the initial columns of \mathbf{A} and the initialization matrix $\tilde{\mathbf{A}}^{(0)}$ are at most θ : $\theta \geq \max_i \arccos(\langle \tilde{\mathbf{A}}^{(0)}, \mathbf{A}^i \rangle)$, in degree.

3.3.1.2 Protocol and performance indicators. We employ the algorithms 7, 9, 5 for solving eq.3.11. They share a similar stopping criterion $\frac{\|\tilde{\mathbf{s}}^{(k)} - \tilde{\mathbf{s}}^{(k-1)}\|_2}{\|\tilde{\mathbf{s}}^{(k)}\|_2} < 10^{-7}$. The FB algorithms in BCD are stopped whenever $\frac{\|\tilde{\mathbf{s}}^{(j)} - \tilde{\mathbf{s}}^{(j-1)}\|_2}{\|\tilde{\mathbf{s}}^{(j)}\|_2} < 10^{-8}$. The maximal numbers of iterations are large enough for the FB, BCD and PALM algorithms and they reach the stopping criterion. It may be not the case for PALS, which is stopped after 2000 iterations. The values of the spectral angles θ will vary between 0 and 45 degrees. The estimated sources are initialized with $\mathcal{S}_\Lambda(\tilde{\mathbf{A}}^{(0)\dagger} \mathbf{X})$, where Λ varies between 1 and 500.

For the different values of Λ and θ , we will study the following performance indicators:

- Maximal angle (MA) made between the recovered and ground truth columns of \mathbf{A} : $\max_{j=1..n} \arccos(\langle \tilde{\mathbf{A}}^j, \mathbf{A}^j \rangle)$, in degree. This illustrates the unmixing performance of the algorithms, with a conservative point of view (only the largest angle).
- The minimal value of the SDR App.C. Even if \mathbf{A} is correctly recovered, the sources can be not precisely estimated. The Signal to Distortion Ratio presented in [131] provides a global criterion on the source estimation. A large value of SDR means that the recovered sources are correctly estimated, with small interferences, artefacts or noise residual.
- The computational time (CT), in second.

3.3.1.3 Ideal setting - Case A.

Numerical Experiments. The performance indicators values obtained by the three algorithms for the ideal case A are displayed in fig.3.2.

The PALM and BCD implementations perform similarly in this experiment. Besides, they are not influenced by the initialization: they reach the same precision for the different values of θ . That is not the case of PALS which is very sensitive to the initialization if Λ is not large. On the overall, the results are more precise if Λ is small (except for PALS if the initialization is far from the ground truth).

Influence of Λ for Case-A. With this perfect setting, the ℓ_1 penalty only helps to discriminate between the sources. In that case, using a very small value for Λ (but not null) would lead to the best results: more precise, and less biased sources. Indeed, the ℓ_1 norm, whose proximal operator is the soft-thresholding operator, biases the sources with a factor Λ . Hence if Λ is large,




Figure 3.2: Performances of BCD, PALM and PALS for varying initialization and regularization parameters values, Case A.

the sources estimates are much more biased, have stronger artifacts fig.3.3.

If the implementation is robust to the initialization such as PALM and BCD in this experiment, and that the sources are exactly sparse with noiseless observations, then setting Λ to small values would return the best results.

However, it appears for PALS, which is not robust to the initialization, that using large values of Λ makes the algorithm robust to the initialization fig.3.2g.

We remark that the largest entries, thanks to sparse modeling, are the most discriminant for the separation [14]. Indeed, with the morphological diversity principle 1.2.2, the largest entries of \mathbf{S} are likely to be significant for only one source. That is, the associated samples of \mathbf{X} provide




Figure 3.3: First samples of the recovered source \mathbf{S}_1 with PALM (we cannot distinguish the results from PALM and BCD), with $\Lambda = 1$ and $\Lambda = 100$ for the ideal case - \mathbf{A} .

already a good estimation of \mathbf{A} , and almost no unmixing is needed fig.3.4.

Besides, the residuals, due to the interferences between the sources while the mixing matrix is not correctly estimated, are thresholded with large values of Λ . On the other hand, with small values, the residues are kept and updated until the mixing matrix is correctly estimated. Therefore, the errors due to the mis-estimations of \mathbf{A} , are less propagated from one update to another if the thresholds are high.

That is why PALS is much more robust when using large values of Λ .




Figure 3.4: On the left: scatter plot of two of the initial observations, \mathbf{X}_1 versus \mathbf{X}_2 . On the right, scatter plot of two 'thresholded' observations $\tilde{\mathbf{X}} = \mathbf{A}\mathcal{S}_{1000}(\mathbf{S})$.

Differences PALM-BCD. One can notice that *BCD and PALM in this experiment perform very similarly in term of precision* (MA fig.3.2a, 3.2d and SDR fig.3.2b,3.2e). Whenever the problem is quite simple (few sources, small condition numbers of \mathbf{S} and \mathbf{A}), BCD and PALM perform generally quite similarly, and well. Indeed, in this experiment, *the two implementations are not influenced by the initialization*: the performances do not depend on the value of θ but only on Λ .

However, one can notice that *PALM is much faster than BCD* fig. 3.2f and 3.2c (when they both perform similarly). Indeed, PALM does not exactly minimize the subproblems, but simply updates the components with one proximal gradient step: it can converge more straightforwardly and more smoothly to a minimum fig.3.5.




Figure 3.5: Illustration of the minimization paths for BCD (blue), and PALM (red), with two components x_1 and x_2 . The cost function is represented by some of its level sets, the yellower, the lower.

If the problem is more challenging, *BCD and PALM can perform differently* because the problems are not convex and their minimization schemes are different. In particular, the components updates in PALM depend on all the current components: for example $\tilde{\mathbf{S}}^{(k)}$ depends on $\tilde{\mathbf{S}}^{(k-1)}$ and $\tilde{\mathbf{A}}^{(k-1)}$ whereas with BCD, $\tilde{\mathbf{S}}^{(k)}$ only depends on $\tilde{\mathbf{A}}^{(k-1)}$.

3.3.1.4 Case B - Noisy setting.

Numerical Experiment. In contrast with the first ideal example, the observations are corrupted by an additional Gaussian noise \mathbf{N} . The performances of the PALM (BCD performs similarly except for the CT.) and the PALS implementations with varying parameters are displayed in fig.3.6.

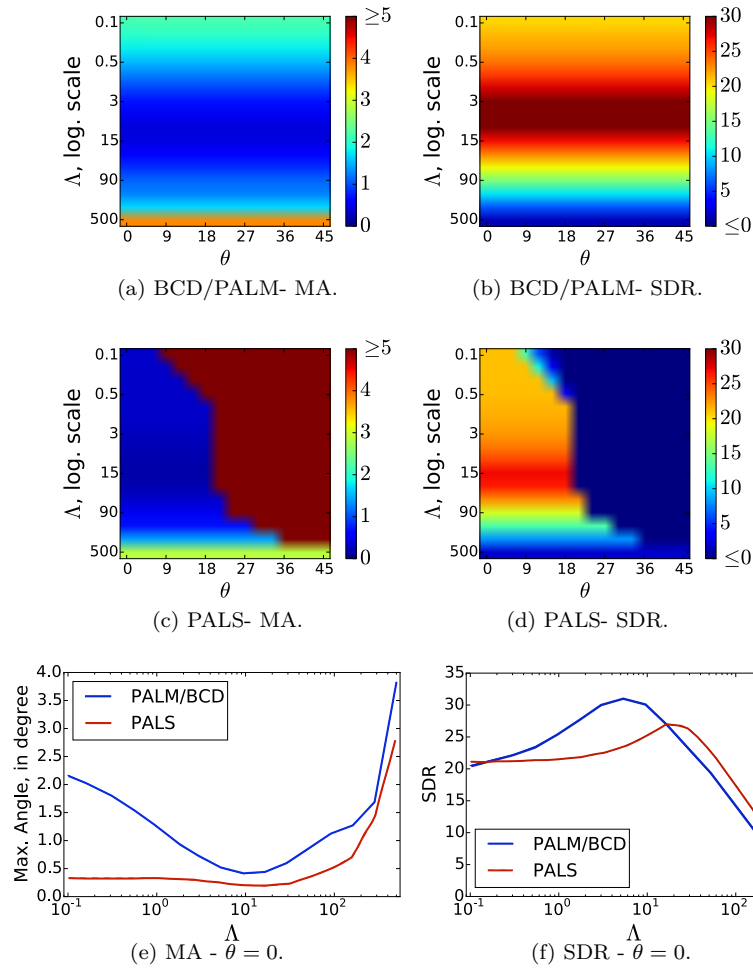


Figure 3.6: Performances of BCD/PALM and PALS for varying initializations and regularization parameters values, Case B.

Similarly to Case A, PALM/BCD and PALS do not perform similarly: PALM/BCD are robust to the initialization fig.3.6a,3.6b whereas PALS is not, except with large values of Λ fig.3.6c,3.6d . In contrast with Case A, the performances of the implementations do not evolve monotonically with Λ : *they reach a maximum of precision for a significant value of Λ* fig.3.6f,3.6e .

Role of Λ in the noisy setting. If the observations are noisy, the ℓ_1 regularization does not only help to discriminate between the sources, but also to denoise the recovered signals. Indeed, we recall that when fixing \mathbf{A} , the subproblem associated with the update of \mathbf{S} is the

following:

$$\operatorname{argmin}_{\mathbf{S}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_2^2 + \|\Lambda \odot \mathbf{S}\|_1.$$

When the sources are exactly and jointly sparse and the observations are noiseless, such as in Case A, setting Λ to a small value, amount to solve a least-square problem, and reaches the best precision. However, if the data are noisy, then it is necessary to denoise the sources and to set accordingly the values of Λ : too small, the sources are still noisy, too large, the sources are unnecessary biased fig.3.6f. Appropriate values of Λ should properly handle the trade-off between the sparsity constraint (numerous source entries set to zero, effective denoising) and the data fidelity term (good fit of the data, less biased but corrupted recovered sources).

In BSS, the mixing matrix should also be learned, and similarly, the most precise \mathbf{A} is recovered when the sources are correctly denoised fig.3.6e.

Differences between PALM/BCD and PALS. We notice that the values of Λ reaching the maximum of precision are different for PALS and PALM/BCD: in contrast with PALM/BCD, PALS does not exactly the cost function (3.11), and is not affected similarly by the noise.

Besides, PALS is based on pseudo-inversion, which makes the algorithm very sensitive to the presence of noise, especially if the mixing matrix has a large condition number. For instance, in fig.3.6f, we can notice that the best SDR value is around 27dB for PALS whereas it is around 32dB for PALM/BCD.

In the presence of noise, using PALM/BCD which do not involve pseudo-inversion lead to a better precision, if one is able to determine a correct value of Λ and if PALM/BCD are robust to the initialization for the given setting.

3.3.1.5 Case C - Approximately sparse sources.

Numerical Experiments. In contrast with the first ideal example A, the sources are approximately sparse. This setting is more realistic. The performances of PALM/BCD (which also perform similarly in this experiment) and the PALS implementations with varying parameters Λ and θ are displayed in fig.3.7.

In contrast with Case A and B, the parameters Λ do not influence similarly PALM/BCD and PALS, and PALS performs better than PALM/BCD (if θ is not too large).

Unmixing with approximately sparse sources - BCD/PALM. The unmixing of approximately sparse sources is more challenging than with perfectly sparse sources, with disjoint supports. From an unmixing point of view, noisy observations or observations of approximately




Figure 3.7: Performances of BCD/PALM and PALS for varying initialization and regularization parameters values, Case C.

sparse sources share the same difficulty: the largest samples of \mathbf{X} are not exactly in the direction of the columns of \mathbf{A} and it is necessary to unmix, even the large entries, to recover \mathbf{A} . By increasing Λ , the influence of the small entries becomes negligible and the unmixing is improved fig.3.7a.

However, the best estimation of the sources (SDR) is obtained with the smallest values of Λ . Indeed, since there is no noise, a small value of Λ leads to less bias fig.3.7b. If \mathbf{A} were exactly known, estimating \mathbf{S} with a least squares would return the best solution.

Since \mathbf{A} and \mathbf{S} are linked, especially when solving exactly (3.11) with PALM/BCD, it becomes challenging a find a trade-off for the value of Λ : a large value leads to a good estimation of \mathbf{A} , whereas the source estimates are more precise with small values.

Unmixing with approximately sparse sources - PALS. In contrast with PALM/BCD, the influence of Λ for PALS is similar to Case A: the smallest Λ is, the more precise (and less robust to the initialization) the results are fig.3.7c, 3.7d.

This can be explained by the fact that, for PALS, the minimization of the data-fidelity term and the ℓ_1 are dissociated (least-squares and then soft-thresholding). By estimating \mathbf{S} first with a least-squares estimation, the sources are 'virtually' estimated with $\Lambda = 0$, and then thresholded with the correct value of Λ . That is why, PALS is much less sensitive to the parameters setting (regardless of the initialization). We point out that, even if the results appear to be closer to the

ground-truths with PALS, the values of the cost-function (3.11) are smaller with PALM/BCD than with PALS. In fig. 3.8a, the objective function has a smaller value with the components returned by PALS/BCD (the values obtained with BCD differ from less than $1e^{-9}\%$ with the ones of PALM in this example) than with the components returned by PALS or the ground truth, even if the SDR obtained by PALS is clearly larger while θ is small (but greatly increases when θ is large). The minimization is, strictly speaking, better handled with PALM/BCD.




Figure 3.8: Variations of the SDR and cost function values for PALS and PALM and $\Lambda = 7$

3.3.2 Influence of the parameters Λ : summary

The regularization parameters Λ play a key role in sparse BSS since their values affect all the properties that an efficient algorithm should meet:

- **Robustness:** in order to be robust to Gaussian noise, the values of Λ should be large enough to denoise the sources.
- **Accuracy:** In the noiseless setting, small values of Λ provide precise estimates of the sources. However, in the presence of noise, the best tuple is obtained with significant values of Λ .
- **Reliability:** Only the large values of Λ make the implementations insensitive to the initialization.

These properties are conflicting, and are difficult to meet altogether if using constant values of Λ . Besides, finding a-priori the values of Λ is not possible without further information.

That is why, instead of considering a constant Λ , the authors in [14] rather propose an adaptive *path* for Λ for PALS. The so-called *coarse to fine strategy* for PALS starts from a large value of Λ , which is progressively decreases towards a final value. The unmixing is fastened and not less accurate. This is reminiscent to the fixed point continuation strategy proposed in [69] within the compressive sensing framework with ℓ_1 regularization.

3.3.3 Adaptive threshold strategy

We propose first to present more precisely this adaptive threshold strategy, for PALS and its extension for PALM and BCD employing proximal calculus and introduced in [112]. Then, we will see how this strategy alleviates the main difficulties arising with fixed thresholds, which have been discussed previously.

3.3.3.1 Updates of the regularization parameters in practice. We then present how, precisely, we have implemented the threshold updates in the different algorithms. We recall that the thresholds should meet the two following properties:

- i) Denoise the sources efficiently, *i.e.* the final thresholding values should be set according to the noise level contaminating the sources.
- ii) Start at large values in order to threshold the non-discriminant entries and only keep the largest samples to determine the orientation of \mathbf{A} during the first iterations.

Parameter updates in PALS. We start by presenting how the final threshold values are set. As stated above, they are chosen in order to denoise the sources efficiently. Let us assume that the mixing matrix \mathbf{A} is recovered. Each estimated source $\tilde{\mathbf{S}}_i$ before the thresholding, is given by $\tilde{\mathbf{S}}_i = \mathbf{S}_i + (\mathbf{A}^\dagger \mathbf{N})_i$, where $(\mathbf{A}^\dagger \mathbf{N})_i$ is the Gaussian noise corrupting the i th source.

The denoising of a sparse signal, here \mathbf{S}_i , from a Gaussian corruption $(\mathbf{A}^\dagger \mathbf{N})_i$ is a well understood problem which can be seen as a detection procedure, see App.B. It amounts to set, for each source i , a threshold value Λ_i which will put to zero the entries corresponding to the noise, and keep the entries larger than Λ_i corresponding to the support of the i th source. This hypothesis testing problem is generally handled with the “ k - σ ” rule: the probability that an entry of $\tilde{\mathbf{S}}_i$ with an amplitude larger than $3\sigma_i$, where σ_i is the standard deviation of the Gaussian noise $(\mathbf{A}^\dagger \mathbf{N})_i$ corrupting the i th source, corresponds to the Gaussian noise indeed is only of 0.4%. The value of k , generally taken between 1 and 3 depends on the applications, [124].

The value of the standard deviation σ_i of the projected noise $(\mathbf{A}^\dagger \mathbf{N})_i$ is generally unknown, but can be approximated using the Median Absolute Deviation (mad) of the signal: $\sigma_i \approx 1.48 \times \text{mad}((\mathbf{A}^\dagger \mathbf{N})_i)$, where $\text{mad}(x) = \text{median}_i(|x_i - \text{median}_i(x_i)|)$. Besides, given that the mad operator is robust to sparse contamination, we have $\text{mad}((\mathbf{A}^\dagger \mathbf{N})_i) = \text{mad}((\mathbf{A}^\dagger \mathbf{N})_i + \mathbf{S}_i) = \text{mad}(\tilde{\mathbf{S}}_i)$. In other words, we can estimate the Gaussian noise level directly from the estimated sources $\tilde{\mathbf{S}}_i$ (before the thresholding) and set the final thresholds $\Lambda_i = k_S \text{mad}(\tilde{\mathbf{S}}_i)$, with $k_S = 1.48 \times k$. Details on this procedure can be found in App.B.

Then, we discuss the decreasing strategy for the thresholds. During the unmixing process, the current estimated sources may vary strongly from one iteration to another while the mixing matrix is not fairly estimated. That is why, the threshold parameters are set directly according to the estimated sources $\tilde{\mathbf{S}} = \tilde{\mathbf{A}}^\dagger \mathbf{X}$ in order to fully take into account their evolution, eq.3.15. We

recall that we aim, during the first iterations, to only consider the largest and most discriminant entries for the separation. For this purpose, we select an increasing percentage of the source entries larger than the final threshold values $k_S \text{mad}(\tilde{\mathbf{S}})$, eq.3.16. More precisely, at the k th iteration given a total number of K loops, the estimation of Λ_i is twofold: first, we only keep the entries of $\tilde{\mathbf{S}}_i$ larger than the final threshold $k_S \text{mad}(\tilde{\mathbf{S}}_i)$; and second, the threshold value is set to the $100 \times \frac{K-k}{K}$ percentile of the kept entries $|\tilde{\mathbf{S}}_i|_{|\tilde{\mathbf{S}}_i| > k_S \text{mad}(\tilde{\mathbf{S}}_i)}$. This ensures that an increasing number of entries are considered, and that the final threshold will indeed be $k_S \text{mad}(\tilde{\mathbf{S}})$. In summary, we propose the following source update permitting a precise control of their sparsity level, at the k th iteration given a total number of K loops:

$$\text{Compute the least-squares estimates: } \tilde{\mathbf{S}} = \tilde{\mathbf{A}}^{(k-1)\dagger} \mathbf{X} \quad (3.15)$$

$$\text{Update the thresholds : } \Lambda_j = \text{pct}(|\tilde{\mathbf{S}}_j|_{|\tilde{\mathbf{S}}_j| > k_S \text{mad}(\tilde{\mathbf{S}}_j)}, 100 \times \frac{K-k}{K}), \forall j = 1, \dots, n \quad (3.16)$$

$$\text{Soft-thresholding of the sources : } \tilde{\mathbf{S}}^{(k)} = \mathcal{S}_\Lambda(\tilde{\mathbf{S}}) \quad (3.17)$$

where k_S is a numerical constant usually between 1 and 5 [124] and $\text{pct}(\mathbf{Y}, k)$ denotes the k th percentile of the set composed of the entries of \mathbf{Y} .

The number of outer loops is generally fixed to few hundreds (300 for the following numerical experiments). This implementation corresponds to the GMCA algorithm, [14].

Parameter updates in BCD and PALM. The parameter update is quite different with BCD based on proximal calculus since we do not have access to the projected sources. In [112], [113], the authors propose to update the regularization parameters before each update of the sources at the outer loop level, based on the current estimated projected noise level.

The update of the thresholds still follows the same strategy: i) final threshold values set according to the noise level, and ii) large threshold values during the first iterations.

Let us assume that the sources and mixing matrix are recovered. At every proximal gradient step, the update of the sources is two-steps: a gradient descent step, and then the soft-thresholding. This situation is similar to the two-steps update of the sources with PALS (except that the process is iterative with BCD): we have a first noisy estimate of the sources $\tilde{\mathbf{S}} = \mathbf{S} + \frac{1}{L_S} \mathbf{A}^T (\mathbf{X} - \mathbf{A}\mathbf{S})$, where L_S is the gradient step, which is thresholded in the second step $\mathcal{S}_{\frac{\Lambda}{L_S}} \left(\mathbf{S} + \frac{1}{L_S} \mathbf{A}^T (\mathbf{X} - \mathbf{A}\mathbf{S}) \right)$.

As discussed previously, we aim to recover a sparse signal \mathbf{S} from its noisy observations $\tilde{\mathbf{S}}$. The noise level corrupting the sources $\tilde{\mathbf{S}}$ in this case is given by $\frac{1}{L_S} \mathbf{A}^T (\mathbf{X} - \mathbf{A}\mathbf{S}) = \frac{1}{L_S} \mathbf{A}^T (\mathbf{N})$. In other words, the noise is projected on the source domain with $\mathbf{A}^T \mathbf{N}$ with PALM or BCD (proper minimization of the LASSO problem), whereas it is projected with $\mathbf{A}^\dagger \mathbf{N}$ with PALS. Thus, by keeping the “k- σ ” rule for setting the final value of the thresholds, we obtain: $\frac{\Lambda_i}{L_S} =$

$k_S \text{mad}(\frac{1}{L_S} \mathbf{A}^T(\mathbf{X} - \mathbf{AS}))$, where $k_S \in (1, 5)$. One can notice that it amounts to set Λ to $\text{mad}(\mathbf{A}^T(\mathbf{X} - \mathbf{AS}))$, the gradient of the quadratic term. More details can be found in App.B.

In order to only select the largest entries of the sources at the beginning of the algorithm, the authors of [112], [113] propose to further select a large value of k_S , which usually belongs to $(1, 5)$ for the BCD implementation. The parameters are updated before each estimation of \mathbf{S} (before the inner loop with ISTA or FISTA):

$$\text{Compute the gradient: } \mathbf{G} = \tilde{\mathbf{A}}^{(k-1)T}(\mathbf{X} - \tilde{\mathbf{A}}^{(k-1)}\tilde{\mathbf{S}}^{(k-1)})$$

$$\text{Update the thresholds : } \Lambda_j = k_S \text{mad}(\mathbf{G}_j), \forall j = 1, \dots, n$$

where k_S decreases linearly from $\min_i(\frac{\|\mathbf{G}^0\|_\infty}{\text{mad}(\mathbf{G}_i)})$ (with $\mathbf{G}^{(0)}$ the gradient at the first outer loop) to the final value in $(1, 5)$ during 80% of the iterations for BCD [112], [113]. The parameters are then fixed for the convergence. This ensures that only the largest entries are considered during the first iterations.

The dynamics of PALM are not similar to BCD or PALS, for which a complete 'inversion' of the components is undertaken for every update. With PALM, only a proximal gradient step is done at every iteration. If PALM is initialized closely to the sought-after solutions, using a $k - \text{mad}$ setting of the parameters (similarly to BCD but with fixed k) works well in practice. However, this strategy is not reliable for the PALM implementation in general.

We underline that these strategies may not be satisfactory for BCD/PALM. Indeed, they are much more sensitive to the parameters setting than PALS. The latter, broadly speaking, only needs to start with large parameters (robustness with respect to the initialization) and then to decrease towards the approximated noise corrupting the sources (robustness against Gaussian noise). On the other hand, the final values of the regularization parameters greatly influence the results returned by BCD and PALM, but the proposed strategy does not yield a precise control of the thresholds.

3.3.3.2 Numerical experiments. With the adaptive thresholds, we will see that GMCA becomes very interesting: precise, reliable and very fast. We present the results obtained by GMCA for the three previous settings ⁵.

We remark in fig.3.9, that the threshold strategy makes the algorithm not influenced by the initialization. Besides, even if the threshold 'paths' are not exactly identical (the initializations are different and the parameters are based on the current values of the estimates), the algorithm

⁵In the framework of this thesis, there is no interest of using BCD because all the regularizations that we will use, have an explicit proximal operator: GMCA can be employed and is much faster than BCD.




Figure 3.9: Performances of GMCA for the 3 data-sets, and varying initializations.




Figure 3.10: Mean value of Λ for Case B during GMCA, and the 16 different initializations.

GMCA returns a very similar solution, for every θ : the parameters Λ stabilize towards similar values fig.3.10.

The mixing matrix \mathbf{A} is on the overall correctly estimated, but the SDR for Case B and Case C are quite low: this strategy permits a correct unmixing, but the estimation of the sources can be improved.

3.3.4 Benefits and drawbacks of the adaptive threshold strategy

Let us start with the benefits of the proposed strategy.

First, the thresholds start with large values. Indeed, while the sources are not correctly un-

mixed, they are not jointly sparse. As a consequence, the estimated noise levels are influenced by the interferences and larger: for example with PALS $\text{mad}(\tilde{\mathbf{A}}^\dagger \mathbf{X}) = \text{mad}(\tilde{\mathbf{A}}^\dagger \tilde{\mathbf{A}} \tilde{\mathbf{S}} + \tilde{\mathbf{A}}^\dagger \mathbf{N}) \geq \text{mad}(\tilde{\mathbf{A}}^\dagger \mathbf{N})$. Hence, the thresholds should be adapted to the level of interferences: a large residue leads to large thresholds, and only the largest and most discriminative samples are kept.

On the other hand, whenever the mixing matrix is correctly estimated, the residue level decreases (no interferences between the sources): the source estimates are refined. If the sources are exactly sparse, or the noise level is significant, and if \mathbf{A} has been fairly retrieved, the estimated noise level is close to the one corrupting the sources.

For these reasons, the proposed strategy is interesting (for PALS): it provides robustness against Gaussian noise and the initialization. However, it also suffers from some drawbacks.

The projected sources with PALS correspond to $\mathbf{A}^\dagger \mathbf{X} = \mathbf{S} + \mathbf{A}^\dagger \mathbf{N}$: the noise is projected, with the pseudo-inverse, on the source subspace. This projection tends to amplify the noise in some directions if \mathbf{A} has a large condition number. A fortiori, the corresponding Λ , computed from the mad of this projection, is also larger for the sources which have been more corrupted (and these sources are more biased). Besides, if the sources are approximately sparse, then their mad is not null: the global mad of the projected sources is larger than the noise level $\text{mad}(\mathbf{A}^\dagger \mathbf{X}) > \text{mad}(\mathbf{A}^\dagger \mathbf{N})$. Last, PALS is on the overall less precise than BCD and PALM since it does not exactly minimize the subproblems.

The adaptive thresholds strategy of PALS is very efficient, and makes it, most of the times, reliable, robust and fast, but does not yield precision.

3.3.5 Enhancing the results of GMCA

The proposed implementation of PALS is most of the times, very efficient to unmix the data. However, it suffers from a lack of precision, due to an over-estimated noise level inducing bias, and the projections which do not rigorously solve the subproblems.

3.3.5.1 Combining the strategies. We propose to combine PALS and PALM to increase the precision and obtain a strategy that is provably convergent:

- Warm-up: the components are initialized with PALS. It provides a fair estimation of the components, is robust with respect to the initialization and Gaussian noise thanks to the adaptive thresholding strategy, and fast. However, it can be quite imprecise if the components have a large condition number. Besides, it is not provably convergent.
- Refinement: the components are refined with PALM, which is faster than BCD. In this step, the thresholds are fixed, and based on the components returned by PALS. It is provably

convergent and more precise because the cost function is rigorously minimized.

We would like to highlight that the need for appropriate heuristics to build reliable matrix factorization procedures has also been pointed out in the framework of NMF in [64].

3.3.5.2 Bias and large dynamics between the sources. The ℓ_1 norm is very efficient in sparse BSS [115, C.2.3]. The main counterpart of the ℓ_1 is that it introduces artifacts due to the source shrinkage.

We propose a reweighting ℓ_1 strategy [26] during the refinement procedure to reduce the bias. It consists in modifying the thresholds given the current values of the components. An entry with a large value is likely to be correctly estimated, and relatively less influenced by the noise: it can be thresholded with a smaller value. On the other hand, an entry whose amplitude is similar to the noise level/threshold does not certainly belong to the signals to be retrieved. As so, its threshold value should remain stable and large. More precisely, we propose to define a weighting matrix $\mathbf{W}_S \in \mathbb{R}^{n \times t}$ such as:

$$\mathbf{W}_S = \frac{\Lambda}{\Lambda + |\tilde{\mathbf{S}}|} \quad (3.18)$$

The sources, sparse in the direct domain, will be then thresholded by $\mathbf{W}_S \odot \Lambda$. That is, a large entry is less thresholded than a small entry. The weights are generally updated few times. We underline that the weighting depends on the data: if the data are approximately sparse, a stronger weighting can help.

In a certain extend, the reweighting ℓ_1 strategy can be seen as a way to manage the variabilities intra-sources. The dynamics inter-sources are automatically handled by the proposed threshold strategy since one regularization parameter is chosen per source.

3.3.5.3 Latest version of GMCA for sparse BSS. The final form of the proposed sparse matrix factorization algorithm (in the direct domain) is presented in Alg.13.

The numerical results obtained with this algorithm are presented in tab.3.2, for $\theta = 45^\circ$.

Algorithm 13 Enhanced GMCA

```

1: procedure GMCA( $\mathbf{X}, \tilde{\mathbf{A}}^{(0)}$ )
2:   for  $k=1, \dots, K$  do ▷ WARM-UP
3:     if  $\text{cond}(\tilde{\mathbf{A}}^{(k-1)}) < 10000$  then ▷ Pseudo-inverse if small condition number
4:        $\tilde{\mathbf{S}}^{(k)} \leftarrow \tilde{\mathbf{A}}^{(k-1)\dagger} \mathbf{X}$ 
5:     else: ▷ Truncated inversion otherwise
6:        $L = \|\tilde{\mathbf{A}}^{(k-1)T} \tilde{\mathbf{A}}^{(k-1)}\|_{2,s}$ 
7:       for  $j=1, \dots, 100$  do:
8:          $\mathbf{S}^{(k)} \leftarrow \mathbf{S}^{(k)} + \frac{1}{L} \tilde{\mathbf{A}}^{(k-1)T} (\mathbf{X} - \tilde{\mathbf{A}}^{(k-1)} \mathbf{S}^{(k)})$ 
9:       Compute  $\Lambda_j = \text{pct}(|\tilde{\mathbf{S}}_j|_{|\tilde{\mathbf{S}}_j| > k_S \text{mad}(\tilde{\mathbf{S}}_j)}, 100 \times \frac{K-k}{K}), \forall j = 1, \dots, n$ 
10:       $\tilde{\mathbf{S}}^{(k)} \leftarrow \mathcal{S}_\Lambda(\tilde{\mathbf{S}}^{(k)})$  ▷ Soft-thresholding
11:      if  $\text{cond}(\tilde{\mathbf{S}}^{(k)}) < 100$  then ▷ Pseudo-inverse if small condition number
12:         $\tilde{\mathbf{A}}^{(k)} \leftarrow \mathbf{X} \tilde{\mathbf{S}}^{(k)\dagger}$ 
13:      else: ▷ Truncated inversion otherwise
14:         $L = \|\tilde{\mathbf{S}}^{(k-1)} \tilde{\mathbf{S}}^{(k-1)T}\|_{2,s}$ 
15:        for  $j=1, \dots, 100$  do:
16:           $\tilde{\mathbf{A}}^{(k)} \leftarrow \tilde{\mathbf{A}}^{(k)} + \frac{1}{L} (\mathbf{X} - \tilde{\mathbf{A}}^{(k)} \mathbf{S}^{(k)}) \tilde{\mathbf{S}}^{(k)T}$ 
17:           $\tilde{\mathbf{A}}^{(k)i} \leftarrow \frac{(\tilde{\mathbf{A}}^{(k)i})^i}{\|\tilde{\mathbf{A}}^{(k)i}\|_2}, \forall i = 1, \dots, n$  ▷ Projection
18:       $\tilde{\mathbf{S}}^{(1)} \leftarrow \tilde{\mathbf{S}}^{(K)}, \tilde{\mathbf{A}}^{(1)} \leftarrow \tilde{\mathbf{A}}^{(K)}$ 
19:      for  $j=1, \dots, J$  do ▷ REFINEMENT
20:        Compute the residue:  $\mathbf{G} = \tilde{\mathbf{A}}^{(j)T} (\mathbf{X} - \tilde{\mathbf{A}}^{(j)} \tilde{\mathbf{S}}^{(j)})$  ▷ Compute the thresholds
21:         $\Lambda_j = k_S \text{mad}(\mathbf{G}_j), \forall j = 1, \dots, n$  ▷ Thresholds
22:         $\Lambda \leftarrow \frac{\Lambda}{\Lambda + |\tilde{\mathbf{S}}^{(j)}|}$  ▷ Reweighting
23:         $\tilde{\mathbf{S}}^{(0)} \leftarrow \tilde{\mathbf{S}}^{(j)}, \tilde{\mathbf{A}}^{(0)} \leftarrow \tilde{\mathbf{A}}^{(j)}$ 
24:        while do not converge do ▷ PALM
25:          Compute  $L_S = \|\tilde{\mathbf{A}}^{(k-1)T} \tilde{\mathbf{A}}^{(k-1)}\|_{s,2}$ 
26:           $\tilde{\mathbf{S}}^{(k)} \leftarrow \mathcal{S}_{\frac{\Lambda}{L_S}} \left( \tilde{\mathbf{S}}^{(k-1)} + \frac{1}{L_S} \tilde{\mathbf{A}}^{(k-1)T} (\mathbf{X} - \tilde{\mathbf{A}}^{(k-1)} \tilde{\mathbf{S}}^{(k-1)}) \right)$  ▷ Proximal gradient step
for  $\mathbf{S}$ 
27:      Compute  $L_A = \|\tilde{\mathbf{S}}^{(k)} \tilde{\mathbf{S}}^{(k)T}\|_{s,2}$ 
28:       $\tilde{\mathbf{A}}^{(k)} \leftarrow \tilde{\mathbf{A}}^{(k-1)} + \frac{1}{L_A} (\mathbf{X} - \tilde{\mathbf{A}}^{(k-1)} \tilde{\mathbf{S}}^{(k)}) \tilde{\mathbf{S}}^{(k)T}$  ▷ Gradient descent for  $\mathbf{A}$ 
29:       $(\tilde{\mathbf{A}}^{(k)})^i \leftarrow \frac{(\tilde{\mathbf{A}}^{(k)})^i}{\max(1, \|(\tilde{\mathbf{A}}^{(0)})^i\|_2)} \forall i = 1, \dots, n$  ▷ Proximal Operator for  $\mathbf{A}$ 
30:       $k \leftarrow k + 1$ 
31:       $\tilde{\mathbf{S}}^{(j+1)} \leftarrow \tilde{\mathbf{S}}^{(k)}, \tilde{\mathbf{A}}^{(j+1)} \leftarrow \tilde{\mathbf{A}}^{(k)}$ 
return  $\tilde{\mathbf{A}}^{(J+1)}, \tilde{\mathbf{S}}^{(J+1)}$ .
```

Case	GMCA			GMCA+ refinement		
	MA	SDR	CT	MA	SDR	CT
A- Ideal	0.04	60	2.4	0.04	60	2.6
B- Noisy	0.3	23.5	2.4	0.1	37.7	2.6
C- Approx. Sparse	3.1	8.4	2.7	3.9	22.4	29

Table 3.2: Performance indicators for GMCA and GMCA+refinement for the 3 illustrative examples.

With the combination, the computational cost is slightly increased (significantly for case - C, which requires more updates of the weights, or a stronger weighting scheme). The sources estimates are much more precise: we have a gain of 14dB for the cases - B and C, tab. 3.2 by adding the refinement step.

3.4 Summary and extensions

In this chapter, we have presented the latest minimization scheme employed for solving standard sparse BSS algorithms. It consists in two phases: first, a fair estimation of the components and the associated regularization parameters are retrieved with a PALS implementation and the adaptive threshold strategy. The results are finally refined with a PALM implementation of the minimization problem.

This strategy meets the following criteria:

- Robustness: Robustness to Gaussian noise is provided by the adaptive threshold strategy.
- Accuracy: The refinement step, based on reweighted ℓ_1 and the PALM implementation yield precise results.
- Reliability: The *coarse-to-fine* threshold strategy [14] during the warm-up aims to return a robust estimate of the mixing matrix.
- Convergence: The decreasing threshold strategy makes PALS empirically stable. The refinement with PALM is provably convergent.
- Easy-to-use: The thresholds are automatically set according to the noise level corrupting the data during the warm-up.

The proposed heuristics (PALS with adaptive thresholding strategy, followed by PALM with reweighting ℓ_1) are still efficient in more complex settings.

Sparsity in a transformed domain. In many applications, such as in imaging, the sources are sparsely represented in a dictionary $\Phi_{\mathbf{S}}$. In that case, the authors of [14] propose rather to estimate directly \mathbf{A} and the expansion coefficients of \mathbf{S} in $\Phi_{\mathbf{S}}$ by solving the sparse BSS problem in $\Phi_{\mathbf{S}}$ directly with PALS:

$$\underset{\mathbf{A}, \alpha_{\mathbf{S}}}{\text{minimize}} \frac{1}{2} \|\mathbf{X}\Phi_{\mathbf{S}}^T - \mathbf{A}\alpha_{\mathbf{S}}\|_2 + \|\Lambda \odot \alpha_{\mathbf{S}}\|_1.$$

This way, we obtain a fast approximation of $\tilde{\mathbf{A}}, \tilde{\mathbf{S}}$, where $\tilde{\mathbf{S}} = \alpha_{\tilde{\mathbf{S}}}\Phi_{\mathbf{S}}$. Then, similarly to the sparse BSS problem in the direct domain, we can use PALM with the reweighted ℓ_1 . The thresholds are based on the residual expressed in $\Phi_{\mathbf{S}}$, and the weighting on $\tilde{\mathbf{S}}\Phi_{\mathbf{S}}^T$, see for instance Chapter 5.

In other words, the proposed heuristics can be quite straightforwardly transposed if \mathbf{S} is sparsely represented in a transformed domain.

Handling composite priors. Whenever the priors on the components cannot be easily handled with the projections (for example non-negativity and sparsity in a transformed domain [113]), PALS cannot be used.

It is possible to use the BCD implementation with the associated threshold strategy [113]. Nonetheless, this approach is time consuming.

Another approach would be to develop an appropriate threshold strategy for PALM, since it is generally faster than BCD (especially if the proximal operator are explicit). In contrast with BCD, we have access to the projected sources at each update (after the gradient descent, and before the proximal operator), similarly to PALS. As so, we may replicate the fine control of the sparsity level of the sources that is powerful for PALS. However, and contrary to PALS, the inversion of the mixing matrix requires several outer loops for PALM, what might bring further difficulties.

Providing an efficient implementation of PALM, with an automatic setting of the regularization parameters would be very interesting. It would provide a very generic algorithm, able to handle complex priors.

Robust BSS. We will see in the following chapters that most of the heuristic strategies can be kept and transposed to the framework of robust BSS: automatic setting of the regularizations parameters, two-steps approach and projected least-squares.

Besides, the explicit estimation of the outliers makes possible the presence of leakages between the sources and the outliers, additionally to the interferences between the sources. Handling the interactions between the sources and the outliers, as well as between the sources themselves, requires the developments of additional heuristics that will be presented in the following chapters.

Chapter 4

Robust BSS in the presence of sparse outliers

Summary

4.1 Robust sparse BSS: recoverability	76
4.1.1 Data Modeling	76
4.1.2 Spectral differences between the source contributions and the outliers	77
4.1.3 Separation between the source contribution and the outliers	78
4.2 Robust GMCA	80
4.2.1 A naive extension of GMCA	80
4.2.2 The rGMCA algorithm	82
4.2.3 Evaluation of rGMCA	84
4.3 Robust AMCA	86
4.3.1 An analogy with partially correlated sources	86
4.3.2 Robust AMCA Algorithm	87
4.4 Numerical Experiments	96
4.4.1 Experimental protocol	96
4.4.2 Simulations for synthetic data	97
4.4.3 Application to NMR spectral unmixing	102
4.5 Summary and future works	104

Sparse BSS methods are sensitive to the presence of spurious large errors, which are unfortunately ubiquitous in some applications 2.1.1. In this chapter, we will focus on the sparse BSS problem in the presence of outliers and sources sharing the same morphology.

We will present two algorithms which explicitly estimate the mixing matrix, sources and outliers by exploiting sparse modeling of the data. Besides, they both further implement a weighting

scheme to penalize the corrupted samples during the unmixing yielding reliable algorithms. For both, the rationale of the weighting schemes, algorithms and numerical experiments are provided, and compared with state-of-the-art robust BSS methods. This chapter has been adapted from [33] and [31].

4.1 Robust sparse BSS: recoverability

4.1.1 Data Modeling

In this chapter, we focus on the presence of outliers and sources sparsely represented in the same dictionary. This situation is encountered in many applications, including in spectrometry, or in astrophysics where spurious point sources and the foreground emissions can be both sparsely represented in the wavelet domain 2.1.1.

Priors on the components. In this chapter, we will assume that the following priors on the components hold true:

- Sources: The sources are assumed to be sparsely represented in a dictionary $\Phi_{\mathbf{S}}$. For instance, in the LC/MS data analysis 2.1.1, the sources are quite sparse in the direct domain, but using 1D wavelets would result in sparser decompositions. The number of sources is supposed to be known.
- Mixing matrix: In order to avoid the scaling indeterminacy between \mathbf{A} and \mathbf{S} , the columns of \mathbf{A} will be normalized. We make no further assumption on \mathbf{A} . The following methods and numerical experiments are influenced by some characteristics of \mathbf{A} , especially its condition number. This will be discussed in the following.
- Outliers: Each row of \mathbf{O} is sparsely represented in $\Phi_{\mathbf{S}}$, similarly to \mathbf{S} : $\mathbf{O} = \alpha_{\mathbf{O}}\Phi_{\mathbf{S}}$ and $\alpha_{\mathbf{O}}$ is composed of few significant entries. We will also assume that the outliers do not significantly cluster in a given direction. For this purpose, we will assume in the beginning of this chapter that some of the non-zero entries of $\alpha_{\mathbf{O}}$ are independently distributed, and some of the columns of $\alpha_{\mathbf{O}}$ are entirely active. We will detail this assumption in the following paragraph.

For the sake of clarity, we will assume that $\Phi_{\mathbf{S}} = \mathbf{I}_{t \times t}$. In other words, the outliers and sources are sparse in the domain of observations. Similarly to [14] and [13], all the results could be generalized to any orthogonal transform $\Phi_{\mathbf{S}}$ and provide a good approximation for redundant sparse representations such as tight frames, which have diagonally dominant Gram matrices (*e.g.* undecimated wavelet transforms, curvelets [124], etc.).

4.1.2 Spectral differences between the source contributions and the outliers

As stated above, we assume that the outliers do not significantly cluster in any direction, Section 2.1.2. The rationale of the proposed robust BSS method is based on the difference of spectral properties (column-space) between \mathbf{AS} and \mathbf{O} :

- The source contribution \mathbf{AS} is clustered along the directions given by the mixing matrix. We remark that this clustering aspect is a consequence of the sparse modeling of the sources, and the morphological diversity principle 1.2.2.
- In contrast, the corrupted samples $\{(\mathbf{O})^k\}_{k=1..t}$ do not cluster in any specific direction fig.4.1a. We will consequently not consider the presence of numerous and independently corrupted entries (more generally, row or row and column sparse outliers), which are clustered along the canonical axes, fig.4.1b. Besides, we remark that independently distributed outliers (column and row sparse) are not frequent in multi/hyperspectral imaging: many anomalies are visible at several wavelengths and thus column-sparse.

We also point out that on the other hand, if the outliers are indeed column and row sparse, then the problem is easier. Besides, the exact separation between the outliers and sources is possible if the columns of \mathbf{A} contain several active entries. In such cases, the separation between the components can be achieved using the problem eq.4.2.

Broadly speaking, robust BSS methods aim to find the n directions which aggregate the most of data samples. If the outliers are also clustered in a specific given direction, it becomes impossible to distinguish between a source contribution and the outliers. Besides, if a corrupted sample \mathbf{X}^k has a direction similar to one of the columns of \mathbf{A} , it will necessarily be estimated as belonging to the source contribution (the elements sharing a same spectral signature belong to a same class, correspond to a same emission).

Thus, if the outliers and sources share the same morphology, the *outliers should be spectrally well spread out in contrast with the source contribution which clusters along the directions given by \mathbf{A}* , in order to recover \mathbf{A} .

We underline that most of the state-of-the-art methods of robust BSS rely on a stronger spectral diversity assumption 2.2.1: they require the outliers to not lie in the subspace spanned by \mathbf{A} and thus cannot handle the determined case $m = n$, if \mathbf{A} is full-rank.




Figure 4.1: Three sparse sources are mixed into 3 observations (sources contribution in blue), corrupted by 0.2% of outliers (red stars). In (a), the outliers are column sparse whereas they are row and column sparse in (b), the amplitudes of the active entries are identical for (a) and (b).

4.1.3 Separation between the source contribution and the outliers

If the outliers do not cluster too strongly, with respect to the source contribution, it is possible to retrieve \mathbf{A} . However, separating \mathbf{AS} from \mathbf{O} is a challenging problem, even if \mathbf{A} is perfectly known.

The separation between the outliers and the sources can be solved pixel-wise, in contrast with the unmixing. Let us recast the data model in the compressive sensing framework [25]:

$$\begin{aligned} \mathbf{X}^k &= \begin{bmatrix} \mathbf{A} & \mathbf{I}_{m \times m} \end{bmatrix} \begin{bmatrix} \mathbf{S}^k \\ \mathbf{O}^k \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{A} & \mathbf{I}_{m \times m} \end{bmatrix} \mathbf{y}^k, \quad \forall k = 1, \dots, t, \end{aligned}$$

where $\begin{bmatrix} \mathbf{S}^k \\ \mathbf{O}^k \end{bmatrix} = \mathbf{y}^k$ is the sought-after vector.

Several characteristics impact this problem: the sensing matrix $\begin{bmatrix} \mathbf{A} & \mathbf{I}_{m \times m} \end{bmatrix}$ (in brief, the correlation between the columns of the sensing matrix), and the sparsity pattern of the signal \mathbf{y}^k to be recovered (we can assume that \mathbf{S} is highly sparse thanks to the MDP, Section 1.2.2, but the entries of \mathbf{O}^k can all be active). Depending on these properties, the previous problem can be studied with different approaches:

- Classical compressive sensing: This particular structure has been studied in [85] and [84] (in a more general framework where the sensing matrix is composed of 2 particular submatrices). The vector \mathbf{y}^k can be recovered using a standard iterative algorithm (e.g. IHT [11]) as long as its number of active entries is small (very small if $n = m$). Another approach proposed in [24] also needs the vector \mathbf{y}^k to be highly sparse. This mandates the outliers to be row-sparse.

- 'Cross and bouquet' model [133]: In [133], the authors show that the vector \mathbf{y}^k can be recovered in the presence of a large level of corruption (more than 50%) if the columns of \mathbf{A} are highly correlated. For given dimensions (m, n) , if the columns of \mathbf{A} are highly correlated, they “span” only a small area of \mathbb{R}^m (together with the sparsity constraint and the MDP, $\mathbf{A}\mathbf{S}$ is close to the cone generated by \mathbf{A}): this can be seen as a “relaxation” of the rPCA framework, Section 2.2.1. However, a highly correlated mixing matrix \mathbf{A} would clearly hamper the unmixing. Besides, completely corrupted columns of \mathbf{X} are still not recoverable.
- Outliers with known spectra: Let us suppose that the spectra of the outliers are known. More precisely, let $\mathbf{O}_D \in \mathbb{R}^{m \times T}$ be composed of these normalized spectra, where T designates the number of corrupted columns, such that $\mathbf{O}_D \boldsymbol{\alpha}_O = \mathbf{O}$ and $\boldsymbol{\alpha}_O \in \mathbb{R}^{T \times t}$ is column and row sparse. If the outliers are uniformly distributed, we end up with the previous data model and $\mathbf{O}_D = \mathbf{I}$.

The data model can be recast as:

$$\mathbf{X}^k = \begin{bmatrix} \mathbf{A} & \mathbf{O}_D \end{bmatrix} \begin{bmatrix} \mathbf{S}^k \\ \boldsymbol{\alpha}_O^k \end{bmatrix}.$$

Assuming that the vector $\begin{bmatrix} \mathbf{S}^k \\ \boldsymbol{\alpha}_O^k \end{bmatrix}$ is $K = k_s + k_o$ -sparse in $\mathbf{D} = \begin{bmatrix} \mathbf{A} & \mathbf{O}_D \end{bmatrix}$ with $k_s = \|\mathbf{S}^k\|_0$ and $k_o = \|\boldsymbol{\alpha}_O^k\|_0$, a sufficient condition for the identifiability of $\begin{bmatrix} \mathbf{S}^k \\ \boldsymbol{\alpha}_O^k \end{bmatrix}$ is given by¹ [47]:

$$K < \frac{1}{2} \left(1 + \frac{1}{\mu_{\mathbf{D}}} \right),$$

where $\mu_{\mathbf{D}}$ designates the so-called mutual coherence of the dictionary \mathbf{D} . The mutual coherence of \mathbf{D} is defined as $\mu_{\mathbf{D}} = \max_{i,j} |\langle \mathbf{d}_i, \mathbf{d}_j \rangle|$ where \mathbf{d}_i stands for an atom of dictionary \mathbf{D} . In this specific case, the mutual coherence is given by:

$$\max \left(\max_{(i,j)} |\langle \mathbf{A}^i, \mathbf{A}^j \rangle|, \max_{(p,q)} |\langle \mathbf{O}_D^p, \mathbf{O}_D^q \rangle|, \max_{(l,k)} |\langle \mathbf{A}^l, \mathbf{O}_D^k \rangle| \right), \quad (4.1)$$

In general terms, the separation between the outliers and the sources requires: i) the mixing matrix \mathbf{A} to have a small coherence (the sources should have clearly different spectra), ii) the spectra of the outliers to not be highly correlated (and so they cannot cluster in close directions), iii) the spectra of the outliers and the sources to not be too correlated as well. Unfortunately, this last condition is not easy to satisfy in the determined case/when m is small. For instance, for the Planck mission 2.1.1, $m = 9$, whereas there are more than a thousand of point source emissions: their emissions laws are necessarily correlated.

¹This general result can be refined using the special structure of the sensing matrix, [84]

In brief, the joint estimation of \mathbf{A} , \mathbf{S} , and column-sparse \mathbf{O} is possible if the supports of the outliers and sources are disjoint, or if the outliers lie in the subspace orthogonal to the span of \mathbf{A} .

Otherwise, one can only expect (if the outliers are in general position - column sparse), to retrieve the support of the outliers not strictly lying in the span of \mathbf{A} , and \mathbf{A} . This is quite similar to the results of Outliers Pursuit [135], 2.2.1. However, we will further aim to retrieve \mathbf{A} and not only its span, including in the determined case.

4.2 Robust GMCA

In the following of this chapter, we will present two robust BSS methods which aim to retrieve \mathbf{A} , by estimating jointly the components \mathbf{A} , \mathbf{S} , and \mathbf{O} . They both build upon sparse modeling of the data to highlight the structural difference between \mathbf{AS} and \mathbf{O} .

More precisely, we start with an algorithm coined robust GMCA (rGMCA), extending GMCA to robust BSS ². It will be assumed that the corrupted samples are composed of both entirely corrupted columns and (few) uniformly distributed corrupted entries.

4.2.1 A naive extension of GMCA

A straightforward strategy to account for the presence of outliers in the framework of GMCA is done by including an extra sparse term \mathbf{O} enforcing the sparsity of the outliers. This approach, named Naive robust GMCA (NrGMCA), can be formulated as:

$$\operatorname{argmin}_{\mathbf{O}, \mathbf{A}, \mathbf{S}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS} - \mathbf{O}\|_2^2 + \sum_{j=1}^n \lambda_j \|\mathbf{S}_j\|_1 + \alpha \|\mathbf{O}\|_1, \quad (4.2)$$

where the first term is the data-fitting term, well suited to deal with the Gaussian noise \mathbf{N} , and the two others terms enforce the sparsity of the sources and the outliers. This formulation was proposed in [5], [1] and [60] (with an $\ell_{2,1}$ norm for the outliers, column-sparse) both in the framework of dictionary learning.

This problem is very effective for column and row sparse outliers. Indeed, let us assume for the sake of clarity that $m = n$ and that the outliers and sources have at most one active entries per column. The outliers can be factorized as $\mathbf{O} = \mathbf{I}_{n \times n} \mathbf{O}$. Hence the norm of the outliers seen as a source contribution (factorized with \mathbf{I}), is still equal to $\|\mathbf{O}\|_1$: there is no gain for factorizing the outliers, fig.4.2b. On the other hand, if \mathbf{A} is quite dense (several entries active per column), then $\|\mathbf{AS}\|_1 = \sum_{j=1}^n \|\mathbf{A}^j\|_1 \|\mathbf{S}_j\|_1$ (for exactly sparse and disjoint sources) is larger than $\|\mathbf{S}\|_1$ since much more entries are active, fig.4.2a. In summary, there is no gain for the cost-function

²This section has been adapted from [33]

to wrongly estimate the outliers as a source contribution nor to estimate a part of \mathbf{AS} as corresponding to some outliers: the sparsest separation should lead to a good estimation of the outliers and \mathbf{AS} .




Figure 4.2: Schematic illustrations of the behavior of the source and outlier contributions in the source domain (\mathbf{S} and \mathbf{O} on the left), and in the observation domain (\mathbf{AS} and \mathbf{O} on the right). The sources are sparse in the direct domain with disjoint support. The outliers are also sparse in the direct domain, and with a support uniformly drawn at random.




Figure 4.3: Schematic illustration of the factorization of a matrix \mathbf{W} , which is column sparse and whose active columns are collinear. This term can be factorized as $\mathbf{W} = \mathbf{UV}$.

However, this first naive approach cannot handle large and column-sparse outliers, especially if \mathbf{A} is not dense (only few entries of \mathbf{A} are active). Indeed, whenever the outliers are column-sparse, there is a gain in term of sparsity, for factorizing the columns of \mathbf{O} sharing a same direction. Let $\mathbf{W} \in \mathbb{R}^{m \times t}$ be a column sparse matrix, whose active columns are collinear. This matrix can be factorized with the product \mathbf{UV} , where $\mathbf{U} \in \mathbb{R}^{m \times 1}$ a normalized vector collinear with the active columns of \mathbf{W} , and $\mathbf{V} \in \mathbb{R}^{1 \times t}$, whose columns are such $\mathbf{V}^k = \text{sign}(\langle \mathbf{U}, \mathbf{W}^k \rangle) \times \|\mathbf{W}^k\|_2$. The term \mathbf{W} can represent the contribution of a source i : $\mathbf{A}^i \mathbf{S}_i$, with $\mathbf{U} = \mathbf{A}^i$, and $\mathbf{V} = \mathbf{S}_i$. It can also be employed to represent some outliers sharing a same spectrum/orientation, see fig.4.3.

In this case, we have $\|\mathbf{W}\|_1 = \|\mathbf{U}\|_1 \times \|\mathbf{V}\|_1$. Thus, we observe a gain when factorizing \mathbf{W} : $\|\mathbf{W}\|_1 - \|\mathbf{V}\|_1 = (\|\mathbf{U}\|_1 - 1) \times \|\mathbf{V}\|_1$:

- $(\|\mathbf{U}\|_1 - 1) \geq 0$: this term depends on the 'density' of \mathbf{U} . Indeed, if the vector \mathbf{U} is exactly 1-sparse, then there is no benefit in term of sparsity, for factorizing the contribution. This is the case for column and row sparse outliers, fig.4.2b. On the other hand, if several entries of \mathbf{U} are active, then the factorization is costless.

Thus, if at least one column of \mathbf{A} is 1-sparse but one column of \mathbf{O} is not, it becomes impossible to avoid the misestimation of the two terms with eq.4.2.

- $\|\mathbf{V}\|_1$: whenever the vector \mathbf{U} is not exactly sparse, the benefit in term of sparsity (ℓ_1) is proportional to $\|\mathbf{V}\|_1$. Consequently, if there is at least one column of \mathbf{O} having more than 1 active entry and an amplitude much larger than the one of the sources, then NrGMCA should factorize this outlier contribution and consider the less energetic source contribution as corresponding to outliers.

That is why, in cases where the outliers are the dominant contribution to the data, NrGMCA erroneously estimates these largest samples as being the most discriminant entries of the sources for the separation and thus fail to estimate \mathbf{A} .

4.2.2 The rGMCA algorithm

The MDP-based algorithms (with ℓ_1 enforcing sparsity) assume that the most discriminant samples for the source separation are the largest ones 1.2.2. This assumption is no longer valid in the presence of large outliers, which mislead the estimation of the mixing matrix, Section 2.1.3. Estimating \mathbf{O} , \mathbf{A} and \mathbf{S} becomes more challenging and requires at least improving the robustness of the estimation of \mathbf{A} against the influence of the outliers. For this purpose, we propose to extend NrGMCA building upon the AMCA algorithm [13] (further details on AMCA will be given in 4.3.2).

In a different context, the AMCA algorithm extends GMCA in the special case of sparse and partially correlated sources, where the MDP does not hold [13]. In brief, this method relies on an iterative weighting scheme that penalizes non-discriminant entries of the sources. Inspired by this approach, we propose to implement a similar weighting scheme to penalize samples that are likely to be contaminated with large outliers, similarly to the weighting scheme proposed in [79]. In the spirit of AMCA, the influence of the corrupted samples are weakened by using a weighting scheme in the mixing matrix update stage with the following cost function:

$$\operatorname{argmin}_{\mathbf{O}, \mathbf{A}, \mathbf{S}} \frac{1}{2} \|(\mathbf{X} - \mathbf{AS} - \mathbf{O})\mathbf{W}\|_2^2 + \sum_{j=1}^n \lambda_j \|\mathbf{S}_j\|_1 + \alpha \|\mathbf{O}\|_1,$$

where \mathbf{W} is the penalizing diagonal matrix of size $t \times t$. The role of the weighting procedure is to penalize the samples of the sources that are likely to be corrupted with outliers. It is therefore natural to define the weights \mathbf{W} based on the current estimate of the outlier matrix $\tilde{\mathbf{O}}$. In the spirit of [13], an efficient weighting procedure consists in defining the weights based on the level of corruption of the columns of the outlier matrix as follows: $\mathbf{W}_{i,i} = \frac{1}{\varepsilon + \|\tilde{\mathbf{O}}^i\|_1}$, where ε stands for $\frac{\text{median}_{|S_{(i,i)}| > 0} |S_{(i,i)}|}{10}$. Subsequently, the penalization of a given data sample will increase with the amplitude of the outliers as well as the number of outliers per data sample (entirely corrupted columns are more difficult to distinguish from the sources contribution, assuming that all the entries of \mathbf{A} are active). Besides, the value ε depends on the source amplitude: the outliers are damaging if they are larger than the source contributions (and so should be penalized according to the source level).

Following the structure of the GMCA algorithm, this problem is solved by using PALS 3.2.2. The structure of the algorithm is presented in Algorithm 14. Instead of estimating alternatively \mathbf{O} , \mathbf{S} and \mathbf{A} , we found that applying GMCA to the current estimate of $\mathbf{X} - \tilde{\mathbf{O}}$ for estimating \mathbf{A} and \mathbf{S} and then to estimate \mathbf{O} from $\mathbf{X} - \tilde{\mathbf{A}}\tilde{\mathbf{S}}$ provides the most effective estimation procedure:

- Estimating jointly $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{S}}$ from $\mathbf{X} - \tilde{\mathbf{O}}$ amounts to estimate the subspace $\tilde{\mathbf{A}}\tilde{\mathbf{S}}$ of rank n , embedding most of the samples, and to unmix the data samples. This is the non-convex part of the problem, but building upon GMCA and its thresholding strategy 3.3.3, we can develop a robust and efficient unmixing.
- Estimating \mathbf{O} from $\mathbf{X} - \tilde{\mathbf{A}}\tilde{\mathbf{S}}$ aims to recover the corrupted data samples, not correlated with \mathbf{A} .

Choice of parameters and initialization. The large outliers are the most damaging as they can severely mislead the estimation of the mixing matrix if they are not estimated as outliers from the start. That is why, the algorithms NrGMCA and rGMCA start by estimating the largest values of \mathbf{X} as outliers. On the other hand, the orientation of the mixing matrix $\tilde{\mathbf{A}}$, which is initialized as a random matrix whose columns are normalized and entries are Gaussian, is deduced from the data cleaned from these large outliers $\mathbf{X} - \tilde{\mathbf{O}}$: our first estimation of $\tilde{\mathbf{O}}$ should not be too conservative to keep the clustering aspect of $\mathbf{X} - \tilde{\mathbf{O}}$. For this purpose, we propose to estimate $\tilde{\mathbf{O}}$ with a soft-thresholding at the value $\tilde{\alpha}^0 = \text{median}_{|\mathbf{X}| > 3\sigma} |\mathbf{X}|$, where σ denotes the standard deviation of the Gaussian noise \mathbf{N} . The thresholding strategy employed during the unmixing was presented in Section 3.3.3. The threshold α decreases linearly towards its final value, 3σ , see Appendix B.³

In practice, the algorithm is stopped whenever the mixing matrix \mathbf{A} is stable.

³The thresholding strategy, especially for the outliers, can be largely improved. This version requires the knowledge of σ , whereas it could be estimated directly from the data.

Algorithm 14 rGMCA Algorithm

```

1: procedure rGMCA( $\mathbf{X}$ ,  $n$ )
2:   Set  $\tilde{\alpha}^0$ ,  $\tilde{\mathbf{O}}^0$ ,  $\tilde{\mathbf{A}}^{0,0}$ ,  $\tilde{\lambda}^0$ ,  $\tilde{\mathbf{S}}^{0,0}$ , and  $\tilde{\mathbf{W}}^0$ .
3:   while  $k < K$  do
4:     while  $j < J$  do
5:        $\tilde{\mathbf{S}}^{k,j} = \mathcal{S}_{\tilde{\lambda}^j} \left( (\tilde{\mathbf{A}}^{k,j-1})^\dagger (\mathbf{X} - \tilde{\mathbf{O}}^{k-1}) \right)$ 
6:        $\tilde{\mathbf{A}}^{k,j} = (\mathbf{X} - \tilde{\mathbf{O}}^{k-1}) \tilde{\mathbf{W}}^k (\tilde{\mathbf{S}}^{k,j} \tilde{\mathbf{W}}^k)^\dagger$ 
7:       Set  $\tilde{\lambda}^{j+1} < \tilde{\lambda}^j$ 
8:        $\tilde{\mathbf{O}}^k = \mathcal{S}_{\tilde{\alpha}^k} (\mathbf{X} - \tilde{\mathbf{A}}^{k,J-1} \tilde{\mathbf{S}}^{k,J-1})$ 
9:       Update  $\tilde{\mathbf{W}}^k$ 
10:      Set  $\tilde{\alpha}^{k+1} < \tilde{\alpha}^k$ 
11:      Set  $\tilde{\mathbf{A}}^{k+1,0} = \tilde{\mathbf{A}}^{k,J-1}$ 
return  $\tilde{\mathbf{S}}^{(k-1)}$ ,  $\tilde{\mathbf{A}}^{(k-1)}$ ,  $\tilde{\mathbf{O}}^{(k-1)}$ .

```

4.2.3 Evaluation of rGMCA

In this subsection, we evaluate rGMCA in the field of the biomedical engineering with simulated data. We propose to separate the different Nuclear Magnetic Resonance spectra of a simulated mixture, which can represent the data provided in NMR spectroscopy. By performing BSS on the mixture of spectra, we should be able to identify the different molecules of the mixture [113].⁴ The estimated NMR spectra of the menthone, the folic acid, the ascorbic acid and the myo-inositol from SDBS⁵ are convolved with a Laplacian kernel of 2-samples width at half maximum (implementation from pyGMCA⁶). The number of observations is set according to fig.4.4. The Gaussian noise \mathbf{N} is drawn from a Gaussian law with a standard deviation of 0.1. The outliers are drawn from a Gaussian law with standard deviation 10^3 , so that 20 columns and 1% of the entries, broadly distributed, are corrupted (the maximal amplitudes of the sources is around 1000 also, fig.4.5).

4.2.3.1 NMR spectrum identification and influence of the number of observations. We compare the performances of rGMCA with standard BSS methods: GMCA [14], PCP+GMCA (the outliers are first estimated with PCP 2.2.1 and then discarded from \mathbf{X} [23]), the minimization of the β -divergence with statistical independence prior 2.2.2 (implementation from [63]) and NrGMCA. The values of the free parameters for PCP (λ in [23]) and the β -divergence minimization algorithm (β) are tuned to return the best results based on several trials. We set $J=500$ and $K=20$ (typical values are a few hundreds and tens respectively). The minimization of the β -divergence only returns an estimate for \mathbf{A} , and thus we propose to compare the methods by using a criterion depending only on \mathbf{A} [14]: $\Delta_A = \frac{\|\mathbf{P}\mathbf{A}^\dagger \mathbf{A} - \mathbf{I}\|_1}{n^2}$, where \mathbf{P} stands for the scale/permutation indeterminacies correction, App.C. The medians of Δ_A ob-

⁴These mixtures have a morphology similar to the LC-MS data introduced in Section 1.1.1. That is why they served as realistic data sets in [115], prior to the utilization of the algorithms on the real LC-MS data.

⁵<http://sdfs.db.aist.go.jp>

⁶<http://www.cosmostat.org/software/gmcalab/>

tained from 80 Monte-Carlo simulations are displayed in fig.4.4.




Figure 4.4: Δ_A versus the number of observations.

The GMCA and minimization of the β -divergence are not able to perform correctly with this setting. Besides, we can observe a clear difference between the results obtained by NrGMCA and rGMCA, illustrating the benefit of the weighting scheme to provide robustness fig.4.4. Despite the effectiveness of rGMCA and PCP+GMCA if $m \gg n$ fig.4.4, none of these algorithms are able to handle the outliers if $m = n$. The algorithm rGMCA is the only one that provides a correct estimate of \mathbf{A} for $m \geq 5$ by means of the weighting scheme. However, rGMCA cannot clearly separate the outliers from the sources: some leakages from the outliers towards the estimated sources are present fig.4.5.




Figure 4.5: Estimates of the menthone's NMR spectrum (original spectrum: top of the images) with rGMCA (middle) and GMCA (bottom), with $m = 6$. Red stars denote samples corrupted by outliers.

4.2.3.2 Limitations of rGMCA. We introduced a novel method to separate sparse sources in the presence of outliers. The proposed method relies on the joint sparsity-based separation of the outliers and the sources. This strategy allowed us to implement a weighting scheme that penalizes corrupted data samples, which showed to highly limit the impact of the outliers during the unmixing process. Numerical experiments demonstrated the good performances of our algorithm to robustly estimate the mixing matrix.

However, the proposed rGMCA is unable to estimate robustly the mixing matrix in the determined case. The latter is particularly difficult to tackle since the outliers also lie in the span of \mathbf{A} . This makes the weighting scheme of rGMCA inefficient. Indeed, the current outliers can be almost null and very difficult to detect since for any mixing matrix $\tilde{\mathbf{A}}$, $\mathbf{X} = \tilde{\mathbf{A}}(\tilde{\mathbf{A}}^\dagger \mathbf{X})$ - the residual $\mathbf{X} - \tilde{\mathbf{A}}\tilde{\mathbf{S}}$ can be very small (and so is $\tilde{\mathbf{O}}$).

Last, we point out that instead of using \mathbf{O} for the weighting, it would have been also efficient to base our weighting scheme on $\mathbf{X} - \tilde{\mathbf{A}}\tilde{\mathbf{S}}$ instead of \mathbf{O} : even if the outliers were wrongly estimated during the initialization, the weighting would have been efficient (for $m > n$).

4.3 Robust AMCA

In this section, we propose a novel algorithm, named robust Adaptive Morphological Component Analysis (rAMCA), that extends the rGMCA algorithm ⁷. Unlike the rGMCA algorithm, the proposed algorithm further relies on two novel elements: i) a refined modeling of the outliers in the source domain based on an analogy with partially correlated sources, and ii) an improved outlier estimation procedure. The proposed method is shown to yield a highly effective estimation of the mixing matrix. It also performs very well when the number of observations is close or equal to the number of sources; a challenging setting for which currently available robust BSS methods fail.

In contrast with the rGMCA framework, we will assume that the outliers are only column-sparse. This sparsity pattern is well suited for multi/hyperspectral images analysis: it corresponds to the presence of anomalies, visible at every frequency, such as the point source emissions in the Planck data 2.1.1.

4.3.1 An analogy with partially correlated sources

In the previous section, it clearly appears that solving robust BSS problems requires discriminating between the outliers and the sources. Since both the outliers and the sources are assumed to

⁷This section is adapted from [31]




Figure 4.6: Three sparse sources are mixed into 4 noisy observations. Fig.(a): scatter plot of two noisy observations, (b): scatter plot of two of the projected data given by $\hat{\mathbf{S}} = \mathbf{A}^\dagger \mathbf{X}$, (c): scatter plot of the estimated sources given by $\tilde{\mathbf{S}} = \tilde{\mathbf{A}}^\dagger \mathbf{X}$, where $\tilde{\mathbf{A}}$ has been estimated with GMCA and is far from the initial \mathbf{A} . The initial source contribution is represented in blue, and the one of the outliers with the red stars.

be sparsely represented in the same domain, sparsity alone cannot be the right separation criterion. Fortunately, both components are assumed to have different distributions: source samples tend to cluster along the canonical axes in the source domain while the samples of the projected outliers $\hat{\mathbf{O}} = \mathbf{A}^\dagger \mathbf{O}$ (*i.e.* projection of \mathbf{O} in the source domain) do not have any preferred clustering direction⁸. This is testified by the difference between the distributions in the source domain of the sources samples (blue dots) and corrupted samples (red stars) in fig.4.6b.

If the mixing matrix \mathbf{A} were perfectly known, the sources would be approximated by projecting the corrupted data onto the span of \mathbf{A} : $\hat{\mathbf{S}} = \mathbf{A}^\dagger \mathbf{X}$. The estimated sources are the linear combination of the clean sources and the projected outliers: $\hat{\mathbf{S}} = \mathbf{S} + \hat{\mathbf{O}}$. Due to the projected outliers contribution which is broadly distributed, some of the largest entries of $\hat{\mathbf{S}}$ are active simultaneously in several sources (*c.f.* the red contribution in fig.4.6b). These shared active samples are reminiscent of the partial correlations of the sources discussed in [13]. Indeed, the samples of partially correlated sources can be similarly divided into two groups: the discriminant samples respecting the MDP (the jointly sparse contribution in blue in fig.4.6b) and the samples corresponding to the partial correlations which active simultaneously in several sources (the broadly distributed contribution in red in fig.4.6b). Unlike the rGMCA algorithm we introduced previously, we propose to exploit the analogy between the impact of the projected outliers and sparse and partially correlated sources, which yields a novel robust BSS algorithm that is described in the following.

4.3.2 Robust AMCA Algorithm

Following the analogy between the impact of outliers and partial correlations, the rAMCA algorithm will be built upon the AMCA algorithm (Adaptive Morphological Component Analysis),

⁸They do not cluster along the canonical axes: they are not jointly sparse and do not respect the MDP. They can 'cluster' in the source domain if \mathbf{A} has a large condition number.

which has been designed to deal with partially correlated sources [13]. In this specific context, it was underlined in [13] that the ability to identify correlated entries is critical to perform the separation. For that purpose, the AMCA algorithm builds upon an adaptive weighting scheme that assigns to each column of the observation coefficients \mathbf{X} a weight, whose goal is to penalize correlated entries in the separation process. Details about the weighting procedure will be given below. The AMCA algorithm performs by minimizing the following problem:

$$\underset{\mathbf{A}, \mathbf{S}}{\text{minimize}} \frac{1}{2} \|(\mathbf{X} - \mathbf{AS}) \mathbf{W}\|_2^2 + \sum_{i=1}^n \lambda_i \|\mathbf{S}_i\|_1, \quad (4.3)$$

where $\mathbf{W} \in \mathbf{R}^{t \times t}$ is the weight matrix.

In the spirit of rGMCA, we propose to estimate jointly \mathbf{A} , \mathbf{S} and \mathbf{O} by exploiting the sparsity of the sources and the outliers. Unlike this first algorithm, we propose to further employ a weighting scheme similar to AMCA, penalizing the correlated entries of the estimated sources. This can be done by substituting the problem in eq. 4.3 with the following one:

$$\underset{\mathbf{S}, \mathbf{A}, \mathbf{O}}{\text{minimize}} \frac{1}{2} \|(\mathbf{X} - \mathbf{AS} - \mathbf{O}) \mathbf{W}\|_2^2 + \sum_{i=1}^n \lambda_i \|\mathbf{S}_i\|_1 + \beta \|\mathbf{O}\|_{2,1}. \quad (4.4)$$

The $\ell_{2,1}$ norm, defined such as $\|\mathbf{O}\|_{2,1} = \sum_{j=1}^m \|\mathbf{O}^j\|_2$, favors solutions \mathbf{O} with few entirely active columns. This regularization term is well suited to capture outliers that are distributed in general position in the data domain.

This problem is non-convex but can be tackled using a minimization procedure such as BCD or PALS 3.2, which sequentially minimize subsets of variables. A natural choice would consist in estimating alternatively the three variables of interest \mathbf{A} , \mathbf{S} , and \mathbf{O} . However, we found that this choice performs poorly in practice since errors are more likely to propagate from one variable to the other during the sequence of minimization steps. Besides, the propagation of errors is strengthened by the joint estimation of the regularization parameters, depending on the current variables. We rather opted for a sequential minimization of two blocks of variables (similarly to rGMCA):

- Joint estimation of (\mathbf{A}, \mathbf{S}) : This is the non-convex part of the problem. This unmixing step is robustified by means of the weighting scheme, penalizing the corrupted samples. Besides, the regularization parameters λ_i can be estimated using the decreasing threshold strategy of GMCA 3.3.3: we can start with large thresholds λ_i to provide robustness, without modifying the estimation of \mathbf{O} since they are fixed (more precisely λ_i can be large without facing strong leakages from the sources towards the estimated outliers).
- Estimation of the outliers matrix \mathbf{O} : The outliers correspond, in this problem, to the contribu-

tions not lying in the column-span of \mathbf{A} (if $m > n$) or to the correlated entries of the estimated sources. They are thus easiest to detect and estimate given a current tuple $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$. As well, the regularization parameters are adapted to discard the samples that appear to be corrupted.

The algorithm is presented in Alg.15.

Procedure 15 rAMCA Algorithm

```

1: procedure RAMCA( $\mathbf{X}, n$ )
2:   Initialize  $\tilde{\mathbf{A}}^{(0)}$  (randomly or with a PCA),  $\tilde{\mathbf{S}}^{(0)} = \mathbf{0}$  and  $\tilde{\mathbf{O}}^{(0)} = \mathbf{0}$ .
3:   while  $k < K$  do
4:     Set  $\tilde{\mathbf{S}}^{(0,k)} \leftarrow \tilde{\mathbf{S}}^{(k-1)}$  and  $\tilde{\mathbf{A}}^{(0,k)} \leftarrow \tilde{\mathbf{A}}^{(k-1)}$ 
5:     while  $i < I$  do ▷ Joint estimation of  $\mathbf{A}$  and  $\mathbf{S}$ 
6:       Update  $\tilde{\mathbf{S}}^{(i,k)}$  from (4.6)
7:       Update  $\tilde{\mathbf{W}}$  from (4.8)
8:       Update  $\tilde{\mathbf{A}}^{(i,k)}$  from (4.7)
9:     Set  $\tilde{\mathbf{S}}^{(k)} \leftarrow \tilde{\mathbf{S}}^{(i-1,k)}$  and  $\tilde{\mathbf{A}}^{(k)} \leftarrow \tilde{\mathbf{A}}^{(i-1,k)}$ 
10:    Update  $\tilde{\mathbf{O}}^{(k)}$  from (4.10) ▷ Estimation of  $\mathbf{O}$ 
  return  $\tilde{\mathbf{S}}^{(k-1)}, \tilde{\mathbf{A}}^{(k-1)}, \tilde{\mathbf{O}}^{(k-1)}$ .

```

4.3.2.1 Estimating the sources and the mixing matrix. Applying the PALS technique to estimate the mixing matrix and the sources amounts to minimizing the problem in Eq. 4.4 assuming \mathbf{O} is fixed:

$$\underset{\mathbf{S}, \mathbf{A}}{\text{minimize}} \frac{1}{2} \|(\mathbf{X} - \mathbf{AS} - \mathbf{O}) \mathbf{W}\|_2^2 + \sum_{i=1}^n \lambda_i \|\mathbf{S}_i\|_1. \quad (4.5)$$

The problem shares similarities with the problem solved by the AMCA algorithm (see (4.3)) with the exception that it applies to the residual $\mathbf{X} - \mathbf{O}$ rather than the raw observations \mathbf{X} . Following the AMCA algorithm, the problem in (4.5) is tackled by minimizing alternately the cost function with respect to \mathbf{A} and \mathbf{S} with the two following steps (PALS based procedure 3.2.2):

- Updating \mathbf{S} assuming \mathbf{A} is fixed : Minimizing (4.5) with respect to \mathbf{S} consists in solving the following convex problem:

$$\underset{\mathbf{S}}{\text{minimize}} \frac{1}{2} \|(\mathbf{X} - \mathbf{AS} - \mathbf{O}) \mathbf{W}\|_2^2 + \sum_{i=1}^n \lambda_i \|\mathbf{S}_i\|_1.$$

Unless \mathbf{A} is orthogonal, the previous problem does not admit a closed form solution. In the spirit of alternated least-square minimization techniques proposed in [13] to rather approximate this step with a projected least-square, which highly limits the computational cost of

the update ⁹:

$$\mathbf{S}_i = \mathcal{S}_{\lambda_i}([\mathbf{A}^\dagger (\mathbf{X} - \mathbf{O})]_i). \quad (4.6)$$

- Updating \mathbf{A} assuming \mathbf{S} is fixed : Minimizing (4.5) with respect to \mathbf{A} amounts to solving the following convex problem:

$$\underset{\mathbf{A}}{\text{minimize}} \frac{1}{2} \|(\mathbf{X} - \mathbf{A}\mathbf{S} - \mathbf{O}) \mathbf{W}\|_2^2.$$

which admits a closed form solution:

$$\mathbf{A} = (\mathbf{X} - \mathbf{O}) \mathbf{W} (\mathbf{S}\mathbf{W})^\dagger. \quad (4.7)$$

In practice, to avoid the balance indeterminacy between \mathbf{A} and \mathbf{S} , we assume that the columns of \mathbf{A} are normalized for the ℓ_2 norm. Similarly to what is done with AMCA, this additional constraint is handled by normalizing the columns of \mathbf{A} after the projected least-squares.

Similarly to the AMCA algorithm, the weights play a central role. In the setting of robust BSS, they help providing robustness to the remaining outliers contribution in the estimated residual $\mathbf{X} - \mathbf{O}$. Following the analogy with partial correlations, the weights aim at penalizing entries of the estimated sources which are in general position rather than clustered along a canonical axis fig.4.6b. The former are more likely related to residuals of outliers while the latter are characteristics of the sources. Following [13], samples in general position can be traced by measuring the sparsity level of the columns of the estimated sources using an ℓ_q norm, $q \in (0, 1)$. For illustrative purpose, we display in fig.4.7 the value of the ℓ_q norm of the projected sources samples: the correlated entries (corresponding to $\mathbf{A}^\dagger \mathbf{O}$) correspond to the entries having the largest ℓ_q norm, especially when q is small.

Therefore, the diagonal elements of the weight matrix \mathbf{W} are defined as follows:

$$\mathbf{W}_{i,i} = \frac{1}{\sqrt{\|\mathbf{S}^i\|_q + \epsilon}}, \quad \forall i = 1, \dots, t \quad (4.8)$$

where \mathbf{S} denotes the normalized sources $\mathbf{S}_i = \frac{\mathbf{S}_i}{\|\mathbf{S}_i\|_2}$ and where ϵ is a scalar typically small used to avoid numerical issues. The sources are normalized in order to take into account large dynamics between the sources: a coefficient can be significant and discriminant for a source with a small amplitude while being smaller than the non-discriminant samples of larger sources. The param-

⁹The weights \mathbf{W} are taken into account in the parameters λ_i which are varying.




Figure 4.7: Left- top to bottom: projected sources $\tilde{\mathbf{S}} = \mathbf{A}^\dagger \mathbf{X}$, then $\|\tilde{\mathbf{S}}^k\|_1$ and $\|\tilde{\mathbf{S}}^k\|_{0.1} \forall k = 1..t$. Right: sources samples projected on the ℓ_q ball.

eter q is chosen in the range $[0, 1]$. More precisely, it is set according to $\frac{0.1^{\frac{1}{k}}}{2}$. It starts with a larger value, to not penalize too strongly the unmixed sources, and then decreases to penalize the non-sparse entries, the projected outliers fig.4.7: starting directly with a too small value of q may prevent any unmixing by forcing the solutions to stay at it is, and on the other hand, very large outliers require a strong penalization obtained with a small value of q .

In fig.4.8, we display an example of the current estimated sources $\tilde{\mathbf{S}}$ and the ones 'perceived' during the unmixing with AMCA. In practice, it amounts to project the normalized sources on the ℓ_q ball, fig.4.7b. Thanks to the weighting scheme, the influence of the large correlated entries of \mathbf{S} (the corrupted entries) is clearly lessened: they do not hinder the unmixing.

4.3.2.2 Estimating the outliers. In the rAMCA algorithm, the estimation of \mathbf{O} given \mathbf{A} and \mathbf{S} is carried out by solving the problem in (4.4):

$$\underset{\mathbf{O}}{\text{minimize}} \frac{1}{2} \|(\mathbf{X} - \mathbf{AS} - \mathbf{O}) \mathbf{W}\|_2^2 + \beta \|\mathbf{O}\|_{2,1}.$$

Given that only the diagonal terms of \mathbf{W} are non-zero, this problem is separable. It amounts to solve for each sample $k \in \{1..t\}$:

$$\underset{\mathbf{O}^k}{\text{minimize}} \frac{1}{2} \left\| \left((\mathbf{X} - \mathbf{AS})^k - \mathbf{O}^k \right) \mathbf{W}_k^k \right\|_2^2 + \beta \|\mathbf{O}^k\|_2.$$

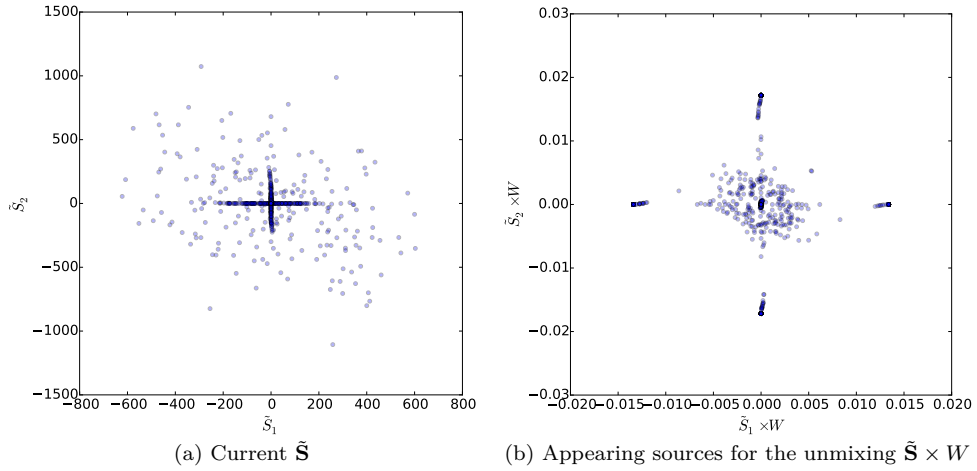


Figure 4.8: Left: scatter plot of two estimated sources $\tilde{\mathbf{S}}$ during AMCA. Right: scatter plot of the 'appearing' sources for the unmixing $\tilde{\mathbf{S}}\mathbf{W}$, for $q = 0.2$

This problem is equivalent to:

$$\text{minimize}_{\mathbf{O}^k} \frac{(\mathbf{W}_k^k)^2}{2} \left\| (\mathbf{X} - \mathbf{AS})^k - \mathbf{O}^k \right\|_2^2 + \beta \left\| \mathbf{O}^k \right\|_2.$$

Then, by setting $\tilde{\beta} = \frac{\beta}{(\mathbf{W}_k^k)^2}$, we end up with:

$$\text{minimize}_{\mathbf{O}^k} \frac{1}{2} \left\| (\mathbf{X} - \mathbf{AS})^k - \mathbf{O}^k \right\|_2^2 + \tilde{\beta} \left\| \mathbf{O}^k \right\|_2.$$

This problem has a closed form solution which has been derived in [83]:

$$\mathbf{O}^k = (\mathbf{X} - \mathbf{AS})^k \times \left(1 - \frac{\tilde{\beta}^k}{\left\| (\mathbf{X} - \mathbf{AS})^k \right\|_2} \right)_+. \quad (4.9)$$

Detecting the outliers. Most sparsity-based thresholding procedures can be interpreted as detection procedures: detecting sparse samples out of dense noise. In that case, it is customary to fix the value of the threshold based on the noise statistics [124], see App.B. Similarly, and according to (4.9), the support (*i.e.* the set of active columns) of \mathbf{O} is defined by the set of columns whose ℓ_2 norm exceeds the threshold $\tilde{\beta}$. Ideally, the columns having an ℓ_2 norm smaller than $\tilde{\beta}$ should correspond to the remaining Gaussian noise. Consequently, the values of $\tilde{\beta}$ should also be fixed based on the Gaussian noise statistics. In that case, only the Gaussian noise contributes to the residual outside the support of \mathbf{O} . Therefore, the samples $\left\{ \left\| (\mathbf{X} - \mathbf{AS})^k \right\|_2 \right\}_{k: \left\| \mathbf{O}^k \right\|_2 = 0}$ follow a χ law with m degrees of freedom. The value of $\tilde{\beta}$ can then be chosen based on the expected value of the χ law: $\sigma \times \sqrt{2} \times \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})}$, where σ corresponds to the standard deviation of \mathbf{N} , App. B.

Nevertheless, relying on the noise statistics only provides a detection procedure that is not reliable in the determined case. Indeed, even if \mathbf{A} is correctly recovered, the outliers are very likely to leak into the estimated sources $\tilde{\mathbf{S}}$ since they also lie in the span of \mathbf{A} : $\tilde{\mathbf{S}} = \mathbf{S} + \mathbf{A}^\dagger \mathbf{O}$, such that $\mathbf{A}\tilde{\mathbf{S}} = \mathbf{A}\mathbf{S} + \mathbf{O}$. An accurate detection of the outliers based on the residual $\mathbf{X} - \mathbf{A}\tilde{\mathbf{S}}$ is then not possible. To overcome this issue, we propose to rather build the detection procedure on a quantity that allows discriminating between the outliers and the sources, especially in the determined case.

We emphasized in Section 4.3.1 that in the source domain the entries of \mathbf{S} are jointly sparse, *i.e.* clustered along the canonical axes, whereas the projected outliers behave as correlated non-sparse entries. In this context, the δ -density, which has been introduced in [125], provides a convenient measure of sample sparsity that permits to discriminate between sparse and non-sparse columns of $\tilde{\mathbf{S}}^{10}$. The δ -density of any j th non-zero sample of the estimated sources is defined as $\delta(\tilde{\mathbf{S}}^j) = \frac{\|\tilde{\mathbf{S}}^j\|_1}{\|\tilde{\mathbf{S}}^j\|_\infty}$. This quantity takes its values between 1 (for one active entry, *i.e.* 1-sparse column vector $\tilde{\mathbf{S}}^j$) and n (for a column whose entries have the same amplitudes, *i.e.* highly non-sparse, 'non-informative' vector $\tilde{\mathbf{S}}^j$). More interestingly, it is independent on the amplitude of the columns and well suited for sparse and approximately sparse signals. Since the corrupted columns of $\tilde{\mathbf{S}}$ are non-sparse, they consequently have large δ -density values. In this framework, detecting the support of \mathbf{O} can be performed by identifying the columns of the estimated $\tilde{\mathbf{S}}$ whose δ -density is larger than a certain threshold α that needs to be determined. This is somehow reminiscent of the outlier detection discussed in [134].

In the general setting, determining an optimal numerical value for α is challenging without an accurate statistical modeling of the sources and the outliers. We propose to use the following statistical modeling:

- the sources are drawn from a generalized Gaussian law with parameter ρ denoted by $\mathcal{G}(\rho)$.
- the amplitudes of the outliers in the sources domain follow a Gaussian law \mathcal{N} , well suited to model samples that are distributed in general position.

Let us notice that the variances of these statistical models do not matter since the δ -density is independent on the amplitude. From this statistical model, the threshold α is derived from a classical hypothesis testing procedure such that, for any random variable X of size n :

$$\Pr(\delta(X) < \alpha | X \sim \mathcal{G}(\rho)) = \Pr(\delta(X) > \alpha | X \sim \mathcal{N}).$$

where $\Pr(\delta(X) < \alpha | X \sim \mathcal{G}(\rho))$ stands for the probability for the δ -density to be smaller than α assuming that every entry of X is distributed according to a centered generalized Gaussian law with unit variance and shape parameter ρ . Figure 4.9 illustrates three different cases with $n = 10$: (a) the case $\rho = 1$, which corresponds to a low sparsity level, and then (b) and (c), the cases $\rho = 0.5$ and $\rho = 0.3$ that correspond to realistic sparsity levels for the coefficients of

¹⁰The δ -density should be performed on normalized sources if the sources do not share the same dynamics.

sparse representations of natural signals. The value of α varies from 3.9 to 3.3. Since we have no precise prior knowledge about the distributions, we derive numerically the value α for the corresponding n from the Laplacian law (the largest possible for sparse sources respecting the MDP). This choice is quite conservative for the sources since only the samples having a δ -density larger than α are estimated as being corrupted.




Figure 4.9: Numerical approximations of the cumulative distribution functions of $\delta(\cdot)$ for different values of ρ and $n = 10$. In blue: fig.(a): $Pr(\delta(X) < \gamma | X \sim \mathcal{G}(1))$, fig.(b) $Pr(\delta(X) < \gamma | X \sim \mathcal{G}(0.5))$, fig.(c) $Pr(\delta(X) < \gamma | X \sim \mathcal{G}(0.3))$. In red: $Pr(\delta(X) > \gamma | X \sim \mathcal{N})$.

According to (4.9), the amplitude of the detected outliers is derived from the estimated residual $\mathbf{X} - \tilde{\mathbf{A}}\tilde{\mathbf{S}}$. Previously, we underline that $\mathbf{X} - \tilde{\mathbf{A}}\tilde{\mathbf{S}}$ is very likely to contain some errors. A more conservative but more effective choice consists in deriving the amplitude of the detected outliers from the data \mathbf{X} .

As a summary, the outliers \mathbf{O} are estimated as follows:

$$\tilde{\mathbf{O}}^k = \begin{cases} 0 & \text{if } \delta(\tilde{\mathbf{S}}^k) < \alpha \\ \mathbf{X}^k \times \left(1 - \frac{\tilde{\beta}}{\|\mathbf{X}^k\|_2}\right) & \text{otherwise,} \end{cases} \quad (4.10)$$

where $\tilde{\beta} = \text{mad}(\mathbf{X} - \mathbf{A}\mathbf{S} - \mathbf{O}) \times \sqrt{2} \times \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})}$ and $\text{mad}(\mathbf{X} - \mathbf{A}\mathbf{S} - \mathbf{O})$ corresponds to a good estimate of the standard deviation of \mathbf{N} if it is not known. Despite the simplicity of the statistical model used to derive a value for α and consequently β , the proposed scheme has performed correctly in the numerical experiments of Section 4.4.2. Furthermore, at each iteration of the rAMCA algorithm 15, the couple (\mathbf{A}, \mathbf{S}) is fully re-estimated, which also makes the algorithm less sensitive to mis-estimations of the outliers \mathbf{O} .

4.3.2.3 Choice of the parameters. Strategy for λ : We adopt the decreasing thresholding strategy that was presented in 3.3.3. More precisely, given the total number of iterations I , the j th projected source $\tilde{\mathbf{S}}_j^{(i,k)} = (\tilde{\mathbf{A}}^{(i-1,k)\dagger}(\mathbf{X} - \tilde{\mathbf{O}}^{(k)}))_j$ is thresholded at the i th iteration by:

$$\lambda_j = \text{pct} \left(|\tilde{\mathbf{S}}_j^{(i,k)}|_{|\tilde{\mathbf{S}}_j^{(i,k)}| > 3\sigma_j}, 100 \times \frac{I-i}{I} \right).$$

Number of inner loops I : The number of iterations is set to $I = 1000$, which turned to be a good compromise in the numerical experiments.

Strategy for β : In the spirit of the decreasing value strategy used for $\tilde{\lambda}$ in AMCA, the number of eligible active samples of the estimated outliers is increased during the algorithm. More precisely, at the k th iteration, we select the outliers among the $5k\%$ largest entries of the residue in order to limit the number of false estimations. We underline that these parameters are also automatically determined: α depends only on the number of sources and β on the number of observations.

Number of outer loops K : Last, the number of outer loops is maximally set to 100. In practice, the algorithm is stopped when $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{O}}$ are jointly stabilized fig.4.10. More precisely, rAMCA stops at the k th outer loop if: $\max_{j=1..n} \langle \tilde{\mathbf{A}}^{(k-1)j}, \tilde{\mathbf{A}}^{(k)j} \rangle < 5^\circ$, and $\text{supp}(\tilde{\mathbf{O}}^{(k-1)}) = \text{supp}(\tilde{\mathbf{O}}^{(k)})$, where $\text{supp}(\mathbf{x})$ denotes the support of the vector \mathbf{x} .

4.3.2.4 Stability of rAMCA. Since the problem (4.4) is not convex, we can only expect to converge to a local minimum of the cost function of interest. Besides, given that the proposed strategy uses varying parameters, the convergence to a critical point, strictly speaking, cannot be proved. However, the stability of the two variables of interest, the support of the corrupted samples and the mixing matrix, is heuristically well motivated.

We propose to minimize the function using the PALS method, to fasten the standard BCD strategy, Chapter 2. It has been shown in [128] that minimizing (4.4) alternately for each variable with fixed parameters and the BCD strategy converges to a stationary point. However, in practice, minimizing (4.4) with the cyclic rule and with fixed parameters performs poorly: it is challenging to estimate jointly the variables and regularization parameters without propagating the errors. That is why, we minimize 4.4 using a sequential minimization alternating between the blocks (\mathbf{A}, \mathbf{S}) and \mathbf{O} , as well as a strategy with decreasing parameters.

Once the detrimental outliers (or the data estimated as being detrimental) have been removed from the observations, the AMCA algorithm, whose stability has been discussed in [13], returns a similar \mathbf{A} from one iteration to another (since the input $\mathbf{X} - \mathbf{O}$ is constant from one outer iteration to another one), fig.4.10a, 4.10b.

For illustrative purpose, we display the maximal angle made between the columns of $\tilde{\mathbf{A}}^{(k)}$ and $\tilde{\mathbf{A}}^{(k+1)}$ (see App. C for the metrics) as well as the percentage of estimated corrupting columns for $n = m = 10$ sources generated according to Section 4.4.1 and 30% of corrupting columns. After few outer loops, the number of estimated columns fig.4.10a and $\tilde{\mathbf{A}}$ almost not vary fig.4.10b (a variation with the maximal order of magnitude of $10^{-3}\%$ is observed for \mathbf{A} due to the projected least squares).

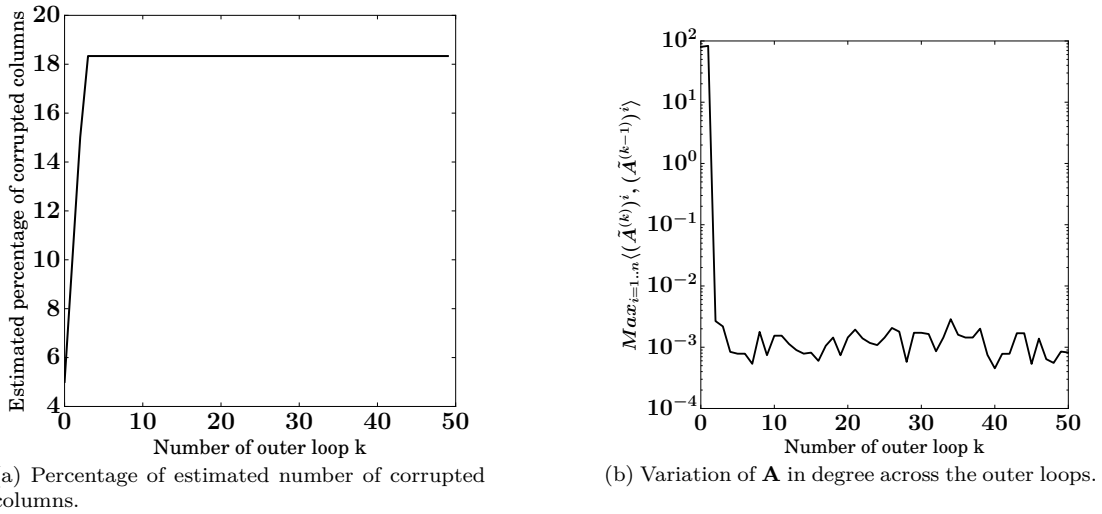


Figure 4.10: Convergence of rAMCA.

4.4 Numerical Experiments

4.4.1 Experimental protocol

In this section, rAMCA is compared with various robust BSS algorithms:

- GMCA 1.3.2: this standard sparse BSS method is used to illustrate the sensitivity of the *non-robust* BSS algorithms to the presence of outliers.
- AMCA [13] whose performances show the benefits of the weighting scheme (difference between AMCA and GMCA) and of the explicit estimation of \mathbf{O} (difference between AMCA and rAMCA).
- rGMCA: the discrepancy between its performances and the ones of rAMCA illustrates the key role of the novel penalization and outliers detection procedure, which are, unlike rGMCA, based on the refined modeling of the outliers in the source domain.
- the robust minimization of the β -divergence Section 2.2.2, [100], (implementation similar to [63]), which assumes that $m = n$ and only estimates the mixing matrix.
- the robust combination PCP+GMCA: the outliers are first estimated with PCP Section 2.2.1, [23] which assumes that $m \gg n$, and then the sources and mixing matrix are estimated with GMCA.

The parameters of PCP+GMCA and of the minimization of the β -divergence are manually tuned. In the first part of this section, their performances are evaluated on various scenarios with synthetic data, which allows performing Monte-Carlo simulations.

4.4.1.1 Performance criteria. We emphasize that the algorithms listed above do not yield a precise estimation of the sources but rather provide a robust estimation of \mathbf{A} .

Therefore, we will focus on assessing the performances of these algorithms with respect to the mixing matrix. More precisely, we propose to evaluate the accuracy of the different algorithms as well as their reliability, which is particularly relevant since BSS problems are non-convex.

The quantity $\Delta_A = \frac{\|\tilde{\mathbf{A}}^\dagger \mathbf{A} - \mathbf{I}\|_1}{n^2}$ is used as a global indicator of the mixing matrix estimation accuracy [14]. Additionally, for every simulation and for each algorithm, we record the number of runs for which \mathbf{A} has been *correctly* recovered (normalized to 1). The mixing matrix is said to be correctly recovered if, for every column of \mathbf{A} , the angle between the estimated and true i th column is smaller than 5° : $\arccos(\langle \tilde{\mathbf{A}}^i, \mathbf{A}^i \rangle) < 5^\circ$, see Appendix C.

4.4.1.2 Data Setting. The comparisons are first carried out on synthetic data in order to illustrate the impact of parameters such as the percentage of corrupted data or the number of observations with Monte Carlo simulations (48 simulations). The data are generated as follows:

- A total of 8 sources (unless otherwise stated) are drawn from a Bernoulli-Gaussian law whose activation rate is fixed to 5%, and the standard deviation of their amplitude σ_S to 100. The number of samples t is fixed to 4096.
- The mixing matrix is drawn according to a normal law with zero mean. The columns of \mathbf{A} are normalized to unit ℓ_2 norm.
- The outliers are generated so as to corrupt at random a low number of columns of \mathbf{X} . The activation of these columns is drawn according to a Bernoulli process with probability ρ , which fixes the average number of corrupted columns to ρt . The amplitude of the outliers is drawn at random from a Gaussian distribution with zero mean and standard deviation σ_O .
- The noise is generated according to a Gaussian distribution with zero mean. Its standard deviation is set to 0.1.

4.4.2 Simulations for synthetic data

4.4.2.1 Influence of the number of observations. We underlined in Section 4.1.3 that the separation of the sources contribution and the outliers is more challenging if m is close to n . This was further illustrated in Section 4.2.3. The ratio $\frac{m}{n}$ is therefore a crucial parameter in BSS, especially in the presence of outliers. In this paragraph, the data are composed of m observations (or bands in hyperspectral imaging). The amplitude of the outliers is fixed to $\sigma_O = 100$ for $n = m$ and then the amplitude ratio between the outliers and the source contribution is kept constant. The percentage of outliers is fixed to 10% with $\rho = 0.1$.

As shown in fig.4.11a, rAMCA tends to be less influenced by the number of observations. The results of all the methods (except the β -divergence minimization algorithm) are better if m is very large: the condition number of \mathbf{A} is smaller and the outliers can be better distinguished from the source contribution since the energy of the outliers lying in the subspace generated by \mathbf{A} is lower when m is large. In this regime, the low-rankness of the term $\mathbf{A}\mathbf{S}$ becomes a valid

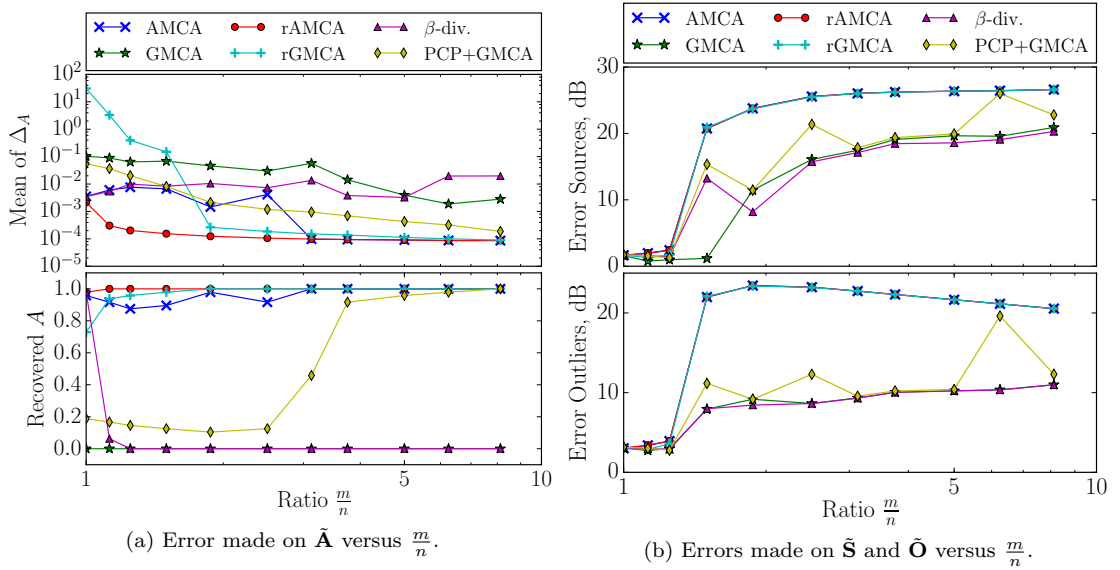


Figure 4.11: Influence of the number of observations on the estimations of \mathbf{A} , \mathbf{S} and \mathbf{O} .

assumption, which makes PCP more efficient [23].

The results are not strictly improved with an increasing number of measurements for the β -divergence algorithm. Since the β -divergence minimization algorithm has been designed for the determined case only ($m = n$), its application to the over-determined case ($m > n$) requires first, a dimension reduction step. This pre-processing step, which is performed by PCA, is also impacted by the presence of outliers and hampers the performances of this algorithm.

In order to further illustrate the impact of the ratio $\frac{m}{n}$, the errors $\frac{\|\mathbf{S}\|_2}{\|\mathbf{s}-\tilde{\mathbf{S}}\|_2}$ and $\frac{\|\mathbf{O}\|_2}{\|\mathbf{O}-\tilde{\mathbf{O}}\|_2}$ are displayed for a single example. Since the minimization of the β -divergence does not explicitly return \mathbf{O} and \mathbf{S} , we (re)-estimate \mathbf{O} and \mathbf{S} by minimizing (4.4) for fixed \mathbf{A} , the mixing matrices estimated by the different algorithms. A good separation of \mathbf{S} and \mathbf{O} is possible if $m \gg n$ because the outliers are less likely to lie in the span of \mathbf{A} ; this is clearly shown in fig.4.11b. Despite an accurate recovery of \mathbf{A} for rAMCA when m is small, the error made on the estimated outliers and sources is large fig.4.11b: the separation is not possible without any additional assumption on the sources and the outliers. Moreover, these errors decrease when the ratio $\frac{m}{n}$ increases whereas the error made on \mathbf{A} remains more stable: the separation benefits from enhanced estimation of \mathbf{A} as well as from a lower contribution of the outliers in the range of \mathbf{A} .

4.4.2.2 Influence of the condition number of \mathbf{A} . The condition number of the mixing matrix \mathbf{A} is a crucial parameter for BSS. Given that the columns of \mathbf{A} are normalized, a large condition number implies that the columns of \mathbf{A} are correlated. In multi/hyperspectral

unmixing it means that some of the components to be separated have quite similar spectral signatures: as so, they are difficult to differentiate.

The condition number of \mathbf{A} is even more important in robust BSS since it influences the behavior of the projected outliers in the sources domain. Indeed, the projected outliers $\mathbf{A}^\dagger \mathbf{O}$ have larger amplitudes, and as so are more detrimental. Indeed, the thresholding has a relatively smaller impact on these data.

In order to illustrate its impact on the unmixing, we will compare the performances of rAMCA, AMCA and GMCA with varying amplitude of the outliers, and varying condition number such as in fig.4.12. The number of sources and observations are fixed to 6, in the presence of 10% of outliers.

One can notice with fig.4.12 that for a given amplitude of the outliers, the methods (especially




Figure 4.12: Median over 20 runs of Δ_A for varying \mathbf{A} and amplitude of the outliers.

GMCA) are hampered with an increasing condition number of \mathbf{A} . Since GMCA is based on the MDP, it fails whenever the projected outliers are larger than the sources, see also Section 2.1.3 (hence, broadly, whenever $\frac{\sigma_S}{\sigma_O}$ is smaller than the condition number of \mathbf{A}). The AMCA algorithm is much robust than GMCA thanks to the weighting scheme, which is still able to detect and penalize the correlated entries/the outliers. However, AMCA may fail when the condition number (with respect to the amplitude of the outliers) is too large: first, the unmixing is itself more challenging, and second, even with the weighting scheme, it is costless to have broadly distributed projected outliers than “clustered” projected outliers when the condition number of \mathbf{A} is large. In particular, if two columns of \mathbf{A} are very close, it is generally costless to approximate them with only one column of the estimated $\tilde{\mathbf{A}}$, and orientate the other column of $\tilde{\mathbf{A}}$ in the direction of the large outliers samples. On the other hand, by further removing the outliers, rAMCA is much reliable.

4.4.2.3 Influence of the number of samples t . In the following experiment, we investigate the influence of the number of samples. In order to observe the impact of this data dimension on the combination PCP+GMCA, we consider that 6 sources are mixed into 30 observations (the low-rank assumption is valid), which are corrupted by $\rho_O = 10\%$ of active

outliers with $\sigma_O = 50$. We set σ_N to 0.1. The number of samples t varies according to the x -axis of fig.4.13.




Figure 4.13: Performance results of the methods versus the number of samples t .

As shown in fig.4.13, all the algorithms are less reliable if only few samples are available since the clustering aspect of the source contributions is not significant (an unmixing, even without outliers, is challenging if only few samples are available). Besides, all the strategies become more and more precise as the number of samples t increases.

Increasing the number of samples has several favorable effects on the unmixing: the number of samples available to unmix the sources becomes sufficient regardless of the presence of outliers, and the clustered aspect of \mathbf{AS} has a greater importance since there are more and more clustered samples in the term \mathbf{AS} but the outliers are still in general position (generating randomly several outliers in a same direction is quite unlikely).

The results would have been different if the proportion of corrupted samples in a given direction had been set constant from one value of t to another. For instance, if one resizes an multi/hyperspectral data cube, these proportions are kept constant, and for the largest image size, few but several outliers are in a same direction. There are some applications (most of the observations of physical processes) where the outliers are not strictly speaking in general position (e.g. a point source emission in the Planck data corresponds to more than only one pixel), but whose contributions are less structured/clustered than the one of the sources. The weighting scheme penalizes the less clustered solutions, and so, still returns \mathbf{A} . That is why AMCA and rAMCA requires less samples than the others methods to perform accurately fig.4.13.

In the following, the impact of two other parameters will be investigated: the percentage of corrupted data and their amplitudes. We will focus on the determined case which is more challenging. Since the low-rankness assumption makes no sense in the determined case, the algorithm PCP+GMCA will not be evaluated.

4.4.2.4 Influence of the amplitude of the outliers. In the following experiments, we consider that 10% of the data samples are corrupted with outliers. Fig.4.14a shows the behavior the algorithms when the amplitude of the outliers σ_O varies.

The figure 4.14a shows that the standard GMCA rapidly fails to correctly recover the mixing

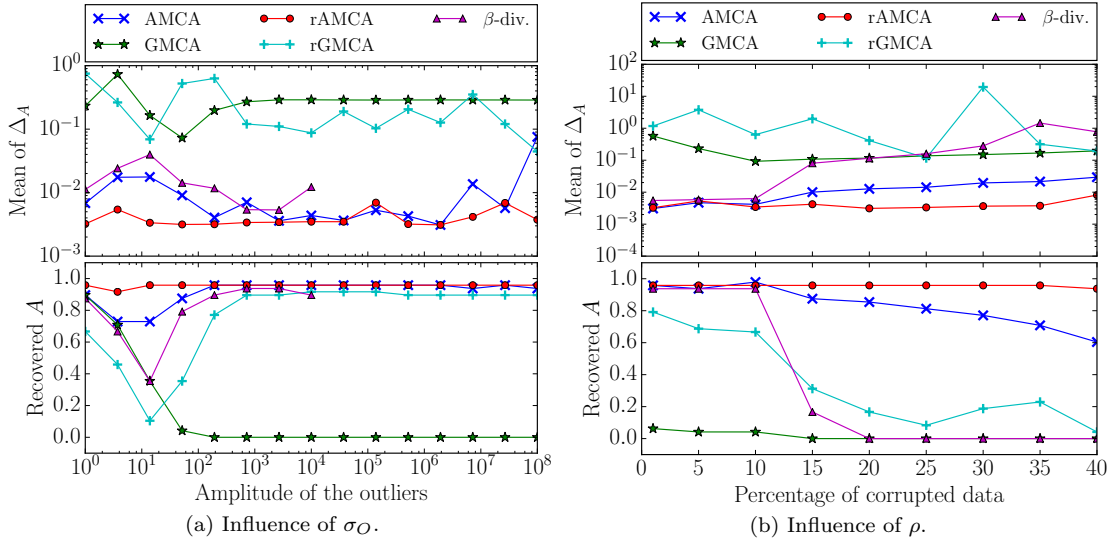


Figure 4.14: Influence of the amplitude and the activation rate of the outliers.

matrix when the amplitude of the outliers increases. In these experiments, the algorithms AMCA and β -divergence minimization algorithms provide very similar results. Interestingly, rAMCA tends to be the least impacted by the amplitude of the outliers, especially when their amplitude is of the order of the source's level or very large. When the amplitude of the outliers and the sources are close, the weighting schemes of AMCA and rGMCA are less effective at penalizing the outliers. Unlike AMCA, the rAMCA algorithm progressively removes a certain level of the outliers' component, which further enhances the separation performances.

4.4.2.5 Influence of the percentage of corrupted data. In this section, the amplitudes of the outliers σ_O is fixed to 100. The figure 4.14b shows the behavior of the BSS algorithms when the percentage of corrupted columns ρ varies according to the values of the x-axis.

As illustrated in fig.4.14b, the β -divergence algorithm is able to recover correctly the mixing matrix when the number of corrupted columns of \mathbf{X} is low (*i.e.* typically below 10%). The rGMCA algorithm is rapidly impacted by an increasing number of corrupted data. On the other hand, the AMCA-based algorithms are less influenced by the percentage of outliers. The rAMCA algorithm provides a significantly better estimate of the mixing matrix when the number

of outliers is larger than 10%.

4.4.3 Application to NMR spectral unmixing

In this section, we propose to compare the different algorithms in a more realistic setting: the separation of Nuclear Magnetic Resonance (NMR) spectra, similarly to Section 4.2.3. The presence of instrumental artifacts is very frequent in spectrometry and makes difficult the interpretation of the data. Such artifacts can be approximated by outliers contaminating entire columns of the data matrix, such as for the LC-MS data, Section 2.1.1.

Following Section 4.2.3, the sources are composed of 6 theoretical NMR spectra of the cholesterol, folic acid, adenosine, oleic acid, menthone and saccharose extracted from the SDBS database¹¹ with $t = 2048$ samples. These spectra are further convolved with a Laplacian kernel of varying width at half maximum (implementation from pyGMCA¹²), which models the resolution of the instrument, fig.4.15a. The set of artificially corrupted data samples is fixed to 10 blocks of 20 consecutive columns. Their amplitudes are drawn according to a Chi-distribution with 1 degree of freedom, and they are further convolved with the same kernel as the sources. The amplitude of the outliers is set so that the energy of each block of outliers corresponds to the average contribution of a source in the observations $\frac{\|\mathbf{O}\|_2}{10} = \frac{\|\mathbf{AS}\|_2}{n}$, fig.4.15c. In the following experiments, the data are made of 10 mixtures computed with a positive mixing matrix their entries are drawn from a Chi-law with 1 degree of freedom and then the columns are normalized) and corrupted also by the presence of the centered Gaussian noise with $\sigma_N = 0.1$.

Given that all the variables are non-negative, we will also compare AMCA and rAMCA with rNMF [58], whose code is online. This method exploits the low-rankness of \mathbf{AS} , the non-negativity and the “sum-to-one” constraint (that is, the amplitudes of each sample of \mathbf{S} sum to one) to differentiate between the low-rank subspace and the outliers. The “sum-to-one” constraint, which is not a valid assumption in this setting, is replaced by the constraint on the columns of \mathbf{A} , which are assumed to be normalized. We use the following inputs for rNMF: the ground truth \mathbf{A} , the projected sources $(\mathbf{A}^\dagger \mathbf{X})_+$ and the non-negative part of the corresponding residue.

The resulting sources admit a sparser distribution in the wavelet domain. Subsequently, the data are transformed with the undecimated wavelet transform [122] prior to applying the BSS algorithms, except for rNMF. Let us notice that a same wavelet transform is used for the outliers and the sources because they have a similar morphology in the present setting. In the previous experiments, we evaluated the separation performances of the algorithms in the case of exactly sparse signals. The NMR sources we consider in this section rather exhibit an approximately sparse distribution in the wavelet domain. We propose to evaluate the behavior of the robust BSS algorithms when both the sources and the outliers follow an approximate rather than exact sparse model. A simple way to evaluate the behavior of the algorithms with respect to the sparse

¹¹<http://sdb.sdb.aist.go.jp>

¹²<http://www.cosmostat.org/software/gmcalab/>




Figure 4.15: Top: illustration of one observation \mathbf{X}_i , without (left) and with outliers (right, corrupted entries are represented with the red dashed line). Bottom: estimated sources with rAMCA and PCP+GMCA for a width of the kernel of 6 (left) and right, performances of the different algorithms versus the width of the kernel (right).

model is to evaluate their performances when the width of the convolution kernel increases. Low width values will make the source model close to the exact sparse model while large values will provide approximately sparse sources.

Figure 4.15d displays the evolution of the mixing matrix criterion when the width of the convolution kernel varies. It is interesting to notice that the minimization of the β -divergence, PCP+GMCA, and the rNMF algorithms do not provide satisfactory separation results. This experience is particularly challenging for these methods since: the low-rank assumption is not valid, 'the sum-to-one' constraint necessary to the separation between \mathbf{AS} and \mathbf{O} for rNMF has been removed, and the outliers are less and less sparse as the width of the kernel increases. As well, the rGMCA provides good separation results when the width is low but it rapidly yields incorrect results when the width of the kernel increases. Indeed, let us recall that the outliers are also approximately sparse, which makes these separation scenarios close to the cases we investigated previously where the number of outliers is very large. This is typically the kind of settings where these methods tend to fail. The rAMCA and AMCA provide the most accurate estimates of the mixing. The discrepancy with respect to the other algorithms is particularly large when

the kernel has a large width. In this regime, the level of correlation between the sources increases, a phenomenon to which the AMCA algorithm is robust [13]. Last, one of the sources estimated by rAMCA and PCP+GMCA is displayed in fig.4.15b. Contrary to PCP+GMCA, the source is correctly recovered by rAMCA outside the support of \mathbf{O} because \mathbf{A} is correctly estimated by rAMCA. However, the leakages from the outliers into the sources estimated by rAMCA are still important: they come from the coarse scale of the wavelet coefficients, which is not sparse and for which we cannot differentiate the two contributions. Taking into account the non-negativity of the signals would limit these leakages, but necessitates the use of proximal algorithms, Section 3.1, if combined with sparsity in a transformed domain [113]. Nonetheless, the weighting scheme of rAMCA and AMCA is sufficient to obtain a robust estimation of \mathbf{A} .

4.5 Summary and future works

In this chapter, we introduce new algorithms for tackling BSS problems in the presence of outliers. The proposed rAMCA algorithm performs by estimating jointly the mixing matrix, the sources and the outliers. Inspired by the AMCA algorithm, it first provides a robust estimation of the sources and the mixing matrix. Additionally, it exploits the difference of structures of the outliers and the sources to provide a robust detection and estimation of the outliers based on their sparsity level in the source domain. Numerical experiments have been carried out on Monte-Carlo simulations with various experimental scenarios, which show that rAMCA yields a robust and reliable estimation of the mixing matrix. It provides the state-of-the-art separation results especially in the highly challenging determined case.

Future works and improvements.

Refinement step: The presented algorithms rGMCA and rAMCA only provide a good estimation of the mixing matrix. Even if a perfect separation between the sources and outliers is not always possible as observed in Section 4.1.3, a refinement step for separating the outliers from the source contribution would be of interest, for fixed \mathbf{A} . The separation between the outliers lying in the subspace orthogonal to the span of \mathbf{A} (in the over-determined case) and the source contribution would be at least improved.

Estimating the number of sources Estimating the number of sources is a topical issue in BSS [97], [141]. That is why, assuming the knowledge of the exact number of sources is a common assumption in BSS since most of the standard BSS methods are very sensitive to a mis-estimation of this number. At best, if the number of sources \tilde{n} (provided) is under-estimated, one can only expect to recover \tilde{n} of the initial n columns of \mathbf{A} and resulting mixed sources. On the other hand, if the number of sources \tilde{n} is over estimated, we can expect that n of the \tilde{n} estimated columns correspond to the initial columns of \mathbf{A} .

In this chapter, we have seen that AMCA-based algorithms are able to retrieve robustly the mixing matrix in the presence of outliers. In other words, they can estimate robustly the n most clustering directions. One can then draw a connection with the estimation of the number of sources: the outliers can be seen as less clustering sources contributions.

In order to investigate the influence of the hypothetical number of sources, we generate 10 sources, mixed into 16 noisy observations corrupted with 10% of outliers with $\sigma_O = 100$, and a Gaussian noise with $\sigma_N = 0.1$. Then, we use AMCA and rAMCA with varying number of estimated sources \tilde{n} (x -axis of fig.4.16). In cases where $\tilde{n} < n$, we check if the \tilde{n} recovered columns correspond to \tilde{n} of the n initial columns of \mathbf{A} . On the other hand, if $\tilde{n} > n$ we check if n among the \tilde{n} recovered columns correspond to the n initial columns of \mathbf{A} .




Figure 4.16: Performances of rAMCA and AMCA versus the number of looked-after sources in the presence of outliers

It can be observed in fig.4.16, that AMCA and rAMCA are robust to an underdetermined and overdetermined number of sources, except if the number \tilde{n} is much smaller than the initial n .

These preliminary results show the potential of AMCA-based methods to estimate the number of sources jointly with the mixing matrix and sources.

Chapter 5

Robust BSS in transformed domains

Summary

5.1 Separation between the outliers and the sources	108
5.1.1 Combining spectral and morphological diversity	108
5.1.2 Robust (non-blind) source separation as a sparse decomposition problem	109
5.2 Robust GMCA in transformed domains	113
5.2.1 Block Coordinate Minimization	113
5.2.2 A prox-linear implementation	114
5.2.3 Warm-up procedure	116
5.3 1D Monte-Carlo Simulations	120
5.3.1 Performance criteria and algorithms for comparison	120
5.3.2 1D Simulations	123
5.4 Application to simulated astrophysical data	129
5.4.1 Simulated data	129
5.4.2 Upgrades of tr-rGMCA	130
5.4.3 Experiments	131
5.5 Summary and future works	136

The methods rAMCA and rGMCA, introduced in the previous chapter, can robustly estimate the mixing matrix in the presence of outliers sharing the morphology of the sources. However, they cannot separate precisely the source contribution from the outliers.

In this chapter, we propose to exploit the difference of morphology/geometrical content between the outliers and the sources to separate precisely the two contributions. This difference of morphology, introduced in Section 1.2.2, is often encountered in imaging problems: striping lines due to malfunctions of captors have a different morphology than natural images

in multi/hyperspectral imaging or point-source emissions have a different geometry than the sought-after signals in the ESA-Planck mission, 2.1.2. By only assuming that the outliers and the sources have a different morphology, the proposed strategy coined tr-rGMCA (robust Generalized Morphological Component Analysis in transformed domains), is able to separate precisely the sources and the outliers, in a wide variety of problems, including in the challenging determined case ($n = m$).

We will start this chapter by focusing on the separation of the outliers from the source contribution for which we explain why the morphological diversity is a powerful assumption for robust BSS. Then, we introduce the tr-rGMCA problem, the associated algorithm and the strategies used for the automatic choice of the parameters. Last, the results of numerical experiments on 1D Monte-Carlo simulations and 2D simulated astrophysics data are displayed for the comparison of tr-rGMCA with standard robust BSS methods.

This chapter corresponds to [32] and its submitted extension.

5.1 Separation between the outliers and the sources

As pointed out in Section 2.2, robust blind source separation can merely be split into two distinct problems: i) the robust estimation of the mixing matrix \mathbf{A} from the data without considering outlier removal and ii) the exact or accurate separation between the outliers \mathbf{O} and the sources \mathbf{AS} . In this section, we discuss how the morphological diversity between the components can help tackling simultaneously these two problems.

5.1.1 Combining spectral and morphological diversity

In this section, we introduce an additional prior that helps differentiating between the sources and the outliers: the morphological diversity between \mathbf{O} and \mathbf{S} , previously introduced in Section 1.2.2 (for transformed domains). While spectral diversity refers to the relative distributions of the sources and the outliers in the column-space, morphological diversity deals with their relative distribution in the row-space.

In a large number of applications, the sources to be retrieved and the outliers share different morphologies, such as in Planck data fig.5.1. In this case, spurious points sources are the perfect example of column sparse outliers. These components are local singularities that are morphologically distinct from more diffuse astrophysical components. Therefore, building upon the concept of morphological diversity, we hereafter propose to reformulate robust BSS as a special case of a multichannel MCA problem (see Section 1.2.2). In the remaining of this chapter, we will make use of the following assumptions:

- **Morphological diversity between the sources and the outliers:** We assume that the




Figure 5.1: Simulated components of the ESA-Planck mission: synchrotron's map (one row of \mathbf{S}) (a) and observation of the point-source contamination at a given frequency (one row of \mathbf{O}).

sources are sparsely represented in the transformed domain or dictionary $\Phi_{\mathbf{S}} \in \mathbb{R}^{d_S \times t}$, whose number of atoms is such that $d_S \geq t$ and that the outliers have a sparse representation in $\Phi_{\mathbf{O}} \in \mathbb{R}^{d_O \times t}$, with $d_O \geq t$:

$$\mathbf{O}_j = \alpha_{\mathbf{O}_j} \Phi_{\mathbf{O}}, \forall j \in \{1, \dots, m\} \quad \text{and} \quad \mathbf{S}_i = \alpha_{\mathbf{S}_i} \Phi_{\mathbf{S}}, \forall i \in \{1, \dots, n\},$$

where $\{\alpha_{\mathbf{O}_j}\}_{j=1, \dots, m}$ and $\{\alpha_{\mathbf{S}_i}\}_{i=1, \dots, n}$ are composed of few significant samples. A toy example is provided in fig.5.2. It highlights the benefits of exploiting the morphological diversity: in $\Phi_{\mathbf{S}}$, the outlier contribution is broadly distributed with a very small amplitude fig.5.2d,5.2f, whereas in $\Phi_{\mathbf{O}}$, they can be easily detected fig.5.2a,5.2c (and reciprocally for the sources samples).

- **Sparse modeling of the outliers:** We also consider that the sparse representations of the outliers corrupt entirely some columns and are broadly distributed in all the directions. For this purpose, we will assume that $\mathbf{O}\Phi_{\mathbf{O}}^T$ is column sparse such as in fig.5.2.

5.1.2 Robust (non-blind) source separation as a sparse decomposition problem

A special case of sparse decomposition in an overcomplete dictionary. The sources' contribution $\mathbf{L} = \mathbf{A}\mathbf{S}$ to the data \mathbf{X} is sparsely represented in the multichannel dictionary: $\mathbf{A} \otimes \Phi_{\mathbf{S}}$, whose atoms are composed of tensor products between the columns of \mathbf{A} and the atoms of $\Phi_{\mathbf{S}}$.

Similarly, the rows of the outlier matrix \mathbf{O} are assumed to be sparse in some dictionary $\Phi_{\mathbf{O}}$ so that $\mathbf{O} = \alpha_{\mathbf{O}}\Phi_{\mathbf{O}}$, where the coefficients $\alpha_{\mathbf{O}}$ are column sparse. Let \mathbf{O}_D be the submatrix made




Figure 5.2: Two sources sparse in DCT are mixed into three observations, corrupted with sparse outliers. (a): scatter plot of the two first rows of \mathbf{O} in $\Phi_{\mathbf{O}}$, (d): scatter plot the same rows of \mathbf{O} in $\Phi_{\mathbf{S}}$, (b): scatter plot of the first two sources in $\Phi_{\mathbf{O}}$, (e): scatter plot of the same sources in $\Phi_{\mathbf{S}}$, (c): scatter plot of the two first corrupted observations in $\Phi_{\mathbf{O}}$ and last, (f): scatter plot of the same observations in $\Phi_{\mathbf{S}}$. The source contribution is represented with the blue dots and the outliers with the red stars.

of the normalized non-zero columns of $\alpha_{\mathbf{O}}$, built so that the k th non-zero column of $\alpha_{\mathbf{O}}$ at the position t equals $\mathbf{O}_D^k = \frac{\alpha_{\mathbf{O}}^t}{\|\alpha_{\mathbf{O}}^t\|_2}$. We then denote $\alpha_{\mathbf{O}'}$, the expansion coefficients of \mathbf{O} in $\mathbf{O}_D \otimes \Phi_{\mathbf{O}}$, such that $\mathbf{O}_D \alpha_{\mathbf{O}'} = \alpha_{\mathbf{O}}$. The matrix $\alpha_{\mathbf{O}'}$ is then column and row sparse and $\|\alpha_{\mathbf{O}'}\|_1 = \|\alpha_{\mathbf{O}}\|_{2,1}$. With this parameterization, the outliers are sparsely represented in the multichannel dictionary $\mathbf{O}_D \otimes \Phi_{\mathbf{O}}$.

The observations are consequently sparsely represented in the multichannel dictionary $\mathbf{D} = [\mathbf{A} \otimes \Phi_{\mathbf{S}}, \mathbf{O}_D \otimes \Phi_{\mathbf{O}}]$:

$$\mathbf{X} = \begin{bmatrix} \mathbf{A} & \mathbf{O}_D \end{bmatrix} \begin{bmatrix} \alpha_{\mathbf{S}} & 0 \\ 0 & \alpha_{\mathbf{O}'} \end{bmatrix} \begin{bmatrix} \Phi_{\mathbf{S}} \\ \Phi_{\mathbf{O}} \end{bmatrix}.$$

Assuming that \mathbf{A} and \mathbf{O}_D are known, estimating the sources \mathbf{S} and the outliers \mathbf{O} from \mathbf{X} boils down to tackling a sparse decomposition problem in the overcomplete multichannel dictionary \mathbf{D} , similarly to Section 4.1.3. In the following, we also make use of the mutual coherence of the dictionary to provide a deeper insight into the proposed robust component separation.

Assuming that the components are K -sparse in \mathbf{D} with $K = \|\alpha_{\mathbf{S}}\|_0 + \|\alpha_{\mathbf{O}'}\|_0$, a sufficient condition for the identifiability of $\alpha_{\mathbf{S}}$ and $\alpha_{\mathbf{O}'}$ [47] is given by (see [7] for a more precise recovery condition):

$$K < \frac{1}{2} \left(1 + \frac{1}{\mu_{\mathbf{D}}} \right),$$

where the mutual coherence $\mu_{\mathbf{D}}$ of \mathbf{D} is defined as $\mu_{\mathbf{D}} = \max_{i,j} |\langle \mathbf{d}_i, \mathbf{d}_j \rangle|$, where \mathbf{d}_i stands for an atom of the multichannel dictionary \mathbf{D} (*i.e.* multichannel atoms are composed of tensor products of atoms from the spectral dictionaries and morphological dictionaries). Furthermore, the same condition also guarantees that the $\alpha_{\mathbf{S}}$ and $\alpha_{\mathbf{O}'}$ can be recovered by solving the following basis pursuit problem [47]:

$$\underset{\alpha_{\mathbf{O}'}, \alpha_{\mathbf{S}}}{\operatorname{argmin}} \|\alpha_{\mathbf{O}'}\|_1 + \|\alpha_{\mathbf{S}}\|_1 \text{ s.t. } \mathbf{X} = \mathbf{A}\alpha_{\mathbf{S}}\Phi_{\mathbf{S}} + \mathbf{O}_D\alpha_{\mathbf{O}'}\Phi_{\mathbf{O}}. \quad (5.1)$$

The term of interest in the above recovery condition is the mutual coherence $\mu_{\mathbf{D}}$, which is equal, in this specific case, to:

$$\max \left(\max_{(i,p) \neq (j,q)} |\langle \mathbf{A}^i, \mathbf{A}^j \rangle| |\langle \Phi_{\mathbf{S}}^p, \Phi_{\mathbf{S}}^q \rangle|, \max_{(m,u) \neq (n,v)} |\langle \mathbf{O}_D^m, \mathbf{O}_D^n \rangle| |\langle \Phi_{\mathbf{O}}^u, \Phi_{\mathbf{O}}^v \rangle|, \max_{(l,c),(k,d)} |\langle \mathbf{A}^l, \mathbf{O}_D^k \rangle| |\langle \Phi_{\mathbf{S}}^c, \Phi_{\mathbf{O}}^d \rangle| \right), \quad (5.2)$$

where the columns of \mathbf{A} , $\Phi_{\mathbf{O}}$ and $\Phi_{\mathbf{S}}$ are normalized to have unit ℓ_2 norm. In this expression, the cross-terms $\max_{(l,c),(k,d)} |\langle \mathbf{A}^l, \mathbf{O}_D^k \rangle| |\langle \Phi_{\mathbf{S}}^c, \Phi_{\mathbf{O}}^d \rangle|$ are the most relevant to discriminate the outliers and the source contribution and provide a different way to re-interpret robust (non-blind) source separation:

- **Spectral diversity or rPCA regime, Sections 2.2.1 and 4.1.3** In case the outliers and sources share a same morphology, (see Section 4.1.3 for a more complete discussion), only the cross-term between the mixing matrix and the outliers columns $\max_{(l,c)} |\langle \mathbf{A}^l, \mathbf{O}_D^k \rangle|$ is relevant for the separation. In the framework of rPCA, whenever the source contribution $\mathbf{A}\mathbf{S}$ has low rank, $\max_{(l,k)} |\langle \mathbf{A}^l, \mathbf{O}_D^k \rangle|$ vanishes when \mathbf{O}_D lies in the subspace that is orthogonal to the span of \mathbf{A} , which naturally ensures the identifiability of both the sources and the outliers. In the general case, assuming that \mathbf{O} has independently and sparsely distributed entries and that \mathbf{A} is broadly distributed such as in the setting of rPCA, leads to spectral dictionaries \mathbf{A} and \mathbf{O}_D with low coherence. This is precisely in this regime that rPCA can ensure the identifiability of the components.
- **Morphological diversity or MCA regime, Section 1.2.2:** When the low-rankness of the observations is not a valid assumption or when the span of \mathbf{A} and \mathbf{O}_D are not incoherent, such as in the determined case, only the morphological diversity can help identifying the components. In that case, the dictionaries $\Phi_{\mathbf{O}}$ and $\Phi_{\mathbf{S}}$ are assumed to be incoherent, which makes $\max_{(c,d)} |\langle \Phi_{\mathbf{S}}^c, \Phi_{\mathbf{O}}^d \rangle|$ the relevant term for the separation. This is precisely in this regime that MCA can ensure the separation between components that can only be identified thanks to their difference of morphologies. In this case, robust component separation can be solved in the determined case.
- **Morpho/Spectral diversity:** In the general case, both the spectral and morphological dic-

Estimation	Diversity	Regime	Methods	Advantages	Weaknesses
AS and O	Morphological	$m \geq n$	MCA [54]	No assumption on the collinearity of O and A .	AS should be sparse in $\Phi_{\mathbf{S}}$. The spectral structure may not be preserved.
	Spectral	$m \gg n$	PCP [23], or refinements e.g. [142], [105]	Proven separability.	$\mathbf{O}\Phi_{\mathbf{O}}^T$ column and row sparse.
			OP [135]	$\mathbf{O}\Phi_{\mathbf{O}}^T$ column sparse.	No identifiability of O .
A , S and O	Spectral	$m \gg n$	rNMF [58]	Well adapted for hyperspectral unmixing.	Non-negativity, sources samples in the simplex and presence of almost pure-pixels.
		$m \geq m$	rAMCA [31]	Estimation of A	No identifiability of O .
	Morphological & spectral	$m \geq n$	tr-rGMCA	Identifiability in all regimes	

Table 5.1: Strategies able to separate **AS** and **O**.

tionaries are incoherent, the relevant coherence term is the product $\max_{(l,c),(k,d)} |\langle \mathbf{A}^l, \mathbf{O}_D^k \rangle| |\langle \Phi_{\mathbf{S}}^c, \Phi_{\mathbf{O}}^d \rangle|$. In this regime, robust component separation benefits from incoherence of both the morphological and spectral dictionaries: $\max_{(l,c),(k,d)} |\langle \mathbf{A}^l, \mathbf{O}_D^k \rangle| |\langle \Phi_{\mathbf{S}}^c, \Phi_{\mathbf{O}}^d \rangle| \leq \min(\max_{(l,k)} |\langle \mathbf{A}^l, \mathbf{O}_D^k \rangle|, \max_{(c,d)} |\langle \Phi_{\mathbf{S}}^c, \Phi_{\mathbf{O}}^d \rangle|)$. This is expected to greatly improve the accuracy of the separation. For instance, in this regime, column-sparse outliers can be identified while methods that only make use of the spectral diversity like Outliers Pursuit [135] can only ensure the identification of the support of the outliers and not their amplitude.

In the framework of robust *blind* source separation, the spectral dictionary $[\mathbf{A} \ \mathbf{O}_D]$ is not known and has also to be learned. For this purpose, we describe in the next section a novel algorithm coined tr-rGMCA that makes use of both spectral and morphological diversities to estimate jointly **A**, **S** and **O** given the two dictionaries $\Phi_{\mathbf{S}}$ and $\Phi_{\mathbf{O}}$ so as to build upon the spectral and the morphological diversities between the components. Based on whether they rely on spectral or morphological diversity, currently available blind separation strategies are summarized in table.5.1, which completes the presentation in 2.3.

5.2 Robust GMCA in transformed domains

In this section, we introduce the tr-rGMCA (rGMCA in transformed domains) algorithm that builds upon both spectral and morphological diversities to estimate simultaneously \mathbf{A} , \mathbf{S} and \mathbf{O} . The tr-rGMCA algorithm performs by minimizing a cost function whose elements are based on the following assumptions:

- Data fidelity term: The data are assumed to be described by the linear mixture model $\mathbf{X} = \mathbf{AS} + \mathbf{O} + \mathbf{N}$. The squared Frobenius norm $\|\mathbf{X} - \mathbf{AS} - \mathbf{O}\|_2^2$ is used as fidelity term to measure a discrepancy between the data and the model. This distance is well suited to account for the additive Gaussian noise \mathbf{N} that usually contaminates the data.
- Penalty term for the sources: The compressibility of \mathbf{S} in $\Phi_{\mathbf{S}}$ is enforced with a weighted ℓ_1 norm of the expansion coefficients of \mathbf{S} : $\|\Lambda \odot \mathbf{S}\Phi_{\mathbf{S}}^T\|_1$. The weighting matrix $\Lambda \in \mathbf{R}^{n \times t}$ includes both the regularization parameters and the weights defined in standard re-weighting ℓ_1 penalization [26], Section 3.3.5.
- Penalty term for the outliers: The sparsity pattern of \mathbf{O} is enforced in the cost function using the composite $\ell_{2,1}$ norm [58, 83]: $\|\Upsilon \odot \mathbf{O}\Phi_{\mathbf{O}}^T\|_{2,1}$. Again the matrix $\Upsilon \in \mathbf{R}^{1 \times t}$ contains the regularization parameters as well as weights in the sense of re-weighting $\ell_{2,1}$. The morphological diversity assumption implies that $\Phi_{\mathbf{O}}$ and $\Phi_{\mathbf{S}}$ are somehow incoherent.
- Scaling indeterminacy: In order to control the scaling indeterminacy between \mathbf{A} and \mathbf{S} , the columns of \mathbf{A} have an energy bounded by 1. The columns of \mathbf{A} are constrained to lie in the ℓ_2 ball with unit radius: $\chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A})$.

Therefore, the algorithm tr-rGMCA estimates jointly \mathbf{A} , \mathbf{O} and \mathbf{S} by minimizing the following cost function:

$$\underset{\mathbf{A}, \mathbf{S}, \mathbf{O}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS} - \mathbf{O}\|_2^2 + \|\Lambda \odot \mathbf{S}\Phi_{\mathbf{S}}^T\|_1 + \|\Upsilon \odot \mathbf{O}\Phi_{\mathbf{O}}^T\|_{2,1} + \chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}). \quad (5.3)$$

The resulting cost function is a multi-convex *non-smooth* optimization problem: it is globally non-convex but subproblems with all variables fixed except one are convex. Hence, it is customary to optimize this type of cost function using Block Coordinate methods, Section 3.2.

5.2.1 Block Coordinate Minimization

Updating each block \mathbf{A} , \mathbf{S} and \mathbf{O} alternately at each iteration can be carried out in different ways. The subproblems associated with the minimization of the problem 5.3, assuming all the

blocks are fixed but one are given by:

$$\mathbf{P}_A : \quad \underset{\mathbf{A}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S} - \mathbf{O}\|_2^2 + \chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}). \quad (5.4)$$

$$\mathbf{P}_S : \quad \underset{\mathbf{S}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S} - \mathbf{O}\|_2^2 + \|\Lambda \odot \mathbf{S}\Phi_S^T\|_1. \quad (5.5)$$

$$\mathbf{P}_O : \quad \underset{\mathbf{O}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S} - \mathbf{O}\|_2^2 + \|\Upsilon \odot \mathbf{O}\Phi_O^T\|_{2,1}. \quad (5.6)$$

These three problems can be written as $\text{argmin}_{\mathbf{Y}} f_Y(\mathbf{Y}) + g_Y(\mathbf{Y})$, where $f_Y(\cdot)$ is related to a differentiable data-fidelity term (whose gradient denoted as ∇f_Y is L_Y -Lipschitz) and $g_Y(\cdot)$ is the proximal regularization associated with the component \mathbf{Y} (see Section 3.1). In general, they do not admit a closed-form solution and therefore require resorting to iterative minimization procedures such the Proximal Forward-Backward Splitting algorithm (FB) [38], [107] presented in Section 3.1.2. Whether it is based on BCD or PALM, it is possible to design a minimizer that provably converges to a local stationary point of the problem 5.3. In this context, either the BCD or the PALM algorithm can be chosen. However, BCD, introduced in Section 3.2.3, yields a computationally intensive minimization strategy. In the sequel, we therefore opted for the prox-linear approach, which is at the origin of the PALM algorithm presented in Section 3.2.4.

5.2.2 A prox-linear implementation

In the framework of PALM, each component is updated with one proximal gradient step eq.5.7 at the k th iteration (Section 3.2.4):

$$\tilde{\mathbf{Y}}^{(k)} \leftarrow \underset{\frac{1}{L_Y} g_Y}{\text{prox}} \left(\tilde{\mathbf{Y}}^{(k-1)} - \frac{1}{L_Y} \nabla f_Y(\tilde{\mathbf{Y}}^{(k-1)}) \right). \quad (5.7)$$

From this generic update, the three steps that compose the tr-rGMCA algorithms are described as follows:

- **Update of the sources.** Assuming Φ_S is orthonormal, the proximal operator of the function $\mathbf{S} \mapsto \|\Lambda \mathbf{S} \Phi_S^T\|_1$ is exactly $\mathbf{S} \mapsto \mathcal{S}_\Lambda(\mathbf{S} \Phi_S^T) \Phi_S$. Therefore, at iteration k of the PALM procedure, the update of $\tilde{\mathbf{S}}^{(k)}$ is given by:

$$\tilde{\mathbf{S}}^{(k+1)} \leftarrow \mathcal{S}_{\frac{\Lambda}{L_S}} \left(\left(\tilde{\mathbf{S}}^{(k)} + \frac{1}{L_S} \tilde{\mathbf{A}}^{(k)T} \left(\mathbf{X} - \tilde{\mathbf{A}}^{(k)} \tilde{\mathbf{S}}^{(k)} - \tilde{\mathbf{O}}^{(k)} \right) \right) \Phi_S^T \right) \Phi_S, \quad (5.8)$$

where the step size L_S is chosen to be equal to the Lipschitz constant of the gradient, *i.e.* the maximal eigenvalue of $\tilde{\mathbf{A}}^{(k)T} \tilde{\mathbf{A}}^{(k)}$.

When Φ_S is not orthonormal, the proximal operator of the function $\mathbf{S} \mapsto \|\Lambda \mathbf{S} \Phi_S^T\|_1$ does not admit a closed form. However, in the next experiments, the dictionaries used of Φ_S will be tight frames (*e.g.* undecimated wavelets) whose Gram matrix is close to the identity matrix. In that specific case, the update (5.8) provides a good approximation for the proximal oper-

ator.

- **Update of the outliers.** Assuming $\Phi_{\mathbf{O}}$ is orthonormal, the update of the outliers at the k th iteration of the PALM procedure is given by:

$$\begin{aligned} \tilde{\mathbf{O}}^{(k+1)} &\leftarrow \tilde{\alpha}_{\mathbf{O}} \Phi_{\mathbf{O}} \text{ where, } \forall j = 1..t \text{ and } \forall i = 1..n : \\ \tilde{\alpha}_{\mathbf{O}}^j &\leftarrow \left(\left((\mathbf{X} - \tilde{\mathbf{A}}^{(k)} \tilde{\mathbf{S}}^{(k)}) \Phi_{\mathbf{O}}^T \right)_i^j \times \max \left(0, 1 - \frac{\Upsilon_i^j}{\left\| \left((\mathbf{X} - \tilde{\mathbf{A}}^{(k)} \tilde{\mathbf{S}}^{(k)}) \Phi_{\mathbf{O}}^T \right)_i^j \right\|_2} \right) \right). \end{aligned} \quad (5.9)$$

In contrast with \mathbf{S} , the proximal gradient step eq.5.9 exactly solves $\mathbf{P}_{\mathbf{O}}$.

Besides, in this chapter, we assume that the outliers corrupt entirely few columns of the observations in their associated transformed domain. However, it would be straightforward to account for row and column sparse outliers in $\Phi_{\mathbf{O}}$ by replacing the $\ell_{2,1}$ norm with the ℓ_1 norm. In this case, (5.9) is simply replaced by: $\tilde{\mathbf{O}} \leftarrow \mathcal{S}_{\Upsilon}((\mathbf{X} - \mathbf{A}\mathbf{S}) \Phi_{\mathbf{O}}^T) \Phi_{\mathbf{O}}$.

- **Update of the mixing matrix.** The proximal gradient step for \mathbf{A} is two step: ((a) corresponds to the gradient step, and (b) to the proximal operator of the characteristic function):

$$\begin{aligned} (a) \quad \tilde{\mathbf{A}}^{(k+1)} &\leftarrow \tilde{\mathbf{A}}^{(k)} + \frac{1}{L_A} (\mathbf{X} - \tilde{\mathbf{A}}^{(k)} \tilde{\mathbf{S}}^{(k)} - \tilde{\mathbf{O}}^{(k)}) \tilde{\mathbf{S}}^{(k)T}, \\ (b) \quad \tilde{\mathbf{A}}^{i(k+1)} &\leftarrow \frac{\tilde{\mathbf{A}}^{i(k+1)}}{\max(1, \|\tilde{\mathbf{A}}^{i(k+1)}\|_2)}, \forall i = 1..n, \end{aligned}$$

where L_A is chosen to be equal to the Lipschitz constant of the gradient, *i.e.* the maximal eigenvalue of $\tilde{\mathbf{S}}^{(k)} \tilde{\mathbf{S}}^{(k)T}$.

Algorithm. The prox-linear minimization of (5.3) can be found in Alg.16.

Limitations of the standard block coordinate minimizers. The proposed prox-linear implementation is sensitive to the setting of the parameters and the initialization, which makes the joint estimation of the regularization parameters and the components highly challenging. As introduced in Section 3.3, the additional heuristics deployed with GMCA (or its extensions AMCA, rGMCA and rAMCA, Chapter 4) permit a robust estimation of the components and regularization parameters. If these heuristics, which are further detailed in Section 5.2.3, yield more robust minimization procedures, they lack provable convergence. Therefore the global optimization strategy used in the tr-rGMCA algorithm is composed of two successive steps:

- **The warm-up step:** a solution of the problem in Equation (5.3) is approximated using a PALS-based algorithm with heuristics. This first step, which is described in Section 5.2.3, aims at

providing a robust first guess of the components for the next, provably convergent, step.

- The refinement step: The goal of this stage, which we described in Alg.16, is to provide a local stationary point of (5.3). The associated regularization parameters are based upon the values of the components returned in the warm-up step, with 5.12.

This two-step strategy was proposed for the basic GMCA in Chapter 3.3.5.

Procedure 16 PALM

```

1: procedure PALM( $\mathbf{X}, \tilde{\mathbf{S}}, \tilde{\mathbf{O}}, \tilde{\mathbf{A}}, \Lambda, \Upsilon$ )
2:   Set  $\tilde{\mathbf{S}}^{(0)} \leftarrow \tilde{\mathbf{S}}, \tilde{\mathbf{O}}^0 \leftarrow \tilde{\mathbf{O}}$  and  $\tilde{\mathbf{A}}^0 \leftarrow \tilde{\mathbf{A}}$ 
3:   while  $p < P$  do
4:     Compute  $L_A$ 
5:      $\tilde{\mathbf{A}}^{(p+1)} \leftarrow \tilde{\mathbf{A}}^{(p)} + \frac{1}{L_A} (\mathbf{X} - \tilde{\mathbf{A}}^{(p)} \tilde{\mathbf{S}}^{(p)} - \tilde{\mathbf{O}}^{(p)}) \tilde{\mathbf{S}}^{(p)T}$ 
6:      $\tilde{\mathbf{A}}^{i(p+1)} \leftarrow \frac{\tilde{\mathbf{A}}^{i(p+1)}}{\max(1, \|\tilde{\mathbf{A}}^{i(p+1)}\|_2)}, \forall i = 1..n$ 
7:     Compute  $L_S$ 
8:      $\tilde{\mathbf{S}}^{(p+1)} \leftarrow \mathcal{S}_{\frac{\Lambda}{L_S}} \left( \left( \tilde{\mathbf{S}}^{(p)} + \frac{1}{L_S} \tilde{\mathbf{A}}^{(p+1)T} (\mathbf{X} - \tilde{\mathbf{A}}^{(p+1)} \tilde{\mathbf{S}}^{(p)} - \tilde{\mathbf{O}}^{(p)}) \right) \Phi_{\tilde{\mathbf{S}}}^T \right) \Phi_{\tilde{\mathbf{S}}}$ 
9:      $\tilde{\mathbf{O}}^{(p+1)} \leftarrow \alpha_{\tilde{\mathbf{O}}^{(p+1)}} \Phi_{\tilde{\mathbf{O}}}$  where  $\forall j = 1..t$  and  $\forall i = 1..m$  :
10:     $\alpha_{\tilde{\mathbf{O}}^{(p+1)}}^j \leftarrow \left( \left( (\mathbf{X} - \tilde{\mathbf{A}}^{(p+1)} \tilde{\mathbf{S}}^{(p+1)}) \Phi_{\tilde{\mathbf{O}}}^T \right)_i^j \times \max \left( 0, 1 - \frac{\Upsilon_i^j}{\|((\mathbf{X} - \tilde{\mathbf{A}}^{(p+1)} \tilde{\mathbf{S}}^{(p+1)}) \Phi_{\tilde{\mathbf{O}}}^T)_i^j\|_2} \right) \right)$ 
  return  $\tilde{\mathbf{S}}^{(P-1)}, \tilde{\mathbf{A}}^{(P-1)}, \tilde{\mathbf{O}}^{(P-1)}$ 

```

5.2.3 Warm-up procedure

In this section, we describe the so-called 'warm-up' stage of the tr-rGMCA algorithm. This procedure aims at providing an approximated solution of the problem (5.3) as well as robustly determining the regularization parameters (similarly to the rAMCA algorithm, in Chapter 4). The proposed strategy builds upon an appropriate choice of the variables to be updated based on either morphological or spectral diversity, that leads to the following PALS procedure:

- **Joint estimation of \mathbf{O} and \mathbf{S} based on morphological diversity.** Jointly estimating \mathbf{O} and \mathbf{S} for fixed \mathbf{A} amounts to solving the following convex optimization problem:

$$\underset{\mathbf{S}, \mathbf{O}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS} - \mathbf{O}\|_2^2 + \|\Lambda \odot \mathbf{S} \Phi_{\tilde{\mathbf{S}}}^T\|_1 + \|\Upsilon \odot \mathbf{O} \Phi_{\tilde{\mathbf{O}}}^T\|_{2,1}. \quad (5.10)$$

which we previously interpreted as some multichannel extension of Morphological Component Analysis. This step, which is detailed in Section 5.2.3.1, essentially exploits the morphological diversity between the outliers and the sources. Besides, we can extend the thresholding strategy proposed for MCA to distinguish between the two contributions, without risking a modification of \mathbf{A} .

- **Joint estimation of \mathbf{A} and \mathbf{S} based on spectral diversity.** Updating \mathbf{A} and \mathbf{S} boils down to tackling the sparse BSS problem from the residual term $\mathbf{X} - \mathbf{O}$:

$$\underset{\mathbf{S}, \mathbf{A}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS} - \mathbf{O}\|_2^2 + \|\Lambda \odot \mathbf{S} \Phi_{\mathbf{S}}^T\|_1 + \chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}). \quad (5.11)$$

While being non-convex, algorithms like GMCA, Section 3.3.5, or AMCA, Section 4.3, provide efficient approximate minimization that yields fair estimates of the tuple $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$. This stage is described in Section 5.2.3.2.

The warm-up procedure alternates between these two problems to minimize (5.3), such as presented in Alg.17. The combination of these two subproblems fully exploits the two types of diversity between the contributions. Besides, the independent reassessment of the sources \mathbf{S} from one subproblem to another limits the propagation of errors and destabilizes local minima, improving the robustness of this initializing scheme.

As it will be described in the remaining of this subsection, the warm-up procedure involves key heuristics that rely on particular parameter strategies and approximations, which are made to fasten the process and improve its robustness with respect to the initialization, noise residual and parameters setting (cf. Section 3.3).

In the numerical experiment section 5.3.2, we provide a comparison between the performances of the warm-up step alone, the refinement (PALM-based) step alone, and the combination of both (tr-rGMCA), showing the robustness of the warm-up procedure as well as the benefit in term of accuracy for using the refinement step.

5.2.3.1 Estimating \mathbf{O} and \mathbf{S} using the morphological diversity. For fixed \mathbf{A} , the outliers \mathbf{O} and the sources \mathbf{S} are the solutions of the problem (5.10). Since, for fixed sources, updating the outliers allows a closed-form expression, we opted for the BCD strategy that alternates between estimations of \mathbf{O} and \mathbf{S} :¹

- **Updating the sources.** The estimation of \mathbf{S} is given by \mathbf{P}_S (5.5). As stated in Section 3.1, \mathbf{P}_S can be solved with ISTA 3.1.2. We point out that in practice, the accelerated FISTA [8] is preferred.

¹The joint estimation of \mathbf{S} and \mathbf{O} can be highly fastened by using a PALS procedure, in the spirit of MCA. However, whenever \mathbf{A} does not have a small condition number or $\Phi_{\mathbf{O}}$ and $\Phi_{\mathbf{S}}$ are not highly incoherent, this process yields an inaccurate separation, which is then very likely to further hamper the unmixing: the errors propagate.

Procedure 17 WarmUp Procedure

```

1: procedure WARMUP( $\mathbf{X}, \tilde{\mathbf{A}}, \tilde{\mathbf{S}}, \tilde{\mathbf{O}}$ )
2:   Initialize  $\tilde{\mathbf{S}}^{(k=0)} \leftarrow \tilde{\mathbf{S}}, \tilde{\mathbf{A}}^{(k=0)} \leftarrow \tilde{\mathbf{A}}$  and  $\tilde{\mathbf{O}}^{(k=0)} \leftarrow \tilde{\mathbf{O}}$ .
3:   while  $k < K$  do
4:     Set  $\alpha_{\mathbf{S}}^{(i=1,k)} \leftarrow \tilde{\mathbf{S}}^{(k-1)} \Phi_{\mathbf{S}}^T, \tilde{\mathbf{A}}^{(i=1,k)} \leftarrow \tilde{\mathbf{A}}^{(k-1)}$ , and
5:      $\alpha_{\mathbf{X-O}} \leftarrow (\mathbf{X} - \tilde{\mathbf{O}}^{(k-1)}) \Phi_{\mathbf{S}}^T$ 
6:     while  $i < I$  do ▷ Joint estimation of  $\mathbf{A}$  and  $\mathbf{S}$ 
7:        $\alpha_{\mathbf{S}}^{(i,k)} \leftarrow \mathcal{S}_{\Lambda}(\tilde{\mathbf{A}}^{(i-1,k)\dagger} \alpha_{(\mathbf{X-O})})$ 
8:        $\tilde{\mathbf{A}}^{(i,k)} \leftarrow \alpha_{(\mathbf{X-O})} \alpha_{\mathbf{S}}^{(i=1,k)\dagger}$ 
9:       Decrease  $\Lambda$ 
10:    Set  $\tilde{\mathbf{S}}^{(k)} \leftarrow \alpha_{\mathbf{S}}^{(i=I-1,k)} \Phi_{\mathbf{S}}$  and  $\tilde{\mathbf{A}}^{(k)} \leftarrow \tilde{\mathbf{A}}^{(I,k)}$ 
11:    Set  $\tilde{\mathbf{S}}^{(\ell=0,j=0,k)} \leftarrow \tilde{\mathbf{S}}^{(k)}$  and  $\tilde{\mathbf{O}}^{(\ell=0,j=0,k)} \leftarrow \tilde{\mathbf{O}}^{(k-1)}$ 
12:    for  $\ell < L$  do ▷ Reweighting Procedure
13:      while  $j < J$  do ▷ Joint estimation of  $\mathbf{S}$  and  $\mathbf{O}$ 
14:        Update  $\tilde{\mathbf{S}}^{(\ell,j,k)}$  with FISTA using the proximal gradient step (5.8)
15:        Update  $\tilde{\mathbf{O}}^{(\ell,j,k)}$  with the closed form (5.9)
16:        Update  $\Lambda$  and  $\Upsilon$  for the reweighting procedure according to (5.12)
17:        Set  $\tilde{\mathbf{S}}^{(\ell+1,0,k)} \leftarrow \tilde{\mathbf{S}}^{(\ell,J,k)}, \tilde{\mathbf{O}}^{(\ell+1,0,k)} \leftarrow \tilde{\mathbf{O}}^{(\ell,J,k)}$ 
    return  $\tilde{\mathbf{A}}^{(K)}, \tilde{\mathbf{S}}^{(L,J,K)}, \tilde{\mathbf{O}}^{(L,J,K)}$ .

```

- **Updating the outliers.** The estimation of \mathbf{O} is given by $\mathbf{P}_{\mathbf{O}}$ (5.6). The corresponding update with the FB algorithm is the closed form eq.5.9.

Parameter updates. In this subproblem, an adapted setting of the parameters Λ and Υ is important to control the leakages between the two components and so achieve a good separation between \mathbf{AS} and \mathbf{O} .

- **Reweighting scheme:** The ℓ_1 and $\ell_{2,1}$ norms introduce some biases, which can be detrimental to the BSS problem in the presence of outliers, or at least lead to inaccurate solutions with artifacts (we refer to the Section 3.3.5). For this reason, a reweighting scheme is implemented [26, 113]: the values of the parameters Λ and Υ depend on the values of the estimated variables. More precisely, we will set $\Lambda = \lambda \odot \mathbf{W}_{\mathbf{S}}$ and $\Upsilon = v \odot \mathbf{W}_{\mathbf{O}}$, where $\lambda \in \mathbf{R}^{n \times 1}$ is a fixed vector setting the sparsity level of each source, and $\mathbf{W}_{\mathbf{S}} \in \mathbf{R}^{n \times t}$ corresponds to the varying weights. Similarly $v \in \mathbf{R}$ is the scalar setting the global sparsity level of the columns of $\mathbf{O} \Phi_{\mathbf{O}}^T$, while $\mathbf{W}_{\mathbf{O}} \in \mathbf{R}^{1 \times t}$ contains the weighting parameters.
- **Fixed parameters λ and v , App.B:** Similarly to the rAMCA algorithm, Section 4.3.2, the values $\{\lambda_i\}_{i=1..n}$ are fixed to $k\sigma_i$, where σ_i are obtained with the mad of $(\mathbf{A}^T (\mathbf{X} - \mathbf{O} - \mathbf{AS}) \Phi_{\mathbf{S}}^T)_i$. The value of v is set so as to limit the impact of the remaining Gaussian noise on the estimation of \mathbf{O} . Outside the support of $\mathbf{O} \Phi_{\mathbf{O}}^T$, the ℓ_2 norm of the columns of the centered Gaussian residual follows a χ -law with m degrees of freedom, whose expectation is given by $\sigma \times \sqrt{2} \times \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})}$, where σ can be estimated with the value of the mad of $(\mathbf{X} - \mathbf{AS} - \mathbf{O}) \Phi_{\mathbf{O}}^T$.

The parameter v is set to $v = k \times \sigma \times \sqrt{2} \times \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})}$.

- Weights \mathbf{W}_S and \mathbf{W}_O . At every iteration $\ell > 1$ such as in Alg.17, the parameters \mathbf{W}_S and \mathbf{W}_O are updated according to the current values of \mathbf{S} and \mathbf{O} respectively with:

$$\mathbf{W}_S = \frac{\lambda}{\lambda + |(\tilde{\mathbf{S}}\Phi_S^T)|} \text{ and } \mathbf{W}_O^q = \frac{v}{v + \left\| (\tilde{\mathbf{O}}\Phi_O^T)^q \right\|_2} \quad \forall q = 1..t. \quad (5.12)$$

We point out that \mathbf{W}_S and \mathbf{W}_O are reset to 1 for $\ell = 1$ so as to limit the propagation of the errors and make full benefit of the new estimation of \mathbf{A} by not enforcing the solutions to be similar to the previous ones.

5.2.3.2 Sparse BSS for the joint estimation of \mathbf{A} and \mathbf{S} . The joint estimation of \mathbf{A} and \mathbf{S} with fixed \mathbf{O} amounts to perform a standard sparse BSS on the current denoised observations $\mathbf{X} - \mathbf{O}$. For that purpose, we will make use of either the GMCA, Section 1.3.2, or the AMCA, Section 4.3.2, algorithms to update these variables.

During the very first iterations of the warm-up stage, a large part of the outliers is very likely to be mis-estimated and still present in the residual $\mathbf{X} - \tilde{\mathbf{O}}$, what will eventually hamper the unmixing process. Choosing the BSS algorithm that is the most robust to this residual will help enhancing the estimation \mathbf{A} . For that purpose either GMCA or AMCA will be used based on the relative choices of Φ_S and Φ_O :

- Highly incoherent dictionaries: If Φ_O and Φ_S are highly incoherent, the outlier residual is likely to be dense in Φ_S , similarly to the case displayed in fig.5.2. Using the standard fast GMCA, which is robust to the presence of Gaussian noise, and more generally to dense noise, is the best choice.
- Mildly incoherent dictionaries: In this case, the algorithm AMCA should be preferred. Indeed, the representations of the outliers and their residues in Φ_S are likely to be mildly sparse. In that case, we showed in Chapter 4 that the AMCA algorithm provides a more robust estimate of \mathbf{A} .
- Additional priors on the sources: Besides the morphology of the residual of the outliers in Φ_S , another additive knowledge on the data may justify the use of a specific sparse BSS algorithm. For example, if the sources are correlated, the algorithm AMCA, which was originally developed to handle partially correlated sources, should be preferred to GMCA, even if the residual of the outliers is dense in Φ_S .

Since AMCA and GMCA only differ by this weighting scheme, we will present the warm-up procedure using GMCA. The AMCA algorithm is implemented by adding the weighting proce-

dure presented in Section 4.3.2.

Component updates. The fast version of GMCA performs the separation directly in the transformed domain $\Phi_{\mathbf{S}}$. If $\Phi_{\mathbf{S}}$ is orthonormal, the returned results are exact, and provide a good approximation if $\Phi_{\mathbf{S}}$ is diagonally dominant [14]. The GMCA algorithm estimates alternatively \mathbf{A} and $\alpha_{\mathbf{S}}$ by minimizing:

$$\underset{\mathbf{A}, \alpha_{\mathbf{S}}}{\text{minimize}} \frac{1}{2} \left\| ((\mathbf{X} - \mathbf{O})\Phi_{\mathbf{S}}^T - \mathbf{A}\alpha_{\mathbf{S}}) \right\|_2^2 + \|\Lambda \odot \alpha_{\mathbf{S}}\|_1.$$

The algorithm estimates alternatively \mathbf{A} and the coefficients $\alpha_{\mathbf{S}}$ with projected least-squares to fasten the unmixing process. The corresponding updates are given in Alg.17 and further details can be found in Sections 1.3.2 and 3.2.2.

Parameter updates. The strategies used for the setting of the parameters involved in GMCA are crucial for the robustness against the noise and local minima. They are presented below:

- The values of $\Lambda = \lambda \odot \mathbf{W}_{\mathbf{S}}$ plays a key role in AMCA and GMCA. In order to adopt the efficient scheme used presented in Section 3.3.3 and to limit the propagation of the errors due to a previous mis-estimation of \mathbf{S} , the weights $\mathbf{W}_{\mathbf{S}}$ are set to 1 during the unmixing process. The final threshold λ_i for each $\alpha_{\mathbf{S}_i}$ is $k\sigma_i$ where σ_i corresponds to the standard deviation of the noise corrupting the coefficients of the i th source, and $k \in (1, 3)$ [14]. The value of σ_i , if not known, can be estimated with the value of the mad of the coefficient $\alpha_{\mathbf{S}_i}$ before the thresholding operation, (see App.B and Section 3.3.3 for further details).

5.3 1D Monte-Carlo Simulations

We compare tr-rGMCA with standard robust BSS methods. These methods as well as the different criteria used to compare the algorithms are presented at the beginning of this section. The different strategies are compared first with simulated data allowing Monte-Carlo simulations (40 runs for each varying parameter). Last, they are compared on realistic simulated data from the ESA-Planck mission in the presence of additional point-sources emissions which act as outliers.

5.3.1 Performance criteria and algorithms for comparison

5.3.1.1 Algorithms for the comparison. Only few methods presented in the literature can handle the considered problem. Most of these strategies require additional assumptions, and consequently, cannot be used without regarding the setting. In this section, we

present the selected strategies for the comparison explaining in which experiments they will be used.

Proposed optimization strategy. In order to highlight the robustness of the proposed minimization strategy, we will compare it with the following other implementations:

- **Oracle with \mathbf{A} known.** In this case, we assume that \mathbf{A} is known, and we separate the outliers from the source contributions using the morphological diversity between the two components. The difference between its results and the ones of tr-rGMCA illustrates the loss of accuracy leaded by the blind unmixing process.
- **The PALM procedure only.** To underline the advantage of using the initialization procedure, we also minimize 5.3 using only the refinement step in Alg.1. Since a reweighted procedure is implemented in tr-rGMCA, the refinement procedure is run three times: first, it is initialized with null \mathbf{S} and \mathbf{O} and the matrix \mathbf{A} used for tr-rGMCA, and for the second and third times, with the current estimates of \mathbf{S} and \mathbf{O} (it is important for the update of the regularization parameters, Section 5.12).
- **The warm-up step only.** The intermediate performances, obtained by the initialization step only, will be also displayed. A difference between these results and the PALM procedure would bring out the robustness of this initialization step, and the dissimilarity with the all process tr-rGMCA would show the gain of using a more precise refinement step.

Methods used for the comparisons.

- **The combination Outliers Pursuit (OP)+GMCA.** The outliers that are not laying in the subspace spanned by \mathbf{A} are first estimated by applying the Outlier Pursuit algorithm on $\mathbf{X}\Phi_{\mathbf{O}}^T$, [135], Section 2.2.1. Then the algorithm GMCA is applied on the denoised observations $(\mathbf{X} - \tilde{\mathbf{O}})\Phi_{\mathbf{S}}^T$. This strategy requires the term \mathbf{AS} to have low-rank, and thus, it will only be used when $m > n$. Given that the parameter proposed in [135] does not return satisfactory results, we choose to tune its value: we select the best $\tilde{\mathbf{A}}$ among the ones obtained from GMCA after the Outlier Pursuit for which we set the parameter λ (notations of [135]) between $\frac{1}{5\sqrt{t}}$ and $\frac{10}{\sqrt{t}}$ with a step-size of $\frac{1}{5\sqrt{t}}$.
- **The rNMF algorithm [58].** This component separation method, presented in Section 2.2.3, estimates jointly \mathbf{A} , \mathbf{S} and \mathbf{O} . It assumes that the sources and the outliers are non-negative, and that the outliers corrupt entirely (in the domain of observations) few columns of \mathbf{X} . It necessitates also the low-rank assumption, and that the data samples lie in the simplex with almost pure pixels. This method will be used in the last experiment (detection of point-source emissions in astrophysics), in which the outliers, sources and mixing matrix are all non-negative. All the conditions required for the rNMF to be efficient will not be valid.

- **ICA-based on a β divergence minimization** [100], Section 2.2.2². This ICA-based method assumes that the sources are mutually independent and that $m = n$. That is why, we will only use this method when $m = n$ since otherwise, a dimension reduction technique is needed (and is challenging in the presence of outliers, see Section 4.4.2). Besides, it only returns \mathbf{A} , and thus, does not perform the separation between the outliers and the sources contribution.

In contrast with the other methods, a strong morphological diversity makes the unmixing more challenging for this method. Indeed, it should be performed in a domain in which few samples are corrupted, and so in $\Phi_{\mathbf{O}}$. However, if the morphological diversity is strong, then the expansion coefficients of the sources in $\Phi_{\mathbf{O}}$ are highly non-sparse (see for example fig.5.2(c), the sources coefficients in $\Phi_{\mathbf{O}}$ almost follow a Gaussian distribution): this is difficult to handle for ICA-based methods. On the other hand, if $\Phi_{\mathbf{S}}$ and $\Phi_{\mathbf{O}}$ are not highly incoherent, then the outliers are likely to not corrupt all the samples in $\Phi_{\mathbf{S}}$. It is then preferable to perform the minimization in $\Phi_{\mathbf{S}}$ since the sources are better represented and the outliers do not corrupt all samples.

Last, setting the value of β is challenging in practice. We select the best \mathbf{A} for the 20 preselected values of β , starting from 10^{-4} to 0.85.

- **GMCA.** It will be performed on $\mathbf{X}\Phi_{\mathbf{S}}^T$. Its results illustrate the sensitivity of the standard (non-robust) methods to the outliers.
- **The combination MCA+GMCA.** Similarly to the combination OP+GMCA, the outliers are first discarded from the observations using MCA presented in 1.2.2, and the unmixing is then performed on the cleaned data. Instead of using the spectral diversity such as done by the OP algorithm, this combination only exploits the morphological diversity to discard the outliers.

It is indeed possible to separate \mathbf{AS} from \mathbf{O} , without regarding the ratio $\frac{n}{m}$, as long as \mathbf{AS} is sparse in $\Phi_{\mathbf{S}}$. We point out that this hypothesis can be valid only in the presence of a small number of sources. Besides, this approach does not take into account the clustered, structural aspect of the product \mathbf{AS} .

5.3.1.2 Performance criteria. In this section, we present the different criteria used to compare the algorithms. In the context of robust BSS, they should assess the unmixing of the sources (recovery of \mathbf{A}), the separation between the outliers and the sources as well as the reliability of the separation (especially because the problem is not convex) - see App.C for further details on the metrics.

Unmixing.

- For each recovered $\tilde{\mathbf{A}}$, the global quantity $\Delta_A = -10 \log 10 \left(\frac{\|\tilde{\mathbf{A}}^\dagger \mathbf{A} - \mathbf{I}\|_1}{n^2} \right)$ is computed.

²python implementation from [63]

- For each recovered $\tilde{\mathbf{A}}$, the maximal angle between the estimated and actual columns of \mathbf{A} is computed: $\max_{j=1,\dots,n} \arccos\langle \tilde{\mathbf{A}}^j, \mathbf{A}^j \rangle$ (in degree).

For every considered parameter, we sum the number of runs for which an algorithm has returned a mixing matrix whose maximal angle is smaller than 5 degrees (normalized to one).

Estimation of the sources and outliers.

- In Section 5.3, we will only display the median over the n sources of the SDR: it provides a global criterion on the precision of the source estimation. In Section 5.4, the medians as well as the minima for the n sources of the SDR, SAR, SIR and SNR will be displayed, so as to describe more precisely the obtained estimations.
- The sources can be erroneously estimated whereas the outliers and \mathbf{AS} are correctly estimated (for the Frobenius norm). To measure the quality of the separation between \mathbf{AS} and the outliers, the two components of interest for MCA and OP, we also compute the following metric for the outliers: $\mathbf{O}_{SE} = -10 \log \frac{\|\tilde{\mathbf{O}} - \mathbf{O}\|_2}{\|\mathbf{O}\|_2}$, where \mathbf{O} denotes the initial outliers, $\tilde{\mathbf{O}}$ the estimated ones.

5.3.2 1D Simulations

We start by comparing the different strategies on 1D data permitting Monte-Carlo simulations, with varying parameters. For this purpose, we will generate two kinds of data sets, which are described in the next part.

5.3.2.1 Data setting.

- **Data setting 1:** we consider n sources whose expansion coefficients are exactly sparse in DCT. They are drawn from a Bernoulli-Gaussian law, with an activation parameter of 5% and a standard deviation of 100. These sources are mixed into m observations, which are corrupted by outliers and an additive Gaussian noise. The outliers are sparse in the direct domain ($\Phi_{\mathbf{O}} = \mathbf{I}$). The support of the active columns of \mathbf{O} follow a Bernoulli law, with varying activation rates fig.5.3. The amplitude of the active entries are drawn from a centered Gaussian distribution, with a standard deviation equal to $100 \times \frac{8}{m}$ (so that for the two considered numbers of observations $m = 8$ and $m = 40$, it will remain quite constant relatively to the amplitude of \mathbf{AS}). The number of samples is fixed to 4096. For the two data-sets, the entries of the mixing matrix are drawn from a Gaussian distribution and the columns of \mathbf{A} are then normalized for the ℓ_2 norm. Besides, \mathbf{A} is generated so as to have a condition number smaller than 100.
- **Data setting 2:** this is a more realistic setting, with a same number of samples $t = 4096$. The sources are first generated from a Bernoulli Gaussian law in the direct domain, with an

activation rate of 2% and a standard deviation of 100. The sources are then convolved with a Laplacian kernel (FWHM equal to 20). They can be sparsely represented using redundant 1D wavelets [122]. The outliers are generated so as to correspond to a high frequency structured noise- approximately column sparse in the DCT domain. First, we generated a 1×4096 vector whose entries are drawn from a generalized Gaussian distribution with scale parameter 0.1. In order to obtain a high frequency texture, the amplitude of the DCT coefficients are scaled (from 10^{-4} for the lowest frequency to 1 for the highest one, with a logarithmic range), and the lowest 500 coefficients are manually set to 0. Last, this vector is multiplied (dot-wise) by a matrix generated from a Gaussian distribution, whose columns are normalized for the ℓ_2 norm, so that $\mathbf{O}\Phi_{\mathbf{O}}^T$ is approximately column sparse.

The first data setting is almost ideal since the expansion coefficients are exactly sparse and the mutual coherence between the DCT and the direct domain is very low. On the other hand, the second one is more realistic: approximately sparse expansion coefficients and the mutual coherence between the wavelets and the DCT is larger than the one between DCT and the direct domain.

5.3.2.2 1D Monte-Carlo simulations - Optimization strategy. In this first set of experiments using the first data-setting for 8 sources and 8 observations, we consider an easy setting to compare different optimization strategies which can be used to minimize (5.3). The SNR for the additive Gaussian noise is set to 30dB.




Figure 5.3: Performance indicators for a varying percentage of corrupted columns in the determined case for different optimization strategies.

First, one can notice in Fig.5.3 that in the presence of very few outliers (percentage of corrupted columns equal to 1%), the different strategies perform similarly in term of precision and reliability. Moreover, their corresponding values of the SDR Fig.5.3 is also close to the one obtained by the oracle: the unmixing task does not hinder the estimation of the sources. However, in the presence of numerous outliers, some disparities appear: the different strategies do not perform similarly and as well as the oracle. The PALM implementation (refinement step

of tr-rGMCA) is more precise than the initialization step for the unmixing (Δ_A has larger values Fig.5.3), but it is not as robust: except when there are only very few outliers, it cannot recover \mathbf{A} for all the runs, contrary to the initialization step (with a percentage smaller than 30%) Fig.5.3. However, adding the refinement step after the initialization step (the proposed strategy for tr-rGMCA) permits a significant gain in term of precision: all the values of the performance indicators are higher with tr-rGMCA than with the initialization step only. Moreover, the SDR and the error for the outliers Fig.5.3 obtained with tr-rGMCA are very close to the ones of the oracle: the unmixing of tr-rGMCA is robust and does not deteriorate the estimation of the sources while the percentage of corrupted columns is smaller than 30%. On the overall, tr-rGMCA is almost not influenced by the percentage of corrupted columns while this one is smaller than 30%. However it quickly fails, similarly to the oracle, in the presence of a larger percentage. Even if the dictionary chosen for \mathbf{O} is not the most adapted one (\mathbf{O} does not have a very sparse representation), the separation between the outliers and the source contribution can be good as long as the components have sparser representation in their associated dictionary than in the other one.

These results support the proposed strategy used for tr-rGMCA. In the following, only the results obtained by the oracle and tr-rGMCA will be displayed.

5.3.2.3 1D Monte-Carlo simulations - Comparison in the determined case. Only few methods able to handle the presence of outliers in the determined case are present in the literature. We propose to compare these methods with tr-rGMCA in this challenging setting.

Influence of the percentage of corrupted columns - 2 Sources. In this experiment, the data are generated with the first data-set with 2 sources and 2 observations. The SNR, for the Gaussian noise, is set to 60dB. In the determined setting, one can envisage using the mini-




Figure 5.4: Performance indicators for a varying percentage of corrupted columns in the determined case for 2 sources.

mization of the β -divergence, and the combination MCA+GMCA (which can also be used in the over-determined setting). The results obtained by the minimization of the β -divergence, fig.5.4 are better than the ones of GMCA but not as reliable or as precise as the ones of tr-rGMCA or MCA+GMCA. However, we explained in the presentation of the different methods used for the experiments, that this setting is challenging for the minimization of the β -divergence. Besides, the parameter β needs to be finely tuned, and we only tried 20 different values for this parameter. The second comment that can be made regarding fig.5.4, is on the impressive performances of the combination MCA+GMCA which performs very similarly to tr-rGMCA and the so called oracle. We will see in the next experiments that the combination MCA+GMCA is nonetheless not able to handle the presence of a larger number of sources.

Influence of the outlier amplitude - 8 Sources . The data are generated from the second data-setting. We consider that 8 sources have been mixed into 8 observations. The SNR for the Gaussian noise is set to 50dB. In this experiment, we observe the influence of the amplitude of the outliers. For this purpose, we define the SOR (signal to outlier ratio), similarly to the SNR: $SOR = 20 \log \frac{\|\mathbf{AS}\|_2}{\|\mathbf{O}\|_2}$. The support of the outliers remains constant for a given run, and only their amplitude is modified, by setting the SOR according to the value of the x-axis of Fig.5.5.

On the overall, the values of the different performance indicators are smaller than with the




Figure 5.5: Performance indicators for a varying amplitude of the outliers in the determined case for 8 sources.

first data-set: the second data-set, more realistic, is indeed more complicated. More specifically, one can note the significant gap for the SDR between the oracle and the other methods fig.5.5, whereas the outlier estimations have a similar precision fig.5.5: the additional unmixing clearly affects the results. The discrepancy between the minimization of the β -divergence and tr-rGMCA is reduced in this setting: the minimization of the β -divergence is performed in $\Phi_{\mathbf{S}}$, which is favorable to the unmixing.

With this data set and this number of sources, MCA fails to separate precisely the outliers from the source contributions, and the consecutive GMCA returns erroneous solutions. It fails because the component \mathbf{AS} is not sparse enough in $\Phi_{\mathbf{S}}$ (the number of sources is too large), and that

it does not take into account the structure, the clustered aspect of the product \mathbf{AS} . This is illustrated in Fig.5.6: the estimation of \mathbf{AS} obtained by MCA+GMCA is fair, but the resulting sources are clearly not correctly estimated. On the other side, the proposed tr-rGMCA is robust to outliers having large amplitudes, at least, much more than the standard BSS method GMCA.

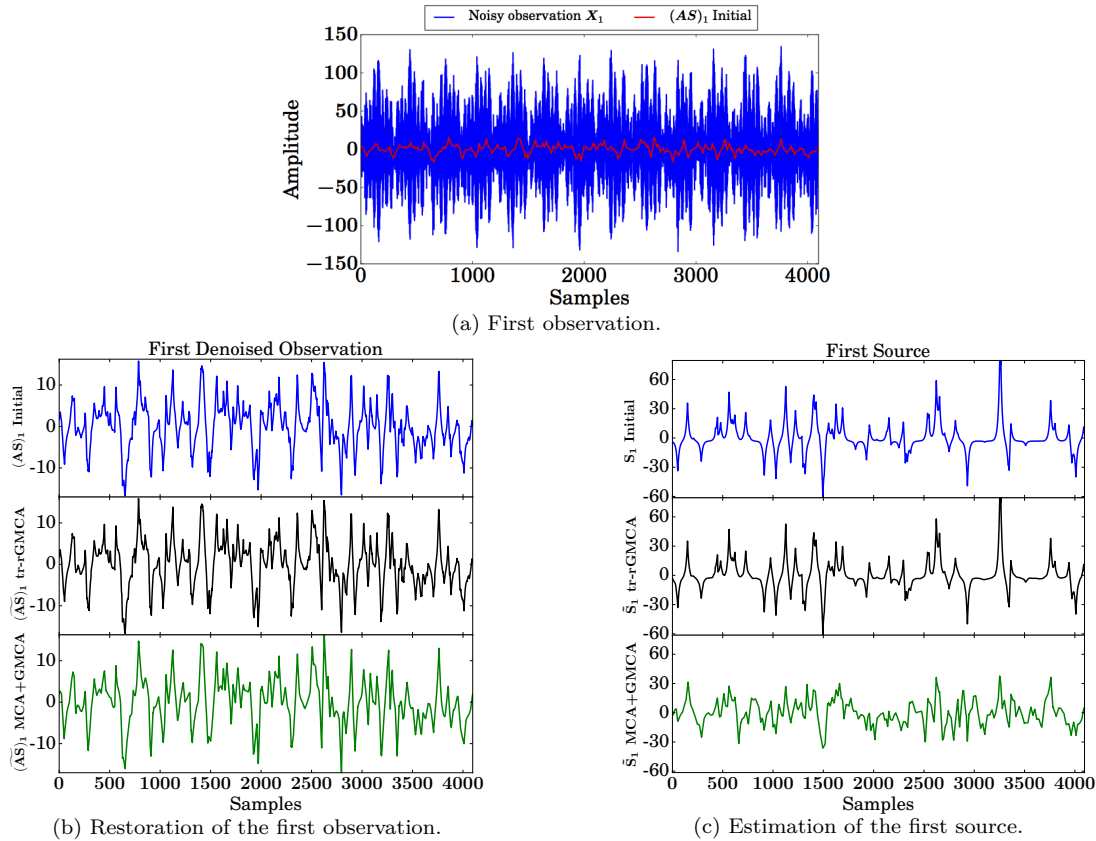


Figure 5.6: Illustrations of the estimated signals for a SOR equal to -10dB . Top: noisy \mathbf{X}_1 and noiseless $(\mathbf{AS})_1$ first observation. On the left, restoration of the first observation, on the right, estimation of the first source. In blue, the initial signals, in black, the ones recovered by tr-rGMCA and in green by the combination MCA+GMCA.

5.3.2.4 1D Monte-Carlo simulations - Comparison in the over-determined case. In Section 5.1.1, we underline the importance of the ratio $\frac{m}{n}$ in robust BSS. To illustrate it, we vary the number of observations, for 6 sources, with the two data-settings. The SNR is fixed to 50 dB and the SOR to -10dB for the 1st data set and 10dB for the second. Besides, the condition number of \mathbf{A} , which plays a crucial role in robust BSS (we refer to the numerical experiments dedicated to the influence of the condition number in Section 4.4.2), is

very likely to decrease with an increasing m . In order to limit the influence of this parameter, the condition number of \mathbf{A} is limited to 5.

First data-set. We start with the first data-set. The results obtained by the different meth-




Figure 5.7: Performance indicators for a varying number of observations, m , for 6 sources for the first data set

ods are improved if $m \gg n$, fig.5.7. Given that the outliers in $\Phi_{\mathbf{S}}$ are broadly distributed, they act similarly to an additive Gaussian noise with a large variance. Most of the methods used for the comparison used GMCA which is robust to the presence of a large Gaussian noise thanks to the thresholding operator whose threshold value varies according to the current noise level. However, this large threshold value leads to the presence of artifacts and biased source coefficients. When the number of observations becomes large, the projection of outliers in the span of \mathbf{A} has a smaller energy, and so, the corresponding apparent noise level becomes also smaller: the artifacts become also smaller, and both \mathbf{A} and \mathbf{S} are more accurate. That is why, most of the methods are able to estimate \mathbf{A} , and \mathbf{S} fairly when $m \gg n$.

It can also be noticed that even if m is close to n , the combination Outlier Pursuit (OP) + GMCA is able to retrieve \mathbf{A} , while GMCA alone cannot. The sources and the outliers are not precisely retrieved, but the results are the second best after tr-rGMCA. With the strong morphological diversity, the outliers are very sparse in $\Phi_{\mathbf{O}}$ while the source contribution is very dense: sparsity is discriminative enough, and OP can discard a part of the outliers. Removing the largest outlier contribution is sufficient, since this data set is very favorable to GMCA.

Second data-set. The different methods are on the all less performant with the second data set, even if an improvement is also noticeable if $m \gg n$, especially for the combination OP+GMCA fig.5.8. The proposed tr-rGMCA is the only method able to estimate precisely and reliably the three variables, including when the number of sources is close to the number of observations. The combination MCA+GMCA struggles to solve the problem because the




Figure 5.8: Performance indicators for a varying number of observations, m , for 6 sources.

number of sources is too large (\mathbf{AS} is not sparse enough). The algorithm Outlier Pursuit (OP) cannot identify precisely the outliers and the term \mathbf{AS} , but can discard efficiently the part of the outliers that are detrimental for the unmixing (the results obtained for \mathbf{A} are fair).

These results strongly differ with the ones obtained by rAMCA in Chapter 4, which was able to estimate the mixing matrix, but not the sources in the determined case.

5.4 Application to simulated astrophysical data

BSS made possible the estimation of high accuracy estimates of the Cosmological Microwave Background (CMB) from multi-wavelength microwave Planck data, Section 1.1.1. However, the presence of point-source emissions and spectral variabilities of some of the galactic foreground emissions is not precisely described by the standard linear mixture model. That is why most of the component separation methods only seek for a partial CMB map, in which the galactic center and the point source emissions of known locations are masked. Since each point source has a specific spectral signature, they cannot be modeled as individual components and are rather considered as outliers, Section 2.1.1. We therefore propose applying the tr-rGMCA algorithm to robustly estimate the galactic emissions (once the CMB is estimated and its contribution discarded from the observations) in the presence of unknown point source emissions.

5.4.1 Simulated data

In the following, we simulate 20 realistic CMB-free observations in the microwave range at the proximity of the galactic center, which have been produced using the Planck Sky Model [44]. These observations correspond to the mixture of 4 galactic emissions, namely, synchrotron, spin dust, free-free, and thermal dust. Since the rank of \mathbf{AS} is 4 and the number of observations is fixed to 20, it will make sense to apply as well separation methods that assume the low-rankness of the source contribution. Each observation is an image of size 128×128 . The signal-to-noise ratio (for the Gaussian noise \mathbf{N}) is set to 60 dB. Ten extra point source emissions with different emission laws are added. The point sources are modeled as Diracs convolved with a

same Gaussian kernel with varying width w , accounting for the point spread function (beam) of the instrument fig.5.9.




Figure 5.9: Top row: the 4 initial emissions. Second row: (left) normalized spectra of the emissions (*i.e.* columns of \mathbf{A}), and then examples of the observations, for a width of the kernel equal to 1.

5.4.2 Upgrades of tr-rGMCA

In contrast with the tr-rGMCA algorithm we used so far, additional properties can be accounted for in the separation:

- Non-negativity of the outliers and the sources. In this application, all the variables are non-negative. Taking into account non-negativity of \mathbf{S} and \mathbf{O} is particularly efficient to limit the leakages and artifacts between the two contributions. Non-negativity is constrained in the version of the tr-rGMCA algorithm that we used in the next experiments.
- Convulsive model for the point sources. The outliers are sparse in the direct domain. However, each one is perfectly described with the convolution of the instrument PSF and a Dirac with unknown position and amplitude. Therefore, the tr-rGMCA algorithm is extended so as to account for this convulsive model.

Consequently, we slightly modify the cost function of tr-rGMCA as follows:

$$\underset{\mathbf{A}, \mathbf{S} \geq 0, \mathbf{O} \geq 0}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S} - \mathbf{H} * \mathbf{O}\|_2^2 + \|\mathcal{Y} \odot \mathbf{O}\|_{2,1} + \|\Lambda \odot \mathbf{S}\Phi_{\mathbf{S}}^T\|_1 + \chi_{\mathbf{Y}: \|\mathbf{Y}^k\|_2 \leq 1, \forall k}(\mathbf{A}), \quad (5.13)$$

where \mathbf{H} denotes the Gaussian kernel. The non-negativity constraints and the deconvolution are taken into account during the joint estimation of \mathbf{O} and \mathbf{S} of the warm-up procedure as well as the refinement step. This does not change the structure of the algorithm and the BSS method that is used to estimate jointly \mathbf{A} and \mathbf{S} (AMCA). Only the updates of \mathbf{O} and \mathbf{S} are changed during their joint estimation in the warm-up and the PALM algorithm.

The cost function of the subproblem associated with the update of \mathbf{O} is composed of one differentiable term, with a Lipschitz gradient, and a regularization term (non-negativity and $\ell_{2,1}$ norm) whose proximal operator has a closed form: the update of \mathbf{O} can be efficiently tackled using the FB algorithm presented in 3.1.2.

On the other hand, the minimization problem associated with the update of \mathbf{S} is also composed of a differentiable term with Lipschitz gradient, and two regularization terms (non-negativity in the direct domain and sparsity in a transformed domain such as in [113]), having both explicit proximal operators. This subproblem is handled with the Generalized Forward Backward Splitting algorithm [110], introduced in 3.1.3.

5.4.3 Experiments

5.4.3.1 A challenging setting. First, we underline that the proposed problem is particularly difficult to tackle:

- It has first been noticed that the large scales of these astrophysical sources are partially correlated [13], which dramatically hampers the performances of standard BSS. This is precisely for this type of sources that the AMCA algorithm [13] has been designed. Therefore, the AMCA algorithm will be used in the warm-up stage to provide robustness with respect to these partially correlations.
- Some features of the thermal dust emission 5.9d have morphologies that are close to the one of the outliers 5.1b. The dictionary $\Phi_{\mathbf{S}}$ should be chosen so that all the sources are well represented, and also so that $\Phi_{\mathbf{O}}$ and $\Phi_{\mathbf{S}}$ are incoherent. More precisely, the astrophysical sources admit an approximately sparse representation in the wavelet domain. The spurious outliers are modeled as the convolution of Dirac functions with the point spread function of the instrument (PSF). More precisely, the convolution kernel is modeled as a Gaussian function $\exp^{-\frac{((x-x_0)^2+(y-y_0)^2)}{w}}$, where (x, y) denotes the position of the pixel, and (x_0, y_0) , the pixel in the center of the image (the kernels are then normalized). In the following, the amplitude of \mathbf{O} is fixed from one experiment to another (and so their energy increases with w). Consequently, this setting makes the particular choice of wavelet functions critical since it will largely impact the coherence between the $\Phi_{\mathbf{S}}$ and $\Phi_{\mathbf{O}}$. On the one hand, highly oscillating wavelet functions (*i.e.* with a large number of vanishing moments) will yield more incoherent dictionaries but at the cost of slightly less sparse representations for the sources. On the other hand, more localized wavelet functions are likely to provide better

sparse representations but at the cost of lowering the morphological diversity between the dictionaries. Therefore, in this particular robust BSS problem, one needs to make a trade-off between the compressibility of the sparse representations, which is essential for source separation, and the morphological diversity between Φ_S and Φ_O , which is of paramount importance for the separation of the sources and the outliers. In the next experiments, Φ_S will be chosen as undecimated Daubechies wavelet transforms with varying vanishing moments.

5.4.3.2 Influence of the dictionary Φ_S . To further highlight the role played by the mutual coherence in the proposed tr-rGMCA algorithm, we propose to investigate the influence of the vanishing moments of the Daubechies wavelet functions used for Φ_S . We only compare the different methods that are influenced by the choice of Φ_S : the so-called oracle, tr-rGMCA, AMCA performed on \mathbf{X} and $\mathbf{X} - \mathbf{O}$ (the combination MCA+AMCA performs so poorly that the influence of Φ_S cannot be commented, and the influence of Φ_S on OP+AMCA can be deduced by the performances of AMCA).

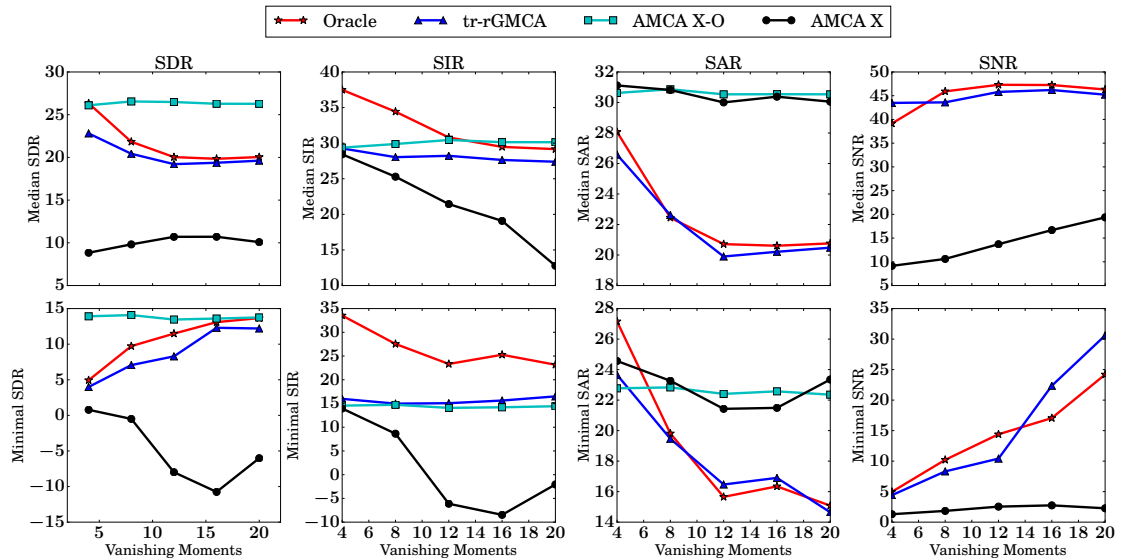


Figure 5.10: Performance indicators for a varying number of vanishing moments.

Vanishing Moments:	4	8	12	16	20
$\frac{\ \mathbf{S}\Phi_{\mathbf{S}}^T\ _1}{\ \mathbf{S}\Phi_{\mathbf{S}}^T\ _2}$	143.36	142.18	142.48	143.06	143.64
$\frac{\ H\Phi_{\mathbf{S}}^T\ _1}{\ H\Phi_{\mathbf{S}}^T\ _2}$	21.97	31.48	38.42	45.13	52.19

Table 5.2: Influence of the number of vanishing moments on the representation of \mathbf{S} and the outliers.

First, the choice of $\Phi_{\mathbf{S}}$ does not significantly impact the AMCA algorithm that is performed by $\mathbf{X} - \mathbf{O}$ fig.5.10: the representation coefficients of the sources are sufficiently sparse, for the different dictionaries, to perform the unmixing (the ratio $\frac{\|\mathbf{S}\Phi_{\mathbf{S}}^T\|_1}{\|\mathbf{S}\Phi_{\mathbf{S}}^T\|_2}$, which somehow measures the level of sparsity of the sources in $\Phi_{\mathbf{S}}$, does not significantly change in table 5.2). However, one of the sources, the thermal dust emission is not very accurately recovered: it is hampered by the correlations between sources and as well as their spectra, which is illustrated in fig.5.9.

The influence of $\Phi_{\mathbf{S}}$ for the oracle and tr-rGMCA are similar. The SIR and SAR decrease when the number of vanishing moments increases. Indeed, the largest scales of the astrophysical sources are partially correlated. Therefore, the most discriminative coefficients in the wavelet domain are located in the finest wavelet scales, while the number of vanishing moments is small. On the other hand, the SNR increases, especially the minimal one: the outliers do not leak towards the estimated sources when the number of vanishing moment is large enough (the outliers are less sparsely represented in $\Phi_{\mathbf{S}}$, table 5.2).

In the following, we will make use of the Daubechies wavelets with 20 vanishing moments so as to recover all the sources fairly while providing an improved separation with respect to the outliers.

5.4.3.3 Influence of the kernel width. In this experiment, $\Phi_{\mathbf{S}}$ is fixed and the width of the Gaussian kernel w , which also tends to alter the coherence between $\Phi_{\mathbf{S}}$ and $\Phi_{\mathbf{O}}$ and as well the morphological diversity between \mathbf{O} and \mathbf{S} , is varying. The kernel width w will vary according to the x -axis of fig.5.11.

Variance - parameter w :	0.1	0.5	1	1.5	2
$\frac{\ \mathbf{S}\Phi_{\mathbf{O}}^T\ _1}{\ \mathbf{S}\Phi_{\mathbf{O}}^T\ _2} *$	105.7	104.5	96.9	79.6	76.9

Table 5.3: Influence of the width of the Gaussian kernel on the representation of \mathbf{S} . * the sources are artificially deconvolved with \mathbf{H} .

As illustrated in fig.5.11, all the methods are impacted by the width of the kernel (*i.e.* the




Figure 5.11: Performance indicators for a varying variance of the Gaussian kernel.

morphology of the outliers).

First, we recall that the unmixing is very difficult when w is small: the outliers contaminate the high frequency content of the sources, which is discriminant for the unmixing (the large scales of the sources are correlated). That is why, the results are overall improved when w increases but is small.

On the other hand, we can notice that the methods are hindered by a large w . In that case, the 'low-frequency' (similar to the kernel) content of the sources, which contains most of their energy, become highly sparse in $\Phi_{\mathbf{O}}$ (tab.5.2): they leak towards the estimated outliers. Consequently, the SAR (fig.5.11) decreases as w increases. This is especially true for the thermal dust emission (associated with the minimal SAR), whose singularities have a morphology very similar to the one of the kernel. The leakages are also reinforced by the fact that the large scales of the sources are correlated: the $\ell_{2,1}$ penalization in $\Phi_{\mathbf{O}}$ is less expensive than the ℓ_1 in $\Phi_{\mathbf{S}}$ for the correlations. Besides, the energy of the outliers on the coarse-scale of $\Phi_{\mathbf{S}}$ increases, but is not thresholded (because it is not sparse): this is clearly hampering the SNR when w is large.

Only the combination OP+AMCA is able to outperform tr-rGMCA in term of highest SDR, while the kernel is not too large. However, only tr-rGMCA is able to fairly recover the thermal dust emission, as well as AMCA performed on $\mathbf{X} - \mathbf{O}$. We underline that the parameter of OP was manually tuned knowing the ground truth, and presumably, if the parameters involved in tr-rGMCA were similarly tuned, its performances would be, at least, similar to the ones of OP+AMCA. The rNMF method, even if it is initialized from the ground truth \mathbf{A} , was not able to correctly unmix the sources and separate the outliers from the source contribution: the fact that the data samples do not lie in the simplex makes this method inefficient in this experiment




Figure 5.12: First row: ground truth signals. Estimated sources with tr-rGMCA (second row), OP+AMCA (third row) and rNMF (fourth row) with $w = 1$ and 20 vanishing moments for the wavelets.

since \mathbf{O} cannot be separated from the \mathbf{AS} .

Illustrative results are provided in fig.5.12, 5.13. Outlier residuals are present in the sources estimated by rNMF and OP+AMCA, fig.5.12. On the other hand, the highest frequency contributions of the sources are not correctly recovered by tr-rGMCA (the SAR are quite low fig.5.12), and have leaked towards the estimated outliers. The spectra recovered by tr-rGMCA are the most precise, the other methods have failed to recover the thermal dust spectrum precisely fig.5.13.




Figure 5.13: Estimated spectra, with $w = 1$ and 20 vanishing moments for the wavelets. Red lines: estimated spectra, black dashed lines: ground truth.

5.5 Summary and future works

In this chapter, we introduced a new solution for the BSS problem in the presence of outliers that permits a robust estimation of the mixing matrix and an accurate separation of the sources and the outliers. The proposed tr-rGMCA algorithm estimates jointly the mixing matrix, the sources and the outliers so as to simultaneously unmix the sources and separate the outliers from the source contribution. Building upon sparse modeling, it first exploits the morpho-spectral diversity between the outliers and source contribution to distinguish between them, including in the challenging determined setting. The tr-rGMCA algorithm builds upon a two-stage optimization procedure: i) a warm-up stage based on heuristics that yield a reliable algorithm with enhanced robustness and ii) a refinement step based on the PALM algorithm that provably converges to a stationary point to the problem. Numerical experiments have been carried out on Monte-Carlo simulations showing the robustness of the proposed approach, which provides state-of-the-art results.

Chapter 6

Hyperspectral unmixing with spectral variabilities

Summary

6.1 Hyperspectral unmixing	138
6.1.1 Linear Mixture Model	138
6.1.2 LMM based methods	139
6.1.3 Spectral variability in the literature	141
6.1.4 Perturbed Linear Mixing Model (PLMM)	143
6.1.5 Constraining the components	144
6.2 Indeterminacy of the models	145
6.2.1 Characterization of the admissible solutions	145
6.2.2 Local endmembers as a deviation from the reference endmembers	146
6.2.3 Regularization of the SV	147
6.2.4 Minimization problems	150
6.2.5 Recoverability of the components	152
6.3 Algorithm	153
6.3.1 Initialization	153
6.3.2 Joint estimation of $\Delta\mathbf{A}$, \mathbf{S} and \mathbf{A}	154
6.3.3 Algorithm improvements	156
6.4 Numerical Experiments	157
6.4.1 Datasets	157
6.4.2 Initialization	158
6.4.3 Spectral variation retrieval	160
6.4.4 Joint Analysis	161
6.5 Summary and perspectives	167

In remote-sensing, hyperspectral images provide a fine spectral description of the observed scene. From these data, it is possible to estimate the composition of the observed region: nature of the elements by identifying their spectral signatures, and their spatial distribution.

In many scenarios, the spectral signatures of the components are spatially variant. These variabilities are not taken into account in the standard linear data model, what hampers the associated unmixing methods. An explicit estimation of the variabilities is thus necessary to recover precisely the composition of the observed region.

In this chapter, we will start by reviewing some of the unmixing methods handling spectral variabilities. We will then extend the robust BSS methods that were introduced in this thesis to estimate the variabilities.

6.1 Hyperspectral unmixing

6.1.1 Linear Mixture Model

LMM. The linear mixture model (LMM) [80] assumes that hyperspectral data can be modeled as followings:

$$\mathbf{X} = \mathbf{A}\mathbf{S} + \mathbf{N},$$

where \mathbf{X} denotes the observations, \mathbf{A} the spectral signatures (also called endmembers), \mathbf{S} the spatial distributions of the elements (designated as abundances), and \mathbf{N} a Gaussian noise.

This model and significances of the components (spatial distribution of the components for \mathbf{S} and their spectral signatures for \mathbf{A}) are exactly the same as for multichannel data analysis 1.1.1. Consequently, blind unmixing (the endmembers are unknown) of hyperspectral data is fundamentally similar to a standard BSS problem with multichannel data.

Nonetheless, the large number of channels (hundreds for hyperspectral versus tens for multichannel) provides a finer spectral description of the observed scene. Consequently, hyperspectral data have some spectral properties which are not shared by multichannel data, e.g. low-rankness ($m \gg n$) or smoothness of the spectral signatures [102].

Simplex assumption. In the special context of 'terrestrial' remote-sensing¹, it is generally further assumed that the variables are non-negative (spectra and spatial distributions are non-negative), and that, for each pixel t , the abundances sum to one: $\sum_{i=1}^n \mathbf{S}_i^t = 1$, [80]. With this simplex assumption, the sources entries indicate the relative proportions of the elements at each

¹In this chapter, we will focus on 'terrestrial' hyperspectral images, for which the simplex assumption holds true (in opposition to astrophysical hyperspectral images, for which the simplex assumption can be not valid). Unless specified otherwise, we will consider that this hypothesis is valid.

pixel. *This is a strong assumption, since the sources are no longer independent:* the j th entry of \mathbf{S} at a pixel k can be inferred from the other entries with $\mathbf{S}_j^k = 1 - \sum_{i=1, i \neq j}^n \mathbf{S}_i^k$.

With these assumptions, the data samples $\{\mathbf{X}^k\}_{k=1..t}$ live in the convex hull of \mathbf{A} fig.6.1: $\{\mathbf{X}^k : \mathbf{X}^k = \sum_{i=1}^n \mathbf{S}_i^k \mathbf{A}^i, \sum_{i=1}^n \mathbf{S}_i^k = 1, \mathbf{S} \geq 0\}$.

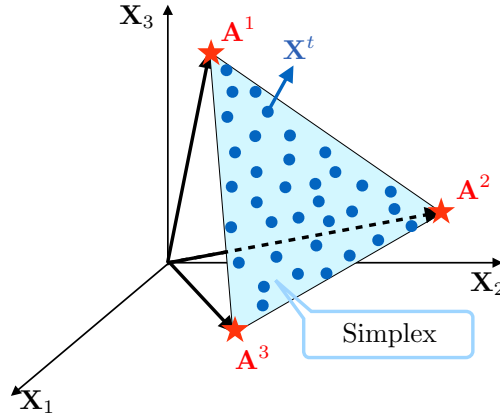


Figure 6.1: Illustration of the simplex defined by 3 endmembers $\mathbf{A}^1, \mathbf{A}^2, \mathbf{A}^3$. The three endmembers (red stars) define the vertices of the simplex (light blue triangle). The data samples (blue circles) lie in the simplex.

We point out that this assumption is not always valid in hyperspectral data analysis e.g., in astrophysics, there may be a region of the observed sky with no observable element.

6.1.2 LMM based methods

Simplex assumption. Several methods have been proposed to solve the blind unmixing problem using the LMM and the simplex hypothesis, see - [10], [98] for a more detailed overview.

Some of these methods assume that, for each component i , there is a pixel t_i such that $\mathbf{X}^{t_i} = \mathbf{A}^i$ (pure pixel assumption) fig.6.2a. It is then possible to estimate the endmembers from the data samples by looking for n data-samples $\{\mathbf{X}^{t_j}, t_j \in (1, \dots, t), j = 1, \dots, n\}$ defining the simplex enclosing all the other data samples. In practice, one can look for n data samples, among $\{\mathbf{X}^j\}_{j=1..t}$, whose associated simplex has the maximal volume, as done by example by VCA [104] or N-FINDR [132].

Another strategy consists in estimating directly \mathbf{A} and \mathbf{S} , without the pure pixel assumption. In that case, looking for (\mathbf{A}, \mathbf{S}) is an all-posed problem without further assumption. The non-




Figure 6.2: Simplex representations of the data for 3 cases. The black arrows symbolize the positions of initial endmembers, the blue circles the data samples, the light blue triangle the simplex defined by the endmembers, and the yellow triangle, the simplex with minimal volume enclosing the data samples.

negativity and simplex assumptions are not sufficient: there is an infinite number of \mathbf{A} such that, \mathbf{A} generates a simplex in which the data samples are enclosed. By assuming that there are well spread out data samples on each facet of the simplex to be recovered fig.6.2b, it is possible to estimate jointly \mathbf{A} and \mathbf{S} by looking for the simplex of minimal volume enclosing the data [93]. On the other side, if there is no pure pixel and if the pixels are too highly mixed, the solution with minimal volume is not the one that we aim to recover fig.6.2c.

Sparsity. If the simplex assumption holds true, the sources have a special structure which can be recovered with sparse BSS, without enforcing the simplex assumption (enforcing the simplex assumption is meaningful only if the sparsity promoting term is not based on an ℓ_1 norm), [65], [108]:

- If there are numerous pure pixels fig.6.3a, 6.2a the sources are nearly jointly sparse. This is the ideal case for sparse BSS.
- The presence of many broadly distributed data samples on the facets of the simplex 6.3b, 6.2b, is also well handled by sparse BSS. Indeed, even if there is no pure pixel, but a large number of not too highly mixed pixels, then sparse BSS will be effective. For instance, any rotation or transformation of the sources in fig.6.3b, results in less sparse sources (in the sense that they will be enclosed in an ℓ_1 ball with larger radius).

Interestingly, the conditions required for the minimal-volume and sparsity-based methods are very similar. Sparse BSS can also be seen as a minimal-volume based technique. Indeed, in the source domain, sparse BSS looks for the tuple (\mathbf{A}, \mathbf{S}) so that \mathbf{S} is enclosed in the ℓ_1 ball with minimal radius.

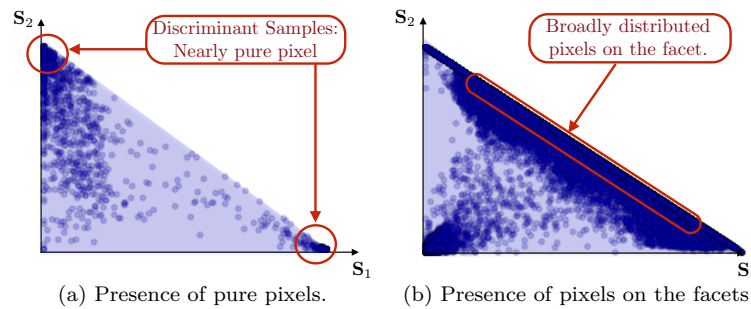


Figure 6.3: Scatter plot of two over three abundances: (a) in the presence of numerous (nearly) pure pixels, (b) in the presence of broadly distributed data-samples on the facets - not highly mixed pixels.

6.1.3 Spectral variability in the literature

The Linear Mixture Model has been widely used for the analysis of hyperspectral data, but cannot account for the spectral variabilities encountered in practice 2.1.1. These spectral variabilities hamper the standard methods based upon the LMM fig.6.4, and are of interest in some situations, for example for crop monitoring [90].




Figure 6.4: Simplex representations of the data. The red stars denote the vertices (endmembers), the dark-blue circles the data sample in the simplex, and the light-blue circles the data samples not living in the simplex with the LMM. The minimal volume simplex enclosing all the data is larger than the one defined by the reference endmembers.

With the recent advances in hyper-sensors, the estimation of the spectral variabilities (SV) has become a burning issue during the last years. More precisely, efforts have been undertaken to model non-linearities and/or SV. The non-linearity can be, in a certain extend, related to the SV if considering that the deviations from the linear model can be recast as spectral perturbations. Reviews of the first strategies can be found in [119], [45], [138], and some are presented below:

- The bilinear/multilinear/post-nonlinear approaches (non-linearity [45]): based on physical con-

siderations (multiple reflections for instance), this class of strategies assumes that each pixel can be expressed as a combination of the spectra, but also multiplicative combination of the spectra, *e.g.* $\mathbf{A}^i \odot \mathbf{A}^j$, $i, j = 1, \dots, n$ or any higher order relations between the spectra - [71], [3], [73] (with a more general framework). These models have been more precisely developed to handle non-linearities, but the residuals from the first order endmember contribution can be interpreted as spectral variabilities. However, it seems in practice that, all the variabilities cannot be expressed precisely with such models unless considering high enough moments such that all the combinations create an (over-)complete spectral dictionary.

- Non-linearity as another component (non-linearity [58], [4]): In [58] and [4], the authors propose to model the non-linearity as a third component, similarly to what was done in the previously presented robust BSS methods with \mathbf{O} . The deviations are assumed to be also non-negative in [58]: it cannot precisely model the SV (it makes more sense to assume that the perturbations from the LMM created by the SV are centered). In [4], the authors propose, in a Bayesian framework, to estimate the three components, without the non-negativity assumption on the outliers.
- Regression with spectral library (spectral variation, [138], [77]): This approach consists in estimating the abundances and endmembers, given that the endmembers belong to a predefined spectral library. In practice, selecting the active endmembers per pixel and associated abundances amounts to solve a compressive sensing problem [77], [25]. The effectiveness of this method strongly relies on the spectral library: there is a trade-off between the capability to represent various spectral variations (many spectra), and the possibility to separate the contribution (low coherence of the mixing matrix [25]).
- Endmembers as realizations of multivariate distributions (spectral variation, [138]): This strategy is more flexible than the method restraining the endmembers to a predefined set. In this case, the endmembers are seen as a realization of a multivariate distribution, whose parameters depend on the pixel and element [72]. Nonetheless, the accuracy of this method crucially depends on the distribution and parameters. Besides, and in contrast with the regression from a spectral library, the user has fewer control on the returned endmembers.
- SV as a perturbation from reference endmembers (spectral variation, [127]): In [127], the authors propose to model the SV as a perturbation from some reference endmembers. We will build upon this approach, and describe it in the following.

6.1.4 Perturbed Linear Mixing Model (PLMM)

In order to account for the SV, we will assume that the data can be well described by the PLMM that was first proposed in [127]:

$$\mathbf{X}^t = \sum_{i=1}^n \mathbf{S}_i^t \times (\mathbf{A}^i + \Delta\mathbf{A}_i^t) + \mathbf{N}^t, \quad (6.1)$$

where \mathbf{X}^t denotes the t th pixel (column vector of size m) of the hyperspectral cube $\mathbf{X} \in \mathbb{R}^{m \times t}$, \mathbf{S}_i^t the t th pixel of the i th source $\mathbf{S} \in \mathbb{R}^{n \times t}$, \mathbf{A}^i the endmember (column vector of $\mathbf{A} \in \mathbb{R}^{m \times n}$) associated with the i th source, $\mathbf{N} \in \mathbb{R}^m \times n$ the Gaussian noise corrupting the observations and last, $\Delta\mathbf{A}_i^t$ the spectral perturbation of the i th element at the t th pixel (column vector of size m), $\Delta\mathbf{A}_i \in \mathbb{R}^{m \times t}$, $\forall i = 1..n$.

This model can be recast with the following matrix form:

$$\mathbf{X} = \underbrace{\mathbf{A}\mathbf{S}}_{\text{best rank-}n \text{ approximation}} + \underbrace{\sum_{i=1}^n \Delta\mathbf{A}_i \odot \mathbf{S}_i}_{\text{local deviations}} + \mathbf{N} \quad (6.2)$$

$$= \sum_{i=1}^n \underbrace{(\mathbf{A}^i + \Delta\mathbf{A}_i)}_{\text{local endmembers}} \odot \mathbf{S}_i + \mathbf{N} \quad (6.3)$$

Properties of the components.

Abundances \mathbf{S} : The sources are non-negative since they correspond to the spatial distribution of the components.

Endmembers \mathbf{A} : The mixing matrix is also non-negative because each column represents the spectral signature of one element.

Spectral variations $\Delta\mathbf{A}_i$: The local spectra should be non-negative: $\mathbf{A}^i + \Delta\mathbf{A}_i \geq 0$, $\forall i = 1, \dots, n$. Indeed, they correspond to the local spectral signature of one component, and as so, should be non-negative.

We will further assume that these elements are centered, *i.e.*, for each component i , its local spectra $(\mathbf{A}^i + \Delta\mathbf{A}_i)$ are centered around the reference endmember \mathbf{A}^i . For illustrative purpose, we display in fig.6.5b the endmembers of pure pixels of water, extracted from a MOFFETT data scene fig.6.5a and associated deviations for one pixel fig.6.5c.

Moreover, the deviations from the LMM, $\sum_{i=1}^n \Delta\mathbf{A}_i \odot \mathbf{S}_i$ are likely to not cluster in a specific direction. From this point of view, they correspond to the outliers term \mathbf{O} which was estimated in the previous chapter.




Figure 6.5: Spectral variations of the water for a MOFFETT scene. The local endmembers $\{\mathbf{A}^i + \Delta\mathbf{A}_i^t\}_t$ (in black in (b) and (c)) have been extracted from pixels estimated as being pure $\{\mathbf{S}_i^t\}_{\mathbf{S}_i^t=1}$ (a). The SV ($\Delta\mathbf{A}_i^t$ in green in (c)) are particularly important for this component.

However, and in contrast with the deviations previously studied, these deviations are not independent from the sources and mixing matrix.

6.1.5 Constraining the components

In this chapter, we will compare the influence of the basic hypotheses made on \mathbf{S} or \mathbf{A} on the recoverability of the SV: data samples living in the simplex, or spectra and local spectra living on the hyper-sphere (standard prior for BSS).

Data samples in a simplex - Model A. This model is the one proposed in [127]. It is assumed that fig.6.6a:

- The data follow the PLMM model 6.1.
- The sources are non-negative and the data samples live in the simplex defined by $\mathbf{S} \geq 0$, and $\mathbf{1}_n \mathbf{S} = \mathbf{1}_t$ ².
- The mixing matrix is non-negative $\mathbf{A} \geq 0$, and the local spectra are also non-negative $\mathbf{A}^i + \Delta\mathbf{A}_i \geq 0$, $\forall i = 1, \dots, n$.

With this model, the endmembers are not normalized, but rather take into account the energy of the components.

Endmembers on the hypersphere - Model B . In many applications, including with hyperspectral data, the source samples do not live in a simplex. For instance, in astrophysics, some regions of the sky do not contain observable elements. That is why, we have preferred another approach, whose hypotheses would be valid for a wider range of applications. Similarly to the robust BSS problems we have studied so far, we will assume that the spectra belong to the hypersphere \mathcal{S}^{m-1} fig.6.6b:

²Strictly speaking, the data samples do not longer belong to a simplex because of the SV. But we will say that the data belong to the simplex if the sources sum-to-one and are non-negative. The use of the Hadamard product is also an abuse of notation since the matrices do not have the same dimensions: the matrices are multiplied entry-wise along their common dimension

- The data follow the PLMM model 6.1.
- The sources are non-negative: $\mathbf{S} \geq 0$.
- The mixing matrix is non-negative $\mathbf{A} \geq 0$, and the local spectra are also non-negative $\mathbf{A}^i + \Delta\mathbf{A}_i \geq 0, \forall i = 1..n$. We will further assume that the estimated endmembers live on \mathcal{S}^{m-1} : $\|\mathbf{A}^i\|_2 = 1, \forall i = 1, \dots, n$, and $\|\mathbf{A}^i + \Delta\mathbf{A}_i^k\|_2 = 1 \forall i = 1, \dots, n, k = 1, \dots, t$. This assumption was used previously to constrain the scaling indeterminacy between \mathbf{A} and \mathbf{S} . It will be also useful for constraining the SV in this application.




Figure 6.6: Illustration of the possible positions of the data samples \mathbf{X}^k given the reference endmembers, for 2 elements and 3 observations.

6.2 Indeterminacy of the models

First, we underline that this problem is particularly ill-posed: even if \mathbf{A} and \mathbf{S} are perfectly known, there is still an infinite number of SV fitting the data for the non-pure pixels. Indeed, without additional assumption, estimating the SV amounts to seek for a local mixing matrix, $\mathbf{A}'_k \in \mathbb{R}^{m \times n}$ such that $\mathbf{A}'_k{}^i = \mathbf{A}^i + \Delta\mathbf{A}_i^k$, and $\mathbf{X}^k = \mathbf{A}'_k \mathbf{S}^k$ at every pixel k . In other words, without effective constrains of the SV, estimating the SV is independent from \mathbf{A} .

6.2.1 Characterization of the admissible solutions

The basic assumptions made by the two models lead to different admissible solutions for the local mixing matrix \mathbf{A}' . We will assume that \mathbf{A} and \mathbf{S} are known in the following of this section.

Possible solutions with Model - A. We point out that with the model A, the energetic contribution of each element is no longer related to the source amplitudes. For instance in fig.6.7a (with a same \mathbf{S}^k), the contributions of each element to the sample \mathbf{X}^k are different for the two set of local endmembers. *Consequently, for fixed \mathbf{S} , neither the energy nor the direction of the local endmembers are known fig.6.7a, even if \mathbf{S} is perfectly known.*




Figure 6.7: Illustration of possible local endmembers contribution for a fixed \mathbf{S}^k , for 2 endmembers and 3 observations.

Possible solutions with Model - B. Given that the local endmembers should belong to the hypersphere, the energetic contributions of the elements are still proportional to the source entries with Model-B. For fixed \mathbf{S} , we can recover the local endmembers *up to a rotation* (axial rotation with axis given by the direction of the data \mathbf{X}^k), with non-negative local endmembers.

6.2.2 Local endmembers as a deviation from the reference endmembers

In order to constrain the SV, we propose to illustrate them with two examples.

Detection of illumination variations. We point out that models A and B are very different, and that Model-A is more sensitive to the presence of deviations fig.6.6. Indeed, a data sample, which does not live in the simplex generated by the reference endmembers \mathbf{A} , is easily detected as being “corrupted” with model A: the SV of this sample are not null and should be consequently adapted. On the other hand, as long as a data sample belongs to the convex cone generated by \mathbf{A} , no spectral variation can be detected easily. The SV contribution is rather accounted by modifying the amplitudes of the sources. This greatly influences the detection and recoverability of the components.

In particular, the model A is sensitive to variations of illumination (if \mathbf{X}^k is on the simplex, $\alpha \mathbf{X}^k$ is not, $\forall \alpha > 0, \alpha \neq 1$). In contrast, a variation of illumination can be directly handled by the model B, by changing the amplitudes of the sources. For instance, fig.6.8a illustrates a change of illumination. The sample \mathbf{X}^k does not lie in the initial simplex (even by changing \mathbf{S}^k): the variations of illumination should be accounted in the perturbation. In this case, the orientation of the endmembers and the abundances are unchanged, and the endmembers are only scaled. The variation of illumination can only be detected with model-B by analyzing the variations of the amplitude of the sources samples.




Figure 6.8: Estimation of the SV. (a) Variation of illumination - model A, (b) SV with modification of the shape and norm of the endmembers - model A, (c) SV with modification of orientation, model B. The red arrows designate the two reference endmembers, the arrows in purple and blue the spectral perturbations, and in dashed lines: the novel simplex associated with the perturbations for model A and the novel directions of the local endmembers with model B.

Variation of direction/shape. Whenever a sample X^t does not lie in the cone generated by A , a modification of the shape/direction of the endmember is required for the Model B (and consequently A).

For example in fig.6.8b, the estimated local spectra are not proportional to the reference endmembers, their norm as well as their direction have changed.

With the model B, the norm of the local spectra is fixed, and only the direction is changed fig.6.8c.

6.2.3 Regularization of the SV

For both models, there is an infinite number of spectral variations fitting the data, even if S is known. Whether the basic model is based on model A or model B, different priors can be added to retrieve meaningful solutions, to ensure that the recovered local spectra are close to the reference endmembers (the retrieved elements belong to n classes, sharing similar spectral signatures).

We will assume that either model A or B is used, and focus on the additional priors that can be added on the SV to constrain the proximity of the local endmembers with the ones of references.

Global penalization of the SV. In [127], the authors propose to penalize the spectral variations with an ℓ_2 norm in order to *minimize the global energy of the SV*:

$$\text{Total Deviation: } \frac{1}{2} \sum_{i=1}^n \|\Delta A_i\|_2^2$$

With this regularization, the problem, for fixed (\mathbf{A}, \mathbf{S}) becomes well-posed: there is a unique solution for models A and B.

Nonetheless, the unique regularization parameter chosen for these deviations cannot account for the difference of energy between the components. Indeed, this regularization assumes implicitly (by seeing this term as a MAP) that the SV are generated according to a same Gaussian centered distribution. Consequently, the different SV are similarly penalized whereas they may have very different energies. This may be particularly inappropriate if using Model A, since the energy of the component is accounted by the endmembers (for example the endmembers of the MOFFETT scene of fig.6.5 are represented in fig. 6.9, Model-A).




Figure 6.9: Endmembers (Model-A) recovered by VCA for a MOFFETT scene.

In order to fully take into account the *inter-variability* that exists between the components, and *intra-variability* between one pixel to another for a same component, the total deviation penalization could be weighted:

$$\text{Total Deviation - multivariate: } \frac{1}{2} \sum_{i=1}^n \|\mathcal{Y}_i \Delta \mathbf{A}_i\|_2^2,$$

where $\mathcal{Y}_i \in \mathbb{R}^{1 \times t}$ contains the regularization parameters for $\Delta \mathbf{A}_i$, $i = 1, \dots, n$.

Pixel-wise penalization. As mentioned in [127], the $\ell_{2,1}$ can also be of interest to promote spatial sparsity of the SV. However, it appears in practice that the SV are not really spatially sparse - only the strongest deviations are indeed localized.

Nevertheless, we point out that the $\ell_{2,1}$ can also be considered as a flexible penalization of the energy:

$$\text{Pixel-wise Deviation: } \sum_{i=1}^n \|\mathcal{Y}_i \Delta \mathbf{A}_i\|_{2,1},$$

where $\mathcal{Y}_i \in \mathbb{R}^{1 \times t}$. With the $\ell_{2,1}$, the SV will be thresholded pixel-wise and not globally.

Nonetheless, there is no obvious advantage of using the $\ell_{2,1}$ norm in comparison to the weighted ℓ_2 squared norm for model A, especially given that the $\ell_{2,1}$ norm is not differentiable. The use of the $\ell_{2,1}$ is however meaningful if one wants only to recover the areas with the strongest deviations, or using a transformed domain to work with sparse representations of the SV. This last point will be discussed later. In the framework of the Model B, we will see that the $\ell_{2,1}$ norm has an interesting interpretation which supports its utilization.

Roles of the penalizations. Depending on the data model that is employed, the penalization has not the same influence.

Model - A: With the model A, the endmembers are representative of the directions/shapes of the spectral signatures and also of the energies of the elements. Indeed, we have observed previously that if \mathbf{S} is perfectly known, the explicit estimation of the SV brings a double indeterminacy: energy and direction. The regularization thus tackles both indeterminacies by looking for the SV with the smallest energy. For instance in fig.6.8, the sample \mathbf{X}^k is less expensively represented with the deviations of fig.6.8b than fig.6.8a (with similar values of the regularization parameters). Hence, even if the problem is no longer ill-posed, looking for the SV with minimal energy may not yield the true SV.

Model - B: With the model B, the endmembers are only representative of the directions/shapes of the spectral signatures. Minimizing the $\ell_{2,1}$ norm of the deviations amounts to seek for the smallest spectral angles fig.6.10. Indeed, for each pixel k and component i , the angular deviation θ_i^k between the local and reference spectrum (non-negative), and the SV $\Delta \mathbf{A}_i^k$ are linked by the formula $\|\Delta \mathbf{A}_i^k\|_2 = \sin\left(\frac{\theta_i^k}{2}\right) \times 2$ leading to:

$$\sum_{i=1}^n \|\Delta \mathbf{A}_i\|_{2,1} = \sum_{i=1}^n \sum_{k=1}^t \sin\left(\frac{\theta_i^k}{2}\right) \times 2. \quad (6.4)$$

Consequently, the minimization of the $\ell_{2,1}$ norm of $\Delta \mathbf{A}_i$ provides a proxy for the minimization of the angular deviations between the local spectrum and the reference, for each pixel.

Last, we recall that for fixed \mathbf{S} , the local spectra are recovered up to a rotation. *By adding the $\ell_{2,1}$ regularization, this indeterminacy is removed, which makes this model very powerful.* However, if the angular deviations are large for a given pixel, the return solution will not correspond to the sought after one.




Figure 6.10: Illustration of a local spectra perturbation and associated local spectrum. The circle represents a slice of the sphere embedding the spectra.

6.2.4 Minimization problems

We present two minimization problems, one for each model, aiming to recover jointly the components $\mathbf{A}, \mathbf{S}, \Delta\mathbf{A}_i$, of the form:

$$\underset{\mathbf{A}, \mathbf{S}, \Delta\mathbf{A}_i}{\text{minimize}} \frac{1}{2} \left\| \mathbf{X} - \mathbf{A}\mathbf{S} - \sum_{i=1}^n \Delta\mathbf{A}_i \odot \mathbf{S}_i \right\|_2^2 + \mathcal{J}_1(\mathbf{A}) + \mathcal{J}_2(\mathbf{S}) + \mathcal{J}_3(\Delta\mathbf{A}_1, \dots, \Delta\mathbf{A}_n), \quad (6.5)$$

where \mathcal{J}_i denote the regularization functions associated with the components to be retrieved.

6.2.4.1 Problem formulation Model A - [127]. The authors of [127] aim to retrieve the components with the following priors:

- **Abundances:** The abundances are assumed to be non-negative and to “sum-to-one”. These properties are enforced using two characteristic functions. The authors have also proposed to promote the spatial smoothness of the abundances with the term: $\frac{1}{2} \|\mathbf{S}\mathbf{H}\|_2^2$, where \mathbf{H} computes the difference between the abundances of a given pixel and its 4 closest neighbors.
- **Endmembers:** The endmembers are non-negative. Similarly to the standard LMM model, the volume of the simplex enclosing the data (not strictly for the PLMM because of the SV) defined by the endmembers should be minimal. The authors propose 3 different penalizations for constraining the volume, which seem to perform quite similarly in the tested numerical experiments in [127]: proximity of the endmembers to the ones estimated by VCA, minimal mutual distance between the endmembers (implicit minimal volume) and minimal volume with pure pixel assumption - (the one used in the numerical experiments of this chapter). For the latter, it consists in minimizing directly the volume of the simplex defined by the projection of the reference endmembers on the subspace spanned by the $n - 1$ principal components of \mathbf{X} , [127] and references therein.

- Spectral variations: As presented previously, the authors propose to minimize the energy of the SV with $\sum_{i=1}^n \|\Delta \mathbf{A}_i\|_2^2$. Besides, the endmembers plus the SV (the local spectra) are non-negative.

The complete cost function is the following [127]

$$\begin{aligned} & \underset{\mathbf{A}, \mathbf{S}, \Delta \mathbf{A}_i}{\text{minimize}} \frac{1}{2} \left\| \mathbf{X} - \mathbf{A}\mathbf{S} - \sum_{i=1}^n \Delta \mathbf{A}_i \odot \mathbf{S}_i \right\|_2^2 + \gamma \|\mathbf{S}\mathbf{H}\|_2^2 + \Upsilon \sum_{i=1}^n \|\Delta \mathbf{A}_i\|_2^2 + \beta |\det(\begin{bmatrix} \mathbf{T} \\ \mathbf{1}_{1,n} \end{bmatrix})|^2 \\ & \text{s.t. } \mathbf{A} \geq 0, \mathbf{S} \geq 0, (\mathbf{A}^i + \Delta \mathbf{A}_i) \geq 0, \mathbf{1}_{1 \times n} \mathbf{S} = \mathbf{1}_{1 \times t}, \forall i = 1..n, k = 1..t, \end{aligned}$$

where \mathbf{T} denotes the projection of \mathbf{A} on the subspace spanned by the $n - 1$ principal components of the observations \mathbf{X} .

The minimization is tackled using an ADMM algorithm and the regularization parameters are chosen by cross-validation.

6.2.4.2 Problem formulation Model B. Based on the model B, we aim to retrieve the components with the following priors:

- Abundances: We consider that the abundances are non-negative (characteristic function) and “sparse” in the direct domain (ℓ_1 norm).
- Endmembers: We assume that the endmembers are non-negative, and belong to the hypersphere.
- Spectral perturbations: We suppose that the local spectra are non-negative, and also live on the hypersphere. Besides, we will penalize the $\ell_{2,1}$ norm of the spectral deviations $\sum_{i=1}^n \|\Upsilon_i \Delta \mathbf{A}_i\|_{2,1}$, where $\Upsilon_i \in \mathbb{R}_+^{1 \times t}$, $\forall i = 1, \dots, n$. It should ensure that among the possible local spectra, the ones recovered are close to the reference endmembers with minimal spectral angle deviations.

Our minimization problem is the following:

$$\underset{\mathbf{A}, \mathbf{S}, \Delta \mathbf{A}_i}{\text{minimize}} \frac{1}{2} \left\| \mathbf{X} - \mathbf{A}\mathbf{S} - \sum_{i=1}^n \Delta \mathbf{A}_i \odot \mathbf{S}_i \right\|_2^2 + \sum_{i=1}^n \|\Upsilon_i \Delta \mathbf{A}_i\|_{2,1} + \|\Lambda \odot \mathbf{S}\|_1 \quad (6.6)$$

$$\text{s.t. } \mathbf{A} \geq 0, \mathbf{S} \geq 0, (\mathbf{A}^i + \Delta \mathbf{A}_i) \geq 0, \mathbf{A}^i \in \mathcal{S}^{m-1}, (\mathbf{A}^i + \Delta \mathbf{A}_i^k) \in \mathcal{S}^{m-1}, \forall i = 1..n, k = 1..t.$$

The associated algorithm and parameters setting are described in the next section.

6.2.5 Recoverability of the components

None of the two problems is able to recover exactly the components in practice. Model A suffers from a lack of priors for estimating the SV (the double indeterminacy in direction and norm cannot be solved using only the minimal energy assumption, and the perturbed data samples lying on the data simplex cannot be detected), and model B cannot detect the perturbed data samples lying in the convex cone generated by \mathbf{A} (model B can fairly estimates the SV only if \mathbf{A} and \mathbf{S} are correctly estimated).

Since the SV contribution is in general position, estimating fairly \mathbf{A} is possible with GMCA or AMCA but recovering correctly the sources is very challenging. Indeed, the SV contribution lying in the convex cone defined by \mathbf{A} should not leak toward the estimated $\tilde{\mathbf{A}}\tilde{\mathbf{S}}$. This problem is quite similar in the spirit to the separation between the outliers and the sources contribution in robust BSS. Likewise for the robust BSS framework, sparse modeling can be exploited to recover a good estimate of \mathbf{S} .

Source thresholding. The source estimates in the direct domain before thresholding are broadly given by $\tilde{\mathbf{S}} \approx \mathbf{S} + \mathbf{A}^\dagger(\sum_{i=1}^n \Delta\mathbf{A}_i \odot \mathbf{S}_i)$. Hence, if it were possible to set the thresholds Λ so that the residuals $\mathbf{A}^\dagger(\sum_{i=1}^n \Delta\mathbf{A}_i \odot \mathbf{S}_i)$ are thresholded, running GMCA or AMCA would provide fair estimates of \mathbf{A} and \mathbf{S} . Unfortunately, this was not possible in practice with the tested data sets for two reasons.

First, the difference of energy between the elements makes the setting of the thresholds quite complicated. For instance, let us consider two elements with very different amplitudes (e.g. the water and the soil in the studied MOFFETT scenes fig.6.9). The projection of the SV contribution associated with the soil onto the water source estimate is similar to the amplitude of the water source itself, and so cannot be thresholded without hampering strongly the water source estimate.

Second, even in the presence of sources with similar dynamics, a significant thresholding of the sources will hamper the sources estimates since the sources are not significantly sparse fig.6.3b. However, if the data are composed only of nearly pure pixels with sources having similar dynamics, it is then possible to recover fairly the sources and mixing matrix by just running AMCA.

Sparse modeling in transformed domains. In the spirit of tr-rGMCA, it would be possible to separate accurately the term \mathbf{AS} from $\mathbf{O} = \sum_{i=1}^n \Delta\mathbf{A}_i \odot \mathbf{S}_i$ if these two terms had different morphologies, or if there were a domain in which the sources were highly (and jointly) sparse and the SV contribution still broadly distributed.

Finding such dictionary is not straightforward with the sought-after abundances. Indeed, given that the sources are not independent, sparsifying the expansion coefficients *may rather highlight*

the dependence between the sources and returns highly correlated expansion coefficients (that is the case with the data sets used in the experiments and the wavelets for instance).

Likewise robust BSS, additional information may be needed to separate the SV from the sources contributions: such as a spatial prior. We remark that even if the SV and the source contributions are linked, they can still have a different morphology, for example if the SV are due to variations of illumination.

For instance, the curvelet transform is of interest to recover the SV along the boundaries between the components in the studied MOFFETT scenes. Indeed, the borders between the components define some smooth curves. Using the curvelets would then bring the morphological diversity (curvelets versus direct domain), and spatial smoothness along the borders. However, let us mention that using a redundant transform for the SV greatly raises the computational cost of the methods (the size of the SV is already $m \times t \times n$).

Our first idea was to extend the previous works, ideally tr-GMCA to correctly recover the components. As explained above, the extension is in fact not possible with the terrestrial data that we have studied: sparse modeling cannot easily be exploited to estimate and separate the different contributions.

6.3 Algorithm

We will detail in this section the algorithmic scheme that we employed to minimize eq.6.6.

The minimization strategy is a 2-step approach, in the spirit of the robust BSS algorithms that were previously presented. The proposed procedure is composed of first: an initialization of \mathbf{A} and \mathbf{S} with AMCA (see Section 4.3) followed by a BCD minimization (see Section 3.2.3) of the cost function for estimating jointly \mathbf{A} , \mathbf{S} and $\Delta\mathbf{A}$.

6.3.1 Initialization

This first step consists in initializing jointly \mathbf{A} and \mathbf{S} , i.e. we look for the best rank- n linear approximation of the data. We will build upon the AMCA algorithm to unmix robustly the data.

Given that the deviations are in general position, we propose to build upon the AMCA algorithm to estimate robustly the mixing matrix, similarly to the previous robust BSS problems. Even if the deviations are not sparse, some areas are more affected by the spectral deviations, and will be more penalized using AMCA. We have also preferred AMCA to GMCA to be robust to the presence of other components which may have not been taken into account but would hamper GMCA.

The sources are not, strictly speaking, sparse in the direct domain. The decreasing threshold strategy introduced in 3.3.3, is in this case particularly efficient: by starting with a large thresh-

old, AMCA estimates the endmembers from the nearly pure pixels; and then by *decreasing the thresholds towards zero*, the source estimates are not biased by the ℓ_1 norm. Besides, the non-negativity of \mathbf{A} and \mathbf{S} is enforced during the AMCA algorithm by projecting \mathbf{A} and \mathbf{S} on the non-negative orthant after every component update.

The sources are then refined with few proximal gradient updates (10 iterations) in order to take into account the non-negativity of the sources more properly. During this refinement, the source samples which contribute for less than 10% of a sample are set to zero. This is a conservative choice, but estimating the SV associated with small coefficients is particularly challenging: it is complicated to differentiate the SV from the influence of the residual.

The initialization scheme is presented in Alg.18.

Procedure 18 Initialization Abundances-Endmembers

```

1: procedure INITIALIZATION( $\tilde{\mathbf{A}}$ )
2:    $\tilde{\mathbf{A}}, \mathbf{S} \leftarrow \text{AMCA}(\mathbf{X})$  ▷  $\Lambda$  decreases towards zero
3:   Compute  $L_A = \|\tilde{\mathbf{A}}^T \tilde{\mathbf{A}}\|_{2,s}$ 
4:   for k=1..10 do ▷ Refinement of the sources
5:      $\tilde{\mathbf{S}} \leftarrow [\mathbf{S} + \frac{1}{L_A} \tilde{\mathbf{A}}^T (\mathbf{X} - \tilde{\mathbf{A}} \tilde{\mathbf{S}})]_+$  ▷ Proximal gradient step with non-negativity
6:      $\tilde{\mathbf{S}} \leftarrow \mathcal{S}_{\Lambda: \Lambda^t = \frac{\|\tilde{\mathbf{s}}^t\|_1}{10}}(\tilde{\mathbf{S}})$  ▷ Conservative threshold
   return  $\tilde{\mathbf{A}}, \tilde{\mathbf{S}}$ 

```

6.3.2 Joint estimation of $\Delta\mathbf{A}$, \mathbf{S} and \mathbf{A}

In this step, we aim at estimating the SV and refine the sources and mixing matrix with a BCD like procedure. The algorithm is presented in Alg.19.

Procedure 19 Joint Estimations \mathbf{A} , \mathbf{S} and $\Delta\mathbf{A}$

```

1: procedure JOINT ESTIMATION( $\tilde{\mathbf{A}}^{(0)}, \tilde{\mathbf{S}}^{(0)}$ )
2:   while do not converge do
3:      $\widetilde{\Delta\mathbf{A}}^{(k)} \leftarrow \text{GFBS}_{eq.6.7}(\mathbf{X}, \tilde{\mathbf{A}}^{(k-1)}, \tilde{\mathbf{S}}^{(k-1)})$ 
4:     Update  $\Lambda$ 
5:      $\tilde{\mathbf{S}}^{(k)} \leftarrow \text{FISTA}(\mathbf{X}, \tilde{\mathbf{A}}^{(k-1)}, \widetilde{\Delta\mathbf{A}}^{(k-1)}, \Lambda)$ 
6:     Recenter  $\mathbf{A}^{(k)}$  from  $\mathbf{A}^{(k-1)} + \widetilde{\Delta\mathbf{A}}^{(k)}$  and compute the associated  $\widetilde{\Delta\mathbf{A}}^{(k)}$ 
7:      $k \leftarrow k + 1$ 
   return  $\tilde{\mathbf{A}}^{(k)}, \tilde{\mathbf{S}}^{(k)}, \widetilde{\Delta\mathbf{A}}^{(k)}$ 

```

Source estimation:. The cost-function, for fixed \mathbf{A} and $\Delta\mathbf{A}$, can be solved pixel-wise and parallelized, $\forall k = 1..t$:

$$\underset{\mathbf{S}^k \geq 0}{\text{minimize}} \frac{1}{2} \|\mathbf{X}^k - \mathbf{A}'_k \mathbf{S}^k\|_2^2 + \|\Lambda^k \odot \mathbf{S}^k\|_1,$$

where $\mathbf{A}'_k = \mathbf{A}^i + \Delta\mathbf{A}_i^k$ represents the local endmembers at the k th pixel. The proximal operator of the combination of the ℓ_1 norm and the non-negative constraint has a closed form [112], see Appendix A. Consequently, the subproblem, for each pixel, is easily minimized with the FB algorithm presented in 3.1.2.

The parameters Λ are set to $3 \text{ mad} \left(\mathbf{A}^T \left(\mathbf{X} - \mathbf{A}\mathbf{S} - \sum_{i=1}^n \Delta\mathbf{A}_i \odot \mathbf{S}_i \right) \right)$ (see Section 3.3.3 and App.B).

Spectral variation estimation:. The estimation of the SV is given by:

$$\begin{aligned} & \underset{\Delta\mathbf{A}}{\text{minimize}} \frac{1}{2} \left\| \mathbf{X} - \mathbf{A}\mathbf{S} - \sum_{i=1}^n \Delta\mathbf{A}_i \odot \mathbf{S}_i \right\|_2^2 + \sum_{i=1}^n \|\mathcal{Y}_i \Delta\mathbf{A}_i\|_{2,1} & (6.7) \\ \text{s.t. } & (\mathbf{A}^i + \Delta\mathbf{A}_i) \geq 0, \mathbf{A}^i \in \mathcal{S}^{m-1}, (\mathbf{A}^i + \Delta\mathbf{A}_i^k) \in \mathcal{S}^{m-1}, \forall i = 1, \dots, n, k = 1, \dots, t, \end{aligned}$$

Similarly to the sources, this problem can be easily parallelized since all the pixels are independent.

The proximal operator of the $\ell_{2,1}$ norm term plus the characteristic functions is not explicit but is efficiently solved using the GBFS algorithm introduced in Section 3.1.3, [110].

The values of \mathcal{Y}_i are set according to the gradient $\nabla \left\| \mathbf{X} - \mathbf{A}\mathbf{S} - \sum_{i=1}^n \Delta\mathbf{A}_i \odot \mathbf{S}_i \right\|_2^2(\Delta\mathbf{A})$ of the differentiable term. In contrast with the sources, it does not yield clearly an estimate of the noise affecting the SV, but has two similar properties:

- It is large if the residual is large (only the largest SV are updated, with a large bias), which is interesting during the first iterations.
- It decreases towards 0 when the residual is nearly null (the small SV are also updated, and there is less bias and so a better fitting of the data).

In the spirit of the thresholding strategy which was proposed for the $\ell_{2,1}$ norm, App.B, the thresholds are set according to this ersatz of noise standard deviation, to which we added a

weighting³:

$$\Upsilon_i = \underbrace{\sqrt{2} \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})} \times \text{mad}(\nabla \left\| \mathbf{X} - \mathbf{A}\mathbf{S} - \sum_{i=1}^n \Delta \mathbf{A}_i \odot \mathbf{S}_i \right\|_2^2(\Delta \mathbf{A}))}_{\text{standard threshold for } \ell_{2,1}} \times \underbrace{\frac{\max(\mathbf{S}_i)}{\max(\mathbf{S})}}_{\text{inter-variability between the sources}} \times \underbrace{\mathbf{W}_i}_{\text{intra-variability of the source } i}.$$

The term $\frac{\max(\mathbf{S}_i)}{\max(\mathbf{S})}$ permits to rescale the values of Υ according to value of the sources (otherwise the residuals would be fit by the sources of largest energy). Indeed, we underline that the values of the gradients are much larger with the largest sources. However the SV represent only directions, not an energy depending of the pixel (the situation would be different with Model-A). That is why, we rescale the different gradients during parameter estimations.

The weighting term is such that $\mathbf{W}_i^t = \min(10, \frac{\|\mathbf{S}_i^t\|_1}{\mathbf{S}_i^t})$, ensures that the most significant source, at the pixel t , will handle the SV, and that a small entry will not be associated with large SV (otherwise, we risk to overfit the residuals with the smallest samples).

Mixing matrix:. The estimation of the mixing matrix is given by:

$$\begin{aligned} & \underset{\mathbf{A}}{\text{minimize}} \frac{1}{2} \left\| \mathbf{X} - \mathbf{A}\mathbf{S} - \sum_{i=1}^n \Delta \mathbf{A}_i \odot \mathbf{S}_i \right\|_2^2 \\ & \text{s.t. } (\mathbf{A}^i + \Delta \mathbf{A}_i) \geq 0, (\mathbf{A}) \geq 0, \mathbf{A}^i \in \mathcal{S}^{m-1}, (\mathbf{A}^i + \Delta \mathbf{A}_i^k) \in \mathcal{S}^{m-1}, \forall i = 1..n, k = 1..t, \end{aligned} \quad (6.8)$$

In contrast with the other components, the mixing matrix is (too) constrained. Once the SV are not null, the constraint $(\mathbf{A}^i + \Delta \mathbf{A}_i^k) \in \mathcal{S}^{m-1}$ makes any update of \mathbf{A} impossible. To alleviate this issue, we recenter 'manually' \mathbf{A} by taking the means of the local spectra, and adapting the SV consequently.

6.3.3 Algorithm improvements

This algorithm and associated numerical experiments are preliminary, but illustrate the interest of constraining the endmembers to live on the hypersphere.

Numerous improvements, and even reformulation of the problem would be of interest. We propose some of them:

³This threshold strategy is the one used in the numerical experiments. We would like to mention another procedure, developed after the first submission of the thesis. Let us suppose that the sources, mixing matrix and deviations are exactly recovered. The first order condition for eq.6.7, at the t th pixel and i th component gives us: $\Delta \mathbf{A}_i^t = \frac{1}{(\mathbf{S}_i^t)^2} \text{prox}_{\Upsilon_i^t \ell_{2,1}}(\Delta \mathbf{A}_i^t \times (\mathbf{S}_i^t)^2 + \mathbf{N}^t \times (\mathbf{S}_i^t))$. Then, one can threshold the deviation whose impact has an energy smaller than the Gaussian noise level with standard deviation σ (estimated with the mad on the residual) with $\Upsilon_i^t = \sigma \times \sqrt{2} \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})} \times \mathbf{S}_i^t$ (non-negative sources). One can prefer to threshold the deviation according to a minimal angle deviation θ by setting $\Upsilon_i^t = 2 \times \sin(\frac{\theta}{2}) \times (\mathbf{S}_i^t)^2$. They both directly, and locally, take into account the dynamics of the sources.

- The proposed formulation cannot be properly minimized. It would be however possible to consider directly the local spectra $\mathbf{A}_{loc} \in \mathbb{R}^{m \times n \times t}$ with $\mathbf{A}_{loc}^i = \mathbf{A}^i + \Delta \mathbf{A}_i$ and to reformulate the problem as follows:

$$\begin{aligned} \underset{\mathbf{A}, \mathbf{A}_{loc}, \mathbf{S}}{\operatorname{argmin}} \quad & \frac{1}{2} \left\| \mathbf{X} - \sum_{i=1}^n \mathbf{A}_{loc}^i \odot \mathbf{S}_i \right\|_2^2 + \sum_{i=1}^n \|\mathcal{Y}_i(\mathbf{A}^i - \mathbf{A}_{loc}^i)\|_{2,1} + \|\Lambda \odot \mathbf{S}\|_1 \\ \text{s.t.} \quad & \mathbf{A} \geq 0, \mathbf{S} \geq 0, (\mathbf{A}_{loc}^i) \geq 0, \mathbf{A}^i \in \mathcal{S}^{m-1}, (\mathbf{A}_{loc}^i) \in \mathcal{S}^{m-1}, \forall i = 1..n. \end{aligned} \quad (6.9)$$

With this formulation, we avoid the trivial local minimum with respect to \mathbf{A} once the $\Delta \mathbf{A}_i$ are not null, and the reference endmembers \mathbf{A} can be properly recentered.

- The setting of the regularization parameters for the $\Delta \mathbf{A}$, \mathcal{Y} , is not completely satisfactory. They start with a large value and then decrease towards zero, which is necessary to avoid a too quick overfitting of the data (without the updates of the sources). However, it seems that using the relation (6.4) between the norms of the $\Delta \mathbf{A}$ and the spectral deviations θ^i would be more efficient. Indeed, it is possible, by setting correctly the entries of \mathcal{Y} to control exactly the spectral deviations.
- The regularization parameters for the sources can mimic the simplex assumption: one can envisage to threshold the mixed pixels so that the corresponding sources entries are not larger than the amplitudes of the sources for the pure pixels. It would limit the leakages from the SV contribution towards the estimated sources.

6.4 Numerical Experiments

In the following, we will display the results obtained from simulated and real hyperspectral data (MOFFETT scenes). For the 3 data sets, we will look for 3 components: soil, vegetation and water.

6.4.1 Datasets

We will consider 3 datasets:

- Dataset 1: A first real hyperspectral data cube, with $m = 189$, and $t = 50 \times 50$ pixels.
- Dataset 2: A second real hyperspectral data cube (same area), with $m = 177$, and $t = 128 \times 128$ pixels.
- Dataset 3: A simulated hyperspectral data cube, generated from the first data set with $m = 189$, and $t = 50 \times 50$ pixels. The observations are generated from the results obtained by VCA [104] on the first data set (\mathbf{A} and \mathbf{S}). Spectral variations are added to the data: the $\Delta \mathbf{A}$

are drawn from a centered Gaussian distribution, and the local spectra are then normalized.

The maximal spectral angle is set to 15° for the water and 5° for the soil and vegetation.

For the 3 datasets, the abundances and spectra recovered from VCA (or the ground truth) are available. Besides, for the first data-set, the results obtained in [127] will be also displayed⁴. Last, the data sets have been multiplied by a factor of 100 for convenience.

6.4.2 Initialization

We first propose to describe the results obtained after the initialization with the dataset #1, and to compare them with the one obtained by VCA. For illustrative purpose, the endmembers estimated by VCA are normalized, and the energy of the abundances is modified accordingly.

Abundances. The abundances retrieved by AMCA and VCA are represented in fig.6.11. The results are on the overall quite similar. The main differences appear for the abundance of the water:

- The amplitude of $\mathbf{S}_{\text{water}}$ is abnormally large on the shore with AMCA fig.6.11c,6.11f. This can be explained by the fact that the data (including the SV contribution) lying in the convex cone are estimated as belonging to the source contribution (whereas it is not the case with VCA thanks to the simplex assumption 6.11i). Given that the energy of the water is much smaller than the other contributions, relatively small residuals from the two other components become very significant when projected on $\mathbf{S}_{\text{water}}$.
- One can also remark the abnormal presence of soil on the lake surface, due also to the projection of SV on the water contribution 6.11d.

Endmembers. The normalized spectra obtained by the initialization and VCA are displayed in fig.6.12.

The spectra of the soil and the vegetation are quite similar. The ones of the water differ more significantly. In fig.6.12b, the spectra of the pure pixels of water are also displayed. One can notice that the spectrum recovered by our initialization is closer to the mean of these spectra: the spectral signatures are robustly estimated by AMCA.

Detection of the largest SV. From the components returned by the initialization, it is already possible to localize the largest deviations.

We set $\tilde{\mathbf{X}} = \tilde{\mathbf{A}}(\max(0, \tilde{\mathbf{A}}^\dagger \mathbf{X}))$, the non-negative projection of \mathbf{X} on the span of the retrieved spectra, with non-negative contribution.

One can notice that the largest residuals are obtained for the water (whose SV are the most

⁴We are grateful to P.-A. Thouvenin for providing their results, as well as the datasets.

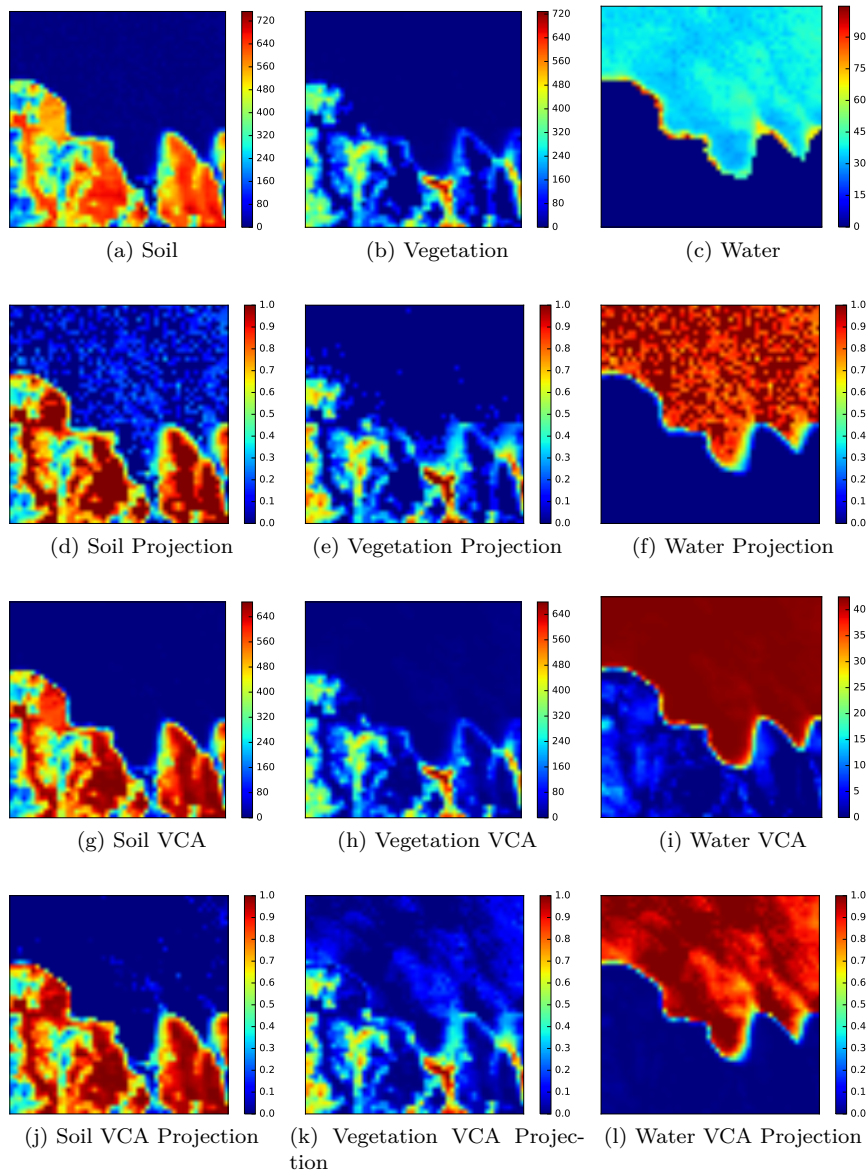


Figure 6.11: Abundance maps for the first data set. First row: after the initialization process, second row: energetic contribution of the element per pixel / projection of the sources with $\frac{\tilde{\mathbf{S}}^k}{\|\tilde{\mathbf{S}}^k\|_1} \forall k = 1, \dots, t$ where $\tilde{\mathbf{S}}$ are associated with normalized spectra, third row: abundances of VCA, fourth row: abundances of VCA similarly projected.

important), and the coastline fig.6.13a: the SV on these areas do not lie in the span of $\tilde{\mathbf{A}}$. Besides, when further looking at the δ -density (introduced in 4.3), we can remark that all the source entries are active, with similar amplitude, at the level of the coastline fig.6.13b: the best

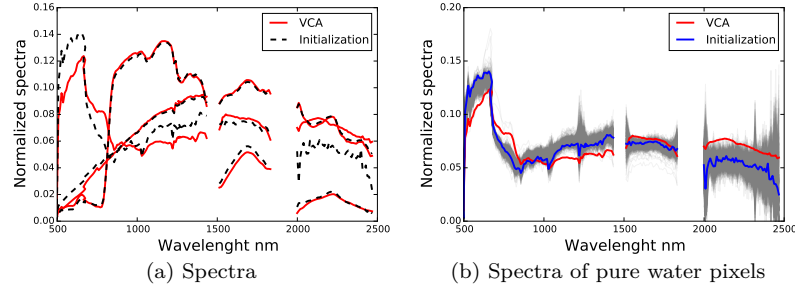


Figure 6.12: Left: normalized spectra recovered by VCA and the initialization. Right: spectral signature of the water recovered by VCA and the initialization, and the spectral signature of the pure water pixels.




Figure 6.13: Left: residual error per pixel $k \frac{\|\tilde{\mathbf{x}}^k - \mathbf{x}^k\|_2}{\|\mathbf{x}^k\|_2}$, in percent. Middle: δ -density of the pixel $k \frac{\|\tilde{\mathbf{s}}^k\|_1}{\max(\tilde{\mathbf{s}}^k)}$. Right: residual error per band: $\|\tilde{\mathbf{X}}_k - \mathbf{X}_k\|_2$, $k = 1..m$

fit (for the ℓ_2) needs to active simultaneously all the sources. Last, we underline that even if some bands have been removed before the analysis, some are still very corrupted/have a very low SNR fig.6.13c. It would be of interest to penalize these bands. This can be achieved for instance, by adding a weighting scheme when estimating the sources in the spirit of the rGMCA algorithm presented in Section 4.2:

$$\mathbf{S} \leftarrow \mathcal{S}_\Lambda((\mathbf{W}\mathbf{A})^\dagger)(\mathbf{W}\mathbf{X}),$$

where \mathbf{W} is a diagonal matrix, such that $\mathbf{W}_{m,m} = \frac{1}{\epsilon + \|\tilde{\mathbf{x}}_m - \mathbf{x}_m\|_2}$. With this kind a penalization, the sources would be estimated directly from the less-noisy bands.

6.4.3 Spectral variation retrieval

In order to illustrate the benefit of constraining the local spectra to live on the hypersphere, we propose to only estimate the spectral variations.

For this purpose, we fix \mathbf{A} and \mathbf{S} with the ground truths using the third data set, and seek for

the corresponding SV.

Data fitting. By estimating the SV, we almost perfectly fit the observations fig.6.14. Whereas the initial relative residual was up to 25% fig.6.14a, it decreases to less than 0.2% after estimating the SV fig.6.14b. The largest deviations are still encountered at the level of the coastline.

However we point out that the residual error is not representative of the accuracy reached for

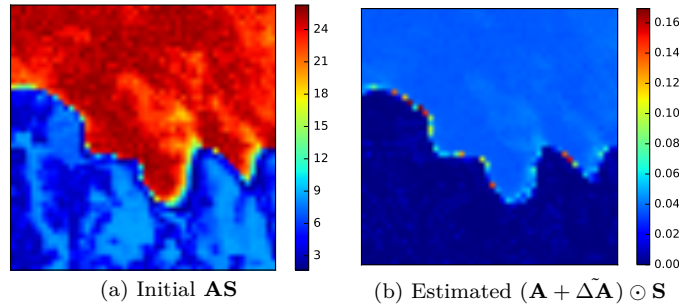


Figure 6.14: Left: residual error per pixel $k \frac{\|(\mathbf{AS})^k - \mathbf{X}^k\|_2}{\|\mathbf{X}^k\|_2}$, in percent. Right: residual error per pixel after estimating the SV $\frac{\|((\mathbf{A} + \Delta\mathbf{A}) \odot \mathbf{S})^k - \mathbf{X}^k\|_2}{\|\mathbf{X}^k\|_2}$, in percent .

the estimation of the SV: the problem is clearly ill-posed and fitting the observations is (too) simple.

SV estimation. We propose to analyze more precisely the accuracy reached for the SV estimation. It appears that the SV of the dominant components (per pixel) are correctly estimated. It is indeed very easy to estimate the SV of pure pixels (whose spectra are directly obtained by normalizing the observations), whereas estimating the SV of a mixture is more challenging. It can be clearly observed in fig.6.15: all the zones with one dominant component are correctly estimated (for example for the water in fig.6.15d, 6.15a), whereas there are significant errors for the coastline.

6.4.4 Joint Analysis

6.4.4.1 Dataset 1 - comparison Model A- Model B. Last, we present the final local spectra and sources returned by our method and compare them with the ones obtained in [127]. In order to clarify the comparison, the results of [127] are normalized as the following: the local spectra are normalized for the ℓ_2 norm, and the scaling factor is added to the sources: $\tilde{\mathbf{S}}_i^k \leftarrow \mathbf{S}_i^k \times \|\mathbf{A}^i + \Delta\mathbf{A}_i^t\|_2$, $\forall i = 1..n, k = 1..t$.




Figure 6.15: Pixel residual per component. Initial contribution of the component i : $x_i = (\mathbf{A}^i + \Delta\mathbf{A}_i) \odot \mathbf{S}_i$, without the SV: $y_i = (\mathbf{A}^i) \odot \mathbf{S}_i$, with estimated SV: $\tilde{x}_i = (\mathbf{A}^i + \Delta\tilde{\mathbf{A}}_i) \odot \mathbf{S}_i$. Top row: initial residual per pixel k : $\frac{\|x_i^k - y_i^k\|_2}{\|x_i^k\|_2}$, in percent. Bottom row: residual with estimated SV: $\frac{\|x_i^k - \tilde{x}_i^k\|_2}{\|x_i^k\|_2}$, in percent.

Source Estimation. The components returned by the two models are similar enough to be compared. We can observe the presence of stronger leakages from the SV contribution towards the estimated water map for Model B fig.6.16. However, the simplex assumption does not completely avoid these leakages: the amplitudes of the water map entries at the border between the water and the soil are much larger than for the water alone. Besides, it seems that Model B was more successful to determine the energy fraction of the components pixel-wise: the pixels returned by model A at the level of the water are only composed of 80% of water fig.6.16 (we recall that the simplex assumption does not longer determine the fraction of energy of each component since the energy is also accounted in the varying $\Delta\mathbf{A}$ with model A). The abundances maps of the soil and vegetation are more similar for models A and B, except for the pure water region.

Spectral Estimation. The estimated references endmembers returned by the two approaches are displayed in fig.6.19. The spectra of the vegetation and the soil are quite similar,

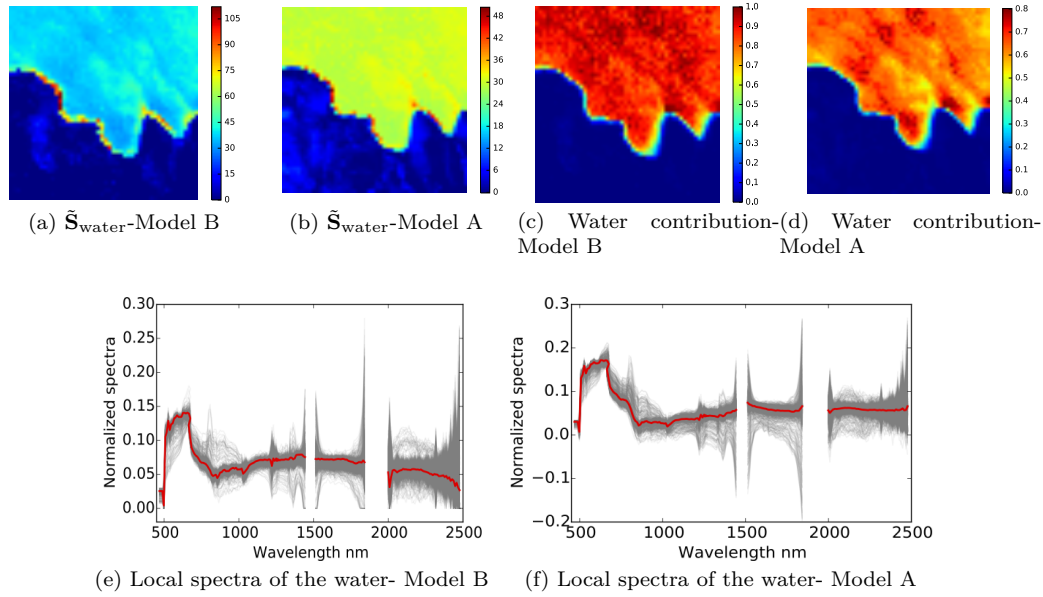


Figure 6.16: Returned components for the water for the real data set 1. Top row, from left to right: estimated sources model B and then A [127], water contribution to the pixel: $\frac{\tilde{\mathbf{S}}_{water}^k}{\|\tilde{\mathbf{S}}^k\|_1}$, $\forall k = 1..t$ for model B and A. Bottom: local spectra, with in red the reference spectra $\tilde{\mathbf{A}}$ with model B (left) and model A.

especially the ones of the soil. However, the estimated spectra for the water differ significantly.

The local spectra of the water present very large deviations for both methods fig.6.17f. This component is less energetic and is thus more sensitive to the presence of noise and residual (even the pure pixels of water present large variabilities fig.6.12b).

We can observe larger deviations for the local spectra of the soil and vegetation with model B than with model A fig.6.18, fig.6.17. This can be explained by the fact that a unique regularization parameter for the different SV was used for model A: 'a same amount of energy' is considered per element whereas the reflectance values of the soil are much higher than the ones of the water fig.6.9. However, it is quite difficult to state on the quality of the retrieved SV without any ground truth.

For the two models, we remark abnormal large deviations at the limit of the discarded bands: this was expected from the large residuals encountered in fig.6.13c.

We have also displayed in fig.6.20, the covariance matrix of the estimated local endmembers, in the spirit of [116]. We can remark that only few local endmembers have been wrongly estimated, in the sense that they are almost all more correlated with the other local spectra of

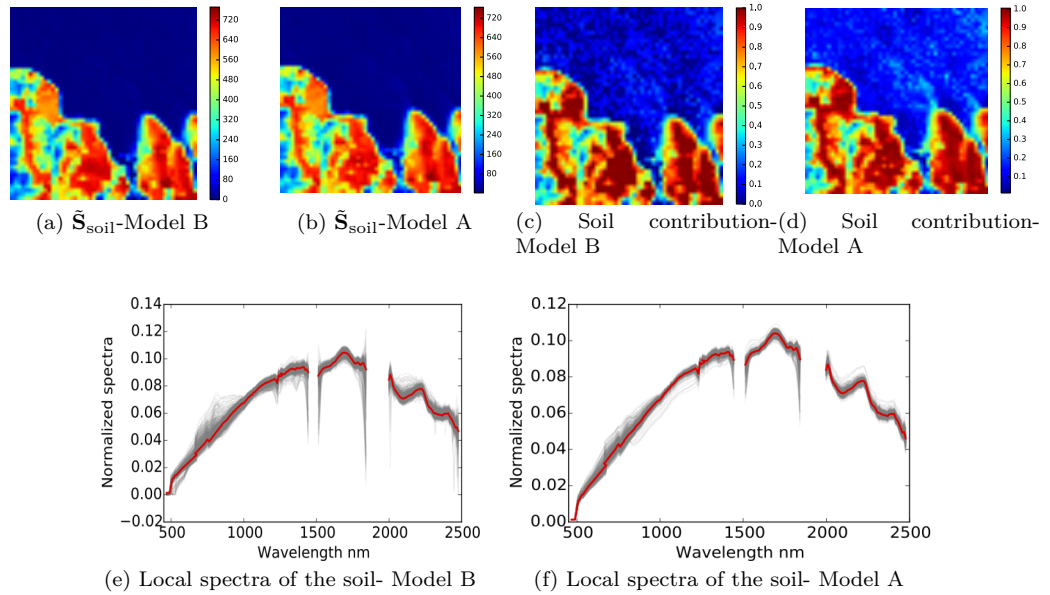


Figure 6.17: Returned components for the soil for the real data set 1. Top row, from left to right: estimated sources model B and then A [127], water contribution to the pixel: $\frac{\mathbf{S}_{soil}^k}{\|\mathbf{S}^k\|_1}$, $\forall k = 1..t$ for model B and A. Bottom: local spectra, with in red the reference spectra \mathbf{A} with model B (left) and model A.

their class (water, soil and vegetation) than the ones belonging to the other classes. Besides, one can observe that the intrinsic variabilities of the spectra of the water are indeed more significant than for the others elements (smallest values for the correlations). Last, even if the soil and the vegetation spectra are inherently quite correlated, it seems that the returned local spectra are not too similar (the separation is fair).

Last, we have observed significant differences for the estimation of the water contribution (abundance maps and spectra). In order to assess the quality of the components retrieved by Model B, we display the normalized spectra of 100 pure pixels of water (namely, 100 pixels from the top of the image, which should be composed only of water) fig.6.21. We can observe that the corresponding local spectra recovered with Model B are much closer to the 'true' local spectra. Consequently, the reference endmember and local spectra recovered by model B are probably more correctly estimated than with model A: this illustrates the possible benefit of the hypersphere constraint.

6.4.4.2 Dataset 2. The results obtained for the second dataset are displayed in fig.6.22.

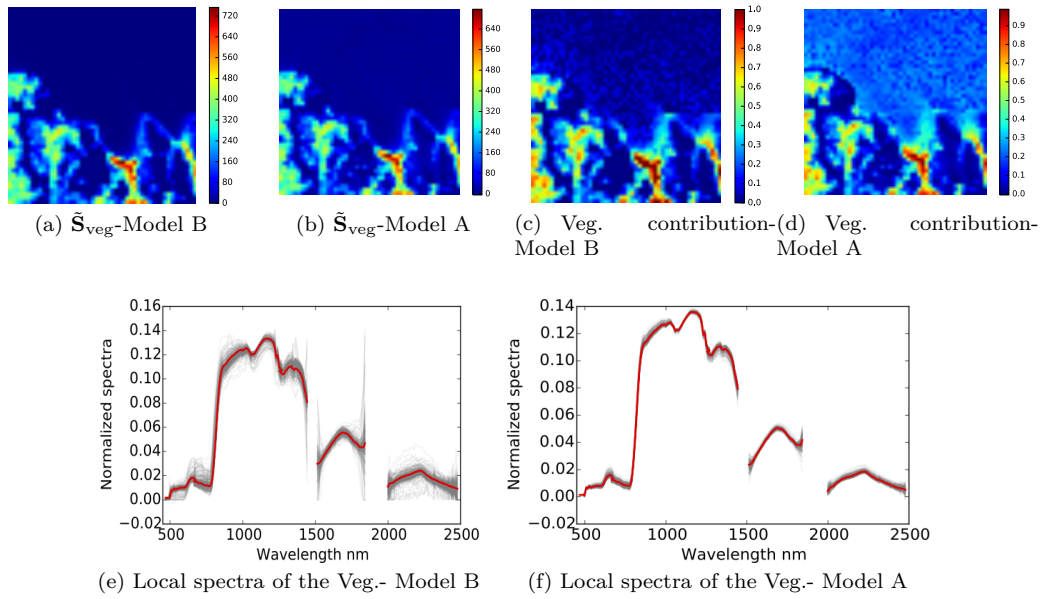


Figure 6.18: Returned components for the vegetation for the real data set 1. Top row, from left to right: estimated sources model B and then A [127], water contribution to the pixel: $\frac{\tilde{\mathbf{S}}_{\text{veg}}^k}{\|\tilde{\mathbf{S}}^k\|_1}$, $\forall k = 1..t$ for model B and A. Bottom: local spectra, with in red the reference spectra $\tilde{\mathbf{A}}$ with model B (left) and model A.




Figure 6.19: Estimated reference endmembers with model A and model B.

Once again, the results obtained by our method or VCA mainly differ for the water estimates. The energy of the water is 20 times smaller than the other components and it has large SV: the impact of the residuals is much more significant for this element fig.6.22a, 6.22b. Notably,




Figure 6.20: Covariance matrix of the returned local endmembers Model B.

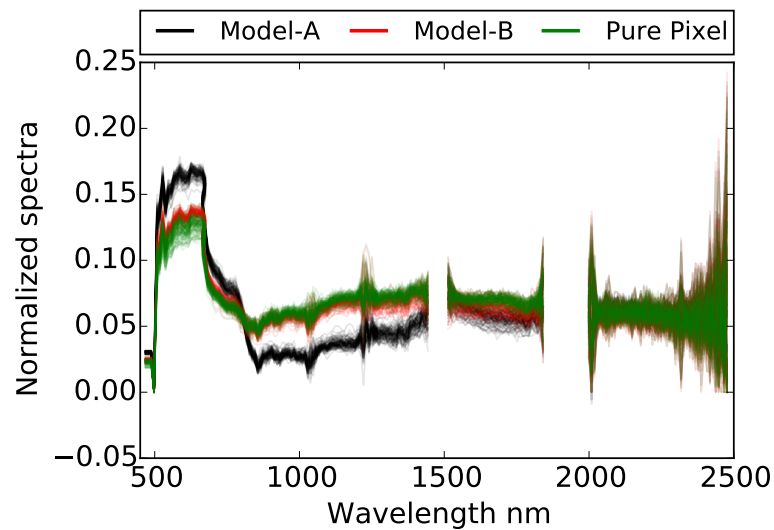


Figure 6.21: Normalized spectra of pure pixels of water and comparison with the corresponding local spectra obtained by Model A and Model B.

one can observe the presence of 2 paths, strongly appearing on the water abundance with our method fig.6.22b, but whose local spectra are very similar to a combination of soil and vegetation with strong deviations. However, the local spectra exhibit acceptable variations (thanks to the

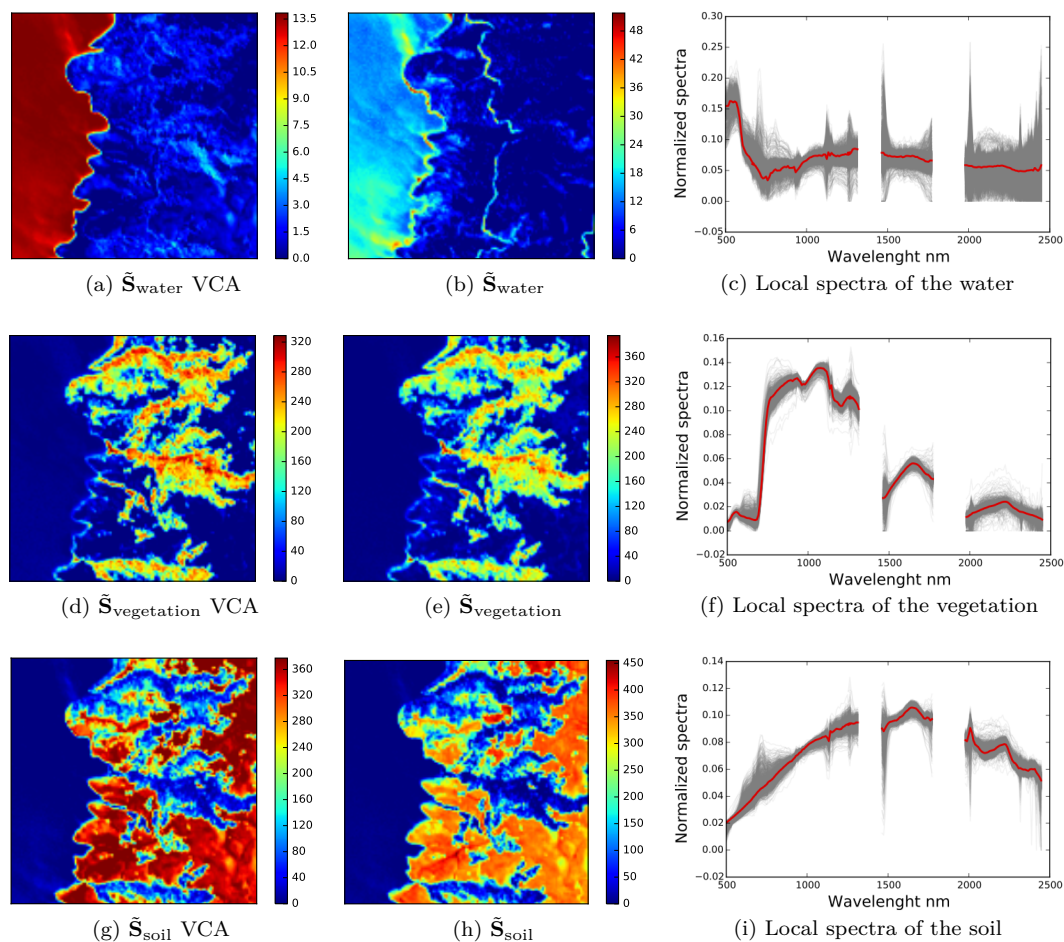


Figure 6.22: Returned components for the real data set 2. Left column: sources obtained from VCA. Middle : estimated sources. Right: local spectra, with in red the reference spectra $\tilde{\mathbf{A}}$.

reweighting, the local spectra of the water are not too affected by the mis-estimation of the abundance) fig.6.22.

6.5 Summary and perspectives

In this chapter, we aimed to extend the robust BSS methods to estimate spectral variations in hyperspectral images. By using the data model proposed first in [127], we introduce another approach to constrain the components, namely by enforcing the local spectra to live on the hypersphere.

In practice, extending the previous methodology developed for robust BSS is not straightforward because the sources of interest, which are not independent in remote sensing, do not admit

highly jointly sparse representations. However, and even if we cannot prevent from the leakages between the spectral variations contribution towards the estimated sources (sparsity cannot be strongly enforced), the preliminary results are encouraging.

In order to improve the estimation of the components for terrestrial hyperspectral images, we should find another prior on the sources to limit the leakages from the spectral variability towards the estimated sources. A prior on smoothness, such as the one proposed in [127] would be of interest.

Future work also includes the applications of the method to astrophysical scenes, for which the simplex assumption does not hold, but whose sources can be (highly) sparsely represented in the wavelet domain. It is an easier framework for this method.

Conclusion

Blind Source Separation is a powerful tool to extract meaningful information from multivalued data. Unfortunately, most BSS methods are hampered by large and structured deviations from the linear mixture model, which are frequently encountered in some domains including astrophysics or remote-sensing. In this thesis, we studied and proposed new robust BSS methods aiming to jointly estimate the mixing matrix, the sources and also these outliers.

The novel methods studied in this thesis rely on a joint estimation of the different components building upon sparse modeling. Emphasis has been placed on the reliability of the proposed methods in a wide range of settings, including the determined case.

In Chapter 4, we focused on the BSS problem in the presence of outliers and sources sharing the same morphology. While the outliers cannot be exactly identified, an accurate estimation of the mixing matrix has been shown to be possible. The robustness of the unmixing process relies on an adaptive detection and penalization of the corrupted samples.

In Chapter 5, we take advantage of the morphological diversity between the deviations and sources which is encountered in many imaging problems. The developed strategy exploits jointly the spatial and spectral diversities between the different components, to precisely separate outliers and sources contributions.

Last, in Chapter 6, we studied the estimation of the spectral variabilities in hyperspectral images. In particular, we compare the influence of two fundamental model assumptions: samples of the abundances 'summing to one' versus spectral signatures belonging to the hypersphere. The sparsity level of the sources of interest does not yield a precise estimation of the components but still leads to comparable results with the state-of-the-art strategies.

We underline that the proposed strategies have been developed to be efficient in a large variety of settings, without requiring a fine tuning of the few parameters. The corresponding codes will be made available at <http://www.cosmostat.org/software/gmcalab>.

Outlook

Several improvements for the proposed works and new perspectives come on the scene during this thesis. Here are some of them.

Algorithmic consideration. Matrix factorization problems using sparse modeling are challenging for two main reasons. First, the problems are non-convex and local minima should be avoided. Second, the parameters should be learned from the data as well. Performing both tasks jointly is not simple. Up to know, the most performing scheme we have developed consists of two steps: i) a broad estimation of the components together with the regularization parameters based on projected least-squares, ii) a refinement step using proximal calculus (PALM).

It would be interesting, in terms of computational time, precision and flexibility, to design an efficient strategy built jointly, on proximal calculus and parameter estimation. This requires a deeper understanding of the dynamics, propagation of the residuals, and sensitivity to the parameters, involved in matrix-factorization using prox-linearization.

Broader variety of deviations. We mainly consider in this work, the presence of deviations well described by an extra linear term, independent of the two variables of interest \mathbf{A} and \mathbf{S} . This model is in particular, well suited to handle the presence of unexpected physical events, but not effective enough to handle more complex deviations, such as component variabilities (spatial or spectral variations).

More generally, the observations can be related to the source contribution with the following generic model:

$$\mathbf{X} = \mathcal{F}(\mathbf{A}, \mathbf{S}),$$

where $\mathcal{F} : \mathbb{R}^{m \times n} \times \mathbb{R}^{n \times t} \rightarrow \mathbb{R}^{m \times t}$ is a degrading mapping.

We mention two special cases of degrading mappings which are frequently encountered in practice but not well handled, yet:

- **Component variabilities.** Several types of component variabilities can be met: spectral (variations of the spectral signature of a component from one pixel to another) or even temporal/spatial (variations of the component itself from one channel to another one - for example due to a delay between the observation).

Various potential improvements can be studied. First, if one keeps the hypersphere constraint for the spectra, deploying adapted tools such as wavelets on the hypersphere, could be relevant. Besides, it would be also interesting to better exploit the fine spectral details of hyperspectral images by using a sparsity constraint on the spectra represented in a transform

domain as well [102]: it would be easier to denoise and to control the shape of the recovered spectra. However, these improvement steps would only be beneficial if another hypothesis (e.g. shape conservation or spatial/spectral smoothness of the variabilities) is meaningful for the data for interest.

In particular, in some applications, the local spectra we aim to recover correspond to shifted versions of the reference spectral signature. In such cases, it would be appropriate to directly seek for these special features by minimizing the displacements, [101].


Also, the underlying physical process associated with the spectral variations is quite well known for the Planck data. More precisely, the emissions depend on some physical parameters, such as the local temperature, or gas densities in the sky. Instead of estimating the emissions, it would be more efficient and precise to look directly for these physical parameters, building upon an improved physical modeling of the observations. That is, the unknown mapping $\mathcal{F}(\mathbf{A}, \mathbf{S})$ can be efficiently replaced by $\mathcal{G}(\mathbf{Y}_1, \dots, \mathbf{Y}_n)$, where $\mathbf{Y}_1, \dots, \mathbf{Y}_n$ are the sought after physical parameters, and \mathcal{G} is a fairly well known mapping [55].

- Other noise process In many imaging applications including in astrophysics, too few photons are reaching the sensors: the noise is not Gaussian. In this case, the observations are better described by a Poisson process, with $\mathbf{X} = \mathcal{P}(\mathbf{A}\mathbf{S})$. This noise, which affects randomly the observations, is not well handled by the standard BSS methods. The combination of the latest developments in optimization with an adaptation of the heuristics employed for sparse BSS in the presence of Gaussian noise may lead to encouraging results.


For illustrative purposes, we mention the observations of the supernova Cassiopeia A, taken by the Chandra telescope⁵, fig.6.23. The noise, at these high energies, mainly originates from a counting process, fig.6.23a⁶. The observations are thus well described by a Poisson process. Besides, the expansion of the supernova leads to spectral variabilities, broadly speaking due to Doppler effect, fig.6.23c. Estimating precisely the components and spectral variabilities from a Poisson process is a challenging task, from which it would be possible to map the different components and infer the geometry of the expansion of the supernova.

⁵<http://chandra.harvard.edu>


⁶Extracted from <https://github.com/facero/BHESS>.



(a) Signatures of some components.



(b) Observation at 5keV



(c) Normalized signatures of the iron

Figure 6.23: Observations of Cassiopeia A by Chandra. Top: spectral signatures of some components of the supernova. Bottom-left: observation at 5keV. Bottom-right: some recovered spectral signatures of the iron (blue and green, reference in red).

Appendix A

Proximal Operators

In the following table, we present the different functions that are used in this thesis and their associated proximal operators.

Similarly to $\|\Lambda \cdot \Phi_S^T\|_1$, we do not find a closed form formulation for $\|\mathcal{Y} \cdot \Phi_O^T\|_{2,1}$ when Φ_O is not

Function	Proximal operator
$\chi_{\mathbf{Y}: \ \mathbf{Y}\ _2 \leq 1}(\mathbf{X})$	$\mathbf{X}' : (\mathbf{X}')^i = \frac{\mathbf{X}^i}{\max(1, \ \mathbf{X}^i\ _2)} \forall i$ [38]
$\ \Lambda \odot \mathbf{X}\ _1$	$S_\Lambda(\mathbf{X})$ [38]
$\ \Lambda \odot \mathbf{X} \Phi_S^T\ _1$	$S_\Lambda(\mathbf{X} \Phi_S^T) \Phi_S$ [124] (exact if Φ_S is orthonormal and good approximation if transformation with diagonally dominant Gram matrix)
$\ \mathcal{Y} \odot \mathbf{X}\ _{2,1}$	$\mathbf{X}' : (\mathbf{X}')^i = \mathbf{X}^i \times \left(1 - \frac{\gamma^i}{\ (\mathbf{X})^i\ _2}\right)_+$, $\forall i$, [83].
$\ \mathcal{Y} \odot \mathbf{X} \Phi_O^T\ _{2,1}$	$\mathbf{X}' \Phi_O : (\mathbf{X}')^i = (\mathbf{X} \Phi_O^T)^i \times \left(1 - \frac{\gamma^i}{\ (\mathbf{X} \Phi_O^T)^i\ _2}\right)_+$, $\forall i$ (exact if Φ_O is orthonormal and good approximation if transformation with diagonally dominant Gram matrix)
$\ \mathcal{Y} \odot \mathbf{X}\ _{2,1} + \chi_{\mathbf{Y}: \mathbf{Y} \geq 0}(\mathbf{X})$	$\mathbf{X}' : (\mathbf{X}')^i = \mathbf{X}_+^i \times \left(1 - \frac{\gamma^i}{\ (\mathbf{X})_+^i\ _2}\right)_+$, $\forall i$, [137, Theorem 1].

orthonormal. In the spirit of the approximation made for the ℓ_1 norm, we propose to threshold the columns of the expansion coefficients, and then come back to the domain of observations. In practice, these approximations made to handle sparsity in a transformed domain give better results than the synthesis formulation, what supports the use of these approximations.

Appendix B

Threshold value

Let us consider the following robust BSS problem:

$$\underset{\mathbf{A}, \mathbf{S}, \mathbf{O}}{\text{minimize}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS} - \mathbf{O}\|_2^2 + \sum_{i=1}^n \lambda_i \|\mathbf{S}_i\|_1 + \beta \|\mathbf{O}\|_p$$

Setting the values of the parameters $\{\lambda\}_{i=1..n}$ and β is crucial. Thanks to sparse modeling, the choice of these parameters can be based on noise removal theory. Indeed, estimating the outliers or the sources, w.r.t the other variables amounts to solve a detection problem as described in the following.

B.1 Noise removal/detection with sparse modeling

Sparse modeling is particularly suitable for Gaussian noise removal [49], [36], [53]. For illustrative purpose, we consider a sparse 1D signal $y \in \mathbb{R}^{1 \times t}$, and the noisy observation $x = y + n$, where n accounts for the presence of Gaussian noise, whose entries are i.i.d., with standard deviation σ , see fig.B.1.

The energy of the sparse signal y is concentrated in few active samples, whereas the one of the Gaussian noise n is broadly distributed over all the samples. In order to retrieve the signal y , one can then select only the largest samples of x , *i.e.* the ones corresponding to the support of y . This can be achieved by minimizing the following problem:

$$\underset{y}{\operatorname{argmin}} \frac{1}{2} \|x - y\|_2^2 + \lambda \|y\|_1,$$

whose closed form solution is given by: $\tilde{y} = \mathcal{S}_\lambda(x)$.

That is, the signal y is recovered from x by thresholding/putting to zero, the entries of x smaller than λ : the proposed value of λ in fig.B.1 would select properly the support of y . The choice




Figure B.1: Illustration of the noisy observation x (in blue) of a sparse signal y (in red). The value of λ should be based on the value of the noise.

of λ can be seen as an hypothesis testing problem [124]: an entry x^t corresponds only to the Gaussian noise contribution n , or not and thus corresponds to a noisy active entry of the signal $y + n$.

The “ k - σ ” rule is commonly used for setting the value of λ [124]: the probability that an entry with an amplitude larger than 3σ corresponds only to the Gaussian contribution is 0.4% only. The value of k depends on the application: a large k is more conservative (less false recoveries) whereas a smaller k induces more detections.

The standard deviation of the Gaussian noise n is generally not known in practice, but can be estimated with the median absolute deviation (mad) of x . Indeed, the mad of n corresponds to a good approximation of σ , and given that the mad operator is almost not influenced by sparse contribution, we end up with $\sigma \approx 1.4826 \times \text{mad}(n) \approx 1.4826 \text{mad}(x)$ since $\text{mad}(y) \approx 0$ if y is sparse enough.

If the signals are sparsely represented in a transformed domain Φ , the rationale and setting of the parameters are kept identical, (performed in Φ). In case of multi-resolution transforms, a threshold value can be chosen on each scale.

B.2 Noise removal for the sources

Estimating the sources, for fixed \mathbf{A} and \mathbf{O} amounts to solve the following problem:

$$\underset{\mathbf{S}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{X} - \mathbf{O} - \mathbf{AS}\|_2^2 + \sum_{i=1}^n \lambda_i \|\mathbf{S}_i\|_1,$$

which is reminiscent of the previous denoising problem. However, we aim to jointly separate and denoise the sources.

In this thesis, it is assumed that the entries of \mathbf{N} are Gaussian i.i.d.. This does not imply that

the Gaussian noise contribution in the source domain (*i.e.* $\mathbf{A}^\dagger \mathbf{N}$) has the same statistics from one source to another. That is why, a threshold value is computed for every source.

With pseudo-inversion

In the fast version of GMCA, or AMCA, the sources are estimated first from the projection: $\tilde{\mathbf{S}} = \mathbf{A}^\dagger (\mathbf{X} - \mathbf{O})$, and then thresholded. If \mathbf{A} is fairly well recovered, each projected source corresponds to $\tilde{\mathbf{S}}_i = \mathbf{S}_i + (\mathbf{A}^\dagger \mathbf{N})_i$, $\forall i = 1, \dots, n$. We recover exactly the setting of B.1, with $y \leftarrow \mathbf{S}_i$ and $n \leftarrow (\mathbf{A}^\dagger \mathbf{N})_i$. For each i , the entries of $(\mathbf{A}^\dagger \mathbf{N})_i$ are indeed Gaussian, but with an unknown associated standard deviation $\sigma \approx 1.4826 \text{mad}(\tilde{\mathbf{S}}_i)$. Consequently, the final threshold value, for each source i , is chosen equal to $k \text{mad}(\tilde{\mathbf{S}}_i)$, and $k \in (1, 5)$ generally.

With proximal gradient updates

A proximal gradient update also is two-steps: (i) a gradient update yielding, broadly, to a noisy estimate of the sources, and (ii) the soft-thresholding. Hence, if the mixing matrix is correctly recovered, the noisy sources, before the soft-thresholding, correspond to $\tilde{\mathbf{S}} = \mathbf{S} + \frac{1}{L} \mathbf{A}^T (\mathbf{X} - \mathbf{O} - \mathbf{AS})$, where $L = \|\mathbf{A}^T \mathbf{A}\|_{2,s}$. These noisy sources can be decomposed with $y \leftarrow \mathbf{S}$ and $n \leftarrow \frac{1}{L} \mathbf{A}^T (\mathbf{X} - \mathbf{O} - \mathbf{AS})$, the projected noise on the sources domain¹. Consequently, the i th source $\tilde{\mathbf{S}}_i$ can be denoised with a threshold equal to $k \text{mad}(\frac{1}{L} [\mathbf{A}^T (\mathbf{X} - \mathbf{O} - \mathbf{AS})]_i)$. This amounts to set $\lambda_i = k \text{mad}([\mathbf{A}^T (\mathbf{X} - \mathbf{O} - \mathbf{AS})]_i)^2$.

We propose also another approach for the derivations of these thresholds. Using BCD or PALM, the minimization problem for fixed \mathbf{A} is properly handled. Consequently, we can use the Karush–Kuhn–Tucker condition:

$$\mathbf{A}^T (\mathbf{X} - \mathbf{AS} - \mathbf{O}) = \lambda \mathbf{P} \text{ with } \mathbf{P} \in \partial \|\mathbf{S}\|_1.$$

The condition $\mathbf{P} \in \partial \|\mathbf{S}\|_1$ implies that $\|\mathbf{P}\|_\infty = 1$, and thus that $\|(\mathbf{A}^T (\mathbf{X} - \mathbf{O} - \mathbf{AS}))_i\|_\infty = \lambda_i$. Since we aim to obtain $\mathbf{X} - \mathbf{O} - \mathbf{AS} \approx \mathbf{N}$, we end up with $\|(\mathbf{A}^T (\mathbf{N}))_i\|_\infty \approx \lambda_i$. By setting $\lambda_i = 3\sigma_i$, where σ_i is the standard deviation of the noise corrupting the i th source $(\mathbf{A}^T \mathbf{N})_i$

¹We point out that one can also use the first order condition to obtain this result. Let $s = \mathbf{S}_j$ be the j th source, $a = \mathbf{A}^j$ its associated column vector, and $\mathbf{R} = \mathbf{X} - \mathbf{O} - \sum_{i=1, i \neq j}^n \mathbf{S}_i \mathbf{A}^i$. The estimation s for fixed \mathbf{R} and a is given by:

$$\underset{s}{\text{argmin}} \frac{1}{2} \|\mathbf{R} - as\|_2^2 + \lambda_i \|s\|_1.$$

This problem has a closed form, which is given by $s = \mathcal{S}_{\lambda_i}(a^T \mathbf{R})$, by using the fact that $a^T a = 1$ since the columns of \mathbf{A} are normalized. If we assume that we know perfectly the other sources and the mixing matrix:

$$\mathbf{R} = \mathbf{AS} + \mathbf{O} + \mathbf{N} - \sum_{i=1, i \neq j}^n \mathbf{S}_i \mathbf{A}^i - \mathbf{O} = as + \mathbf{N}. \text{ Hence, the closed form solution can be written as } s = \mathcal{S}_{\lambda_i}(s + a^T \mathbf{N}),$$

or equivalently, the sources are corrupted by $\mathbf{A}^T \mathbf{N}$.

²We point out that if \mathbf{A} is orthonormal, the value obtained for each source i with the two noise projections $\text{mad}(\mathbf{A}^T ((\mathbf{X} - \mathbf{O}) - \mathbf{AS}))_i$ or $\text{mad}((\mathbf{A}^\dagger (\mathbf{X} - \mathbf{O}))_i)$ are equivalent.

(approximated with the mad operator on $\mathbf{A}^T(\mathbf{X} - \mathbf{O} - \mathbf{AS})$), we indeed ensure that more than 99% of the corrupted entries will be set to zero, which provides an approximation of $\|\mathbf{A}^T(\mathbf{N})_i\|_\infty$.

B.3 Noise removal for the outliers

The update of the outliers, for fixed \mathbf{A} and \mathbf{S} is given by:

$$\operatorname{argmin}_{\mathbf{O}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS} - \mathbf{O}\|_2^2 + \beta \|\mathbf{O}\|_p,$$

where p depends on the sparsity pattern of the outliers: $p = (2, 1)$ for column-sparse outliers, and $p = 1$ for row and column sparse outliers.

B.3.1 With the ℓ_1 norm

In contrast to the sources, the outliers are not mixed and are thus estimated directly with $\mathcal{S}_\beta(\mathbf{X} - \mathbf{AS})$. Since we assumed that \mathbf{N} is i.i.d, the noise level should be constant from one observation to another. Hence, the threshold β is chosen to be equal to $k \operatorname{mad}(\mathbf{X} - \mathbf{AS} - \mathbf{O})$ similarly to B.1, since $\operatorname{mad}(\mathbf{X} - \mathbf{AS} - \mathbf{O}) \approx \operatorname{mad}(\mathbf{N})$.

B.3.2 With the $\ell_{2,1}$ norm

The $\ell_{2,1}$ norm promotes column-sparse outliers. In that case, the detection process is performed on the norm of the samples: this better preserves the structure of the component and makes the detection easier see fig.B.2.

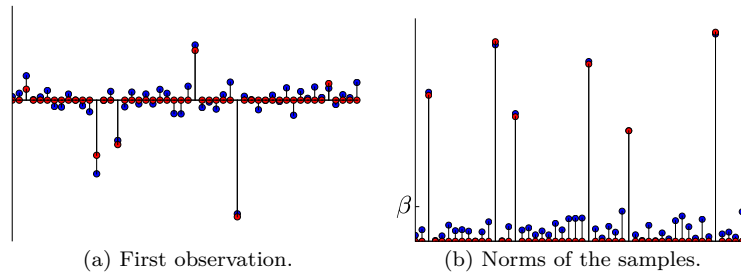


Figure B.2: Illustration of noisy observations (in blue) of column sparse outliers (in red). On the left: observation of only one channel, on the right: observation of the norms of the samples. The value of β should be chosen according to the contribution of the Gaussian noise.

The estimation of the outliers is performed by minimizing:

$$\operatorname{argmin}_{\mathbf{O}} \frac{1}{2} \|\mathbf{X} - \mathbf{AS} - \mathbf{O}\|_2^2 + \beta \|\mathbf{O}\|_{2,1}.$$

Every sample $k \in 1, \dots, t$ of the closed form solution is such that $(\mathbf{X} - \mathbf{AS})^k \times \left(1 - \frac{\beta}{\|(\mathbf{X} - \mathbf{AS})^k\|_2}\right)_+$. Hence, the value of β should be chosen so that, with a high probability, the samples having an energy larger than β correspond to the outliers, the signal of interest, and the ones with an energy smaller than β to Gaussian noise only.

The Gaussian noise contribution (the ℓ_2 norm of a column vector of \mathbf{N}) follows a Chi distribution with m degrees of freedom. Similarly to what was done in the previous sections, it would be possible to set the value of β so that, the probability (chosen) that a sample with a norm larger than β correspond to a Chi distribution with m degrees of freedom is very small.

We have rather preferred using a more conservative threshold, related to the $k\sigma$ residual of the sources. The expected value of a Chi distribution with m degrees of freedom generated from a Gaussian law with standard deviation σ is $\sigma \times \sqrt{2} \times \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})}$. The value of β in this thesis is $k \times \sigma \times \sqrt{2} \times \frac{\Gamma(\frac{m+1}{2})}{\Gamma(\frac{m}{2})}$, where $\sigma \approx 1.4826 \text{ mad}(\mathbf{X} - \mathbf{AS} - \mathbf{O})$.

Appendix C

Performance criteria

Assessing the quality of the separation and estimation of the sources is necessary for comparing different strategies. Many criteria have been proposed in the literature, which do not provide similar insights on the performances. In this appendix, we present the criteria that have been used in this thesis along with a description of the features of the algorithms they evaluate.

We underline that in BSS, the algorithms should yield an accurate unmixing (estimation of \mathbf{A}) and a precise recovery of the sources (estimation of \mathbf{S}). These raise two remarks: first, a correct estimation of the mixing matrix may not lead to a precise estimation of the sources (in the presence of outliers or significant Gaussian noise), and second, a fair recovery of \mathbf{AS} does not mean that \mathbf{A} and \mathbf{S} are correctly recovered. For instance, in the noiseless determined case, any invertible matrix $\tilde{\mathbf{A}}$, and corresponding sources $\tilde{\mathbf{A}}^\dagger \mathbf{X}$ perfectly fit the data, whereas the unmixing can be completely wrong.

Unmixing.

- For each recovered $\tilde{\mathbf{A}}$, the quantity $\Delta_A = \left(\frac{\|\tilde{\mathbf{A}}^\dagger \mathbf{A} - \mathbf{I}\|_1}{n^2} \right)$ is computed [14] (or the inverse of this quantity in dB in Chapter 5). This provides a global criterion for the quality of the unmixing. However, a fair but inaccurate \mathbf{A} cannot be easily distinguish from a mixing matrix \mathbf{A} , where all but one column have been very accurately estimated.
- For each recovered $\tilde{\mathbf{A}}$, the maximal angle made between each of its column and the ones of the initial \mathbf{A} is return: $\max_{j=1..n} \arccos\langle \tilde{\mathbf{A}}^j, \mathbf{A}^j \rangle$ (in degree).

We sum the number of runs for which an algorithm has returned a mixing matrix whose maximal angle is smaller than 5 degrees. This quantity, normalized to 1, provides a good indicator of the reliability of the algorithms. Besides, it allows for a better interpretation of the quantity Δ_A by differentiating between an inaccurate (all the columns are fairly estimated) and a 'wrong' (all but few columns are precisely estimated) unmixing.

Estimation of the sources.

- In [131], the authors decompose each retrieved source s as the sum:

$$s = s_{target} + s_{interference} + s_{noise} + s_{artifacts},$$

where s_{target} denotes the projection of the retrieved source on the sought-after one, $s_{interference}$ the residue due to the interferences with the other sources, s_{noise} accounts for the part due to the presence of noise (the outliers in our case), and last, $s_{artifacts}$, represents the remaining artifacts. This decomposition is used to derive the following indicators [131]:

- Signal to Distortion Ratio $SDR(s) = 20 \log \left(\frac{\|s_{target}\|_2}{\|s_{interference} + s_{noise} + s_{artifacts}\|_2} \right)$.
- Signal to Interference Ratio $SIR(s) = 20 \log \left(\frac{\|s_{target}\|_2}{\|s_{interference}\|_2} \right)$.
- Signal to Noise Ratio $SNR(s) = 20 \log \left(\frac{\|s_{target} + s_{interference}\|_2}{\|s_{noise}\|_2} \right)$.
- Signal to Artifact Ratio $SAR(s) = 20 \log \left(\frac{\|s_{target} + s_{interference} + s_{noise}\|_2}{\|s_{artifacts}\|_2} \right)$.

If not stated otherwise, the median over the n sources are displayed.

Appendix D

Publications

Refereed articles

- C.Chenot and J.Bobin. Blind Source Separation with outliers in transformed domains. *Submitted*
- C.Chenot and J.Bobin. Blind separation of sparse sources in the presence of outliers. *Signal Processing*, Elsevier, vol.138, 2017.
- C.Chenot, J.Bobin and J.Rapin. Robust sparse blind source separation. *SPL*, IEEE, vol.22, 2015.

Conferences

- C.Chenot and J.Bobin. BSS with Corrupted Data in Transformed Domains. *LVA-ICA 2017*, Grenoble, France.
- C.Chenot and J.Bobin. Sparse BSS with corrupted data in transformed domains. *iTwist'16*, Aalborg, Denmark.
- C.Chenot, J.Bobin and J.Rapin. Sparse BSS in the presence of outliers. *SPARS 2015*, Cambridge, UK.

Bibliography

- [1] A. ADLER, M. ELAD, Y. HEL-OR, AND E. RIVLIN, Sparse coding with anomaly detection, in *Machine Learning for Signal Processing (MLSP)*, 2013 IEEE International Workshop on, IEEE, 2013, pp. 1–6. 80
- [2] Y. ALTMANN, Nonlinear unmixing of hyperspectral images, PhD thesis, Institut National Polytechnique de Toulouse-INPT, 2013. 22
- [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*, 21 (2012), pp. 3017–3025. 142
- [4] Y. ALTMANN, S. MCCLAUGHLIN, AND A. HERO, Robust Linear Spectral Unmixing Using Anomaly Detection, *IEEE Transactions on Computational Imaging*, 1 (2015), pp. 74–85, doi:10.1109/TCI.2015.2455411. 36, 142
- [5] S. AMINI, M. SADEGHI, M. JONEIDI, M. BABAIE-ZADEH, AND C. JUTTEN, Outlier-aware dictionary learning for sparse representation, in *Machine Learning for Signal Processing (MLSP)*, 2014 IEEE International Workshop on, IEEE, 2014, pp. 1–6. 80
- [6] H. ATTOUCH, J. BOLTE, P. REDONT, AND A. SOUBEYRAN, Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the Kurdyka-Lojasiewicz inequality, *Mathematics of Operations Research*, 35 (2010), pp. 438–457. 52, 54, 56
- [7] C. AUBEL, C. STUDER, G. POPE, AND H. BOLCSKEI, Sparse signal separation in redundant dictionaries, in *Information Theory Proceedings (ISIT)*, 2012 IEEE International Symposium on, IEEE, 2012, pp. 2047–2051. 110
- [8] A. BECK AND M. TEOULLE, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, *SIAM Journal on Imaging Sciences*, 2 (2009), pp. 183–202. 54, 55, 117
- [9] 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*, 1 (2013), pp. 6–36. 3

- [10] 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, *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of*, 5 (2012), pp. 354–379. 2, 20, 139
- [11] T. BLUMENSATH AND M. E. DAVIES, Iterative hard thresholding for compressed sensing, *Applied and computational harmonic analysis*, 27 (2009), pp. 265–274. 78
- [12] J. BOBIN, Y. MOUDDEN, J.-L. STARCK, AND M. ELAD, Morphological diversity and source separation, *IEEE Signal Processing Letters*, 13 (2006), pp. 409–412. 10
- [13] J. BOBIN, J. RAPIN, A. LARUE, AND J.-L. STARCK, Sparsity and Adaptivity for the Blind Separation of Partially Correlated Sources, *Signal Processing, IEEE Transactions on*, 63 (2015), pp. 1199–1213, [doi:10.1109/TSP.2015.2391071](https://doi.org/10.1109/TSP.2015.2391071). 76, 82, 83, 87, 88, 89, 90, 95, 96, 104, 131
- [14] J. BOBIN, J.-L. STARCK, J. FADILI, AND Y. MOUDDEN, Sparsity and morphological diversity in blind source separation, *Image Processing, IEEE Transactions on*, 16 (2007), pp. 2662–2674. 7, 11, 12, 13, 15, 16, 17, 19, 25, 45, 56, 57, 59, 65, 67, 73, 76, 84, 97, 120, 179
- [15] J. BOBIN, J.-L. STARCK, J. FADILI, Y. MOUDDEN, AND D. DONOHO, Morphological Component Analysis: An adaptive thresholding strategy, *IEEE Trans. On Image Processing*, 16 (2007), pp. 2675 – 2681. 10, 11
- [16] J. BOBIN, F. SUREAU, J.-L. STARCK, A. RASSAT, AND P. PAYKARI, Joint Planck and WMAP CMB map reconstruction, *A&A*, 563 (2014). 2, 3, 7
- [17] J. BOLTE, S. SABACH, AND M. TEBoulLE, Proximal alternating linearized minimization for nonconvex and nonsmooth problems, *Mathematical Programming*, 146 (2014), pp. 459–494. 43, 51, 55, 56
- [18] T. BOUWMANS, A. SOBRAL, S. JAVED, S. K. JUNG, AND E.-H. ZAHZAH, Decomposition into low-rank plus additive matrices for background/foreground separation: A review for a comparative evaluation with a large-scale dataset, *Computer Science Review*, 23 (2017), pp. 1 – 71. 31
- [19] S. BOYD AND L. VANDENBERGHE, Convex optimization, Cambridge university press, 2004. 40
- [20] P. H. CALAMAI AND J. J. MORÉ, Projected gradient methods for linearly constrained problems, *Mathematical programming*, 39 (1987), pp. 93–116. 42
- [21] E. CANDÉS AND B. RECHT, Exact matrix completion via convex optimization, *Communications of the ACM*, 55 (2012), pp. 111–119. 30

- [22] E. J. CANDÈS, Ridgelets: theory and applications, PhD thesis, Stanford University, 1998. 9
- [23] E. J. CANDÈS, X. LI, Y. MA, AND J. WRIGHT, Robust Principal Component Analysis?, *Journal of the ACM (JACM)*, 58 (2011), p. 11. 27, 29, 30, 32, 33, 37, 38, 84, 96, 98, 112
- [24] E. J. CANDÈS AND T. TAO, Decoding by linear programming, *IEEE transactions on information theory*, 51 (2005), pp. 4203–4215. 78
- [25] E. J. CANDÈS AND M. B. WAKIN, An introduction to compressive sampling, *IEEE signal processing magazine*, 25 (2008), pp. 21–30. 78, 142
- [26] E. J. CANDÈS, M. B. WAKIN, AND S. P. BOYD, Enhancing sparsity by reweighted ℓ_1 minimization, *Journal of Fourier analysis and applications*, 14 (2008), pp. 877–905. 71, 113, 118
- [27] J.-F. CARDOSO, Blind signal separation: statistical principles, *Proceedings of the IEEE*, 86 (1998), pp. 2009–2025. 6
- [28] J.-F. CARDOSO, Dependence, correlation and gaussianity in independent component analysis, *Journal of Machine Learning Research*, 4 (2003), pp. 1177–1203. 7
- [29] V. CHANDRASEKARAN, S. SANGHAVI, P. A. PARRILO, AND A. S. WILLSKY, Rank-sparsity incoherence for matrix decomposition, *SIAM Journal on Optimization*, 21 (2011), pp. 572–596. 30, 33, 38
- [30] P. CHARBONNIER, L. BLANC-FÉRAUD, G. AUBERT, AND M. BARLAUD, Deterministic edge-preserving regularization in computed imaging, *IEEE Transactions on image processing*, 6 (1997), pp. 298–311. 35
- [31] C. CHENOT AND J. BOBIN, Blind separation of sparse sources in the presence of outliers, *Signal Processing*, 138 (2017), pp. 233 – 243. 76, 86, 112
- [32] C. CHENOT AND J. BOBIN, BSS with corrupted data in transformed domains, in *International Conference on Latent Variable Analysis and Signal Separation*, Springer, 2017, pp. 542–552. 108
- [33] C. CHENOT, J. BOBIN, AND J. RAPIN, Robust Sparse Blind Source Separation, *Signal Processing Letters, IEEE*, 22 (2015), pp. 2172–2176. 25, 76, 80
- [34] E. CHOZENOUX, J.-C. PESQUET, AND A. REPETTI, A block coordinate variable metric forward-backward algorithm, *Journal of Global Optimization*, (2016), pp. 1–29. 55
- [35] A. CICHOCKI, S. CRUCES, AND S.-I. AMARI, Generalized alpha-beta divergences and their application to robust nonnegative matrix factorization, *Entropy*, 13 (2011), pp. 134–170. 34, 37

- [36] R. R. COIFMAN AND D. L. DONOHO, Translation-invariant denoising, Springer, 1995. 9, 174
- [37] P. COLLABORATION, Planck: The scientific program, ESA publication ESA-SCI (2005)/01, (2005). 3
- [38] P. L. COMBETTES AND V. R. WAJS, Signal recovery by proximal forward-backward splitting, *Multiscale Modeling & Simulation*, 4 (2005), pp. 1168–1200. 42, 49, 114, 173
- [39] P. COMON, Independent Component Analysis, a new concept?, *Signal processing*, 36 (1994), pp. 287–314. 6, 33
- [40] P. COMON AND C. JUTTEN, Handbook of Blind Source Separation: Independent component analysis and applications, Academic press, 2010. 6, 7, 19, 33
- [41] S. CRUCES AND A. CICHOCKI, Blind source extraction in gaussian noise, in *Proc. Int. Workshop Independent Component Analysis and Blind Signal Separation*, 2000, pp. 63–68. 19
- [42] I. DAUBECHIES ET AL., Ten lectures on wavelets, vol. 61, SIAM, 1992. 9
- [43] M. DAVIES, Identifiability issues in noisy ICA, *IEEE Signal processing letters*, 11 (2004), pp. 470–473. 7
- [44] J. DELABROUILLE AND ET AL., The pre-launch Planck Sky Model: a model of sky emission at submillimetre to centimetre wavelengths, *Astronomy & Astrophysics*, 553 (2013), A96, p. A96, [doi:10.1051/0004-6361/201220019](https://doi.org/10.1051/0004-6361/201220019), [arXiv:1207.3675](https://arxiv.org/abs/1207.3675). 129
- [45] N. DOBIGEON, J.-Y. TURNERET, C. RICHARD, J. C. M. BERMUDEZ, S. MCLAUGHLIN, AND A. O. HERO, Nonlinear unmixing of hyperspectral images: Models and algorithms, *IEEE Signal Processing Magazine*, 31 (2014), pp. 82–94. 22, 141
- [46] D. DONOHO AND V. STODDEN, When does non-negative matrix factorization give a correct decomposition into parts?, in *Advances in Neural Information Processing Systems 16*, MIT Press, 2004. 7
- [47] D. L. DONOHO, M. ELAD, AND V. N. TEMLYAKOV, Stable recovery of sparse overcomplete representations in the presence of noise, *IEEE Transactions on information theory*, 52 (2006), pp. 6–18. 79, 110, 111
- [48] D. L. DONOHO AND X. HUO, Uncertainty principles and ideal atomic decomposition, *IEEE Transactions on Information Theory*, 47 (2001), pp. 2845–2862. 16
- [49] D. L. DONOHO AND I. M. JOHNSTONE, Ideal spatial adaptation by wavelet shrinkage, *biometrika*, (1994), pp. 425–455. 174

- [50] L. DU, X. LI, AND Y.-D. SHEN, Robust nonnegative matrix factorization via half-quadratic minimization, in Data Mining (ICDM), 2012 IEEE 12th International Conference on, IEEE, 2012, pp. 201–210. 35
- [51] S. EGUCHI AND Y. KANO, Robustifying maximum likelihood estimation. 34
- [52] M. ELAD, Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing, Springer, 2010. 7, 10
- [53] M. ELAD AND M. AHARON, Image denoising via sparse and redundant representations over learned dictionaries, IEEE Transactions on Image processing, 15 (2006), pp. 3736–3745. 174
- [54] M. ELAD, J.-L. STARCK, P. QUERRE, AND D. L. DONOHO, Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA), Applied and Computational Harmonic Analysis, 19 (2005), pp. 340–358. 10, 112
- [55] H. ERIKSEN, J. JEWELL, C. DICKINSON, A. BANDAY, K. GÓRSKI, AND C. LAWRENCE, Joint Bayesian component separation and CMB power spectrum estimation, The Astrophysical Journal, 676 (2008), p. 10. 171
- [56] C. FÉVOTTE, N. BERTIN, AND J.-L. DURRIEU, Nonnegative matrix factorization with the Itakura-Saito divergence: With application to music analysis, Neural computation, 21 (2009), pp. 793–830. 34
- [57] C. FÉVOTTE AND N. DOBIGEON, Nonlinear hyperspectral unmixing with robust nonnegative matrix factorization, arXiv preprint arXiv:1401.5649, (2014). 36, 37, 38
- [58] C. FEVOTTE AND N. DOBIGEON, Nonlinear Hyperspectral Unmixing With Robust Nonnegative Matrix Factorization, Image Processing, IEEE Transactions on, 24 (2015), pp. 4810–4819, doi:10.1109/TIP.2015.2468177. 20, 36, 102, 112, 113, 121, 142
- [59] C. FÉVOTTE AND J. IDIER, Algorithms for nonnegative matrix factorization with the β -divergence, Neural computation, 23 (2011), pp. 2421–2456. 34
- [60] P. A. FORERO, S. SHAFER, AND J. D. HARGUESS, Sparsity-driven Laplacian-regularized outlier identification for dictionary learning, IEEE Transactions on Signal Processing, 65 (2017), pp. 3803–3817. 80
- [61] J. FRIEDMAN, T. HASTIE, AND R. TIBSHIRANI, A note on the group lasso and a sparse group lasso, arXiv preprint arXiv:1001.0736, (2010). 42, 43
- [62] N. GADHOK AND W. KINSNER, Rotation sensitivity of independent component analysis to outliers, in Electrical and Computer Engineering, 2005. Canadian Conference on, IEEE, 2005, pp. 1437–1442. 25, 33

- [63] N. GADHOK AND W. KINSNER, An Implementation of β -Divergence for Blind Source Separation, in *Electrical and Computer Engineering, 2006. CCECE'06. Canadian Conference on, IEEE, 2006*, pp. 1446–1449. 84, 96, 122
- [64] N. GILLIS AND A. KUMAR, Exact and heuristic algorithms for semi-nonnegative matrix factorization, *SIAM Journal on Matrix Analysis and Applications*, 36 (2015), pp. 1404–1424. 71
- [65] N. GILLIS AND R. J. PLEMMONS, Sparse nonnegative matrix underapproximation and its application to hyperspectral image analysis, *Linear Algebra and its Applications*, 438 (2013), pp. 3991–4007. 140
- [66] R. GRIBONVAL, Should penalized least squares regression be interpreted as maximum a posteriori estimation?, *IEEE Transactions on Signal Processing*, 59 (2011), pp. 2405–2410. 7
- [67] R. GRIBONVAL, R. JENATTON, AND F. BACH, Sparse and spurious: dictionary learning with noise and outliers, *IEEE Transactions on Information Theory*, 61 (2015), pp. 6298–6319. 26
- [68] R. GRIBONVAL AND S. LESAGE, A survey of sparse component analysis for blind source separation: principles, perspectives, and new challenges, in *ESANN'06 proceedings-14th European Symposium on Artificial Neural Networks, d-side publi., 2006*, pp. 323–330. 7
- [69] E. T. HALE, W. YIN, AND Y. ZHANG, Fixed-point continuation applied to compressed sensing: implementation and numerical experiments, *Journal of Computational Mathematics*, (2010), pp. 170–194. 65
- [70] A. HALIMI, Y. ALTMANN, G. S. BULLER, S. McLAUGHLIN, W. OXFORD, D. CLARKE, AND J. PIPER, Robust unmixing algorithms for hyperspectral imagery, in *Sensor Signal Processing for Defence (SSPD), 2016, IEEE, 2016*, pp. 1–5. 25
- [71] A. HALIMI, Y. ALTMANN, N. DOBIGEON, AND J.-Y. TOURNERET, Nonlinear unmixing of hyperspectral images using a generalized bilinear model, *IEEE Transactions on Geoscience and Remote Sensing*, 49 (2011), pp. 4153–4162. 142
- [72] A. HALIMI, N. DOBIGEON, AND J.-Y. TOURNERET, Unsupervised unmixing of hyperspectral images accounting for endmember variability, *IEEE Transactions on Image Processing*, 24 (2015), pp. 4904–4917. 142
- [73] A. HALIMI, P. HONEINE, AND J. M. BIOUCAS-DIAS, Hyperspectral unmixing in presence of endmember variability, nonlinearity, or mismodeling effects, *IEEE Transactions on Image Processing*, 25 (2016), pp. 4565–4579. 142

- [74] A. B. HAMZA AND D. J. BRADY, Reconstruction of reflectance spectra using robust nonnegative matrix factorization, *Signal Processing, IEEE Transactions on*, 54 (2006), pp. 3637–3642. 35
- [75] K. S. HE, D. ROCCHINI, M. NETELER, AND H. NAGENDRA, Benefits of hyperspectral remote sensing for tracking plant invasions, *Diversity and Distributions*, 17 (2011), pp. 381–392. 3
- [76] P. J. HUBER ET AL., Robust estimation of a location parameter, *The Annals of Mathematical Statistics*, 35 (1964), pp. 73–101. 35
- [77] M.-D. IORDACHE, J. M. BIUCAS-DIAS, AND A. PLAZA, Sparse unmixing of hyperspectral data, *IEEE Transactions on Geoscience and Remote Sensing*, 49 (2011), pp. 2014–2039. 142
- [78] C. JUTTEN AND J. HERAULT, Blind separation of sources, part I: An adaptive algorithm based on neuromimetic architecture, *Signal processing*, 24 (1991), pp. 1–10. 6
- [79] Q. KE AND T. KANADE, Robust ℓ_1 norm factorization in the presence of outliers and missing data by alternative convex programming, in *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 1, IEEE, 2005, pp. 739–746. 25, 35, 82
- [80] N. KESHAVA AND J. F. MUSTARD, Spectral unmixing, *IEEE signal processing magazine*, 19 (2002), pp. 44–57. 3, 22, 138
- [81] D. KONG, C. DING, AND H. HUANG, Robust nonnegative matrix factorization using $\ell_{2,1}$ -norm, in *Proceedings of the 20th ACM international conference on Information and knowledge management*, ACM, 2011, pp. 673–682. 35, 37
- [82] Y. KOREN, R. BELL, AND C. VOLINSKY, Matrix factorization techniques for recommender systems, *Computer*, 42 (2009). 2
- [83] M. KOWALSKI, Sparse regression using mixed norms, *Applied and Computational Harmonic Analysis*, 27 (2009), pp. 303–324. 45, 92, 113, 173
- [84] P. KUPPINGER, G. DURISI, AND H. BOLCSKEI, Uncertainty relations and sparse signal recovery for pairs of general signal sets, *IEEE Transactions on Information Theory*, 58 (2012), pp. 263–277. 78, 79
- [85] J. N. LASKA, M. A. DAVENPORT, AND R. G. BARANIUK, Exact signal recovery from sparsely corrupted measurements through the pursuit of justice, in *Signals, Systems and Computers, 2009 Conference Record of the Forty-Third Asilomar Conference on*, IEEE, 2009, pp. 1556–1560. 78

- [86] S. M. LEACH, J.-F. CARDOSO, C. BACCIGALUPI, R. BARREIRO, M. BETOULE, J. BOBIN, A. BONALDI, J. DELABROUILLE, G. DE ZOTTI, C. DICKINSON, ET AL., Component separation methods for the Planck mission, *Astronomy & Astrophysics*, 491 (2008), pp. 597–615. 3, 20
- [87] D. D. LEE AND H. S. SEUNG, Learning the parts of objects by non-negative matrix factorization, *Nature*, 401 (1999), pp. 788–791. 7, 34
- [88] D. D. LEE AND H. S. SEUNG, Algorithms for non-negative matrix factorization, in *Advances in neural information processing systems*, 2001, pp. 556–562. 34
- [89] T.-W. LEE, M. S. LEWICKI, M. GIROLAMI, AND T. J. SEJNOWSKI, Blind source separation of more sources than mixtures using overcomplete representations, *IEEE signal processing letters*, 6 (1999), pp. 87–90. 7
- [90] C. C. LELONG, P. C. PINET, AND H. POILVÉ, Hyperspectral imaging and stress mapping in agriculture: a case study on wheat in Beauce (France), *Remote sensing of environment*, 66 (1998), pp. 179–191. 22, 141
- [91] Q. LI, H. LI, Z. LU, Q. LU, AND W. LI, Denoising of Hyperspectral Images Employing Two-Phase Matrix Decomposition, *Selected Topics in Applied Earth Observations and Remote Sensing*, *IEEE Journal of*, 7 (2014), pp. 3742–3754, [doi:10.1109/JSTARS.2014.2360409](https://doi.org/10.1109/JSTARS.2014.2360409). 20, 31
- [92] Q. LI, Z. LU, Q. LU, H. LI, AND W. LI, Noise reduction for hyperspectral images based on structural sparse and low-rank matrix decomposition, in *Geoscience and Remote Sensing Symposium (IGARSS), 2013 IEEE International*, IEEE, 2013, pp. 1075–1078. 31
- [93] C.-H. LIN, W.-K. MA, W.-C. LI, C.-Y. CHI, AND A. AMBIKAPATHI, Identifiability of the simplex volume minimization criterion for blind hyperspectral unmixing: The no-pure-pixel case, *IEEE Transactions on Geoscience and Remote Sensing*, 53 (2015), pp. 5530–5546. 140
- [94] C.-J. LIN, Projected gradient methods for nonnegative matrix factorization, *Neural computation*, 19 (2007), pp. 2756–2779. 42
- [95] W. LIU, P. P. POKHAREL, AND J. C. PRÍNCIPE, Correntropy: Properties and applications in non-gaussian signal processing, *IEEE Transactions on Signal Processing*, 55 (2007), pp. 5286–5298. 35
- [96] G. LU AND B. FEI, Medical hyperspectral imaging: a review, *Journal of biomedical optics*, 19 (2014), pp. 010901–010901. 3
- [97] Q. LV AND X.-D. ZHANG, A unified method for blind separation of sparse sources with unknown source number, *IEEE Signal Processing Letters*, 13 (2006), pp. 49–51, [doi:10.1109/LSP.2005.860540](https://doi.org/10.1109/LSP.2005.860540). 104

- [98] W.-K. MA, J. BIOCAS-DIAS, T.-H. CHAN, N. GILLIS, P. GADER, A. PLAZA, A. AMBIKAPATHI, AND C.-Y. CHI, A signal processing perspective on hyperspectral unmixing: Insights from remote sensing, *Signal Processing Magazine, IEEE*, 31 (2014), pp. 67–81, doi:10.1109/MSP.2013.2279731. 3, 139
- [99] J. MAIRAL, F. BACH, J. PONCE, ET AL., Sparse modeling for image and vision processing, *Foundations and Trends® in Computer Graphics and Vision*, 8 (2014), pp. 85–283. 7
- [100] M. MIHOKO AND S. EGUCHI, Robust blind source separation by beta divergence, *Neural computation*, 14 (2002), pp. 1859–1886. 33, 34, 37, 96, 122
- [101] H. MORTADA, V. MAZET, C. SOUSSEN, AND C. COLLET, Separation of delayed parameterized sources, in *EUSIPCO*, Kos Island, Greece, Aug. 2017, <https://hal.archives-ouvertes.fr/hal-01497248>. 171
- [102] Y. MOUDDEN AND J. BOBIN, Hyperspectral BSS using GMCA with spatio-spectral sparsity constraints, *IEEE Transactions on Image processing*, 20 (2011), pp. 872–879. 4, 46, 138, 171
- [103] J. M. NASCIMENTO AND J. M. DIAS, Does independent component analysis play a role in unmixing hyperspectral data?, *IEEE Transactions on Geoscience and Remote Sensing*, 43 (2005), pp. 175–187. 7
- [104] J. M. NASCIMENTO AND J. M. DIAS, Vertex component analysis: A fast algorithm to unmix hyperspectral data, *IEEE transactions on Geoscience and Remote Sensing*, 43 (2005), pp. 898–910. 139, 157
- [105] T.-H. OH, Y.-W. TAI, J.-C. BAZIN, H. KIM, AND I. S. KWEON, Partial sum minimization of singular values in Robust PCA: Algorithm and applications, *IEEE transactions on pattern analysis and machine intelligence*, 38 (2016), pp. 744–758. 31, 112
- [106] P. PAATERO AND U. TAPPER, Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values, *Environmetrics*, 5 (1994), pp. 111–126. 7, 47, 56
- [107] N. PARIKH, S. P. BOYD, ET AL., Proximal algorithms., *Foundations and Trends in optimization*, 1 (2014), pp. 127–239. 40, 41, 42, 43, 55, 114
- [108] Y. QIAN, S. JIA, J. ZHOU, AND A. ROBLES-KELLY, Hyperspectral unmixing via $\ell_{1/2}$ sparsity-constrained nonnegative matrix factorization, *IEEE Transactions on Geoscience and Remote Sensing*, 49 (2011), pp. 4282–4297. 140
- [109] J. QIN, K. CHAO, M. S. KIM, R. LU, AND T. F. BURKS, Hyperspectral and multispectral imaging for evaluating food safety and quality, *Journal of Food Engineering*, 118 (2013), pp. 157–171. 3

- [110] H. RAGUET, J. FADILI, AND G. PEYRÉ, A generalized forward-backward splitting, SIAM Journal on Imaging Sciences, 6 (2013), pp. 1199–1226. 43, 131, 155
- [111] H. RAGUET AND L. LANDRIEU, Preconditioning of a generalized forward-backward splitting and application to optimization on graphs, SIAM Journal on Imaging Sciences, 8 (2015), pp. 2706–2739. 55
- [112] J. RAPIN, J. BOBIN, A. LARUE, AND J.-L. STARCK, Sparse and non-negative BSS for noisy data, IEEE Transactions on Signal Processing, 61 (2013), pp. 5620–5632. 16, 48, 57, 66, 67, 68, 155
- [113] J. RAPIN, J. BOBIN, A. LARUE, AND J.-L. STARCK, NMF with Sparse Regularizations in Transformed Domains, SIAM Journal on Imaging Sciences, 7 (2014), pp. 2020–2047. 7, 46, 67, 68, 74, 84, 104, 118, 131
- [114] J. RAPIN, A. SOULOUMIAC, J. BOBIN, A. LARUE, C. JUNOT, M. OUETHRANI, AND J.-L. STARCK, Application of Non-negative Matrix Factorization to LC/MS data, Signal Processing, (2015), p. 8. 2, 20
- [115] J. R. RAPIN, Sparse decompositions for advanced data analysis of hyperspectral data in biological applications, theses, Université Paris Sud - Paris XI, Dec. 2014, <https://tel.archives-ouvertes.fr/tel-01128517>. 5, 20, 71, 84
- [116] C. REVEL, Y. DEVILLE, V. ACHARD, AND X. BRIOTTET, Inertia-constrained pixel-by-pixel nonnegative matrix factorisation: a hyperspectral unmixing method dealing with intra-class variability, arXiv preprint arXiv:1702.07630, (2017). 163
- [117] S. SARDY, P. TSENG, AND A. BRUCE, Robust wavelet denoising, IEEE Transactions on Signal Processing, 49 (2001), pp. 1146–1152. 36
- [118] B. SHEN, L. SI, R. JI, AND B. LIU, Robust nonnegative matrix factorization via ℓ_1 norm regularization, arXiv preprint arXiv:1204.2311, (2012). 36
- [119] B. SOMERS, G. P. ASNER, L. TITS, AND P. COPPIN, Endmember variability in spectral mixture analysis: A review, Remote Sensing of Environment, 115 (2011), pp. 1603–1616. 22, 141
- [120] J.-L. STARCK, E. J. CANDÈS, AND D. L. DONOHO, The curvelet transform for image denoising, IEEE Transactions on image processing, 11 (2002), pp. 670–684. 9
- [121] J.-L. STARCK, M. ELAD, AND D. DONOHO, Redundant multiscale transforms and their application for morphological component separation, Advances in Imaging and Electron Physics, 132 (2004), pp. 287–348. 10

- [122] J.-L. STARCK, J. FADILI, AND F. MURTAGH, The undecimated wavelet decomposition and its reconstruction, *IEEE Transactions on Image Processing*, 16 (2007), pp. 297–309. 9, 102, 124
- [123] J.-L. STARCK, Y. MOUDDEN, J. BOBIN, M. ELAD, AND D. DONOHO, Morphological component analysis, in *Optics & Photonics 2005*, International Society for Optics and Photonics, 2005, pp. 59140Q–59140Q. 10
- [124] J.-L. STARCK, F. MURTAGH, AND J. M. FADILI, Sparse image and signal processing: wavelets, curvelets, morphological diversity, Cambridge University Press, 2010. 7, 10, 12, 66, 67, 76, 92, 173, 175
- [125] C. STUDER, Recovery of Signals with Low Density, arXiv preprint arXiv:1507.02821, (2015). 93
- [126] F. SUREAU, J.-L. STARCK, J. BOBIN, P. PAYKARI, AND A. RASSAT, Sparse point-source removal for full-sky CMB experiments: application to WMAP 9-year data, *Astronomy & Astrophysics*, 566 (2014), p. A100. 20, 31
- [127] P.-A. THOUVENIN, N. DOBIGEON, AND J.-Y. TOURNERET, Hyperspectral unmixing with spectral variability using a perturbed linear mixing model, *IEEE Transactions on Signal Processing*, 64 (2016), pp. 525–538. 142, 143, 144, 147, 148, 150, 151, 158, 161, 163, 164, 165, 167, 168
- [128] P. TSENG, Convergence of a block coordinate descent method for nondifferentiable minimization, *Journal of optimization theory and applications*, 109 (2001), pp. 475–494. 49, 50, 56, 95
- [129] P. G. VAN DOKKUM, Cosmic-ray rejection by laplacian edge detection, *Publications of the Astronomical Society of the Pacific*, 113 (2001), p. 1420. 20
- [130] E. VINCENT, N. BERTIN, R. GRIBONVAL, AND F. BIMBOT, From blind to guided audio source separation: How models and side information can improve the separation of sound, *IEEE Signal Processing Magazine*, 31 (2014), pp. 107–115, doi:10.1109/MSP.2013.2297440. 2
- [131] E. VINCENT, R. GRIBONVAL, AND C. FÉVOTTE, Performance measurement in blind audio source separation, *Audio, Speech, and Language Processing*, *IEEE Transactions on*, 14 (2006), pp. 1462–1469. 58, 180
- [132] M. E. WINTER, N-findr: An algorithm for fast autonomous spectral end-member determination in hyperspectral data, in *SPIE’s International Symposium on Optical Science, Engineering, and Instrumentation*, International Society for Optics and Photonics, 1999, pp. 266–275. 139

- [133] J. WRIGHT AND Y. MA, Dense error correction via ℓ_1 -minimization, IEEE Transactions on Information Theory, 56 (2010), pp. 3540–3560. 79
- [134] J. WRIGHT, A. Y. YANG, A. GANESH, S. S. SASTRY, AND Y. MA, Robust face recognition via sparse representation, IEEE transactions on pattern analysis and machine intelligence, 31 (2009), pp. 210–227. 93
- [135] H. XU, C. CARAMANIS, AND S. SANGHAVI, Robust PCA via Outlier Pursuit, in Advances in Neural Information Processing Systems, 2010, pp. 2496–2504. 31, 32, 33, 37, 80, 112, 121
- [136] Y. XU AND W. YIN, A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion, SIAM Journal on imaging sciences, 6 (2013), pp. 1758–1789. 38, 42, 49, 51, 52, 54, 55, 56
- [137] Y.-L. YU, On decomposing the proximal map, in Advances in Neural Information Processing Systems, 2013, pp. 91–99. 173
- [138] A. ZARE AND K. HO, Endmember variability in hyperspectral analysis: Addressing spectral variability during spectral unmixing, IEEE Signal Processing Magazine, 31 (2014), pp. 95–104. 22, 141, 142
- [139] H. ZHANG, W. HE, L. ZHANG, H. SHEN, AND Q. YUAN, Hyperspectral Image Restoration Using Low-Rank Matrix Recovery, Geoscience and Remote Sensing, IEEE Transactions on, 52 (2014), pp. 4729–4743, doi:10.1109/TGRS.2013.2284280. 31
- [140] L. ZHANG, Z. CHEN, M. ZHENG, AND X. HE, Robust non-negative matrix factorization, Frontiers of Electrical and Electronic Engineering in China, 6 (2011), pp. 192–200. 36
- [141] G. ZHOU, Z. YANG, S. XIE, AND J. M. YANG, Mixing matrix estimation from sparse mixtures with unknown number of sources, IEEE Transactions on Neural Networks, 22 (2011), pp. 211–221, doi:10.1109/TNN.2010.2091427. 104
- [142] T. ZHOU AND D. TAO, Godec: Randomized low-rank & sparse matrix decomposition in noisy case, in Proceedings of the 28th International Conference on Machine Learning (ICML-11), 2011, pp. 33–40. 31, 112
- [143] Z. ZHOU, X. LI, J. WRIGHT, E. CANDÉS, AND Y. MA, Stable Principal Component Pursuit, in Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on, IEEE, 2010, pp. 1518–1522. 31
- [144] M. ZIBULEVSKY AND B. PEARLMUTTER, Blind source separation by sparse decomposition in a signal dictionary, Neural Computation, 13 (2001), pp. 863–882, doi:10.1162/089976601300014385. 7

Titre : Parcimonie, diversité morphologique et séparation robuste de sources.

Mots clés : SSA, parcimonie, robustesse, données aberrantes, factorisation de matrices, diversité morphologique.

Résumé : Cette thèse porte sur le problème de Séparation Aveugle de Sources (SAS) en présence de données aberrantes. La plupart des méthodes de SAS sont faussées par la présence de déviations structurées par rapport au modèle de mélange linéaire classique: des événements physiques inattendus ou des dysfonctionnements de capteurs en sont des exemples fréquents. Nous proposons un nouveau modèle prenant en compte explicitement les données aberrantes. Le problème de séparation en résultant, mal posé, est adressé grâce à la parci-

monie. L'utilisation de cette dernière est particulièrement intéressante en SAS robuste car elle permet simultanément de démêler les sources et de séparer les différentes contributions. Ces travaux sont étendus pour l'estimation de variabilité spectrale pour l'imagerie hyperspectrale terrestre. Des comparaisons avec des méthodes de l'état-de-l'art montrent la robustesse et la fiabilité des algorithmes associés pour un large éventail de configurations, incluant le cas déterminé.

Title : Sparse modeling, morphological diversity and robust source separation.

Keywords : BSS, sparse modeling, robustness, outliers, matrix factorization, morphological diversity.

Abstract : This manuscript addresses the Blind Source Separation (BSS) problem in the presence of outliers. Most BSS techniques are hampered by the presence of structured deviations from the standard linear mixing model, such as unexpected physical events or malfunctions of sensors. We propose a new data model taking explicitly into account the deviations. The resulting joint estimation of the components is an ill-posed problem, tackled using sparse modeling. The latter is particularly efficient

for solving robust BSS since it allows for a robust unmixing of the sources jointly with a precise separation of the components. These works are then extended for the estimation of spectral variability in the framework of terrestrial hyperspectral imaging. Numerical experiments highlight the robustness and reliability of the proposed algorithms in a wide range of settings, including the full-rank regime.