Unsupervised separation of sparse multivalued components with applications in astrophysics

Imane El Hamzaoui

To cite this version:

Imane El Hamzaoui. Unsupervised separation of sparse multivalued components with applications in astrophysics. Signal and Image Processing. Université Paris-Saclay, 2020. English. NNT : 2020UPASG015 . tel-03150357

HAL Id: tel-03150357
https://theses.hal.science/tel-03150357

Submitted on 23 Feb 2021

Thèse de doctorat

NNT : 2020UPASG015

Unsupervised separation of sparse multivalued components with applications in astrophysics

Thèse de doctorat de l'Université Paris-Saclay

École doctorale n° 580, Sciences et technologies de l'information et de la communication (STIC)
Spécialité de doctorat: Traitement du signal et des images
Unité de recherche: Université Paris-Saclay, CEA, Département d'Electronique des Détecteurs et d'Informatique pour la Physique, 91191, Gif-sur-Yvette, France
Référent: Faculté des sciences d'Orsay

Thèse présentée et soutenue en visioconférence totale, le 9 décembre 2020, par

Imane El hamzaoui

Composition du jury:

Cédric Richard Professeur, Laboratoire Lagrange Université Côte d'Azur	Président
Nicolas Dobigeon Professeur, INP-ENSEEIHT Toulouse	Rapporteur & Examinateur
Said Moussaoui Professeur, École Centrale de Nantes	Rapporteur & Examinateur
Caroline Chaux Chargée de recherche CNRS, Institut de mathématiques de Marseille	Examinatrice
Pascal Larzabal Professeur, École Normale Supérieure	Examinateur
Jérôme Bobin Chercheur, CEA Paris-Saclay	Directeur

Remerciements

Je remercie tout d'abord M. Cédric Richard, M. Said Moussaoui, M. Nicolas Dobigeon, Mme Caroline Chaux et M. Pascal Larzabal de m'avoir fait l'honneur de faire partie de mon jury de thèse. Je remercie également mon directeur de thèse M. Jérôme Bobin pour son aide précieuse, sa patience, sa disponibilité et nos échanges scientifiques qui m'ont été très instructifs. Je remercie également les membres du DEDIP de m'avoir accueillie au sein de leur unité de recherche ainsi que les membres du Cosmostat.

Je suis reconnaissante envers mon entourage proche qui m'a accompagnée et soutenue durant ces trois dernières années.

Enfin, c'est tout naturellement que je dédie cette thèse à mon grand frère, ma mère et mon père que j'aime du plus profond de mon être.

Résumé

Les dernières décennies ont connu un développement sans précédent des techniques d'imagerie multivaluée. Celles-ci permettent d'acquérir les images d'une scène selon différents points de vue. Un exemple trivial de ce type de données se trouve dans les photographies en couleur qui sont obtenues en enregistrant simultanément une même scène selon trois bandes de fréquence à savoir le rouge, le vert et le bleu. De façon générale, les données multivaluées sont modélisées mathématiquement par une matrice $\mathbf{X} \in \mathbb{R}^{m \times t}$ où m désigne le nombre d'observations (trois dans le cas des photographies en couleur) et t le nombre d'échantillons (soit le nombre de pixels en ce qui concerne les images).

La prolifération de données multivaluées s'est accompagnée d'avancées dans le domaine de l'analyse et du traitement des données. Dans ce cadre, la Séparation de Sources Aveugles (que l'on désignera par SSA et dont l'acronyme anglais est BSS pour Blind Source Separation), introduite dans Comon and Jutten [2010] et largement utilisée depuis, permet d'apprendre des décompositions physiquement interprétable des données multivaluées.

En SSA, on considère que les données multivaluées sont constituées de la somme des contributions élémentaires de n composantes dites *sources*. L'objectif est alors de retrouver les composantes physiques de chacune de ces sources. Le formalisme mathématique majoritairement adopté pour modéliser ce problème est le Modèle de Mélange Linéaire (ou LMM en anglais pour Linear Mixture Model): celui-ci suppose que les observations proviennent d'un mélange linéaire pondéré inconnu des sources.

La SSA consiste donc à retrouver les sources ainsi que les poids du mélange à partir de la matrice de données. Cependant, ce problème est mal posé et nécessite l'ajout d'informations à priori pour restreindre le champ des solutions possibles, par exemple supposer l'indépendance statistique des sources ou encore la positivité des coefficients des deux matrices. Nos travaux s'inscrivent dans le cadre d'une toute autre classe de pénalisation à savoir le principe de parcimonie. Cette dernière désigne que la représentation des sources (potentiellement dans un domaine transformé) est composée en grande majorité de coefficients nuls ou du moins de faible amplitude.

De nombreuses méthodes de séparation de sources avec contrainte de parcimonie ont été proposées au cours des dernières décennies (telles que l'algorithme GMCA Bobin et al. [2007]) et se sont montrées très efficaces dans de nombreux domaines d'applications (dans l'étude de données astrophysique en micro-ondes Collaboration et al. [2014] par exemple ou encore l'analyse de données en spectrométrie pour la santé Rapin [2014]). Cependant, ces méthodes sont inadaptées dans le cas des données astrophysiques en hautes énergies. En effet, ces dernières présentent des particularités généralement omises dans les algorithmes de sBSS classiques, à savoir:

- Les données astrophysiques en hautes énergies proviennent d'un processus de comptage de photons. Ainsi, le bruit contaminant les données est de nature poissonnienne et doit être considéré comme tel pour permettre une modélisation fidèle des données. Or, il est largement admis dans la plupart des algorithmes de BSS que le bruit suit une distribution Gaussienne. Une nouvelle méthode de séparation de sources avec prise en compte du bruit de Poisson est nécessaire et utile pour de nombreuses autres applications encore.[1]

- En astrophysique, les colonnes de la matrice de mélange représentent les spectres des sources présentes dans l'image. Du fait de phénomènes physiques (l'effet Doppler notamment), il arrive que la distribution spectrale d'un même élément soit modifiée en fonction du pixel considéré dans l'image. Ce phénomène, communément appelé variabilités spectrales n'est que très peu pris en compte dans la littérature et nullement dans les méthodes de séparation de sources parcimonieuses dédiées aux applications astrophysiques.

Séparation de Sources Aveugle de données contaminées par du bruit de Poisson

Nous proposons une nouvelle approche de séparation de sources parcimonieuses (dans un domaine transformé) adaptée aux observations contaminées par du bruit de Poisson. La substitution du modèle de mélange

---

[1] On retrouve ainsi ce problème dans le cas des données de spectroscopie en rayons γ par exemple (Xu et al. [2019]).

linéaire par le modèle poissonien présente deux écueils principaux. Tout d'abord, le terme d'attache aux données n'est pas lisse en 0. Ensuite, il devient très mal conditionné dans un régime de faibles statistiques. Pour pallier à ces deux principaux problèmes, nous proposons d'utiliser une approximation lisse de ce terme via la technique de Nesterov (Nesterov [2005]) et de minimiser la fonction de coût résultante par le Block Coordinate algorithm (une méthode d'optimisation multi-convexe Tseng [2001] que nous détaillerons dans le chapitre 2). Enfin, nous présenterons l'algorithme développé au cours de la thèse, le pGMCA (poisson-Generalized Morphological Component Analysis), étudierons ses performances vis-à-vis des algorithmes de l'État-de-l'Art et conclurons par les résultats de son application à des données réalistes Chandra.

Séparation de Sources Aveugle en présence de variabilités spectrales

Dans le chapitre 5, nous présenterons les travaux menés au cours de la thèse sur la Séparation de Sources Aveugle en présence de variabilités spectrales. Nous choisirons de modéliser ce phénomène par une matrice de mélange spatialement dépendante. Ainsi, au lieu d'avoir pour tous les échantillons une unique matrice de mélange $\mathbf{A}$, nous proposons de considérons un champ de matrices de mélange $\{\mathbf{A}[k]\}_{k=1..t}$. Le problème de SSA étant alors très mal posé en l'état, nous nous sommes inspirés de la linéarisation des matrices de mélange proposée par Thouvenin et al. [2015] pour produire un modèle adapté à nos données. Celui-ci consiste à décrire la matrice de mélange locale en un échantillon k comme la somme d'une matrice de mélange de référence (spatialement indépendante) et d'une faible perturbation qui fait office de variabilité spectrales. Étant donné que les variabilités étudiées présentent une variation lisse, nous imposons au terme de faible perturbation une contrainte de régularité à travers la norme $\ell_{2,1}$ dans un domaine adapté à la structure des variablités. Ainsi, la diversité morphologique entre les sources et les variabilités nous permet de séparer leurs contributions. Nous introduirons l'algorithme conçu au cours de la thèse et nous en étudierons le comportement, les performances et les limitations. Nous présenterons également les résultats de son application aux données Chandra.

Notations, definitions and acronyms

General notations

- x: a scalar;
- $\mathbf{x}$: a vector;
- $\mathbf{X}$: a matrix;
- $\mathbf{X}_i$: the $i-th$ line of the matrix $\mathbf{X}$;
- $\mathbf{X}^j$: the $j-th$ column of the matrix $\mathbf{X}$;
- $\mathbf{X}_i^j$: the $(i,j)-th$ entry of the matrix $\mathbf{X}$;
- $\mathbf{X}_{i=1..n}$: the lines of the matrix $\mathbf{X}$ indexed by the subset $I=[1,n]$;
- $\{\mathbf{X}^j\}_{j=1..t}$: : the columns of the matrix $\mathbf{X}$ indexed by the subset $J=[1,t]$;
- $\mathbf{X}^\dagger$: Moore-Penrose pseudo-inverse of $\mathbf{X}$;
- $\mathbf{X}^T$: transpose of $\mathbf{X}$;
- $\mathbf{X}^{(k)}$: estimation of the matrix $\mathbf{X}$ at the $k-th$ iteration in iterative algorithms;
- $\mathrm{Diag}(x_1,...,x_n) \in \mathbb{R}^{n\times n}$: Diagonal matrix whose diagonal coefficients are $x_1,...,x_n$. If we define the vector $\mathbf{x} \in \mathbb{R}^n$ such that $\mathbf{x}=[x_1,...,x_n]$, then $\mathrm{Diag}(\mathbf{x}) = \mathrm{Diag}(x_1,...,x_n)$;

- $OB(m)$: the oblique set defined by all matrices $\mathbf{X}$ whose columns have unit ℓ_2 norms, $OB(m) = \{\mathbf{X} : \forall j, \|\mathbf{X}^j\|_2 = 1\}$;
- K^+: the non-negative orthant *i.e.* the set composed of matrices with non-negative entries, $K^+ = \left\{\mathbf{X} : \forall i,j, \mathbf{X}_i^j \geq 0\right\}$. For the sake of simplicity, $\mathbf{X} \in K^+$ will be equivalently written as $\mathbf{X} \geq 0$;
- $\mathcal{S}^{m-1}$: $(m-1)$-dimensional hypersphere;
- D_β: β-divergence defined for some matrices $\mathbf{U}$ and $\mathbf{V}$ as

$$D_\beta\left(\mathbf{U},\mathbf{V}\right) = \sum_{i,t} \frac{1}{\beta(\beta-1)} \left(\mathbf{U}_i[t] + (\beta-1)\mathbf{V}_i[t]^\beta + \beta \mathbf{U}_i[t]\mathbf{V}_i[t]^{\beta-1}\right).$$

Operators

- $\odot$: element-wise or Hadamard product;
- $\oslash$: element-wise division;
- $\iota_\mathcal{C}$: indicator function of the set $\mathcal{C}$. For every $x \in \mathcal{C}$, $\iota_\mathcal{C}(x) = 0$ if $x \in \mathcal{C}$ and $+\infty$ otherwise;
- ∇f: gradient of f;
- $\mathcal{S}_\Lambda(\mathbf{X})$: soft-thresholding operator defined for each $(i,j)-th$ coefficient by $\mathcal{S}_\Lambda(\mathbf{X})_{i,j} = \max(0, \mid \mathbf{X}_{i,j} \mid -\Lambda_{i,j}) \times \mathrm{sign}(\mathbf{X}_{i,j})$;

- MAD: Median Absolute Deviation;
- prox_f: proximal operator of f;
- $\mathcal{P}$: probability;
- $\Pi_{\cdot \geq 0}$: projection onto the positive orthant;
- $\Pi_{\|.\|_2=1}(\mathbf{X})$: projection of the columns of $\mathbf{X}$ onto the ℓ_2 unit hypersphere.

- $\|\mathbf{X}\|_F$: the Frobenius norm of the matrix $\mathbf{X}$ is defined as $\|\mathbf{X}\|_F = \sqrt{\text{Trace}\,(\mathbf{X}\mathbf{X}^T)}$;
- $\|\mathbf{X}\|_p$, $p \in \mathbb{R}_+$: ℓ_p (quasi-) norm of a matrix $\mathbf{X} \in \mathbb{R}^{m\times t}$ seen as a vector, $\|\mathbf{X}\|_p = (\sum_{i=1}^m \sum_{j=1}^t \mid \mathbf{X}_i^j \mid^p)^{\frac{1}{p}}$;
- $\|\mathbf{X}\|_{s,2}$: spectral norm of $\mathbf{X}$ *i.e.* largest eigenvalue of $\mathbf{X}$;
- $\|\mathbf{X}\|_{2,1}$: $\ell_{2,1}$ norm of $\mathbf{X} \in \mathbb{R}^{m\times t}$ *i.e.* $\|\mathbf{X}\|_{2,1} = \sum_{k=1}^t \|\mathbf{X}\|_2$;
- $\|\mathbf{X}\|_\infty$: maximum norm of $\mathbf{X}$ *i.e.* maximum absolute value of the coefficients of $\mathbf{X}$.

Abbreviations

- BSS: Blind Source Separation;
- sBSS: sparse Blind Source Separation;
- SV: spectral variabilities;
- ICA: Independent Component Analysis;
- NMF: Non-negative Matrix Factorization;
- DCT: Discrete Cosine Transform;
- MDP: Morphological Diversity Principle;
- LMM: Linear Mixture Model;
- PET: Positron Emission Tomography;
- TAC: Time-Activity-Curve;
- ASC: abundance sum-to-one constraint;
- ANC: abundance non-negative constraint;
- PP: pure pixel;
- MAP: Maximum A Posteriori;
- MCMC: Markov Chain Monte-Carlo;
- CMB: Cosmic Microwave Background;
- ESA: the European Space Agency;

- LC-MS: the Liquid Chromatography - Mass Spectrometry;
- ELMM: the Extended Linear Mixture Model;
- PLMM: the Perturbed Linear Mixture Model;
- ALMM: the Augmented Linear Mixture Model;

- GMCA: Generalized Morphological Component Analysis;
- svGMCA: spectral variatibilities Generalized Morphological Component Analysis;
- pGMCA: Poisson Generalized Morphological Component Analysis;
- L-GMCA: Local Generalized Morphological Component Analysis;
- FB: Forward Backward Splitting algorithm;
- GFB: Generalized Forward Backward algorithm;
- BCD: Block Coordinate Descent algorithm;
- PBC: Proximal Block Coordinate algorithm;
- PALM: Proximal Alternating Linearized Minimization algorithm;
- pALS: Projected Alternating Least square algorithm;
- CSR: Constrained Sparse Regression;
- FCLS: Fully Constrained Least Square;
- MESMA: Multiple Endmember Spectral Mixture Analysis;
- ADMM: Alternating Direction Method Descent algorithm;

- GMSE: General Mean Square Error;
- MSE: Mean Square Error;
- SAD: Spectral Angular Deviation;
- RE: Reconstruction Error;
- SIR, SAR, SDR, SNR: Signal to interference-artefact-distorsion-noise ratio;

Specific naming

- n: number of sources;
- m: number of observations;
- t: number of samples;
- $\mathbf{X} \in \mathbb{R}^{m \times t}$: the observation matrix;
- $\mathbf{A} \in \mathbb{R}^{m \times n}$: the mixing matrix;
- $\mathbf{S} \in \mathbb{R}^{n \times t}$: the sources matrix;
- $\mathbf{N} \in \mathbb{R}^{m \times t}$: the noise matrix;
- $\mathbf{E} \in \mathbb{R}^{m \times n}$: the endmember matrix;
- $\hat{\mathbf{A}} \in \mathbb{R}^{n \times t}$: the abundance matrix;
- $\mathbf{Y} \in \mathbb{R}^{m \times t}$: the observation matrix of remotely sensed data.
- $\{\mathbf{A}[k] \in \mathbb{R}^{m \times n}\}_{\{k=1..t\}}$: sample-dependent mixing matrices;
- $\mathbf{\Delta}^i \in \mathbb{R}^{m \times t}$: spectral variabilities matrix affecting the $i-th$ source;
- $\theta^i \in \mathbb{R}^t$: Angular variability affecting the $i-th$ source;

Contents

Chapter 1

Blind Source Separation and sparse modeling

The unprecedented development of multivalued data sensors over the last years calls for performing analysis tools capable of retrieving meaningful information for these data. In this context, this thesis is aimed at introducing novel methods of unsupervised analysis of multivalued high-energy astrophysical data.

In this chapter, we propose examples of multivalued data, before presenting sparse Blind Source Separation (sBSS), an efficient analysis tool that is at the core of this thesis. Then, we introduce the problems encountered in applying sBSS to high-energy astrophysics and finally we present the outline of the rest of this thesis.

1.1 Multivalued data analysis with BSS

Multivalued data designate observations where a scene is recorded simultaneously with different modalities. A popular example is color digital images a the one depicted in figure 1.1: the photographed scene is acquired with three bands: the red band, the green band and the blue band (figure 1.2). multivalued data are modeled using a matrix $\mathbf{X} \in \mathbb{R}^{m \times t}$ where m designates the number of observations and t the number of samples of the acquired scene.

A key element of multivalued data is the inter-observation structure: in our example, each seed of spice has a fixed spectrum, the values taken by the samples in the seed location over the acquisition range are interdependent. Therefore, multivalued data have a double structure: a spatial structure (corresponding to the coherence of the $i-th$ line of $\mathbf{X}$ taken as an individual signal) and a spectral one.

Blind Source Separation focuses on learning meaningful decomposition of multivalued data by exploiting their structural features.

Figure 1.1: Example of color digital image.

In this chapter, we start by giving examples of application of BSS in different scientific domains. We introduce afterwards three classes of data modelling among which BSS with sparsity that is the core of our thesis.

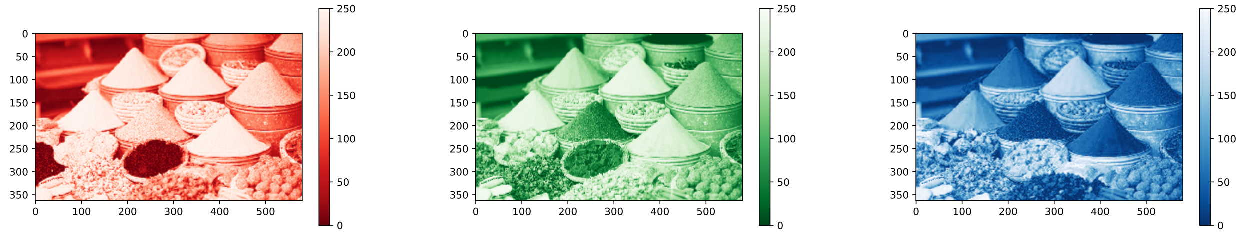

Red band X^1 Green band X^2 Blue band X^3

Figure 1.2: Decomposition of the color image in its three bands: red, green and blue.

1.1.1 Examples of applications

Extended sources in X-ray astronomy

The observation of the sky in X-ray offers valuable information about the Milky Way, Galactic Sources such as supernovae remnants and extra-galactic sources such as clusters of galaxies. Chandra telescope [1] is an example of spectro-imager observing the X-ray sky in the range $1-8$keV. It provides multivalued images of extended source whose spatial and spectral properties are of great interest. More specifically, observable sources are composed of entangled physical components that we aim at retrieving. Further details on Chandra telescope data are provided in section 1.3.

We display in figure 1.3 observations of the supernova remnant Cassiopeia A acquired at different frequencies by the ACIS-S instrument in 2004. In this example, data consist of the mixture of the synchrotron continuum emission and the iron emission at two redshift. The spatial distribution of these three elements is given in figure 1.4. Each one of them participates to the mixture with weights depending on their spectrum over the acquisition range. The goal of BSS is to retrieve both spatial and spectral distributions of the physical components.

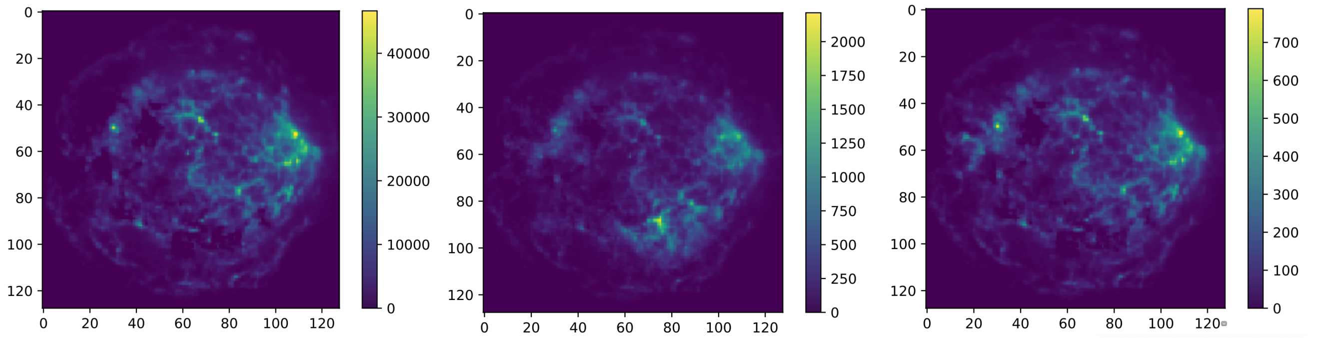

Figure 1.3: Examples of observations of the supernova Cassiopeia A with Chandra satellite.

Estimation of the CMB and foreground emissions from ESA-Planck data

The Cosmological Microwave Background (CMB) is the oldest observable of the Universe. It carries fundamental information for the cosmologists since most cosmological parameters can be derived from it. Many experiments have been dedicated to its observation during the last years, such as WMAP satellite Bennett et al. [2013] or the Planck Mission launched by ESA in 2009 Collaboration et al. [2014]. For example, the Planck mission consisted of observing the full sky at 9 frequencies on the range of the microwave from 30 to 857 GHz. However, CMB cannot be observed on its own since it is mixed with galactic foregrounds (*e.g.* synchrotron)

---

[1] https://chandra.harvard.edu

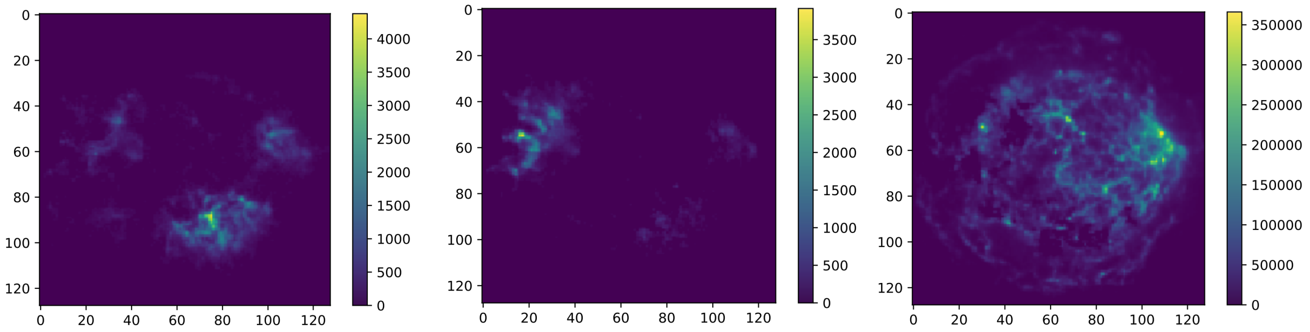

Figure 1.4: Spatial distribution of blue-shifted Fe (left), synchrotron (center) and red-shifted Fe (right).

and point sources. Retrieving CMB map requires disentangling in Planck data the CMB map itself and the contributions of the galactic foregrounds observed. Finally, it boils down to a Blind Source Separation problem. The figure 1.5 shows Planck data represented as a 2D-1D data cube.

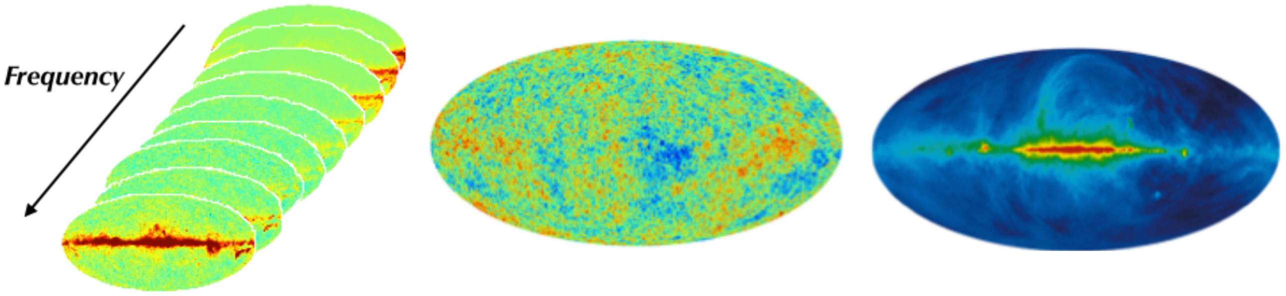

Figure 1.5: from left to right: 2D-1D data cube provided by Planck mission with two spatial dimensions and one spectral dimension; CMB map; Synchrotron map (figures from Chenot [2017]).

LC-MS data analysis

Let us consider a fluid of which we want to know the composition and the concentration of its constituents. The Liquid Chromatography - Mass Spectrometry (LC-MS) Rapin [2014] technique consists of driving the fluid through a chromatography column. Depending on their properties, the physical components of the fluid leave the column with a certain speed. At given time samples, the output of the liquid chromatography is analyzed with a mass spectrometer, providing its mass to charge ratio as featured in figure 1.6.

Finally, we obtain multivalued data where the observations consist of the mass to charge ratio. This example is interesting since it emphasizes the variety of the class of multivalued data that does not reduce to multispectral data. On the opposite of the previous examples, LC-MS data do not consist of a scene that is observed with different frequencies, however the principle remains the same. The problem of the recovery of the spectra and the concentrations of the constituents of the liquid is a BSS problem.

Spectral unmixing from hyperspectral terrestrial images

During last decades, high-quality multi-sensors have flourished and proven to be interesting tools of observations in various application fields. In addition to the applications mentioned above, one can think of crop monitoring for agriculture (Lelong et al. [1998]) or remote sensing (Dobigeon et al. [2009b]) accessed through hyperspectral imaging of the surface of the earth.

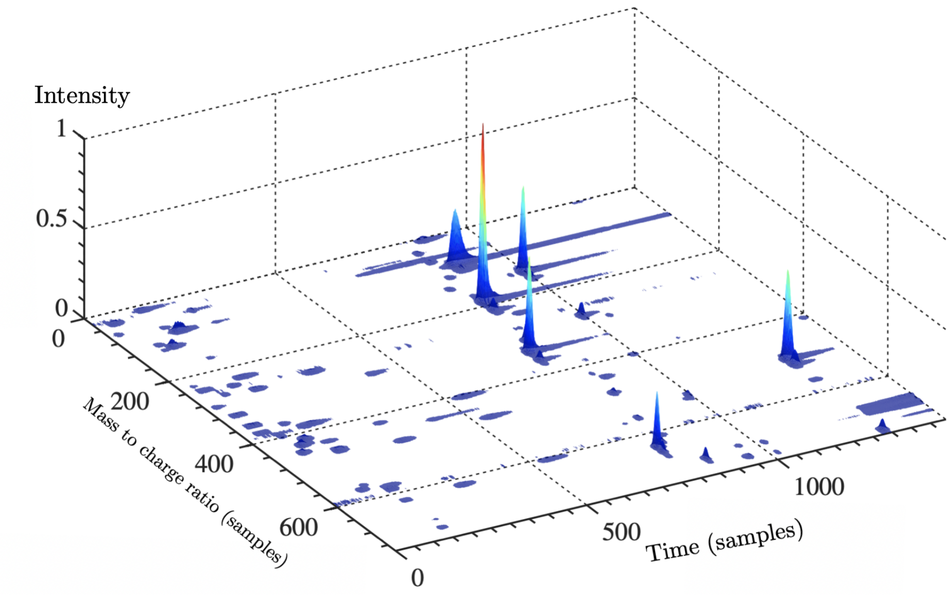

Figure 1.6: Example of observations provided by LC-MS. Figure from Rapin [2014]

1.1.2 Mixing Model

The simplest and most popular mathematical modelling of multivalued data is to consider that the observations correspond to the weighted linear combination of the n elements called sources and denoted $\mathbf{S}_{i=1..n}$:

$$\mathbf{X} = \sum_{i=1}^{n} \mathbf{A}^i \mathbf{S}_i. \tag{1.1}$$

For example, in Chandra data as described in figure 1.7, the sources correspond to the spatial distributions of the elements (synchrotron, Fe1 and Fe2), and the weights $\{\mathbf{A}^i\}_{i=1..n}$ are the spectra of these elements.

This model, known as the **Linear Mixture Model** (LMM) can be recast in the following matrix form:

$$\mathbf{X} = \mathbf{AS}, \tag{1.2}$$

i.e. the multivalued data matrix $\mathbf{X} \in \mathbb{R}^{m\times t}$ called **observation matrix** is the product of **the mixing matrix** $\mathbf{A} \in \mathbb{R}^{m\times n}$ and the **source matrix** $\mathbf{S} \in \mathbb{R}^{n\times t}$. Each column i of the mixing matrix $\mathbf{A}^i \in \mathbb{R}^{m\times 1}$ contains the linear mixing coefficients (spectral distributions of the elements in 1.7) corresponding to the $i-th$ source $\mathbf{S}_i \in \mathbb{R}^{1\times t}$ (spatial distributions in 1.7).

In order to take into account the deviations of the data to the LMM and the noise due to measurements, we introduce a noise term so that the equation (1.2) is rewritten as

$$\mathbf{X} = \mathbf{AS} + \mathbf{N}, \tag{1.3}$$

where $\mathbf{N} \in \mathbb{R}^{m\times t}$.

As mentioned above, multivalued data have a double structure: the structure of each observation considered as an individual signal, and the structure inter-observation or inter-canal. While modelling the data as the product of a mixing matrix and a source matrix, the inter-observation coherence lies fully in the mixing matrix whereas the spatial structure corresponds to the structure of each source (*i.e.* each line of $\mathbf{S}$). Such a model induces that the sources $\mathbf{S}_{i=1..n}$ do not share one with the other any structure.

The aim of Source Separation is to extract the source matrix $\mathbf{S}$ and the mixing matrix $\mathbf{A}$ (in the blind case) from the unique knowledge of the observation matrix $\mathbf{X}$ and the number of sources present in the mixture. It

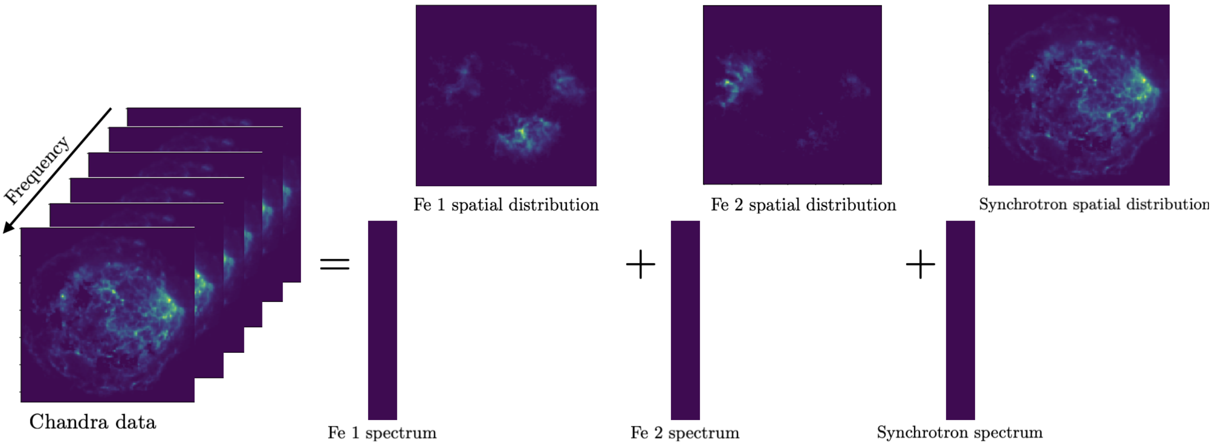

Figure 1.7: Illustration of LMM with Chandra data.

is a matrix factorization problem. In the rest of this thesis, we will consider, as it is usually the case in BSS problem, that we are in the over-determined case meaning that the number of sources to recover is lower than the number of observations, $n \leq m$.

Recovering the couple $(\mathbf{A}, \mathbf{S})$ from $\mathbf{X}$ is an ill-posed problem since there is an infinity of solutions. For example, the couple $(\mathbf{MA}, \mathbf{M}^{-1}\mathbf{S})$ is also solution for any invertible matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$. However, the goal of BSS is not finding any couple that will fulfill the equation (1.2) but recovering the mixing matrix and the source matrix that have physical meaning, for instance the couple that will represent spectra and spatial distribution for Chandra data. To incorporate this prior knowledge, regularization terms are added enabling the discrimination between the infinite possibility of solutions.
We can gather these priors in three families that yielded three classes of BSS: the Independent Component Analysis (ICA, Comon and Jutten [2010]), the Non-Negative Matrix Factorization (NMF, Gillis and Glineur [2012]) and the Sparse Blind Source Separation (sBSS, Zibulevsky [2003]).

1.1.3 Blind Source Separation: ICA, NMF and sparse BSS

Independent Component Analysis

Independent Component Analysis is a statistical method where the sources are supposed to be generated from random variables (random signals).

As mentioned before, the sources should not share similar structures one with another i.e. they are incoherent. Nevertheless, this incoherence principle is merely a notion that requires to be "translated" in terms of statistical criterion.

In Comon [1994] where the principle of ICA is introduced, the authors show that an efficient way of capturing the incoherence between the sources in the statistical framework is the measure of their **statistical independence**. More specifically, it has been proven (Comon [1994] and Darmois [1953]) that if the following assumptions are met:

1. The number of observations is equal to the number of sources (yielding an invertible mixing matrix with $\det(\mathbf{A}) > 0$ if $\mathbf{A}$ is a square matrix);

2. At most one source follows a Gaussian distribution;

3. The sources are mutually independent;

Then there exists a unique unmixing matrix $\mathbf{B}$ such that $\mathbf{S} = \mathbf{BX}$ up to a scaling and permutation indeterminacy.

In spite of the strong and simple theoretical background, the application in practice of ICA is not straightforward since it requires the measure of the so-called statistical independence. Several implementations have been proposed over the years: measures on the mutual difference through Kullback-Leibler divergence (Lee et al. [2000]) or the entropy of the sources (Bell and Sejnowski [1995]) to only mention two.

However, measuring the statistical independence between the sources systematically requires the knowledge of the probability density function of the sources. While approximations may very likely hamper the unmixing, estimation of the probability functions without prior information drastically complexifies the problem.

In addition to that, an other limitation of ICA is the lack of robustness to noise. ICA framework, built without accounting for possible perturbation to the model, is not well-suited to the presence of noise. The author of Davies [2004] has proven that the model becomes degenerate in the presence of noise.

Most importantly, the very basis of ICA is the statistical independence between the sources and, as we will see for remotely sensed data for example, this is far from being true.

The reader is invited to refer to Comon and Jutten [2010] for a complete overview of these methods.

Non-Negative Matrix Factorization

NMF is based upon a very common assumption which is the positivity of the sources and the mixing matrix (Lee and Seung [2001], Paatero and Tapper [1994]). This constraint is met in numerous applications:

- In Chandra data, the sources and the mixing matrix represent the spatial and spectral distributions which are by definition non-negative;
- In remotely sensed data, the coefficients of the source or abundance matrix are by definition comprised between 0 and 1, whereas the columns of $\mathbf{A}$ model physical components spectra;
- In LC-MS, the mixing matrix represents the elution profiles of the fluid components which has non-negative values.

The wide applicability of NMF comes at a cost: the non-negativity constraint can be insufficient to perform the unmixing. For example, it has been shown in Donoho and Stodden [2004] that the identifiability of the model is ensured in the presence of at least one pure pixel per source (*i.e.* pixel where only one source is active).

Given the ill-posedness nature of NMF problem, the authors of Kim and Park [2008], Eggert and Korner [2004] and Hoyer [2002] have proposed to add a sparsity-promoting penalization to the sources. This regularization, physically motivated in a very large number of applications, has been successively exploited in various works such as Cichocki and Zdunek [2007], Cherni et al. [2019] or Rapin [2014] to only cite a few.

Sparse Blind Source Separation

Sparse BSS, introduced in Zibulevsky and Pearlmutter [2001] constitutes the framework of this thesis. We present it in details in the following.

1.2 Sparse Blind Source Separation

1.2.1 The concept of sparsity

Definition of sparsity

A signal $x \in \mathbb{R}^t$ is said to be exactly sparse if a very few number of its t entries are non-zero. For example, the signal depicted in figure 1.8a is exactly sparse since 98.8% of the entries take a value of 0.

The support of a sparse vector is defined as the set of indices of the non-zero entries, $\mathrm{supp}(x) = \{i / x[i] \neq 0\}$. Therefore, we have $\mathrm{card}(\mathrm{supp})(x) = \|x\|_0$. The notion of sparse signals support will be useful in the following.

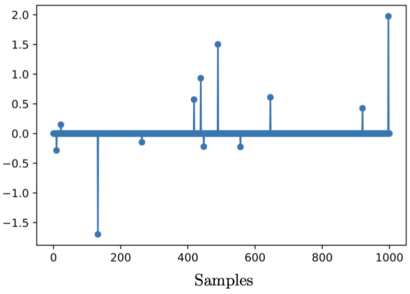

(a) Example of exactly sparse signal

Exact sparsity is generally not met in real-world data. We can however characterize a signal by its approximate sparsity. In this case, only a few entries have high amplitude. An approximately sparse signal can be reconstructed with a high accuracy by keeping only its high amplitude coefficients. A typical example of approximately sparse signal is a realization of a Generalized Gaussian distribution[2] of parameter $\rho \leq 1$ since the probability of having large entries is small. An example is displayed in figure 1.9. The corresponding histogram shows that the wide majority of entries have a value around 0.

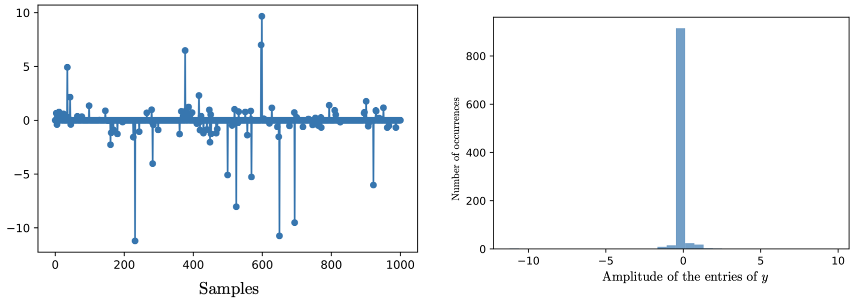

Figure 1.9: Example of approximately sparse signal (left figure) and its corresponding histogram (right.)

While sparse signals can be met in LC/MS data analysis, it is more likely to encounter natural signals with "specific structure" enabling their sparse representation in an adequate transformed domain.

Sparsity in a transformed domain

Let us consider a non-sparse signal $y \in \mathbb{R}^t$ and a dictionary $\mathbf{\Phi} \in \mathbb{R}^{T \times t}$ with $T \geq t$. The signal y can be decomposed as

$$y = \alpha \mathbf{\Phi},$$

where $\alpha \in \mathbb{R}^{1 \times T}$ contains the expansion coefficients of the input signal y in $\mathbf{\Phi}$. y is said to be sparse in the dictionary $\mathbf{\Phi}$ if most of the coefficients of α are zero.

The sparsity level of a signal representation highly depends on the transform. The selection of the appropriate basis or frame is driven by the structure (or *morphology*) of the signal. For example, let us suppose that we seek at storing the information of the signal shown in figure 1.10 at the least cost. In other words, we aim at encoding the energy of the signal in the least number of atoms possible. This signal can be decomposed as the sum of a piecewise constant signal and a cosine of frequency f. The representation in DCT domain is not sparse because of the piecewise continuous component. The Haar transform Haar [1910] is not well-suited neither for our problem due to the cosine. Nevertheless, the transform of the signal in the dictionary built as the union of the Haar transform and DCT transform is exactly sparse and answers to our problem. The constructed transform

---

[2] The probability distribution function of a Generalized Gaussian is given by $\frac{\rho}{2\Gamma(1/\rho)}\exp^{-|x|^{\rho}} \forall x \in \mathbb{R}$ where ρ is the shape parameter and Γ the gamma function.

is an over-complete/redundant dictionary. Despite its interesting properties, it requires careful attention since sparsity can be formulated in two non-equivalent ways, namely the analysis and the synthesis. The reader is invited to refer to Rapin [2014] for more information.

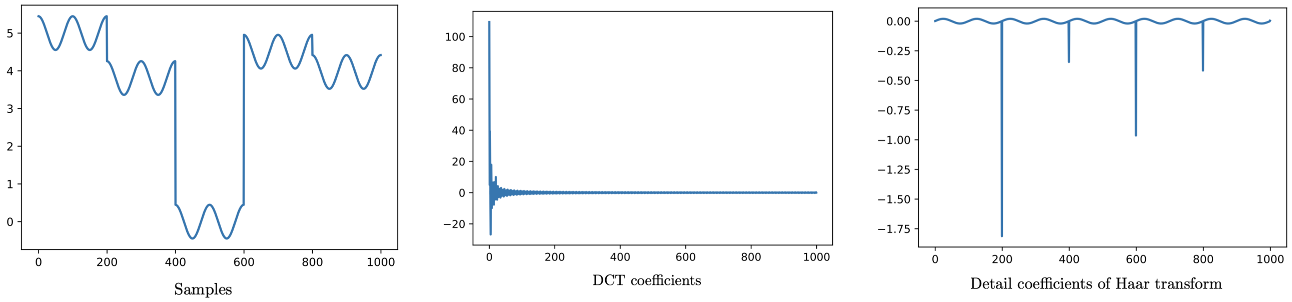

Figure 1.10: From left to right: signal x built as the sum of a cosine and piecewise constant signal; DCT transform of the signal x; Details coefficients of the Stationary Haar transform of x with a level of resolution of $J = 1$.

Over the years, many transforms have been proposed to sparsify signals of different morphologies. Among the most widely employed ones, we can find:

- DCT for oscillating signals;
- Wavelets for punctual singularities (Daubechies [1992], Starck et al. [2007]);
- Curvelets and ridgelets for lines and curvilinear structures such as edges (Starck et al. [2002], Candes [1998]) to only cite a few.

Additionally, sparse representation of a signal can also be obtained using dictionary learning methods.

1.2.2 The principle of morphological diversity

Sparsity has been extensively exploited in various domains among which image restoration or the separation of overlapping cartoon and texture. In this context, the authors of Elad et al. [2005] have developed the sparsity-based concept of morphological diversity.

Morphological diversity for signals of different structures

The concept of morphological diversity has been introduced in the context of image decomposition. It is assumed that an image can be decomposed into a structural part (cartoon layer that consists of large objects) and a textural part (containing the periodical fine components of the image). For example, the image depicted in 1.11 is the linear combination of two morphological components: the periodic textures and the cartoon shown in 1.11.

The image denoted $\mathbf{x}$ can be decomposed as

$$\mathbf{x} = \mathbf{x_1} + \mathbf{x_2}.$$

The two morphological components of the original image have different structures/morphologies. Therefore, each one of them is sparse in a different domain. We will denote $\mathbf{\Phi}_1$ (respectively $\mathbf{\Phi}_2$) the sparsity domain of $\mathbf{x_1}$ (respectively $\mathbf{x_2}$). For instance, the cartoon part can be appropriately sparsified in wavelet domain as it is the case for the wide majority of cartoons. The textured component shows a periodic structure, therefore the DCT transform is well-suited for sparse representation.

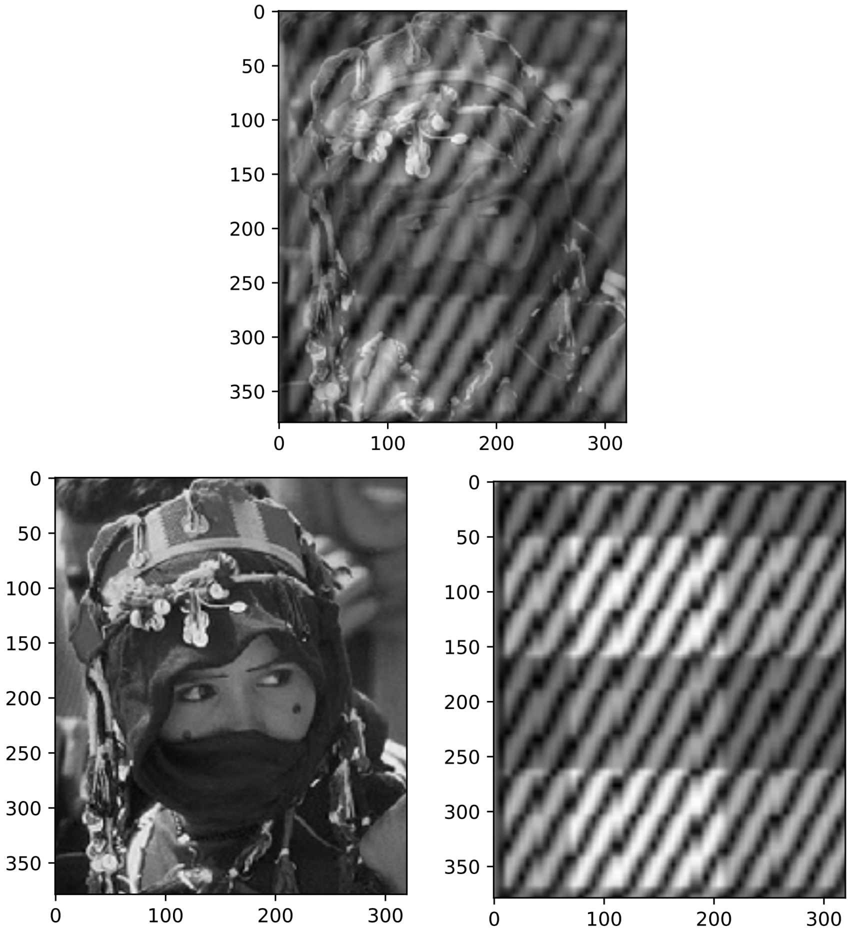

Figure 1.11: From left to right: synthetic image, decomposition of the original images into two morphological components: cartoon layer and periodical textured layer

The Morphological Diversity Principle is based on a very natural idea: the morphological components have different structures which confers them sparse representation in different transforms, and supposing a proper choice of the dictionaries, each of the component has the sparsest representation in its own dictionary. This can be written as:

$$\|\mathbf{x_i}\boldsymbol{\Phi}_i\|_0 < \|\mathbf{x_i}\boldsymbol{\Phi}_{j\neq i}\|_0. \tag{1.4}$$

In our example, the cartoon is less sparse in DCT domain than it is in wavelet domain and the texture has, in the same way, a sparser representation in DCT domain than in wavelet domain. Therefore, *the sparsity domain of the components act as discriminants between the components.*

Morphological diversity applied to Component separation Introduced in Starck et al. [2004], the Morphological Component Analysis algorithm exploits the morphological diversity principle in order to decompose an observation into its structural components.

Let $\mathbf{x} \in \mathbb{R}^t$ be a signal decomposable into the sum of morphological diverse signals:

$$\mathbf{x} = \sum_{i=1}^{n} \mathbf{x_i}. \tag{1.5}$$

We suppose that the following conditions are met:

- Each morphological component $\mathbf{x_i}$ admits a sparse representation in a well chosen dictionary $\boldsymbol{\Phi}_i \in \mathbb{R}^{t \times L_i}$;

- The morphological diversity principle is respected: the decomposition of any morphological component $\mathbf{x_i}$ in the dictionary $\mathbf{\Phi}_i$ is sparser than the decomposition of $\mathbf{x_i}$ in $\mathbf{\Phi}_j$ with $j \neq i$.

Supposing that the signal is contaminated by Gaussian noise, the decomposition can be formulated as the following optimization problem:

$$\min_{\mathbf{x_1},\mathbf{x_2},\ldots,\mathbf{x_n}} \|\mathbf{x} - \sum_{i=1}^{n} \mathbf{x_i}\|_F^2 + \sum_{i=1}^{n} \lambda_i \|\mathbf{x_i}\mathbf{\Phi}_i^T\|_0. \tag{1.6}$$

In eq. (1.6), the first term corresponds to the data-fidelity term which is formulated as Frobenius norm since we suppose the noise additive Gaussian. The second term is the ℓ_0 sparsity promoting term of the morphological components in their respective dictionaries with $\{\lambda_i\}_{i=1..n}$ the regularization parameters.

The optimization problem (1.6) being non-convex because of the ℓ_0 norm, the authors propose to substitute it by the closest convex norm adequate for sparsity, the ℓ_1 norm.
The problem (1.6) becomes then

$$\min_{\mathbf{x_1},\mathbf{x_2},\ldots,\mathbf{x_n}} \|\mathbf{x} - \sum_{i=1}^{n} \mathbf{x_i}\|_2^2 + \sum_{i=1}^{n} \lambda_i \|\mathbf{x_i}\mathbf{\Phi}_i^T\|_1. \tag{1.7}$$

In the following of this thesis, we will employ systematically the ℓ_1 norm to promote sparsity.

Morphological diversity for signals of similar structures

Morphological diversity has been introduced for the separation of different types of signal. The transform domains act as discriminants for the morphological components. A natural question arises: how to separate signals that have the same structure but vary in their composition?
Let us consider the two images $\mathbf{x_1}$ and $\mathbf{x_2}$ of figure 1.12: they share the same morphologies and have therefore sparse representations in the same transform domain denoted $\mathbf{\Psi}$. However, the human eye is able to distinguish between both of these photographs, notably because of the amplitudes, the emplacement of the fine-detailed structures etc.

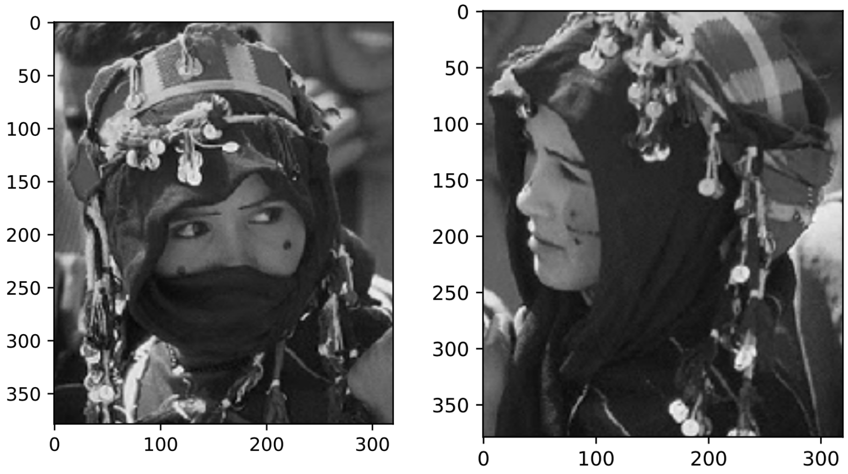

Figure 1.12: Two natural images with the same structures

The extension of Morphological Diversity introduced in Bobin et al. [2007] is based on the difference of information between the two images rather than their difference of morphology. Given that sparse representation of the two signals consist of a few active coefficients encoding most of the energy, it is very likely that the transform of $\mathbf{x_1}$ and $\mathbf{x_2}$ in their sparsity domain has different supports.

In this case, *it is the highest amplitude entries that play the role of discriminants* between the two signals, whereas in the case of morphological diversity as introduced in Elad et al. [2005], the distinction was made by the transformed domains. The scatter plot of $\mathbf{x_1}\mathbf{\Psi}^T$ versus $\mathbf{x_2}\mathbf{\Psi}^T$ is very characteristic of sparsity and morphological diversity in the same transformed domain:

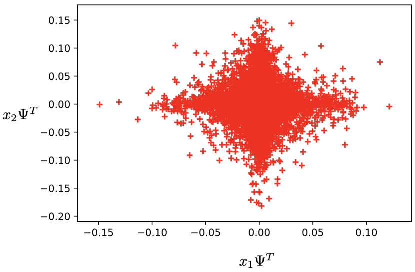

Figure 1.13: Scatter plot of the wavelet diagonal coefficients of the first image $\mathbf{x_1}$ versus the wavelet diagonal coefficients of $\mathbf{x_2}$

- The highest coefficients are aligned along the axes due to the sparsity of each individual signal;
- There are very few coefficients of high amplitudes for the two components simultaneously, which is descriptive of their morphological diversity.

Morphological Component Analysis algorithm has been extended to multivalued data. The main idea is to account for spatial and spectral morphologies. The method, coined mMCA (Bobin et al. [2005]) paved the way to morphological diversity and sparsity applied to BSS.

1.2.3 Sparsity and morphological diversity applied to BSS

The objective of sparse Blind Source Separation is to recover the mixing matrix and the sources while assuming that the sources are sparse in a possibly transformed domain and respect the Morphological Diversity Principle. The observation matrix consists of m observations or canals, each one having t samples:

$$\mathbf{X} = \mathbf{A}\mathbf{S} \tag{1.8}$$

where $\mathbf{S}$ is the concatenation in lines of sources sparse in a same domain $\mathbf{\Psi}$.

The classical problem of Blind Source Separation under sparsity assumption can thus be formulated as the following optimization problem:

$$\min_{\mathbf{A},\mathbf{S}} \|\mathbf{X}-\mathbf{A}\mathbf{S}\|_F^2 + \iota_{\{X:\forall i=1..n,\ \|X^i\|_2^2=1\}}(\mathbf{A}) + \|\mathbf{R}\odot\mathbf{S}\Psi^T\|_1, \tag{1.9}$$

where

- the first term corresponds to the data-fidelity term under LMM assumption considering the noise additive Gaussian;
- the second term enforces the normalization of the columns of the mixing matrix to avoid scaling indeterminacy;
- the third term is the sparsity enforcing penalization function applied to the sources in a transformed domain Ψ.

The figure 1.14 illustrates an example of multivalued data. We have considered the two sources of figure 1.12 to be mixed by a (5×2) mixing matrix. We do not take into account for now contamination with noise. The sources, as we have seen in the previous part, respect the Morphological Diversity Principle. Therefore, the plot of their transform has the particular cross shape characteristic of morphologically diverse components. Interestingly, the scatter plot of two observations in the domain $\mathbf{\Psi}$ (figure 1.16) has a rotated cross shape. Indeed, the observations correspond to the mixture of the sources accordingly to the directions given by the columns of the mixing matrix.

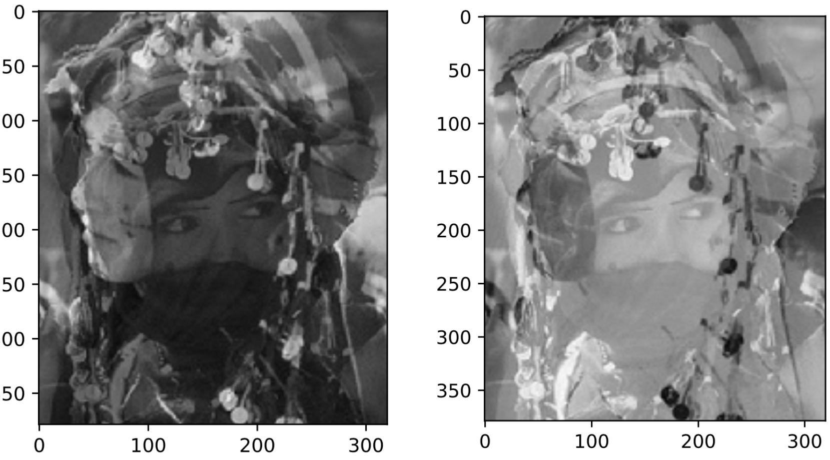

Figure 1.14: Two observations obtained by the mixture model (1.8)

Let $\hat{\mathbf{A}}$ be a matrix with the same span as the ground truth mixing matrix *i.e.* $\mathbf{X} = \hat{\mathbf{A}}\hat{\mathbf{S}}$ with $\hat{\mathbf{S}} = \hat{\mathbf{A}}^{\dagger}\mathbf{X}$. The couple$(\hat{\mathbf{A}}, \hat{\mathbf{S}})$ is therefore a possible solution to the ill-posed BSS problem. As illustrated in figure 1.16, the expansion coefficients of the decomposition in wavelets of $\hat{\mathbf{S}}_1$ and $\hat{\mathbf{S}}_2$ are not aligned on the canonical axes. The transforms of the obtained sources are not very sparse, and, a fortiori, they do not respect the morphological diversity principle.

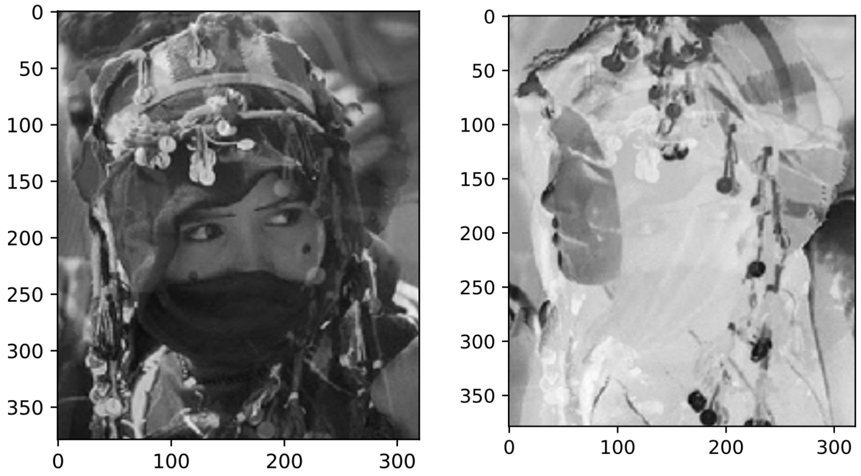

Figure 1.15: Sources $\hat{\mathbf{S}}$ estimated without accounting for sparsity

The loss of compressibility is interpretable geometrically: the ℓ_1 ball containing all the wavelet coefficients $\mathbf{S}_i\mathbf{\Psi}^T$ has a much smaller radius than the one enclosing the coefficients $\hat{\mathbf{S}_i}\mathbf{\Psi}^T$.

Taking into account sparsity and morphological diversity principles boils down to seeking for the solution of

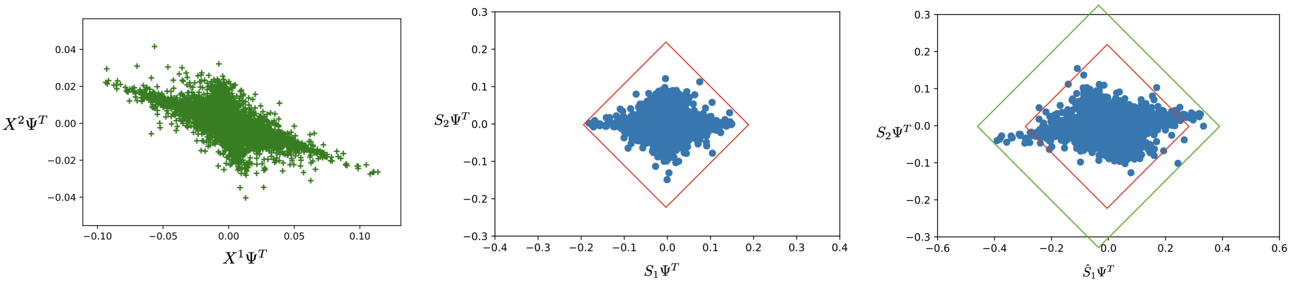

Figure 1.16: Left: Scatter plot of $X^1\boldsymbol{\Psi}^T$ versus $X^1\boldsymbol{\Psi}^T$. Middle: Scatter plot of $\mathbf{S}_1\boldsymbol{\Psi}^T$ versus $\mathbf{S}_2\boldsymbol{\Psi}^T$. Right: Scatter plot of $\hat{\mathbf{S}_1}\boldsymbol{\Psi}^T$ versus $\hat{\mathbf{S}_2}\boldsymbol{\Psi}^T$.

smallest ℓ_1 radius. Therefore, the range of solution of the ill-posed problem is much restrained and by recovering the sources, one can avoid interferences and therefore have a good estimation of the mixing matrix.

Sparsity and Morphological diversity act as discriminants between all the possible solutions not only for the source matrix but also for the mixing matrix.

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

- The matrices $\mathbf{A}$ and $\mathbf{S}$ are **full-rank**.
- The number of sources n is known.
- We place ourselves in the over-determined case and we consider that we have more samples than observations: $n \leq m \leq t$.
- The **columns of $\mathbf{A}$ are normalized** in the sense of ℓ_2 norm in order to avoid scaling indeterminacy.
- The sources are **sparse in the same domain** and **respect the Morphological Diversity principle**.

1.3 Sparse BSS for high-energy astrophysical data

The observation of the sky in high-energy is of uttermost importance to address some of the major challenges in modern astrophysics (*e.g.* the origin of Cosmic rays or the mechanisms of Supernovae explosion). Missions such as Chandra or Athena are telling examples of the ever growing interest in X-ray and Gamma-ray astronomy.

High-energy satellites provide multispectral images of the sky in which physical components are entangled. Deriving physical parameters from these data requires the unmixing of the elementary components observed, which is a BSS problem.

The authors of Picquenot et al. [2019b] have applied sparse BSS to the CasA Supernova Remnant acquired by Chandra telescope. In the energy range selected $(5 - 8 \text{ keV})$, the entangled physical components are known to be synchrotron continuum and iron emission line. The spectrum of Fe line is shifted by Doppler effect due to the expansion of supernova. It varies across the line of sight. The authors of Picquenot et al. [2019b] propose to simplify the problem by considering only the two extreme spectra of Fe that are observed in the image, namely the blue-shifted and the red-shifted spectrum. The figure 1.17 shows the three spectra considered in the mixture process.

The sparse BSS problem tackled in Picquenot et al. [2019b] is written as the following optimization problem:

$$\min_{\mathbf{A},\mathbf{S}} \frac{1}{2} f(\mathbf{A},\mathbf{S}) + h(\mathbf{S}) + g(\mathbf{A}),$$

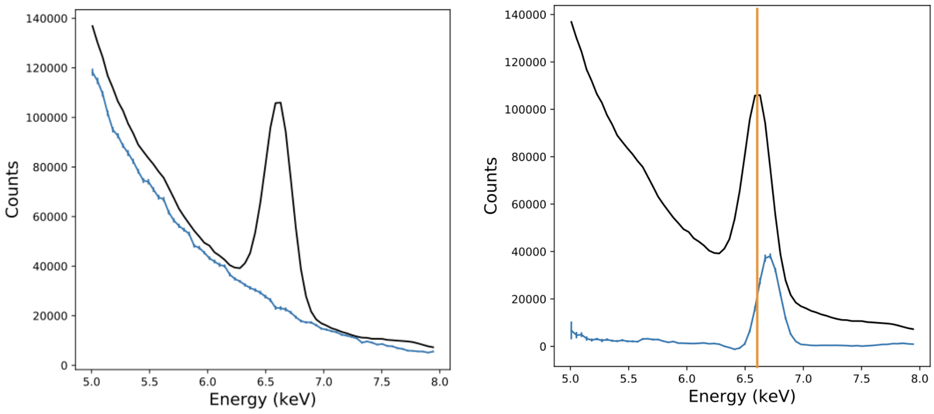

(a) Synchrotron emission

(b) blue-shifted Fe line

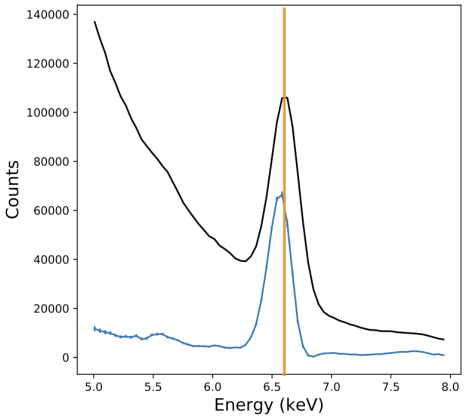

(c) red-shifted Fe line

Figure 1.17: Blue: Recovered spectra by Picquenot et al. [2019b]; Black: total spectral signature; Orange: position of the iron spectrum

where

- The mixing matrix $\mathbf{A}$ is composed of the spectra of the three physical components (Fe1, Fe2 and synchrotron);
- The lines of the source matrix $\mathbf{S}$ constitute the spatial distribution of the components;
- The function f is the data-fidelity term that measures the discrepancy between the model and the observations. The authors of Picquenot et al. [2019b] have considered that the contaminating noise is additive Gaussian, which leads to the use of the squared Frobenius norm,

$$f(\mathbf{A}, \mathbf{S}) = \| \mathbf{X} - \mathbf{A}\mathbf{S} \|_F^2.$$

- The penalization function h enforces the sparsity of the sources in the Starlet undecimated wavelet transform (Starck et al. [2015]), well-suited for astrophysical applications,

$$h(\mathbf{S}) = \sum_{i=1}^{3} \lambda_i \| \mathbf{S}_i \mathbf{\Phi}^T \|_1,$$

  where $\mathbf{\Phi} \in \mathbb{R}^{T \times t}$ denotes the Starlet transform and $\lambda_i \in \mathbb{R}^T$ stands for the regularization parameter corresponding to the $i-th$ source;

- The function g is employed to enforce the oblique constraint on the mixing matrix.

However, the proposed modelling is an **oversimplification** of the reality for two main reasons:

- The noise contaminating the observations is Poissonian, considering that the observations follow a Gaussian distribution can lead to erroneous results.
- The Doppler effects does not generate two Fe emission lines, it is in fact the same component that is affected by spectral variabilities.

1.3.1 Poisson noise

The authors of Picquenot et al. [2019b] propose to consider the classical Linear Mixture Model assuming the data are contaminated by additive Gaussian noise $\mathbf{X} = \mathbf{AS} + \mathbf{N}$.

However, this model does not describe high-energy astrophysical data that derive from photon count process. Instead, one should account for the exact statistics of the data that are corrupted with shot noise (Badenes [2010]).

The measurements can be defined statistically from the probability for a sample to take the value $\mathbf{X}_i[k]$ knowing the value of the pure mixture $\mathbf{AS}$:

$$\mathcal{P}(\mathbf{X}_i[t] = |[\mathbf{AS}]_i[t]) = \frac{e^{-[\mathbf{AS}]_i[t]}\,[\mathbf{AS}]_i[t]^{\mathbf{X}_i[t]}}{\mathbf{X}_i[t]!},$$

where $\mathbf{X}_i[k]$ (resp. $\mathbf{AS}]_i[t]$) is the sample of the observation matrix (resp. "pure mixture" matrix) at the $i-th$ line and $k-th$ column.

Despite being useful for various applications (NMR spectroscopy Sun and J.Xin [2011], Cherni et al. [2020], γ-ray spectroscopy Xu et al. [2019] and X-ray astrophysical imaging Badenes [2010]), performing sparse BSS while accounting for Poisson noise has never been studied to the best of our knowledge.

We propose in chapter 3 a novel method that specifically tackles the problem of sBSS for Poisson measurements and an algorithm, pGMCA, that we have applied to realistic simulations of Chandra observations of CasA supernova remnant.

1.3.2 Spectral variabilities

Due to the Doppler effect, the expansion of supernovae remnants generates spectral variabilities on components deriving from shocked ejecta such as Fe and Si emission lines. In order to truly account for these spatially varying spectra, one should consider not one mixing matrix per image but *sample-dependent mixing matrices* $\mathbf{A} \in \mathbb{R}^{m \times t \times n}$.

Consequently the generation process of the measurements does not follow the Linear Mixture Model but it can be modeled as

$$\mathbf{X}[k] = \sum_{i=1}^{n} \mathbf{A}^i[k]\mathbf{S}_i[k]$$

where $\mathbf{X}[k] \in \mathbb{R}^m$ designates the spectral signature of the $k-th$ pixel, $\mathbf{A}^i[k] \in \mathbb{R}^m$ indicates the spectral distribution of the $i-th$ element at the $k-th$ pixel and $\mathbf{S}_i[k]$ corresponds to the sample of the source matrix at the $i-th$ line and the $k-th$ column.

To the best of our knowledge, L-GMCA algorithm (Bobin et al. [2013]) is the only sparse BSS algorithm applied to astrophysics that tackles the problem of spectral variabilities. However, it does not propose a model explicitly accounting for variabilities neither it exploits their structure. Therefore, it is unable to provide a fine reconstruction of the SV (applications of L-GMCA to Chandra telescope data can be found in 5.5).

In the specific context of high-energy astrophysics, some non-BSS analysis tools have also been proposed to handle with spectral variabilities. In a nutshell, they consist of selecting "regions of interest" through human eye [3] and from which spectra are then extracted and fitted independently. The spatial information is not considered at all in the separation of the components which yields to poor disentanglement. Some analysis tools dedicated to GeV imaging propose to consider spatial structure of the components as well by fitting parametric shapes such as disks and Gaussians. However, these models are very simplistic and have no physical motivations. On the overall, State-of-the-Art analysis tools applied to astrophysical data with variabilities neither perform spectral unmixing of the sources (by exploiting the multi-dimensional aspect of the data as in BSS) nor propose reliable models.

To the best of our knowledge, the few models accounting for spectral variabilities are all intended to be applied to remote sensing. Consequently, we propose in chapter 5 a model that fully takes into account spectral variabilities and the specifics of astrophysical data. The developed algorithm, coined svGMCA has been applied to realistic X-ray astrophysical observations.

---

[3] The authors of Iwasaki et al. [2019] have recently proposed an automatic detection of the regions of interest.

Chapter 2

Algorithmic framework for Sparse Blind Source Separation

2.1 The optimization problem behind sparse BSS

The problem of Blind Source Separation under sparsity assumption can be formulated as the following optimization problem:

$$\min_{\mathbf{A},\mathbf{S}} \mathcal{F}(\mathbf{A},\mathbf{S}) + \mathcal{G}(\mathbf{A}) + \mathcal{H}(\mathbf{S}). \tag{2.1}$$

The cost function of equation (2.1) is constituted of three terms:

- $\mathcal{F}$ corresponds to the data-fidelity term which will be chosen identical to the negative likelihood of the variables $\mathbf{A}$ and $\mathbf{S}$ in this thesis. For instance, if we consider that the noise is additive Gaussian and that the LMM holds, this term would consist of the square Frobenius norm

  $$\mathcal{F}(\mathbf{A},\mathbf{S}) = \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_F^2,$$

  where $\mathbf{AS}$ corresponds to the noiseless mixture matrix under LMM assumption.
  Nevertheless, the LMM assumption does not hold true in X-ray imaging because of the presence of spectral variabilities that should be taken into account.
  Moreover, whenever data stem from a counting process such as Chandra telescope observations, the approximation of noise statistics by a normal distribution is no more valid. Consequently, the function $\mathcal{F}$ should be formulated as the antiloglikelihood of the shot noise (Poisson noise):

  $$\mathcal{F}(\mathbf{A},\mathbf{S}) = \mathbf{A}\mathbf{S} - \mathbf{X} \odot \log(\mathbf{A}\mathbf{S}),$$

  where $\odot$ stands for the Hadamard product.

  Regardless of the modelling of the data, we consider that the function $\mathcal{F}$ is **multi-convex**, meaning **it is not globally convex but convex towards each block individually**. Moreover, we assume that $\mathcal{F}$ is **differentiable with a Lipschitz gradient**.
  These assumptions hold true for the remainder of this thesis and more specifically in the case of shot noise where the Poisson antiloglikelihood is non-smooth: we propose in chapter 3 an approach using a differentiable approximation of the negative likelihood as data-fidelity term $\mathcal{F}$.

- The function $\mathcal{G}(\mathbf{A})$ is used to describe the penalizations on the mixing matrix:
  - It is customary in sparse BSS to consider at least one constraint on $\mathbf{A}$: the columns normalization employed in order to avoid scaling indeterminacy. It is enforced using the oblique constraint

    $$\mathcal{G}(A) = \iota_{\{X:\forall i=1..n,\ \|X^i\|_2^2=1\}}(\mathbf{A}),$$

    where ι_X stands for the characteristic function of the convex set X.
  - Since the oblique constraint is not a properly convex function, it can be substituted by enforcing the ℓ_2 norm of the columns to be smaller than 1:

    $$\mathcal{G}(A) = \iota_{\{X:\forall i=1..n,\ \|X^i\|_2^2\leq 1\}}(\mathbf{A}),$$

  - The positivity constraint, physically justified when the columns of the mixing matrix stand for spectra of physical components (C.F. chapters 3 and 5), can be added to the oblique constraint:

$$\mathcal{G}(\mathbf{A}) = \iota_{\{X:\forall i=1..n,\ \|X^i\|_2^2=1\}}(\mathbf{A}) + \iota_{.\geq 0}(\mathbf{A}).$$

- The term $\mathcal{H}(\mathbf{S})$ is employed to enforce the penalizations on the source matrix, for instance:

  - it can be used to promote sparsity of the sources in a domain $\Psi \in \mathbb{R}^{T\times t}$

$$\mathcal{H}(\mathbf{S}) = \|\mathbf{R} \odot \mathbf{S}\Psi^T\|_1,$$

    where $\odot$ stands for the Hadamard product. The matrix $\mathbf{R} \in \mathbb{R}^{n\times T}$ stands for the regularization parameters, controlling the trade-off between the data-fidelity term and the sparsity promoting term. It can be decomposed as $\mathbf{R} = \mathbf{DG}$ where $\mathbf{D} \in \mathbb{R}^{n\times n}$ is a diagonal matrix composed of the regularization parameters for each source and $\mathbf{G} \in \mathbf{R}^{n\times t}$ contains the individual penalization parameters if a reweighting ℓ_1 scheme Candes et al. [2008] is employed (otherwise, all coefficients of $\mathbf{G}$ are equal to 1);

  - Non-negativity can also be enforced on the source matrix, therefore the penalization function becomes:

$$\mathcal{H}(\mathbf{S}) = \iota_{.\geq 0}(\mathbf{S}).$$

  Throughout this thesis, we may employ other penalization function $\mathcal{G}$ and $\mathcal{H}$ depending on the problem considered. Both of them will be chosen **proper, lower semi-continuous, convex but not necessarily smooth**. We provide a list of these functions in appendix A.

Performing sparse BSS boils down to minimizing the cost function (2.1). However, this is a challenging task for two main reasons:

1. The classical gradient-based minimization methods cannot be employed in the case of function (2.1) because of the non-smoothness of the penalization terms $\mathcal{G}$ and $\mathcal{H}$.

   Nevertheless, these functions are closed proper convex and, consequently, admit a **proximal operator**. It is then possible to employ the so-called **proximal algorithms** to minimize such non-smooth convex functions Combettes and Pesquet [2011], Parikh et al. [2014].

2. Proximal algorithms enable the minimization towards each of the variables, however the minimization towards both of them which is required in BSS problems is tricky because of the non-convexity of the term $\mathcal{F}$. Indeed, the problem is not globally convex but has the particularity of being convex towards one variable while the other is fixed. Fortunately, optimization-based techniques have flourished over the last years offering a large framework of minimization strategies among which some are well adapted to our problem (Paatero and Tapper [1994], Bolte et al. [2014], Xu and Yin [2014], Tseng [2001] to only cite a few).

In this chapter, we will detail the above mentioned methods that can be used in sparse BSS problems. In the first part, we will explain the principle of proximal operator and how it can be used to minimize non-smooth functions. In the second part, we will review some popular minimization strategies employable to perform the minimization of the problem (2.1). Finally, we will discuss the applications of these methods to *standard* sparse BSS problems. This chapter is aimed at providing optimization tools that will be useful for the problems tackled in the next chapters of this thesis.

2.2 Quick introduction to non-smooth analysis

We propose a quick review of elementary concepts of non-smooth analysis.

Let $\mathbf{x} \in \mathbb{R}^t$ be a $t-$dimensional vector and f an extended real-valued function of $\mathbf{x}$.

- The effective domain of f is defined and denoted as $\operatorname{dom}(f) := \{\mathbf{x} \in \mathbb{R}^t, f(\mathbf{x}) < +\infty\}$.
- A set $\mathcal{C} \in \mathbb{R}^t$ is called convex if $\forall(\mathbf{x}, \mathbf{y}) \in \mathcal{C}^2$ and $\forall\lambda \in [0, 1]$,
$$\lambda\mathbf{x} + (1 - \lambda)\mathbf{y} \in \mathcal{C}.$$
- f is called convex if its effective domain is convex and $\forall(\mathbf{x}, \mathbf{y}) \in \operatorname{dom}(f)^2$ and $\forall\lambda \in [0, 1]$,
$$f(\lambda\mathbf{x} + (1 - \lambda)\mathbf{y}) \leq \lambda f(\mathbf{x}) + (1 - \lambda)f(\mathbf{y}).$$
It is strictly convex if its effective domain is convex and $\forall\mathbf{x} \neq \mathbf{y} \in \operatorname{dom}(f)^2$ and $\forall\lambda \in [0, 1]$,
$$f(\lambda\mathbf{x} + (1 - \lambda)\mathbf{y}) < \lambda f(\mathbf{x}) + (1 - \lambda)f(\mathbf{y}).$$
- The convex function f is called proper if its effective domain is nonempty and it never reach $-\infty$ *i.e.* there exists as least one $\mathbf{x} \in \mathbb{R}^t$ s.t. $f(\mathbf{x}) < \infty$ and $\forall\mathbf{x} \in \mathbb{R}^t$ $f(\mathbf{x}) > \infty$.
- f is called lower semi-continuous if its lower level sets defined as $\{\mathbf{x} \in \mathbb{R}^t : f(\mathbf{x}) \leq y\}$ are closed for every $a \in \mathbb{R}$.
- If f is a convex function and $\mathbf{x}$ is a stationary point of f, then $\mathbf{x}$ is a minimizer of f.
- A proper lower semi-continuous f admits a proximal operator defined as
$$\operatorname{prox}_f := \underset{\mathbf{u}\in\mathbb{R}^t}{\operatorname{argmin}} f(\mathbf{u}) + \frac{1}{2}\|\mathbf{u} - \mathbf{x}\|_2^2.$$

2.3 Minimization of convex non-smooth functions

2.3.1 Introduction

Let $\Gamma_0(\mathbb{R}^t)$ denote the class of closed proper convex functions $f \colon \mathbb{R}^t \to \mathbb{R} \cup \{\infty\}$ (that can also be called proximable functions) and let us consider the following problem:

$$\underset{\mathbf{x}\in\mathbb{R}^t}{\operatorname{minimize}} f(\mathbf{x}). \tag{2.2}$$

(2.2) is a non-smooth (convex) optimization problem for which classical gradient descent methods are excluded. It can however be solved by building a smooth approximation of the cost function f and minimizing the latter. The Moreau-Yosida envelope, defined as $M_f^t(\mathbf{x}) := \inf_{\mathbf{u}\in\mathbb{R}^t}(f(\mathbf{u}) + \frac{1}{2t}\|\mathbf{u} - \mathbf{x}\|^2)$ is a regularized (smooth) approximation of the function f with the same minimizers. Therefore, one can solve the optimization problem 2.2 by solving its exact smooth reformulation

$$\underset{\mathbf{x}\in\mathbb{R}^t}{\operatorname{minimize}} M_f(\mathbf{x}),$$

where the gradient of the Moreau envelope is given by $\nabla M_f(\mathbf{x}) = \mathbf{x} - \operatorname{prox}_f(\mathbf{x})$.

The proximal operator is thus the gradient step for minimizing a smooth approximation of f. In other words, **the proximal operator can be viewed as an analogy to gradient descent for non-smooth function**.

The minimization of the subproblem (2.1) towards one block ($\mathbf{A}$ or $\mathbf{S}$) can be efficiently performed by exploiting the proximal operator of the penalization terms $\mathcal{G}$ and $\mathcal{H}$.

The proximal operators used in this thesis are provided in appendix A.

2.3.2 Proximal algorithms

We present in this part two algorithms aiming at solving non-smooth optimization problems. They have been widely used in this thesis and are frequently mentioned in this manuscript.

Forward Backward Splitting algorithm

Let us consider the following minimization problem:

$$\min_{\mathbf{x}} f(\mathbf{x}) + g(\mathbf{x}), \tag{2.3}$$

where f is a differentiable function whose gradient is β-Lipschitzian and g a closed proper convex function.

Equation (2.3) is encountered when considering the minimization towards $\mathbf{S}$ of standard sparse BSS problem for example. We assume that the mixing process follows the LMM and that the sources are sparse in a transformed domain $\Psi \in \mathbb{R}^{T \times t}$. The minimization towards the source matrix is a convex non-smooth problem formulated as

$$\min_{\mathbf{S}} \frac{1}{2} \|\mathbf{X} - \mathbf{A}\mathbf{S}\|_F^2 + \left\|\Lambda \odot \mathbf{S}\Psi^T\right\|_1 . \tag{2.4}$$

The cost function consists of the sum of two terms: on the one hand the data-fidelity term expressed as the squared Frobenius norm and therefore differentiable and on the other hand the sparsity enforcing ℓ_1 norm on the sources in the transformed domain Ψ which is proximable but non-smooth.
Forward Backward Splitting algorithm (FB, Combettes and Wajs [2005]) is an easy-to-implement method aiming at solving non-smooth optimization problems of type (2.3). A general formulation can be found in algorithm 1.

Algorithm 1: Forward Backward algorithm

```
Initialization: $\gamma \in (0, \frac{2}{\beta})$.
while not converge do
    $\mathbf{x}^{(j+1)} = \text{prox}_{\gamma g}(\mathbf{x}^{(j)} - \gamma \nabla f(\mathbf{x}^{(j)}))$
    $j \leftarrow j + 1$
end
```

The FB algorithm can be split into two stages:

- A forward stage: it is the gradient descent operation involving the differentiable term of the functional f. In order to ensure the convergence of the algorithm, the stepsize γ should be chosen as described in algorithm 1;
- A backward stage: the proximal operator acts as a gradient step for the non-smooth function g. If f is null, the FB algorithm boils down to the proximal point algorithm.

The adaptation of FB algorithm to the problem (2.4) is known as ISTA algorithm (Parikh et al. [2014]) and can be accelerated using the FISTA variant (Beck and Teboulle [2009]).

If the sum of the proximable functions does not have a straightforward proximal operator but each ones does separately, we can solve the optimization function with the generalization of FB algorithm presented below.

Generalized Forward Backward Splitting algorithm

Let us consider the following optimization problem

$$\min_{\mathbf{x}} f(\mathbf{x}) + \sum_{i=1}^{n} g_i(\mathbf{x}), \tag{2.5}$$

with f being a proper convex and differentiable function of β-Lipschitzian gradient and the functions $g_1, .., g_n$ proximable functions. We also consider that FB algorithm would not be efficient to solve this problem *i.e.* $\sum_{i=1}^{n} g_i$ does not have an easy-to-compute proximal operator.

The so-called Generalized Forward Backward Splitting algorithm (GFB, Raguet et al. [2013]) that has been introduced to conveniently minimize the functional (2.5) is as follows:

Algorithm 2: Generalized Forward Backward algorithm

```
Initialization: $\gamma \in (0, \frac{2}{\beta})$, $\mu \in (0, \min(\frac{3}{2}, \frac{1+\frac{2}{\beta\gamma}}{2}))$, $w_i \in (0,1)$ s.t. $\sum_{i=1}^{n} w_i = 1$, $\mathbf{y}_i^{(0)} \leftarrow \mathbf{x}^{(0)}\ \forall i = 1..n$.
while not converge do
    for $i = 1...n$ do
        $\mathbf{y}_i^{(j+1)} \leftarrow \mathbf{y}_i^{(j)} + \mu\, \mathrm{prox}_{\frac{\gamma}{w_i} g_i}(2\mathbf{x}^{(j)} - \mathbf{y}_i^{(j)} - \gamma \nabla f(\mathbf{x}^{(j)})) - \mu \mathbf{x}^{(j)}$
    end
    $\mathbf{x}^{(j+1)} \leftarrow \sum_{i=1}^{n} w_i \mathbf{y}_i^{(j+1)}$
    $j \leftarrow j + 1$
end
```

GFB algorithm alternates between an explicit (forward) gradient descent operation on the differentiable function f and an implicit (backward) step of application of the proximal operators of the g_i functions onto intermediate variables $\mathbf{y}_i$. Its convergence is ensured when the two hyperparameters γ and μ are chosen as in algorithm 2.

We have seen that, when optimizing standard BSS problem towards $\mathbf{S}$, the Forward Backward algorithm can be used if we enforce the sparsity constraint in direct domain. This is not the case when we suppose the sources positive in direct domain and sparse in a transformed domain (even if the transform is orthogonal) since the sum of the two constraints does not have a straightforward proximal operator. This problem falls into the scope of application of GFB algorithm but not FB algorithm.

2.4 Multi-convex optimization

We have seen in the previous section how the concept of proximal operator can be employed to solve non-smooth optimization problems.

We propose now to review the major strategies that will be extensively used in this thesis to tackle sparse BSS problems in a global manner. These approaches are all based on optimization approaches and built upon the specific multi-convex structure of the cost function.

2.4.1 Algorithmic framework

Optimization problem

The minimization of the sparse BSS problem of interest, that we have formulated in (2.1) as:

$$\min_{\mathbf{A},\mathbf{S}} \mathcal{F}(\mathbf{A},\mathbf{S}) + \mathcal{G}(\mathbf{A}) + \mathcal{H}(\mathbf{S}),$$

is very challenging due to the multi-convexity of the data-fidelity term $\mathcal{F}$. Moreover, the two blocks $\mathbf{A}$ and $\mathbf{S}$ are coupled since $\mathcal{F}$ measures the divergence between the data matrix $\mathbf{X}$ and the pure mixture $\mathbf{AS}$. This problem calls for specific minimization strategies.

Let us consider the following general formulation of multi-convex problems that comprises the problem (2.1) of interest:

$$\min_{Y_1 \in \mathbb{R}^{m_1 \times t_1},\ Y_N \in \mathbb{R}^{m_N \times t_N}} F(Y_1, ..., Y_N) := f(Y_1, ..., Y_N) + \sum_{I=1}^{N} g_i(Y_i). \tag{2.6}$$

In equation (2.6), the involved functions and variables are such that:

- $f : \mathbb{R}^{m_1 \times t_1} \times ... \times \mathbb{R}^{m_N \times t_N} \to \mathbb{R}$ is a block multi- convex function meaning it is not necessarily globally convex but convex towards one block. If we define for each of these N blocks Y_i the function f_i s.t.

  $$\begin{aligned} f_i \colon \mathbb{R}^{m_i \times t_i} &\to \mathbb{R} \\ Y &\to f(Y_1, ..., Y_{i-1}, Y, Y_{i+1}, ..., Y_N); \end{aligned}$$

  The functions $f_{i=1..n}$ are convex and differentiable with β_i-Lipschitz gradient.
  In (2.1), the data-fidelity term $\mathcal{F}$ plays the role of the function F.

- Each function g_i is proper, lower semi-continuous and convex but not necessarily smooth.
  In sparse BSS formulation, there are two blocks ($\mathbf{A}$ and $\mathbf{S}$) and two proximable functions ($\mathcal{G}$ and $\mathcal{H}$).

In the remainder of this chapter, we will denote F_i the following function:

$$\begin{aligned} F_i \colon \mathbb{R}^{m_i \times t_i} &\to \mathbb{R} \\ Y &\to F(Y_1, ..., Y_{i-1}, Y, Y_{i+1}, ..., Y_N). \end{aligned}$$

The idea behind the approaches that we present in this part is a partition of the coordinates into N blocks ($Y_{i=1..n}$ in equation (2.6), and $\mathbf{A}$ and $\mathbf{S}$ in equation (2.1)) and a minimization of the cost function with respect to one block while the others are held fixed. The algorithms that will be presented in the following are all part of the Gauss-Seidel family and share the same structure, their difference lies in the update scheme of each block.

In this section, we present the BCD algorithm (Tseng [2001]) that will be used as optimization strategy in the next chapters, PBC algorithm (Attouch et al. [2010]) which has interesting properties that merit to be mentioned and finally PALM (Bolte et al. [2014] and pALS (Paatero and Tapper [1994]) algorithms that have been studied in sparse BSS framework.

2.4.2 Block Coordinate Descent Algorithm

Principle

The update scheme of the Block Coordinate Descent algorithm (BCD, Tseng [2001]) is cyclic and consists in the exact minimization of each subproblem towards the corresponding block while the others are fixed. A description of BCD is provided in algorithm 3. The structure of the subproblems and the conditions required for f and g allow to use proximal algorithms for each iteration; for example one can use the FB algorithm (or an accelerated version) or its generalization if necessary.

Algorithm 3: Block Coordinate Descent Algorithm.

```
Initialization: $Y_1^0, ..., Y_N^0$.
while not converge do
    for $i = 1...N$ do
        $Y_i^{(j)} \leftarrow \mathrm{argmin}_{Y_i}( f_i^{(j)}(Y_i) + g_i(Y_i))$
    end
    $j \leftarrow j + 1$
end
```

Application to sparse BSS

For more clarity, we present the application of BCD algorithm to a standard sparse BSS problem formulated as

$$\min_{\mathbf{A},\mathbf{S}} \frac{1}{2}\|\mathbf{X} - \mathbf{A}\mathbf{S}\|_F^2 + \iota_{OB(m)}(\mathbf{A}) + \|\Lambda_{\mathbf{S}} \odot \mathbf{S}\Phi^T\|_1. \tag{2.7}$$

In this equation, the Frobenius norm stands for the data-fidelity term assuming the noise is Gaussian. The ℓ_1 norm on the sources with the corresponding regularization parameters $\Lambda_{\mathbf{S}}$ is the sparsity-promoting term and the indicator function enforces the oblique constraint on $\mathbf{A}$. Both of these terms act as the proximable functions $g_{i=1..n}$ present in eq. (2.6).

The adaptation of BCD to the standard sparse BSS problem given in eq. (2.7) consists of a sequential optimization over each variable while the other is held fixed. It is provided in algorithm 4.

The subproblem towards the sources corresponds to the sum of a quadratic term and ℓ_1 sparsity-promoting constraint; it can therefore be efficiently solved with ISTA or its accelerated version FISTA. The update towards the mixing matrix consists of a simple FB algorithm since the proximal operator of the oblique constraint is explicit.

It must be noted that the cyclic estimation of the subproblems comes with a high computational cost. In general, BCD can be an interesting choice to solve sparse BSS problems, the exact minimization of each subproblem providing stable solutions with the sought structures. For example, Rapin [2014] showed that it is an appealing approach when positivity is enforced for the couple $(\mathbf{A}, \mathbf{S})$. However, particular attention must be drawn to the choice of regularization parameters and initialization, to which BCD is extremely sensitive unlike other algorithms (C.F. pALS algorithm, Paatero and Tapper [1994], Berry et al. [2007], section 2.4.5).

Convergence conditions

The author of Tseng [2001] has shown that each cluster point of iterates generated by BCD method is a stationary point provided certain conditions. In the case of eq. (2.7), the convergence of BCD is proven since the conditions stated in theorem 4.1.b are met:

- If the starting couple $(\mathbf{A}^{(0)}, \mathbf{S}^{(0)})$ respects the condition given by the indicator function, then the level set $X^0 = (\mathbf{A}, \mathbf{S}) : f(\mathbf{A}, \mathbf{S})$ is compact and f is continuous on X^0.
- Since f is multiconvex, it is convex in $(S_1^1, ..., S_n^t)$ (with $\mathbf{S} = [S_j^k]_{j=1...n,k=1...t} \in \mathbb{R}^{n \times t}$) and convex in $\mathbf{A}$.
- The domain of the function f^0 defined as $f^0(\mathbf{A}, \mathbf{S}) = \frac{1}{2}\|\mathbf{X} - \mathbf{A}\mathbf{S}\|_F^2$ is open and f^0 is differentiable on $\mathrm{dom}(f^0)$, therefore, according to Lemma 3.1 concerning assumption A.1 of Tseng [2001], f^0 is regular on X^0.

Algorithm 4: BCD applied to sparse BSS

```
Result: $\mathbf{A}, \mathbf{S}$
Initialization: $\mathbf{A}^{(0)}$ , $\mathbf{S}^{(0)}$
while not converge do
    • Update $\mathbf{S}^{(k)}$ assuming $\mathbf{A}^{(k)}$ is fixed using FB algorithm:
    Compute $L_{\mathbf{S}} = \|\mathbf{A}^{(k)T}\mathbf{A}^{(k)}\|_{s,2}$
    $\mathbf{S}^{(0)} \leftarrow \mathbf{S}^{(k)}$
    while not converge do
        $\mathbf{S}^{(j+1)} \leftarrow \mathcal{S}_{\frac{\Lambda_{\mathbf{S}}}{L_{\mathbf{S}}}}(\mathbf{S}^{(j)} + \frac{1}{L_{\mathbf{S}}}\mathbf{A}^{(k)T}(\mathbf{X} - \mathbf{A}^{(k)}\mathbf{S}^{(j)}))$
        $j \leftarrow j + 1$
    end
    $\mathbf{S}^{(k+1)} \leftarrow \mathbf{S}^{(j)}$
    • Update $\mathbf{A}^{(k)}$ assuming $\mathbf{S}^{(k+1)}$ is fixed using FB algorithm:
    Compute $L_{\mathbf{A}} = \|\mathbf{S}^{(k+1)}\mathbf{S}^{(k+1)T}\|_{s,2}$
    $\mathbf{A}^{(0)} \leftarrow \mathbf{A}^{(k)}$
    while not converge do
        $\mathbf{A}^{(j+1)} \leftarrow \mathbf{A}^{(j)} + \frac{1}{L_{\mathbf{A}}}(\mathbf{X} - \mathbf{A}^{(j)}\mathbf{S}^{(k+1)})\mathbf{S}^{(k+1)T}$
        for i=1...n do
            $\mathbf{A}^{i(j+1)} \leftarrow \frac{\mathbf{A}^{i(j+1)}}{\max(1, \|\mathbf{A}^{i(j+1)}\|_2)}$
        end
        $j \leftarrow j + 1$
    end
    $\mathbf{A}^{(k+1)} \leftarrow \mathbf{A}^{(j)}$
    $k \leftarrow k + 1$
end
```

- We use BCD with a cyclic rule.

Unless stated otherwise, all of the minimization problems studied in this thesis meet the convergence requirements of BCD algorithm. The reader is invited to refer to theorem 4.1.b of Tseng [2001] for more details.

2.4.3 Proximal Block Coordinate Algorithm

Principle

In Attouch et al. [2010], the authors propose a sequential minimization of the proximal regularization of the subproblems. The update of the block coordinate X_i at the iteration k is performed by solving the following problem

$$\min_{Y_i^{(k+1)} \in \mathbb{R}^{m_i \times t_i}} F_i^{(k)}(Y_i^{(k)}) + \frac{1}{2L_i^{(k)}}\|Y_i - Y_i^k\|_2^2, \tag{2.8}$$

where $L_i^{(k)}$ a numerical constant corresponding to the i-th block at the k-th iteration, it is strictly lower bounded by 0 and upper bounder. The addition of the proximity operator allows to rewrite eq. (2.8) equivalently as

$$\mathrm{prox}_{L_i^{(k)}(F_i^{(k)})}(Y_i(k)). \tag{2.9}$$

Therefore, the minimization of the subproblem is performed locally by definition of the proximal operator which allows for faster computational time in comparison with BCD method.
The structure of the PBC algorithm is as follows:

Algorithm 5: Proximal Block Coordinate Algorithm.

```
Initialization: $Y_1^0, ..., Y_N^0$.
while not converge do
    for $i = 1...N$ do
        $Y_i^{(j+1)} \leftarrow \mathrm{prox}_{L_i^{(k)}(F_i^{(k)})}(Y_i^{(j)})$
    end
    $j \leftarrow j + 1$
end
```

2.4.4 Proximal Alternating Linearized Minimization (PALM) algorithm

Principle

The idea behind PALM algorithm comes from the combination of the "explicit (forward)/ implicit (backward)" principle of FBS on non-smooth functions and the proximal regularization of multiconvex problems proposed in PBC. As said before, the authors of Attouch et al. [2010] have proposed an interesting minimization of a local regularization of the Gauss-Seidel scheme. However, its application is hampered by computational pitfalls. In order to alleviate such problems, the authors of Bolte et al. [2014] propose an other approximation of the coordinate subproblems, the proximal linearization. The k^{th} update of a given block consists then in optimizing the proximal regularization of the smooth part linearized at the last iterate, *i.e.*

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

or equivalently

$$Y_i^{(k)} = \mathrm{prox}_{\frac{g_i}{L_i}}(Y_i^{(k-1)} - \frac{1}{L_i^{(k-1)}}\nabla f_i^{(k-1)}(Y_i^{(k-1)})),$$

where $L_i^{(k-1)}$ is a numerical constant that we take equal to the Lipschitz constant of the smooth function.
PALM algorithm can be seen as an alternating minimization approach of the FBS, since the update on each block of coordinate is composed of one gradient (forward) step on the smooth part and a proximal (backward) step on the non-smooth part. The complete method is described in algorithm 6.

Algorithm 6: Proximal Alternating Linearized Minimization Algorithm

```
Initialization: $Y_1^0, ..., Y_N^0$.
while not converge do
    for $i = 1...N$ do
        $Y_i^{(j+1)} \leftarrow \text{prox}_{\frac{g_i}{L_i^{(k)}}}(Y_i^{(j)} - \frac{1}{L_i^{(k)}} \nabla f_i^{(k)}(Y_i^{(j)}))$
    end
    $j \leftarrow j + 1$
end
```

The prox-linearization approach enhances computational time in comparison with BCD, provided the proximal operator of the non-smooth part is explicit. Moreover, it does not requires, as it is the case for PBC, to have an explicit proximal operator of the whole cost function or to encumber the method with subroutines possibly yielding to errors.

Application to sparse BSS

PALM has been studied and applied in many BSS problems. In particular, its performances have been empirically investigated in Kervazo [2019] which has shown that PALM is very sensitive to its initialization and the regularization parameters of the cost function.
The application of PALM algorithm to standard sparse BSS (2.7) given in algorithm 7 shows a major difference with BCD: the update of the block coordinate depends on the current estimates.

Algorithm 7: PALM applied to sparse BSS

```
Result: $\mathbf{A}, \mathbf{S}$
Initialization: $\mathbf{A}^{(0)}$ , $\mathbf{S}^{(0)}$
while not converge do
    • Update of $\mathbf{S}$:
    Compute $L_{\mathbf{S}} = \|\mathbf{A}^{(k)T}\mathbf{A}^{(k)}\|_{s,2}$
    $\mathbf{S}^{(k+1)} \leftarrow \mathcal{S}_{\frac{\Lambda_{\mathbf{S}}}{L_{\mathbf{S}}}}(\mathbf{S}^{(k)} + \frac{1}{L_{\mathbf{S}}}\mathbf{A}^{(k)T}(\mathbf{X} - \mathbf{A}^{(k)}\mathbf{S}^{(k)}))$
    • Update of $\mathbf{A}$:
    Compute $L_{\mathbf{A}} = \|\mathbf{S}^{(k+1)}\mathbf{S}^{(k+1)T}\|_{s,2}$
    $\mathbf{A}^{(k+1)} \leftarrow \mathbf{A}^{(k)} + \frac{1}{L_{\mathbf{A}}}(\mathbf{X} - \mathbf{A}^{(k)}\mathbf{S}^{(k)})\mathbf{S}^{(k)T}$
    for i=1...n do
        $\mathbf{A}^{i(k+1)} \leftarrow \frac{\mathbf{A}^{i(k+1)}}{\max(1, \|\mathbf{A}^{i(k+1)}\|_2)}$
    end
end
$k \leftarrow k + 1$
```

Convergence conditions

Convergence analysis of PALM algorithm relies on the assumption that the cost function to minimize satisfies the Kurdyka-Łojasiewicz properties (Kurdyka [1998], Lojasiewicz [1963]). Roughly speaking, the interest on the

so-called KL functions lies in the fact that they are sharp near a critical point $\bar{u}$ up to a reparametrization of its values. In other words, there exists a function ϕ (called desingularizing function) such that the subgradients of $u \rightarrow \phi \circ (\sigma(u) - \sigma(\bar{u}))$ has a norm greater than 1 on its whole range.
The convergence of PALM algorithm for minimizing the cost function (2.7) is guaranteed since the following requirements are met:

- The cost function F is semi-algebric and therefore it is a KL function. Indeed, it is the sum of a real polynomial function (the data-fidelity term), an indicator function of a semi-algebric set (constraint on the mixing matrix) and a ℓ_1 norm (constraint on the sources).
- Each of the constraints is a proper and lower semicontinuous function.
- The data-fidelity term f_0 is a C^2 function.
- The partial gradients of f_0 towards each of the two block coordinates ($\nabla_x f_0$ and $\nabla_y f_0$) are globally Lipschitzian.
- Both of the Lipschitz constants of $\nabla_x f_0$ and $\nabla_y f_0$ are bounded and non-zero.
- PALM is implemented with cyclic rule.

Unless stated otherwise, all of the minimization problems studied in this thesis meet the convergence requirements of PALM algorithm. The reader is invited to refer to theorem 1 of Bolte et al. [2014] for more details.

2.4.5 Projected Alternating Least square Algorithm

Principle

The Projected Alternating Least square (pALS) Algorithm, introduced in Paatero and Tapper [1994] and Berry et al. [2007] differs from the previously cited algorithms since it does not minimize the Gauss-Seidel scheme or its regularization. In fact, it is not aimed at properly minimizing the cost function but rather proposing an efficient and fast method to solve the problem.
pALS algorithm consists of an alternation of two steps:

- a minimization of the smooth part f_i, which boils down to a least-squares estimate;
- a projection of the solution onto the non-smooth part g_i using the proximal operator.

The complete pALS scheme is given in algorithm 8. This method is computationally fast and easy to implement provided the proximal operator is analytical. It is widely used in NMF or sparse BSS (CF 2.5). However, it loses of its performances if the constraints are slightly less "simple"; for example, the author of Rapin [2014] has shown that it is much less efficient than BCD algorithm when sparsity and non-negativity constraint are simultaneously enforced. Moreover, there is no mathematical guarantee that it would converge towards a stationary point or even minimize the cost function.

Algorithm 8: Projected Alternating Least Squares algorithm

```
Initialization: $Y_1^0, ..., Y_N^0$.
while not converge do
    for $i = 1...N$ do
        $Y_i^{(j+1)} \leftarrow \text{argmin}_{Y_i} f_i^{(j)}(Y_i)$
        $Y_i^{(j+1)} \leftarrow \text{prox}_{g_i}(Y_i^{(j+1)})$
    end
    $j \leftarrow j + 1$
end
```

2.5 Heuristics and optimization strategy for sparse BSS

The recent breakthroughs in optimization provide a wide range of optimization strategies adapted to the minimization of various cost-functions. In particular, they have been of great use for the minimization of standard sparse BSS problems that are of interest for us in this thesis. In this part, we propose to introduce the approaches that have been adopted in sparse BSS and in this thesis.

2.5.1 Introduction

In this section, we focus on a *standard* sparse BSS problem formulated as

$$\min_{\mathbf{A},\mathbf{S}} \frac{1}{2}\|\mathbf{X} - \mathbf{A}\mathbf{S}\|_F^2 + \iota_{OB(m)}(\mathbf{A}) + \|\Lambda_{\mathbf{S}} \odot \mathbf{S}\|_1. \tag{2.10}$$

This cost function encompasses a large number of sparse BSS problems since it is built upon a very few number of assumptions:

- The data are supposed to follow the Linear Mixture Model and to be corrupted with additive Gaussian noise;
- The sources are supposed to be sparse in the direct domain;
- The only penalization applied to the mixing matrix is the oblique constraint used to alleviate scaling indeterminacy.

The minimization of this problem requires, on the one hand, to select among the optimization approaches the most adapted to the problem, and on the other hand to set properly the regularization parameters. In the following, we start by introducing a standard strategy for the minimization of the problem (2.10) that we have extensively used in this thesis, namely the Generalized Morphological Component Analysis (GMCA, Bobin et al. [2007]). The efficiency of this algorithm finds its roots in the automatic choice of regularization parameters it proposes, that we will detail in a second part.

2.5.2 GMCA as an enhanced pALS algorithm

GMCA is a built-upon pALS algorithm for sparse BSS whose update scheme is provided in algorithm 9.

Algorithm 9: pALS applied to sparse BSS

```
Initialization: $\mathbf{A}^{(0)}$ , $\mathbf{S}^{(0)}$
while not converge do
    • Update of $\mathbf{S}$:
    $\mathbf{S}^{(k+1)} \leftarrow \mathcal{S}_{\frac{\Lambda_\mathbf{S}}{L_\mathbf{S}}}(\mathbf{A}^{(k)\dagger}\mathbf{X})$ $\mathbf{A}^{(k+1)} \leftarrow \mathbf{X}\mathbf{S}^{(k+1)\dagger}$
    • Update of $\mathbf{A}$:
    for i=1...n do
        $\mathbf{A}^{i(k+1)} \leftarrow \frac{\mathbf{A}^{i(k+1)}}{\max(1,\|\mathbf{A}^{i(k+1)}\|_2)}$
    end
end
$k \leftarrow k+1$
```

A key element in GMCA algorithm is the setting of the regularization parameter Λ. This latter controls the trade-off between the data-fidelity term and the sparsity penalization. Therefore, a wise choice is required since the solution tends to be very sensitive to the threshold.

The author of Chenot [2017] has shown that the minimization of the problem (2.10) with pALS algorithm is much affected by the choice of Λ:

- A high threshold is required to ensure *robustness* to noise since it enhances the denoising process and provides better results with noisy data. Moreover, it is less sensitive to initialization.
- A low threshold provides more accurate results in the case of noiseless data with fair initialization.

Thus, fixing regularization parameters is a very tricky issue since it depends on the dataset, the initialization and prior information on noise level that are generally unknown.

Adaptive thresholding strategy

In order to cope with a difficult choice of parameters, the authors of Bobin et al. [2007] have proposed a **decreasing thresholding strategy**. The idea behind this **coarse-to-fine** strategy is to gather robustness and insensitivity to the initialization provided by high thresholding at the first iterations, and refinement and accuracy of the results thanks to low final thresholds.

Let us consider that the matrix $\Lambda \in \mathbb{R}^{n \times t}$ is a diagonal matrix, meaning that there is only one threshold per source. The update of the estimated source $\mathbf{S}_i$ in pALS algorithm corresponds to a thresholded version of the least-squares estimate:

$$\mathbf{S}_i^{(k+1)} = \mathcal{S}_{\Lambda_i}(\mathbf{A}^{i(k)\dagger}\mathbf{X}),$$

which, at convergence of the algorithm becomes:

$$\mathbf{S}_i^{(k+1)} = \mathcal{S}_{\Lambda_i}(\mathbf{S}_i^* + \mathbf{A}^{i*\dagger}\mathbf{N}),$$

where $\mathbf{S}^*$ is the true mixing matrix and $\mathbf{A}^{i*\dagger}\mathbf{N}$ is the Gaussian noise contaminating the source i. The term Λ_i corresponds to the threshold applied the $i-th$ source.

Setting the final threshold for each source accordingly to noise level boils down to a simple detection procedure: it should be higher than noise level in order to set at zero noise entries, and lower than largest entries that are very likely to correspond to the support of the true sources.

This is a well-known problem in the case of Gaussian noise which is solved by using the "$\kappa - \sigma$" rule. It amounts to fixing the threshold at $3-\sigma_i$ with σ_i the standard deviation of the noise $\mathbf{A}^{i(k)\dagger}\mathbf{X}$ since the probability that an entry of amplitude higher than the $3-\sigma_i$ threshold belongs to the noise is less than 0.4% (the noise follows a Gaussian distribution).

Since noise statistics are generally unknown, it is customary to use the Median Absolute Deviation (MAD) operator as an approximation of standard deviation:

$$\sigma_i \simeq 1.48\mathrm{MAD}(\mathbf{A}^{i(k)\dagger}\mathbf{X}),$$

with $\forall y \in \mathbb{R}^t,\ \mathrm{MAD}(y) = \mathrm{median}_{i=1..t} \mid y_i - \mathrm{median}_{i=1..t}(y_i) \mid)$. Besides, the MAD operator is very well-suited to our problem since it is insensitive to sparse contamination. Therefore, the selection of the threshold Λ_i boils down to

$$\Lambda_i = 1.48 \times 3 \times \mathrm{MAD}(\mathbf{A}^{i*\dagger}\mathbf{N}).$$

In order to propose a decreasing strategy that accounts for the evolution of the sources and the mixing matrix during the iterations, the thresholds are computed at each iteration (k) according to the current estimates $\Lambda_i = \mathrm{MAD}(\mathbf{A}^{i(k)\dagger}\mathbf{X})$. Thus, the threshold at final iterations will be computed on sources and mixing matrix close to the true ones.

Interpretation in terms of Morphological Diversity

The success of GMCA thresholding strategy can be interpreted in terms of Morphological Diversity. As we have mentioned in chapter 1, Morphological Diversity states that the coefficients of highest amplitudes are the most discriminant since they are very likely to correspond to one source only. GMCA exploits the prominent role of high amplitude coefficients by considering high thresholds at the first iterations. This allows not only to perform component separation but also to be robust to noise. Indeed, the highest coefficients are less corrupted by noise and therefore, the estimation based on such coefficients will be most accurate.

However, if the thresholds are set to high values all along the algorithm, the estimation of the sources may be based on not enough statistics (if there are very few high-amplitude coefficients) that can lead to erroneous results. In addition to that, such a strategy neglects the possibility of partially correlated sources notably for high amplitudes. Finally, setting to zero low amplitude coefficients can be counterproductive in the case where morphological diversity is merely an approximation of the dataset. The decrease of the threshold allows for low amplitude coefficients to be accounted for gradually.

Concerning the mixing matrix, according to the Morphological Diversity Principle, the highest coefficients are aligned according to the mixing directions. By selecting these coefficients as a first step, GMCA is able to sketch the mixing matrix. The estimation is then refined with the introduction of lower amplitude coefficients thanks to lower thresholds.

In conclusion, the adaptive thresholding strategy has a crucial role on two important aspects of BSS:

1. Enhancement of the quality of separation: starting from high threshold, GMCA proposes from the very first iterations an accurate estimation of the different mixing directions and the corresponding sources;

2. Robustness to noise: final thresholds are selected accordingly to the noise contaminating the sources, enabling an efficient denoising and a refinement of the sources estimates.

2.5.3 Exploiting GMCA heuristics in other approaches

GMCA algorithm proposes an easy-to-implement decreasing strategy that is enabled by the sources update of pALS algorithm. Nevertheless, it does not exactly minimize the problem (2.10) and provides no mathematical guarantee of convergence on the opposite of PALM and BCD algorithms.

This section focuses on the attempts that have been made by Chenot [2017] and Kervazo [2019] to apply GMCA adaptive thresholding strategy to BCD and PALM algorithms, the objective being to gather the robustness to noise and insensitivity to initialization provided by GMCA and the mathematical guarantees of BCD and PALM algorithms.

GMCA heuristics for PALM algorithm

When employed to minimize the cost function (2.10), PALM algorithm is generally used with fixed regularization parameters during the whole unmixing process. In addition to the absence in the literature of automatic regularization parameter choice, the author of Kervazo [2019] has shown that PALM is very sensitive to the choice of thresholds. In other words, it is of crucial importance to provide PALM algorithm appropriate regularization parameters but no automatic heuristic is proposed in the literature to the best of our knowledge.

The author of Kervazo [2019] has proposed to extend GMCA automatic thresholding to PALM algorithm. The idea was to compute the $\kappa - \mathrm{MAD}$ operator on the update of the sources

$$\forall i = 1..n \Lambda_i = \kappa \times L_{\mathbf{S}} \times \mathrm{MAD}(\mathbf{S}^{(k)} + \frac{1}{L_{\mathbf{S}}}\mathbf{A}^{(k)T}(\mathbf{X} - \mathbf{A}^{(k)}\mathbf{S}^{(k)})),$$

where $L_{\mathbf{S}} = \| \mathbf{A}^{(k)T}\mathbf{A}^{(k)} \|_{s}, 2.$
On the opposite of GMCA algorithm, the threshold is not computed on the least-square estimates of the sources but on a gradient descent step which, at convergence, boils down to the following:

$$\forall i = 1..n \ \Lambda_i = \kappa \times L_{\mathbf{S}} \times \mathrm{MAD}(\mathbf{A}^{*T}\mathbf{N} - \mathbf{A}^{*T}\mathbf{A}^{*}s),$$

where s stands for the errors on the sources (for more information on this result, the reader is invited to refer to Kervazo [2019]). The sources update expression presents a term of interferences mixed with a term of errors on the sources. Therefore, if the initial couple $(\mathbf{A}^{(0)}, \mathbf{S}^{(0)})$ is far from the true matrices, high thresholding generates more artefacts (*i.e.* higher s), leading to more interferences on the mixing matrix (*i.e.* due to the term $\mathbf{A}^{*T}\mathbf{A}^{*}$). The estimation of the mixing matrix is therefore deteriorated which hampers in turn the estimation of the sources at the next iteration and so on. **Unless starting from reliable initialization, adapting GMCA heuristics to PALM algorithm results in poor separation**.

The author of Kervazo [2019] has also presented preliminary results on the application of $\kappa-$MAD to BCD algorithm that show it suffers from the same difficulties as PALM algorithm.

2-step strategy

In order to cope with the difficult setting of regularization parameters in PALM and BCD algorithms, the authors of Chenot [2017] and Kervazo [2019] have proposed a hybrid strategy where GMCA algorithm is employed as a warm-up stage to PALM (or BCD).

This approach benefits from GMCA fair estimations of the couple $(\mathbf{A}, \mathbf{S})$ used as initializations for PALM (or BCD) algorithm. Therefore, the $\kappa-$MAD strategy can be employed in a second stage in PALM updates (or BCD updates) since the level of interferences is very likely to be less than the level of noise. Thus, the two-step strategy provides automatic thresholding strategies with better results than GMCA and mathematical guarantees of convergence.

The two-step (or two-stage) approach with BCD algorithm will be detailed in chapter 3 and chapter 5 in the context of BSS with Poisson noise and BSS in the presence of spectral variabilities.

Conclusion

We have presented in this chapter some popular proximal algorithms and minimization schemes adapted to multi-convex problems. We have focused on optimization strategies for sparsity-enforcing BSS and, more specifically, the automatic setting of regularization parameters introduced in GMCA algorithm (Bobin et al. [2007]) and that will be useful in the rest of this thesis.

CHAPTER 3

Blind Source Separation from Poisson measurements

In this chapter, we develop a sparsity-based Blind Source Separation to tackle the analysis of data following Poisson statistics. To a large extent, it reports upon the results published in Bobin et al. [2020].

Let us recall that the problem of sparse Blind Source Separation (sBSS) has been extensively studied when the noise is additive and Gaussian. This is however not the case when the measurements follow Poisson or shot noise statistics, which is customary with counting-based measurements. To that purpose, we introduce a sparse BSS algorithm coined pGMCA (poisson-Generalized Morphological Component Analysis) that specifically tackles the blind separation of sparse sources from measurements following Poisson statistics. The proposed algorithm builds upon Nesterov's smoothing technique to define a smooth approximation of sparse BSS, with a data fidelity term derived from the Poisson likelihood. This allows to design a block coordinate descent-based minimization procedure with a simple choice of the regularization parameter.

3.1 Context

According the standard instantaneous linear mixture , each observation is described as a linear combination of n elementary sources or components $\mathbf{S}_j$.
Let us recall that, for a single observation i at sample t, the linear mixture model leads to the following formulation:

$$\mathbf{X}_i[t] = \sum_{j=1}^{n} a_{ij}\mathbf{S}_j[t] + \mathbf{N}_i[t], \tag{3.1}$$

where a_{ij} stands for the mixing weight that quantifies the contribution of source $\mathbf{S}_j$ in the i-th observation. It is generally more conveniently recast into the following matrix form:

$$\mathbf{X} = \mathbf{AS} + \mathbf{N}, \tag{3.2}$$

where $\mathbf{X} \in \mathbb{R}^{m\times t}$ is the observation or data matrix, $\mathbf{S} \in \mathbb{R}^{n\times t}$ the source matrix, $\mathbf{A} \in \mathbb{R}^{m\times n}$ the mixing matrix and $\mathbf{N} \in \mathbb{R}^{m\times t}$ the noise contribution. In this context, BSS aims at recovering both the mixing matrix $\mathbf{A}$ and the sources $\mathbf{S}$ from the data $\mathbf{X}$ only. This is essentially an unsupervised matrix factorization. Being an ill-posed inverse problem, it calls for additional assumptions about the sources and/or the mixing matrix, such as statistical independence Comon and Jutten [2010], non-negativity Févotte and Dobigeon [2015] or sparsity Zibulevsky and Pearlmutter [2001], Bobin et al. [2006] to only name three very traditional approaches.

However, the above linear mixture model does not describe precisely the kind of data that originate from counting processes, which are customary in NMR spectroscopy Sun and J.Xin [2011], Cherni et al. [2020], γ-ray spectroscopy Xu et al. [2019] , X-ray imaging Badenes [2010], Picquenot et al. [2019a] or low-photon count optics to only cite a few. In these settings, it is essential to accurately account for the exact statistics of the measurements, which precisely follow a Poisson distribution. The data $\mathbf{X}$ are however only defined statistically from the *noiseless* mixtures $\mathbf{AS}$. More precisely, the probability for a given sample to take the value $\mathbf{X}_i[t]$ is then given by the Poisson law:

$$\mathcal{P}(\mathbf{X}_i[t] = |[\mathbf{AS}]_i[t]) = \frac{e^{-[\mathbf{AS}]_i[t]}\,[\mathbf{AS}]_i[t]^{\mathbf{X}_i[t]}}{\mathbf{X}_i[t]!}, \tag{3.3}$$

where $[\mathbf{AS}]_i[t]$ is the sample of the matrix $\mathbf{AS}$ located at the i-th row and t-th column. In this case, the observations do not rigorously follow the linear mixture model since the Poisson stochastic process does not preserve linearity: *only holds on average*. For a single stochastic realisation, fitting for the linear mixture model without accounting for the exact measurement stochastic distribution is very likely to yield biased estimates of the sources and the mixing matrix.

In the ICA and NMF frameworks

To tackle BSS from Poisson measurements, a straightforward approach consists in maximizing the likelihood of the mixture variables. In the case of Poisson statistics, this amounts to minimizing the Kullback-Leibler divergence between the data $\mathbf{X}$ and the mixture model $\mathbf{AS}$ with respect to $\mathbf{A}$ and $\mathbf{S}$. This has been investigated both in the scope of Independent Component Analysis (ICA - Mihoko and Eguchi [2002]) and Non-negative Matrix Factorization (NMF - Févotte and Dobigeon [2015], Févotte and Idier [2011]), where it generally refers to robust BSS. In this setting, the mixing matrix and the sources are estimated by minimizing the β-divergence D_β between the data $\mathbf{X}$ and the model $\mathbf{AS}$ under the assumption that both $\mathbf{A}$ and $\mathbf{S}$ are non-negative:

$$\min_{\mathbf{A}\geq 0, \mathbf{S}\geq 0} D_\beta\left(\mathbf{X}, \mathbf{AS}\right), \tag{3.4}$$

where D_β denotes the β-divergence. Multiplicative algorithms Cherni et al. [2020] are customarily used as they implicitly encode the non-negativity of the mixture parameters. For $\beta = 1$, the β-divergence is exactly the Kullback-Leibler divergence. In this case, the above problem is similar to a maximum likelihood estimate of $\mathbf{A}$ and $\mathbf{S}$ assuming that the data are distributed according to a Poisson distribution.

Imposing the sparsity of the sought-after sources in the framework of NMF has further been investigated with Euclidean distances Eggert and Korner [2004], Hoyer [2002] as well as β-divergence Le Roux et al. [2015] data fidelity terms. The main limitation of these algorithms is that sparsity is enforced in the sample domain, which is not suitable for a large variety of applications, where sparsity needs to be enforced in a transformed domain.

In the sparse BSS framework

In the scope of sparse BSS, the sources are assumed to admit a sparse distribution in some signal representation $\mathbf{\Phi}$. More precisely, each source $\{\mathbf{S}_j\}_{j=1,\cdots,n}$ can be described in some basis, waveform dictionary or signal representation $\mathbf{\Phi}$ as $\mathbf{S}_j = \alpha_j \mathbf{\Phi}$. Then, sparse BSS algorithms seek the mixing matrix and the sources so that the latter are mutually the sparsest. It has long been highlighted that sparse modeling allows to compressively encode the information content of the sources to be estimated. In this context, improving the contrast between these sources eventually leads to enhanced separation processes Zibulevsky [2003], Bronstein et al. [2005], Li et al. [2006], Bobin et al. [2007] as well as increased robustness with respect to noise Bobin et al. [2007]. However, in the case of Poisson noise, the noise variance is proportional to the square root of the Poisson mean, *i.e.* the product $\mathbf{AS}$. Since sparse BSS is mainly sensitive to the most salient features of the sources Bobin et al. [2007, 2015], the Poisson nature of the noise will largely impact the performances of standard sparse BSS methods.

To the best of our knowledge, the blind separation of sources that are sparse in a transformed domain has not been investigated when the data follow a Poisson distribution. The proposed approach first consists in performing the separation by putting together sparse regularization in a transformed domain and a data fidelity term derived from the Poisson likelihood. This raises two important challenges: i) most modern-day algorithms for sparse matrix factorization require the data fidelity term to be smooth Xu and Yin [2014], which is not the case for the Poisson likelihood about 0, ii) in the low statistics regime (*i.e.* when the Poisson mean takes small values), the Poisson likelihood, taken as a function of either the sources $\mathbf{S}$ or the mixing matrix $\mathbf{A}$, becomes highly ill-conditioned. This largely hampers the separation quality as well as the speed of convergence of gradient-based algorithms. In this chapter, these challenges are addressed by developing of a BCD-based algorithm (Block-coordinate descent - Tseng [2001]) along with a smooth approximation of the Poisson likelihood based on Nesterov's smoothing technique Nesterov [2005]. The proposed method coined Poisson Generalized Morphological Component Analysis (pGMCA) is detailed in Section 3.2. Numerical experiments are carried out on simulated and realistic X-ray astrophysical data in Section 3.3.

3.2 Sparse BSS from Poisson measurements

3.2.1 From additive Gaussian noise to Poisson statistics

In the next, we assume that each source admits a sparse representation in some signal representation $\mathbf{\Phi}$, which is described with the synthesis sparse signal model: $\forall j = 1, \cdots, n,\ \mathbf{S}_j = \alpha_j \mathbf{\Phi}$, where the expansion coefficients vector α_j is of size T. In the GMCA algorithm, the sparsity of the expansion coefficients α_j is enforced by minimizing its re-weighted ℓ_1 norm Candes et al. [2008] with $p \leq 1$:

$$\min_{\mathbf{A} \in OB(m), \mathbf{S}} \|\mathbf{\Lambda} \odot \alpha\|_{\ell_p} + \frac{1}{2} \|\mathbf{X} - \mathbf{A}\alpha\mathbf{\Phi}\|_F^2, \tag{3.5}$$

where the matrix α is of size $n \times T$; each of its rows j is defined by the expansion coefficients α_j of the j-th source $\mathbf{S}_j$. In practice, the choice $p = 1$ is adopted. The term $\mathbf{\Lambda}$ is a matrix of size $n \times T$ that contains the regularization parameters, which will be described in Section 3.2.4. The last term is the data fidelity term, which would be identical to the negative loglikelihood of the variables $\mathbf{A}$ and $\mathbf{S}$ for additive Gaussian noise. The oblique constraint alleviates the standard scale indeterminacy between the mixing matrix and the sources.
Switching from additive Gaussian to Poisson noise naturally requires substituting the quadratic Frobenius norm with the Poisson negative loglikelihood:

$$\min_{\mathbf{A} \in \mathcal{C}, \mathbf{S} \geq 0} \left\| \mathbf{\Lambda} \odot \mathbf{S}\mathbf{\Phi}^T \right\|_{\ell_1} + \mathcal{L}(\mathbf{X}|\mathbf{A}, \mathbf{S}), \tag{3.6}$$

where $\mathcal{L}(\mathbf{X}|\mathbf{A}, \mathbf{S})$ is defined as follows:

$$\begin{aligned} \mathcal{L}(\mathbf{X}|\mathbf{A}, \mathbf{S}) &= \sum_{i,t} [\mathbf{A}\mathbf{S}]_i[t] - \mathbf{X}_i[t] \log([\mathbf{A}\mathbf{S}]_i[t]) \qquad (3.7) \\ &= \mathbf{A}\mathbf{S} - \mathbf{X} \odot \log(\mathbf{A}\mathbf{S}) \qquad (3.8) \end{aligned}$$

The mixing matrix is further imposed to have non-negative entries. This entails that the mixing matrix belongs to $\mathcal{C}$, which is defined as the intersection between the oblique set and the non-negative orthant: $\mathcal{C} = OB(m) \cap K^+$. As well, the sources are constrained to have non-negative entries.

3.2.2 A smooth approximation of the data fidelity term

Finding a solution to the problem in eq.(3.6) raises major challenges. Firstly, this problem is not globally convex but convex with respect to each variable $\mathbf{A}$ and $\mathbf{S}$ assuming the other one is fixed. This class of optimization problems are dubbed multi-convex problems. In Bobin et al. [2007, 2015], the GMCA algorithm has been built as a projected ALS algorithm (Alternate Least-Squares), which has been showed to provide a reliable (*i.e.* robust with respect to initialization) and effective proxy to approximate solution of the problem in eq.(3.5). This choice was in part motivated by its low computational cost. However, when the non-negativity of the sources and the mixing matrices needs to be enforced, it has been emphasized in Rapin et al. [2014] that an ALS-based approach like GMCA is not appropriate. Furthermore, a similar minimisation procedure is well adapted when the data fidelity is quadratic, but cannot be applied to the problem in eq.(3.6). Several alternatives have been proposed to tackle multi-convex problems such as the Block-Coordinate-Descent (BCD - Tseng [2001], Xu and Yin [2014]) or the Proximal Alternating Linear Minimization (PALM - Bolte et al. [2014], Xu and Yin [2014]). However, the Poisson negative loglikelihood described in eq.(3.7) is not differentiable about 0. This excludes using methods based on proximal gradient descent such as the PALM.

The second major issue raised by sparsity-enforcing regularization is the choice of the regularization parameter $\mathbf{\Lambda}$, which has a significant impact on the separation quality Kervazo et al. [2020]. Fine-tuning the regularization parameters is quite challenging in general as it depends on the unkwnown sources to be estimated. As an alternative, we have proposed a automatic procedure to fix the values of $\mathbf{\Lambda}$ Rapin [2014], Kervazo et al. [2020] but it only applies to methods that build upon proximal gradient descent.

To alleviate these bottlenecks, we propose to rather substitute the Poisson negative loglikelihodd in eq.(3.6) with a smooth approximation, which allow to use standard methods for multi-convex optimization as well and extending automatic regularization parameter tuning Rapin et al. [2014], Kervazo et al. [2020] to the present setting. For that purpose, a straightforward way would amount to replacet $\mathcal{L}\left(\mathbf{X}|\mathbf{A},\mathbf{S}\right)$ in eq.(3.6) with a quadratic approximation of the form:

$$\mathcal{L}\left(\mathbf{X}|\mathbf{A},\mathbf{S}\right) \simeq \|(\mathbf{X}-\mathbf{AS}) \oslash \sqrt{\mathbf{X}+\epsilon^2}\|_F, \tag{3.9}$$

where the scalar ϵ would avoid the singularities of the Poisson negative log-likelihood about 0. However, this would yield a rather rough and inaccurate approximation, especially at low statistics, when the Poisson average $\mathbf{AS}$ takes small values. Building a smooth approximation of $\mathcal{L}\left(\mathbf{X}|\mathbf{A},\mathbf{S}\right)$ with a well-controled accuracy can be done by implementing the smoothing technique introduced by Nesterov in Nesterov [2005].

Following Nesterov's approach, a smooth approximate $\mathcal{L}_\mu$ of the convex function $\mathcal{L}$ can be built from its Fenchel dual $\mathcal{L}^\star$ Rockafellar [1970] as follows:

$$\mathcal{L}_\mu(\mathbf{X}|\mathbf{Y}) = \text{inf}_{\mathbf{U}} \, \langle \mathbf{Y},\mathbf{U}\rangle - \mathcal{L}^\star(\mathbf{X}|\mathbf{U}) - \mu g(\mathbf{U}), \tag{3.10}$$

where $\langle \mathbf{Y},\mathbf{U}\rangle$ stands for the scalar product between the matrices $\mathbf{Y}$ and $\mathbf{U}$. The smoothing parameter μ is a strictly positive real-valued scalar. The function g is strongly convex; in practice it is simply chosen as $g(\mathbf{U}) = \|\mathbf{U}\|_F^2$. It can be shown that the dual function of the Poisson neg-loglikelihood takes the following form:

$$\mathcal{L}^\star(\mathbf{X}|\mathbf{U}) = \mathbf{X}\odot(1-\log(\mathbf{X})) - \mathbf{X}\odot(1-\log(\mathbf{U}))$$

Following the definition in Equation (3.10), $\mathcal{L}_\mu$ is differentiable for any non-negative entry and admits a Lipschitz gradient $\nabla\mathcal{L}_\mu$.

Expression of the gradient According to Nesterov's smoothing technique, computing a smooth approximation of the negative likelihood $\mathcal{L}(\mathbf{X}|\mathbf{Y})$ can be built by altering its Fenchel dual. The negative log-likelihood of the Poisson distribution is separable according to each sample:

$$\mathcal{L}(\mathbf{X}|\mathbf{Y}) = \sum_{i,t} \mathcal{L}(\mathbf{X}_i[t]|\mathbf{Y}_i[t])$$

Consequently, we hereafter consider deriving a smooth approximation $\mathcal{L}(x|y) = y + x\log(y)$ on a single sample $x = \mathbf{X}_i[t]$. It's Fenchel dual is defined as:

$$L^\star(u) = \max_y uy - y + x\log(y),$$

which takes the following form:

$$L^\star(u) = x(\log(x)-1) - x\log(1-u)$$

A smooth approximation of the above negative log-loglikelihood can be built by adding a strongly convex term to the Fenchel dual. A straightforward approach consists in adding a quadratic term as follows:

$$\mathcal{L}_\mu(x|y) = \max_u uy - L^\star(u) - \frac{\mu}{2}\|u\|_2^2 \tag{3.11}$$

It has been showed in Nesterov [2005] that such a function is differentiable and its gradient is given by:

$$\nabla_y \mathcal{L}_\mu(x|y) = \text{argmax}_u uy - L^\star(u) - \frac{\mu}{2}\|u\|_2^2 \tag{3.12}$$

Simple algebra allows to define the solution $\nabla_y \mathcal{L}_\mu(x|y)$ as being one of the roots u of:

$$u^2 - \frac{y}{\mu}u + \frac{1}{\mu}(y-x) - 1 = 0$$

The only root that is lower than 1 is given by

$$\nabla_y \mathcal{L}_\mu(x|y) = \frac{y+\mu}{2\mu}\left[1 - \sqrt{1 - 4\mu\frac{y-x}{(y+\mu)^2}}\right]$$

Going back to the matrix formulation leads to:

$$\begin{aligned} \nabla \mathcal{L}_\mu(\mathbf{X}|\mathbf{Y}) \quad &= \quad \frac{1}{2\mu}(\mathbf{Y}+\mu) \odot \cdots & (3.13)\\ &\quad\; \cdots\left[1 - \sqrt{1 - 4\mu(\mathbf{Y}-\mathbf{X}) \oslash (\mathbf{Y}+\mu)^2}\right] & (3.14) \end{aligned}$$

The square root and power of two are also taken entry-wise. The choice of the smoothing parameter μ will be discussed in Section 3.2.4.

3.2.3 pGMCA: a BCD-based algorithm

The general aproximate optimisation problem to be solved is the following one:

$$\min_{\mathbf{A}\in\mathcal{C}, \alpha\mathbf{\Phi}\geq 0} \|\mathbf{\Lambda} \odot \alpha\|_{\ell_1} + \mathcal{L}_\mu\left(\mathbf{X}|\mathbf{A}, \alpha\mathbf{\Phi}\right), \tag{3.15}$$

where $\mathbf{S} = \alpha\mathbf{\Phi}$.
To minimize this problem, either the BCD or the PALM algorithms can be implemented. Following Rapin et al. [2014], the BCD allows to limit potential propagation errors by iteratively and sequentially updating $\mathbf{A}$ and $\mathbf{S}$ with a complete minimisation. The resulting pGMCA algorithm (poisson Generalized Morphological Component Analysis) is described in

Description of the update of $\mathbf{S}$

For a fixed current estimate $\mathbf{A}^{(k-1)}$ of $\mathbf{A}$, the sources are updated by solving the problem:

$$\min_\alpha \|\mathbf{\Lambda} \odot \alpha\|_{\ell_1} + \iota_{K^+}\left(\alpha\mathbf{\Phi}\right) + \mathcal{L}_\mu\left(\mathbf{X}|\mathbf{A}, \alpha\mathbf{\Phi}\right), \tag{3.16}$$

This convex problem is composed of a smooth data fidelity term and two non-smooth regularization functional, which makes the framework of proximal algorithms Parikh et al. [2014] perfectly adapted to build a minimizer. In eq. (3.16), the function regularization parameters have well-defined proximal operator but it does not admit a closed-form expression. This precludes the use of standard proximal algorithms such as the Forward-Backward Splitting algorithm (FBS - Combettes and Pesquet [2011]). In this context, the Generalized Forward-Backward Splitting algorithm (GFB -Raguet et al. [2013]) is way more adapted since it allows to call independently the proximal operators of individual regularization terms. Details of the GFB implementation are given below:

Algorithm 10: GFB implementation

```
Initialization of $\mathbf{A}$ and $\mathbf{S}$ using the GMCA algorithm.
while $l < L_{outer}$ and $\|\mathbf{A}^{(l+1)} - \mathbf{A}^{(l)}\|_F / \|\mathbf{A}^{(l)}\|_F > \epsilon$ do
    • Update $\mathbf{S}$ following the procedure described in Section 3.2.3:
    $\mathbf{S}^{(k+1)} = \text{Argmin}_{\alpha} \mathbf{\Phi} \geq 0 \, \|\mathbf{\Lambda} \odot \alpha\|_{\ell_1} + \mathcal{L}_{\mu}\left(\mathbf{X} | \mathbf{A}^{(k)}, \alpha \mathbf{\Phi}\right)$
    • Update $\mathbf{A}$ following the procedure described in Section 3.2.3:
    $\mathbf{A}^{(k+1)} = \text{Argmin}_{\mathbf{A} \in \mathcal{C}} \mathcal{L}_{\mu}\left(\mathbf{X} | \mathbf{A}, \mathbf{S}^{(k+1)}\right)$
end
```

- Gradient expression: the gradient of the data fidelity term with respect to the sources is given by: $\mathbf{g_S} = \mathbf{A}^T \nabla \mathcal{L}_{\mu}(\mathbf{X}, \mathbf{AS})$. It is Lipschitz with constant $L_{\mathbf{S}} = \|\mathbf{A}\mathbf{A}^T\|_2 / \mu$.

- Sparse regularization: The proximal operator of the sparse regularization $\|\mathbf{\Lambda} \odot (\,.\,)\|_{\ell_1}$ is well-known to be the soft-thresholding operator that applies independently on each entry of its input argument $\mathbf{Y}_i[t]$: $\text{prox}_{\mathbf{\Lambda} \odot (\,.\,)\|_{\ell_1}}(\mathbf{Y}_i[t]) =$
$$\begin{cases} \mathbf{Y}_i[t] - \mathbf{\Lambda}_i[t] \, \text{sign}(\mathbf{Y}_i[t]) & \text{if } |\mathbf{Y}_i[t]| > \mathbf{\Lambda}_i[t] \\ 0 & \text{otherwise} \end{cases}$$

- *Positivity constraint:* Depending on the properties of the signal representation $\mathbf{\Phi}$, the proximal operator of the sparse regularization $\iota_{(\,.\,)\mathbf{\Phi} \geq 0}(.)$ can take different expressions.
  When $\mathbf{\Phi}$ is orthogonal (*e.g.* Discrete Cosine Transform, Biorthogonal Wavelets, etc.), it is straightforwardly defined as:
  $$\text{prox}_{\iota_{(\,.\,)\mathbf{\Phi} \geq 0}}(\mathbf{Y}) = [\mathbf{Y}\mathbf{\Phi}]_+ \mathbf{\Phi}^T$$
  If $\mathbf{\Phi}$ is a tight frame, for which $\mathbf{\Phi}^T \mathbf{\Phi} = \mathbf{I}$ (*e.g.* Undecimated wavelets Starck et al. [2007], Curvelets Starck et al. [2002], etc.), it can be shown to take the following closed-form formulation Rapin et al. [2014]:
  $$\text{prox}_{\iota_{(\,.\,)\mathbf{\Phi} \geq 0}}(\mathbf{Y}) = \mathbf{Y} + [-\mathbf{Y}\mathbf{\Phi}]_+ \mathbf{\Phi}^T$$

The algorithm is described in 11.

The hyperparameters of the GFB include weights $\{\rho_i\}_{i=1,2}$ that impose a relative weight between the two intermediate variables $\mathbf{U}_{i=1,2}$, and therefore the related penalisations. In the next, these weights have both been fixed to $1/2$; other choices did not lead to significant differences of the speed of convergence. The stepsize γ is fixed to $\gamma = \frac{\mu}{\|\mathbf{A}^T\mathbf{A}\|_2}$.

Description of the update of $\mathbf{A}$

For a fixed current estimate $\mathbf{S}^{(k)}$ of $\mathbf{S}$, the mixing matrix is updated by solving the following problem:

$$\min_{\mathbf{A} \in \mathcal{C}} \mathcal{L}_{\mu}(\mathbf{X} | \mathbf{A}, \mathbf{S}) \,. \tag{3.17}$$

Since the set $\mathcal{C}$ is the intersection of the oblique set OB(m) and the non-negative orthant, the above problem can be rewritten by introducing the characteristic function ι of these sets:

$$\min_{\mathbf{A}} \iota_{(\,.\,) \in \mathcal{C}}(\mathbf{A}) + \mathcal{L}_{\mu}(\mathbf{X} | \mathbf{A}, \mathbf{S}) \,. \tag{3.18}$$

Algorithm 11: GFB implementation for the update of $\mathbf{S}$

```
Fix the step size $\gamma = \frac{2\mu}{\|\mathbf{A}^{(k)}\|_2^2}$, $\{\rho_i\}_{i=1,\cdots,2} = 0.5$, $\{\mathbf{U}_i\}_{i=1,\cdots,2} = \mathbf{S}^{(k)}$
while $l < L_{inner}$ and $\|\mathbf{S}^{(l+1)} - \mathbf{S}^{(l)}\|_F / \|\mathbf{S}^{(l)}\|_F > \epsilon$ do

end
• Gradient of the data fidelity term: $\mathbf{g}_{\mathbf{S}} = {\mathbf{A}^{(k)}}^T \nabla \mathcal{L}_\mu \left(\mathbf{X}, \mathbf{A}^{(k)}\mathbf{S}^{(l)}\right)$
• Updating the sparsity-enforcing intermediate variable:
$\mathbf{g}U_1^{(l+1)} = \mathbf{g}U_1^{(l)} + \text{prox}_{\gamma \mathbf{\Lambda} \odot (\,.\,)\|_{\ell_1}} \left(2\mathbf{g}S^{(l)} - \mathbf{g}U_1^{(l)} - \gamma \mathbf{g}_{\mathbf{S}}\right) - \mathbf{g}S^{(l)}$
Updating the positivity-enforcing intermediate variable:
 $\mathbf{g}U_2^{(l+1)} = \mathbf{g}U_2^{(l)} + \text{prox}_{i_{(\,.\,)\mathbf{\Phi} \geq 0}} \left(2\mathbf{g}S^{(l)} - \mathbf{g}U_2^{(l)} - \gamma \mathbf{g}_{\mathbf{S}}\right) - \mathbf{g}S^{(l)}$
• Update of $\mathbf{S}$: $\mathbf{g}S^{(l+1)} = \sum_{r=1}^{2} \mu_i \mathbf{g}U_r^{(t+1)}$
```

This problem can be solved using the Forward-Backward splitting algorithm Combettes and Pesquet [2011], or more precisely one of its accelerated variant FISTA Beck and Teboulle [2009]. Details of the implementation of the FISTA implementation are given below:

- Gradient expression: the gradient of the data fidelity term with respect to the sources is given by: $\mathbf{g}_{\mathbf{A}} = \nabla \mathcal{L}_\mu (\mathbf{X}, \mathbf{AS})\, \mathbf{S}^T$. It is Lipschitz with constant $L_{\mathbf{A}} = \|\mathbf{S}^T\mathbf{S}\|_2/\mu$.

- Constraints: The mixing matrix is constrained to belong to the set $\mathcal{C}$, which is the intersection of the oblique set and the non-negative orthant.
  Strictly speaking, the oblique set is not convex and does not rigorously admit a proximal operator. However, if none of the columns of some matrix $\mathbf{Y}$ vanishes, its orthogonal projection onto $OB(m)$ is uniquely defined. The *proximal* operator of its characteristic function $\iota_{(\,.\,)\in OB(m)}$ will be defined as: $\text{prox}_{(\,.\,)\in OB(m)}(\mathbf{Y}) = [\mathbf{Y}^k / \|\mathbf{Y}^k\|_{\ell_2}]$. The proximal operator of the positivity constraint boils down to the orthogonal projection onto the positive orthant can be shown to take the following closed-form formulation Rapin et al. [2014]: $\text{prox}_{(\,.\,)\in K^+}(\mathbf{Y}) = [\mathbf{Y}]_+$.
  From the above definitions, it can be showed that the *proximal* operator of its characteristic function $\iota_{(\,.\,)\in\mathcal{C}}$ can be defined as the composition of $\text{prox}_{(\,.\,)\in K^+}$ and $\text{prox}_{(\,.\,)\in OB(m)}$:

$$\text{prox}_{(\,.\,)\in\mathcal{C}} = \text{prox}_{(\,.\,)\in OB(m)} \circ \text{prox}_{(\,.\,)\in K^+}$$

The update of the mixing matrix $\mathbf{A}$ is summarized as follows:
The iterations stop whenever the relative variation between consecutive iterations of $\mathbf{A}$ is below $\epsilon = 10^{-6}$.

3.2.4 Implementation details

Initialization: The problem in eq. (3.15) being multi-convex, the BCD algorithm turns to be highly sensitive to the initial point. In the meantime the GMCA algorithm has long been advocated as an algorithm that is robust with respect to the initialization Kervazo [2019]. Therefore, the pGMCA algorithm will use the output of GMCA as first guess estimate for the mixing matrix and the sources.

Regularization and smoothing parameters: In practice, the BCD turns out to be quite sensitive to the choice of the regularization parameters. In the current setting, the role played by the matrix $\mathbf{\Lambda}$ is twofold: i) it implements a re-weighted scheme by imposing a relative weighting of the entries of $\mathbf{S}$ in the signal representation

Algorithm 12: FBS implementation for the update of $\mathbf{A}$

```
Fix the step size $\gamma = \frac{2\mu}{\|\mathbf{S}^{(k+1)}\|_2^2}$
while $l < L_{inner}$ and $\|\mathbf{A}^{(l+1)} - \mathbf{A}^{(l)}\|_F / \|\mathbf{A}^{(l)}\|_F > \epsilon$ do
    • Gradient of the data fidelity term:
        $\mathbf{g_A} = \nabla \mathcal{L}_\mu \left( \mathbf{X}, \mathbf{A}^{(l)} \mathbf{S}^{(k+1)} \right) \mathbf{S}^{(k+1)^T}$
    • Proximal step:
        $\mathbf{A}^{(k)} = \text{prox}_{(.)\in\mathcal{C}} \left( \mathbf{A}^{(k-1)} - \gamma \mathbf{g_A} \right)$
end
```

$\mathbf{\Phi}$ (see Candes et al. [2008]), and ii) it has also to be set so as to correctly reject non-sparse contributions, which are more likely to be noise dominated. Consequently, the matrix $\mathbf{\Lambda}$ is designed based on the first guess $\bar{\alpha}$ (*e.g.* the GMCA estimate of the sources in the signal representation) so that for each entry (i, j):

$$\forall i, j;\ \mathbf{\Lambda}_i[j] = \lambda_i w_i[j] \tag{3.19}$$

$$= \lambda_i \frac{\nu}{\nu + \frac{|\bar{\alpha}_i[j]|}{\max_j |\bar{\alpha}_i[j]|}}. \tag{3.20}$$

This formulation allows to split the regularization parameter into two terms: i) the re-weighting contribution $w_i[j]$ that takes its values in the range $[0, 1]$ and $\nu = 1e-3$, and ii) the noise-related part λ_i that is only source-dependent. Hence, the re-weighting term is defined as a modulation with respect to a standard noise-based thresholding Kervazo et al. [2020].

It is important to point out that $\mathbf{\Lambda}$ eventually plays the role of a threshold in the signal representation. Enforcing the sparsity of the solution therefore depends on its ability to throw away non-sparse contributions in the GFB update of $\mathbf{S}$. Since thresholding applies on a gradient descent update of the sources as displayed in , the thresholds $\{\lambda_i\}_{i=1,n}$ are chosen based on the gradient about the initial point:

$$\lambda_i = \tau.\text{MAD} \left(\nabla_{\mathbf{S}} \mathcal{L}_\mu (\mathbf{A}^{(0)}, \mathbf{S}^{(0)}) \mathbf{\Phi}^T \right), \tag{3.21}$$

where MAD stands for the Median Absolute Deviation. This approach is standard when the noise is additive and Gaussian. Indeed, in this case, it yields an empirical estimate of the standard deviation of the measured on the gradient in the sparse domain $\mathbf{\Phi}$. The paramter τ is generally chosen between 1 and 3 Kervazo et al. [2020].

However, it has to be noticed that, while the statistics of the observations are distributed according to a Poisson distribution, the distribution of the gradient according to $\mathbf{S}$ described in eq.(3.21) is way much harder to derive. In this context, the proximal step more generally aim at rejecting non-sparse contribution. A second element that needs to be accounted for is that the gradient with respect to $\mathbf{S}$ is asymmetric, which is likely to hamper the effectiveness of the above thresholding strategy. As shown in Figure 3.1, the asymmetry increases as the smoothing parameter μ decreases.

Finally, it has to be noted that the smoothing parameter μ has a direct impact on the quality of approximation of $\mathcal{L}$, and its gradient $\nabla\mathcal{L}$. When μ tends to 0, $\mathcal{L}_\mu$ converges towards $\mathcal{L}$. A small value for μ should therefore be favored to build a precise smooth approximation of $\nabla\mathcal{L}$.

Altogether, this suggests that there exists a trade-off between the precision of the approximation $\mathcal{L}_\mu$ and the effectiveness of the thresholding procedure. As an illustration Fig.3.2 features the evolution of the mean square error of the recovered sources $\mathbf{S}$ when the mixing matrix is fixed to the true value $\mathbf{A}^\star$. In this experiments, the

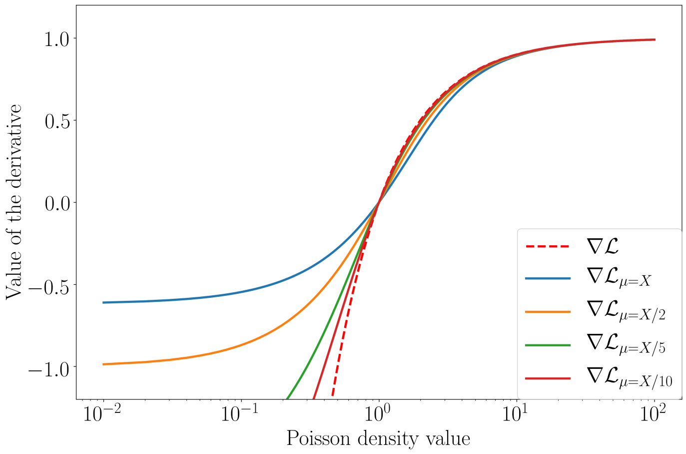

Figure 3.1: Approximations of $\nabla\mathcal{L}$ for different values of μ.

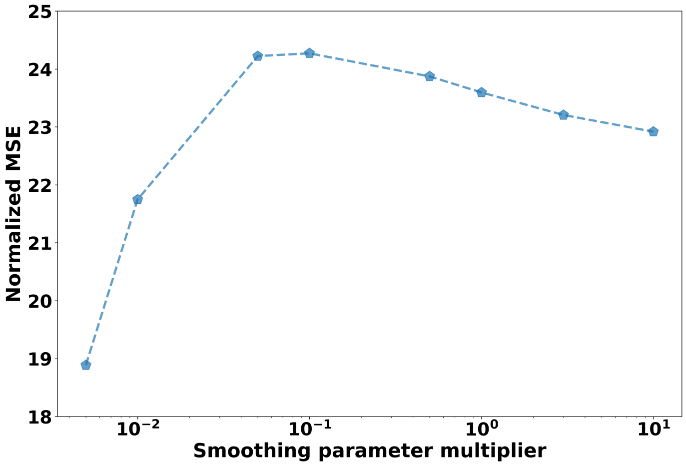

Figure 3.2: Normalized mean squared error as a function of smoothing parameter multiplier; the actual value of μ is defined as the value multiplied by the mean of the data $\langle \mathbf{X} \rangle$.

sources are generated according to the $1D$ model described in Section 3.3.2. Since Equation (3.14) reveals that μ mainly applies as an offset on $\mathbf{Y} = \mathbf{AS}$ in the expression of the gradient of $\mathcal{L}_\mu$, μ has to be directly compared to the value of the Poisson average of the mixture $\mathbf{AS}$ (or average number of counts). In the next paragraph, we discuss how it can be chosen along with the regularization parameters.

Indeed, the distribution of the noise contribution in the sparse representation will highly depend on the approximation of the gradient of $\mathcal{L}_\mu$ with respect to $\mathbf{S}$. For small values of μ – typically smaller than the mean value of the data $\bar{\mathbf{X}}$– Fig.3.1 shows that the gradient will be highly asymmetric about 0. It will tend to be more symmetric for $\mu \simeq \bar{\mathbf{X}}$. Furthermore, Fig.3.3 displays the histogram of the gradient term in the signal representation $A^{*T}\nabla\mathcal{L}_\mu\mathbf{\Phi}^T$ (which is computed about the true mixing matrices and sources) for two values $\mu = \bar{\mathbf{X}}$ and $\mu = \bar{\mathbf{X}}/10$ along with their Gaussian best fit. This figure emphasizes that a more symmetric smooth approximation tends to have a more Gaussian distribution in the transformed domain, which makes the proposed MAD-based strategy efficient in this regime. Consequently, $\mu = \bar{\mathbf{X}}$ is set for the next experiments.

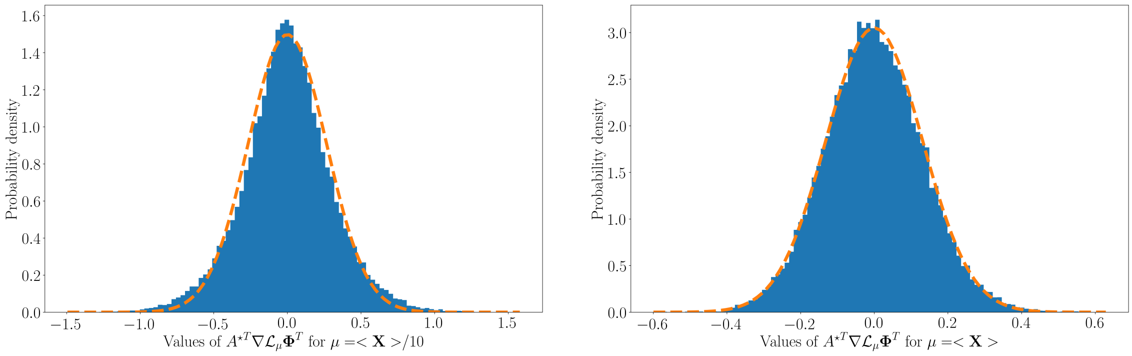

Figure 3.3: Histogram of the gradient of $\mathcal{L}_\mu$ with respect to $\mathbf{S}$ about the true input (mixing matrix and sources) and their Gaussian b best fit.

Smoothing parameter: The smoothing parameter μ has a direct impact on the quality of approximation of $\mathcal{L}$, and its gradient $\nabla\mathcal{L}$. Obviously, when μ tends to 0, $\mathcal{L}_\mu$ converges towards $\mathcal{L}$. More specifically, Equation (3.14) reveals that μ mainly applies as an offset on $\mathbf{Y}$ in the expression of the gradient of $\mathcal{L}_\mu$. Therefore μ has to be directly compared to the value of the Poisson average of the mixture $\mathbf{AS}$ (or average number of counts). In the next paragraph, we discuss how it can be chosen along with the regularization parameters.

Convergence: According to Tseng [2001], Xu and Yin [2013], the convergence of the BCD algorithm to a local critical point of the problem in Equation (3.15) is guaranteed. However, the speed of convergence is directly connected to the choice of the smoothing parameter μ. As detailed above, either it is for the update of the sources or the mixign matrix, the gradient steps are proportional to μ. This entails that a small value for μ will lead to slower convergence. To further illustrate this point, Figure 3.4 shows the number of iterations required to reach convergence (*i.e.* have the convergence criterion fulfilled) as a function of μ. This experiment focus on the update of $\mathbf{S}$ with the mixing matrix fixed to the true value $\mathbf{A}^\star$. The sources are generated according to the unidimensional model described in Section 3.3.2. This figure clearly shows that the required number of iterations rapidly increases when μ is typically lower than the average number of counts in the data $\langle X \rangle$. It has to be noticed that the performance of a similar algorithm without smoothing the Poisson loglikelihood would virtually be obtained for μ tending to 0. In addition to the thresholding strategy discussed in the previous paragraph, using Nesterov's smoothing technique further allows to limit the computational cost of gradient descent based algorithms in the present setting. For these reasons, the smoothing parameter will be fixed to $\langle X \rangle$, which provides a good trade-off between approximation accuracy, separation performance and computational cost.
The number of outer iterations L_{outer} is fixed 100; the number of inner iterations L_{inner} (*i.e.* number of iterations to update $\mathbf{S}$ and $\mathbf{A}$ for each outer iteration) is fixed to 1000. The overall algorithm stops when the Frobenius norm between two consecutive estimates of the mixing matrix $\left\|\mathbf{A}^{(k+1)} - \mathbf{A}^{(k)}\right\|_F$ is lower than $\epsilon = 10^{-6}$.

3.3 Numerical experiments

3.3.1 Experimental set-up

In the following experiments, our prime intention is to investigate the performance of the poisson GMCA algorithm with respect to standard matrix factorization algorithms. To that end, we first make use of synthetic data that mimic uni-dimensional spectroscopic data, which allows to perform Monte-Carlo simulations to assess the robustness of the different methods. The data are described as synthesized as follows:

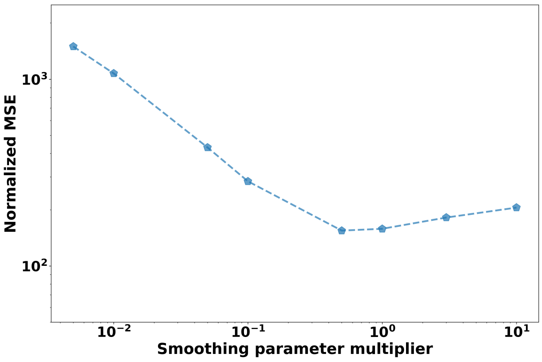

Figure 3.4: Number of iterations for convergence as a function of smoothing parameter multiplier; the actual value of μ is defined as the value multiplied by the mean of the data $\langle \mathbf{X} \rangle$.

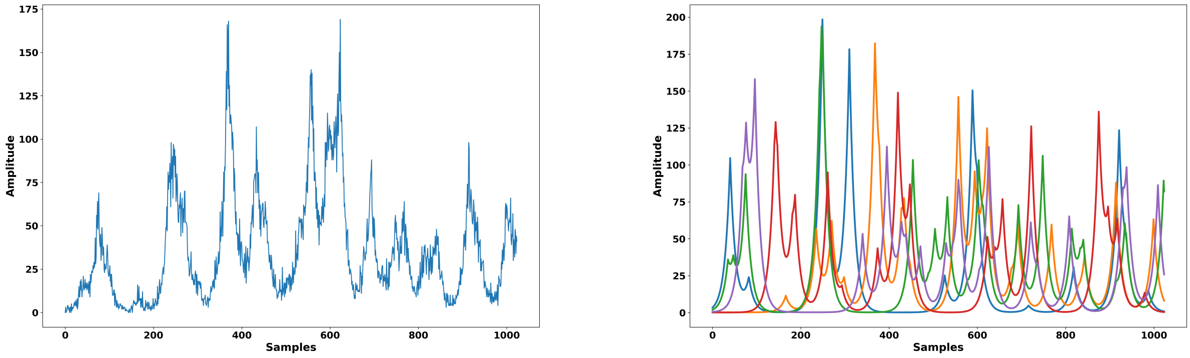

Figure 3.5: Examples of an observation (right panel) built from the input sources (left panel). The average number of counts (Poisson average) is equal to 250.

- The sources $\{s_j\}_{j=1,\cdots,n}$ are K-sparse signals; each source s_j has only K non-zero entries out of T entries. The activation of the entries of the sources will be distributed according to a Bernoulli process π with parameter $\rho = 0.02$ (*i.e.* the probability for a given sample to be non-zero is equal to ρ). Each source is made of $T = 1024$ samples.
- the sources are convolved with a Laplacian kernel of full width at half height (FWHM) equal to 8.
- The mixing matrix will be picked at random from a uniform distribution and processed so as to be non-negative with a pre-defined condition number.

The simulated sources (*resp.* an example of mixture) are displayed in the right (*resp.* right) panel of Figure 3.5. Such sources admit a sparse representation in a frame of undecimated (translation invariant) 1D wavelets Starck et al. [2004].

3.3.2 Methods' description

In addition to the proposed pGMCA algorithm, comparisons will be carried our with the following algorithms:

- **ICA based on a β divergence minimization Mihoko and Eguchi [2002]** [1]. This ICA algorithm estimates an unmixing matrix $\mathbf{B} \in \mathbb{R}^{n \times n}$ so that the estimated sources $\hat{\mathbf{S}} = \mathbf{B}\mathbf{X}$ are mutually independent.

---

[1] Based on a python implementation from Gadhok and Kinsner [2006]

Statistical independence is quantified with the β-divergence $\mathbb{D}_\beta$ between the product of their marginal $\prod_{i=1}^n p_{\mathbf{S}}(\hat{S}_i)$ and their joint distribution $p_{\mathbf{S}}(\hat{\mathbf{S}})$. For $\beta = 1$, this corresponds to the Kullback-Leibler divergence, which we show is perfectly adapted for Poisson measurements. In that case, this quantity, known as mutual information, vanishes if and only if the sources are independent.

- **The sparse NMF algorithm Leroux et al. [2015].** Leroux *et al.* proposed a fresh look at sparse NMF problems. Precisely, they introduce an improved multiplicative update algorithm that correctly impose both the non-negativity of the factors and the sparsity of the sources in the sample domain. The problem solved by the sNMF algorithm reads as: $\min_{\mathbf{A}\geq 0, \mathbf{S}\geq 0} \lambda \|S\|_{\ell_1} + D_\beta \left(\mathbf{X}, \mathbf{AS}\right)$. In the forthcoming numerical experiments, the simulated sources are mildly sparse in the sample domain; therefore it won't be a totally unreasonable assumption. The value of β is fixed to 1 to minimizer the Kullback-Leibler divergence. As well, the best sparsity parameter has been searched in the range $[0.01, 0.5] \times \|\mathbf{S}^\star\|_\infty$, where $\mathbf{S}^\star$ stands for the input sources.
- **The beta-sparse NMF algorithm Cherni et al. [2020].** Similarly to Leroux et al. [2015]. This method minimizes the β-divergence between the data and the model $\mathbf{AS}$ by further enforcing the non-negativity and the sparsity of the estimated sources in the sample domain. The last assumption is however not valid in the following experiments. Instead, it is imposed that the columns of the estimated mixing matrix have unit ℓ_2 norm. We found that setting the parameter β to 1 did not necessarilly yield the best solution; for that purpose, the best value has been optimized in the range $[1, 2.5]$ through a grid search. The sparsity parameter is chosen as described above.
- **The HALS algorithm Gillis and Glineur [2012]** The Hierarchical Alternate Least-Squares (HALS) algorithm is a NMF algorithm that minimizes a standard quadratic data fidelity term under positivity constraints. In constrast to other ALS-based methods, each source and its corresponding column of the mixing matrix are updated individually and sequentially, which leads to a simple and effective minimization scheme.
- **The BetaNMF algorithm Févotte and Dobigeon [2015].** Originating from the hyperspectral imaging community, this method minimizes the β divergence between the data and the model $\mathbf{AS}$ by further requiring the sources to be non-negative, and their samples lie in the simplex with almost pure pixels. The last assumption is however not valid in the following experiments. Instead, it is imposed that the columns of the estimated mixing matrix have unit ℓ_2 norm. The parameter β of the β-divergence is set to 1 to minimize a Kullback-Leibler divergence, which is well suited for Poisson statistics.

- **GMCA Bobin et al. [2007]**.This a standard sparse BSS algorithm. Similarly to the pGMCA algorithm, the signal representation $\mathbf{\Phi}$ is chosen as the isotropic undecimated wavelet transform.

Unless stated differently, each single experimental result will be given as the mean over 25 Monte-Carlo simulations with different mixing matrix, sources and noise realisations.

3.3.3 Comparison criteria

A standard approach to evaluate the performances of matrix factorization algorithms is based on the decomposition of the estimated sources into different error terms Vincent et al. [2006]:

$$s^{\text{est}} = s_{\text{target}} + s_{\text{interf}} + s_{\text{artefacts}},$$

which can be interpreted as:

- s_{target} is the projection of s^{est} on the true source, which corresponds to the sought-after contribution

- s_{interf} quantifies the leakage between sources.
- $s_{\text{artefacts}}$ represents the remaining artifacts, which do not originate from interferences.

In contrast to the standard decomposition, the noise-related term has been left out since in the present case it is data-dependent. Therefore, it cannot be identified as a distinct term. The Source to Distortion Ratio (SDR) is a combined quantity that gathers information from the different terms of the decomposition:

$$\text{SDR}(s^{\text{est}}) = 10\ \log_{10}\left(\frac{\|s_{\text{target}}\|_2^2}{\|s_{\text{interf}} + s_{\text{artefacts}}\|_2^2}\right). \tag{3.22}$$

Since these methods do no impose similar regularization on the sources, they are more fairly compared based on the quality of estimation of the mixing matrix. For that purpose , we make use of the mean and maximal spectral angular distance (SAD) between the estimated column $[\hat{\mathbf{A}}]^j$ and input column $[\mathbf{A}^o]^j$ of the mixing matrix:

$$\text{SAD} = \sum_{j=1}^{n} \cos^{-1}(|\langle[\hat{\mathbf{A}}]^j, [\mathbf{A}^\star]^j\rangle|)$$

Results

Figure 3.6 displays the results for a single source given by the GMCA, pGMCA and HALS for a Poisson average of 250. The mixing matrix has been generated at random with a condition number of 2. In this illustration, both the GMCA and pGMCA algorithms provide reasonable results. It can however be highlighted that the pGMCA algorithm yields a slightly less noisy estimate with better reconstructed peaks. This is particularly visible on peaks around 800th sample.

Evolution with respect to the mean number of counts δ In this experiment, the performances of the BSS methods are evaluated with respect to the Poisson average (*i.e.* mean number of counts) in range that goes from 5 to 5000. The condition number of the generated mixing matrix is set to 2. The left panel of Fig. 3.7 features the evolution of the mean SAD. The HALS and β-sparse NMF provide results that do not vary much across the range which we considered. In contrast, sparse regularization (sNMF , GMCA and pGMCA) clearly helps improving the separation quality for an average δ larger than 50. As testified by the results provided by GMCA and pGMCA, the ability to model sparsity in the wavelet domain provide an extra decrease of the SAD that can reach more than one order of magnitude for $\delta \geq 500$. The pGMCA yields significantly better results for $\delta > 100$, which highlights that the ability of the proposed algorithm to account for the Poisson nature of the noise. For low values of δ, a significant number of data samples vanish, which explains the poor separation quality of most methods that lead to a mean SAD of about 10°.
Similar conclusions can be drawn by measuring the SDR of the estimated sources as shown in the right panel of Fig. 3.7. For $\delta < 100$, all the methods yield a solution with a mean SDR that is below 15dB. For larger Poisson average, the pGMCA algorithm provides an improvement of about 2dB with respect to the GMCA algorithm.

Evolution with respect to the condition number In contrast to the standard additive Gaussian noise, the condition number of the mixing matrix plays an even more important role. Indeed, in the case of Poisson noise, the noise statistics is data dependent and its variance is equal to the noiseless mixture $\mathbf{AS}$. When the condition number of the mixing matrix increases, the noise variance will therefore be dominated by the data contribution spanned by the singular vectors of $\mathbf{A}$ that correspond to the most dominant singular values. This will considerably limit the ability to discriminate between the sources.
To evaluate this effect, the left panel of Fig.3.8 shows the SAD as a function of the condition number of the mixing matrix $\mathbf{A}$ in the range $[1.2, 9]$. It first shows that β-spars, HALS and sNMF provide decent results only when the mixing matrix is close to orthogonality (*i.e.* condition number lower than 2). Both GMCA and pGMCA seem to be more robust with respect to conditioning of the mixing matrix with a SAD below 1° when

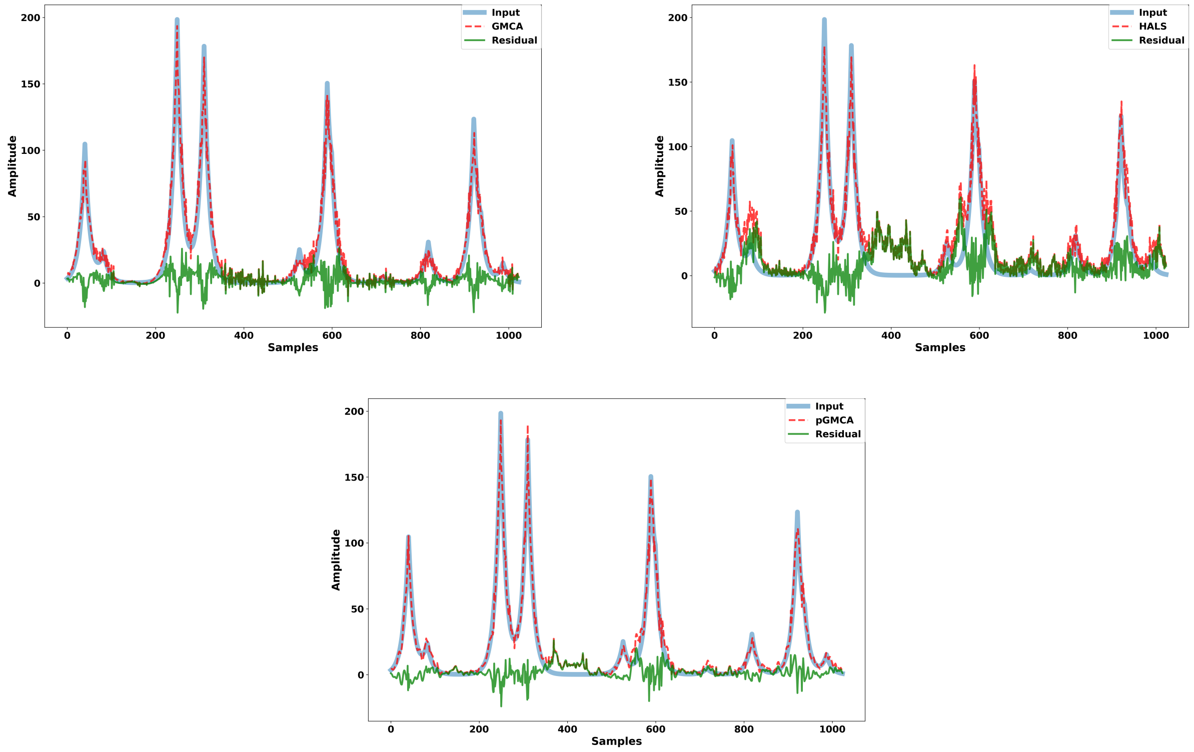

Figure 3.6: Illustration of separation results. *Solid blue:* true source, *dashed red:* estimated source and *solid green:* residual error

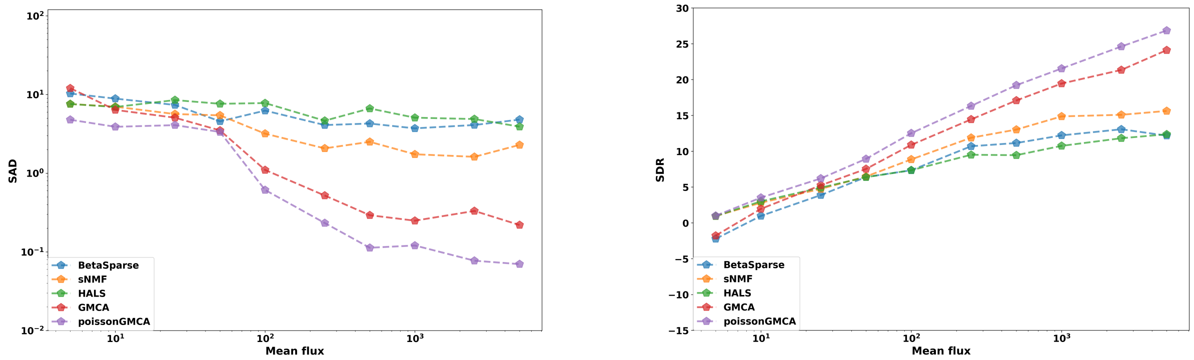

Figure 3.7: Evolution of the SAD (left panel) and SDR (right panel) as a function of the average of the Poisson distribution.

the condition number is lower than 5. All the methods yield poor results when the mixing matrix condition number is larger than 8. The right panel displays the evolution of the SDR of the estimated sources. This Fig. first confirms that when the condition number is larger than 2, the pGMCA algorithm leads to a gain of about 10dB with respect to NMF-based methods and 2dB with respect to the GMCA algorithm.

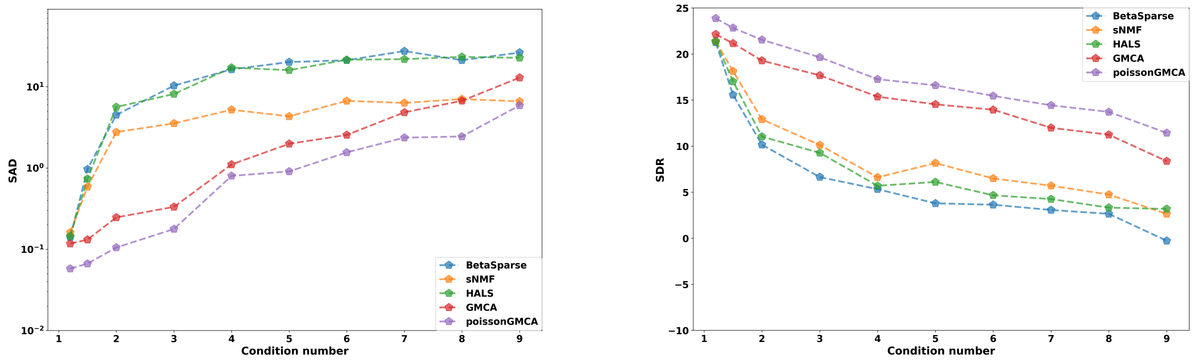

Figure 3.8: Evolution of the SAD (left panel) and SDR (right panel) as a function of the mixing matrix condition number.

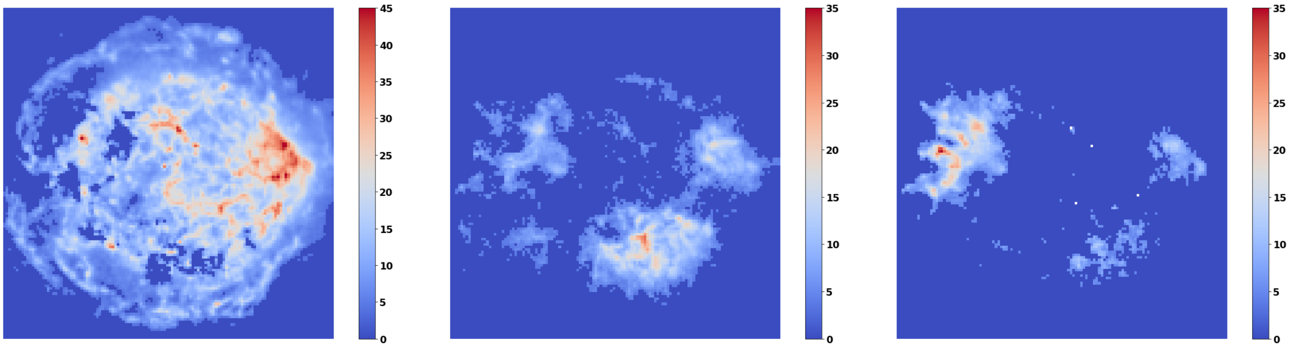

Figure 3.9: Input sources: synchrotron emission and the 2 iron emissions from top to bottom.

3.4 Application in Astrophysics

Description of the data

In this section, numerical experiments are carried out on simulations of astrophysical data that have been generated from real Chandra observations of the Cassiopea A supernova remnants. These data are composed of a linear combination of 3 astrophysical components: synchrotron emission, and 2 redshifted iron (Fe) emission lines. Mathematically, these emissions lines correspond to kronecker elements that are convolved with the telescope impulse response, which is Gaussian-shaped. These lines are centered about different energy values, which depend on the relative speed of propagation of each iron component as dictated by the Doppler effect. The synchrotron component has a power emission law. These components are representative of typical supernovae remnants in the energy band $5000-6000$ eV (electron-volt). In the next experiments, the data are composed of either 50 or 12 observations of size 128×128 pixels. The sources and their spectra are displayed in Fig. 3.9. Examples of mixtures are features in Fig.3.10.

3.4.1 Illustration

In the next experiments, the separation performances of the various methods will be evaluated based on the mean number of counts (*i.e.* average of the Poisson distribution) of the observed data $\mathbf{X}$, the number of observations m, and the relative intensity of the iron components with respect to the synchrotron emission. The data are generated so that the average of the Poisson measurements can be described with the following

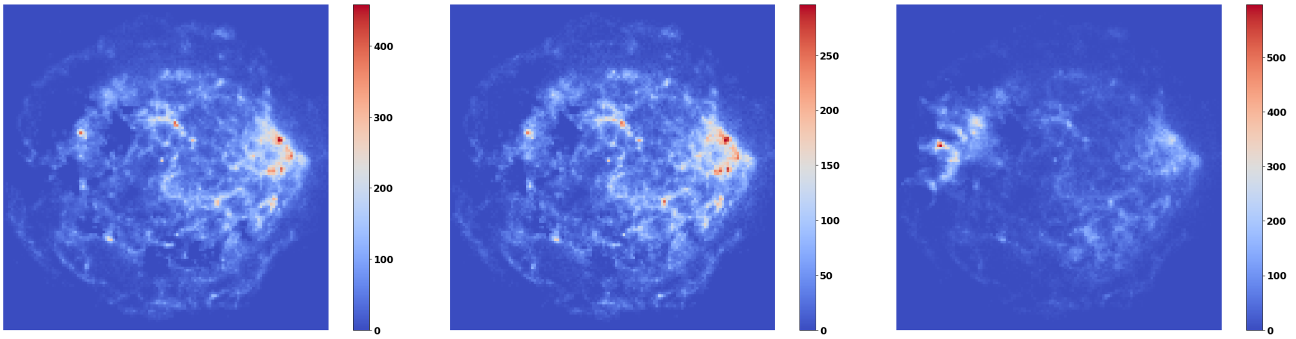

Figure 3.10: Examples of mixtures at 5279, 5548 and 5817eV.

formula:

$$\mathbf{X}^\star[i,t] = \frac{\rho\left(\mathbf{A}^\star_{\text{Fe}}\mathbf{S}^\star_{\text{Fe}} + \phi\mathbf{A}^\star_{\text{sync}}\mathbf{S}^\star_{\text{sync}}\right)}{mt^2\sum_{i,t}\left[\left[\mathbf{A}^\star_{\text{Fe}}\mathbf{S}^\star_{\text{Fe}} + \phi\mathbf{A}^\star_{\text{sync}}\mathbf{S}^\star_{\text{sync}}\right][i,t]\right]},$$

where the scalar ρ quantifies the mean number of counts per pixel in the observed data, and ϕ fixes the relative intensity between the synchrotron emission and the two Fe components. Studying the impact of these two parameters allows to consider various scenarios:

- Impact of the total intensity: making the total intensity ρ vary allows to investigate the performances of the BSS methods in the low number of counts regime, where standard methods are less likely to provide good results, as well as larger intensity regimes, where the standard additive Gaussian noise assumption can be a good approximation.

- Impact of the relative intensity: in the current astrophysical setting, the synchrotron emission is a rather diffuse component that mainly plays the role of a background component with respect to the Fe components. Letting their relative intensities to change allows to investigate different observational regimes. More specifically, when the synchrotron emission dominates, the noise statistics are prominently related to this component due to the Poisson nature of the observations, which makes the recovery of the Fe components much more challenging.

In this paragraph, comparisons between the three components are performed with $\phi = 5$ (*i.e.* the total intensity of the synchrotron emission is five times larger than the one of the other two components) and $\rho = 5$ (*i.e.* the Poisson average per pixel is 5). The three panels of Fig.3.9 feature the input and estimated spectra of the three components (*i.e.* the columns of the mixing matrix). The top panels of 3.11 (*resp.* 3.11 and 3.13) show the sources estimated with the GMCA (*resp.* pGMCA and HALS) algorithm. The bottom panels of these figures feature the residuals with respect to the true sources. To make visual inspection easier, the color scales are similar for each method. These results first show that these three methods visually provide reasonable estimated sources. The GMCA algorithm yield a significant higher residual for the three sources. This is particularly true for the synchrotron emission that exhibits a leakage from the first iron source. Similarly, the HALS algorithm shows a significant amount of synchrotron contamination in both the two redshifted iron sources. In contrast, the pGMCA algorithms leads to the lowest residues that visually correlates with the synchrotron emission. Since the synchrotron emission is dominant in this experiment, it is expected to be dominant as well in the noise. The noisy synchrotron imprint is therefore likely to be remainings of noise in the sources estimated with the pGMCA algorithm.

Figure 3.14 displays the estimated spectra for this experiment. This first confirms that the spectra of the two iron sources computed with the GMCA algorithm are not perfectly estimated; this is particularly the case for the first iron source, whose spectrum show a significant leakage of synchrotron at low energy. While the iron

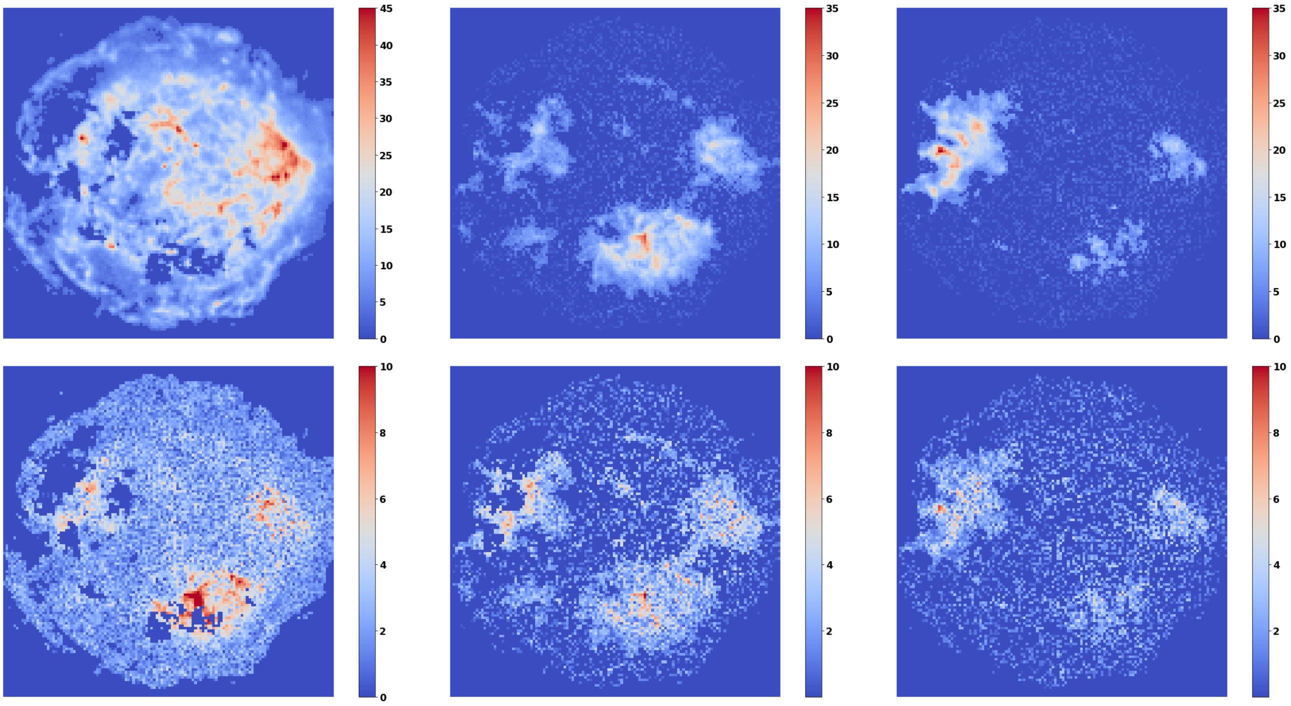

Figure 3.11: Sources estimated with GMCA.

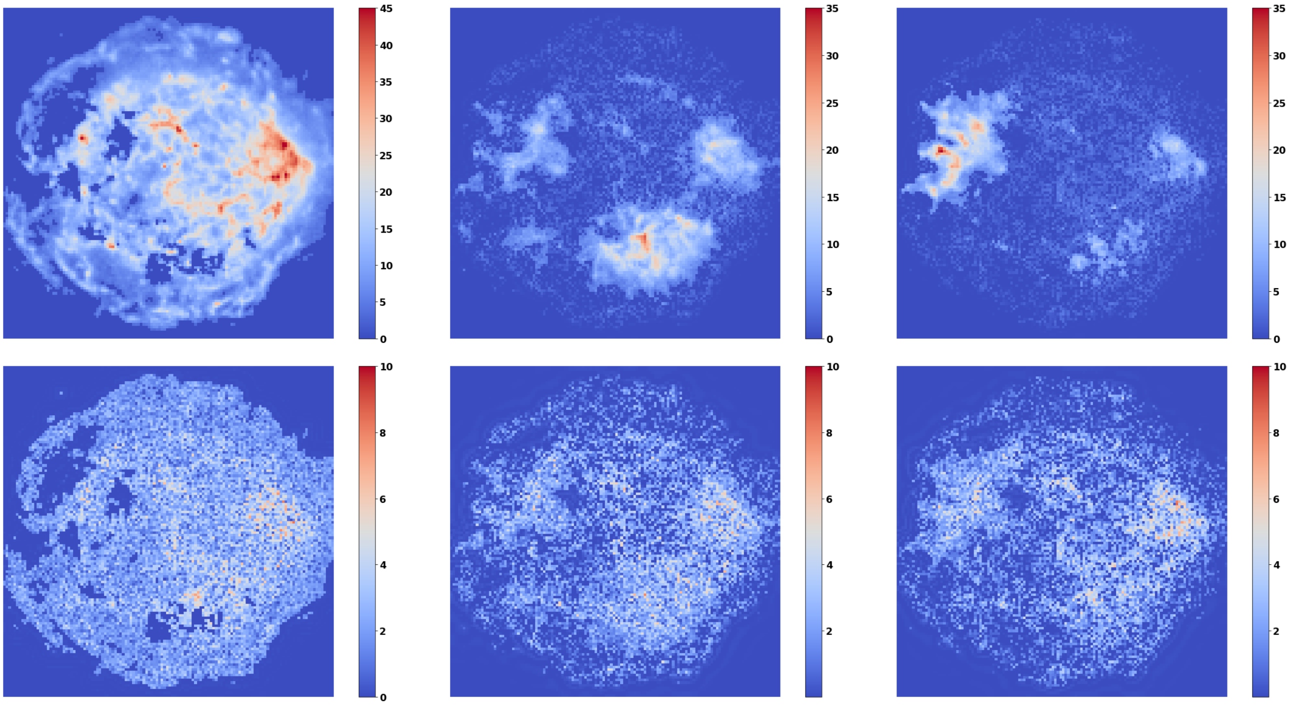

Figure 3.12: Sources estimated with pGMCA.

spectra estimated with the HALS seem to be very well estimated, it has to be highlighted that the synchrotron spectrum exhibits a significant bias in the range $[5600, 6000]eV$, where these sources peak. Keeping in mind that the synchrotron emission is dominant, this yields a much larger bias on the estimated sources as shown in Fig. 3.13. Finally, the pGMCA algorithm leads to the more robust estimates of the three spectra.

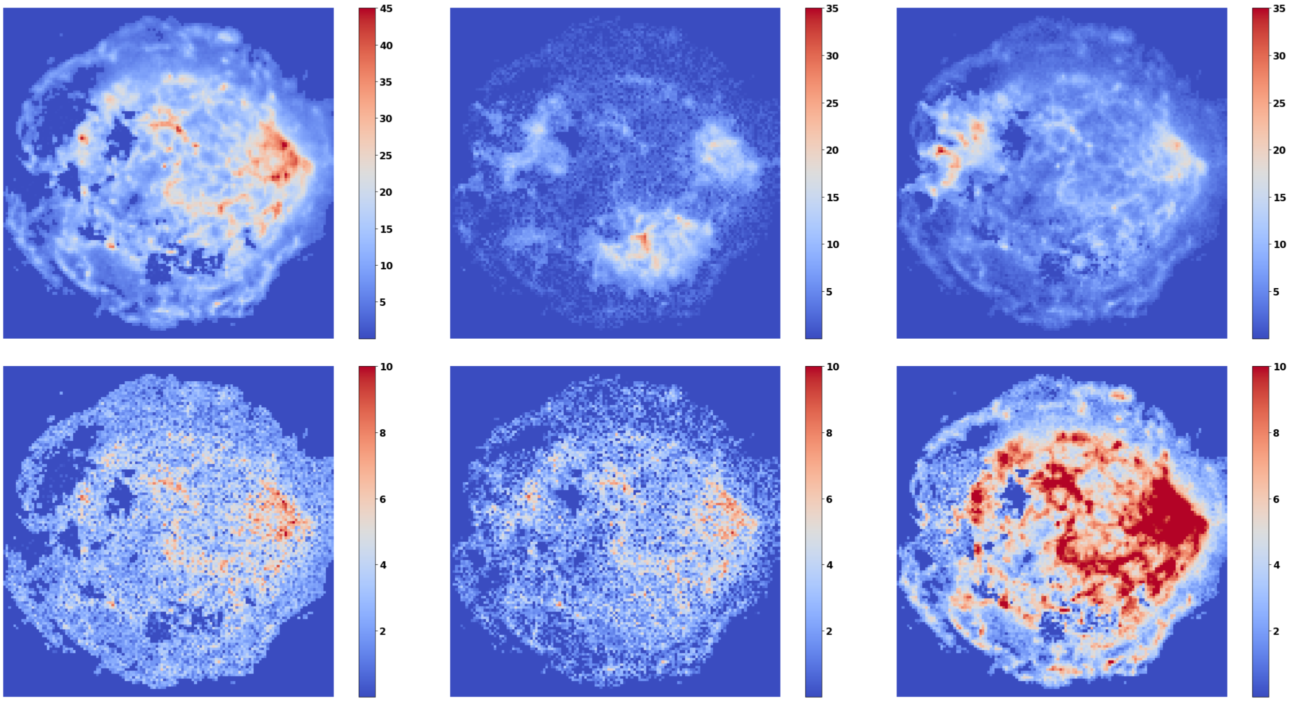

Figure 3.13: Sources estimated with HALS.

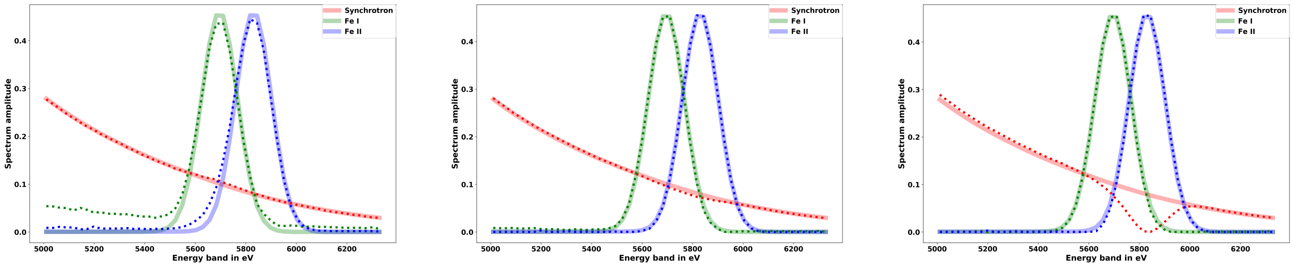

Figure 3.14: Input (solid lines) and estimated spectra (dashed lines) with GMCA, pGMCA and HALS.

3.4.2 Comparisons with respect to the mean flux

In this section, we investigate the performances of BSS when the mean flux evolves. In these experiments the relative intensity parameter is fixed to $\phi = 2$ and the number of observations is fixed to 12. Figure 3.15 displays the evolution of the mean and maximum SAD when the mean flux varies from 0.5 to 100.
In this experiment, the sNMF algorithms do not perform correctly while the β-sparse NMF algorithm provide very good results, and can compete with the GMCA algorithm higher mean flux. To that respect, the differences between sNMF with β-sparse NMF HALS is probably explained by the very good robustness of the HALS algorithm with respect to the initialization.
It can be pointed out that modeling the sparsity of the sources in the wavelet domain largely help improving the separation as shown by the results of GMCA and pGMCA in Fig. 3.15. Further accounting for the Poisson statistics of the observations yields SAD values lower than 0.1° for mean number of counts larger than 8. Interestingly, the results of the GMCA algorithm deteriorate significantly for values larger than 10. Since the measurements follow Poisson statistics, the signal-to-noise ratio (SNR) for a given pixel t and observation i is equal to $\sqrt{[\mathbf{AS}]_i[t]}$. As a consequence, large amplitude data samples have a smaller SNR. Since sparse BSS methods are more sensitive to the samples with the largest amplitudes Bobin et al. [2015], such methods will be

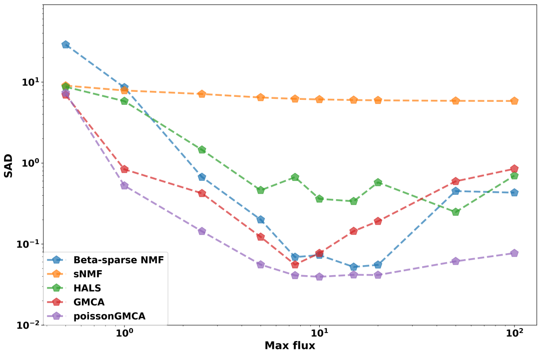

Figure 3.15: Spectral angular distance as a function of the mean flux.

hampered by the Poisson nature of the measurements even for low number of counts. In that case, the ability to account for the exact statistics of the noise allows to mitigate this effect; this leads to a SAD that is one order of magnitude below than the other methods for a mean number of counts of 100.

3.4.3 Evolution of the relative intensity of the sources

In this paragraph, we investigate the impact of the relative intensity of the sources on the separation process. More precisely, in the present astrophysical setting, the synchrotron emission can be regarded as a contamination with respect to the two redshifted iron sources. The accuracy of the estimated sources will therefore depend on the level of the synchrotron. In contrast to the standard additive Gaussian noise, this is particularly challenging when the observations are generated according to a Poisson distribution since the resulting stochasticity (or Poisson noise) will be dependent on the dominating source. This will hinder the separation process to a large extent. As a consequence, we evaluate the ability of the different BSS methods to accurately estimate the two iron components when the relative intensity of the synchrotron evolves.

Numerical experiments have been carried out with two different values for the number of observations: $m = 12$ and $m = 50$. In fact, for a fixed mean number of counts, estimating a fixed number of sources from a larger number of sources implies an increasing dimension reduction factor during the separation process. In the Poisson statistics case, this entails that the noise that contaminates the sources will tend to "Gaussianize" as stated by the law of large numbers when the number of observations soars. This should be a more favorable scenario for standard BSS methods.

Case $m = 12$ The left panel of Fig. 3.16 features the evolution of the SAD when the parameter ϕ varies between 0.1 and 10 for a fixed number of observations $m = 12$. In this experiment, the Poisson average evolves in the range $[10, 655]$. First, the HALS and β-sparse NMF methods provide decent results with an average SAD of about 1°. Interestingly, the sNMF algorithm leads to rather reasonable values of the SAD when the level of synchrotron emission if low (*i.e.* $\rho < 0.5$): in this regime, the two iron sources, which turn to be the sparsest ones, are dominant.

Both the GMCA and pGMCA algorithms lead to their best results for $\rho = 0.3$ when the three sources share quite similar amplitudes. At high synchrotron level, both methods perform similarly with a SAD of 1° when the synchroton emission is ten times larger ($\rho = 10$). In the range $[0.1, 2]$, the pGMCA algorithm provides significantly cleaner separation results with a SAD that is more than one order than the other methods, which highlights the gain obtained by accounting for the exact Poisson nature of the measurements.

Case $m = 50$ The right panel of Fig. 3.16 shows the value of the SAD for $\phi \in [0.5 \text{ and } 10]$. In this experiment, the mean number of counts evolves in the range $[3, 334]$. For a larger number of observations $m = 50$, classical BSS methods perform slightly better than for $m = 12$ while the average number of counts is two times lower.

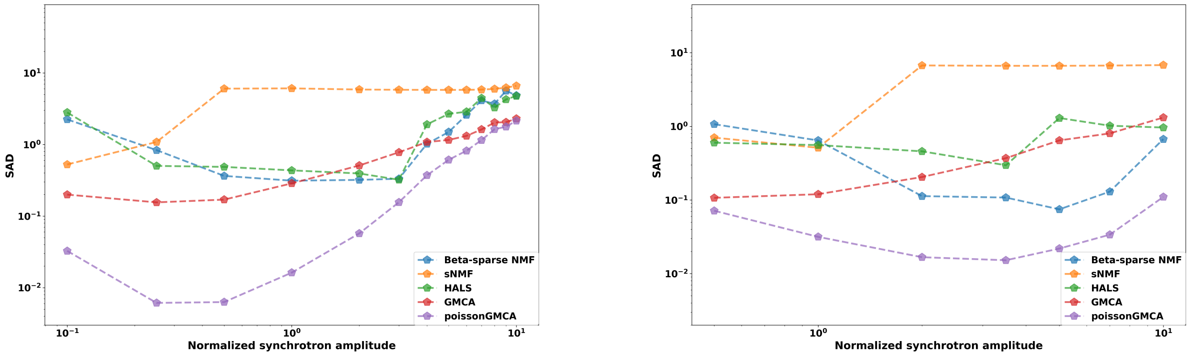

Figure 3.16: Spectral angular distance as a function of the synchrotron level - *left:* $m = 12$, *right:* case $m = 50$.

This is especially visible for large levels of the synchrotron emission. More interestingly, compared to GMCA, the pGMCA algorithm provides slightly worse results, especially for $\rho = 2$. Indeed, as we pointed out previously, as the number of observations increase, standard BSS methods are less hampered by the Poisson statistics of the measurements. Still, the pGMCA yields separation results with a SAD gain up to one order of magnitude for $\rho = 3.5$.

Conclusion

In this chapter, we investigated a new sparsity enforcing method to tackle Blind Source Separation problems from measurements that follow Poisson statistics. Switching from the classical additive Gaussian noise to Poisson statistics entails key challenges: i) as a non-convex problem, the optimization strategy has a strong impact on the efficiency and robustness of the separation process especially when sparse regularization and Poisson likelihood lead to a non-smooth problem, ii) automatic parameter tuning becomes highly difficult in the case of Poisson statistics. For that purpose, the proposed pGMCA algorithm builds upon a BCD-like minimization scheme to optimize an approximate problem with a smooth data fidelity term. For that purpose, it is proposed to make use of Nesterov's smoothing technique to build a differentiable approximation of the Poisson log-likelihood. We further show that the proposed approach allows to define a simple and robust strategy to automatically tune the regularization parameters. Numerical experiments are carried out on both synthetic 1D spectroscopic data and realistic 2D simulations of astrophysical X-ray images. The results show a clear improvement of the separation quality with respect to existing methods in these different settings. We would like to highlight that the pGMCA algorithm has been used with success to analyse supernovae remnants in Picquenot et al. [2020].

Chapter 4

Source Separation in the presence of spectral variabilities

The rapid increase of multi-wavelength sensors has raised specific attention to the analysis of multivalued data. The last decades have witnessed the proliferation of data processing tools such as Blind Source Separation seeking for a more realistic representation of the objects in the light of sight. In this context, the variation of spectra of materials or physical elements to be recovered, known as spectral variabilities phenomenon, is of greater interest to have a faithful modelling of the data. In this chapter, we will show, building onto examples, that spectral variabilities are inherent to multispectral data, whether it is because of the natural phenomena studied or the characteristics of the imagers. Afterwards, we focus on remote sensing that has witnessed the proliferation of multiple methods aiming at tackling spectral variabilities. We present the models introduced in the remote sensing field that explicitly account for spectral variabilities. Finally, we focus on the only Blind Source Separation method present in the literature (to the best of our knowledge) tackling spectral variabilities problem in astrophysics.

4.1 Position of the problem

4.1.1 Spectral variabilities in astrophysical X-ray imaging and beyond

Blind or semi-blind Source Separation has taken the lion's share in providing a robust and efficient analysis framework for astrophysical data. For instance, it has proven to be the case for the Planck mission results (Bobin et al. [2008]). The authors of Picquenot et al. [2019b] have proposed an application of GMCA to the supernova Cassiopeia A Chandra observations. In their paper, they have considered the iron constituent of the supernova as two distinct components to retrieve, whose spectra are Gaussian-shaped with shifted central frequencies. This is a rough approximation since *it is the same component whose spectrum is shifted because of the Doppler effect*. Taking into account this spatial spectral variability would have yielded a better component separation, more accurate source distribution map and would have brought new information (e.g. Fe redshift map). To the best of our knowledge, no BSS method has been proposed in the literature to tackle the problem of spectral variabilities for high-energy astrophysical data.

Variabilities affecting observations acquired by Chandra telescope is an example out of many of spectral variabilities in multivalued data:

- Dynamic Positron Emission Tomography (PET) is an imaging technique aimed at measuring the metabolic activities of an organ through the concentration of a radiotracer injected in the body of the patient. The quantification of this concentration is provided by the extraction of tissue Time-Activity-Curves (TACs) from images acquired at different time-frames, which can be viewed as a BSS problem. However, classical BSS methods are not well-adapted to this application since the TAC associated to one tissue (the specific binding tissue) is spatially variant (in addition to temporally) Cavalcanti et al. [2018].

- Materials in remotely sensed images may present more than one spectrum throughout the line-of-sight of the imagers and not accounting for it is the major source of error in terrestrial data analysis (Somers et al. [2011]).

More generally, the presence of spectral variabilities is ubiquitous to natural scenes (for example because of illumination changes Castrodad et al. [2011]).

Neglecting such an important phenomenon raises two problems. First of all, the risk of propagating errors on the source matrix and second of all, the loss of information due to the neglect of those spectral variabilities.

As mentioned in 1.3.2, there is very few literature in BSS on addressing spectral variabilities in the context of astrophysical data. The majority of State-of-the-Art techniques find their roots in remote sensing.
In the following of this chapter, we start by reminding some basic notions on remote sensing and compare its specifics to the ones that are of interest in this thesis. In a second part, we present some successful spectral variabilities addressing methods in remote sensing and their limitations when applied to X-ray astrophysical imaging and finally, we present a BSS method used to tackle variabilities and applied to astrophysics.

4.2 Spectral unmixing in remote sensing: a framework to addressing spectral variabilities

To the best of our knowledge, the vast majority of methods addressing the problem of spectral variabilities find their roots in remotely sensed data.

Therefore, before presenting methods accounting for spectral variabilities, it appears to us that it would be important to have an overview of the specifics of remote data processing, its assumptions, its similarities with the ones we use in this thesis and its specific methods that are used as basis for spectral variabilities addressing methods.

4.2.1 The Linear Mixture Model in remote sensing

Multispectral imaging permits the monitoring of processes taking place on the surface of earth. Remotely sensed data consist of the observation of a scene in various spectral bands. The materials that one seeks at monitoring through spectral imaging are substances of macroscopic homogeneous characteristics present in the scene. However performing the imager might be, its spatial resolution is limited and therefore, each pixel of the resulting image cannot correspond to only one material. The composite nature of pixels is even more present with hyperspectral images since the gain on spectral information comes with a loss on spatial resolution.
Blind Source Separation or the so-called spectral unmixing has proven to be a powerful tool in remotely sensed signal modeling Keshava and Mustard [2002]. It consists in considering the measured spectrum as the mixture of constituent spectra. In the following, we will name the materials constitutive of the observed scene *endmembers* as well as their characteristic spectra denoted $\mathbf{E}$. The spatial distribution of each material in the image will be called *abundances* and denoted $\hat{\mathbf{A}}$.
The analogy with source separation as we apply it in astrophysics is straightforward: the mixing matrix, whose columns represent spectra constitutive of the line of sight on the sky, is the equivalent of the endmember matrix. On the other hand, the source matrix plays the role of the abundance matrix.

As in astrophysics, the most popular modelling in remote sensing is the Linear Mixture Model: the mixing process is supposed linear.
Let $\mathbf{E} \in \mathbb{R}^{m \times n}$ denote the collection of n observed spectra $e_{i,1 \leq i \leq n}$. Let $\mathbf{Y} \in \mathbb{R}^{m \times t}$ stand for the observations, $\hat{\mathbf{A}} \in \mathbb{R}^{n \times t}$ for the abundance matrix and $\mathbf{N} \in \mathbb{R}^{m \times t}$ for the model imperfections and the noise supposed additive Gaussian. The signal is described as the linear sum of the constituent contributions, which can be written in the following matrix form:

$$\mathbf{Y} = \mathbf{E}\hat{\mathbf{A}} + \mathbf{N}, \tag{4.1}$$

The endmember matrix is comprised of the components spectra in its columns and the abundance matrix is the concatenation of the abundances line vectors. For the sake of simplicity, the conventions on the notations of the number of sources, observations and samples will be kept in this part *i.e.* (m, n, t).

Despite the popularity of this model, remotely sensed data are not exempt of nonlinearities. As explained in

Keshava and Mustard [2002], the model (4.1) is only valid when the materials appear in a spatially segregated pattern and the mixing scale is macroscopic.

- If the materials are in a multilayered configuration, the photons interact with more than one material, causing the radiance[1] acquired by the sensor to be nonlinear to the radiance of each material. A typical example of this phenomenon would be the observation of plant canopies areas: the light refracted by the ground materials interacts with the canopy (Keshava and Mustard [2002]).

- The second case of nonlinearities in remote sensing appears when the materials interact at a microscopic level. This shows a crucial subtlety in the definition of the endmembers: if we aim at recovering the mixture of the materials interacting at microscopic level as one endmember, the problem becomes linear on the opposite of recovering each one separately Bioucas-Dias et al. [2012].

On the latter, we consider a "check-board type scene" with materials interacting at a macroscopic level.

In remote sensing, abundances represent not so much spatial distribution as fractions of the present endmembers. The vision of data unmixing is different from the one we are accustomed to in astrophysics: in remote sensing, each pixel spectrum is seen as the combination of the endmembers spectrum **weighted by their proportions in that pixel**. In astrophysics, the source matrix is seen as the spatial distribution of the elements in the image.
We would like to point out that considering the abundance of a material in a pixel as representative of its proportion is a very common approximation but may not be always accurate. For example, a highly reflective material in low proportion will have a very high abundance when using reflectance[2] to collect data. In this case, considering the abundance as a quantification of the physical proportion of the material is erroneous (Bioucas-Dias et al. [2012]). In the remainder, we will neglect such particular cases.

4.2.2 Spectral unmixing based upon geometrical approach

Recovering the couple $(\mathbf{E}, \hat{\mathbf{A}})$ from the data matrix $\mathbf{Y}$ (and possibly the number of components n) is an ill-posed problem that calls for additional assumptions.

In the context of remotely sensed data, it is customary to account for two specific constraints:

- The abundances relative to each pixel, given that they stand for proportions of the constituent endmembers, must have a sum equal to one.
  This assumption is known as the **abundance sum-to-one constraint (ASC)**:

  $$\forall k = 1..t, \sum_{i=1}^{n} a_i[k] = 1, \tag{4.2}$$

  with $a_i[k]$ the abundance of $i - th$ endmember in the pixel k.

- The second classical remote sensing assumption is the so-called **abundance non-negative constraint (ANC)** which constraints the positivity of the abundances as they stand for endmembers fractions:

  $$\forall i = 1..n, \forall k = 1..t, \ a_i[k] \geq 0. \tag{4.3}$$

Therefore, it appears that the data live in the convex hull spanned by the endmembers characterized by the two aforementioned constraints. It can be written as

$$\mathcal{C} \equiv \{ \ \mathbf{X} = \mathbf{E}\hat{\mathbf{A}} + \mathbf{N} : \ A_i^T \mathbf{1} = 1, \ A_i \leq \mathbf{0}, \ i = 1.....n \ \}, \tag{4.4}$$

---

[1] Spectral radiance is the spectral flux that reaches the instrument per unit area and per unit of solid angle, perpendicular to the surface of the detector.

[2] The spectral reflectance is a measure of the proportion of light reflected by a surface at a given wavelength.

with A_i the $i-th$ line of the abundance matrix and $\mathbf{1}$ (*resp.* $\mathbf{0}$) the vector of size t whose elements are equal to 1 (*resp.* 0).
Assuming that the endmembers are linearly independent, the set $\mathcal{C}$ is the $(n-1)$ **simplex in** $\mathbb{R}^m$.

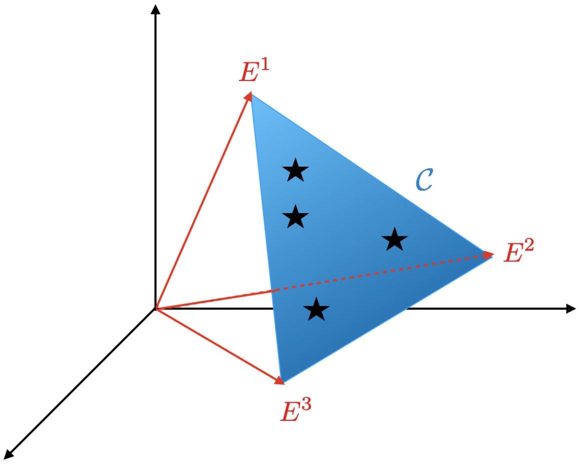

Figure 4.1: Illustration of the simplex $\mathcal{C}$ (in blue) in which live the data represented by the black stars. The three endmembers are depicted in red.
If the ASC is relaxed, the data live in the cone spanned by the endmembers E^1, E^2 and E^3.

The ASC constraint is very restrictive since it implies the presence of at least one component on each pixel, which is not the case for astrophysical X-ray imaging or some terrestrial hyperspectral data Chenot [2017].

Geometrical LMM based methods

We propose a quick overview of Geometrical approaches in remote sensing given that some spectral variabilities addressing algorithms are built upon such methods.
The simplex in which spectral data live is fully characterized by the endmembers since they correspond to its vertices. Reciprocally, one can theoretically recover the endmembers from the simplex. The geometrical methods are built upon the idea that one can infer the simplex characteristics (*e.g.* its volume) from the dataset. However, characterizing the simplex from the data points can be very challenging, due the high indeterminacy of the model. The ANC and ASC assumptions are not sufficient to ensure the uniqueness of the solution. Therefore, some geometrical methods propose a third constraint which consists in considering that there exists at least one "pure" data point for each endmember. It is known as the pure pixel assumption.

Pure pixel (PP) methods Let us suppose that there is in the captured scene at least one pixel where there is no mixture, in other words only one material is observed. In spectral dimensions, this pixel defines one of the vertices of the simplex we aim at recovering. If there is for each material such a pure pixel, the simplex is then given by the maximal volume enclosing all the spectral data and its extremal points correspond exactly to the sought endmembers. Many iterative algorithms have been proposed to detect such pure pixels; the authors of Winter [1999] propose to inflate a simplex inside the data until obtaining the one of maximal volume; the Pure Pixel index algorithm (Boardman [1993]) consists of projecting data points onto skewers until finding the extremal ones. One can also think of the Vertex Component Analysis algorithm (Nascimento and Dias [2005b]), the Iterative Error Analysis algorithm (Neville et al. [1999]) and many others.

Assuming the presence of pure pixels allows easy interpretation of the methods and efficient algorithms. However, this assumption is very strong in terms of constraints. It implies that the observed scene contains areas constituted of a unique element; and that these areas are wide enough so that the instrument captures it as pixels whatever its resolution is. The first constraint inhibits the generalization of PP methods to a tremendous number of data and the second one is very hard to fulfill particularly when exploiting hyperspectral images (the upgrading of spectral resolution comes with a deterioration of the spatial one). Therefore, geometrical methods relaxing the PP assumption have been proposed.

Simplex facet methods By simplex facet methods we refer to geometrical approaches that require the pixels not to be highly mixed. For instance, we can consider that there is at least a few data points that come from the weighted sum of only two endmembers. These latter lie therefore in a facet of the simplex in the spectral space. The main idea of this class of methods is seeking for the simplex enclosing all pixels and of minimal value. Therefore, this approach encompasses Pure Pixel algorithms since in that case, the simplex of maximal volume is confused with the one of minimal volume.
The relaxation of pixels purity comes with a complexification of the objective function to minimize. As for the PP assumption, iterative methods have been developed, consisting at evaluating the volumes of each possible simplex and selecting the one of minimal value (Craig [1994]). Other methods build upon optimization techniques to minimize a cost function that enforces the simplex to have a minimal volume. Supplementary constraints have been proposed such as penalization on the violation of the ANC constraint (to handle noise perturbations, *e.g.* Li and Bioucas-Dias [2008], Bioucas-Dias [2009], Chan et al. [2009]) or non-negativity of the couple $(\mathbf{E}, \hat{\mathbf{A}})$ (*e.g.* Miao and Qi [2007]). Given the non-convexity of the volume term, the authors of Berman et al. [2004] proposed replacing it by the sum of squared distances between the simplex vertices.

Geometrically based approaches have straightforward interpretation and the algorithms are generally easy to compute. However, most of these methods are based on very strong requisites that restrain their application range. Not only they cannot be applied to datasets that do not fulfill the simplex constraint, they are also inefficient for highly mixed data in the simplex framework.

Geometrical approaches cannot be applied to X-ray imaging since they are built upon the simplex constraint. However, both of the Pure Pixel assumption and the Simplex facets methods induce sparsity of the abundances that can be captured with sparse BSS (Chenot [2017]). Indeed, the presence of pure pixels amounts to sparse abundances that nearly respect the Morphological Diversity principle and the Simplex facets imply a low mixture of sparse sources. However, sparsity prior is not sufficient to raise the indeterminacies of the problem and additional assumptions should be considered (such as non-negativity of the spectra and the abundances).

4.2.3 Statistical approaches and sparse regression for remote sensing

The spectral unmixing approaches can be recast as three main categories (Bioucas-Dias et al. [2012]): the geometrical approaches that take roots in the interpretation of the simplex geometry, the statistical approaches that consider the problem as a statistical inference one and the sparse regression approaches that are of interest in our work.

Statistical approaches

The main idea of statistical approaches is to consider spectral unmixing as an inference problem addressed with Bayesian estimation theory. The methods detailed below propose a joint posterior estimation of the couple $(\mathbf{E}, \hat{\mathbf{A}})$ through the Maximum A Posteriori (MAP) estimator. By denoting P_X (resp. $P_{X|Y}$) the probability density function of X (resp X given Y), the estimation is computed using the Bayes paradigm:

$$P(\mathbf{E}, \hat{\mathbf{A}}) \propto P_{\mathbf{Y}|\mathbf{E},\hat{\mathbf{A}}}(\mathbf{Y}|\mathbf{E}, \hat{\mathbf{A}}) P_{\mathbf{E}}(\mathbf{E}) P_{\hat{\mathbf{A}}}(\hat{\mathbf{A}}), \tag{4.5}$$

$P_Y(Y)$ being a normalization constant.

If we consider the noise to be Gaussian i.i.d. with covariance matrix $\sigma^2 I$, the joint maximum a posteriori corresponds to the joint minimization of the corresponding criterion

$$(\hat{\mathbf{E}}, \hat{\hat{A}}) = \arg\min_{\mathbf{E},\hat{\mathbf{A}}} \frac{1}{2\sigma^2} \left\| \mathbf{Y} - \mathbf{E}\hat{\mathbf{A}} \right\|_F^2 - \log P_{\mathbf{E}}(\mathbf{E}) - \log P_{\hat{\mathbf{A}}}(\hat{\mathbf{A}}).$$

The analogy with the convex optimization framework is straightforward; the first term corresponds to the data-fidelity term and the priors antiloglikelihood play the role of the penalizations on each variable. For example, the authors of Arngren et al. [2011] developed a Bayesian tool that can be seen as the stochastic version of the geometrical MVC-NMF approach mentioned above (Miao and Qi [2007]).

The majority of stochastic approaches are based on the posterior estimators. The difference between the methods usually appear in the chosen prior distributions.For example, the authors of Moussaoui et al. [2006] have assigned to the priors of $\mathbf{E}$ and $\hat{\mathbf{A}}$ distributions that are hyperparameter-dependent, and to these hyperparameters non-informative priors. Given the complexity of such hierarchical bayesian modeling, the authors exploit the Markov Chain Monte-Carlo (MCMC). This is a simulation tool that, in a nutshell generate random variables that asymptotically will follow the sought posterior distribution. The authors of Dobigeon et al. [2009a] and Dobigeon et al. [2009b] follow the same hierarchical bayesian modeling with differences on the applications and on the selected priors.

The stochastic framework allows for a wider application range, however it comes with a very important computational complexity which may be extremely limiting for large datasets.

Sparse regression

The last but not least class of approach we present is the one of most interest in this chapter since it is very close to the approach that we use in the rest of the manuscript.

When considering endmembers as deterministic sets and not probabilistic distributions, the aforementioned techniques (geometrical approaches) suffer of several drawbacks. Firstly, the presence of pure pixel or at least not highly mixed spectral pixels is not guaranteed (and generally not met), secondly these methods rely on the ability of the algorithm to detect such pure spectral signatures while the spatial resolution may be insufficient. In this context, the authors of Iordache et al. [2011], Rogge et al. [2006] and Iordache et al. [2010] have proposed to unmix observed scenes in a semi-supervised fashion. They make use of the availability of spectral libraries *i.e.* libraries containing pure signatures generally acquired on the field by spectro-radiometer. Therefore, spectral unmixing boils down to identifying which endmembers of the spectral libraries are present in the observed pixels. Let us denote the spectral libraries $\mathbf{e}$, it is composed of numerous endmembers $\mathbf{e} = [\mathbf{E}^1, \mathbf{E}^2 ... \mathbf{E}^L] \in \mathbb{R}^{m \times L}$. The number of active endmembers in the scene is very low in comparison with the number of available spectral signatures, $n \ll L$. Thus, the recovery of these active endmembers can be done through a sparsity enforcing term on the endmembers.

The problem of finding the optimal subset of $\mathbf{e}$ to model the acquired data can be written as:

$$\min_{\hat{\mathbf{A}}} \left\| \hat{\mathbf{A}} \right\|_0 s.t. \left\| \mathbf{Y} - \mathbf{e}\hat{\mathbf{A}} \right\|_2 \leq \epsilon, \hat{\mathbf{A}} \geq \mathbf{0}, \tag{4.6}$$

or in its convexified form by replacing the ℓ_0 norm by the ℓ_1 norm:

$$\min_{\hat{\mathbf{A}}} \left\| \hat{\mathbf{A}} \right\|_1 s.t. \left\| \mathbf{Y} - \mathbf{e}\hat{\mathbf{A}} \right\|_2 \leq \epsilon, \hat{\mathbf{A}} \geq \mathbf{0}. \tag{4.7}$$

This problem, known as linear sparse regression can be solved using greedy algorithm such as the Orthogonal Matching Pursuit (Pati et al. [1993]) or convex-based methods such as the Basis Pursuit Denoising (Chen et al. [2001]).

Eq. (4.7) can be equivalently written as

$$\min_{\hat{\mathbf{A}}} \left\| \mathbf{Y} - \mathbf{e}\hat{\mathbf{A}} \right\|_2 + \lambda \left\| \hat{\mathbf{A}} \right\|_1 s.t. \hat{\mathbf{A}} \geq \mathbf{0}, \tag{4.8}$$

known as the constrained sparse regression (CSR) problem where λ stands for a trade-off between the data-fidelity term and the sparsity of the abundances (Bioucas-Dias and Figueiredo [2010]).

If the ASC constraint is considered in addition to the ANC, CSR changes into the Fully Constrained Least Square problem (FCLS), mathematically formulated as

$$\min_{\hat{\mathbf{A}}} \left\|\hat{\mathbf{A}}\right\|_1 s.t. \left\|\mathbf{Y} - e\hat{\mathbf{A}}\right\|_2 \leq \epsilon, \hat{\mathbf{A}} \geq \mathbf{0}. \tag{4.9}$$

Therefore, we obtain the number of present endmembers present in the scene (which was unknown), the nature of these endmembers (corresponding to the ones of non-null abundance vector) and the composition of each pixel of the image.

Despite of the simplicity and the efficiency of such an approach, it presents two major drawbacks: on the one hand it supposes being able to obtain such libraries in the field (impossible for obvious reasons in astrophysics), and on the other hand; one cannot be ensured that the data acquisition of spectral libraries will be done under the same conditions as with airborne data which requires reliable calibration pre-processing.

4.3 Addressing spectral variabilities in spectral unmixing and Blind Source Separation

4.3.1 Main approaches to spectral variabilities

We have seen the principal classes of methods of spectral unmixing in remote sensing. All of these methods are based upon two crucial elements:
(i) we suppose a linear combination of the endmembers at each pixel,
(ii) we assume only one spectral signature per endmember.
However, both of these assumptions are not met in real-world data.

Whereas nonlinear spectral mixture analysis has flourished in the literature to address endmember nonlinearities, spectral variabilities is still very poorly studied. Indeed, in most spectral unmixing methods, we confound the nature of the endmember and its spectral signature, even though it is more than likely that one material/constituent of the observed scene appears in the image through more than one spectral signature/column of the endmember matrix $\mathbf{E}$. Indeed, endmembers (or spectral signatures) of a material may vary within the extent of an image due to two principal causes in remote sensing: either the material itself presents an intrinsic variability either the variabilities are consequences of environmental changes. These changes are mostly due to different conditions of illumination of the scene or atmospheric changes when capturing images with airborne sensors.

The authors of García-Haro et al. [2005] showed that neglecting the SV phenomenon induces possibly highly erroneous estimation of abundances and the nature of elements captured. Regarding the spectral variabilities due to atmospheric changes, Gao et al (Gao et al. [2009]) proved that the classical pre-processing step of atmospheric correction is not sufficient to eliminate the effect of the SV.

Therefore, in the following, we will make the distinction between **the materials** meaning by it the components present in the scene and **the endmembers**, which are multiple for one material when it is affected by spectral variabilities.

Given that spectral variabilities appear for reasons as common as a change of illumination during acquisition time, remote sensing community has developed some specific tools to this phenomenon.
We can classify these methods in three families:

- methods deriving from the statistical approaches seen in 4.2.3,
- methods addressing the SV as deterministic sets,
- modification of the Linear Mixture Model to explicitly account for SV.

Considering that the aim of this thesis is to propose an unsupervised decomposition model for multivalued data in the presence of spectral variabilities, we propose to present briefly the statistical and set-based approaches while we will focus on the modelization of the SV that are of interest for our work.

4.3.2 Statistical approach

The main idea of this class of approaches is to consider that the spectral signatures of one material are samples of the same multivariate statistical distribution. If we denote e_i the endmember of the element i, we have (Zare and Ho [2013]):

$$e_i \sim \mathcal{F}(.|\theta_i), \tag{4.10}$$

where θ_i stands for the parameters of the distribution $\mathcal{F}$ associated with the $i-th$ element.
The spectral vectors $\mathbf{Y}$ are generated according to the stochastic mixture model:

$$\mathbf{Y}[k] = \sum_{i=1}^{n} e_i \hat{\mathbf{A}}_i[k] \; \forall k = 1..t, \tag{4.11}$$

where e_i is distributed according to eq. (4.10), and where the data are supposed noiseless.
Two main statistical models have been studied in the literature: The Normal Composition Model (NCM, Stein [2003]) which describes endmembers as random variables of normal distribution and the Beta Composition Model (BCM, Xu and Yin [2017]) which, by assuming a beta distribution to the endmembers assures the sum-to-one constraint.
Given that statistical approaches are not at the core of this thesis, we suggest to the reader to refer to Zare and Ho [2013] for a detailed overview.

4.3.3 Set-based approaches

The literature on the study of spectral variabilities when considered as deterministic sets showws three main classes:

1. Extracting endmembers variabilities present in the scene from spectral libraries known in advance;
2. Generating spectral bundles from the dataset itself;
3. Taking fully into account spectral variabilities through an optimization problem.

We propose a more detailed overview of these methods in the next section.

Known spectral libraries

The authors of Roberts et al. [1998] proposed in 1998 one of the first methods to take into account spectral variabilities. They proposed the nowadays widely used Multiple Endmember Spectral Mixture Analysis (MESMA) which consists of an exhaustive search on a spectral library of the endmembers that describe each pixel the best ("the best" meaning here with the lowest mean square error between the true pixel and its reconstruction). Many extensions have been developed over the years such as the AutoMCU (Asner and Lobell [2000]) that selects randomly endmembers from the spectral library. Despite their simplicity, these cyclic methods present very high computational cost and are dependent on libraries which, in itself presents many drawbacks as we have seen in 4.2.3 (the library samples are not acquired in the same conditions as the data to analyze). We can however note that the authors of Fu et al. [2016] have proposed a technique to allevate the mismatch between the endmembers of the library and the ones observed in the scene by taking it into account in a CSR-based algorithm.

Spectral bundles

Somers et al proposed to automatically extract spectral bundles from the dataset (Somers et al. [2012]). The process chain consists of running geometrical-based approaches on randomly selected subsets, which generates a spectral library from the dataset itself, using a clustering algorithm to assign to each endmember the corresponding material and afterwards apply MESMA algorithm to characterize each spectral pixel.
Many other methods have been proposed to extend the classical spectral unmixing techniques to the presence of spectral variabilities. For example, the authors of Castrodad et al. [2011] propose to account for spectral variability by exploiting sparse coding and dictionary learning.

4.3.4 Proper modelling of the SV through a modification of the LMM

During these last years, methods have been proposed not only to allievate misestimations due to the neglect of spectral variabilities but to properly model this phenomenon.

These methods and the corresponding algorithms have been extensively exploited because of their accuracy to model real-world remotely sensed data and the efficiency of the algorithm. In the next section, we propose to present these models and also to **interpret them in the scope of X-ray astrophysics**, which is the core of this thesis.

4.4 Modeling spectral variabilities

The methods presented below propose to take into account spectral variabilities by a **modification of the Linear Mixture Model**:

$$\mathbf{Y} = \mathcal{L}(\mathbf{E}, \hat{\mathbf{A}}, \mathbf{V}),$$

where $\mathbf{V}$ is a term describing the spectral variabilities and $\mathcal{L}$ stands for the mixture model chosen to account for the SV. Retrieving the endmembers, the abundances and the spectral variabilities obviously requires penalizations on each of these terms.

Given that the methods presented below have been developed in the context of remote sensing, they rely on the **simplex constraint** meaning that the data respect:

- the sum-to-one constraint: for each sample, the sum of the abundances contribution is equal to one,
- the non-negativity constraint: the endmembers and the abundances are supposed to be non-negative.

Moreover, these approaches rely on the assumption of pure pixel (or at least that the pixels are not too highly mixed) which is excluded for the astrophysical data we study. Instead, for astrophysical applications, a fair modelling would be to assume that the endmembers **live on a hypersphere**:

- Endmembers normalization: The endmembers (or columns of the mixing matrix) are supposed to live in the hypersphere $\mathcal{S}^{m-1}$,
- Non-negativity constraint: since the endmembers (respectively the abundances) stand for spectral (respectively spatial) distribution of the observed constituents, it is natural to enforce their non-negativity.

In each of these frameworks (remote sensing or X-ray astrophysics), the aforementioned constraints are not sufficient to recover the SV because of the indeterminacies of the models. Consequently, the methods presented below propose **SV penalizations** to cope with this problem. The effect of SV penalizations differs drastically if we consider remote sensing framework or astrophysics framework.

We present in this section the models accounting for SV with their most relevant algorithmic details. **We also focus on the proposed penalization on the spectral variabilities and their entailment in X-ray astrophysics framework**

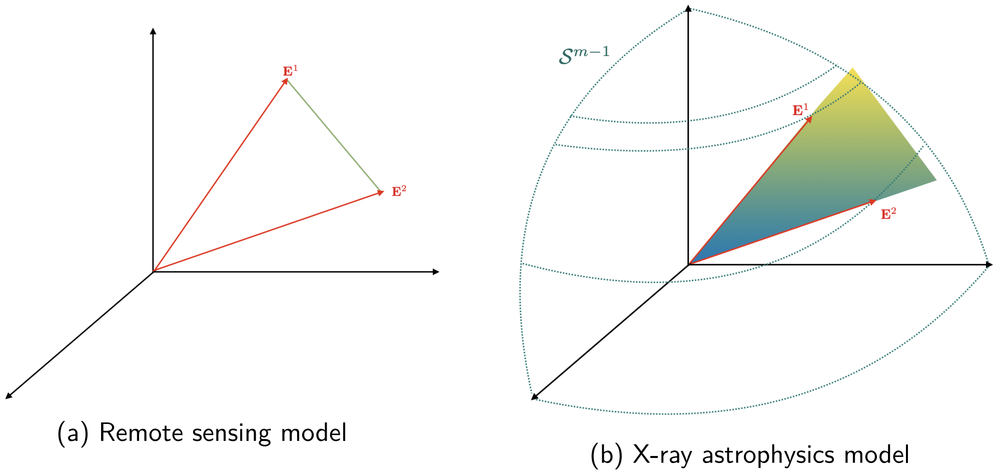

(a) Remote sensing model

(b) X-ray astrophysics model

Figure 4.2: Example of the data space for the two models with two endmembers: (a) data live in the simplex (in green) generated by $\mathbf{E}^1$ and $\mathbf{E}^2$, (b) data live in the convex cone (in green) generated by the endmembers $\mathbf{E}^1$ and $\mathbf{E}^2$ living in the hypersphere $\mathcal{S}^{m-1}$.

4.4.1 The Extended Linear Mixture Model

Terrestrial images are often prone to perturbations predominantly due to variable illumination and atmospheric conditions of the scene, causing variability in the spectral signatures of the endmembers. As explained in Nascimento and Dias [2005a], this special type of variations is characterized by shape invariance: it is the energy of the spectra that differs. As proposed in Nascimento and Dias [2005a] and Nascimento and Dias [2005b] and taken over in Veganzones et al. [2014], this class of spectral variabilities can be modeled by a positive scaling factor varying across the pixels.

Based on this rationale, the authors of Veganzones et al. [2014] propose a modification of the Linear Mixture Model coined **Extended Linear Mixture Model** (ELMM) that describes a pixel $\mathbf{Y}[k] \in \mathbb{R}^t$ as

$$\mathbf{X}[k] = \sum_{i=1}^{n} \mathbf{\Psi}_{ik} \bar{\mathbf{E}}^i \hat{\mathbf{A}}_i[k] + \mathbf{N}[k], \tag{4.12}$$

where $\bar{\mathbf{E}}^i$ is the $i-th$ column of the reference endmember matrix and $\mathbf{\Psi}_{ik}$ is a scaling factor describing the change of amplitude of the $i-th$ endmember at the $k-th$ pixel.
The formulation (4.12) can be written in the following matrix form:

$$\mathbf{Y} = \bar{\mathbf{E}}(\mathbf{\Psi} \odot \hat{\mathbf{A}}) + \mathbf{N}, \tag{4.13}$$

where

- $\bar{\mathbf{E}} \in \mathbb{R}^{m \times n}$ is the reference endmember matrix;
- $\hat{\mathbf{A}} \in \mathbb{R}^{n \times t}$ stands for the abundance matrix;
- The term $\mathbf{\Psi} \in \mathbb{R}^{n \times t}$ contains all the scaling factors that are endmember-dependent in Drumetz et al. [2016];
- The term $\mathbf{N} \in \mathbb{R}^{m \times t}$ stands for the noise and the model imperfections;
- The notation $\odot$ stands for the Hadamard product.

Variation of illumination due to surface topography is easily detected by the remote sensing model: in that case, data do not lie anymore in the simplex but rather in the positive hypercone spanned by the reference endmembers as it can be seen in figure 4.3. The coefficients of the endmembers $\mathbf{\Psi}_{ik}\hat{\mathbf{A}}_i[k]$ are still non-negative but do not sum to one even though the abundances respect the ASC constraint.

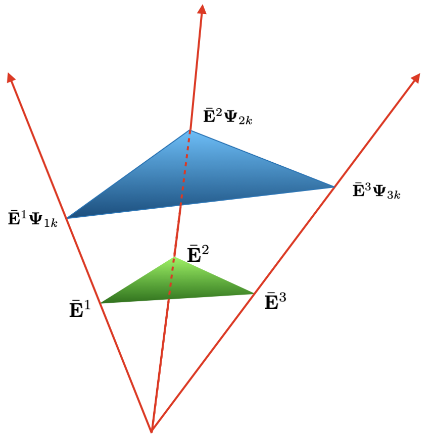

Figure 4.3: Illustration of the positive hypercone in which live the data when considering ELMM framework. In the absence of spectral variabilities, the data would lie in the simplex depicted in green. (Figure inspired by Drumetz et al. [2016])

Regarding the model we are accustomed to in astrophysics, differential illumination and atmospheric conditions are not accounted for as spectral variabilities since the model we consider only allows shape-based variabilities. As suggested in Chenot [2017], the contribution of illumination variations is modeled in this case by a change of amplitude of the sources since the energetic contribution of the data lie in there. Detecting variation of illumination in this model can only be performed by decoupling the effective sources contribution and the illumination contribution from the global sources term. This is a very tricky operation that would require supplementary information.

Problem formulation

In their original paper Veganzones et al. [2014], the authors propose to consider illumination variations that are affecting all spectral bands equally *i.e.*

$$\forall i = 1..n,\ \forall k = 1..t,\ \mathbf{\Psi}_{ik} = \mathbf{\Psi}_k.$$

The equation (4.12) becomes then

$$\mathbf{Y}[k] = \sum_{i=1}^{n} \bar{\mathbf{E}}^i \mathbf{w}_i[k] + \mathbf{N}[k],$$

where the coupled term $\mathbf{w}_i[k] = \mathbf{\Psi}_k\hat{\mathbf{A}}_i[k]$ plays the role of weighting factor. The problem can thus be solved by an approximation of the CLS (constrained least square).

In Drumetz et al. [2016], they propose to extend their model to blind hyperspectral unmixing. The key idea of their method is to enforce the spatial coherence of the fluctuations on illumination and atmospheric conditions. They propose to minimize the following optimization problem:

$$\min_{\mathbf{E},\ \mathbf{\Psi},\ \hat{\mathbf{A}}} \frac{1}{2}\sum_{k=1}^{t} \left\|\mathbf{Y}[k] - \mathbf{E}[k]\hat{\mathbf{A}}_k\right\|_F^2 + \alpha\left\|\mathbf{E}[k] - \bar{\mathbf{E}}\mathbf{\Psi}_k\right\|_F^2 + f_{\hat{\mathbf{A}}}(\hat{\mathbf{A}}) + f_{\mathbf{\Psi}}(\mathbf{\Psi}). \tag{4.14}$$

- The two first terms as modelling terms: the first one corresponds to the data-fidelity term under Gaussian noise while the second one enforces the ELMM. Their decoupling allows small perturbations of the model depending on the value of the hyperparameter α. Moreover, it helps raising the scaling indeterminacy present in (4.13) between the abundances term $\hat{\mathbf{A}}$ and the scaling factor term $\mathbf{\Psi}$.
- The function $f_{\mathbf{\Psi}}$ promotes spatial coherence of the scaling factors through the penalization of horizontal and vertical gradients.
- The abundances are also supposed to exhibit spatial coherence. More specifically, they are assumed to have channel-by-channel spatial regularization which is incorporated through the $\ell_{2,1}$ norm of the gradients. As it is customary in remote sensing, the ANC and ASC are enforced for physical considerations. The ASC also helps raising scaling indeterminacy between the scaling factors and the abundances since it plays the role of calibration of these latters.
- The reference endmembers $\bar{\mathbf{E}}$ are extracted from a classical EEA.

Practical implementation:

Given the multi-convexity of the problem, the authors make use of an alternating non-negative least square algorithm (ANLS Kim and Park [2008]).

Unfortunately, no discussion on the parameter settings has been included, and the displayed results of the the algorithm have been obtained with regularization parameters chosen empirically as the ones that lead to the best performances.

4.4.2 The Perturbed Linear Mixture Model

The PLMM (Thouvenin et al. [2015]) aims at taking into account spatially variant spectral variabilities without any prior knowledge, making it well-suited to a very broad number of applications. The idea is to model the observed varying endmember as *a constant spectral signature affected by a small additive perturbation*.

By denoting $\mathbf{E}^i[k]$ the endmember of the $i-th$ material at the position k, we have:

$$\mathbf{E}^i[k] = \bar{\mathbf{E}}^i + \mathbf{dE}^i[k], \tag{4.15}$$

where

- $\bar{\mathbf{E}}^i \in \mathbb{R}^m$ models the unique "pure" spectral signature relative to the $i-th$ material;
- the variations to the spectral signature are explained with the small additive term $\mathbf{dE}^i[k] \in \mathbb{R}^m$.

The observations can therefore be written as:

$$\mathbf{Y} = \hat{\mathbf{A}}\bar{\mathbf{E}} + \mathbf{\Delta} + \mathbf{N},$$

where

- The term $\hat{\mathbf{A}}$ stand for the abundances;
- $\mathbf{\Delta} = [\mathbf{dE}[1]\hat{\mathbf{A}}_1, ..., \mathbf{dE}[t]\hat{\mathbf{A}}_t]$ where the matrix $dE[k] \in \mathbb{R}^{m\times n}$;
- The term $\mathbf{N}$ stands for the noise and the model imperfections.

This model is very interesting since it enables to account for a wide number of spectral variabilities by a simple additive term.

The adaptation of PLMM to the framework we are accustomed to in X-ray astrophysics is all of the more interesting. Indeed, it amounts to considering that:

- *data are generated according to the PLMM:*

$$\forall k = 1..t,\ \mathbf{X}[k] = \sum_{i=1}^{n} \mathbf{A}^i[k]\mathbf{S}_i[k] + \mathbf{N}[k], \tag{4.16}$$

with

$$\forall i = 1..n, \forall k = 1...t,\ \mathbf{A}^i[k] = \bar{\mathbf{A}}^i + \Delta^i[k], \tag{4.17}$$

and

$$\left\|\Delta^i\right\|_2 \ll \left\|\bar{\mathbf{A}}^i\right\|_2.$$

- *Each mixing direction belongs to the hypersphere i.e.*

$$\forall i = 1..n, \forall k = 1...t,\ \bar{\mathbf{A}}^i \in \mathcal{S}^{m-1} \text{ and } \mathbf{A}^i[k] \in \mathcal{S}^{m-1}.$$

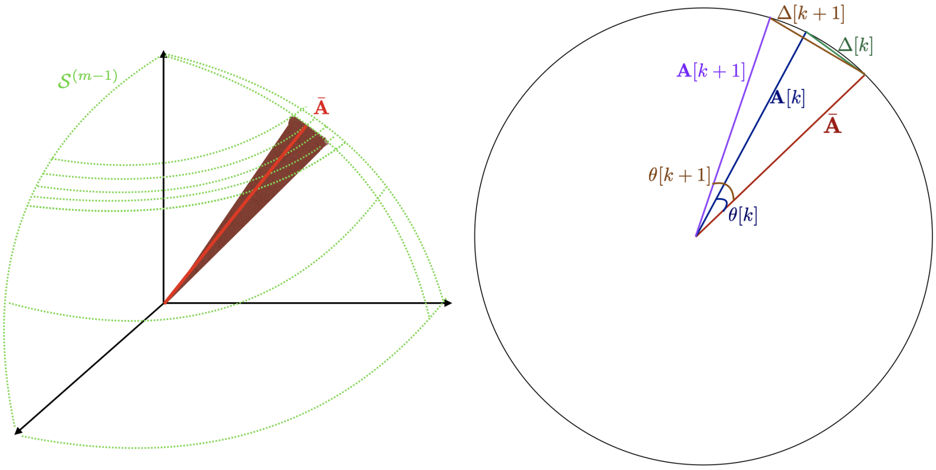

(a) The brown space represents the spectral variabilities distribution.

(b) Projection of $\bar{\mathbf{A}}$ and mixing vectors of one source at two samples k and $k+1$ onto a slice of the hypersphere.

Figure 4.4: Geometrical interpretation of spectral variabilities as angular variabilities

The figure 4.4 shows the parallel between the PLMM in remote sensing framework and in astrophysics framework. We can observe that, under this latter, the SV is strongly related to the angle between the reference spectrum and the one affected by the perturbation.

Indeed, as it has been noticed in Chenot [2017] we have

$$\|\Delta^i[k]\|_2 = 2\sin(\frac{\theta^i[k]}{2}) \simeq \theta^i[k] \tag{4.18}$$

since the variability energy is supposed to be small.

The oblique constraint sheds new light on the interpretation of spectral variabilities: it can be seen as the angular variation between the reference matrix and the mixing matrices.

This interesting property motivates our choice to inspire from PLMM to tackle spectral variabilities in X-ray astrophysics (C.F. chapter 5).

Problem formulation

Returning to the method proposed in Thouvenin et al. [2015], the authors propose to solve the following constrained optimization problem which brings into play penalization functions aiming at raising model indeterminacies:

$$\min_{\bar{\mathbf{E}},\ \mathbf{dE},\ \hat{\mathbf{A}}} \frac{1}{2}\left\|\mathbf{Y}-\bar{\mathbf{E}}\hat{\mathbf{A}}-\boldsymbol{\Delta}\right\|_2+\alpha f_{\hat{\mathbf{A}}}(\hat{\mathbf{A}})+\beta f_{\mathbf{S}}(\mathbf{S})+\gamma f_{\mathbf{dM}}(\mathbf{dM}). \tag{4.19}$$

where (α,β,γ) are scalars that control the trade-off between the data-fidelity term and the penalization functions $(f_{\hat{\mathbf{A}}}, f_{\mathbf{S}}, f_{\mathbf{dM}})$ required because of the ill-posedness of the problem. These penalization functions are chosen to be separable (in terms of pixels).

- As for any remote sensing dedicated approach, the ASC and ANC are enforced on the abundances, the endmembers and the spectral variabilities as well.
- The abundances are supposed to be spatially smooth, which is accounted for by a minimization of the differences between each abundance and its four neighbors.
- The authors propose three types of penalization on the endmembers, all widely used in the remote sensing community: minimizing the volume of the simplex, minimizing the mutual distance between the endmembers (as an approximation to the volume simplex), or minimizing the distance with an a priori known spectral library. The reader may refer to 4.2.2 for more details on classical endmember penalizations.
- Finally, the energy of the SV is minimized since the authors assume that there is generally no available information on the variabilities.

Practical implementation:

The cost function (4.19) is then minimized using a Block Coordinate Descent algorithm (BCD) and each block is minimized using an Alternating Direction Method Descent algorithm (ADMM) Boyd et al. [2011]). One can notice that the authors have not engaged a discussion and, a fortiori, an automatisation of the choice of the hyperparameters (α,β,γ).They are selected by cross-validation.

PLMM applied to medical imaging

We have mentioned briefly in 4.1 the principle of PET imaging analysis. In a nutshell, PET data consist of images acquired at different time-frames and each voxel TAC can be viewed as the weighted linear combination of elementary TACs that we aim at retrieving.

Therefore, the TAC at the $k-th$ voxel $x_k \in \mathbb{R}^m$, with m being the number of time-frames, can be written using the LMM as:

$$x_n = \sum_{i=1}^{n} m^i a_i[k]. \tag{4.20}$$

The TAC corresponding to the $i-th$ tissue is modeled as $m^i \in \mathbb{R}^m$ and is equivalent to spectral signature in remote sensing, and its proportion in the $k-th$ voxel is represented by the scalar $a_i[k]$ (equivalent to the abundances in remote sensing). As for remote sensing, data are supposed to live in the simplex and the classical ANC and ASC are assumed to be verified. The authors of Cavalcanti et al. [2018] have pinpointed that the TAC associated to specific binding is expected to be affected by spatially variant variabilities. Therefore, they have proposed a modification of (4.20) based upon PLMM. This novel model, coined specific binding linear mixture model (SLMM) can be formulated for each observed TAC as:

$$x_n = a_1[n](\bar{m}^1 + dm^1[k]) + \sum_{i=2}^{n} m^i a_i[k]. \tag{4.21}$$

In (4.21), specific binding is considered as the first component, its TAC is formulated as a reference TAC $\bar{m}^1 \in \mathbb{R}^m$ affected by a spatially variant perturbation $dm^1 \in \mathbb{R}^{m \times t}$. The SLMM unmixing method can be viewed as an adaptation of PLMM to the PET images.

The PLMM is very interesting to account for a wide range of spectral variabilities. However, as pinpointed by Hong et al. [2018], this modelling is unable to account for shape-invariant variabilities that would require scaling factors. Moreover, the use in PLMM of a squared Frobenius norm as penalization function for the variabilities implies that they follow a Gaussian distribution whereas the authors of Hong et al. [2018] consider that the true distribution of SV is more likely to be a mixed Gaussian.

4.4.3 The Augmented Linear Mixture Model

Variabilities affecting spectral signatures in remote sensing are of two types:

- shape-invariant variabilities due to variable illumination and atmospheric conditions that affect the endmembers by change of energy,
- amplitude-invariant variabilities caused by varying local temperatures or intrinsic variabilities of the material (as it would be the case for vegetation affected by the variation of chlorophyll concentration). In this scenario, it is the shape of the spectra that changes.

The Extended Linear Mixture Model models only shape-based variabilities and does not account for the second type of VS. While the Perturbed Linear Mixture Model aims at proposing a generalized formulation of the variabilities, the authors of Drumetz et al. [2016] have shown that it is unable to account for shape-based variabilities since they are not incoherent with the endmembers dictionary as it is induced in the PLMM.

Very recently, the authors of Hong et al. [2018] proposed a general framework aimed at encompassing both kinds of spectral variabilities encountered in remote sensing.

Building upon the previous works of Nascimento and Dias [2005a], Nascimento and Dias [2005b] and Drumetz et al. [2016], they represent shape-invariant SV with scaling factors on the spectral signatures. They restrain their modelling to SV due to changes of illumination that affect all spectral bands equally, leading to scaling factors independent of the endmembers.
So far, we have only considered the so-called intra-class variabilities, meaning the variabilities affecting a material independently of the others Somers et al. [2011], Zhang et al. [2006]. The authors of Hong et al. [2018] propose to account for inter-class variabilities as well, meaning the similarities between the endmembers of the materials. The novelty of their method lies in the introducing of a spectral variability dictionary that models both types of variabilities.

The proposed model, coined Augmented Linear Mixture Model (ALMM) can be written as

$$\mathbf{Y} = \bar{\mathbf{E}}\hat{\mathbf{A}}\mathbf{\Psi} + \mathbf{dE}\ \mathbf{B} + \mathbf{N}, \tag{4.22}$$

where

- the diagonal matrix $\mathbf{\Psi} \in \mathbb{R}^{t \times t}$ whose diagonal terms correspond to the scaling factors on each pixel. The choice of shared scaling factors for all endmembers is motivated, in addition to the physical considerations mentioned above, by the ill-posedness of the problem. Considering the same scaling factor allows to re-constraint the problem;
- The term $\bar{\mathbf{E}} \in \mathbb{R}^{m \times n}$ stands for the endmembers dictionary while $\hat{\mathbf{A}} \in \mathbb{R}^{n \times t}$ represents the abundances;
- The additive term $\mathbf{dE}\ \mathbf{B}$ is supposed to represent the SV affecting the shape of the spectral signatures. It is composed of the so-called spectral variabilities dictionary $\mathbf{dE} \in \mathbb{R}^{m \times L}$ where L is the number of basis vectors of $\mathbf{dE}$ and the coefficient matrix $\mathbf{B}$;

- Noise and residual errors are gathered in the matrix $\mathbf{N} \in \mathbb{R}^{m\times t}$.

The authors of Hong et al. [2018] have stressed out that the SV additive term is aimed at capturing not only (intra-class) SV but also deviations to the LMM (*e.g.* spectral signatures that have not been captured by the initial model).

This type of modelling shares similarities with the problem of robust Blind Source Separation. By denoting $\mathbf{O} \in \mathbb{R}^{m\times t}$ the outliers term used in robust BSS modelling, the problem can be written as

$$\mathbf{Y} = \bar{\mathbf{E}}\hat{\mathbf{A}} + \mathbf{O} + \mathbf{N}.$$

If we consider that the outliers admit a certain representation in a domain $\mathbf{\Phi_O} \in \mathbb{R}^{d_o\times t}$ as it is customary in sparse BSS (in that case, the representation is sparse, Chenot [2017]), the previous equation becomes:

$$\mathbf{Y} = \bar{\mathbf{E}}\hat{\mathbf{A}} + \alpha_{\mathbf{O}}\mathbf{\Phi_O} + \mathbf{N},$$

where $\alpha_{\mathbf{O}} \in \mathbb{R}^{m\times d_o}$ is the matrix containing the expansion coefficients of $\mathbf{O}$ in $\mathbf{\Phi_O}$. The correspondence between the product $\alpha_{\mathbf{O}}\mathbf{\Phi_O}$ with the SV additive term of ALMM $\mathbf{dE}\ \mathbf{B}$ is straightforward.

Problem formulation

The method adopted by Hong et al. [2018] is a data-driven dictionary learning method that can be expressed in terms of optimization problem:

$$\min_{\hat{\mathbf{A}},\mathbf{B},\mathbf{\Psi},\mathbf{dE}} \frac{1}{2}\left\|\mathbf{Y} - \bar{\mathbf{E}}\hat{\mathbf{A}}\mathbf{\Psi} + \mathbf{dE}\ \mathbf{B}\right\|_2 + \alpha f_{\hat{\mathbf{A}}}(\hat{\mathbf{A}}) + \beta f_{\mathbf{B}}(\mathbf{B}) + \gamma f_{\mathbf{dE}}(\mathbf{dE}), \text{s.t.} \hat{\mathbf{A}} \geq 0, \mathbf{\Psi} \geq 0. \tag{4.23}$$

- The endmember dictionary $\bar{\mathbf{E}}$ is supposed known a priori;
- $f_{\hat{\mathbf{A}}}(\hat{\mathbf{A}})$ is the regularization function with respect to the abundances, it consists of promoting spatial sparsity of the abundances map since the authors suppose that few materials are observed in one pixel in terrestrial imaging. The abundances are also assumed positive;
- The only assumption made on the scaling factors $\mathbf{\Psi}$ is their positivity;
- The key idea of Hong et al. [2018] is to suppose the low-coherency between the endmembers dictionary and the SV dictionary. This assumption is enforced in the penalization function $f_{\mathbf{dE}}$ through the term $\frac{1}{2}\|\bar{\mathbf{E}}^T\mathbf{dE}\|_F^2$. In addition to that, the authors impose the basis vectors of the SV dictionary to be orthogonal;
- Finally, the regularization term $f_{\mathbf{B}}$ consists simply of a penalization of the energy of the coefficients given the lack of priors on this type of variabilities.

Practical implementation:

According to the available prior information, the authors propose two multi-block ADMM-based algorithm.

The first one solves the spectral unmixing problem resulting from (4.23) provided the SV dictionary (in addition to the endmember dictionary). The second method solves exactly the optimization problem formulated in (4.23): it performs the SV dictionary learning simultaneously with spectral unmixing. No study on the selection of the hyperparameters has been undertaken, which may hinder the performances of the algorithm.

ALMM is a method aiming at providing a general framework for spectral variabilities that are met in remote sensing. However, it introduces a variabilities dictionary on which there is very few information which makes the problem very ill-posed, requiring therefore the knowledge of at least the endmembers matrix. Moreover, it does not specifically focus on the intra-class VS whose knowledge would require a supplementary step.

On the opposite to bundles or dictionary aided methods, PLMM, ELMM and ALMM have the advantage of explicitly accounting for variabilities and proposing unmixing approaches that are mathematically sound. However, the proposed modelling cannot be applied to astrophysical data that interest us in this thesis.

4.5 Accounting for spectral variabilities in astrophysics

Almost all methods explicitly modeling SV are found in remote sensing framework. However, in astrophysical applications which are not exempt from spectral variabilities, only one BSS method has dealt with this issue to the best of our knowledge, namely the local GMCA (or L-GMCA Bobin et al. [2013]).

Many missions dedicated to the study of the CMB have been launched over the last decades including the Planck Mission. The CMB data brought by this survey consist of observations of the full sky at different frequencies where the CMB power spectrum is known to be present in addition to other important astrophysical components. In the literature of BSS applied to Planck observations (Bobin et al. [2008]), data are generally modeled using the LMM:

$$\mathbf{X} = \mathbf{A}\mathbf{S} + \mathbf{N}$$

where $\mathbf{X}$ is the data matrix, $\mathbf{A}$ the mixing matrix, $\mathbf{S}$ the source matrix and $\mathbf{N}$ stands for the noise. This is an ill-posed problem for which the authors of Bobin et al. [2008] propose to exploit the morphology of the astrophysical components present in the mixture. Without going into details, we can say that they present specific structures that can be captured by the wavelet domain in which they have a sparse representation. Therefore, disentangling the CMB contribution from the galatic foregrounds (*i.e.* the remaining sources present in CMB data) can be viewed as a sparse BSS problem.

However, this Linear Mixture Model is an oversimplification of the reality. It is far from optimal for two principal reasons:

- All observations do no have the same resolution and the beams can be anisotropic and spatially variant;
- Most of the sources present in CMB data are known to have varying emissivities across the sky. For example, dust emission presents significant variations across the sky. We are in presence of **spectral variabilities**.

The first issue, on which we shall not linger in this chapter, is tackled in Bobin et al. [2013] using a multiscale GMCA: it consists of adapting the wavelet decomposition such that the observations share the same resolution at a given wavelet scale. To deal with spectral variabilities, the authors consider that the LMM still stands locally *i.e.* data observed on small patches can be described by linear combination of the sources. The method consists of running, for each location k, GMCA algorithm at different scales μ on patches of different sizes L_μ. Therefore a set of mixing matrices $\mathbf{A}^{\mu}_{L_\mu}[k]$ is assigned to each location k. The "optimal" patchsize is then selected through an automatic strategy described in Bobin et al. [2013]. This spatial adaptive strategy allows to have different patches accounting for the local content of the observed scene: small patches for high fluctuations and large patches for stationary emissivities. L-GMCA has proven itself to be very efficient for the Planck mission Bobin et al. [2013]. However, it does not propose an explicit modelling of spectral variabilities neither it exploits their spatial structure.

Conclusion

The wide majority of Component Separation techniques tackling the problem of spectral variabilities find their roots in remote sensing. Unfortunately, the assumptions on which rely these methods make them inadequate to astrophysical applications that are of interest for us in this thesis. Therefore, we propose in the next chapter a new BSS method that explicitly accounts for spectral variabilities and that can be applied in various application fields among which high-energy astrophysics.

Chapter 5

Sparse Blind Source Separation in the presence of spectral variabilities

The spectral variabilities phenomenon is ubiquitous when modeling natural scenes and should be taken into account when performing component separation. This raises the question of finding, on the one hand, a model as realistic as possible that fully accounts for variabilities, and on the other hand, a method capable of estimating both the source matrix and the sample-dependent mixing matrices. This chapter is aimed at introducing a novel modeling of spectral variabilities in sparse BSS, discussed in 5.1. Once the model established, we present the optimization problem in 5.2 and the chosen minimization scheme in 5.3. We then propose in 5.4 an empirical study of the detection and identifiability of the variabilities and we end this chapter by an application of the algorithm to realistic X-ray astrophysical data in 5.5.

5.1 Modelling spectral variabilities in sparse BSS

5.1.1 Ill-posedness of the problem

A natural formalism to account for spectral variabilities would be to model the contribution of each component as the product of the mixing matrix at the position k with the corresponding source. The $k-th$ sample of the data matrix can thus be written as

$$\mathbf{X}[k] = \mathbf{A}[k]\mathbf{S}[k] = \sum_{i=1}^{n} \mathbf{A}^i[k]\mathbf{S}_i[k], \tag{5.1}$$

where $\mathbf{X}[k] \in \mathbb{R}^m$ corresponds to the $k-th$ column of the data matrix, $\mathbf{A}^i[k] \in \mathbb{R}^m$ is the column of the mixing matrix observed at the sample k for the $i-th$ source and $\mathbf{S}_i[k]$ is the sample of the source matrix at the $i-th$ line and the $k-th$ column.

This problem is particularly ill-posed in comparison with the classical modelling of data without variabilities. For illustrative purposes, let us consider observations obtained from the mixture of two sources under two scenarios:

- First scenario (figure 5.1a): the mixing matrix is invariant, the generation process can be modeled by the classical Linear Mixture Model. Each element of the $j-th$ line of the data matrix can be written as

  $$\forall k = 1..t,\ \mathbf{X}^j[k] = \sum_{i=1}^{n} \mathbf{A}^i_j \mathbf{S}_i[k],$$

  involving the same elements of the mixing matrix, namely the spectra of the mixture components at the $j-th$ observation $\mathbf{A}^i_j$. Therefore, *the information on the $j-th$ observation of the mixing matrix $\mathbf{A}_j$ is provided by t samples of the data matrix.*

- Second scenario: (figure 5.1b): We suppose data affected by spectral variabilities, therefore the mixing matrix varies across the samples. Unlike the first case, *each element of the mixing matrix $\mathbf{A}^i_j[k]$ is estimated based on only one observation of the data matrix* :

  $$\forall k = 1..t, \mathbf{X}^j[k] = \sum_{i=1}^{n} \mathbf{A}^i_j[k]\mathbf{S}_i[k].$$

As illustrated in figure 5.2, even if we consider the source matrix known, there is an infinite number of admissible solutions. The problem remains very ill-posed.

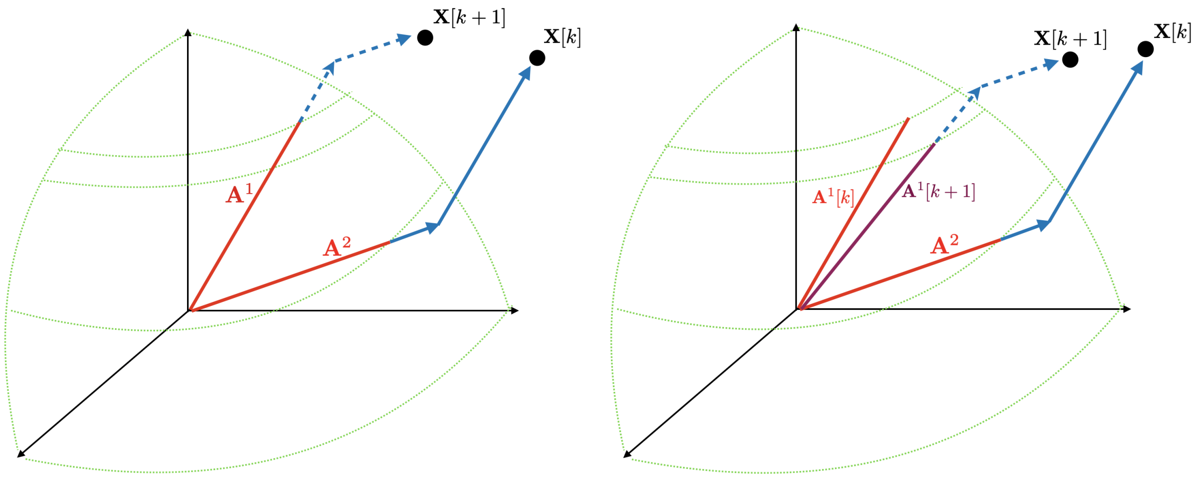

(a) Scenario 1: invariant mixing matrix

(b) Scenario 2: Spatially variant mixing matrix

Figure 5.1: Illustration of mixture model with (a) invariant mixing directions and (b) in the presence of spectral variabilities for the first source.

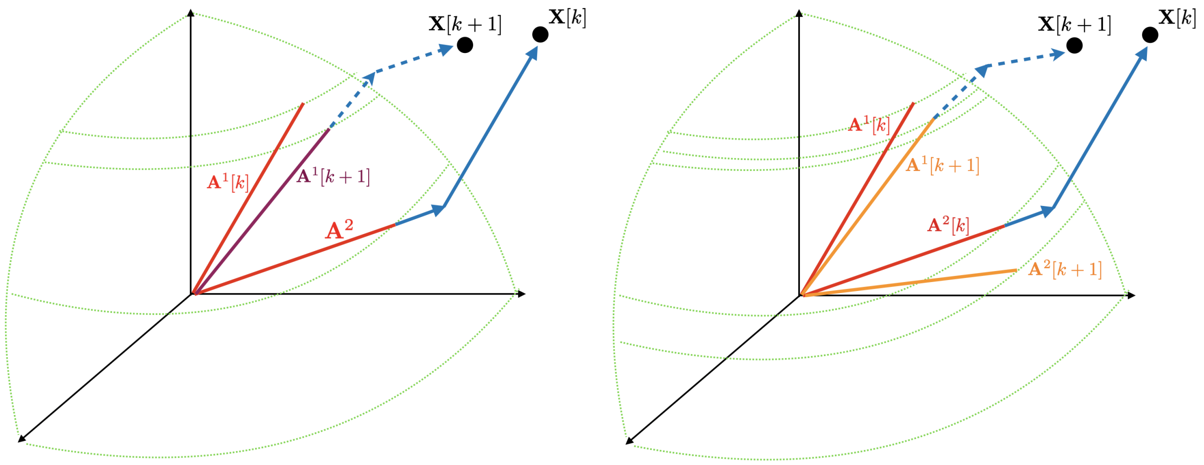

(a) Source 1

(b) Source 2

Figure 5.2: Illustration of the degeneracy of the SV recovery problem. With the same data point $X[k+1]$ and the same source contribution, different scenarios are possible: (a) only the first source is perturbed by VS, (b) both of the sources are affected by VS.

Solving the problem formulated in eq. (5.1) requires additional constraints on the spectral variabilities to raise its indeterminacies.

5.1.2 Linearization and regularization of the spectral variabilities

Regularization of the problem

Recovering sample-dependent mixing matrices from the data and the source matrix amounts to solving

$$\min_{\mathbf{A}\in OB(m)} \sum_{k=1}^{t} \left\| X[k] - \sum_{i=1}^{n} \mathbf{A}^i[k]\mathbf{S}_i[k] \right\|_F^2 \tag{5.2}$$

where the only available prior is the belonging of the columns of$\mathbf{A}$ to the hypersphere. As seen in geometrical examples of the figure 5.1b, this constraint is not sufficient for solving eq. (5.2).

Two classes of regularizations can be adopted to constrain the problem:

- *Regularization on spatial information:* This first type of regularization would consist of exploiting the spatial smoothness of the variations. This is physically justified in X-ray astrophysical applications (5.5).
- *Regularization on shape variation:* in this case, we consider that the spectral variabilities have a particular shape that is known in advance. This regularization is the most constraining: it completely raises degeneracy of the problem. However, it requires prior information that is not trivial.

In this chapter, we chose to focus on the first type of regularization.

Spatial smoothness of the variabilities can be recast mathematically as finding a transformed domain (*e.g.* wavelets) in which the field of sample-dependent mixing matrices $\{\mathbf{A}[k] \in \mathbb{R}^{m\times n}\}_{\{k=1..t\}}$ has a sparse representation, imposing a constraint that drastically restrains the range of admissible solutions. However, applying a transform to the mixing matrices raises other problems since the transform is nonlinear.

An interesting alternative would be to **linearize** the variabilities which allows to isolate the SV term and enforce spatial regularity constraint on it. This is the approach adopted in this chapter.

Linearization of the variabilities based on the PLMM approach

Based on the discussion undertaken in chapter 4 and inspired by Thouvenin et al. [2015], we propose to model the sample-dependent mixing matrices as a reference vector affected by very small perturbations.

Therefore, we can decompose, for each contribution i, the corresponding mixing matrix $\mathbf{A}^i \in \mathbb{R}^{m\times t}$ as the sum of a reference matrix spatially invariant $\bar{\mathbf{A}}^i \in \mathbb{R}^{m\times t}$ and a SV matrix $\mathbf{\Delta}^i \in \mathbb{R}^{m\times t}$:

$$\begin{cases} \forall i = 1..t,\ \mathbf{A}^i = \bar{\mathbf{A}}^i + \mathbf{\Delta}^i, \\ \|\mathbf{\Delta}^i\|_2 \ll \left\|\bar{\mathbf{A}}^i\right\|_2 \\ \forall i = 1..n, \forall k = 1...t,\ \bar{\mathbf{A}}^i \in \mathcal{S}^{(m-1)} \text{ and } \mathbf{A}^i[k] \in \mathcal{S}^{(m-1)} \end{cases}$$

As said in chapter 4, due to the hypersphere constraint, the spectral variability is tantamount to an **angular variability** *i.e.*

$$\|\mathbf{\Delta}^i[k]\|_2 = 2\sin(\frac{\theta^i[k]}{2}) \simeq \theta^i[k] \tag{5.3}$$

assuming the variability energy is small enough.

Regularity constraint on the angular variability

We suppose in the rest of this thesis that the SV exhibit spatially smooth patterns. This assumption boils down to considering that the angular variability has a sparse representation in a well-chosen dictionary Ψ.

Therefore, the regularization constraint can be formulated with the sparsity-promoting ℓ_1 norm on the angular variability

$$\sum_{i=1}^{n}\sum_{k=1}^{t}\|\theta^i[k]\Psi_{\mathbf{A}}^T\|_1,$$

which can be rewritten by small angle approximation as the following penalization term

$$\sum_{i=1}^{n}\sum_{k=1}^{t}\|\|\mathbf{A}^i[k]-\bar{\mathbf{A}}^i\|_2\Psi_{\mathbf{A}}^T\|_1.$$

5.2 Problem formulation

The problem we tackle in this chapter aims at recovering the source matrix and the sample-dependent mixing matrices from the observations. It can be recast as the following optimization problem:

$$\min_{\mathbf{A}\in\mathbb{R}^{m\times t\times n},\mathbf{S}\in\mathbb{R}^{n\times t}}\mathcal{F}(\mathbf{A},\mathbf{S})+\mathcal{G}(\mathbf{A})+\mathcal{H}(\mathbf{S}),\tag{5.4}$$

where $\forall k=1..t,\ \mathbf{A}[k]\in\mathbb{R}^{m\times n}$ denotes the mixing matrix observable at the $k-th$ sample, $\mathbf{S}$ stands for the source matrix, $\mathcal{F}$ represents the data-fidelity term and $\mathcal{G}$ (respectively $\mathcal{H}$) is the penalization function applied to the mixing matrices (respectively the source matrix).

5.2.1 The data-fidelity term

We assume that the noise is additive Gaussian and that the Linear Mixture Model holds locally *i.e.* each column of the data matrix can be written as

$$\forall k=1..t,\mathbf{X}[k]=\sum_{i=1}^{n}\mathbf{A}^i[k]\mathbf{S}_i[k]+\mathbf{N}[k],$$

where $\mathbf{N}[k]$ is the $k-th$ column of the matrix standing for noise and model imperfections.

Therefore, we assign to the data-fidelity term $\mathcal{F}$ the distance at each sample between the observations and the local mixture model:

$$\mathcal{F}(\mathbf{A},\mathbf{S})=\sum_{k=1}^{t}\|\mathbf{X}[k]-\sum_{i=1}^{n}\mathbf{A}^i[k]\mathbf{S}_i[k]\|_F^2.\tag{5.5}$$

5.2.2 Mixing matrix penalization

We choose to consider three constraints on the mixing matrix:

- It is customary in sparse BSS problems to impose the oblique constraint on $\mathbf{A}$. We remind that it permits to avoid scaling indeterminacy.

  $$\mathcal{G}_1(\mathbf{A})=\iota_{\|.\|_2=1}(\mathbf{A})$$

  The function $\mathcal{G}_1$ corresponds to the characteristic function of the hypersphere $\mathcal{S}^{(m-1)}$. It is convex except for the case $\mathbf{A}=\mathbf{0}$, which is never met in practice.
  To obtain convexity on every point, we can relax this constraint by imposing the norm of the columns to be inferior or equal to 1. Its minimum coincides with the minimum of $\mathcal{G}_1$: if the couple $(\mathbf{A},\ \mathbf{S})$ is the solution of normalized columns, then $(\alpha\mathbf{A},\frac{1}{\alpha}\mathbf{S})$, with $\alpha\in(0,1]$ may also be a solution since the mixing matrix belongs to the unit ball. However, since we impose sparsity on the sources, the optimal α tends to be the highest possible, *i.e.* $\alpha=1$. In the rest of this chapter, we will keep $\mathcal{G}_1$ and suppose its convexity.

- The model we propose in this chapter is intended to be applicable for various applications. We can notice that in the application fields for variabilities mentioned in chapter 4, the mixing matrix either stands for spectral signatures of components in an observed scene Thouvenin et al. [2015], or spectra of mixed elements in extended sources such as supernovae remnants (e.g. Chandra data) or time-activity curves of biological tissues (Cavalcanti et al. [2018]). Given these physical interpretations of $\mathbf{A}$, we can legitimately consider that all its coefficients are positive, which is mathematically formulated as the convex characteristic function of the positive set
$$\mathcal{G}_2(\mathbf{A}) = \iota_{.\geq 0}(\mathbf{A}).$$
- **Spectral variabilities:** As stated in 5.1, we impose on the variabilities to be smooth which, in our case boils down to a constraint on the angular variabilities. We impose therefore for each source i the sparsity of θ_i on a transformed domain $\Psi_{\mathbf{A}} \in \mathbb{R}^{T_A \times t}$. Using the approximation (5.3), the penalization function can be written as:
$$\mathcal{H}(\mathbf{A}) = \sum_{i=1}^{n} \mathbf{R}_{\mathbf{A}}[i] \odot \left\| \left\| \mathbf{A}^i - \bar{\mathbf{A}}^i \right\|_2 \Psi_{\mathbf{A}}^T \right\|_1 ,$$
which is a convex function. We approximate it in the following by the $\ell_{2,1}$ norm of the transform of the SV, since this latter has an explicit proximal operator. However, this penalization does not have an explicit operator. Therefore, we propose to approximate it by the $\ell_{2,1}$ norm in the transformed domain $\Psi_{\mathbf{A}}$:
$$\mathcal{G}_3(\mathbf{A}) = \sum_{i=1}^{n} \mathbf{R}_{\mathbf{A}}[i] \odot \left\| (\mathbf{A}^i - \bar{\mathbf{A}}^i) \Psi_{\mathbf{A}}^T \right\|_{2,1} ,$$
with $\odot$ being the Hadamard product and $\mathbf{R}_{\mathbf{A}} \in \mathbb{R}^{n \times T_A}$ the regularization matrix.

The total penalization function on the mixing matrix can then be written as
$$\mathcal{G}(\mathbf{A}) = \iota_{\|\cdot\|_2=1}(\mathbf{A}) + \iota_{.\geq 0}(\mathbf{A}) + \sum_{i=1}^{n} \mathbf{R}_{\mathbf{A}}[i] \odot \left\| (\mathbf{A}^i - \bar{\mathbf{A}}^i) \Psi_{\mathbf{A}}^T \right\|_{2,1} . \tag{5.6}$$

5.2.3 Sources penalization

We suppose that the sources are sparse in a well-chosen domain $\Phi_{\mathbf{S}} \in \mathbb{R}^{T_{\mathbf{S}} \times t}$ and that the Morphological Diversity Principle is respected: the different contributions of the sources may have the same sparsity domain but they do not share the same highest-amplitude coefficients. This principle permits an efficient disentanglement, which is a key element for performing sparse blind source separation.

Strictly speaking, sparsity is enforced with ℓ_0 quasi-norm, that limits the number of non-null coefficients. However, we choose to work with its convexified form, the ℓ_1 norm, easier to minimize (Donoho and Huo [2001]). The corresponding regularization function is:
$$\mathcal{F}(\mathbf{S}) = \left\| \mathbf{R}_{\mathbf{S}} \odot \mathbf{S} \Phi_{\mathbf{S}}^T \right\|_1 .$$
$\mathbf{R}_{\mathbf{S}} \in \mathbb{R}^{n \times T_{\mathbf{S}}}$ contains the regularization parameters and $\odot$ stands for the Hadamard product.

5.2.4 Overall cost function

Finally, the optimization function we aim at minimizing can be written as
$$\min_{\mathbf{A},\mathbf{S}} \sum_{k=1}^{t} \|\mathbf{X}[k] - \sum_{i=1}^{n} \mathbf{A}^i[k] \mathbf{S}_i[k]\|_F^2 + \left\| \mathbf{R}_{\mathbf{S}} \odot \mathbf{S} \Phi_{\mathbf{S}}^T \right\|_1 + \sum_{i=1}^{n} \mathbf{R}_{\mathbf{A}}[i] \odot \left\| (\mathbf{A}^i - \bar{\mathbf{A}}^i) \Psi_{\mathbf{A}}^T \right\|_{2,1} + \iota_{.\geq 0}(\mathbf{A}) + \iota_{\|\cdot\|_2=1}(\mathbf{A}). \tag{5.7}$$

5.3 Proposed approach

5.3.1 Minimization strategy

The optimization problem associated with eq.(5.7) is multi-convex: it is globally non-convex but the minimization sub-problems towards each of the variables $\mathbf{A}$ and $\mathbf{S}$ are convex. Several minimization strategies have flourished during last years to solve such multi-convex problems or approximations of them. We have carried out discussions about some of them in chapter 2.

Given the complexity of the problem (5.7) and the necessity of having a reliable method that avoids error propagation, we have chosen to use BCD algorithm (Tseng [2001]) since it proposes an exact minimization of each sub-problem and it is mathematically provably convergent.

PALM algorithm (Bolte et al. [2014]) could also have been used since its convergence is mathematically motivated. However, it does not minimize the sub-problems but an approximation of them which is not desirable for such a complex cost-function.

5.3.2 Estimation of the source matrix

The minimization of eq. (5.7) towards the source matrix amounts to solving:

$$\min_{\mathbf{S}} \sum_{k=1}^{t} \|\mathbf{X}[k] - \sum_{i=1}^{n} \mathbf{A}^i[k]\mathbf{S}_i[k]\|_F^2 + \left\|\mathbf{R_S} \odot \mathbf{S}\Phi_{\mathbf{S}}^T\right\|_1. \tag{5.8}$$

The equation (5.8) can be solved efficiently through an Iterative Soft-Threshold Algorithm (ISTA proposed by Parikh et al. [2014]). The scheme of implementation to our problem follows the steps mentioned below:

1. Choice of the step-size: The step-size τ_s is chosen as $\tau_s = 1/\max_{k=1..t}\|\mathbf{A}[k]^T\mathbf{A}[k]\|_2$ in order to ensure convergence over all the samples.
2. Gradient descent step: The data-fidelity term being separable over the samples, we compute the gradient descent term $\tilde{\mathbf{S}}$ as:
$$\forall k = 1..t,\ \tilde{\mathbf{S}}[k] = \mathbf{S}[k] + \tau_{\mathbf{S}}\mathbf{A}[k]^T(\mathbf{X}[k] - \sum_{i=1}^{n}\mathbf{A}^i[k]\mathbf{S}_i[k]).$$
3. Proximal step: It consists of the soft-thresholding of $\alpha_{\tilde{\mathbf{S}}} \in \mathbb{R}^{n\times T_S}$ which is the transform of the gradient descent term $\tilde{\mathbf{S}}$ in $\Phi_{\mathbf{S}} \in \mathbb{R}^{T_S\times t}$.

The matrix $\mathbf{R_S} \in \mathbb{R}^{n\times T_S}$ contains the regularization parameters that control the trade-off between the data-fidelity term and the sparsity constraint.
We define it as

$$\mathbf{R_S} = \mathrm{Diag}(\lambda_1, .., \lambda_n)\mathbf{W_S},$$

where λ_i stands for the threshold associated to the $i-th$ source and the matrix $\mathbf{W_S} \in \mathbb{R}^{n\times T_S}$ contains the reweighting coefficients that help reducing bias errors.
For each source i, the individual thresholds are computed with the MAD operator associated to the gradient descent term:

$$\lambda_i = 1.48\kappa\mathrm{MAD}(\alpha_{\tilde{\mathbf{S}}_i}),$$

where $\kappa = 3$ since the noise is supposed additive Gaussian. The reweighting scheme is implemented using the current estimate of the sources:

$$\mathbf{W_S}_i = \frac{\epsilon}{\epsilon + \frac{\|\mathbf{S}_i\Phi_{\mathbf{S}}\|_2}{\|\mathbf{S}_i\Phi_{\mathbf{S}}\|_\infty}},$$

where $\epsilon = 10^{-3}$. The limitation of the biases, enabled by the reweighting scheme, is crucial for the estimation of the spectral variabilities as the experiments in appendix D will show.

We implement an accelerated version of the ISTA algorithm, namely the FISTA algorithm (Beck and Teboulle [2009]. The final updating scheme is provided in 13.

5.3.3 Estimation of the mixing matrix

The optimization problem towards the sample-dependent mixing matrices is:

$$\min_{\mathbf{A}} \sum_{k=1}^{t} \left\| \mathbf{X} - \sum_{i=1}^{n} \mathbf{A}^i[k]\mathbf{S}_i[k] \right\|_{\mathrm{F}}^2 + \sum_{i=1}^{n} \mathbf{R}_{\mathbf{A}}[i] \odot \left\| (\mathbf{A}^i - \bar{\mathbf{A}}^i)\Psi_{\mathbf{A}}^T \right\|_{2,1} + \iota_{.\geq 0}(\mathbf{A}) + \iota_{\|.\|_2=1}(\mathbf{A}), \tag{5.9}$$

It can be split as the sum of a differentiable function and three convex simple non-differentiable terms. All three terms have an explicit proximal operator but their sum does not admit a closed form solution. We can however notice that the sum $\iota_{.\geq 0}(\mathbf{A}) + \iota_{\|.\|_2=1}(\mathbf{A})$ admits an explicit proximal operator.

Two alternatives can be adopted to solve such an optimization problem:

- The Generalized Forward Backward algorithm (Raguet et al. [2013]): this approach proposes a minimization of the optimization function by implementing at each iteration the proximal operator of each constraint separately. However, this can be very time consuming, particularly since the constraints $\sum_{i=1}^n \mathbf{R}_{\mathbf{A}}[i] \odot \left\| (\mathbf{A}^i - \bar{\mathbf{A}}^i)\Psi_{\mathbf{A}}^T \right\|_{2,1}$ and $\iota_{.\geq 0}(\mathbf{A}) + \iota_{\|.\|_2=1}(\mathbf{A})$ project the solution of the gradient descent step onto two very distinct sets.

- We can also combine the classical FISTA algorithm with a subroutine computing the proximal operator of the sum of the non-differentiable terms. In our case, the functions g_1 and g_2 defined as $g_1 = \iota_{.\geq 0} + \iota_{\|.\|_2=1}$ and $g_2 = \sum_{i=1}^n \mathbf{R}_{\mathbf{A}}[i] \odot \left\| \ . \ \Psi_{\mathbf{A}}^T \right\|_{2,1}$ both have analytical proximal operators (the reader is invited to refer to appendix A for more information). We propose to use the Dykstra-like algorithm introduced in Bauschke and Combettes [2008] to have a fast and efficient computation of the proximal operator of $g_1 + g_2$.

Finally, the principal steps of the minimization scheme towards the sample-dependent mixing matrices consist of:

1. Choice of the step-size: As for the estimation of the sources, the step-size is chosen as the minimal step-size over all the samples $\tau_A = 1/\max(\mathbf{S}^2)$.

2. Gradient descent step: The gradient descent term $\tilde{\mathbf{A}}$ is computed as:

$$\forall i = 1..n, \ \tilde{\mathbf{A}}^i = \mathbf{A}^i + \tau_A(\mathbf{X} - \sum_{i=1}^{n} \mathbf{A}^i\mathbf{S}_i)\mathrm{Diag}(\mathbf{S}_i),$$

   where $\mathrm{Diag}(\mathbf{S}_i) \in \mathbb{R}^{t \times t}$ denotes the matrix whose diagonal elements are the coefficients of the vector $\mathbf{S}_i \in \mathbb{R}^t$.

3. Proximal step: It consists of the application of the proximal operator provided by Dykstra-like algorithm to the gradient-descent term. For each source i, the regularization parameter $\mathbf{R}_{\mathbf{A}}[i]$ applies to

$$\alpha_A = \| [\mathbf{A}^i - \bar{\mathbf{A}}^i + \tau_A(\mathbf{X} - \sum_{i=1}^{n} \mathbf{A}^i\mathbf{S}_i)\mathrm{Diag}(\mathbf{S}_i)]\Psi_{\mathbf{A}}^T \|_2.$$

Setting of regularization parameter $\mathbf{R}_{\mathbf{A}}$

Numerous sparse BSS algorithms such as GMCA (Bobin et al. [2007]) adopt a thresholding strategy that aims at ensuring robustness of the algorithm to noise. This is done through an explicit dependence of the threshold to noise statistics.
In the same vein, we propose a thresholding procedure than can be viewed as a detection problem: the threshold should be higher than noise level in order to set at zero noise entries, and lower than highest amplitude coefficients that are very likely to belong to the support of the variabilities. For now, we consider no reweighting procedure, *i.e.* $\mathbf{R}_{\mathbf{A}} = \mathrm{Diag}(\gamma_1, ..., \gamma_n)$.

We base the choice of the threshold on a fixed point argument: assuming that we have reached convergence for the mixing matrices and the source matrix (*i.e.* $\mathbf{S} = \mathbf{S}^{(\infty)}$ and $\mathbf{A} = \mathbf{A}^{(\infty)}$), the threshold of each source i applies to:

$$\begin{aligned}\alpha_A^{(\infty)} =& \|[\mathbf{A}^{i(\infty)} - \bar{\mathbf{A}}^{i(\infty)} + \tau_A(\mathbf{X} - \sum_{i=1}^{n} \mathbf{A}^{i(\infty)}\mathbf{S}_i^{(\infty)})\mathrm{Diag}(\mathbf{S}_i^{(\infty)})]\Psi_{\mathbf{A}}^T\|_2 \\ =& \|[\mathbf{A}^{i(\infty)} - \bar{\mathbf{A}}^{i(\infty)} + \tau_A\mathbf{N}\mathrm{Diag}(\mathbf{S}_i^{(\infty)})]\Psi_{\mathbf{A}}^T\|_2 \\ =& \|\mathbf{\Delta}^{i(\infty)}\Psi_{\mathbf{A}}^T + \tau_A\mathbf{N}\mathrm{Diag}(\mathbf{S}_i^{(\infty)})\Psi_{\mathbf{A}}^T\|_2,\end{aligned}$$

where $\mathbf{\Delta}^{i(\infty)}$ denotes the variabilities term we aim to recover and $\mathbf{N}\mathrm{Diag}(\mathbf{S}_i^{(\infty)})\Psi_{\mathbf{A}}^T$ with $\Psi_{\mathbf{A}} \in \mathbb{R}^{T_A \times t}$ corresponds to the back-projected noise contribution. Consequently, the threshold should be set higher than the back-projected noise level $\|\mathbf{N}\mathrm{Diag}(\mathbf{S}_i^{(\infty)})\Psi_{\mathbf{A}}^T\|_2$ and lower than the spectral variabilities contribution.

For each sample k, the coefficients of the column vector $\mathbf{N}\mathrm{Diag}(\mathbf{S}_i^{(\infty)})\Psi_{\mathbf{A}}^T)[k]$ are m independent and normally distributed random variables (the multiplication of the noise matrix by $\mathrm{Diag}(\mathbf{S}_i^{(\infty)})\Psi_{\mathbf{A}}^T$ creates a dependency between the lines but not the columns). Therefore, the $k{-}th$ element of the vector $\|\mathbf{N}\mathrm{Diag}(\mathbf{S}_i)\Psi_{\mathbf{A}}^T\|_2[k] \in \mathbb{R}^{T_A}$ is the realization of a χ distribution of m degrees of freedom derived from a normal distribution of standard deviation $\sigma_N\|\mathrm{Diag}(\mathbf{S}_i^{\infty})\Psi_{\mathbf{A}}^T[k]\|_2$ with σ_N the standard deviation of the noise $\mathbf{N}$.
Similarly to what has been proposed in GMCA algorithm, we can fix the threshold for each source i as

$$\gamma_i = \tau_A \mathrm{median}_{k=1..T_A}(\mathrm{Q}_{0.996}(\sigma_N\|Diag(\mathbf{S}_i^{\infty})[k]\|_2))) \tag{5.10}$$

with $\mathrm{Q}_p(Z)$ the quantile function associated to a central χ distribution of scale Z for a probability p. In this setting, the probability of an entry of amplitude higher than the threshold to belong to the noise is less than 0.4%.

However, this threshold setting presents three major drawbacks:

- it requires to know the source matrix when the algorithm converges before launching it, moreover it involves that the resulting source matrix $\mathbf{S}^{(\infty)}$ is close to the true one *i.e.* $\mathbf{S}^*$;
- it requires the knowledge of the variance of the noise $\mathbf{N}$;
- If we denote $\hat{\Delta}^i$ the estimated spectral variabilities matrix for the $i-th$ source, it can be decomposed as:

  $$\hat{\Delta}^i = \Delta^{*\ i} + \Delta^i_{\mathrm{noise}} + \Delta^i_{\mathrm{interf+artefacts}},$$

  where

  - $\Delta^{*\ i}$ is the true SV matrix,
  - $\Delta^i_{\mathrm{noise}}$ represents the errors due to noise contamination,
  - $\Delta^i_{\mathrm{interf+artefacts}}$ stand for the errors due to the interferences (related to incorrect unmixing of the sources) and artefacts (that can be caused by the propagation of errors on the source matrix estimation, C.F. appendix D).

The errors due to noise contamination can be handled by the threshold proposed in 5.10. Regarding the interferences and artefacts, they can be removed with the threshold 5.10 only if their energy is smaller than the noise energy, which is generally not the case as will show the experiments in section 5.4 and appendices B and C.

An alternative consists on using the Median Absolute Deviation (MAD) operator, a standard deviation estimator for Gaussian noise. We propose to implement it in the current estimate of α_A, which boils down to fixing the threshold at each iteration l as:

$$\gamma_i^{(l)} = \mathrm{Q}_{0.996}(1.5\mathrm{MAD}([\mathbf{A}^{i(l)} - \bar{\mathbf{A}}^{i(l)} + \tau_A(\mathbf{X} - \sum_{i=1}^{n} \mathbf{A}^{i(l)}\mathbf{S}_i^{(l)})\mathrm{Diag}(\mathbf{S}_i^{(l)})]\Psi_{\mathbf{A}}^T). \tag{5.11}$$

This approach is interesting since:

- it does not require the knowledge *a priori* of the noise statistics or the true source matrix,
- it allows to have high thresholds at the first iterations, enhancing the separation between the components and therefore reducing interferences;
- it captures the errors due to artefacts such as the propagation error resulting from biased sources since it is computed at each iteration on the current estimate of the SV matrix;
- at stabilization of the algorithm, it boils down to noise-dependent threshold as defined in eq. (5.10).

Refinement step

We propose to implement a reweighting procedure at the final iterations of the algorithm in order to avoid biases on the spectral variabilities. The final thresholds $\mathbf{R}_{\mathbf{A}} \in \mathbb{R}^{n \times T_A}$ can then be written as

$$\forall i = 1..n, \ \mathbf{R}_{\mathbf{A}}[i] = \gamma_i \times \frac{\epsilon}{\epsilon + \frac{\|(\mathbf{A}^i - \bar{\mathbf{A}}^i)\Psi_{\mathbf{A}}^T\|_2}{\|(\mathbf{A}^i - \bar{\mathbf{A}}^i)\Psi_{\mathbf{A}}^T\|_\infty}}.$$

The number of iterations before the refinement step and the number of iterations of the refinement step are provided in 5.4.3.

5.3.4 Initialization

Initialization is a key element for BCD-built algorithms. Indeed, even though theoretical convergence of the BCD algorithm is proven, it may converge towards a false minimum particularly in the case of our problem that is very ill-posed. State-of-the-Art methods tackling spectral variabilities problem in astrophysics are very few. We have discussed in chapter 4 the L-GMCA algorithm that has proven successful in the study of Planck data. We propose to initialize our algorithm with an adaptation of L-GMCA. It consists on applying the GMCA algorithm to patches of same size. The mixing matrix is then obtained by concatenating the mixing matrices recovered on each patch and filtering the results. The number of patches depends on the studied data. We explain the numerical settings used for our experiments in 5.4.3. The correction of order permutation is done using a "reference mixing matrix" that results from applying GMCA over all the samples. This technique allows to have a good first guess that initializes the BCD algorithm.

5.3.5 Setting the reference mixing matrix

Our model is based on the linearization of the variabilities which boils down to writing the sample-dependent mixing matrices as the sum of a reference matrix and a small perturbation that accounts for the variabilities $\forall i = 1..n, \ \mathbf{A}^i = \bar{\mathbf{A}}^i + \mathbf{\Delta}^i$. Therefore, the reference mixing matrix should be wisely chosen in order to enable the angular approximation.

Given that L-GMCA allows for a good initialization, we propose to set the reference mixing matrix for each source as the Fréchet mean of the initialization matrix $A^{i(0)}$.

5.3.6 Minimization scheme

The minimization strategy, coined svGMCA algorithm is provided in 13.

Algorithm 13: svGMCA algorithm

```
Initialization: $\mathbf{A}^{(0)} \leftarrow \mathbf{A}_{\textbf{L-GMCA}}$ , $\mathbf{S}^{(0)} \leftarrow \mathbf{S}_{\textbf{L-GMCA}}$
while not converge do
    • Update of $\mathbf{S}^{(k)}$ assuming $\mathbf{A}^{(k)}$ is fixed
    Initialization: $\mathbf{S}^{(0)} \leftarrow \mathbf{S}^{(k)}$, $t_1 = 1$, $\mathbf{y}^{(1)} \leftarrow \mathbf{S}^{(0)}$
    Compute $\tau_{\mathbf{S}} = 1/\max_{k \leq t} \left\| \mathbf{A}^T[k]\mathbf{A}[k] \right\|_2$
    while not converge do
        $\tilde{\mathbf{y}} \leftarrow \mathbf{y}^{(j-1)} + \tau_{\mathbf{S}} \mathbf{A}^{(k)} \otimes$[a]$(\mathbf{X} - \sum_{i=1}^{n} \mathbf{A}^i \mathbf{y}_i^{(j-1)})$
        $\mathbf{y}^{(j)} \leftarrow \mathcal{S}_{\tau_S \mathbf{R}_\mathbf{S}}(\tilde{\mathbf{y}} \Psi_{\mathbf{S}}^T) \Psi_{\mathbf{S}}$
        $t_{j+1} \leftarrow \frac{1+\sqrt{1+4t_j^2}}{2}$
        $\mathbf{S}^{(j+1)} \leftarrow \mathbf{y}^{(j)} + (\frac{t_j - 1}{t_{j+1}})(\mathbf{y}^{(j)} - \mathbf{y}^{(j-1)})$
        $j \leftarrow j + 1$
    end
    $\mathbf{S}^{(k+1)} \leftarrow \mathbf{S}^{(j)}$
    • Update of $\bar{\mathbf{A}}^{(k)}$ assuming $\mathbf{S}^{(k+1)}$ and $\mathbf{\Delta}^{(k)}$ is fixed
    • Update of $\mathbf{A}^{(k)}$ assuming $\mathbf{S}^{(k+1)}$ and $\bar{\mathbf{A}}^{(k+1)}$ is fixed
    Initialization: $\mathbf{A}^{(0)} \leftarrow \mathbf{A}^{(k)}$, $t_1 = 1$, $\mathbf{x}^{(1)} \leftarrow \mathbf{A}^{(0)}$
    Compute $\tau_{\mathbf{A}} = 1/\max(\mathbf{S}^{(k+1)})$
    while not converge do
        $\tilde{\mathbf{x}} = \mathbf{x}^{(k-1)} + \tau_{\mathbf{A}}(\mathbf{X} - \sum_{i=1}^{n} \mathbf{x}^{i(k-1)} \mathbf{S}_i^{(k+1)}) \circledcirc$[b]$\mathbf{S}^{(k+1)}$
        $\mathbf{x}^{(j)} \leftarrow \mathrm{prox}_{\mathrm{Dykstra}} \tilde{\mathbf{x}}$[c]
        $t_{j+1} \leftarrow \frac{1+\sqrt{1+4t_j^2}}{2}$
        $\mathbf{A}^{(j+1)} \leftarrow \mathbf{x}^{(j)} + (\frac{t_j - 1}{t_{j+1}})(\mathbf{x}^{(j)} - \mathbf{x}^{(j-1)})$
        $j \leftarrow j + 1$
    end
end
```

---

[a] We define the product $\otimes$ as $Z = X \otimes Y$ with $X \in \mathbb{R}^{m \times t \times n}$, $Y \in \mathbb{R}^{m \times t}$, $Z \in \mathbb{R}^{n \times t}$ and $\forall k = 1..t,\ Z^k = X^{kT} Y^k$

[b] We define the product $\circledcirc$ as $Z = X \circledcirc Y$ with $X \in \mathbb{R}^{m \times t}$, $Y \in \mathbb{R}^{n \times t}$, $Z \in \mathbb{R}^{m \times t \times n}$ and $\forall i = 1..n,\ Z^i = XY_i$

[c] This step consists of the application of the proximal operator computed by Dykstra-like algorithm as explained in 5.3.3.

5.4 Numerical experiments on synthetic data

The proposed algorithm svGMCA is aimed at tackling the problem of Blind Source Separation in the presence of spectral variabilities. In this section, we propose to understand the behaviour of svGMCA by assessing both qualitatively and quantitatively its limitations and performances for varying degrees of complexity of datasets.

5.4.1 Data description

To conduct a thorough investigation, we chose to work on data we simulate ourselves in order to have a complete control on their generation process and the involved parameters on the one hand, and to assess the performances through metrics computed on ground truth components on the other hand.

We conduct the study on two types of configuration for the sources and two types of configuration for the variabilities, that do not have necessarily physical interpretation but are very well suited to infer the impact of certain parameters.
The considered datasets in our experiments are the following:

1. Spectral variabilities sparse in DCT domain, this configuration allows us to have a complete control on crucial parameters such as the frequency of the variabilities or their sparsity degrees.

2. Piecewise constant spectral variabilities simulating variabilities that can be observed in remote sensing for example.

3. Sources sparse in the direct domain, it allows us to see the impact of the source distribution on the identification and recovery of the variabilities.

4. Source distribution sparse in starlet domain, which enables to see the effect of morphological diversity between SV and sources sparsity domain.

The experiments are carried out on simplistic datasets that allow us to estimate the impact of different parameters of the SV (*e.g.* their frequency, their amplitude...), of the sources (*e.g.* their sparsity degree) and of the mixing matrix (*e.g.* the shape of the mixed spectra). Because of the vast number of undertaken tests and for the sake of simplicity, we restrain ourselves in this part to 1D simulations $n = 2$ sources (unless otherwise stated), each one having $t = 1500$ samples.

5.4.2 Data generation

Synthetic data generation is detailed below.

The sources: They are drawn according to a Bernoulli-Gaussian distribution of activation parameter $\rho_{\mathbf{S}}$ with a non-negativity constraint. The amplitudes of the generated peaks have a variance of 1.
If the source distribution is chosen to be sparse in starlet domain, we convolve the sources described above with a Laplacian kernel of width at half maximum denoted $w_{\mathbf{S}}$. By way of analogy, the Laplacian kernel convolution models the instrumental spectral resolution of NMR imagery.
In the first configuration (no convolution), the sources present an exactly sparse distribution in the direct domain whereas in the second one (convolution with a Laplacian kernel), the starlet domain Starck et al. [2015] is more suited to enhance their sparse feature.

The mixing matrix: We choose to work with a toy example inspired by Chandra observations. Therefore, each of the column of the mixing matrix is representative of a physical spectrum acquired over $m = 5$ frequencies. The spectra we used are two theoretical spectra observed on supernovae remnants such as Cassiopeia A: the first one corresponds to an iron line emission: it is Gaussian-shaped with a width of 0.1, its central position is expected to vary through the line of sight, causing spectral variabilities. The second component is the synchrotron continuum emission that presents a power law shape of varying index.
The figures 5.4 and 5.5 show spectra as the ones used in the experiments for varying positions (in the case of emission lines) and indices (in the case of power laws).

The spectral variabilities The spectral variabilities are generated by the variation of the position of the Gaussian-shaped line in the case of emission line, and the index of the power law in the case of synchrotron emission. Depending on the experiment, the morphology of the SV can be of two kinds.

1. SV sparse in DCT domain, we generate in the DCT domain a vector of size $t = 1500$ from a Bernoulli-Gaussian distribution of activation parameter ρ_{VS} that we multiply with a Gaussian kernel of standard deviation σ_{VS} centered around a central frequency denoted f. The amplitude of the SV is fixed to L. Therefore, the active coefficients of the DCT transform of the variabilities are localized around the Gaussian kernel, enabling a complete control over the frequency of the VS. The multiplication with a Gaussian kernel also allows to have smooth variabilities as it can be observed in 5.7. The generation processes are presented in figure 5.3 and described below.

2. SV piecewise constant, in this second configuration, SV are generated as piecewise constant on blocks of minimal size 60. More specifically, given a number N of blocks, we assign to each block a random value between 0 and 1. Depending on the experiment, they may be convolved with a Gaussian kernel of given standard deviation.

5.4.3 Experimental setup

A good initialization for svGMCA is required given the non-convexity of the ill-posed problem tackled. It is brought by L-GMCA which consists of two steps:

1. we apply GMCA on blocks of the data matrix which size is fixed to 500 samples (unless stated otherwise). This blocksize allows a trade-off between statistical consistency (GMCA requires sufficient statistics) and accuracy towards the variation of the SV (small blocks).

2. Unless we work with piecewise constant spectral variabilities, the mixing matrices obtained with GMCA are smoothened with a Fréchet mean[1].

As stated in the algorithm presentation, svGMCA consists of two steps: the first one without a reweighting procedure on the SV and the second one acts as a refinement and involves a reweighting scheme. The most relevant implementation details are recalled below:

Stage 1

- Update of the sources: The threshold of the sources is obtained with the classical $\kappa\mathrm{MAD}$ operator with $\kappa = 3$ and involves a reweighting procedure with parameter $\epsilon = 10^{-3}$. The tolerance of the stopping criterion, computed on the relative difference between two successive values of the sources in direct domain, is set to $\epsilon_S = 10^{-6}$.

- Update of the mixing matrix: The proximal operator is computed with a Dykstra algorithm of stopping criterion 10^{-6} and the FISTA algorithm is stopped when the relative difference between two iterates is less than 5×10^{-7}.

- This first stage comes to an end when the relative difference between two successive values of the results $(\mathbf{A}, \mathbf{S})$ of the BCD algorithm is less than 10^{-4}.

Stage 2

The second stage that implements a reweighting procedure on the spectral variabilities ends when the tolerance criterion of the BCD algorithm is less than 10^{-6}.

---

[1] A definition of the Fréchet mean is provided in appendix E.

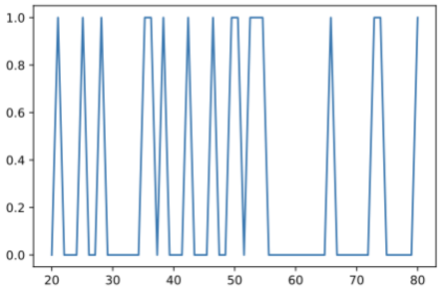

Vector drawn from a Bernoulli-Gaussian distribution of activation parameter $\rho = 0.3$

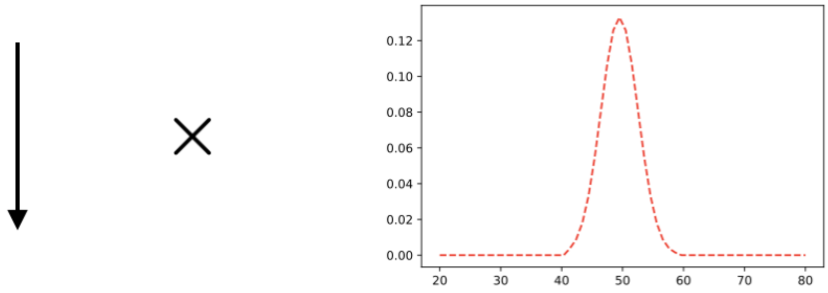

Gaussian kernel of standard deviation $\sigma = 3$

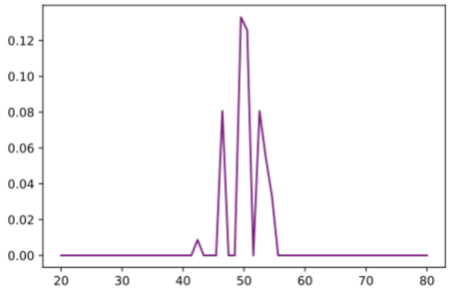

Spectral variability vector in DCT domain

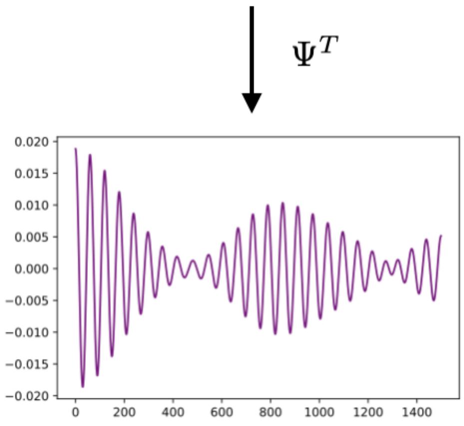

Spectral variability vector in direct domain

Figure 5.3: Illustration of the generation process of SV sparse in DCT domain. We have only illustrated the SV on the range of samples that goes from 20 to 80.

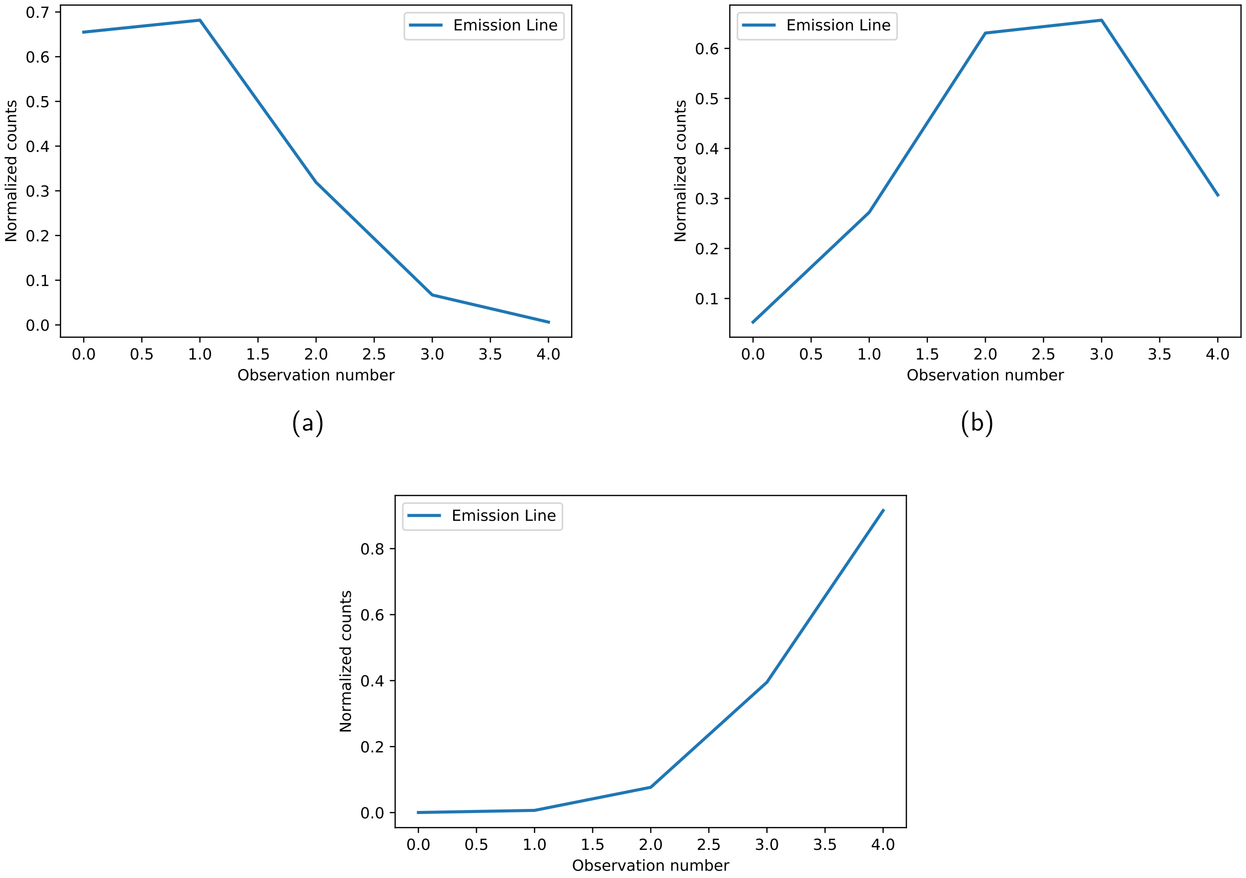

(a)

(b)

(c)

Figure 5.4: emission line for different central positions: 0.1 (a), 0.5 (b) and 0.9 (c).

For each study, we vary the parameter of interest through a range that will be specified and we carry out 10 Monte-Carlo realizations for each one.

The results of svGMCA are compared with L-GMCA (as it has been described above), which, as far as we know, both tackles the problem of BSS with variabilities and can be directly applicable to data as we have generated it.

5.4.4 Performance criteria

The performances of svGMCA and L-GMCA have been assessed through both general metrics (GMSE and SAD) and metrics that we have constructed specifically for spectral variabilities. In the following, $\hat{\mathbf{A}}$ (respectively $\hat{\mathbf{S}}$) denote the ground truth mixing matrix (respectively the ground truth sources).

Overall metrics

1. General Mean Square Error (GMSE): metric for mixing matrices reconstruction. It allows us to measure the quality of reconstruction of the spectral variabilities as well as the reference matrix.

$$GMSE(\mathbf{A}, \hat{\mathbf{A}}) = \frac{1}{mnt}\|\mathbf{A} - \bar{\mathbf{A}}\|_F^2$$

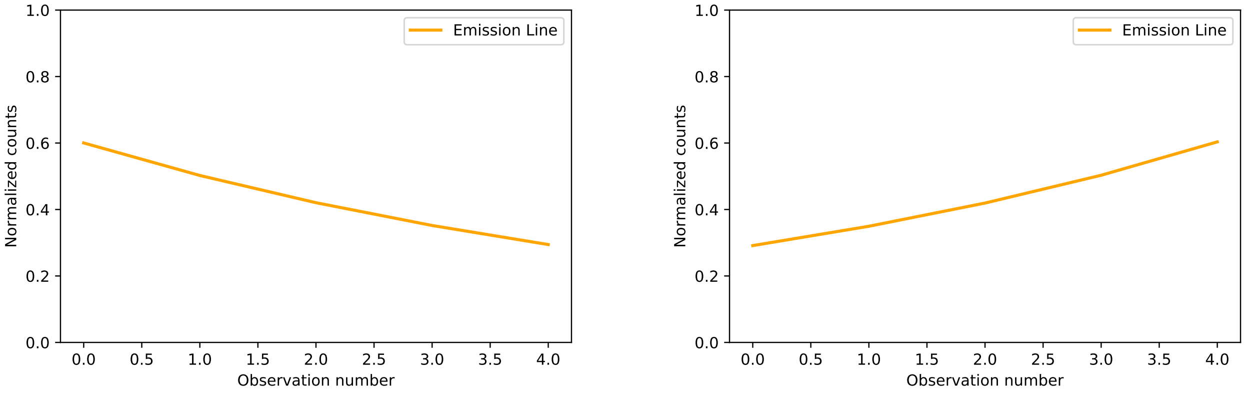

(a)

(b)

Figure 5.5: Synchrotron Emission for indices of -0.9 (a) and 0.9 (b).

2. Spectral Angular Deviation (SAD) on the mixing matrices: metric that computes the angular difference between the columns of mixing matrices (*i.e.* the spectra). This criterion allows us to have an estimation of the angular variability reconstruction.

$$SAD(\mathbf{A},\hat{\mathbf{A}}) = \frac{1}{n\times t}\sum_{i=1}^{n}\sum_{k=1}^{t}\arccos(<\mathbf{A}^i[k],\hat{\mathbf{A}}^i[k]>)$$

3. GMSE on the sources to assess the quality of reconstruction of the source matrix:

$$GMSE(\mathbf{S},\hat{\mathbf{S}}) = \frac{1}{n\times t}\|\mathbf{S}-\hat{\mathbf{S}}\|_F^2$$

Specific metrics

1. MSE on the spectral variabilities in their sparsity domain without accounting for their average value: this metric allows us to have an criterion on the reconstruction of the spectral variabilities without accounting for their average value.

$$\mathrm{GMSE}_{\boldsymbol{\Delta}}(\mathbf{A}^i,\hat{\mathbf{A}}^i) = \frac{1}{n\times T_{\mathbf{A}}}\sum_{i=1}^{n}\| \|(\mathbf{A}^i-\bar{\mathbf{A}}^i)\Phi^T\|_F - \|(\hat{\mathbf{A}}^i-\hat{\bar{\mathbf{A}}}^i)\Phi^T\|_F\|_F^2$$

2. MSE on the angular variability: this metric allows to have an estimation of the recovery of the angular variability without taking account the errors due to the approximation. For eac

$$\mathrm{GMSE}_\theta(\mathbf{A}^i,\hat{\mathbf{A}}^i) = \frac{1}{n\times t}\sum_{i=1}^{n}\|\ \theta^i-\hat{\theta}^i\|_F^2 = \frac{1}{t}\|\arccos(<\mathbf{A}^i,\bar{\mathbf{A}}^i>)-\arccos(<\hat{\mathbf{A}}^i,\hat{\bar{\mathbf{A}}}^i>)\ \|_F^2$$

5.4.5 Conventions

In the following, we will illustrate the behaviour of the performance criteria and their interpretation with instances of ground truth SV and the reconstructed ones. Regarding the displayed figures, we will refer by "*sought variabilities*" the spatial evolution of the ground truth spectral variabilities affecting one component for each observation. In the same way, " *svGMCA variabilities*" (and respectively "*L-GMCA variabilities*") allude to the SV resulting from svGMCA algorithm (respectively L-GMCA algorithm).

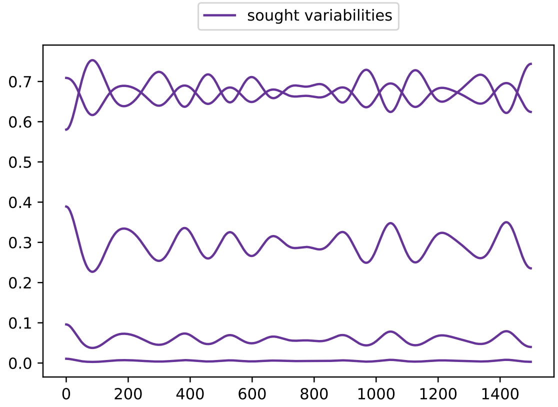

Figure 5.6: Illustration of the conventions chosen in this section: by denoting $\mathbf{\Delta}^i$ the spectral variabilities affecting the $i-th$ component, each of the purple line represents the vector $\mathbf{\Delta}^i_j$ (with $1 \leq j \leq m$) of size $t = 1500$.

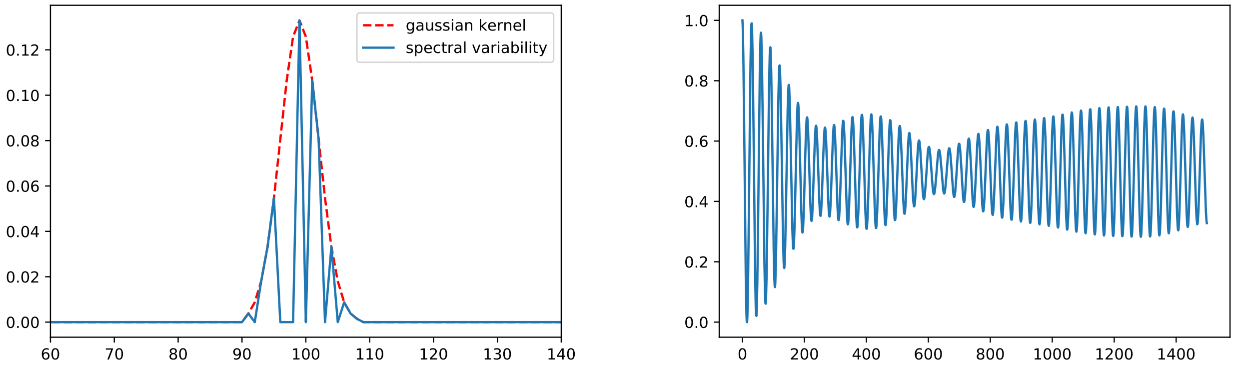

Figure 5.7: The left figure is an example of spectral variabilities of activation parameter $\rho_{VS} = 0.1$ centered around a frequency of 100 and convolved with a Gaussian kernel (dotted red line) of standard deviation $\sigma_{VS} = 3$. The right figure shows the same SV in direct domain.

5.4.6 1D Monte-Carlo simulations with a simplistic case: mixture of two sources and spectral variabilities sparse in DCT domain

In this first series of experiments, we consider spectral variabilities that are exactly or approximately sparse in DCT domain. The studies described below consist of observing and analyzing the behaviour of the algorithms towards the variation of a parameter of interest, depending on the nature of the spectra and the morphology of the sources.

The datasets used in the following experiments are as follows:

- We consider the mixture of two components: a power law and an emission line of width 0.1.
- The sources are generated from a Bernoulli-Gaussian distribution of activation parameter $\rho_{\mathbf{S}}$.
- The spectral variabilities are sparse in DCT domain. The generation process is provided in 5.4.2. As stated before, three parameters on the generation of the SV can be tuned:
  1. The activation parameter of the Bernoulli-Gaussian distribution ρ_{VS} that provides the sparsity degree

of the variabilities;

2. The central frequency of the filtering kernel which describes the speed of variation of the VS;
3. The width of the filtering kernel given by σ_{VS}.

An example of SV is provided in 5.7.

- The noise is generated according to a Gaussian distribution of zero mean. The level noise is chosen so that the signal-to-noise ratio is $\mathrm{SNR} = 65\mathrm{dB}$.

To avoid redundancy, we restrict the description of the datasets to the numerical values of the above-mentioned parameters.

A. Influence of the sparsity degree of the sources:

Sources	Spectral variabilities			emission line	power law
ρ_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
-	0.2	5	20	$[0.85, 0.95]$	$[-0.9, -0.8]$

Table 5.1: Parameter settings for the study of the impact of the sparsity of the sources

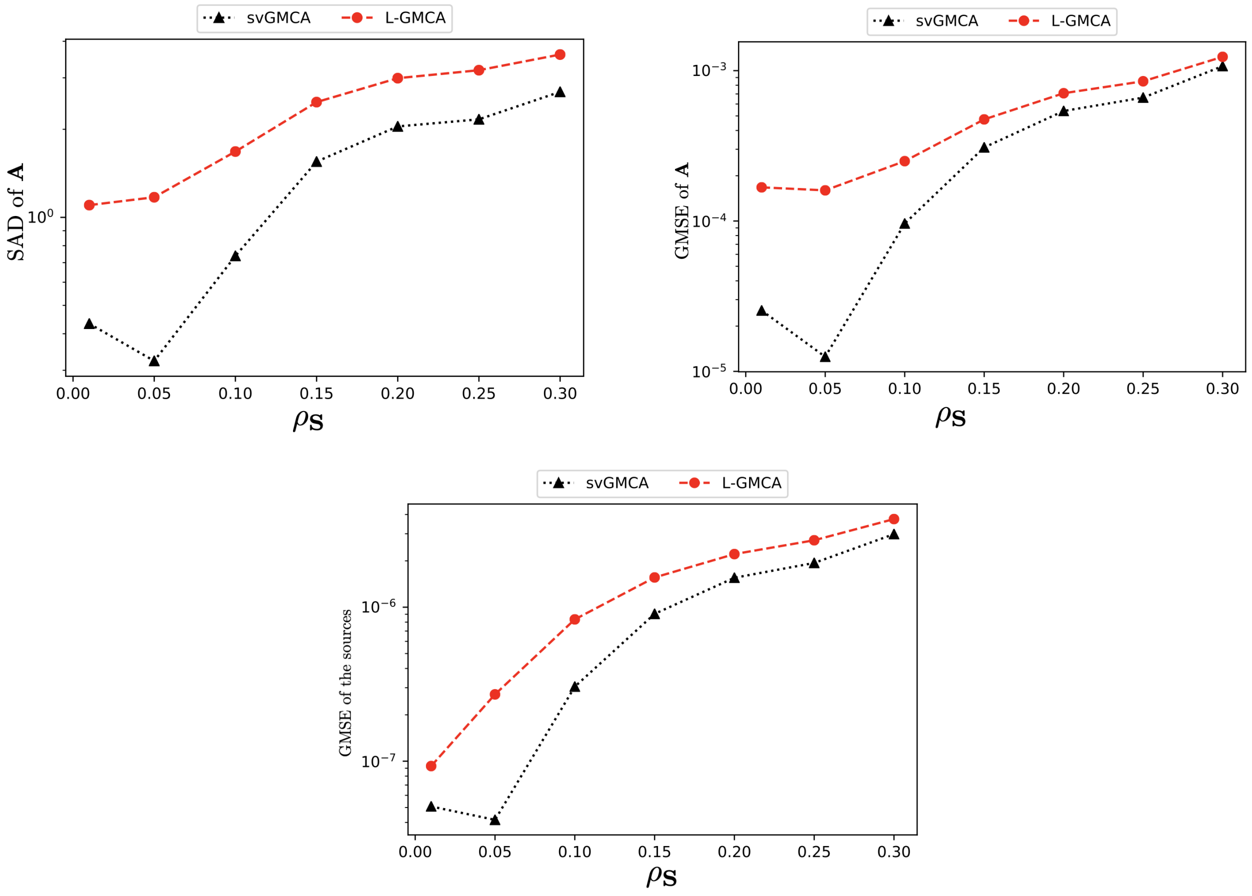

Figure 5.8: Influence of the sparsity degree of the sources

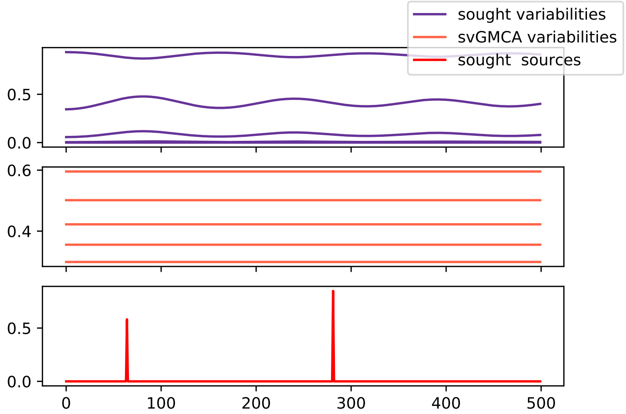

Figure 5.9: Illustration of spectral variabilities and the corresponding source distribution for $\rho_\mathbf{S} = 0.01$.

Figure 5.8: $\rho_\mathbf{S}$ is an indicator of sources sparsity in direct domain, the lower $\rho_\mathbf{S}$ the sparser the sources. Sources with a very low activation parameter have very few active coefficients (no more than a few dozens). Such a distribution amounts to poor spatial information on the spectral variabilities since each coefficient of the SV is estimated based on only one observation of the data matrix. For example, in the extremal case displayed in figure 5.9, the estimation of the whole mixing matrix relies upon only two observation. This lack of "statistics" yields svGMCA to stay in the local minimum brought by the initialization. The amelioration of the results when $\rho_\mathbf{S}$ increases confirms the important role of the sources distribution for the SV.
The deterioration on the estimation of the variabilities when sparsity degree increases hints that the bias inherent to the soft-thresholding of the sources degrades the recovery of the SV. In other words, when $\rho_\mathbf{S}$ increases, the recovery of the sources is of less quality, which, by error propagation, deteriorates the recovery of the variabilities.

B. Influence of the amplitude of the spectral variabilities

Sources	Spectral variabilities			emission line	power law
$\rho_\mathbf{S}$	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
14%	0.4	3	20	$[0.05, 0.25]$	$[0.95 - l, 0.95]$

Table 5.2: Parameter settings for the study of the impact of the amplitude of spectral variabilities

Figure 5.10: The recovery of the spectral variabilities by L-GMCA can be seen as a recovery of their average value on a limited number of patches. The greater the amplitude of the variabilities, the further their variation from their average, and therefore the results of L-GMCA deteriorate. Concerning svGMCA, it is built upon the angular approximation that requires the angle to be small. High amplitudes yield to invalidating this assumption. In addition to that, if the amplitude of the peaks of the SV are high, so are the amplitudes of the biases which participates to the deterioration of the results.

C. Influence of the sparsity degree of the spectral variabilities

Figure 5.11: In this analysis, it is the maximum norm (ℓ_∞) of the variabilities that is fixed while varying their sparsity degree. The figure 5.12 shows the input variabilities for two extremal values of ρ_{VS}.
svGMCA imposes the $\ell_{2,1}$ sparsity constraint on the SV which makes it more efficient when the variabilities are highly sparse. On the opposite, when ρ_{VS} is high, the variabilities tend to "flatten". For example, the variability

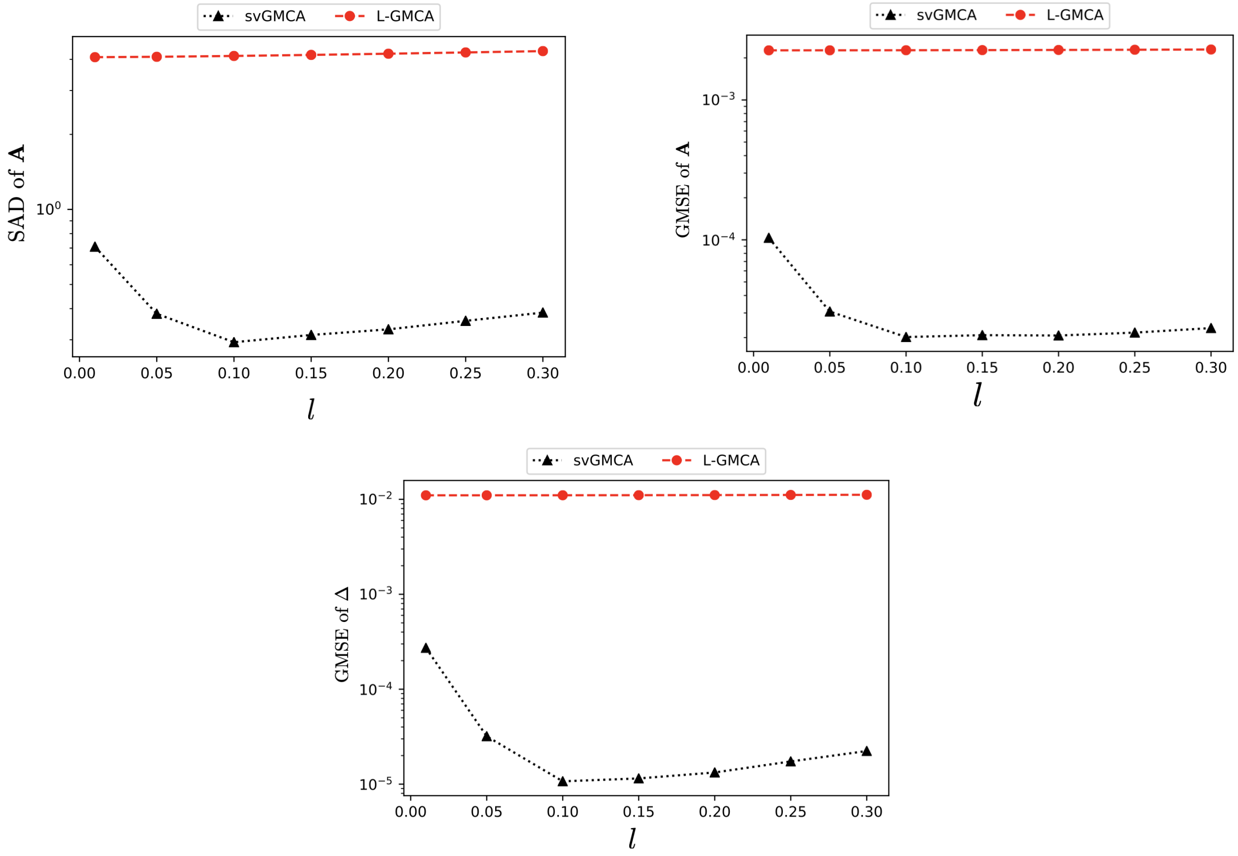

Figure 5.10: Influence of the amplitude of the spectral variabilities.

Sources	Spectral variabilities			emission line	power law
ρ_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
14%	-	5	20	$[0.85, 0.95]$	$[-0.05, 0.05]$

Table 5.3: Parameter setting for the study of the impact of ρ_{VS}

acquired for $\rho_{VS} = 0.9$ (figure 5.12) has a small amplitude starting the 200^{th} sample. Taking into account this property and what was previously stated concerning the influence of the amplitude, it is understandable that L-GMCA is almost as performing as svGMCA for high ρ_S.

D. Influence of the average angle between the two mixing directions

Sources	Spectral variabilities			emission line	power law
ρ_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
14%	0.3	3	20	$[0.85, 0.95]$	$[T, T + 0.1]$

Table 5.4: Parameter settings for the study of the average angle between the mixing directions $\bar{\theta}$.

The two emission lines being normalized, they belong to the ℓ_2 ball and can be characterized by the angle between them. This angle, that we denote $\bar{\theta}$ is of crucial importance since it marks the similarity or dissimilarity between the mixing directions.

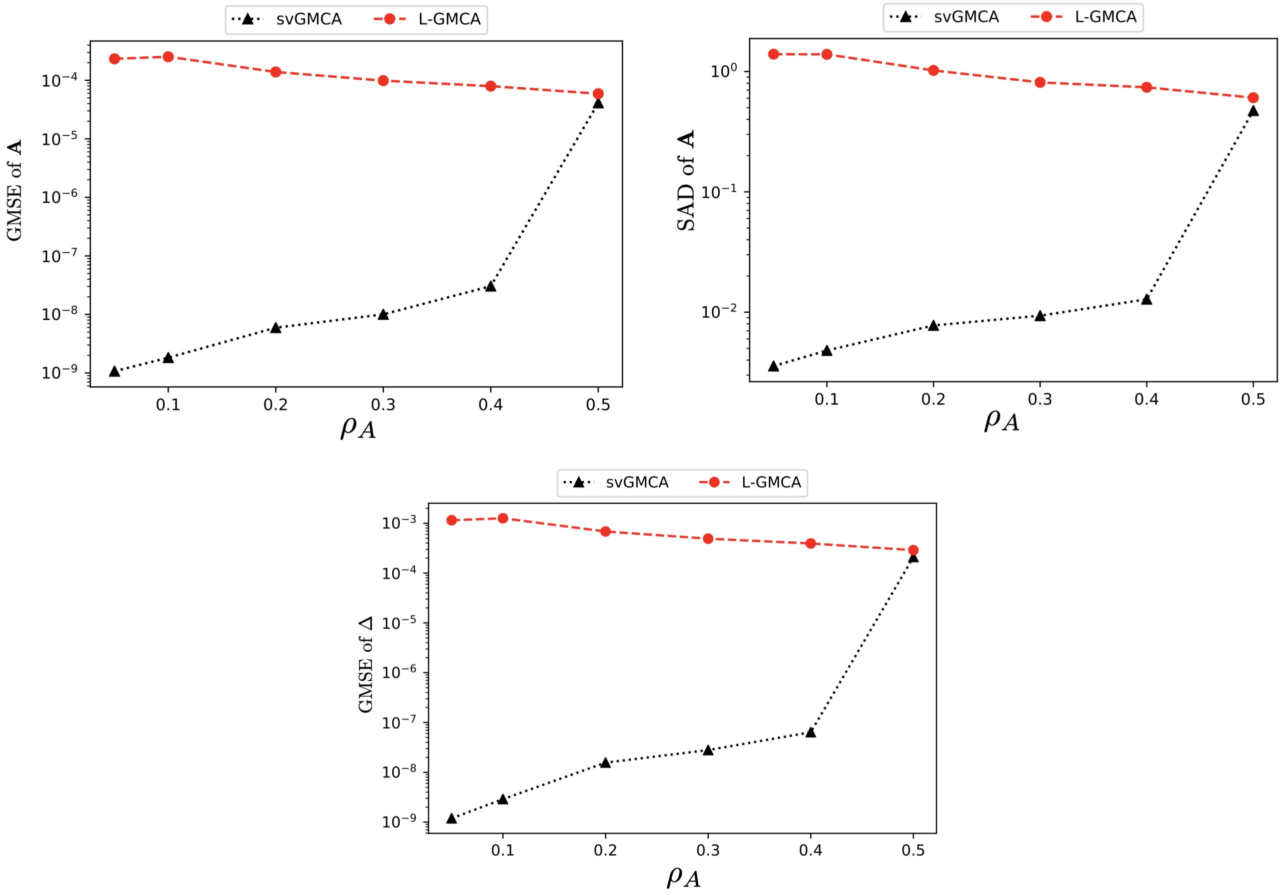

Figure 5.11: Influence of the sparsity degree of the SV

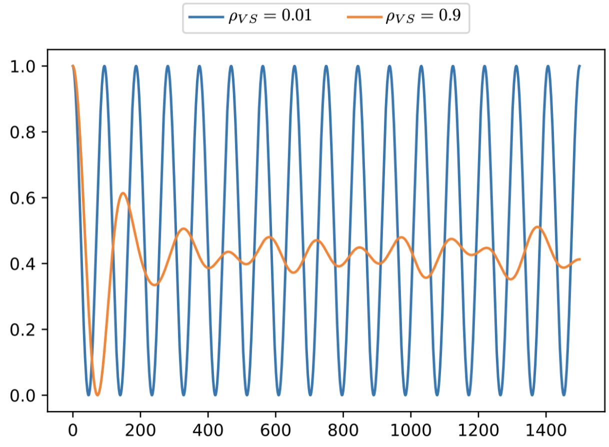

Figure 5.12: Example of spectral variabilities for two sparsity degrees.

Figure 5.13: If the two mixture directions are close or almost colinear, it becomes extremely complicated to unmix the data. The role of this angle is all the more important in the presence of spectral variabilities since, even if the reference mixing directions are far from each other, local mixing directions (*i.e.* mixing directions at a given sample affected by spectral variabilities) can be close if the amplitude of the variabilities is high enough. More generally, it is both the average angle and the amplitude of the SV that should be small "enough" for any BSS algorithm to perform data unmixing.

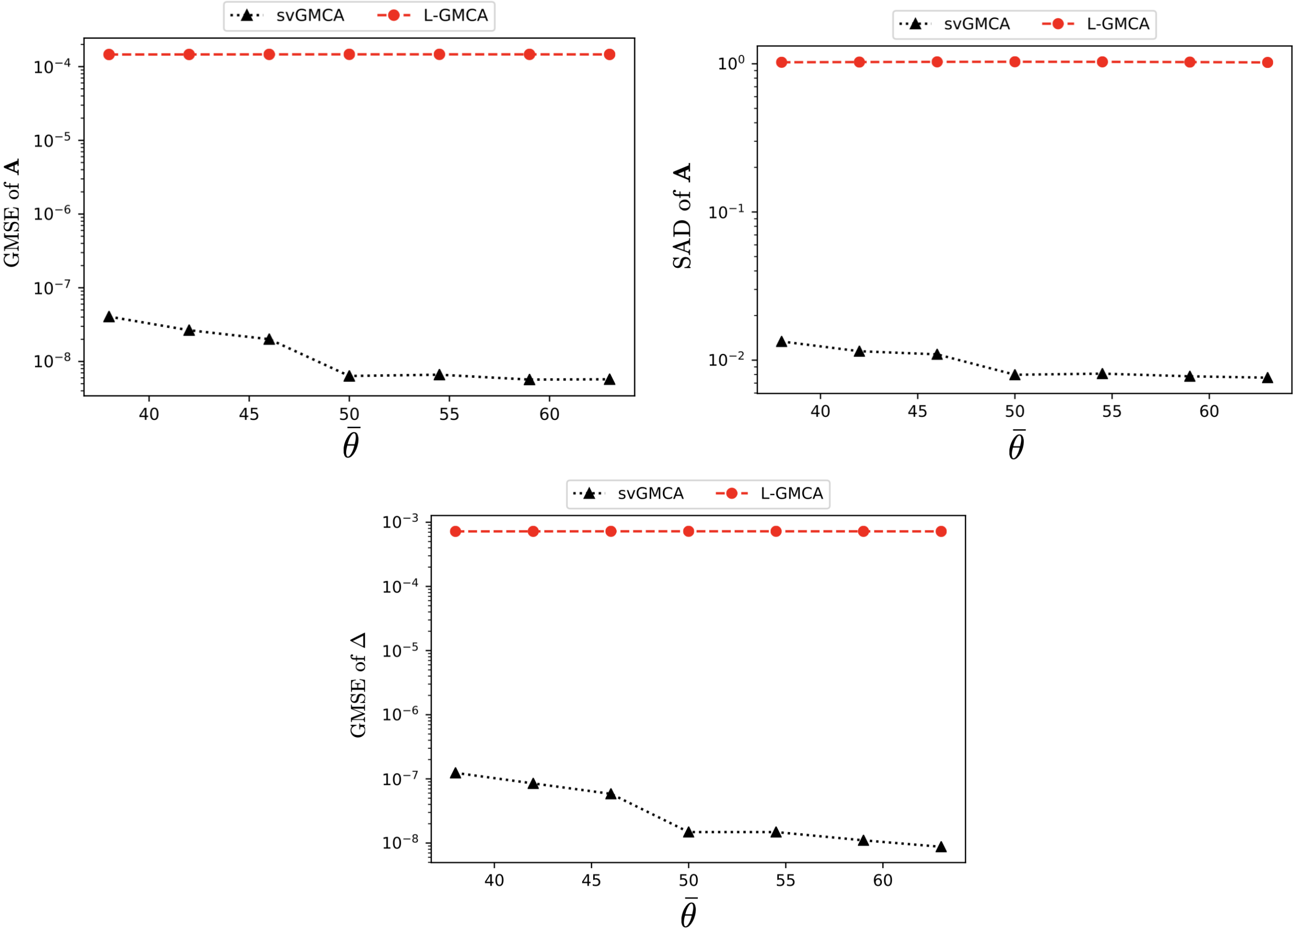

Figure 5.13: Influence of the average angle between the two mixing directions.

E. Influence of the speed of variation of the spectral variabilities

Sources	Spectral variabilities			emission line	power law
$\rho_\mathbf{S}$	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
10%	0.1	3	-	$[0.8, 0.9]$	$[-0.9, -0.8]$

Table 5.5: Parameter settings for the study of the impact of the central frequency f.

Figure 5.14: We observe a deterioration of the performances of svGMCA and L-GMCA when the SV speed of variation gets higher, which is natural since it boils down to recovering a large amount of information (the spectral variabilities) from few statistics (the source distribution). However, we can note that svGMCA still manages to recover fairly the SV with high frequencies (figure 5.15) since the spatial information is encoded in very few samples in DCT domain.

F. Summary of the study

In order to have a global understanding of the behaviour and the performances of svGMCA algorithm, more experiments have been conducted on the following datasets:

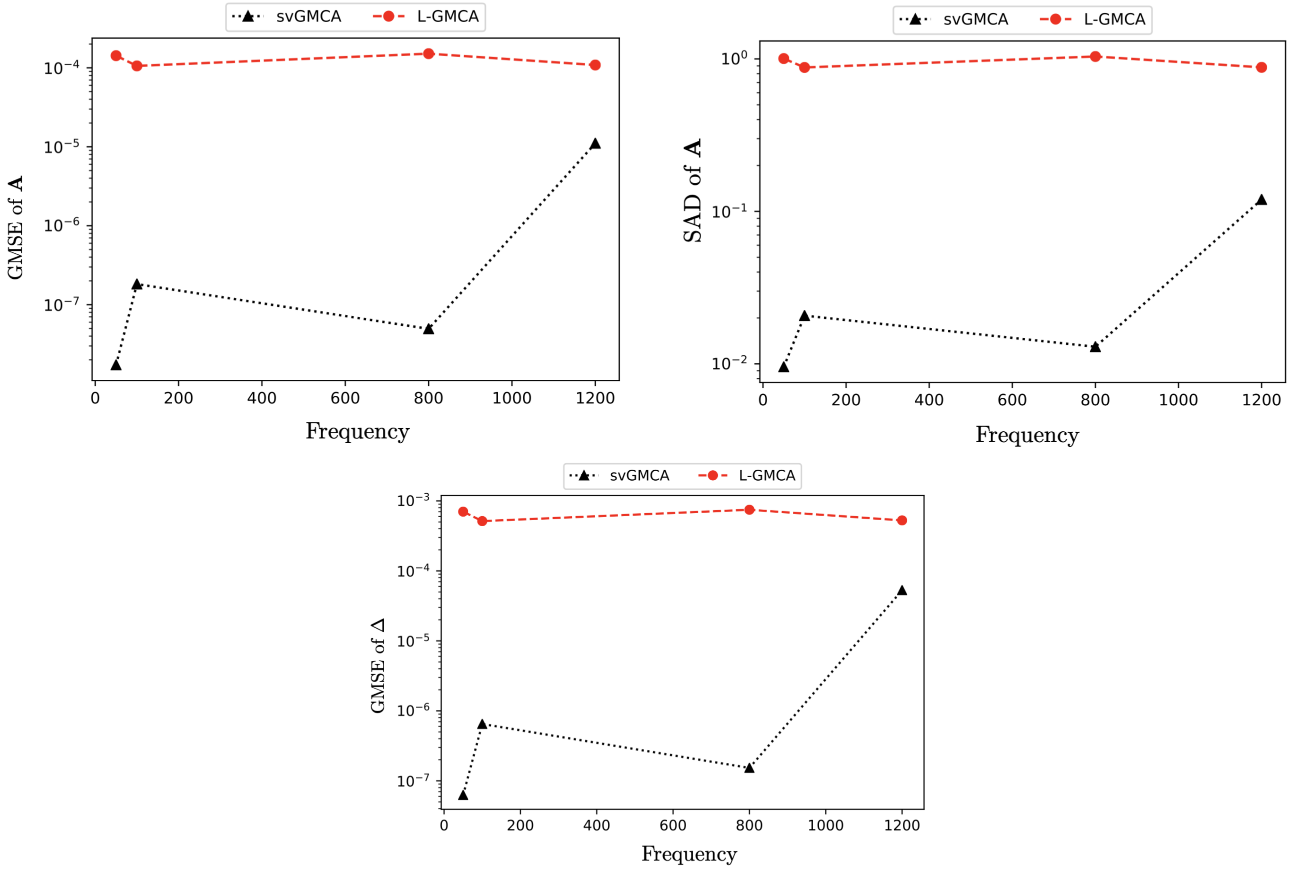

Figure 5.14: Influence of the speed of variation of the SV.

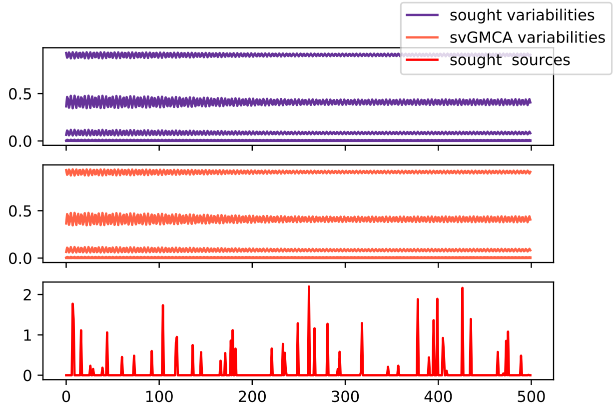

Figure 5.15: Example of sought and recovered spectral variabilities generated with a frequency $f = 800$ and the corresponding sources (truncated to 500 samples over 1500 for illustrative purpose). svGMCA manages to recover the variabilities even though they have a high variation speed in comparison with the sources.

- Mixture of two emission lines with sources sparse in starlet domain,
- Mixture of two emission lines with sources sparse in direct domain,
- Mixture of an emission line and a power law with sources sparse in starlet domain.

For each of these datasets, we have carried out 10 Monte-Carlo realizations for each value of the following parameters while the others are fixed:

- activation parameter ρ_S of the Bernoulli-Gaussian distribution generating the sources (which provides an understanding of the impact of the sparsity degree of the sources),
- amplitude l of the spectral variabilities,
- activation parameter ρ_{VS} of the Bernoulli-Gaussian distribution generating the spectral variabilities (which provides the sparsity degree of the SV in DCT domain),
- the Position Law index range or the emission line position range in the case of the mixture of two iron emission lines (as an indicator of the average angle $\bar{\theta}$),
- The position of the central peak of the SV in DCT domain (which allows a variation of the speed of variation of the SV).

The results of the Monte-Carlo simulations are provided in the appendix B.
We propose in 5.6 a synthesis table of the influence of the parameters.

5.4.7 1D Monte-Carlo simulations in the case of mixture of two sources and piecewise constant spectral variabilities

We investigate now the performances of the algorithm in the case of piecewise constant spectral variabilities. The following experiments aim at observing the influence of a parameter while the others are fixed. The datasets used for the simulations are as follows:

- We consider the mixture of two sources drawn according to a Bernoulli-Gaussian distribution of activation parameter $\rho_{\mathbf{S}}$. The generated peaks are convolved with a Laplacian kernel of full width at half maximum $w_{\mathbf{S}}$.
- Both of the columns of the mixing matrix are inspired from iron emission lines with a width of 0.1.
- The spectral variabilities are piecewise constant. Their generation process follows the description given in 5.4.2. Three parameters are of interest for us in the following studies:
    1. The number of blocks N: by changing the number of blocks on which the SV have a constant value, we are able to change the speed of variation of the SV,
    2. The full width at half maximum (FWHM) of the Gaussian kernel used in case of SV filtering,
    3. Unless stated otherwise, the minimal size of blocks m is fixed at 60.

  An illustration of piecewise constant SV is provided in figure 5.16.
- The noise is drawn from a Gaussian distribution with zero mean. The signal-to-noise ratio is provided for each experiment.

Given the morphology of the spectral variabilities, we have chosen Haar transform as sparsity domain since it is well-suited for piecewise constant or continuous functions. More specifically, we have worked with undecimated Haar transform provided by pywt rather than decimated transform that would lead to a reconstruction of lower quality (Starck et al. [2007]). Unless otherwise stated, the level of resolution is $J = 2$.

To avoid redundancy, we restrict the description of the datasets to the numerical values of the above-mentioned parameters.

Parameter	Impact on svGMCA
Sparsity degree of the sources	Each element of the mixing matrix is estimated based on only one observation. Highly sparse source distribution boils down to very little spatial information on the spectral variabilities, which hampers their detection. On the other hand, low sparsity degree amounts to deterioration on the recovery of the sources (the cost function minimized by svGMCA involves a sparsity constraint on the sources) and, by error propagation, on the variabilities. Moreover, we have shown in appendix D that the bias inherent to the softthresholding of the sources drastically hampers the identification of the variabilities.
Amplitude of the SV	svGMCA is built upon the angular approximation that requires small angle amplitudes. High values of SV amplitudes jeopardize this approximation and make the spectral variabilities term $\mathbf{\Delta}$ less sparse. On the other hand, low-amplitude variabilities are more contaminated by noise.
Sparsity degree of the SV	Since svGMCA enforces the sparsity of the SV, it performs better for highly sparse variabilities.
Angle between the reference mixing directions	In any BSS algorithm, if the average angle is low (large condition number) the mixing directions are correlated and therefore harder to separate and the sources harder to unmix. This is all the more detrimental in the presence of spectral variabilities since the angle between local mixing directions (*i.e.* columns of the mixing matrix affected by spectral variabilities) may be lower in some samples than the average angle.
Speed of variation of the SV	The problem we tackle being very ill-posed, the variabilities require for their identification to have a lower speed of variation than the sources in order to ensure enough observations of the SV (*i.e.* enough statistics).

Table 5.6: Synthesis of the studies carried out on SV sparse in DCT.

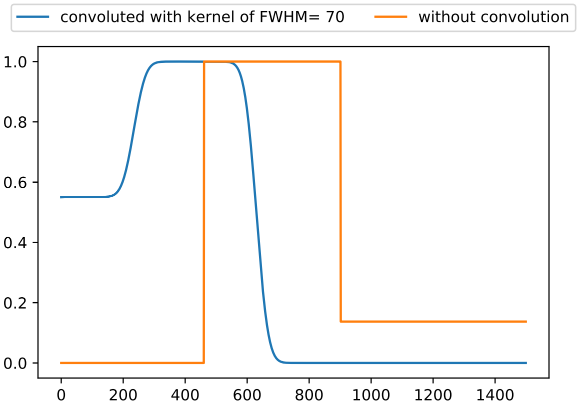

Figure 5.16: Two examples of 3-block piecewise constant spectral variabilities with a minimal size of block fixed to 200. The first illustrated SV (blue) has been generated with a convolutive kernel of FWHM of 70, whereas the second one is kept without filtering.

A. Influence of the input noise

Sources			Spectral variabilities			emission line 1	emission line 2	SNR (in dB)
Convolution	ρ_S	w_S	FWHM	m	N	Position range	Position range	
Yes	4%	3	70	200	3	$[0.05, 0.15]$	$[0.41, 0.49]$	-

Table 5.7: Parameter settings for the study of the SNR

Figure 5.17: To quantitatively assess the influence of the noise on the estimation of the variabilities and the robustness of our algorithm, we have added to our datasets as described above Gaussian noise with varying signal-to-noise ratios (SNR).

The problem being multi-convex, it requires strong regularization which involves high thresholds .Given that the implemented thresholding procedure is dependent of the noise, a higher noise level boils down to a stronger regularization and therefore a shrinkage of the range of solutions. The objective function is less likely to drop in local minima.

Therefore, if the SNR is low, svGMCA is more robust to local minima. However, the performances are hampered by high noise contamination. On the opposite, above $\mathrm{SNR} = 45\mathrm{dB}$, the performances of svGMCA deteriorate since the threshold is low.

It is interesting to notice that L-GMCA exploits the robustness of GMCA and performs fairly and similarly on the whole range. However, it remains less effective than svGMCA since it does not properly account for spectral variabilities.

B. Influence of the amplitude of the SV

Sources			Spectral variabilities			emission line 1	emission line 2	SNR (in dB)
Convolution	ρ_S	w_S	FWHM	m	N	Position range	Position range	
Yes	4%	3	70	200	3	$[0.05, 0.05 + l]$	$[0.9, 0.9 - l]$	45

Table 5.8: Parameter settings for the study of the impact of the amplitude of the SV.

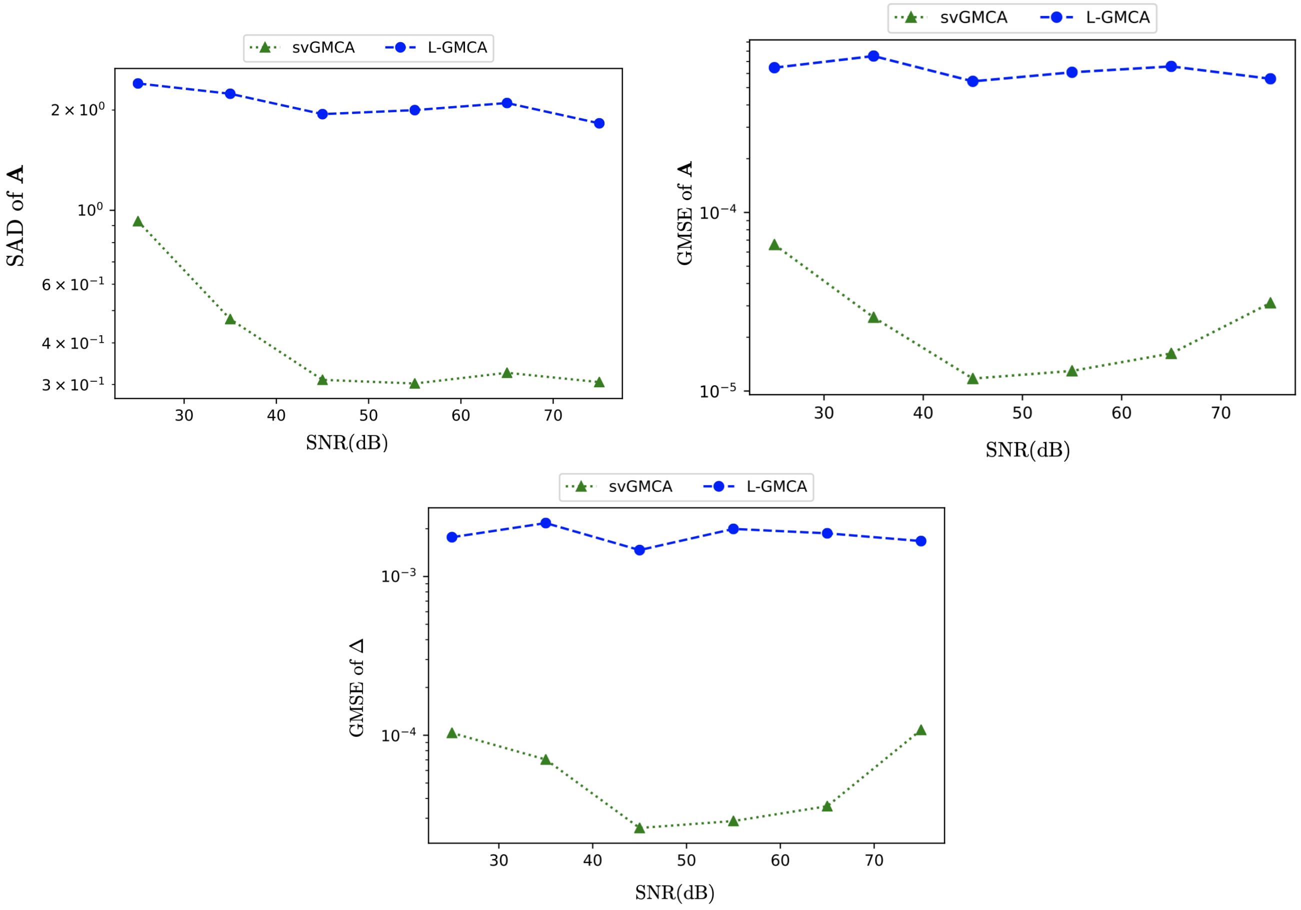

Figure 5.17: Influence of the input noise for the second dataset

Figure 5.18: Spectral variabilities are piecewise constant and the amplitude between the lowest value and the highest one plays a crucial role in their estimation. In this part, we aim at observing the influence of this parameter on the estimation of the variabilities in the case where the sources are sparse in a transformed domain. We observe that both of the algorithms perform badly when the amplitude of the SV affecting the second source is high. In the case of svGMCA, this is due to the angular approximation that requires the amplitude to be low. Regarding L-GMCA, it does not estimate the variations of the variabilities but their "average value" on blocks of the samples. Therefore, high amplitude of the variabilities amounts to larger discrepancy with their average value.

C. Influence of the average angle between the two mixing directions

Sources			Spectral variabilities			emission line 1	emission line 2	SNR (in dB)
Convolution	ρ_S	w_S	FWHM	m	N	Position range	Position range	
Yes	4%	3	70	200	3	$[0.05, 0.25]$	$[0.05 + T, 0.25 + T]$	45

Table 5.9: Parameter settings for the study of the impact of the average angle between the two mixing directions $\bar{\theta}$.

Figure 5.19: The separation between the mixing directions is a key element in BSS and it is of even greater

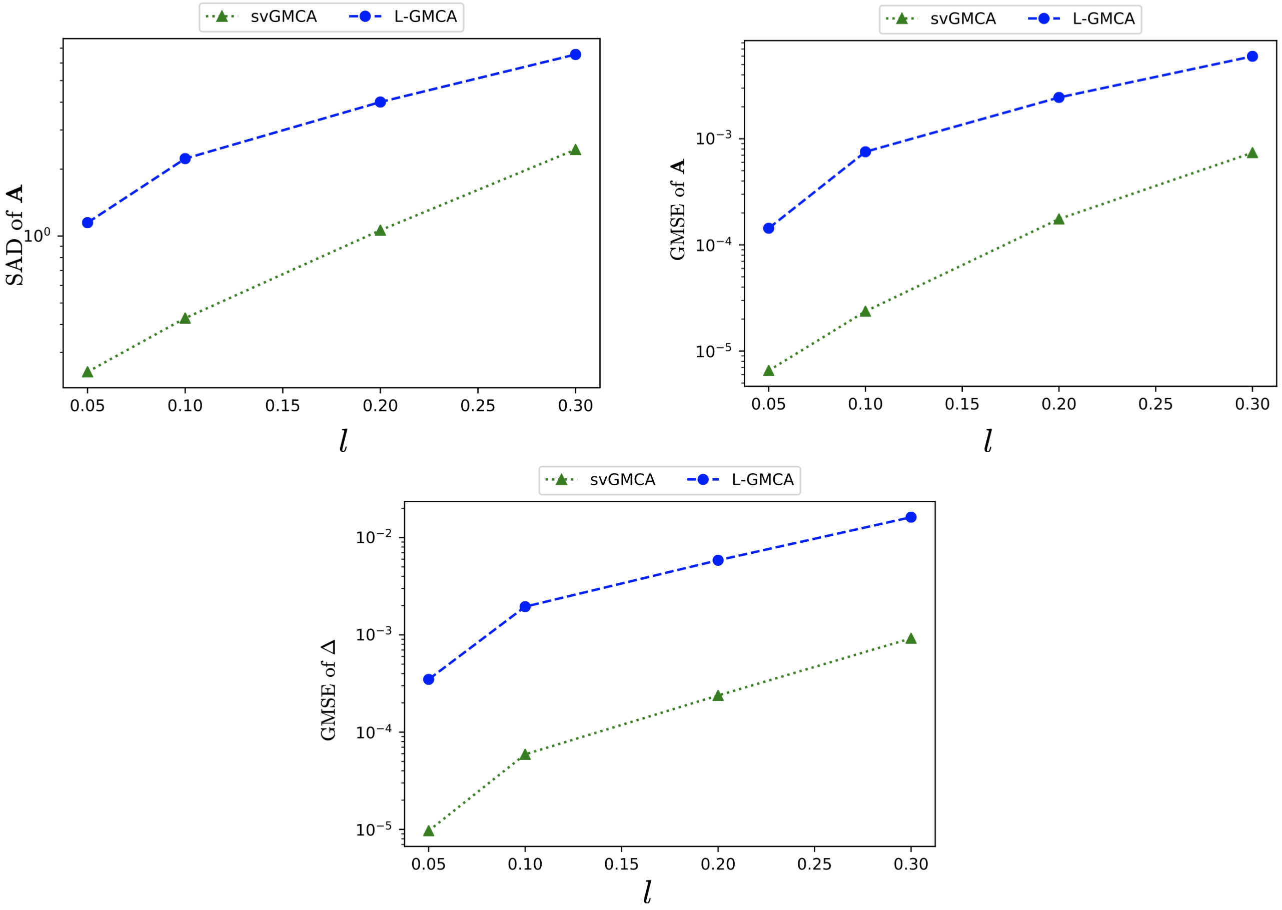

Figure 5.18: Influence of the amplitude of the SV for the second dataset

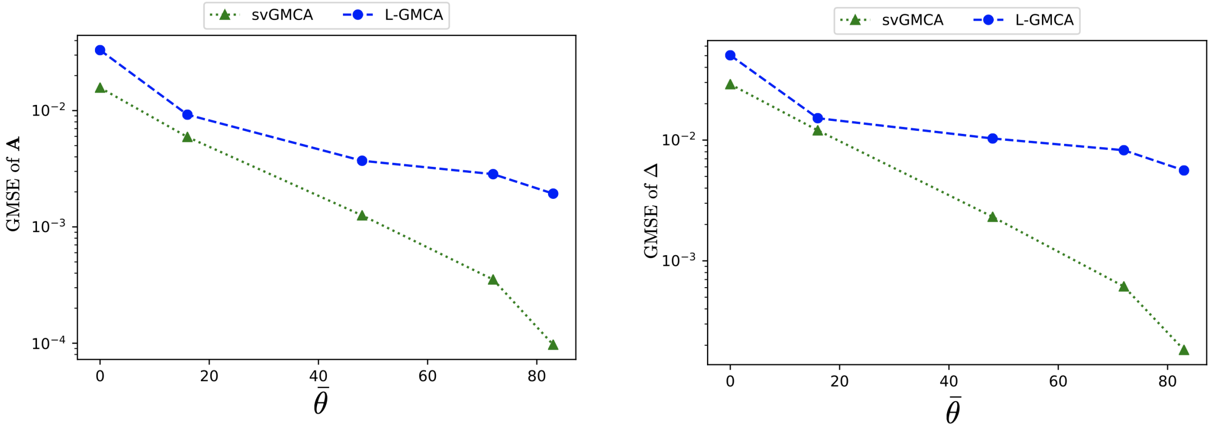

Figure 5.19: Influence of the average angle between the two mixing directions $\bar{\theta}$.

importance in the presence of spectral variabilities. The goal of this study is to observe and analyze how this separation impacts the disentanglement of the sources and the detection of the piecewise constant VS. svGMCA and L-GMCA are similarly impacted by an increase of the average angle. Not only it allows a better source separation but it also helps for the detection of the variabilities as it is shown in the figure 5.19 corresponding to the GMSE of $\mathbf{\Delta}$.

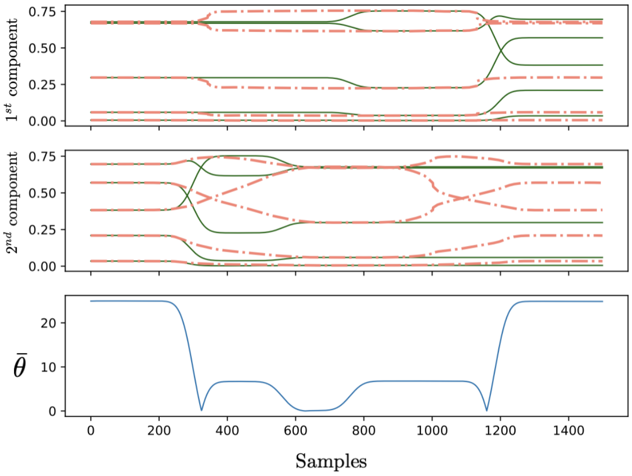

Figure 5.20: This figure shows the sought variabilities (in green) and the recovered ones (in orange) for an instance of the first dataset with $\bar{\theta} = 0$. The bottom figure displays the angular difference between the mixing directions sample-per-sample.

The figure 5.20 illustrates the behaviour of the algorithm when the mixing directions are very close. The angular difference below the 200^{th} sample highlights very clearly the area where svGMCA manages to separate the two sources. Beyond this value, the two average directions are two entangled for svGMCA to correct the permutation error (the SV attributed to the first component beyond the 1200^{th} sample corresponds in fact to the second component and vice-versa)

D. Influence of the speed of variation of the SV

Figure 5.21: By speed of variation of the SV we refer to the number of blocks on which the SV has a constant value. To avoid extremal and non interpretative cases, we propose a minimal size of block set at $m = 60$ samples.

Sources			Spectral variabilities			emission line 1	emission line 2	SNR (in dB)
Convolution	ρ_S	w_S	FWHM	m	N	Position range	Position range	
Yes	4%	3	70	60	-	$[0.05, 0.15]$	$[0.85, 0.95]$	45

Table 5.10: Parameter settings for the study of the impact of the speed of variation of the SV.

As shown in the figures 5.21, the results deteriorate when the number of blocks increases. A high variation speed of the SV amounts, on the one hand to detect each block with less samples, and on the other to a representation of the SV less sparse in the transformed domain.

E. Influence of the filtering of the SV with patches of same size and sources sparse in direct domain

Sources			Spectral variabilities			emission line 1	emission line 2	SNR (in dB)
Convolution	ρ_S	w_S	FWHM	m	N	Position range	Position range	
No	14%	0	-	200	3	$[0.05, 0.25]$	$[0.75, 0.95]$	45

Table 5.11: Parameter settings for the study of the impact of the filtering of the SV.

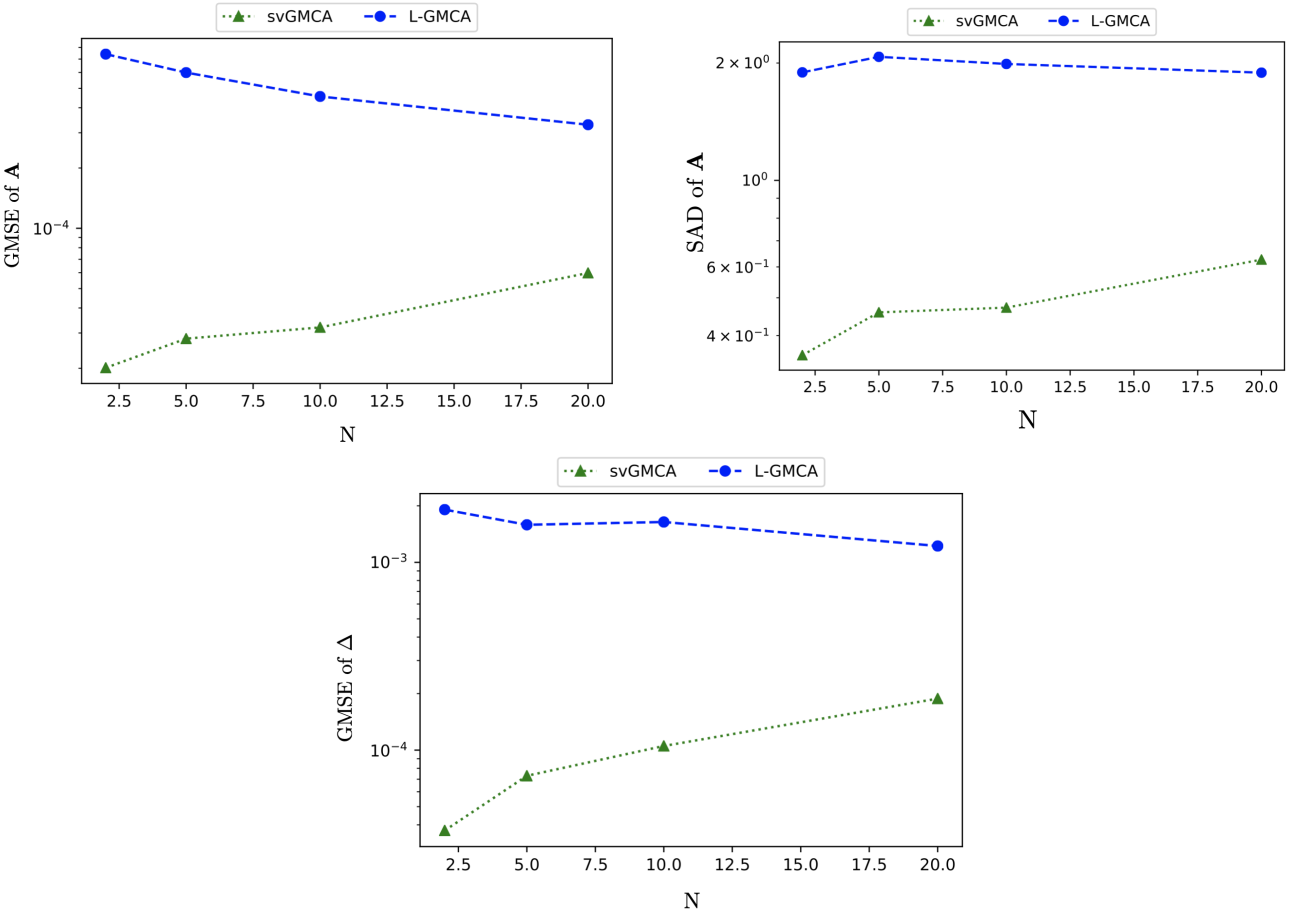

Figure 5.21: Influence of the number of blocks for the second dataset.

The considered spectral variabilities can incorporate a filtering process that changes their morphology. A comparison of SV with and without filtering is given in figure 5.16. The goal of this study is to infer the impact of these changes of morphology due to the filtering kernel and to compare it to two variants of L-GMCA:

- GMCA per patch: it is the algorithm of reference used as initialization and comparison for SV piecewise constant. It consists of simply implementing GMCA on patches of predefined sizes and concatenate the results. In this particular study, we have generated the ground truth variabilities on three blocks and we implement L-GMCA considering this a priori information is known,
- GMCA per patch with filtering: it boils down to implement a Fréchet mean on the SV recovered by L-GMCA. We use this algorithm for comparison since it smoothens the result of GMCA per patch and may be more appropriate for variabilities with a high filtering (high σ).

The figure 5.22 illustrates the performances of the three algorithms assessed through three criteria: GMSE and SAD implemented on the sample-dependent mixing matrices and the GMSE on the spectral variabilities only. As it can be noticed, when there is no filtering scheme (*i.e.* $\sigma = 0$), GMCA per patch achieves the best results (up to 10^2 difference with svGMCA). Given that we consider the size of the blocks known, it simply amounts to implementing GMCA on the three blocks. This method is effective provided we have enough samples, which is ensured by the distribution of the sources (Generalized Gaussian with an activation parameter of 14%). The figure 5.23 is provided to highlight the difference between the results of the algorithm in the case where no filtering is implemented for the SV ($\sigma = 0$).

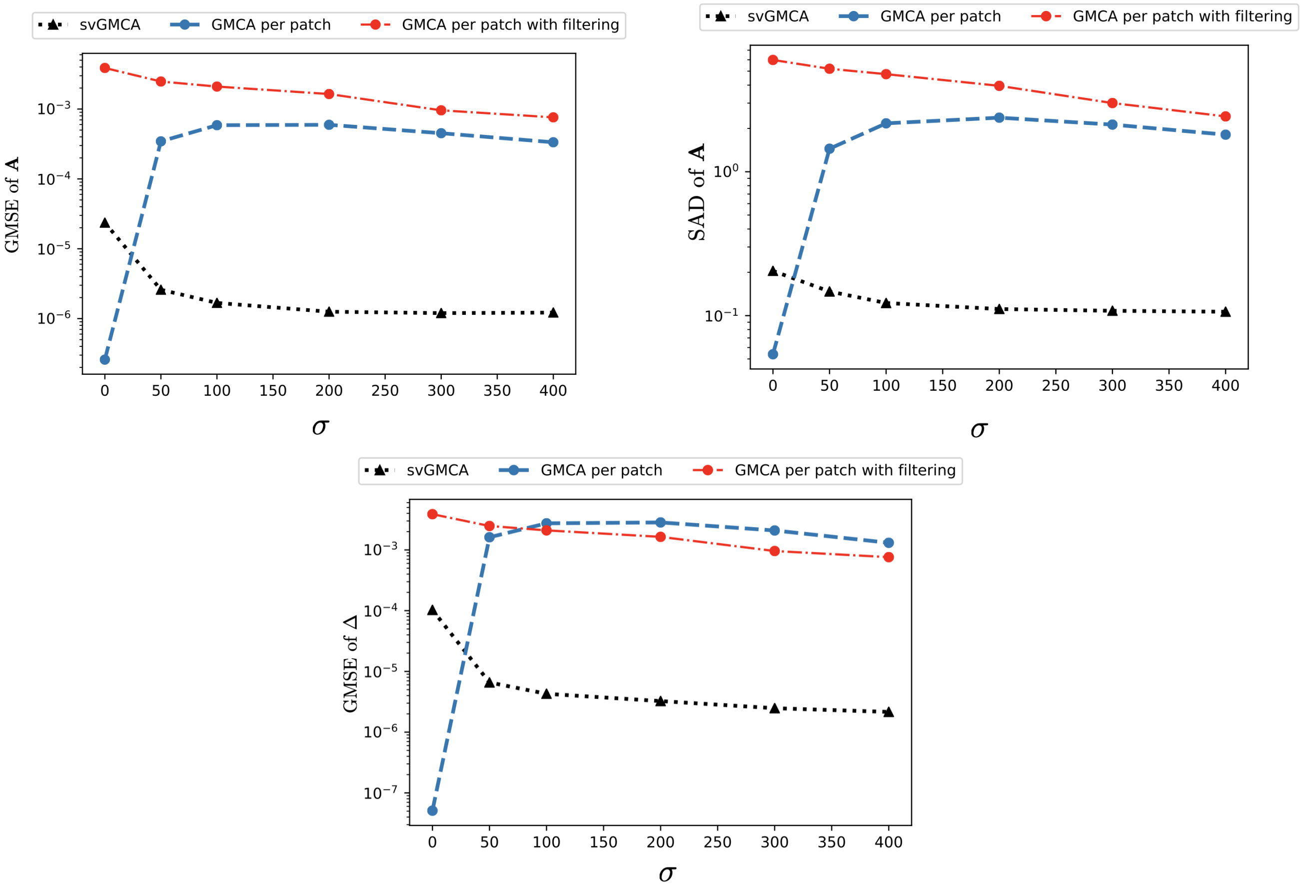

Figure 5.22: Influence of the filter of the SV

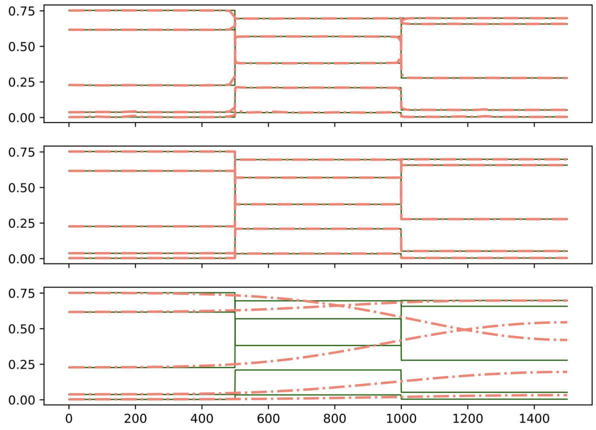

Figure 5.23: Comparison of the sought mixing spectral variabilities without filtering (in green lines) compared to algorithms results: from top to bottom, svGMCA, GMCA per patch and GMCA per patch + filtering

F. Summary of the study

We have carried out other experiments with the mixture of two emission lines with source sparse in direct domain and SV piecewise constant. We have conducted 10 Monte-Carlo simulations for each value of the following parameters while the other ones are fixed:

- signal-to-noise ratio,

- amplitude of the spectral variabilities,
- position range of the emission line (that accounts for a variation of $\bar{\theta}$),
- number of blocks of the spectral variabilities N (representing the speed of variation of the SV),
- FWHM of the Gaussian kernel used for SV filtering.

The results are provided in appendix C. **The key difference between the scenario "SV sparse in DCT domain" and "piecewise constant SV" is the lack of morphological diversity diversity between the sources and the variabilities in that last case.** Indeed, the sources have a sparse distribution in the SV sparsity domain (undecimated Haar domain). This causes the backprojected noise on the variabilities to have a sparse representation and act as outliers. From an algorithmic viewpoint, it amounts to having a cost function with several spurious local minima. An overview table of the results is provided in table 5.12.

Parameter	Impact on svGMCA
Noise level	The noise plays an interesting role in the dataset considered: the increase of noise level allows to have a better regularization of the problem through higher regularization parameter of the $\ell_{2,1}$ norm penalization of the SV. Therefore, it allows to smoothen the cost function. On the opposite, a "too" low SNR would naturally deteriorate the results since the data are noisier.
Amplitude of the SV	As discussed in 5.6, svGMCA requires low-amplitude SV for the angular approximation to hold. However, in the case of sources sparse in direct domain (C.F. appendix C), the backprojected noise has also a sparse representation on Haar transform. Therefore, the detection of the variabilities is hampered if their amplitude is lower than the amplitude of the backprojected noise peaks.
Angle between the reference mixing directions	As said in 5.6, small $\bar{\theta}$ amounts to close mixing directions which yields to poor separation of the components.
Speed of variation of the SV	The spectral variabilities detection requires that they have a lower speed of variation than the sources (C.F. 5.6).
Filtering of the SV	Wider filtering of the SV enhances the morphological diversity between the sources and the variabilities.

Table 5.12: Synthesis of the studies carried out on piecewise constant SV.

5.4.8 1D Monte-Carlo simulations: mixture of three sources and spectral variabilities sparse in DCT domain

We propose to investigate the influence of the condition number in the case of a mixture of more than two sources. To this end, the study is carried out on a dataset where a power law and two emission lines are entangled. The principal setting parameters are provided in 5.13. The variation of the position range of the

Sources	Spectral variabilities			emission line 1	emission line 2	power law
ρ_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Position range	Index range
14%	0.3	3	20	$[0.05, 0.15]$	$[T, T+0.1]$	$[-0.85, -0.75]$

Table 5.13: Parameter settings for the study of the condition number

second emission line allows to generate mixing matrices with different condition numbers. The results provided by svGMCA and L-GMCA feature in figure 5.24. The condition number is related to the correlation of the mixing directions. Since the spectra are normalized, a high condition number implies close mixing matrices which explains the deterioration of the results of both of the algorithms. More particularly, we observe that for a condition number around 18, svGMCA performs as poorly as L-GMCA since the components are too correlated to be unmixed.

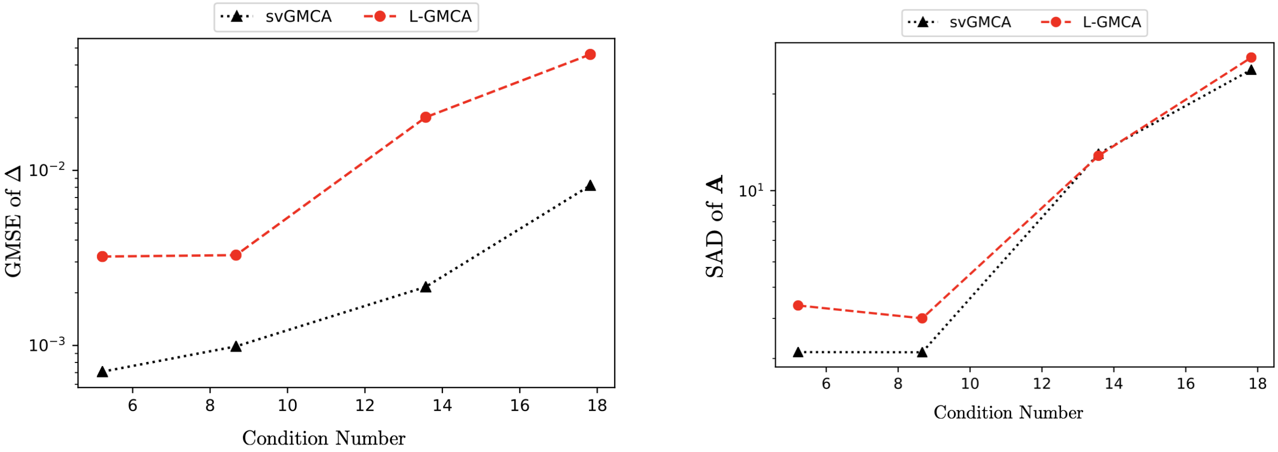

Figure 5.24: Evolution of the performance criteria for different condition number of the mixing matrix.

5.4.9 Comparison with other State-of-the-Art methods

We have presented in chapter 4 three spectral unmixing algorithms that explicitly account for spectral variabilities in remote sensing, PLMM (Thouvenin et al. [2015]), ELMM (Veganzones et al. [2014]) and ALMM (Hong et al. [2018]). Despite their interesting properties for terrestrial images, the algorithms ELMM and ALMM cannot be applied to our datasets (C.F. 4.4.1 and 4.4.3).

Concerning the Perturbed Linear Mixture Model (Thouvenin et al. [2015]), it is not adequate for most of the scenarios we have studied. Indeed, PLMM method requires pure pixel assumption (or at least not too highly mixed pixels) which is not verified for the vast majority of data we have considered. Nevertheless, in order to carry out comparisons of our algorithm with PLMM method, we propose to consider the following dataset that draws near to the scope of application of PLMM method:

- We consider 1D piecewise constant spectral variabilities that mimic the 2D SV generally encountered in terrestrial imaging;
- The sources are highly sparse in direct domain which allows to have samples where only one source is observable.
- The number of samples is $t = 1444$, the number of observations is $m = 5$ and the number of sources is $n = 2$.

The parameter settings are provided in 5.14. The authors of Thouvenin et al. [2015] propose three types

Sources	Spectral variabilities			emission line	power law
$\rho_{\mathbf{S}}$	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
-	0.2	5	20	$[0.85, 0.95]$	$[-0.9, -0.8]$

Table 5.14: Parameter settings

of constraints on the endmember matrix: the penalization on the distance between the endmembers and the spectral signatures brought by VCA algorithm, the volume constraint and the penalization of the mutual distance between the endmembers.
The first two constraints assume the presence of "pure pixels" in the observations, which is excluded in the case of our applications. Therefore, we propose to compare svGMCA to PLMM with spatial smoothness constraint on the endmembers. The PLMM is run with the parameters proposed in Thouvenin et al. [2015] for real data. We have tested the PLMM algorithm with other parameters but they do not show an improvement of the results.

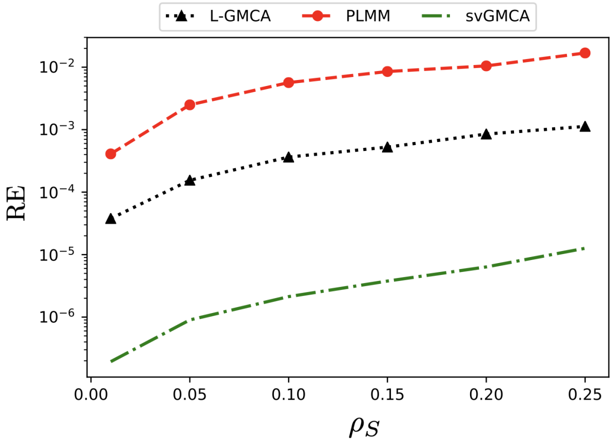

Figure 5.25: Evolution of the performances of PLMM, L-GMCA and svGMCA for varying ρ_S.

We provide in 5.25 a comparison between the three algorithms for varying sparsity degrees of the sources. Given that our method and L-GMCA impose the normalization of the columns on the mixing matrix whereas PLMM enforces the sum-to-one assumption on the abundance matrix, we cannot use a criterion on the spectral variabilities for the comparison. Instead, we assess the performances of the algorithms through the reconstruction error

$$RE = \frac{1}{mt} \|\mathbf{X} - \hat{\mathbf{X}}\|_F^2,$$

where $\hat{\mathbf{X}}$ is the ground truth data matrix and $\mathbf{X}$ the data matrix obtained with the algorithms.

First, we observe that svGMCA outperforms PLMM with a factor up to 10^3. The performance gap is due to the fact that PLMM does not exploit the particular structure of the variabilities. The regularization adopted by PLMM on the SV consists of the penalization of their energy enforced by the minimization of their squared norm. This constraint implies that the SV follow a Gaussian distribution from a Bayesian viewpoint, which is not verified for these data. On the contrary, svGMCA exploits the morphology of the SV through sparsity in undecimated Haar domain. Moreover, the sum-to-one constraint imposed in PLMM on the abundances is not valid for the studied dataset.

L-GMCA, based upon GMCA algorithm that exploits sparsity of the sources, is more performing for low activation parameter. Similarly, svGMCA results deteriorate for high $\rho_{\mathbf{S}}$. Moreover, less sparse source distribution boils down to more bias on the sources which by error propagation hampers the detection of the SV for svGMCA algorithm. Regarding the PLMM method, it is based upon the "pure pixel" assumption which is not verified when the sources are less sparse and therefore more likely to be active simultaneously.

5.5 Application to X-ray images in Astrophysics

Description of the simulated Chandra data

As described in Chapter 1, the Chandra observatory is a X-ray telescope that provides images of the sky in the energy range $[1-8]$ keV. In the present thesis, and in collaboration with A. Picquenot and F.Acero from CEA/DAp, we apply the svGMCA algorithm to supernovae remnants (SNR), and more precisely to the Cassiopea A remnant. This astronomical object is particularly interesting to test the recovery of spectral variabilities since a simple model can already provide a quick description of the observed physical phenomena.
More specifically, at high energy, the observed physics can be simplified to the combination of the synchrotron emission and the single line emission of iron. Furthermore, since SNR originates from a star's explosion, the observed line emission with be either redshifted or blueshifted based on the actual velocity of the observed Fe cloud due to the Doppler effect. This effect, is generally blurred by the integration along the line of sight, which continuously mixes up the total line emissions. However, in the particular case of the Cassiopea A SNR, the iron emission can mainly be decomposed into two slightly overlapping blueshifted and redshifted clouds, as featured in Figure 5.26.
Due to the Doppler effect, the actual shift of the iron emission line directly relates to the velocity field of the gas in the SNR. In that case, and assuming that the SNR explosion is isotropic, the velocity field is spherical, as displayed in Figure 5.28. The synchrotron emission is assumed to have no spectral variabilities. The spectra of the synchrotron and Fe emissions (with maximal blue/redshifts) are displayed in Figure 5.27.
In these preliminary application of the svGMCA algorithm to simulated astrophysical data, Gaussian noise is added. Extending the svGMCA to the case of Poisson noise is left for future investigations.

Setting-up the svGMCA algorithm for the Chandra data

We previously highlighted that a key assumption to be made to apply the svGMCA algorithm is the incoherence of the domain in which the sources are sparse and the domain in which the spectral variabilities are sparse. In the present application, domain incoherence mainly stems from the slowly varying spectral variabilities of the Fe emission, which are directly related to the shift field displayed in Figure 5.28, and the starlet domain, in which the sources are sparse.
Modelling the morphology of the Fe SV could have included a straight modelling of the spherical velocity field of the two redshifted/blueshifted components. In this first application of the svGMCA algorithm to realistic simulation, we rather decided to apply a blind approach with respect to the morphology of the spectral variabilities, assuming that the spectral variabities of the total Fe component are sparse in a 2D discrete cosine transform.
In contrast to the 1D experiments we carried out in the first part of this chapter, an extra difficulty is that the Fe emission is not observed at all in a large portion of the observed field of view, as displayed in Figure 5.26. The estimated SV will therefore be evaluated in regions where the Fe emission is actually observed.
As showed in Figure 5.27, the amplitude of the Fe emission is rather low compared to the synchrotron emission. Therefore, it is key to account for some physical knowledge to provide a reasonable first guess initialization for the svGMCA algorithm. For that purpose, we followed the following step to build such physics-driven initialization:

- **Synchrotron emission :** the synchrotron emission is assumed to have no spectral variabilities. Its contribution in the datacube can therefore be described as the product of a constant spectrum with a spatial amplitude map. Furthermore, it is well established that the synchrotron spectral shape is a power law. A last ingredient is that the synchrotron emission is predominant at low energy. Therefore, the spectrum of the synchrotron can be computed by fitting a power law function from the total spectrum of the datacube (i.e. obtained by summing over all the pixels) for energy lower than 6 keV. The amplitude of the synchrotron emission will be fitted along with the iron emission, as described next.

- **Fe emission :** the Fe emission line stems from the convolution of the spectral instrumental response of the Chandra telescope with a Dirac. A reasonable first order approximation for such instrumental response is a Gaussian kernel, as displayed in Figure 5.27. Therefore, initializing the iron emission is made by fitting a Gaussian-shaped spectrum, with unknown width, mean and amplitude, as well as a power law spectrum with known shape (i.e. the synchrotron spectrum computed at lower energy) but unknown amplitude. This is made more formal in the following equation:

$$s_{\text{observed}} = a_{\text{Fe}} \mathcal{G}(p, \sigma) + a_{\text{Sync}} s_{\text{Sync}} \tag{5.12}$$

where a_{Fe} is the amplitude of the Fe emission, $\mathcal{G}(p, \sigma)$ is a normalized Gaussian shaped signal with mean p and standard deviation σ, a_{Sync} is the synchrotron emission amplitude and s_{Sync} the corresponding normalized spectral shape derived independently at low energy. The parameters are fitted locally on spatial patches of size 8×8 pixels by minimizing standard least-squares problem:

$$\min_{a_{\text{Fe}}, p, \sigma, a_{\text{Sync}}} \left\| s_{\text{observed}} - a_{\text{Fe}} \mathcal{G}(p, \sigma) - a_{\text{Sync}} s_{\text{Sync}} \right\|_2^2 \tag{5.13}$$

The parameters of the svGMCA algorithms are similar to those we fixed for the 1D simulations, with the exception that no reweighting has been applied for these experiments. A second difference is that, since the sources and the spectral variabilities are less incoherent, the threshold we apply to the spectral variabilities in the DCT is slightly larger (i.e. $k = 7$ rather that 3) to limit error propagation from the sources to the SV, and to further impose sparser SV.

Experimental results

In the following experiments, the svGMCA has been tested with three different levels of additive Gaussian noise, from 0 to 20dB. Figure 5.29 displays the input and estimated synchrotron amplitude for a noise level of 10dB. A first visual inspection hardly allows to highlight any difference between the estimated and input synchrotron amplitudes, which suggest that the synchrotron emission amplitude is reasonably estimated. Figure 5.30 further shows the estimated and input synchrotron spectrum, which display a slight difference.
Figure 5.31 features the estimated total Fe amplitude, which exhibits a slight difference in the highest cloud in the field of view, which is mainly associated with the redshifted iron component.
Figure 5.32 shows the estimated spectra across all pixels, as well as the residual with respect to the input Fe spectra. These plots show that the global shape and variabilities of the Fe spectrum are recovered. However, it clearly highlights that some accuracy is lost during the separation process.
Figure 5.33 depicts the estimated iron spectral variabilities map for three different energy bands. It is very interesting to notice that the large-scale variations are accurately captured for these three energy bands. One has further to point out that the main differences can be observed at $6,5773$ keV and $6,6941$ keV in the highest of the three clouds, where finer variations are smoothed out and not correctly recovered. It has to be noticed that these details are mainly related to the redshift component of the Fe contribution, which explains why this is not observed at higher energy. Furthermore, the redshifted component is also the dimmest and therefore the most challenging to capture, as displayed in Figure 5.27.
The following table provides more quantitative results by measuring the MSE of the estimated Fe and synchrotron emissions as well as their SAD. Similarly to the criteria we used in the first part of this chapter, the SAD for the iron emission is taken as the median SAD across all pixels. As expected, this table shows that the recovery results for either the synchrotron emission or the Fe emission are quite reasonable for a SNR of 10 and 20 dB. When the noise level increases, the performances of the svGMCA algorithm rapidly deteriorate. This is particularly clear for the Fe emission.

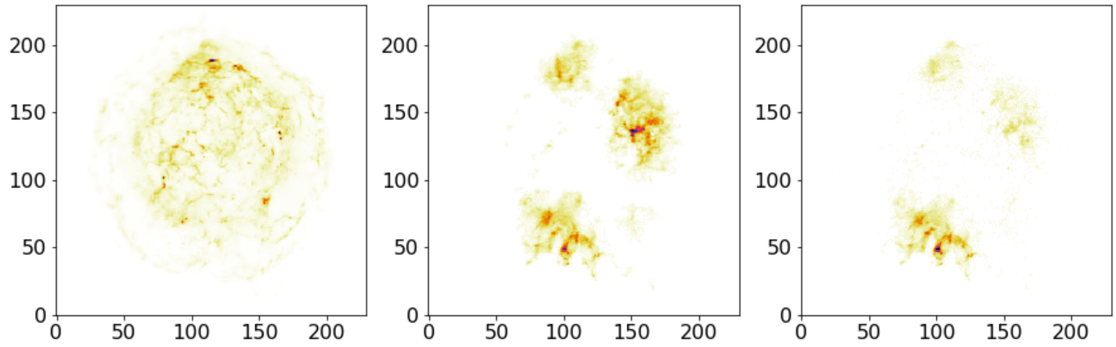

Figure 5.26: Input sources in the Chandra simulated data

	0dB	10dB	20dB
MSE iron amplitude	7.25dB	30.3dB	37.6dB
MSE iron SV	1.58dB	11.7dB	11.8dB
median SAD iron SV	24.8°	2.3°	2.3°
MSE synchrotron amplitude	16.6dB	36.0dB	36.2dB
median SAD synchrotron SV	0.3°	0.17°	0.13°

Conclusion

In this chapter, we have studied the problem of sparse Blind Source Separation in the presence of spectral variabilities. To the best of our knowledge, this is first method explicitly accounting for variabilities in astrophysics. We have proposed to exploit spatial smoothness of the variabilities through a linearization of the variabilities (inspired by the work of Thouvenin et al. [2015])and a sparsity-promoting penalization on the SV term. This regularization allowed us to exploit morphological diversity between the sources and the variabilities to perform source separation and variabilities reconstruction. The designed algorithm, built upon BCD minimization scheme, has been extensively tested and compared with State-of-the-Art techniques. The results of these experiments have shown an improvement on separation quality and helped understanding the behaviour of our algorithm. Finally, we have tested successfully svGMCA on 2D realistic Chandra data.

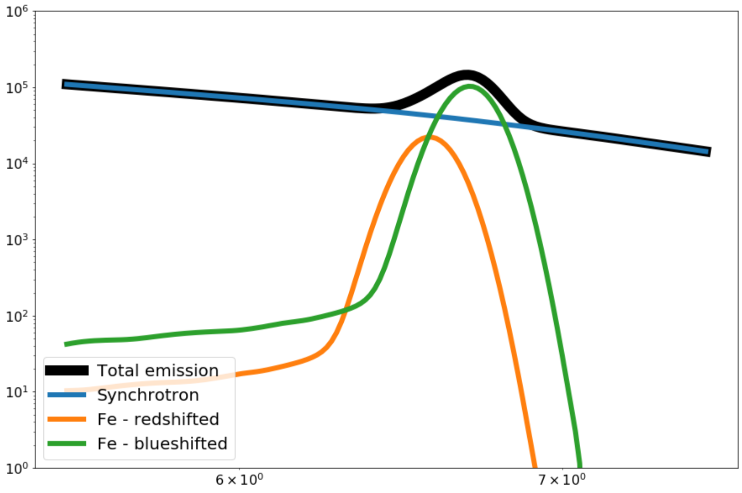

Figure 5.27: Input spectra for the Chandra simulated data. The iron emission lines are represented with the maximum Doppler redshift and blueshift.

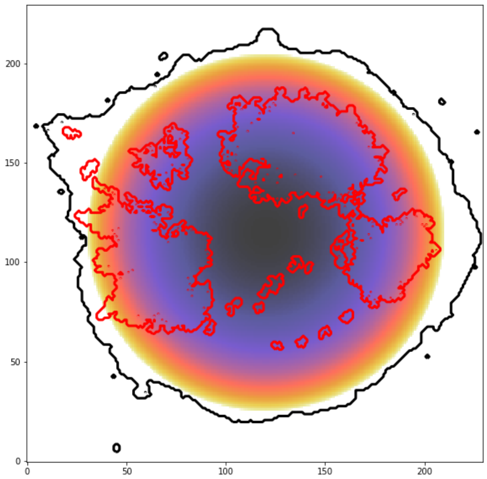

Figure 5.28: Sketch of the spherical spectral variabilities of the Fe emission in the simulated Chandra data. Assuming the SNR velocity field is isotropic, the displayed shift of the iron emission line is spherical. Contour lines display the location of the synchrotron (in black) and Fe (in red) spatial emissions.

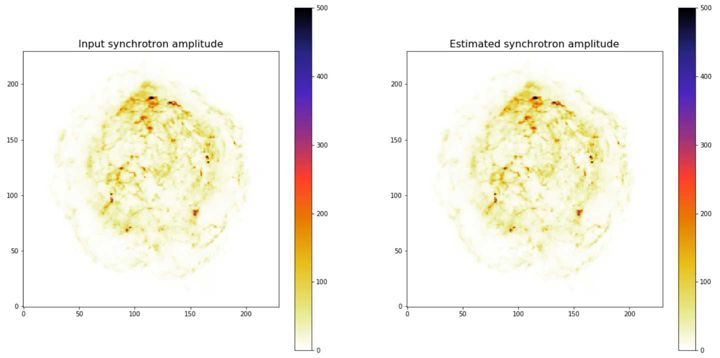

Figure 5.29: Estimated synchrotron emission amplitude - case with $SNR = 10$dB.

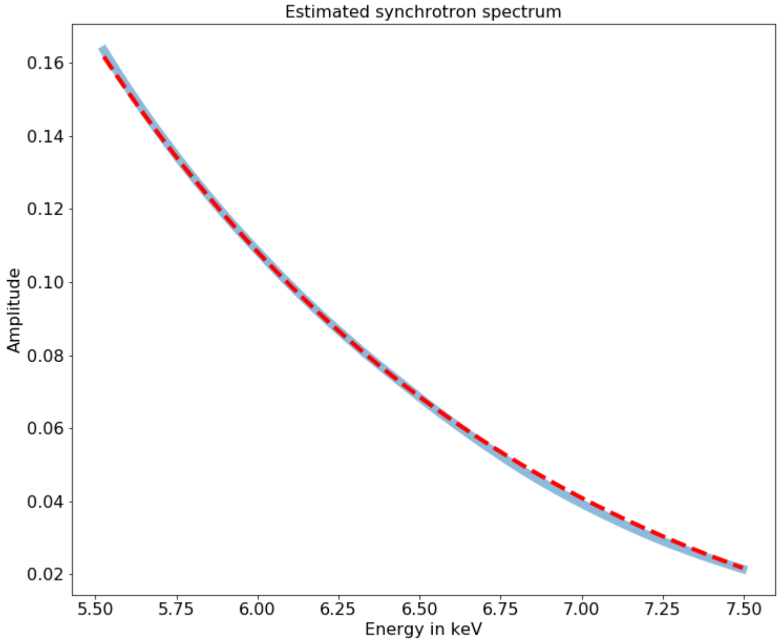

Figure 5.30: Estimated synchrotron emission spectrum - case with $SNR = 10$dB.

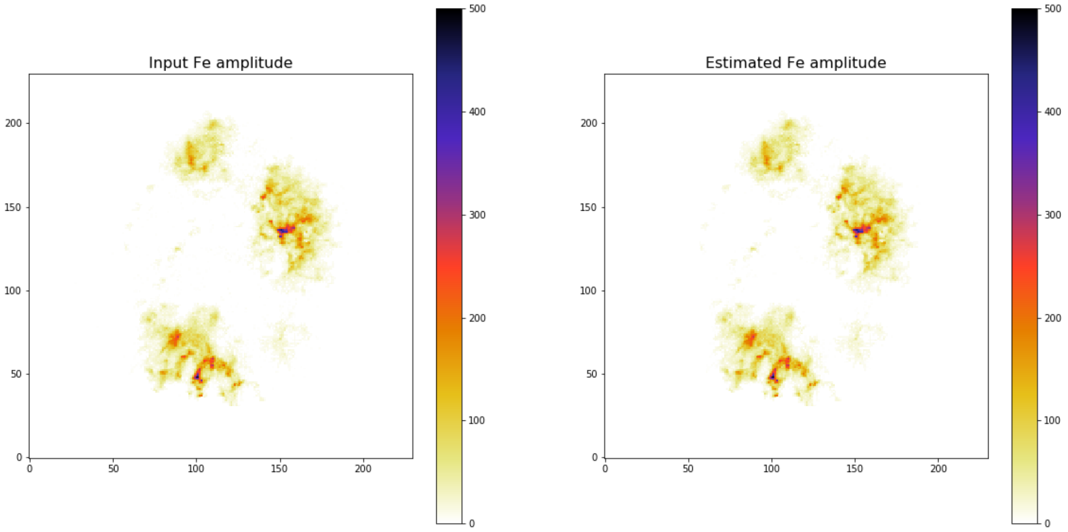

Figure 5.31: Estimated iron emission amplitude - case with $SNR = 10$dB.

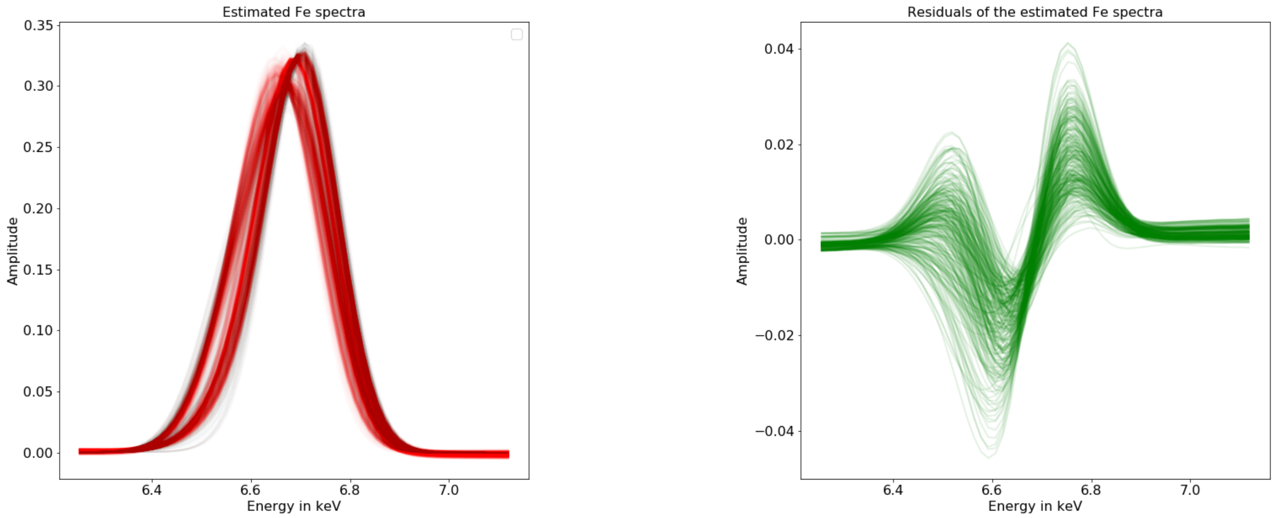

Figure 5.32: Estimated iron emission spectra (left) and residual (right) - case with $SNR = 10$dB.

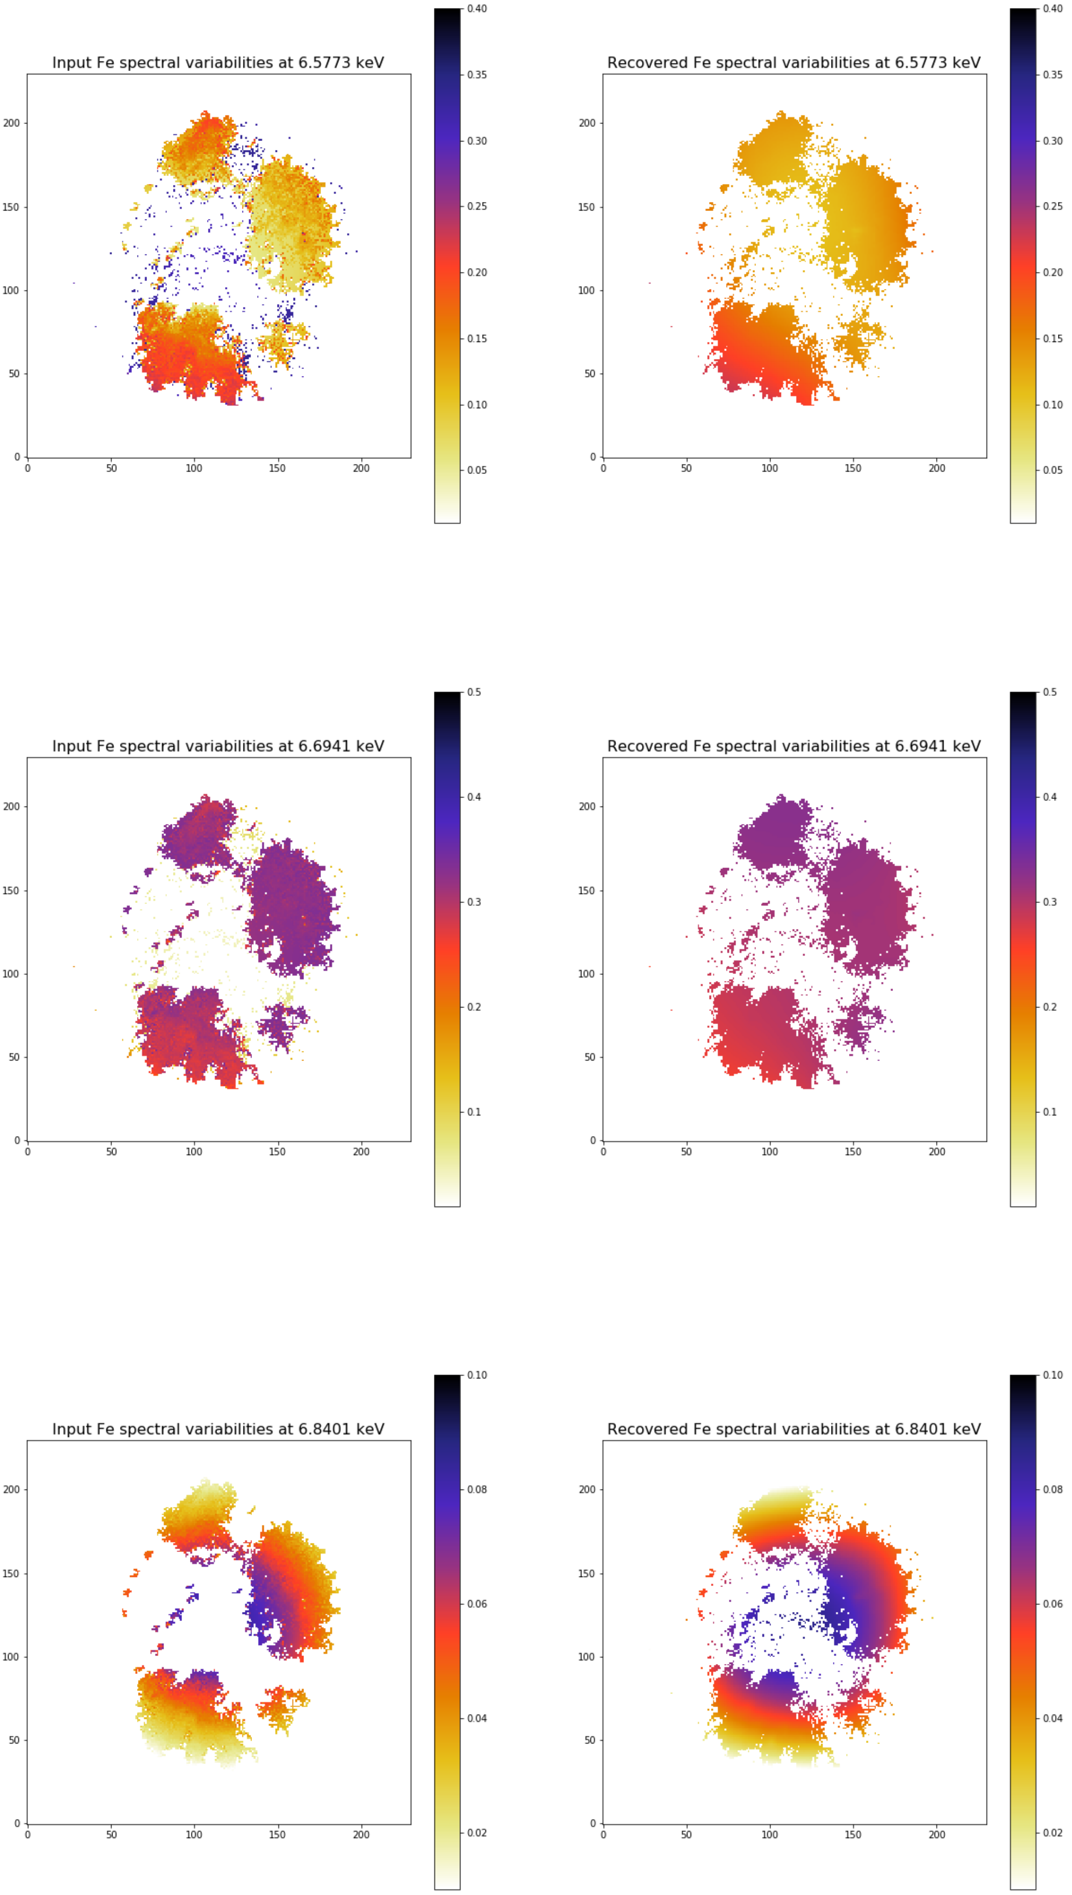

Input Fe spectral variabilities at 6.5773 keV

Recovered Fe spectral variabilities at 6.5773 keV

Input Fe spectral variabilities at 6.6941 keV

Recovered Fe spectral variabilities at 6.6941 keV

Input Fe spectral variabilities at 6.8401 keV

Recovered Fe spectral variabilities at 6.8401 keV

Chapter 6

Conclusion and perspectives

General conclusion

We have entered an new era of proliferation of astrophysical data, requiring efficient analysis tools tailored to their specifics. In this context, Blind Source Separation has proven to be a powerful tool to extract meaningful information from these data. Unfortunately, most BSS methods are inadequate to handle high-energy observations. More specifically, they fail at accounting for the poissonian nature of noise contaminating high-energy observations on the one hand, and for the variabilities affecting the components spectra on the other. In this thesis, we have proposed and studied novel BSS methods aimed at overcoming these pitfalls.

In chapter 3, we have presented a novel sparsity-enforcing BSS algorithm in the presence of Poisson noise or shot noise. The proposed method is based upon an approximation of the optimization problem with Nesterov technique and a minimization of the latter with BCD strategy. The numerical experiments carried out on synthetic 1D spectroscopic data and 2D realistic astrophysical data highlight the ability of the proposed algorithm to separate the components and reconstruct them. Numerical comparison with State-of-the-Art techniques have shown the improvement of separation quality. Our algorithm has already been used with success in the analysis of supernovae remnants by Picquenot et al. [2020]. The work presented in this thesis on sparse BSS with Poisson noise has been reported in an accepted paper Bobin et al. [2020].

The problem of Blind Source Separation in the presence of spectral variabilities has been tackled in chapter 5. We have proposed to exploit the spatial smoothness of the variabilities. Through a linearization of the latter inspired by the perturbed linear mixture model (Thouvenin et al. [2015]), we were able to isolate the SV term and enforce spatial regularization constraint through sparsity of the columns in a transformed domain. Thus, the introduced svGMCA algorithm exploits the morphological diversity between the angular variabilities and the sources to perform source unmixing. We have conducted a thorough study of svGMCA which enabled us to understand its behaviour and limitations. The numerical experiments carried on X-ray astrophysical data have shown its capacity to separate the sources and recover the variabilities on realistic observations.

The methods introduced in this thesis have been designed to be reliable for a wide variety of data and applicable for other fields than astrophysics. For the purpose of participating to reproductible research, the codes will be made available at http://www.cosmostat.org/software.

Perspectives

In this paragraph, we would like to highlight which are the starting points for future investigations on spectral variabilities. These can be decomposed into the two following lines of research:

- **Beyond the linear perturbation mixture model:** In chapter 5, the model we opted for is a standard perturbed linear mixture model (PLMM), which has mainly been investigated in the field of remote sensing. Let us recall that the PLMM boils down to assuming that the actual mixing matrix at a given pixel k is a linear perturbation about some reference mixing matrix $\bar{\mathbf{A}}$:

$$\mathbf{A}[k] = \bar{\mathbf{A}} + \mathbf{\Delta}[k] \tag{6.1}$$

where $\mathbf{\Delta}$ is a perturbation that we assumed to be small in norm. In the propose svGMCA algorithm, this perturbative term is assumed to be spatially sparse in some domain. However, the effectiveness of such sparse regularization strongly relies on the actual sparsity of $\mathbf{\Delta}$. As we illustrated in chapter 5, if the spectral variabilities are generated with sparse angular variabilities, the linearization makes $\mathbf{\Delta}$ less sparse, especially when large angular deviations are observed. Consequently, it would be theoretically more effective to avoid using the PLMM, to stick the non-stationary mixture model. However, this would require to design a sparse modelling dedicated to the field of mixing matrices. This is quite challenging since as these matrices live on the oblique ensemble, which is a property that such a dedicated sparse signal modelling must account for.

A tentative extension would be to make use of the multiresolution representations for manifold-valued data introduced in Rahman et al. [2005] and further extended in Starck et al. [2009]. In a nutshell, such representations built upon a decomposition/reconstruction process that is similar in nature to the a wavelet transform, with the exception that wavelet-like approximations (i.e. low-pass filtering) are computed on the manifold. In Starck et al. [2009], the latter is generally defined as a Fréchet barycenter on the underlying manifold (i.e. the hypersphere in this case), with weights that are defined by standard wavelet-based coefficients. This allows to build a wavelet-like multiresolution representation of manifold-valued data.

If we define by $\mathcal{W}(\,.\,)$ such a multiresolution transform on the oblique ensemble, the cost function of the svGMCA algorithm could be replaced with the following one:

$$\min_{\mathbf{A},\mathbf{S}} \sum_{k=1}^{t} \left\| \mathbf{X}[k] - \sum_{i=1}^{n} \mathbf{A}^i[k]\mathbf{S}_i[k] \right\|_F^2 + \left\| R_S \odot \mathbf{S}\Phi_S^T \right\|_1 + \sum_{i=1}^{n} \gamma_i \|\mathcal{W}(\mathbf{A}^i)\|_{\ell_1} + \iota_{.\geq 0}(\mathbf{A}), \tag{6.2}$$

where the sparsity of the spectral variabilities now directly applies to the mixing matrices. It could be enforced using a standard ℓ_1. Please notice that, there would be no need to impose the sought-after mixing matrices to belong to the oblique ensemble (i.e. to have their columns to belong to the hypersphere). However, this approach would raise difficult challenges; the main one is related to the non-linearity of the transform $\mathcal{W}$, which raises the matter of the convexity of the cost function (when the sources are fixed) as well as which strategy to be deployed to minimize such a cost function.

- **Regularizing the shape of the spectral signatures:** based on the PLMM, the svGMCA algorithm mainly consider spectral variabilities as angular varibilities. As such, imposing the sparsity of the perturbative linear term $\mathbf{\Delta}$ eventually boils down to imposing some sparsity on the angular variations of the spectral variabilities. However, the resulting sparse regularization is not effective enough to avoid or limit the leakage between the sources and their spectra. To that respect, the case of the Chandra data we considered in chapter 5 is particularly illustrative. The model we considered was rather simple; in a more realistic setting more sources will be present, exhibiting a significant amount correlations between them as showed in Picquenot et al. [2019b]. A key element is that these sources have significantly different spectra. In this setting, the main source of diversity between the components are much less their spatial morphologies, and much more the shape of their spectra. A promising approach would therefore to reconsider the regularization of the mixing matrix by rather introducing some constraint onto the shape of their spectra and their variations across pixel.

  In general - as in the Chandra data - physical spectra are likely to live on a low-dimensional manifold. In the scope of Chandra, line emissions are well approximated by Gaussian-shaped spectra with varying position and width - as seen in chapter 5. Since the position and the width are defined according to a one-to-one relationship, the resulting spectral belong to a one-dimensional manifold.

  Accounting for the shape of the spectra can either be made by accounting for some parametric model when available or more generally by making use of some machine learning approach when the underlying manifold is unknown. In such a case, the svGMCA cost function could be replaced by the following one:

$$\min_{\mathbf{A},\mathbf{S}} \sum_{k=1}^{t} \left\| \mathbf{X}[k] - \sum_{i=1}^{n} \mathbf{A}^i[k]\mathbf{S}_i[k] \right\|_F^2 + \left\| R_S \odot \mathbf{S}\Phi_S^T \right\|_1 + \sum_{i=1}^{n} \iota_{\mathcal{M}_i}(\mathbf{A}^i), \tag{6.3}$$

where the constraint $\iota_{\mathcal{M}_i}(\mathbf{A}^i)$ would impose each component to has its spectra living on some known or learnt manifold $\mathcal{M}_i$.

Publications

- J. Bobin, I. El Hamzaoui, A. Picquenot and F. Acero. Sparse BSS from Poisson Measurements. *Accepted in IEEE, Transactions on Image Processing, 2020*. Bobin et al. [2020].
- I. El Hamzaoui and J. Bobin, Sparse BSS with spectral variabilities. *Spars2019, Toulouse, France.*
- I. El hamzaoui and J. Bobin, Sparse component separation from Poisson measurements. *iTwist'18, Marseille, France.* El Hamzaoui and Bobin [2018].
- I. El hamzaoui and J. Bobin, Component Separation from Poisson measurements. *Cosmo21, Valence, Espagne.*

Appendices

APPENDIX A
Proximal operators

The wide majority of proximal operators used in this thesis are explicit. We propose in the table below the most widely used penalization functions in this thesis and their associated proximal operators.

Penalization functions on the source matrix

Penalization function	Proximal operator
$\Vert \mathbf{R} \odot \mathbf{S} \Vert_1$	$\mathcal{S}_{\mathbf{R}}(\mathbf{S})$
$\Vert \mathbf{R} \odot \mathbf{S}\Phi^T \Vert_1$	$\mathcal{S}_{\mathbf{R}}(\mathbf{S}\Phi^T)\Phi$
$\iota_{.\geq 0}(\mathbf{S})$	$\Pi_{.\geq 0}(\mathbf{S})$

Table A.1: Proximal operators associated to common penalization functions applied on $\mathbf{S}$.

Penalization functions on the mixing matrix

Penalization function	Proximal operator
$\iota_{\Vert . \Vert_2 = 1}(\mathbf{A})$	$\Pi_{\Vert . \Vert_2 = 1}(\mathbf{A})$
$\iota_{\Vert . \Vert_2 = 1}(\mathbf{A}) + \iota_{\geq 0}(\mathbf{A})$	$\Pi_{\Vert . \Vert_2 = 1}(\Pi_{.\geq 0}(\mathbf{A}))$
$\Vert \mathbf{R} \odot \mathbf{A}\Phi^T \Vert_{2,1}$	$\hat{\mathbf{A}}\Phi / (\hat{\mathbf{A}})^j = (\mathbf{A}\Phi^T)^j \times (1 - \frac{\mathbf{R}^j}{\Vert (\mathbf{A}\Phi^T)^j \Vert_2})_+$
$\Vert \mathbf{R} \odot (\mathbf{A} - \bar{\mathbf{A}})\Phi^T \Vert_{2,1}$	$\hat{\mathbf{A}}\Phi + \bar{\mathbf{A}} / (\hat{\mathbf{A}})^j = [(\mathbf{A} - \bar{\mathbf{A}})\Phi^T]^j \times (1 - \frac{\mathbf{R}^j}{\Vert [(\mathbf{A} - \bar{\mathbf{A}})\Phi^T]^j \Vert_2})_+$

Table A.2: Proximal operators associated to common penalization functions applied on $\mathbf{A}$.

Appendix B

Monte-Carlo simulations for DCT sparse SV

In order to study empirically the behaviour of svGMCA algorithm presented in chapter 5, we have conducted experiments on synthetic data with Monte Carlo simulations (10 simulations for each parameter). We investigate the influence of the following parameters:

- The sparsity degree of the sources $\rho_\mathbf{S}$;
- The amplitude of the spectral variabilities;
- The sparsity degree of the spectral variabilities;
- The average angle $\bar{\theta}$ between the mixing directions;
- The speed of variation of the spectral variabilities.

The results shown below are consistent with the ones presented in 5.4.6 and the interpretation of the impact of the studied parameters summarized in 5.6 has been derived from the experiments presented in 5.4.6 and the ones that follow. To avoid redundancy, we only highlight the elements that have not been raised in 5.6

B.1 Scenario 1: Mixture of two Emission Lines

In this first part, we consider data generated as follows:

- 2 sources are drawn from a Bernoulli-Gaussian distribution of activation parameter $\rho_\mathbf{S}$. In order to further illustrate the influence of the parameters in different configurations, we will also consider sources with peaks convolved with a Laplacian kernel of full width at half maximum denoted $w_\mathbf{S}$. The number of samples t is fixed at 1500;
- The columns of the mixing matrix are inspired by the iron physical spectrum observed by Chandra telescope. We consider a number of observations of $m = 5$;
- The position range of the spectra are affected by perturbations, generating spectral variabilities. The transform of the SV in DCT domain consists of peaks drawn according to a Bernoulli-Gaussian distribution of activation parameter ρ_{VS} and convolved by a Gaussian kernel of standard deviation σ_{VS} centered around a frequency denoted f.
- The noise is generated according to a Gaussian distribution with zero mean and has a level set so that $\mathrm{SNR} = 65\mathrm{dB}$.

B.1.1 Influence of the sparsity degree of the sources:

Interestingly, the optimal sparsity degree of the sources for the same configuration is higher with sources sparse in direct domain ($\rho_\mathbf{S} = 12\%$) than with sources sparse in starlet domain ($\rho_\mathbf{S} = 4\%$). For the same value of $\rho_\mathbf{S}$, sources with convolutive kernels provide more spatial information on the variabilities than sources that consist exclusively of peaks, enhancing therefore the detection of the SV.

	Sources			Spectral variabilities			Emission Line 1	Emission Line 2
	Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Position range
Dataset 1	No	-	0	0.3	3	20	$[0.05, 0.15]$	$[0.8, 0.9]$
Dataset 2	Yes	-	3	0.3	3	20	$[0.05, 0.15]$	$[0.8, 0.9]$

Table B.1: Parameter setting for the experiment on the influence of the sparsity degree of the sources

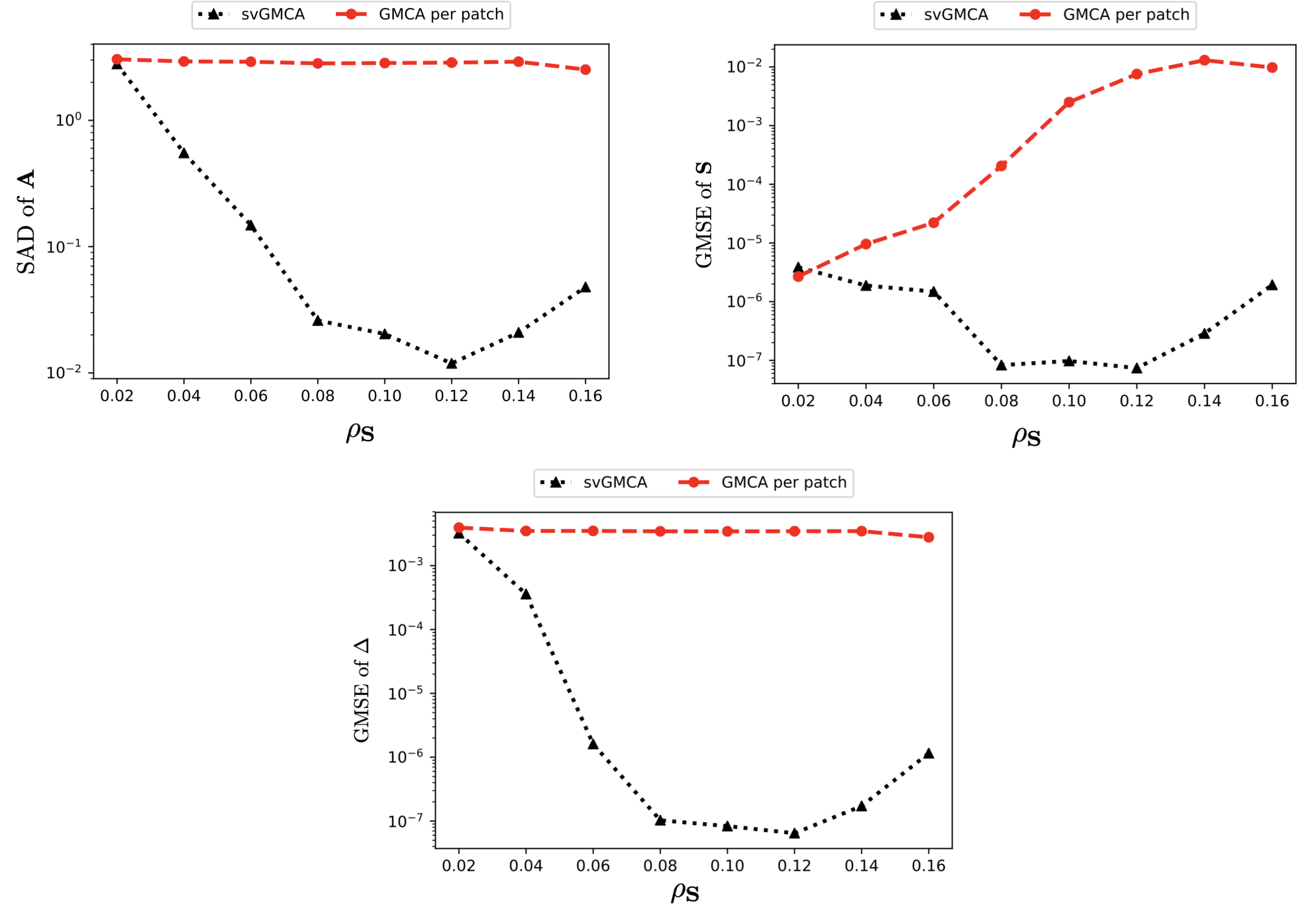

Figure B.1: Influence of the sparsity degree of the sources for the first dataset

B.1.2 Influence of the amplitude of the spectral variabilities

	Sources			Spectral variabilities			Emission Line 1	Emission Line 2
	Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Position range
Dataset 1	No	14%	0	0.3	3	20	$[0.05, 0.05 + l]$	$[0.95 - l, 0.95]$
Dataset 2	Yes	8%	3	0.3	3	20	$[0.05, 0.05 + l]$	$[0.95 - l, 0.95]$

Table B.2: Parameter setting for the experiment on the influence of the amplitude of the SV

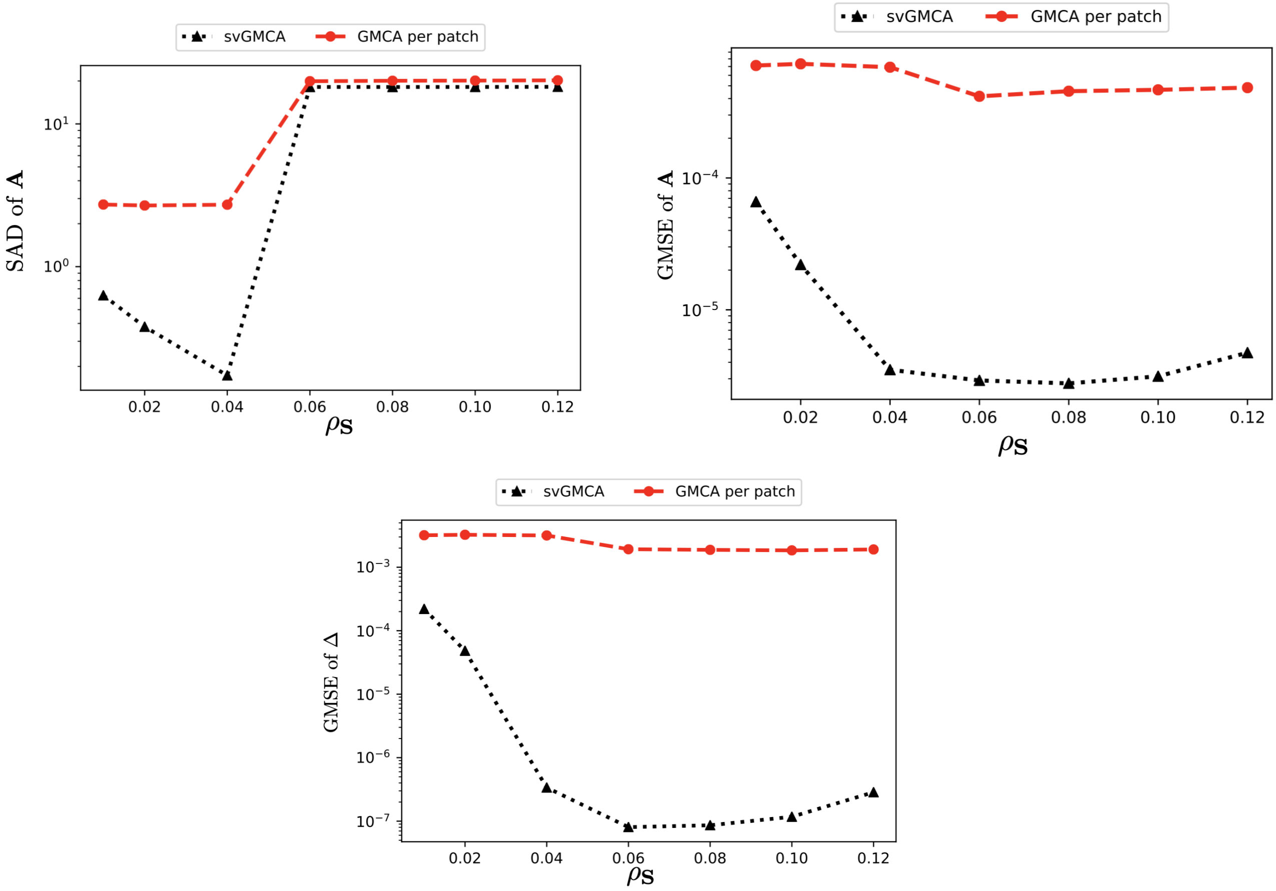

Figure B.2: Influence of the sparsity degree of the sources for the second dataset

B.1.3 Influence of the sparsity degree of the spectral variabilities

	Sources			Spectral variabilities			Emission Line 1	Emission Line 2
	Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Position range
Dataset 1	No	14%	0	-	5	20	$[0.05, 0.15]$	$[0.8, 0.9]$
Dataset 2	Yes	2%	3	-	10	20	$[0.05, 0.15]$	$[0.8, 0.9]$

Table B.3: Parameter setting for the experiment on the influence of the sparsity degree of the SV

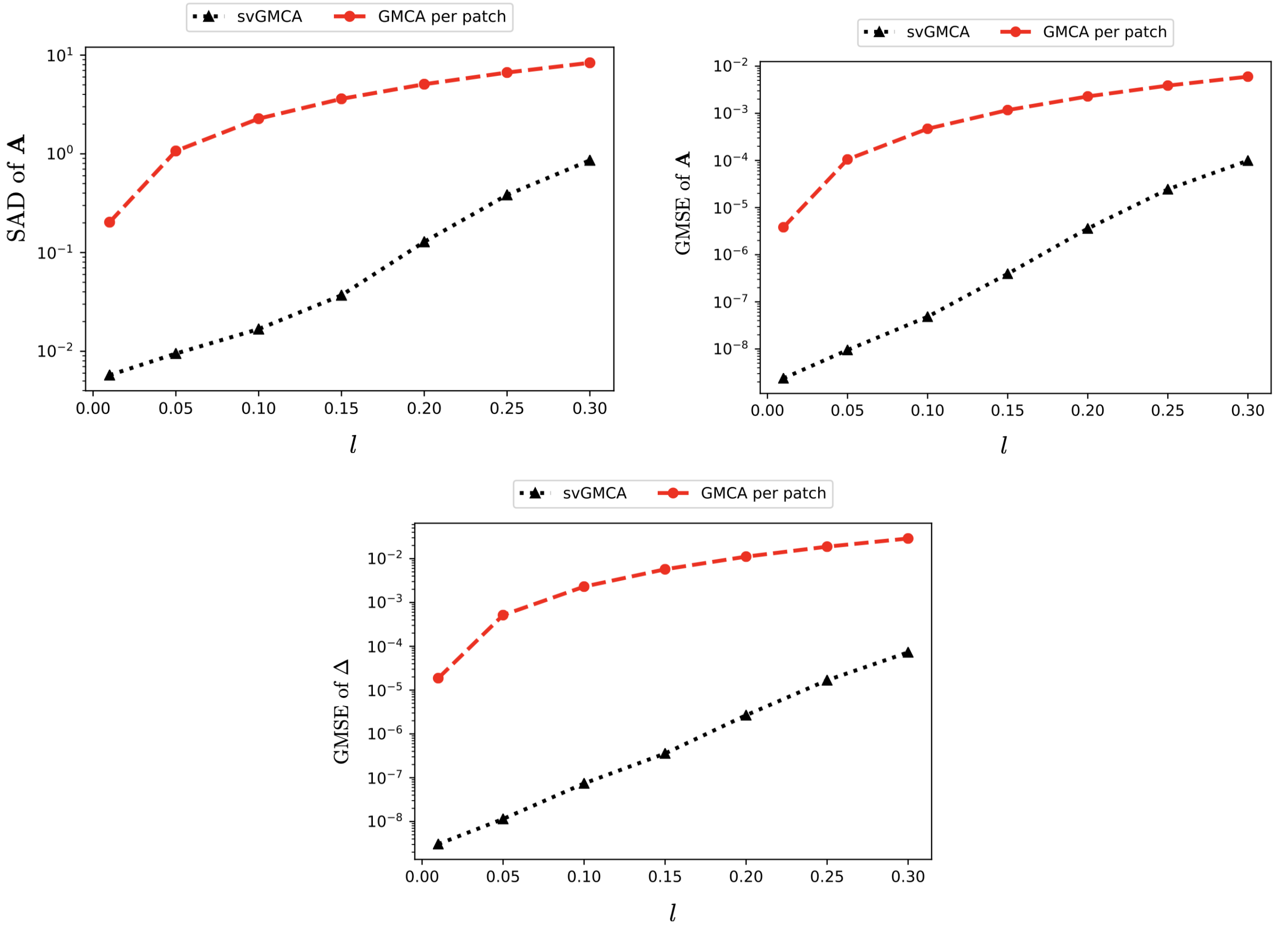

Figure B.3: Influence of the amplitudes of the spectral variabilities for the first dataset

B.1.4 Influence of the average angle between the two mixing directions

	Sources			Spectral variabilities			Emission Line 1	Emission Line 2
	Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Position range
Dataset 1	No	14%	0	0.3	3	20	$[0.05, 0.15]$	$[T, T+0.1]$
Dataset 2	Yes	6%	3	0.3	3	20	[0.05, 0.15]	[T, T+0.1]

Table B.4: Parameter setting for the experiment on the influence of $\bar{\theta}$

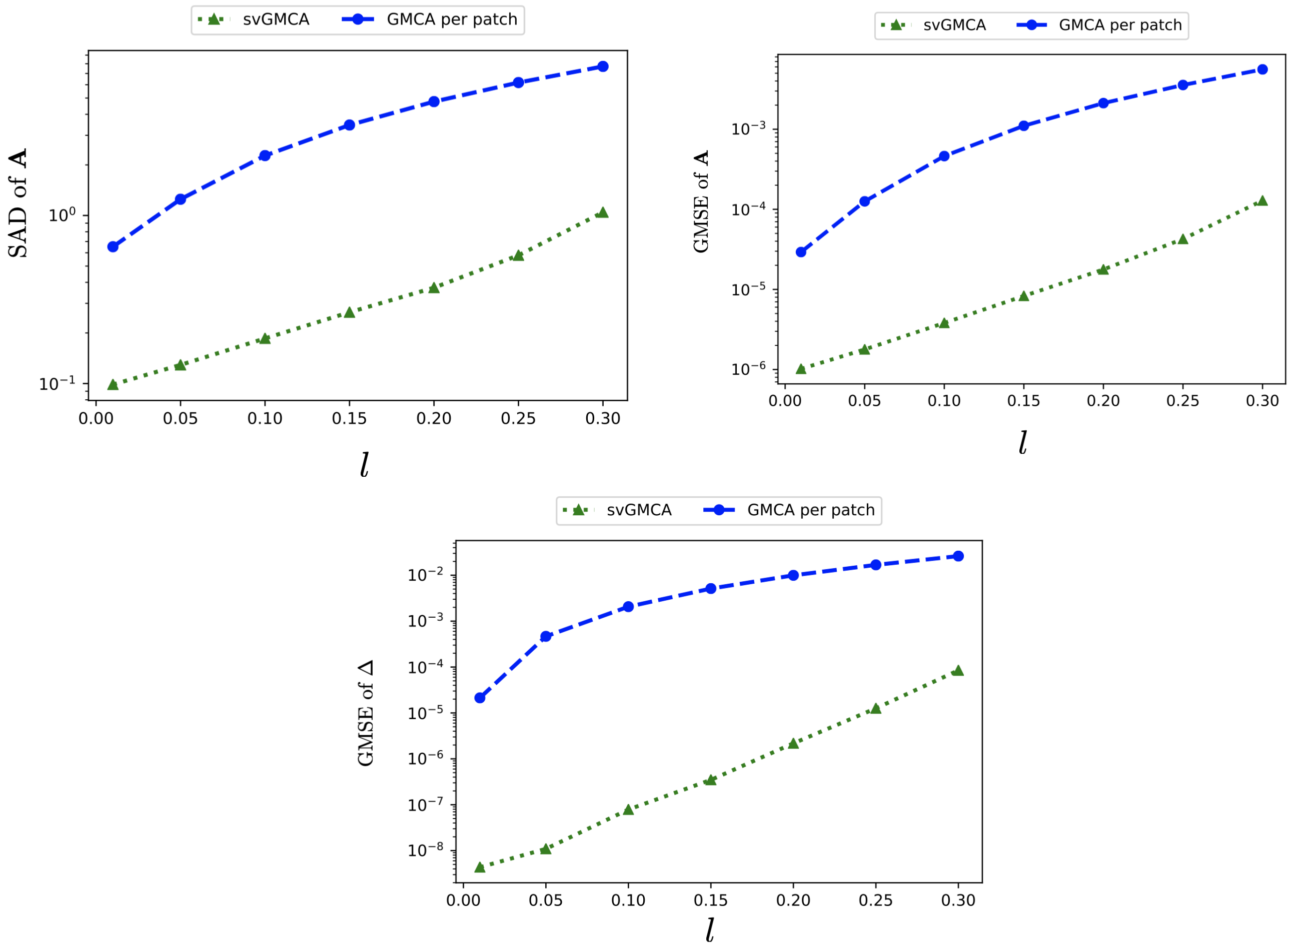

Figure B.4: Influence of the amplitudes of the spectral variabilities for the second dataset

B.1.5 Influence of the speed of variation of the spectral variabilities

	Sources			Spectral variabilities			Emission Line 1	Emission Line 2
	Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Position range
Dataset 1	No	14%	0	0.3	3	-	$[0.05, 0.15]$	$[0.8, 0.9]$
Dataset 2	Yes	6%	3	0.3	3	-	$[0.05, 0.15]$	$[0.41, 0.49]$

Table B.5: Parameter setting for the experiment on the influence of the speed of variation of the SV

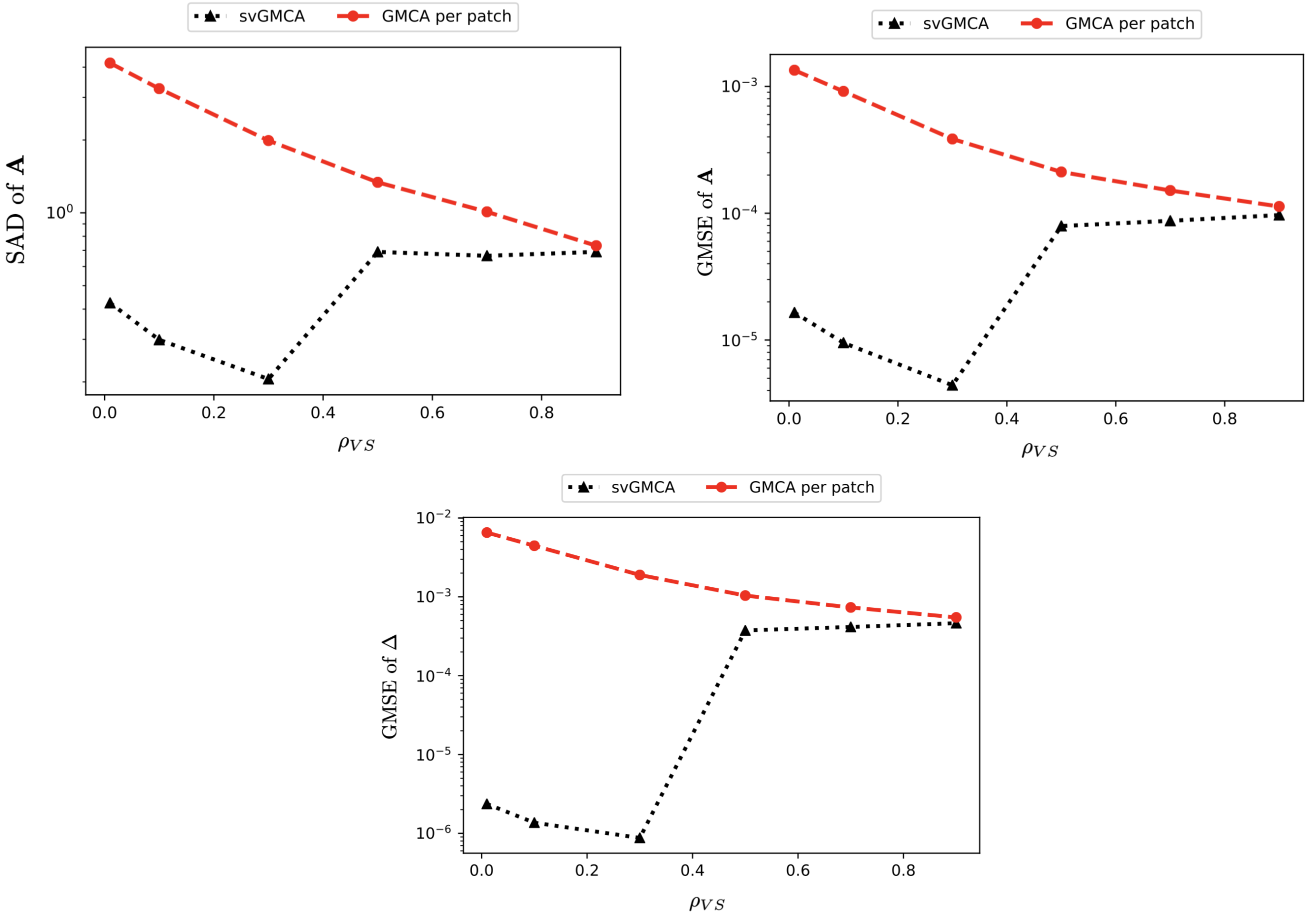

Figure B.5: Influence of the sparsity degrees of the spectral variabilities for the first dataset

B.2 Scenario 2: Mixture of one Emission Line and one Power Law

In this second part, we propose to reconsider the mixing matrix generation process provided in 5.4.6 but this time with sources sparse in starlet domain. Thus, the data are generated as follows:

- We consider the mixture of 2 sources. They are drawn according to a Bernoulli-Gaussian distribution of activation parameter $\rho_{\mathbf{S}}$, the generated peaks are convolved with a Laplacian kernel of full width at half maximum denoted $w_{\mathbf{S}}$. The number of samples t is fixed at 1500;
- The first component is inspired by the iron emission line and the second one by the synchrotron Power Law. The angular variability affects the index of the Power Law and the position range of the emission line. We consider a number of observations of $m = 5$;
- The SV are sparse in DCT domain, their transform consists of peaks drawn according to a Bernoulli-Gaussian distribution of activation parameter ρ_{VS} and convolved by a Gaussian kernel of standard deviation σ_{VS} centered around a frequency denoted f.
- The noise is once again generated according to a Gaussian distribution with zero mean and has a level set so that $\mathrm{SNR} = 65\mathrm{dB}$.

Similarly to what was proposed in the first scenario, we study the impact of the following parameters:

- The sparsity degree of the sources $\rho_{\mathbf{S}}$;

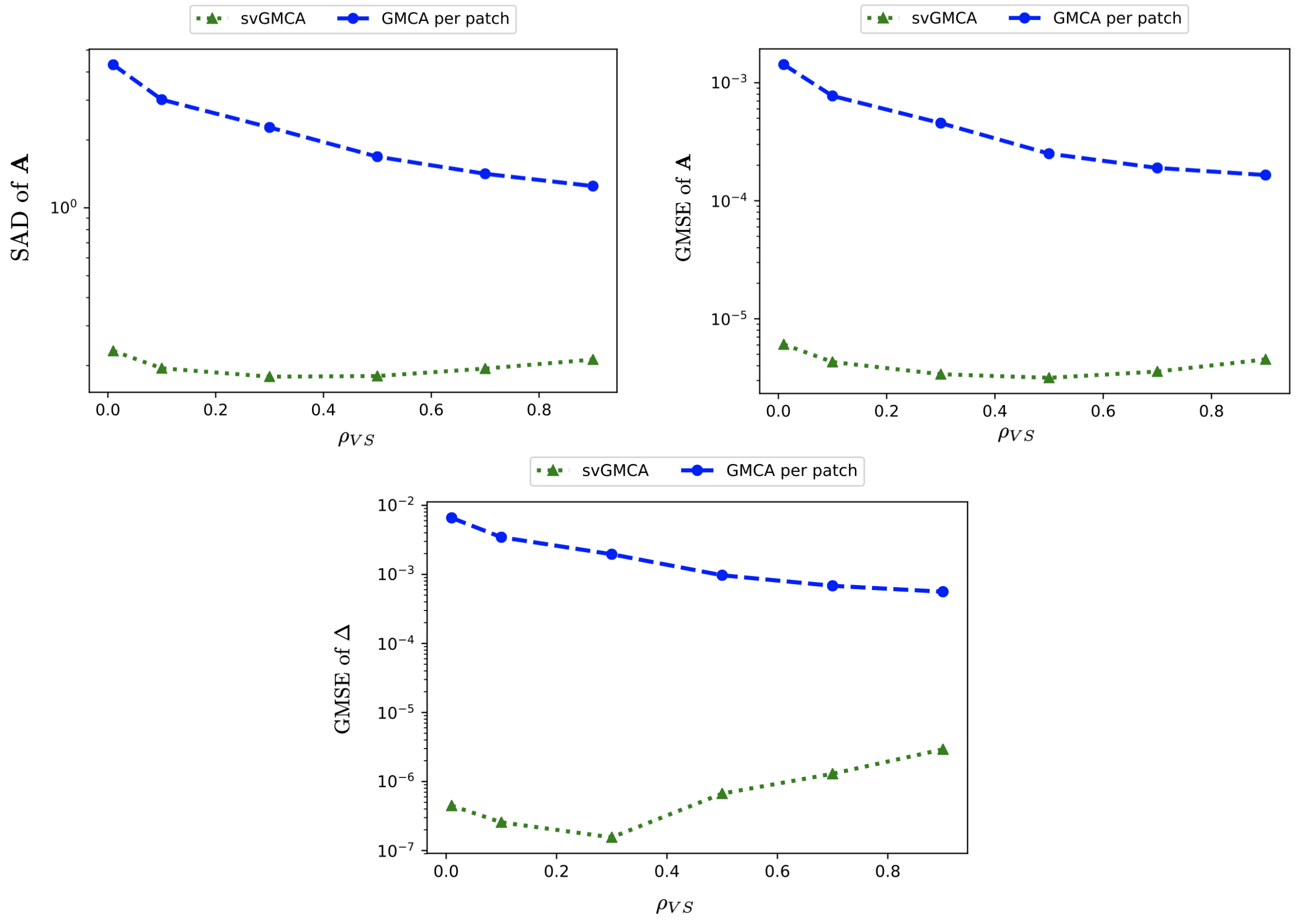

Figure B.6: Influence of the sparsity degrees of the spectral variabilities for the second dataset

- The amplitude of the spectral variabilities;
- The sparsity degree of the spectral variabilities;
- The average angle $\bar{\theta}$ between the mixing directions;
- The speed of variation of the spectral variabilities.

B.2.1 Influence of the sparsity degree of the sources:

Sources			Spectral variabilities			Emission Line	Power Law
Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
Yes	-	2	0.5	5	20	$[0.85, 0.95]$	$[-0.9, -0.8]$

Table B.6: Parameter settings for the study of the impact of the sparsity of the sources

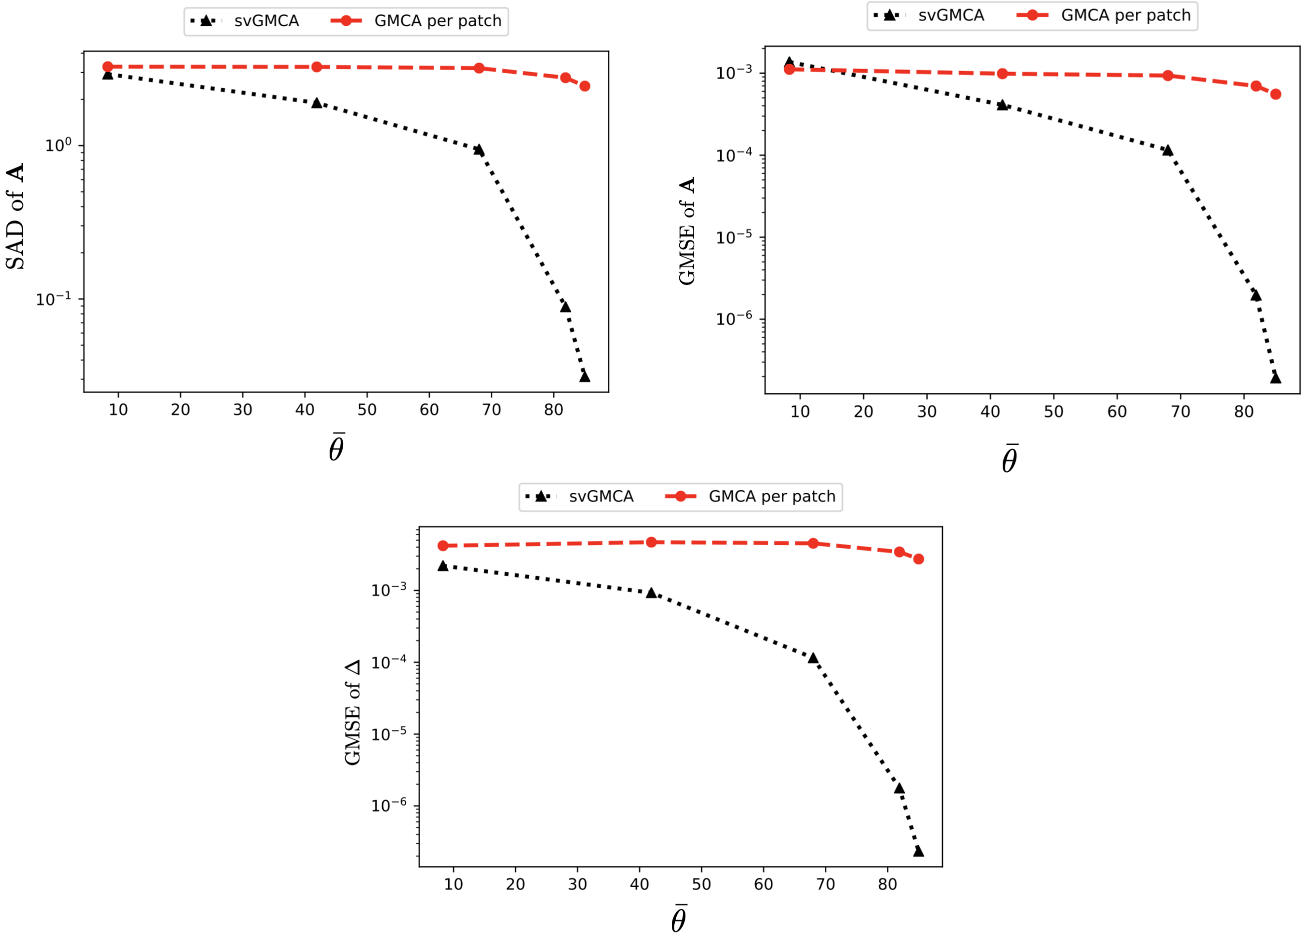

Figure B.7: Influence of the average angles between the two mixture directions $\bar{\theta}$ for the first dataset

B.2.2 Influence of the amplitude of the spectral variabilities

Sources			Spectral variabilities			Emission Line	Power Law
Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
Yes	4%	3	0.3	3	20	$[0.05, 0.25]$	$[0.95 - l, 0.95]$

Table B.7: Parameter settings for the study of the impact of the amplitude of spectral variabilities

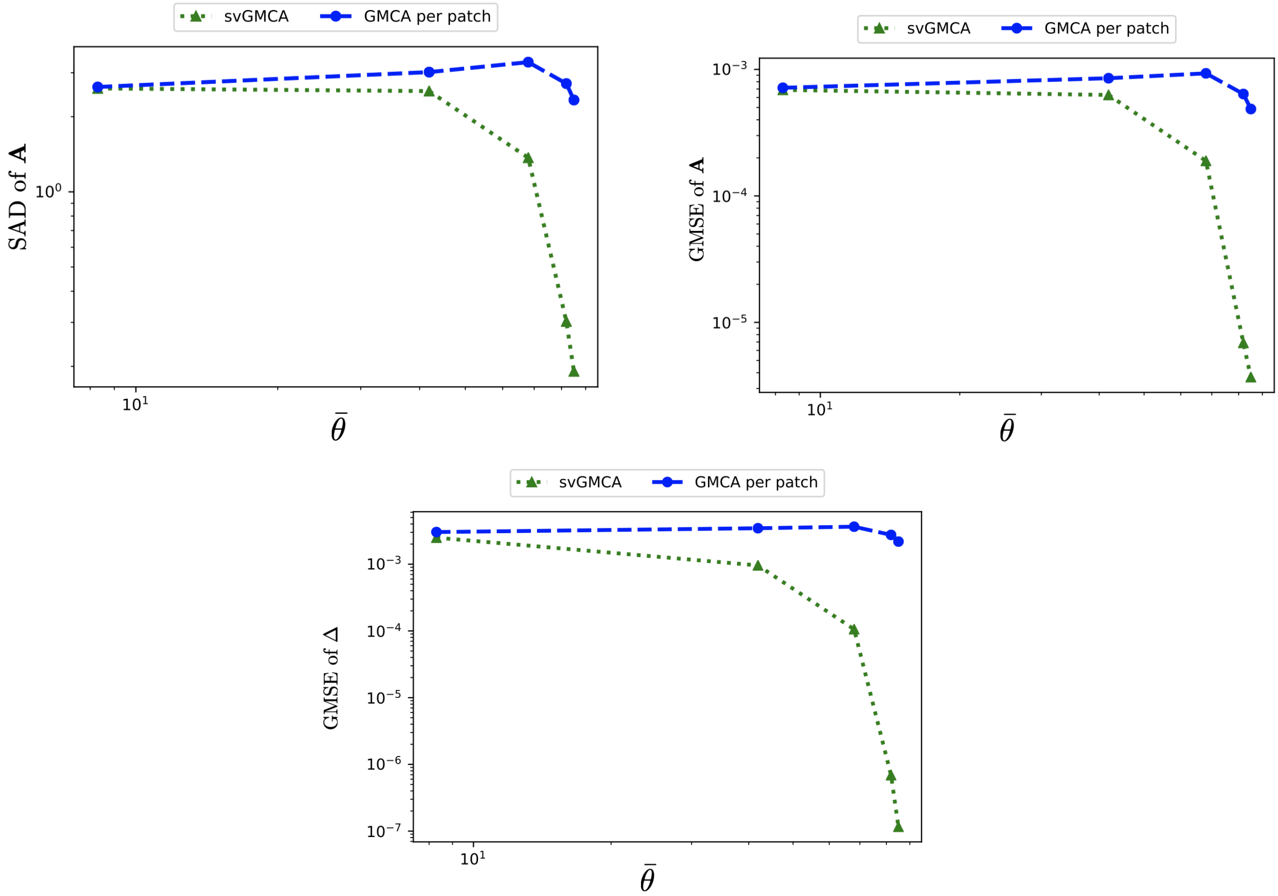

Figure B.8: Influence of the average angles between the two mixture directions $\bar{\theta}$ for the first dataset

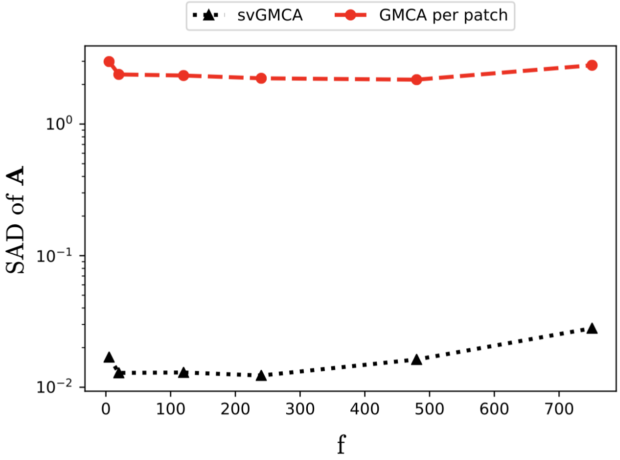

Figure B.9: Influence of the frequency of the spectral variabilities for the first dataset

B.2.3 Influence of the sparsity degree of the spectral variabilities

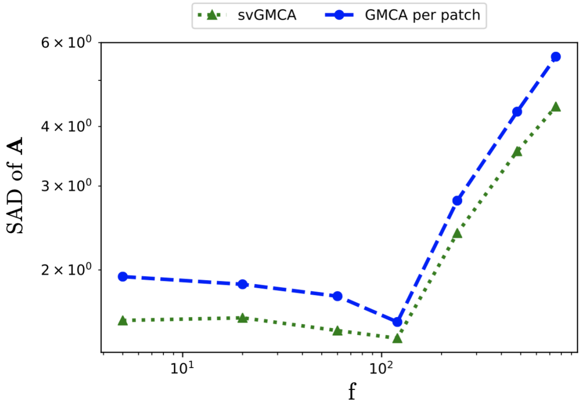

Figure B.10: Influence of the frequency of the spectral variabilities for the second dataset

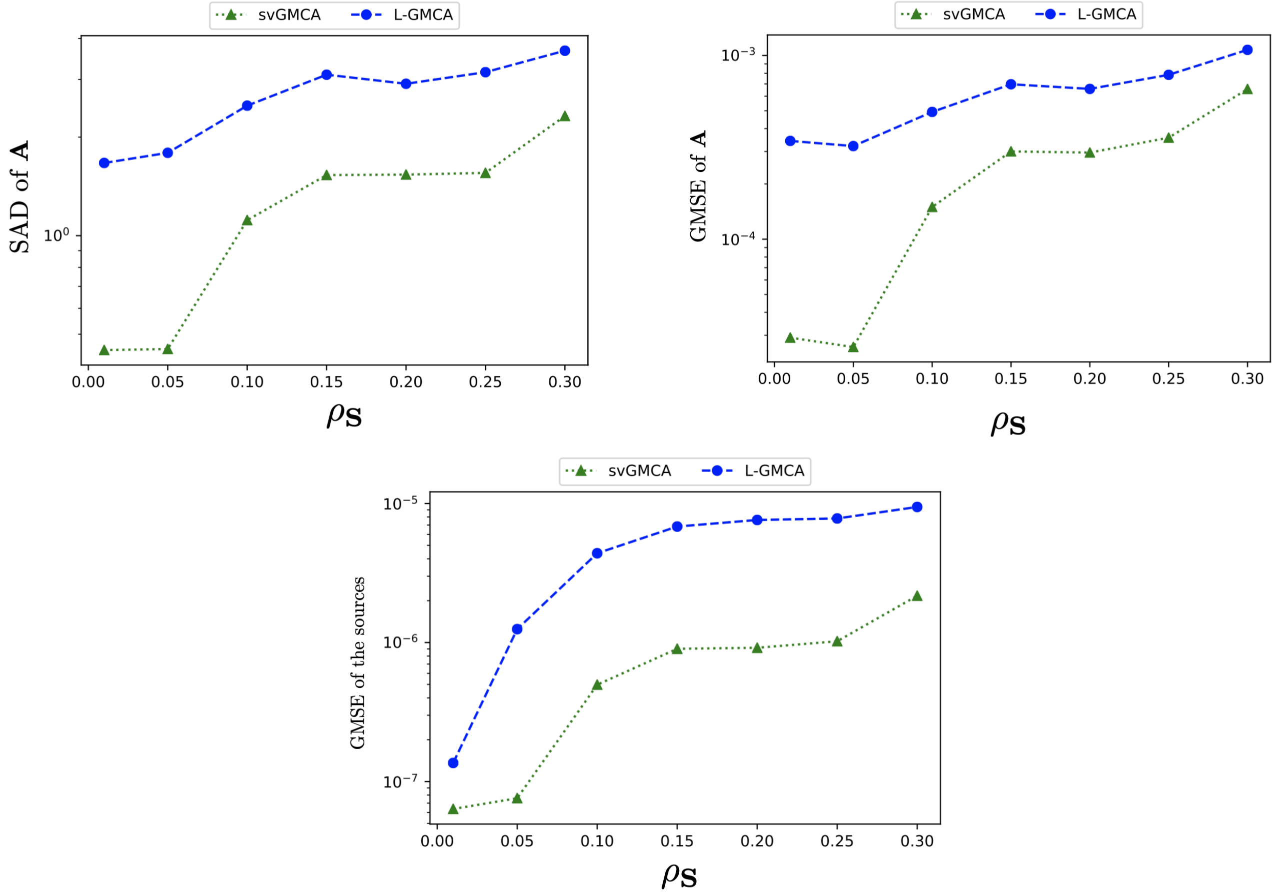

Figure B.11: Influence of the sparsity degree of the sources.

Sources			Spectral variabilities			Emission Line	Power Law
Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
Yes	4%	3	-	5	20	$[0.85, 0.95]$	$[-0.9, -0.8]$

Table B.8: Parameter setting for the study of the impact of ρ_{VS}

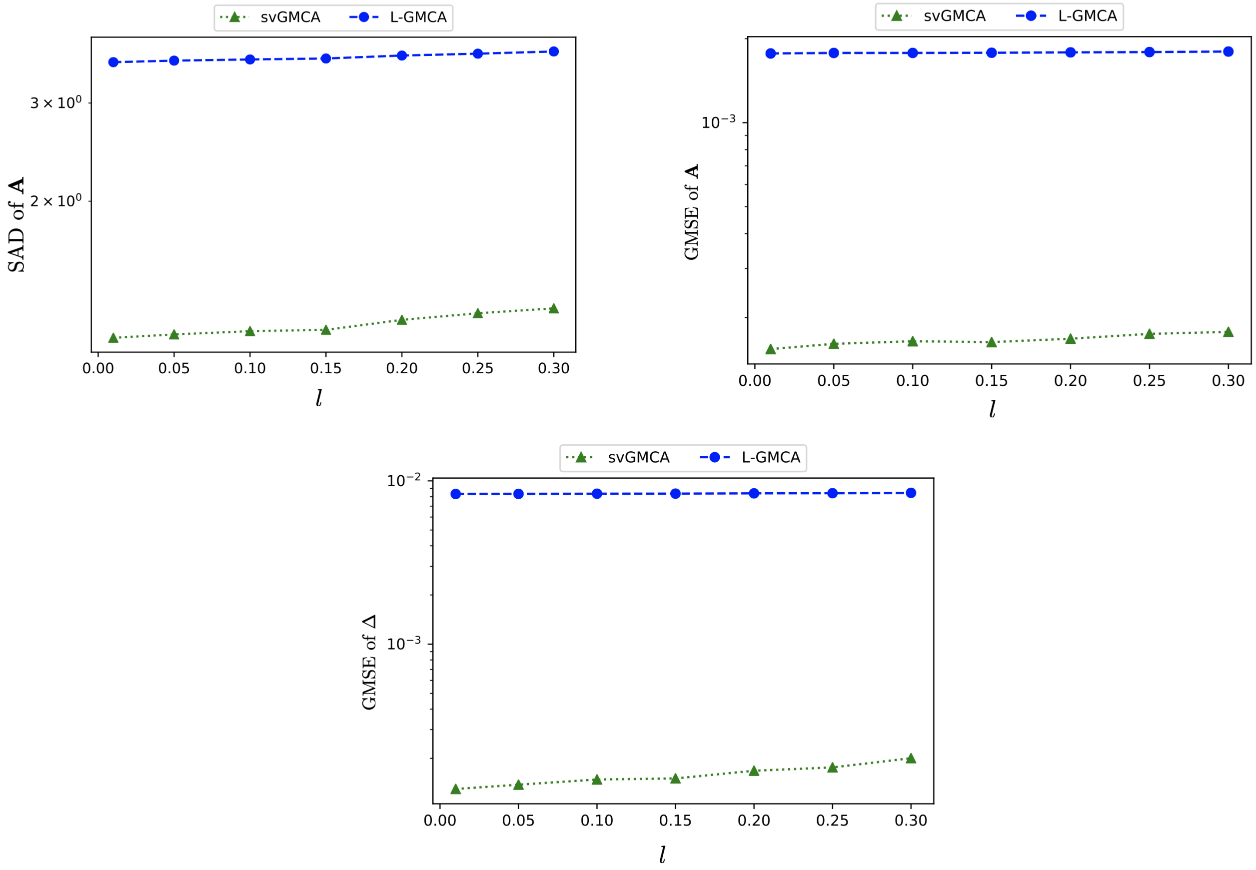

Figure B.12: Influence of the amplitude of the spectral variabilities.

B.2.4 Influence of the average angle between the two mixing directions

Sources			Spectral variabilities			Emission Line	Power Law
Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
Yes	4%	3	0.3	3	20	$[0.85, 0.95]$	$[T, T + 0.1]$

Table B.9: Parameter setting for the experiment on the influence of $\bar{\theta}$

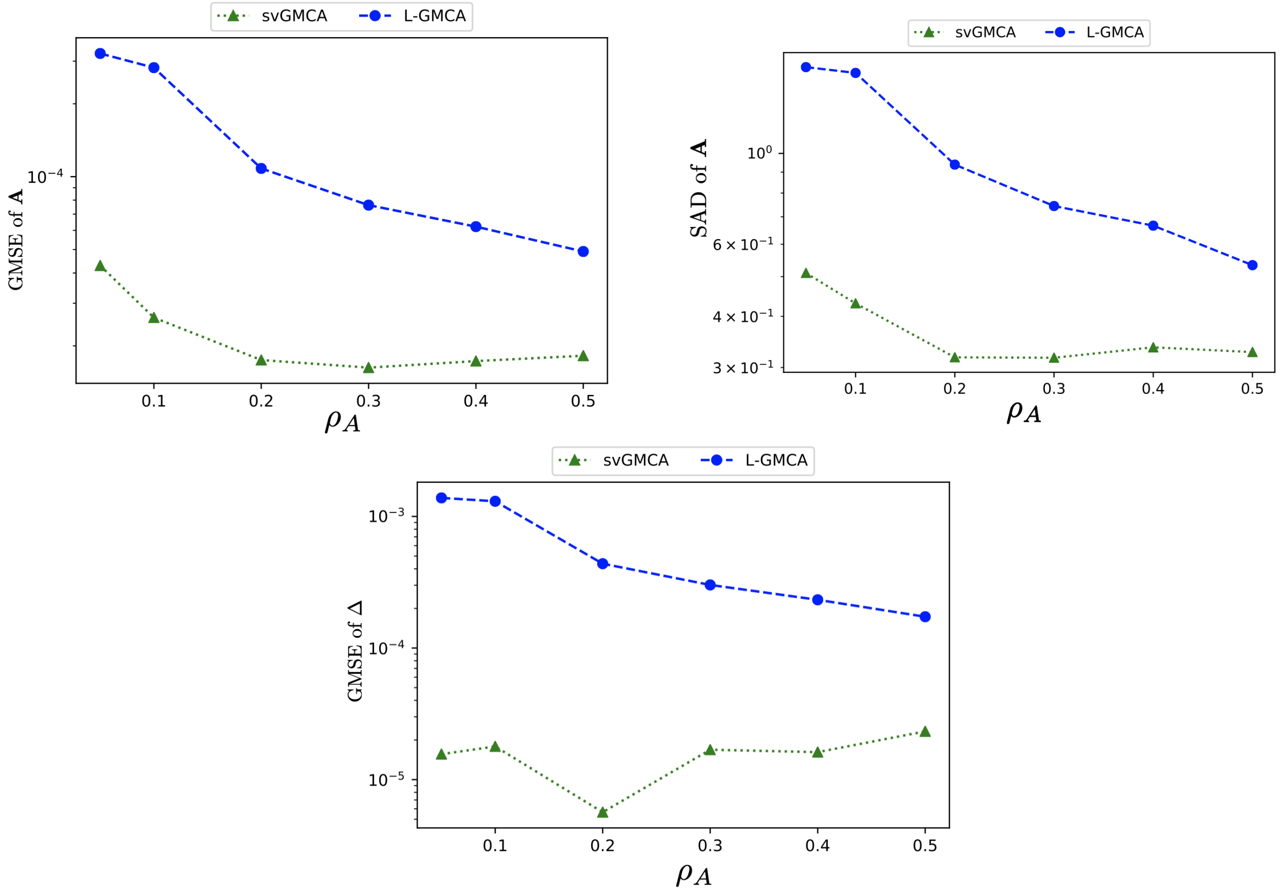

Figure B.13: Influence of the sparsity degree of the SV.

B.2.5 Influence of the speed of variation of the spectral variabilities

Sources			Spectral variabilities			Emission Line	Power Law
Convolution	ρ_S	w_S	ρ_{VS}	σ_{VS}	Position of central peak	Position range	Index range
Yes	4%	3	0.5	3	-	$[0.8, 0.9]$	$[-0.9, 0.8]$

Table B.10: Parameter setting for the experiment on the influence of the speed of variation of the SV

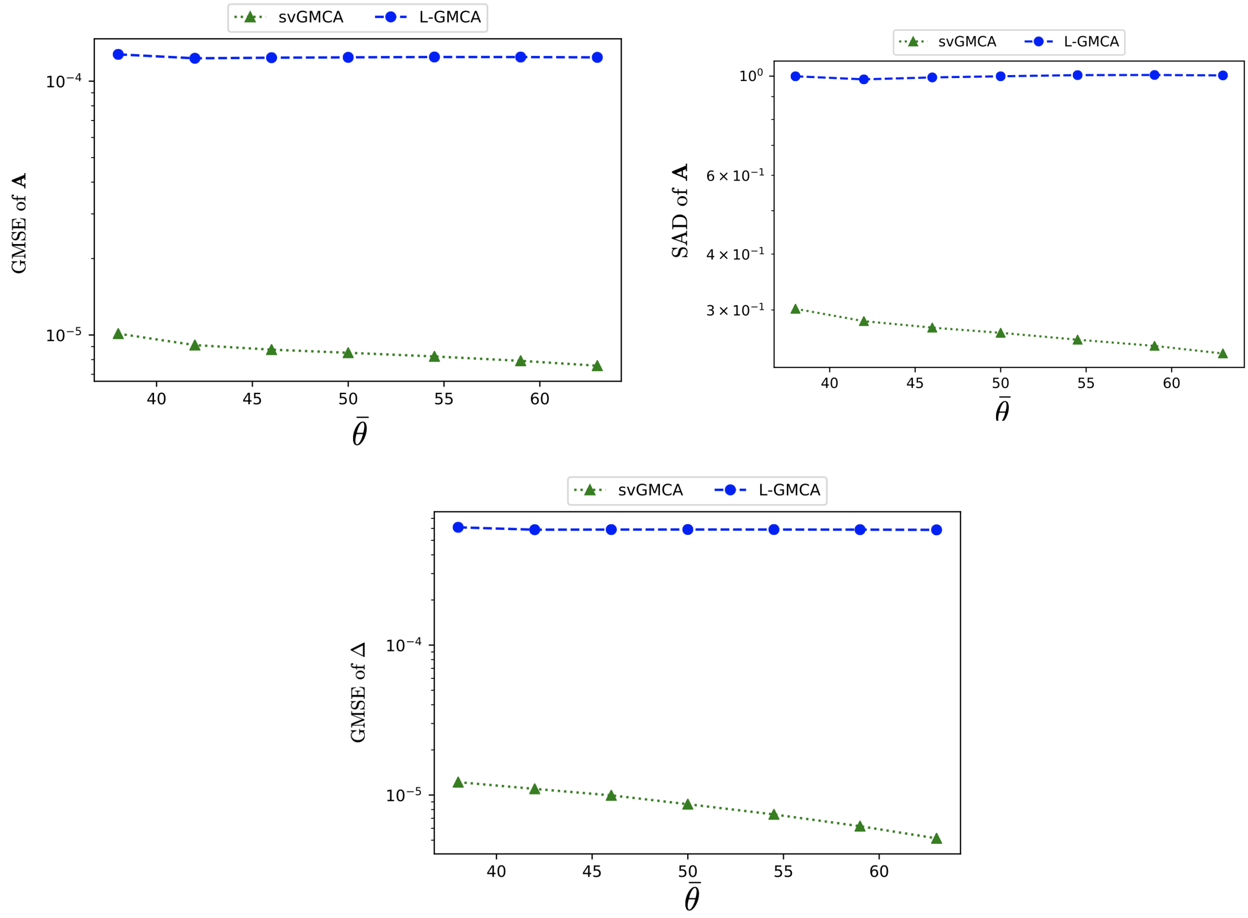

Figure B.14: Influence of the average angle between the two mixing directions.

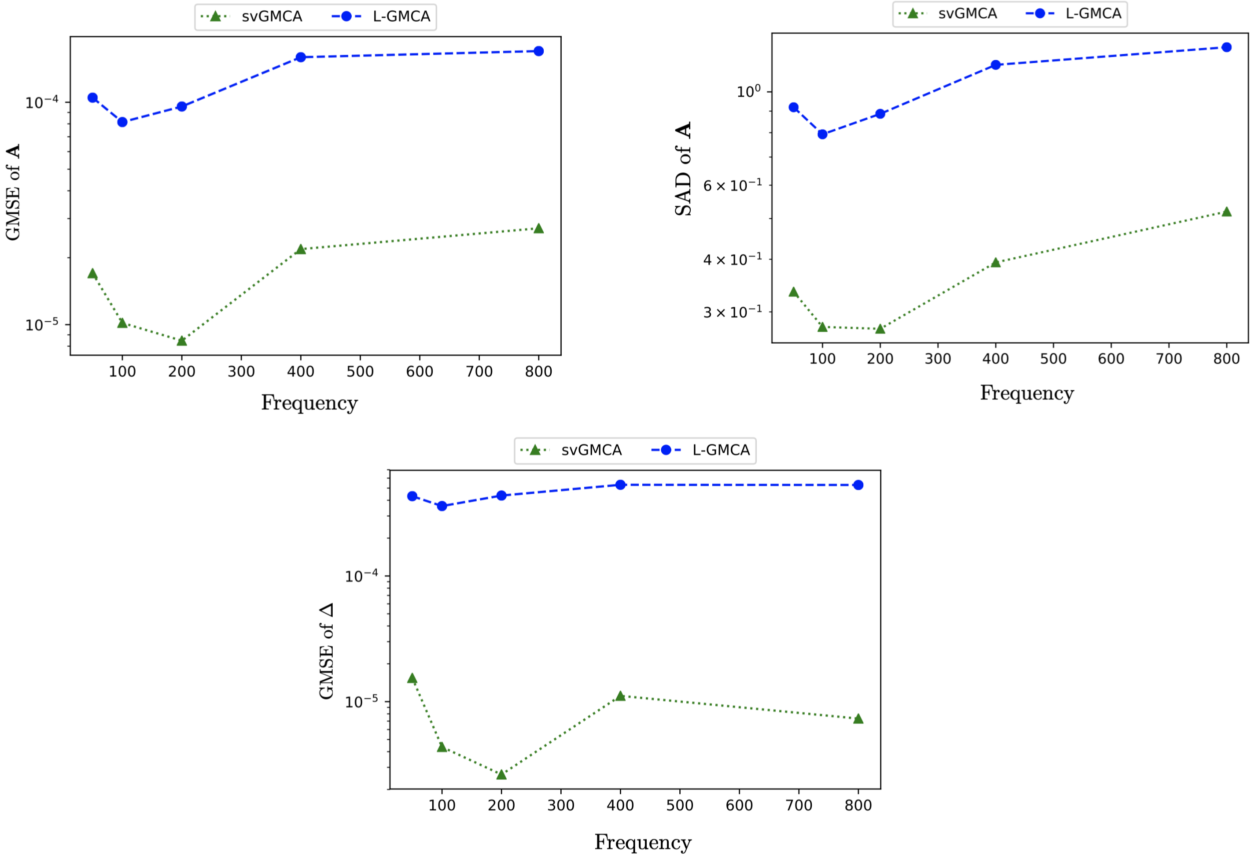

Figure B.15: Influence of the speed of variation of the SV.

APPENDIX C

Monte-Carlo simulations for piecewise constant SV

In order to investigate the performances and limitations of svGMCA algorithm with piecewise constant SV, we have conducted Monte-Carlo simulations with two configuration, sources sparse in direct domain and source sparse in starlet domain (the results of the latter are displayed in 5.4.7). These experiments have led us to the conclusions provided in 5.12.

We provide in this appendix the results of the Monte-Carlo simulations with sources sparse in direct domain. We remind that the dataset is generated as follows:

- We consider the mixture of two emission lines with a width of 0.1.
- The sources are generated according to a Bernoulli-Gaussian distribution of activation parameter $\rho_{\mathbf{S}}$.
- The spectral variabilities are piecewise constant. Given a number N of blocks, we assign to each block a random value between 0 and 1. The blocks are then convolved with a Gaussian kernel. In this study, we focus on the following parameters:
  - The number of blocks N which is an indicator of the speed of variation of the SV,
  - The full width at half maximum (FWHM) of the Gaussian kernel used in case of SV filtering,
- The noise is drawn from a Gaussian distribution with zero mean. The signal-to-noise ratio is provided for each experiment.

As for the configuration where the sources are sparse in starlet domain, we propose to consider the impact of the following parameters:

- The noise level;
- The amplitude of the variabilities;
- The average angle between the two mixing directions;
- The speed of variation of the SV.

C.1 Influence of the input noise

Sources			Spectral variabilities			Emission Line 1	Emission Line 2	SNR (in dB)
Convolution	$\rho_{\mathbf{S}}$	$\sigma_{\mathbf{S}}$	FWHM	m	N	Position range	Position range	
No	14%	0	70	200	3	$[0.05, 0.15]$	$[0.85, 0.95]$	-

Table C.1: Parameter setting for the experiment on the influence of the noise

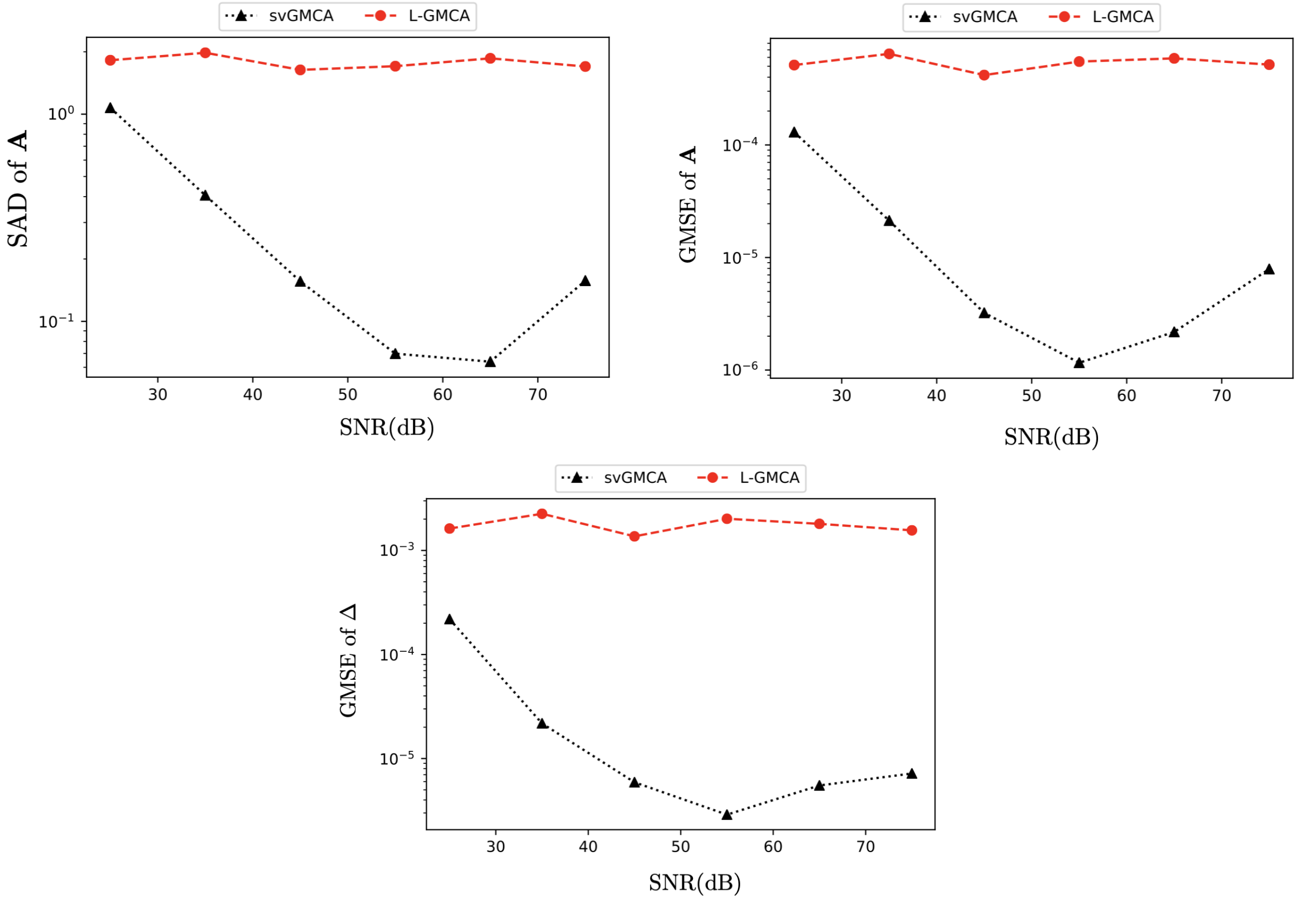

Figure C.1: Influence of the input noise.

C.2 Influence of the amplitude of the SV

Sources			Spectral variabilities			Emission Line 1	Emission Line 2	SNR (in dB)
Convolution	ρ_S	σ_S	FWHM	m	N	Position range	Position range	
No	14%	0	70	200	3	$[0.05, 0.05 + l]$	$[0.9 - l, 0.9]$	45

Table C.2: Parameter setting for the experiment on the influence of the amplitude of the SV

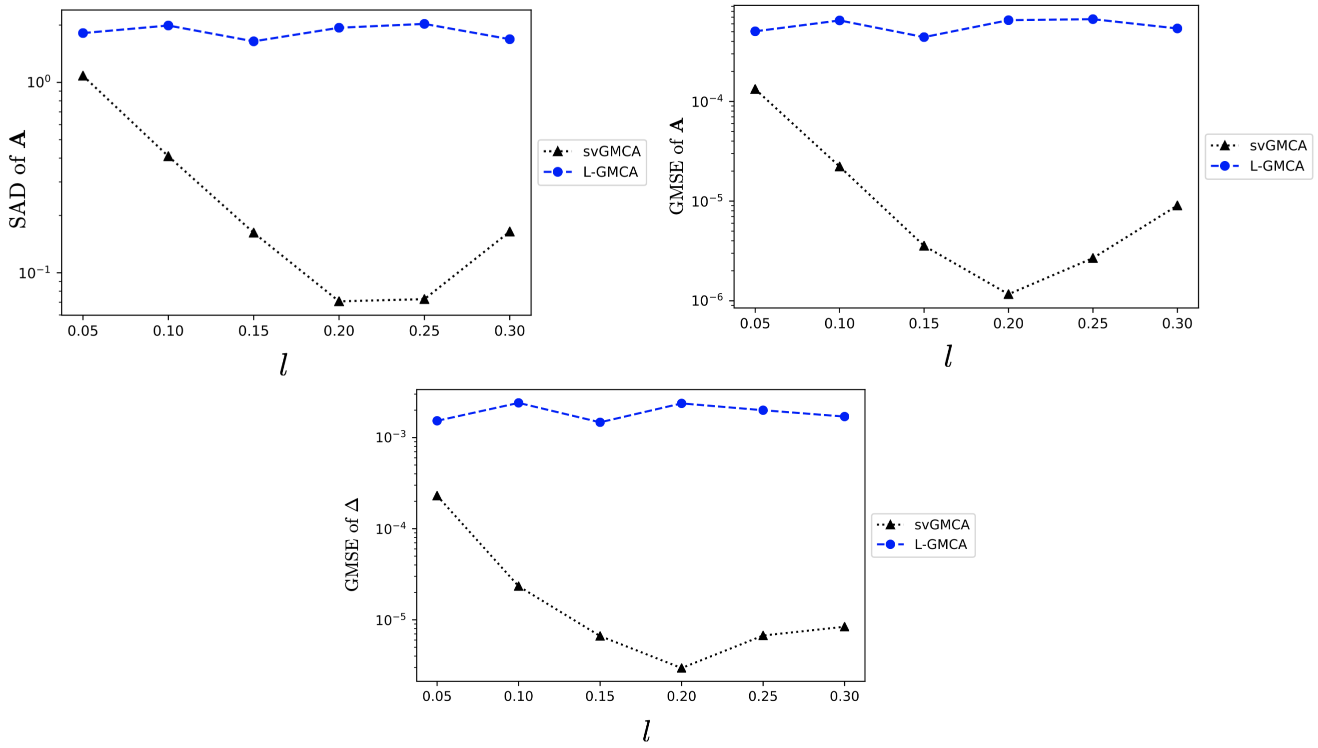

Figure C.2: Influence of the amplitude of the VS.

C.3 Influence of the average angle between the two mixing directions

Sources			Spectral variabilities			Emission Line 1	Emission Line 2	SNR (in dB)
Convolution	ρ_S	σ_S	FWHM	m	N	Position range	Position range	
No	14%	0	70	200	3	$[0.05, 0.25]$	$[0.05+T, 0.25+T]$	45

Table C.3: Parameter setting for the experiment on the influence of $\bar{\theta}$

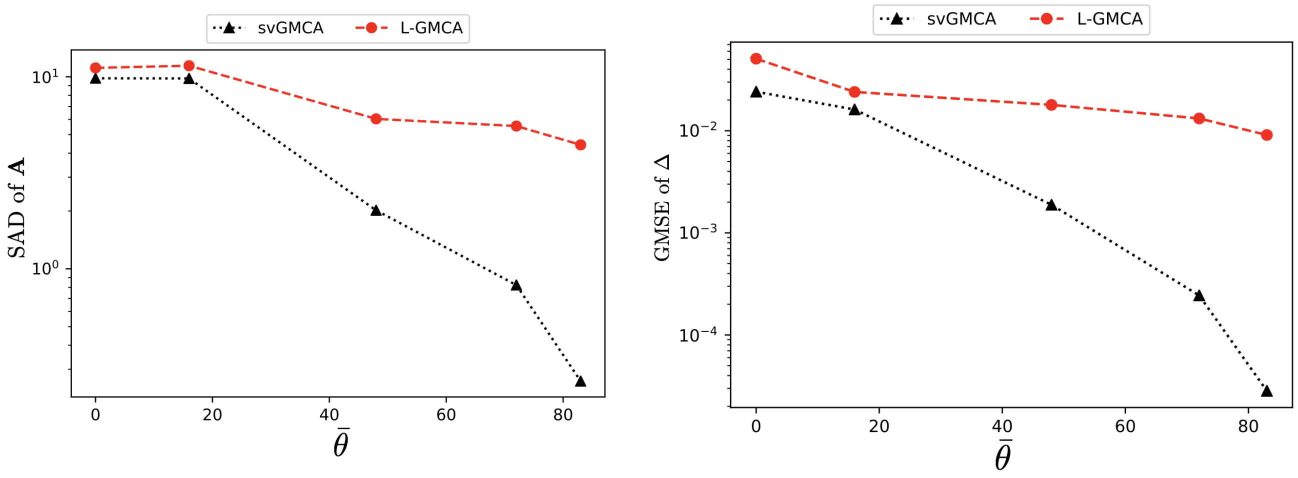

Figure C.3: Influence of the average angle between the two mixing directions $\bar{\theta}$.

C.4 Influence of the speed of variation of the SV

Sources			Spectral variabilities			Emission Line 1	Emission Line 2	SNR (in dB)
Convolution	ρ_S	σ_S	FWHM	m	N	Position range	Position range	
No	14%	0	70	60	-	$[0.05, 0.15]$	$[0.85, 0.95]$	45

Table C.4: Parameter setting for the experiment on the influence of the speed of variation of the SV

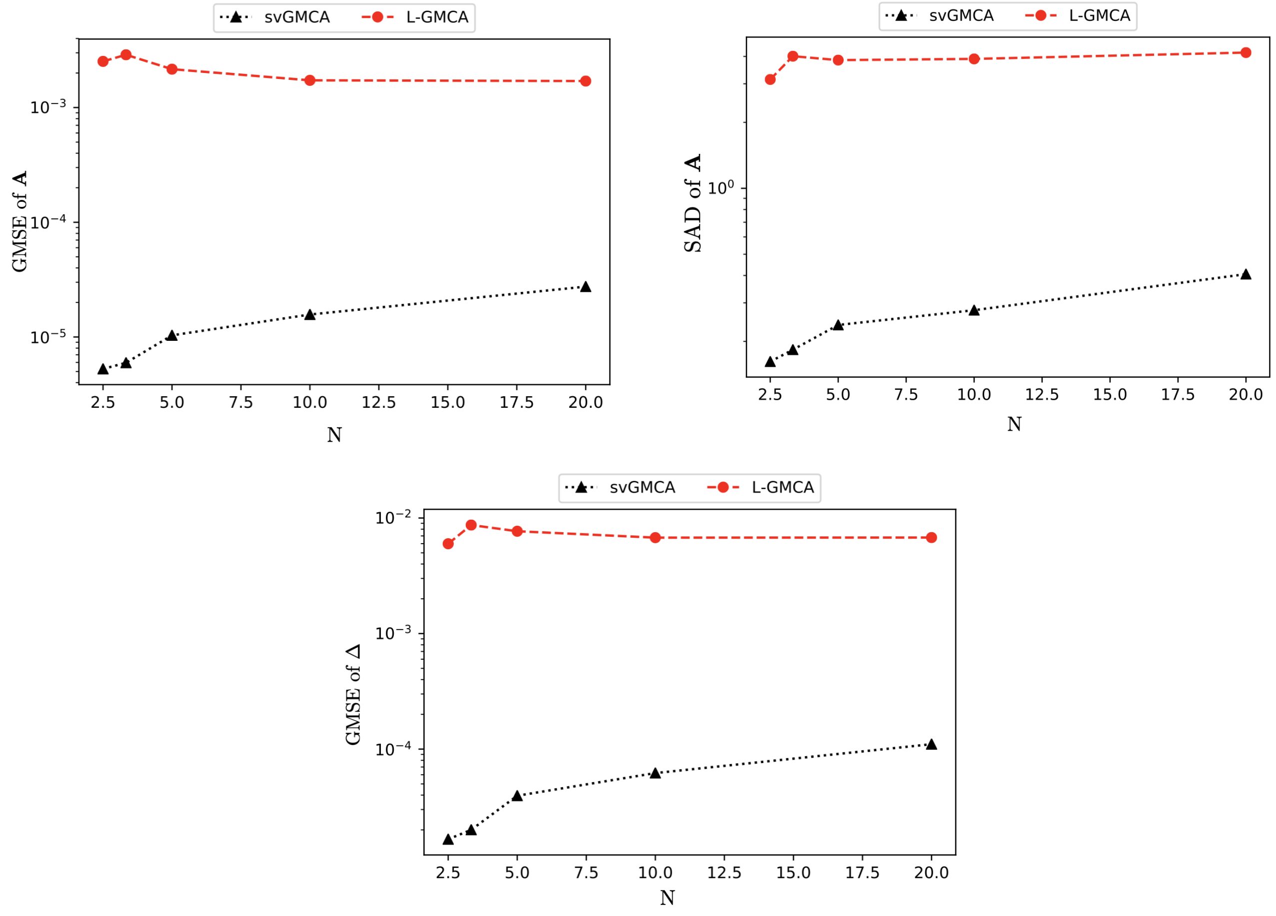

Figure C.4: Influence of the number of blocks.

Appendix D

Impact of the sources distribution in the presence of SV

We have introduced in chapter 5 a novel BSS algorithm accounting for spectral variabilities. In this appendix, we aim at showing empirically the sensitivity of the variabilities estimation to the sources bias.

To that purpose, we propose to study the following datasets:

- We consider the mixture of two sources;
- The spectra are Fe emission lines inspired by Chandra observations. They are Gaussian-shaped with a width of 0.1 and their central frequencies are affected by spectral variabilities;
- The spectral variabilities affecting each Fe line are sparse in DCT domain. Using the same notations as in chapter 5, their characteristic features are given by
    - $\rho_{VS} = 0.3$,
    - $\sigma_{VS} = 3$,
    - central peak at the 20^{th} coefficient.
- The sources are sparse in direct domain. They are generated from a Generalized Gaussian distribution of activation parameter $\rho_{\mathbf{S}}$. We have chosen to perform the test for two activation parameter $\rho_{\mathbf{S}} = 10\%$ and $\rho_{\mathbf{S}} = 24\%$.

In order to investigate the impact of the bias on the sources, we propose the following procedure:

```
Algorithm 14: Experiment procedure
Initialization: $\mathbf{A}^{(0)} \leftarrow \mathbf{A}^*$ , $\mathbf{S}^{(0)} \leftarrow \mathbf{S}_{\text{L-GMCA}}$
• Estimation of the source matrix following the minimization scheme proposed in 5.3.2 with
  regularization parameter $\mathbf{R}_{\mathbf{S}}$ defined as $\mathbf{R}_{\mathbf{S}} = \mathrm{Diag}(\lambda_1, ..., \lambda_n)$ and $\lambda_i$ as
      $\forall i = 1..n,\ \lambda_i = \kappa 1.48 \mathrm{MAD}(\alpha_{\tilde{\mathbf{S}}_i})$
• Estimation of the mixing matrix following the minimization scheme proposed in 5.3.3.
```

This scheme is applied to each dataset with different values of κ, which amounts to different thresholds applied on the sources, and therefore different biased sources used for the final estimation of $\mathbf{A}$. We highlight that no reweighting procedure has been implemented on the sources.

For each dataset (*i.e.* $\rho_{\mathbf{S}} = 24\%$ and $\rho_{\mathbf{S}} = 10\%$), we have carried out 10 Monte-Carlo simulations for each value of κ. The results are provided in D.1 and D.2.

The figures D.1 and D.2 show a deterioration of the results with the increase of the threshold on the sources. This confirms that **the bias generated by the sources thresholding does hamper the estimation of the**

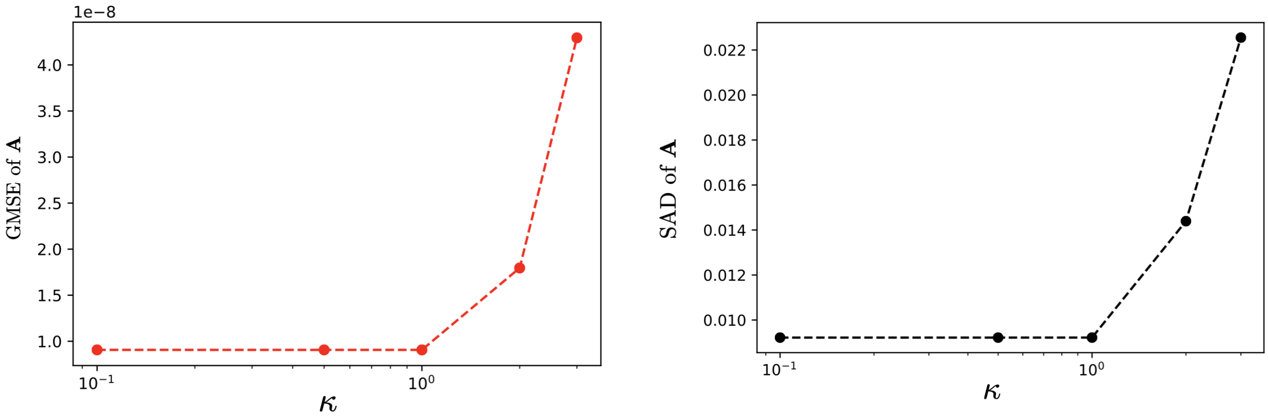

Figure D.1: Sources sparse in direct domain with $\rho = 24\%$

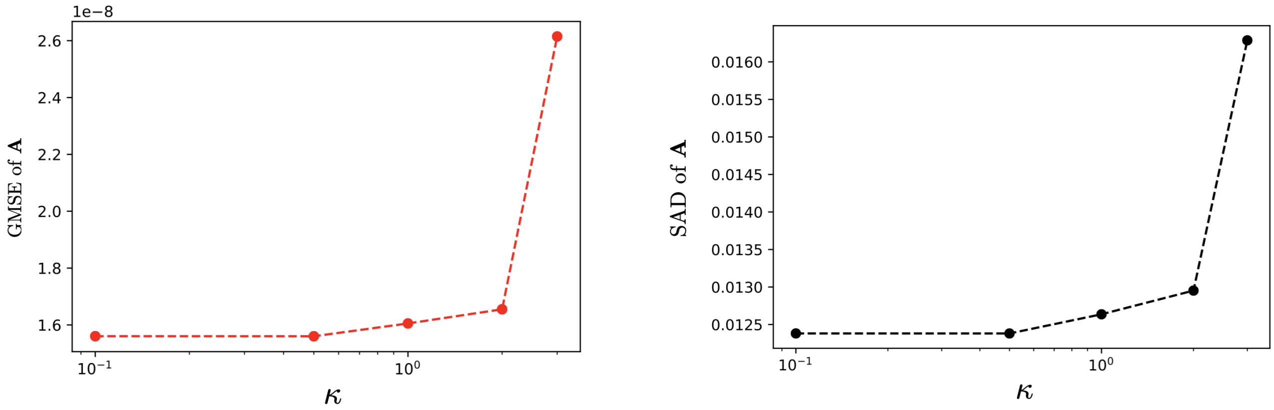

Figure D.2: Sources sparse in direct domain with $\rho = 10\%$

variabilities. Moreover, we observe that in the first dataset ($\rho = 24\%$), the variabilities are more affected by this error propagation. This can be explained by the fact that high thresholding is inadequate to less sparse sources and generates therefore more error propagation on the variabilities.

APPENDIX E

Fréchet mean on the hypersphere

In this thesis, we enforce the oblique constraint on the mixing matrix which implies that the columns belong to the hypersphere $\mathcal{S}^{(m-1)}$.
Let us consider we have a field of mixing matrices $\{\mathbf{A}[j], j = 1..J\}$ and we aim at defining the barycenter $\bar{\mathbf{A}}$ of these matrices. The problem can be written as:

$$\bar{\mathbf{A}} = \underset{\mathbf{A}}{\operatorname{argmin}} \sum_{j=1}^{J} w_j d(\mathbf{A}[j], \mathbf{A}),$$

where

- $\{w_j\}_{j=1...J}$ denote the barycenter weights, therefore they are positive and sum to one $\sum_{j=1}^{J} w_j = 1$;
- d is the distance used for defining the barycenter.

The choice of d as euclidean distance is excluded since we consider that the columns of the mixing matrix belong the hypersphere $\mathcal{S}^{(m-1)}$.
Instead, an appropriate choice would be to define the distance d as the geodesic distance on the hypersphere (that is a Riemann manifold, Absil et al. [2009]). The barycenter defined with such a distance is known as **Fréchet mean**.
The mixing matrix $\bar{\mathbf{A}}$ can therefore be obtained by computing the Fréchet mean using the algorithm proposed in Afsari et al. [2013].

Bibliography

P. Absil, R. Mahony, and R. Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.

B. Afsari, R. Tron, and R. Vidal. On the convergence of gradient descent for finding the riemannian center of mass. *SIAM Journal on Control and Optimization*, 51(3):2230–2260, 2013.

M. Arngren, M. Schmidt, and J. Larsen. Unmixing of hyperspectral images using bayesian non-negative matrix factorization with volume prior. *Journal of Signal Processing Systems*, 65(3):479–496, 2011.

G. Asner and D. Lobell. A biogeophysical approach for automated swir unmixing of soils and vegetation. *Remote sensing of environment*, 74(1):99–112, 2000.

H. Attouch, J. Bolte, P. Redont, and A. Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the kurdyka-łojasiewicz inequality. *Mathematics of operations research*, 35(2):438–457, 2010.

C. Badenes. X-ray studies of supernova remnants: A different view of supernova explosions. *PNAS*, 107(16), 2010.

H. Bauschke and P. Combettes. A dykstra-like algorithm for two monotone operators. *Pacific Journal of Optimization*, 4(3):383–391, 2008.

A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.

A. Bell and T. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159, 1995.

C. Bennett, D. Larson, J. Weiland, N. Jarosik, G. Hinshaw, N. Odegard, K. Smith, R. Hill, B. Gold, M. Halpern, et al. Nine-year wilkinson microwave anisotropy probe (wmap) observations: final maps and results. *The Astrophysical Journal Supplement Series*, 208(2):20, 2013.

M. Berman, H. Kiiveri, R. Lagerstrom, A. Ernst, R. Dunne, and J. Huntington. Ice: A statistical approach to identifying endmembers in hyperspectral images. *IEEE transactions on Geoscience and Remote Sensing*, 42(10):2085–2095, 2004.

M. Berry, M. Browne, A. Langville, P. Pauca, and R. Plemmons. Algorithms and applications for approximate nonnegative matrix factorization. *Computational statistics & data analysis*, 52(1):155–173, 2007.

J. Bioucas-Dias. A variable splitting augmented lagrangian approach to linear spectral unmixing. In *2009 First workshop on hyperspectral image and signal processing: Evolution in remote sensing*, pages 1–4. IEEE, 2009.

J. Bioucas-Dias and M. Figueiredo. Alternating direction algorithms for constrained sparse regression: Application to hyperspectral unmixing. In *2010 2nd Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing*, pages 1–4. IEEE, 2010.

J. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader, and J. Chanussot. Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches. *IEEE journal of selected topics in applied earth observations and remote sensing*, 5(2):354–379, 2012.

J. Boardman. Automating spectral unmixing of aviris data using convex geometry concepts, 1993.

J. Bobin, Y. Moudden, J-L. Starck, and M. Elad. Multichannel morphological component analysis. *Proceedings of Spars05*, pages 103–106, 2005.

J. Bobin, Y. Moudden, J-L. Starck, and M. Elad. Morphological diversity and source separation. *IEEE Signal Processing Letters*, 13 (7):409–412, 2006.

J. Bobin, J-L. Starck, J. Fadili, and Y. Moudden. Sparsity and morphological diversity in blind source separation. *IEEE Transactions on Image Processing*, 16(11):2662–2674, 2007.

J. Bobin, Y. Moudden, J-L Starck, J. Fadili, and N. Aghanim. Sz and cmb reconstruction using generalized morphological component analysis. *Statistical Methodology*, 5(4):307–317, 2008.

J. Bobin, J-L. Starck, F. Sureau, and S. Basak. Sparse component separation for accurate cosmic microwave background estimation. *Astronomy & Astrophysics*, 550:A73, 2013.

J. Bobin, J. Rapin, A. Larue, and J-L. Starck. Sparsity and adaptivity for the blind separation of partially correlated sources. *IEEE Transanctions on Signal Processing*, 63(5):1199–1213, 2015.

J. Bobin, I. El Hamzaoui, A. Picquenot, and F. Acero. Sparse bss from poisson measurements. *accepted in IEEE Tr. on Image Processing*, 2020.

J. Bolte, S. Sabach, and M. Teboulle. Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2):459–494, 2014.

S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine learning*, 3(1):1–122, 2011.

A. Bronstein, M. Bronstein, M. Zibulevsky, and Y. Zeevi. Sparse ICA for blind separation of transmitted and reflected images. *International Journal of Imaging Science and Technology*, 15/1, 2005. 84–91.

E. Candes. *Ridgelets: theory and applications*. PhD thesis, Stanford University Stanford, 1998.

E. Candes, M. Wakin, and S. Boyd. Enhancing sparsity by reweighted l1 minimization. *Journal of Fourier analysis and applications*, 14(5-6):877–905, 2008.

A. Castrodad, Z. Xing, J. Greer, E. Bosch, L. Carin, and G. Sapiro. Learning discriminative sparse representations for modeling, source separation, and mapping of hyperspectral imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 49(11):4263–4281, 2011.

Y. Cavalcanti, T. Oberlin, N. Dobigeon, S. Stute, M. Ribeiro, and C. Tauber. Unmixing dynamic pet images with variable specific binding kinetics. *Medical image analysis*, 49:117–127, 2018.

T-H. Chan, C-Y. Chi, Y-M. Huang, and W-K. Ma. A convex analysis-based minimum-volume enclosing simplex algorithm for hyperspectral unmixing. *IEEE Transactions on Signal Processing*, 57(11):4418–4432, 2009.

S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM review*, 43(1):129–159, 2001.

C. Chenot. Parcimonie, diversité morphologique et séparation robuste de sources, 2017.

A. Cherni, E. Piersanti, and C. Chaux. Nmf-based sparse unmixing of complex mixtures. 2019.

A. Cherni, S. Anthoine, and C. Chaux. β-nmf and sparsity promoting regularizations for complex mixture unmixing, application to 2d hsqc nmr. In *ICASSP 2020, Barcelona, Spain*, volume hal-02457468, 2020.

A. Cichocki and R. Zdunek. Regularized alternating least squares algorithms for non-negative matrix/tensor factorization. In *International Symposium on Neural Networks*, pages 793–802. Springer, 2007.

Planck Collaboration et al. Planck 2013 results: Xii. diffuse component separation. *Astronomy and Astrophysics*, 571:A12, 2014.

P. Combettes and J-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.

P. Combettes and V. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4): 1168–1200, 2005.

P. Comon. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994.

P. Comon and C. Jutten. *Handbook of Blind Source Separation: Independent component analysis and applications*. Academic press, 2010.

M. Craig. Minimum-volume transforms for remotely sensed data. *IEEE Transactions on Geoscience and Remote Sensing*, 32(3): 542–552, 1994.

G. Darmois. Analyse générale des liaisons stochastiques: etude particulière de l'analyse factorielle linéaire. *Revue de l'Institut international de statistique*, pages 2–8, 1953.

I. Daubechies. *Ten lectures on wavelets*. SIAM, 1992.

M. Davies. Identifiability issues in noisy ica. *IEEE Signal processing letters*, 11(5):470–473, 2004.

N. Dobigeon, S. Moussaoui, M. Coulon, J-Y. Tourneret, and A. Hero. Joint bayesian endmember extraction and linear unmixing for hyperspectral imagery. *IEEE Transactions on Signal Processing*, 57(11):4355–4368, 2009a.

N. Dobigeon, S. Moussaoui, J-Y. Tourneret, and C. Carteret. Bayesian separation of spectral sources under non-negativity and full additivity constraints. *Signal Processing*, 89(12):2657–2669, 2009b.

D. Donoho and X. Huo. Uncertainty principles and ideal atomic decomposition. *IEEE Transactions on Information Theory*, 47(7): 2845–2862, 2001.

D. Donoho and V. Stodden. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in neural information processing systems*, pages 1141–1148, 2004.

L. Drumetz, M-A. Veganzones, S. Henrot, R. Phlypo, J. Chanussot, and C. Jutten. Blind hyperspectral unmixing using an extended linear mixing model to address spectral variability. *IEEE Transactions on Image Processing*, 25(8):3890–3905, 2016.

J. Eggert and E. Korner. Sparse coding and nmf. In *2004 IEEE International Joint Conference on Neural Networks (IEEE Cat. No. 04CH37541)*, volume 4, pages 2529–2533. IEEE, 2004.

I. El Hamzaoui and J. Bobin. Sparse component separation from poisson measurements. *arXiv preprint arXiv:1812.04370*, 2018.

M. Elad, J-L. Starck, P. Querre, and D. Donoho. Simultaneous cartoon and texture image inpainting using morphological component analysis (mca). *Applied and computational harmonic analysis*, 19(3):340–358, 2005.

C. Févotte and N. Dobigeon. Nonlinear hyperspectral unmixing with robust nonnegative matrix factorization. *IEEE Transactions on Image Processing*, 24(12):4810–4819, 2015.

C. Févotte and J. Idier. Algorithms for nonnegative matrix factorization with the β-divergence. *Neural computation*, 23(9):2421–2456, 2011.

X. Fu, W-K. Ma, J. Bioucas-Dias, and T-H. Chan. Semiblind hyperspectral unmixing in the presence of spectral library mismatches. *IEEE Transactions on Geoscience and Remote Sensing*, 54(9):5171–5184, 2016.

N. Gadhok and W. Kinsner. An implementation of ß-divergence for blind source separation. In *2006 Canadian Conference on Electrical and Computer Engineering*, pages 1446–1449. IEEE, 2006.

B-C. Gao, M. Montes, C. Davis, and A. Goetz. Atmospheric correction algorithms for hyperspectral remote sensing data of land and ocean. *Remote Sensing of Environment*, 113:S17–S24, 2009.

F. García-Haro, S. Sommer, and T. Kemper. A new tool for variable multiple endmember spectral mixture analysis (vmesma). *International Journal of Remote Sensing*, 26(10):2135–2162, 2005.

N. Gillis and F. Glineur. Accelerated multiplicative updates and hierarchical als algorithms for nonnegative matrix factorization. *Neural computation*, 24(4):1085–1105, 2012.

A. Haar. Zur theorie der orthogonalen funktionensysteme. *Mathematische Annalen*, 69(3):331–371, 1910. doi: 10.1007/BF01456326. URL `https://doi.org/10.1007/BF01456326`.

D. Hong, N. Yokoya, J. Chanussot, and X. Zhu. An augmented linear mixing model to address spectral variability for hyperspectral unmixing. *IEEE Transactions on Image Processing*, 28(4):1923–1938, 2018.

P. Hoyer. Non-negative sparse coding. In *Neural Networks for Signal Processing, 2002. Proceedings of the 2002 12th IEEE Workshop on*, pages 557–565. IEEE, 2002.

M-D. Iordache, A. Plaza, and J. Bioucas-Dias. On the use of spectral libraries to perform sparse unmixing of hyperspectral data. In *2010 2nd Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing*, pages 1–4. IEEE, 2010.

M-D. Iordache, J. Bioucas-Dias, and A. Plaza. Sparse unmixing of hyperspectral data. *IEEE Transactions on Geoscience and Remote Sensing*, 49(6):2014–2039, 2011.

H. Iwasaki, Y. Ichinohe, and Y. Uchiyama. X-ray study of spatial structures in tycho's supernova remnant using unsupervised deep learning. *Monthly Notices of the Royal Astronomical Society*, 488(3):4106–4116, 2019.

C. Kervazo. *Optimization framework for large-scale sparse blind source separation*. PhD thesis, Paris Saclay, 2019.

C. Kervazo, J. Bobin, C. Chenot, and F. Sureau. Use of palm for ℓ_1 sparse matrix factorization: Difficulty and rationalization of an heuristic approach. *Digital Signal Processing*, 97, February 2020.

N. Keshava and J. Mustard. Spectral unmixing. *IEEE signal processing magazine*, 19(1):44–57, 2002.

H. Kim and H. Park. Nonnegative matrix factorization based on alternating nonnegativity constrained least squares and active set method. *SIAM journal on matrix analysis and applications*, 30(2):713–730, 2008.

K. Kurdyka. On gradients of functions definable in o-minimal structures, 1998.

J. Le Roux, F. Weninger, and J. Hershey. Sparse nmf–half-baked or well done? *Mitsubishi Electric Research Labs (MERL), Cambridge, MA, USA, Tech. Rep., no. TR2015-023*, 11:13–15, 2015.

D. Lee and S. Seung. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*, pages 556–562, 2001.

T-W. Lee, M. Girolami, A. Bell, and T. Sejnowski. A unifying information-theoretic framework for independent component analysis. *Computers & Mathematics with Applications*, 39(11):1–21, 2000.

C. Lelong, P. Pinet, and H. Poilvé. Hyperspectral imaging and stress mapping in agriculture: A case study on wheat in beauce (france). *Remote sensing of environment*, 66(2):179–191, 1998.

J. Leroux, F. Weninger, and J. Hershey. Sparse nmf – half-baked or well done? Technical report, Mitsubishi Electric Research Laboratories Technical Report, 2015.

J. Li and J. Bioucas-Dias. Minimum volume simplex analysis: A fast algorithm to unmix hyperspectral data. In *IGARSS 2008-2008 IEEE International Geoscience and Remote Sensing Symposium*, volume 3, pages III–250. IEEE, 2008.

Y. Li, S. Amari, A. Cichocki, D. Ho, and S. Xie. Underdetermined blind source separation based on sparse representation. *IEEE Transactions on Signal Processing*, 54(2):423–437, 2006.

S. Lojasiewicz. Une propriété topologique des sous-ensembles analytiques réels. *Les équations aux dérivées partielles*, 117:87–89, 1963.

L. Miao and H. Qi. Endmember extraction from highly mixed data using minimum volume constrained nonnegative matrix factorization. *IEEE Transactions on Geoscience and Remote Sensing*, 45(3):765–777, 2007.

M. Mihoko and S. Eguchi. Robust blind source separation by beta divergence. *Neural computation*, 14(8):1859–1886, 2002.

S. Moussaoui, C. Carteret, D. Brie, and A. Mohammad-Djafari. Bayesian analysis of spectral mixture data using markov chain monte carlo methods. *Chemometrics and Intelligent Laboratory Systems*, 81(2):137–148, 2006.

J. Nascimento and J. Dias. Does independent component analysis play a role in unmixing hyperspectral data? *IEEE Transactions on Geoscience and Remote Sensing*, 43(1):175–187, 2005a.

J. Nascimento and J. Dias. Vertex component analysis: A fast algorithm to unmix hyperspectral data. *IEEE transactions on Geoscience and Remote Sensing*, 43(4):898–910, 2005b.

Y. Nesterov. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.

R. Neville, K. Staenz, T. Szeredi, J. Lefebvre, P. Hauff, K. Neville, P. Hauff, and R. Neville. Automatic endmember extraction from hyperspectral data for mineral exploration, 1999.

P. Paatero and U. Tapper. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2):111–126, 1994.

N. Parikh, S. Boyd, et al. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014.

Y. Pati, R. Rezaiifar, and P. Krishnaprasad. Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proceedings of 27th Asilomar conference on signals, systems and computers*, pages 40–44. IEEE, 1993.

A. Picquenot, F. Acero, and J. Bobin. Revisiting the ejecta asymmetries in casa with a novel method for component separation in x-rays. *sros*, page 197, 2019a.

A. Picquenot, F. Acero, J. Bobin, P. Maggi, J. Ballet, and G. Pratt. A novel method for component separation of extended sources in x-ray astronomy. *arXiv preprint arXiv:1905.10175*, 2019b.

A. Picquenot, F. Acero, T. Holland-Ashford, L.A. Lopez, and J. Bobin. Threedimensional morphological asymmetries in the ejecta of cassiopeia a using a component separation method in x-rays. *A&A - submitted*, 2020.

H. Raguet, J. Fadili, and G. Peyré. A generalized forward-backward splitting. *SIAM Journal on Imaging Sciences*, 6(3):1199–1226, 2013.

I. Rahman, I. Drori, V. Stodden, D. Donoho, and P. Schröder. Multiscale representations for manifold-valued data. *Multiscale Modeling & Simulation*, 4(4), 2005. doi: 10.1137/050622729. URL `http://link.aip.org/link/?MMS/4/1201/1`. 1201-1232.

J. Rapin. Décompositions parcimonieuses pour l'analyse avancée de données en spectrométrie pour la santé, 2014.

J. Rapin, J. Bobin, A. Larue, and J-L. Starck. Nmf with sparse regularizations in transformed domains. *SIAM journal on Imaging Sciences*, 7(4):2020–2047, 2014.

D. Roberts, M. Gardner, R. Church, S. Ustin, G. Scheer, and R. Green. Mapping chaparral in the santa monica mountains using multiple endmember spectral mixture models. *Remote sensing of environment*, 65(3):267–279, 1998.

T. Rockafellar. *Convex analysis*. Princeton Landmarks in Mathematics and Physics. Princeton University Press, 1970.

D. Rogge, B. Rivard, J. Zhang, and J. Feng. Iterative spectral unmixing for optimizing per-pixel endmember sets. *IEEE Transactions on Geoscience and Remote Sensing*, 44(12):3725–3736, 2006.

B. Somers, G. Asner, L. Tits, and P. Coppin. Endmember variability in spectral mixture analysis: A review. *Remote Sensing of Environment*, 115(7):1603–1616, 2011.

B. Somers, M. Zortea, A. Plaza, and G. Asner. Automated extraction of image-based endmember bundles for improved spectral unmixing. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 5(2):396–408, 2012.

J-L. Starck, E. Candès, and D. Donoho. The curvelet transform for image denoising. *IEEE Transactions on image processing*, 11 (6):670–684, 2002.

J-L. Starck, D. Donoho, and M. Elad. Redundant multiscale transforms and their application for morphological component separation. Technical report, CM-P00052061, 2004.

J-L. Starck, J. Fadili, and F. Murtagh. The undecimated wavelet decomposition and its reconstruction. *IEEE transactions on image processing*, 16(2):297–309, 2007.

J.-L. Starck, Y. Moudden, and J. Bobin. Polarized wavelets and curvelets on the sphere. *Astronomy and Astrophysics*, 497:931–943, April 2009. doi: 10.1051/0004-6361/200811343.

J-L. Starck, F. Murtagh, and J. Fadili. *Sparse image and signal processing: Wavelets and related geometric multiscale analysis*. Cambridge university press, 2015.

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

Y. Sun and J.Xin. A blind source separation method for nearly degenerate mixtures and its applications to nmr spectroscopy. *arXiv:1110.1676*, 2011.

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(2):525–538, 2015.

P. Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of optimization theory and applications*, 109(3):475–494, 2001.

M. Veganzones, L. Drumetz, G. Tochon, M. Dalla Mura, A. Plaza, J. Bioucas-Dias, and J. Chanussot. A new extended linear mixing model to address spectral variability. In *2014 6th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, pages 1–4. IEEE, 2014.

E. Vincent, R. Gribonval, and C. Févotte. Performance measurement in blind audio source separation. *IEEE transactions on audio, speech, and language processing*, 14(4):1462–1469, 2006.

M. Winter. N-findr: An algorithm for fast autonomous spectral end-member determination in hyperspectral data. In *Imaging Spectrometry V*, volume 3753, pages 266–275. International Society for Optics and Photonics, 1999.

J. Xu, J. Bobin, A. de Vismes Ott, and C. Bobin. Sparse spectral unmixing for activity estimation in γ-ray spectrometry applied to environmental measurements. *Applied Radiation and Isotopes, in press*, 2019.

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(3):1758–1789, 2013.

Y. Xu and W. Yin. A globally convergent algorithm for nonconvex optimization based on block coordinate update. *arXiv preprint arXiv:1410.1386*, 2014.

Y. Xu and W. Yin. A globally convergent algorithm for nonconvex optimization based on block coordinate update. *Journal of Scientific Computing*, 72(2):700–734, 2017.

A. Zare and K. Ho. Endmember variability in hyperspectral analysis: Addressing spectral variability during spectral unmixing. *IEEE Signal Processing Magazine*, 31(1):95–104, 2013.

J. Zhang, B. Rivard, A. Sánchez-Azofeifa, and K. Castro-Esau. Intra- and inter-class spectral variability of tropical tree species at la selva, costa rica: Implications for species identification using hydice imagery. *Remote Sensing of Environment*, 105(2):129 – 141, 2006. ISSN 0034-4257. doi: https://doi.org/10.1016/j.rse.2006.06.010. URL `http://www.sciencedirect.com/science/article/pii/S0034425706002409`.

M. Zibulevsky. Blind source separation with relative newton method. In *Proc. ICA*, volume 2003, pages 897–902, 2003.

M. Zibulevsky and B. Pearlmutter. Blind source separation by sparse decomposition in a signal dictionary. *Neural computation*, 13 (4):863–882, 2001.

université
PARIS-SACLAY

ÉCOLE DOCTORALE
Sciences et technologies
de l'information et de
la communication (STIC)

Titre: Décomposition non-supervisée de données multivaluées avec applications en astrophysique

Mots clés: Séparation de Sources Aveugle, Représentations parcimonieuses, Factorisation de matrices, Bruit de Poisson, Variabilités spectrales, Diversité morphologique.

Résumé: L'imagerie multi-spectrale multi-temporelle requiert la mise en place d'outils d'analyses adaptés aux données multi-valuées. Les missions Chandra et Athena en astrophysique sont des exemples parlants des défis passés et à venir dans le traitement des données multi-spectrales. Cette thèse a pour objectif de proposer des modèles d'analyse de données astrophysiques en rayons X et des algorithmes permettant d'extraire les informations utiles pour l'astrophysicien. Il s'agit d'introduire une extension des techniques de séparation de composantes dans le but, d'une part, d'avoir des modèles mathématiques capables de décrire des données multi-valuées contaminées par du bruit de Poisson, et d'autre part, d'estimer les variabilités spectrales très répandues dans les jeux de données astrophysiques en hautes énergies. Les outils numériques développés au cours de cette thèse sont appliqués aux données du télescope Chandra.

Title: Unsupervised separation of sparse multivalued components with applications in astrophysics.

Keywords: Blind Source Separation, Sparse representations, Matrix factorisation, Poisson noise, Spectral variabilities, Morphological Diversity.

Abstract: The rapid increase of multispectral-multitemporal imagers in various application fields requires new data analysis tools particularly suitable for multivalued data. In high-energy astronomy, missions such as Chandra or Fermi are telling examples of signal processing challenges past or to come. This thesis is aimed at proposing new models to analyze X-ray astrophysical data and introducing efficient algorithms to retrieve meaningful information from these data. More specifically, the goal of this thesis is to extend component separation techniques in order to propose models that faithfully describe measurements contaminated with shot noise and that fully account for spectral variabilities ubiquitous in high-energy astrophysical images. The numerical tools developed in this thesis will be applied to X-ray Chandra telescope data.

Université Paris-Saclay
CEA Paris-Saclay, Département d'Astrophysique
Orme des Merisiers, F-91191 Gif-sur-Yvette, France