

Joint deconvolution and semi & unsupervised component separation, with application to radio-interferometry

Rémi Carloni Gertosio

▶ To cite this version:

Rémi Carloni Gertosio. Joint deconvolution and semi & unsupervised component separation, with application to radio-interferometry. Signal and Image processing. Université Paris-Saclay, 2022. English. NNT: 2022UPASG064 . tel-03881392

HAL Id: tel-03881392 https://theses.hal.science/tel-03881392

Submitted on 1 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Joint deconvolution and semi & unsupervised component separation, with application to radio-interferometry

Séparation de composantes semi & non supervisée et déconvolution conjointe, application à la radio-interférométrie

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 Graduate School : informatique et sciences du numérique. Référent : faculté des sciences d'Orsay

Thèse préparée au Département d'Electronique des Détecteurs et d'Informatique pour la Physique (Université Paris-Saclay, CEA) sous la direction de **Jérôme BOBIN**, directeur de recherche

Thèse soutenue à Paris-Saclay, le 8 septembre 2022, par

Rémi CARLONI GERTOSIO

Composition du Jury

Cédric RICHARD	Drécident & Deprestour
Professeur, Université Côte-d'Azur	President & Rapporteur
Jocelyn CHANUSSOT	
Professeur, Université Grenoble-Alpes	Rapporteur & Examinateur
Aurélia FRAYSSE	Fuencia etaiee
Maître de conférences, HDR, Université Paris-Saclay	Examinatrice
Julien GIRARD	Fuencia eterra
Maître de conférences, Université Paris Cité	Examinateur
Jérôme BOBIN	Dive starry de thèse
Directeur de recherche, CEA, Université Paris-Saclay	Directeur de these

NNT : 2022UPASG064

HESE DE DOCTORAT

ÉCOLE DOCTORALE



Sciences et technologies de l'information et de la communication (STIC)

Titre : Séparation de composantes semi & non supervisée et déconvolution conjointe, application à la radiointerférométrie.

Mots clés : séparation aveugle de sources, problème inverse, parcimonie, apprentissage automatique, radiointerférométrie.

Résumé : Avec le développement des radiotélescopes de taille continentale, l'analyse des données radio-interférométriques multifréquences devient un enjeu majeur tant en traitement du signal qu'en astrophysique. A cet égard, les méthodes de séparation de composantes constituent un outil adéquat car elles permettent de séparer les données multifréquences en composantes physiques élémentaires.

Cependant, les algorithmes standard de séparation de composantes ne sont pas adaptés aux données radio-interférométriques à plus d'un titre : *(i)* cellesci sont composées de mesures incomplètes dans le domaine de Fourier, dit des visibilités, et sont en outre potentiellement détériorées par des effets instrumentaux ou non coplanaires, et *(ii)* les signaux recherchés peuvent être sévèrement noyés dans le bruit ou dans d'autres émissions.

En conséquence, la prise en compte de la réponse instrumentale des télescopes requiert le développement d'algorithmes dédiés pour résoudre un problème de séparation et de déconvolution conjoint. De plus, la récupération de signaux faibles nécessite la conception d'approches précises qui exploitent des modèles physiques connus a priori ; cela motive le développement de méthodes de séparation semi supervisées qui intègrent des techniques par apprentissage automatique.

Title: Joint deconvolution and semi & unsupervised component separation, with application to radio-interferometry.

Keywords: blind source separation, inverse problem, sparsity, machine learning, radio-interferometry.

Abstract: With the development of continent-wide radio telescopes, the analysis of multi-frequency radio-interferometric data has become a significant challenge in signal processing and astrophysics. In this respect, component separation methods are an adequate tool as they allow the separation of multi-frequency data into elementary physical components.

However, standard component separation algorithms are not adapted to radio-interferometric data for several reasons: *(i)* they are constituted of incomplete measurements in the Fourier domain, also known as the visibility domain, and are potentially further deteriorated by instrumental or non-coplanar effects, and *(ii)* the sought-after signals can be severely drowned in noise or other emissions.

Consequently, accounting for the telescope instrumental response requires developing dedicated algorithms to solve a joint deconvolution and separation problem. Furthermore, the recovery of weak signals calls for the design of accurate approaches that make use of physical models known a priori; this motivates the development of semi-supervised separation methods that integrate machine learning techniques.

Remerciements

Je remercie tout d'abord Jocelyn Chanussot et Cédric Richard d'avoir accepté d'expertiser ce manuscrit et d'en être rapporteurs. Toute ma gratitude va également aux membres du jury Aurélia Fraysse et Julien Girard pour leur participation au jury et leur présence à la soutenance.

Je tiens ensuite à remercier très sincèrement mon directeur de thèse Jérôme Bobin. C'est un encadrant passionné et disponible, qui a installé entre nous une relation de confiance. En travaillant avec lui, j'ai bénéficié de son expérience et de son expertise; ce faisant, il m'a incité à approfondir mes connaissances et à améliorer la qualité de ma démarche scientifique.

Je remercie mes collègues du LSN pour leur accueil chaleureux en cours de thèse. Je n'oublie pas les membres de CosmoStat qui ont créé un environnement de travail convivial durant la première moitié de thèse. J'ai une pensée particulière pour mes camarades doctorants rencontrés durant ces trois années; ce fut un réel plaisir de partager avec eux de grands moments de complicité au bureau ou à l'extérieur!

Je tiens finalement à remercier mes parents de leurs précieux conseils et de leur soutien constant.

A mon père

Le silence éternel de ces espaces infinis m'effraie. — Blaise Pascal, Pensées.

Contents

Re	emer	cieme	nts	iii
Sy	vnthè	ese		xi
N	otati	ons an	d acronyms	xix
In	trod	uction		1
1	A q	uick r	eview of blind source separation	7
	1.1	1.1 Mixture model		
	1.2	Indep	endent component analysis	9
	1.3	Cost-f	unction-based methods	9
		1.3.1	General framework	9
		1.3.2	Non-negative matrix factorization	10
		1.3.3	Sparse blind source separation	13
			1.3.3.1 Assumptions on the cost function \ldots \ldots \ldots \ldots	13
			1.3.3.2 Proximal operators of the regularization terms	15
			1.3.3.3 Algorithms for sparse blind source separation	16
			1.3.3.4 A particular pALS algorithm: Generalized morpho-	
			logical component analysis	19
	1.4	Learn	ing-based techniques	22
	1.5	Perfor	mance metrics for BSS	25
		1.5.1	Mixing matrix	25
		1.5.2	Sources	26
2	Intr	roducii	ng learned constraints in blind source separation	29
	2.1	Introduction		
	2.2	Towar	ds a semi-blind source separation framework	30
		2.2.1	Constraining the spectra	30
		2.2.2	Machine learning and blind source separation	30
	2.3	nterpolatory autoencoder	32	

		2.3.1	Principle		32
			2.3.1.1 Learning to travel on a manifold		32
			2.3.1.2 Preliminary: affine hulls		33
			2.3.1.3 Architecture of the interpolatory autoencoder		34
		2.3.2	Application		34
2.4 The sGMCA algorithm			GMCA algorithm		36
		2.4.1	Principle		37
		2.4.2	Minimization scheme		38
		2.4.3	Source update		40
		2.4.4	Mixing matrix update		40
			2.4.4.1 Adaptation of the barycentric span projection for s	pec-	
			tra		40
			2.4.4.2 Spectrum identification and model-to-spectrum m	ap-	
			ping		40
			2.4.4.3 Constraint application		41
	2.5	Nume	ical experiments		42
		2.5.1	Learning of the generative models		44
		2.5.2	Results		46
			$2.5.2.1 \text{Overall results} \dots \dots \dots \dots \dots \dots \dots \dots \dots $		46
			2.5.2.2 Impact of the experimental parameters		47
			2.5.2.3 Impact of the generative modeling		51
	2.6	Conch	lsion		50
		0 0 11 0 1	131011	• •	52
0	ЛЛ	•		•••	32
3	Mea	asuring	; and imaging the radio sky with a view to perfor	nin	02 g
3	Mea blin	asuring d sour	and imaging the radio sky with a view to perfor ce separation	nin	52 g 55
3	Mea blin 3.1	asuring d sour Measu	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions	 min	52 g 55 55
3	Mea blin 3.1	asuring ad sour Measu 3.1.1	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions	••• min •••	52 g 55 55 55 56
3	Mea blin 3.1	asuring Id sour Measu 3.1.1 3.1.2 Single	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions	 min 	g 55 55 55 55 56 50
3	Mea blin 3.1 3.2	asuring d sour 3.1.1 3.1.2 Single 3.2.1	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions	min 	g 55 55 55 56 59 59
3	Mea blin 3.1 3.2	asuring d sour 3.1.1 3.1.2 Single 3.2.1 3.2.2	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions	min 	g 55 55 55 56 59 59 59
3	Mea blin 3.1 3.2	asuring d sour 3.1.1 3.1.2 Single 3.2.1 3.2.2 Interfe	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions	min 	g 55 55 56 59 59 59 59 60
3	Mea blin 3.1 3.2 3.3	asuring d sour 3.1.1 3.1.2 Single 3.2.1 3.2.2 Interfe 3.3.1	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions	min 	5 2 5 5 55 56 59 59 59 59 60 60
3	Mea blin 3.1 3.2 3.3	asuring d sour 3.1.1 3.1.2 Single 3.2.1 3.2.2 Interfe 3.3.1 3.3.2	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions Characterization of radio sources Radio antennas dish imaging Small fields of view Large fields of view rometric imaging A basic twin interferometer The visibility function	min 	5 2 5 5 55 55 56 59 59 59 60 60 60 62
3	Mea blin 3.1 3.2 3.3	asuring d sour 3.1.1 3.1.2 Single 3.2.1 3.2.2 Interfe 3.3.1 3.3.2 3.3.3	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions	min 	5 2 5 5 55 56 59 59 59 60 60 60 62 64
3	Mea blin 3.1 3.2 3.3	asuring d sour 3.1.1 3.1.2 Single 3.2.1 3.2.2 Interfe 3.3.1 3.3.2 3.3.3 3.3.4	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions Characterization of radio sources Radio antennas dish imaging Small fields of view Large fields of view rometric imaging A basic twin interferometer The visibility function Towards the forward observation model	min 	5 2 5 5 5 5 5 5 5 6 5 9 5 9 5 9 6 0 6 0 6 2 6 4 6 7
3	Mea blin 3.1 3.2 3.3	asuring Measu 3.1.1 3.1.2 Single 3.2.1 3.2.2 Interfe 3.3.1 3.3.2 3.3.3 3.3.4 3.3.5	g and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions Characterization of radio sources Radio antennas dish imaging Small fields of view Large fields of view rometric imaging A basic twin interferometer The visibility function Coordinate systems for imaging Towards the forward observation model	min 	5 2 5 5 5 5 5 5 5 6 5 9 5 9 5 9 6 0 6 0 6 2 6 4 6 7 6 7 6 7
3	Mea blin 3.1 3.2 3.3	asuring d sour 3.1.1 3.1.2 Single 3.2.1 3.2.2 Interfe 3.3.1 3.3.2 3.3.3 3.3.4 3.3.5	and imaging the radio sky with a view to perfor ce separation ring astronomical radio emissions Characterization of radio sources Radio antennas dish imaging Small fields of view Large fields of view rometric imaging A basic twin interferometer The visibility function Coordinate systems for imaging Towards the forward observation model 3.3.5.1	min 	5 2 5 5 5 5 5 5 5 6 5 9 5 9 5 9 6 0 6 0 6 2 6 4 6 7 6 7 6 7 6 7 6 7
3	Mea blin 3.1 3.2 3.3	asuring Measu 3.1.1 3.1.2 Single 3.2.1 3.2.2 Interfe 3.3.1 3.3.2 3.3.3 3.3.4 3.3.5	g and imaging the radio sky with a view to perforce separationring astronomical radio emissionsCharacterization of radio sourcesRadio antennasdish imagingSmall fields of viewLarge fields of viewA basic twin interferometerThe visibility functionCoordinate systems for imaging	min 	5 2 5 5 5 5 5 6 5 9 5 9 5 9 6 0 6 0 6 2 6 4 6 7 6 7 6 7 6 9
3	Mea blin 3.1 3.2 3.3	asuring d sour 3.1.1 3.1.2 Single 3.2.1 3.2.2 Interfe 3.3.1 3.3.2 3.3.3 3.3.4 3.3.5 Tackli	and imaging the radio sky with a view to perfor ce separationring astronomical radio emissionsCharacterization of radio sourcesRadio antennasMali imagingSmall fields of viewLarge fields of viewA basic twin interferometerThe visibility functionCoordinate systems for imaging3.3.5.1The coplanar case	min 	5 2 5 5 5 5 5 5 5 6 5 9 5 9 5 9 6 0 6 0 6 2 6 4 6 7 6 7 6 7 6 7 6 9 7 1

		3.4.2	Towards	blind source separation	72
4	Joir serv	nt deco vations	onvolutio	on and blind source separation for wide-field ob	- 75
	4.1	Introd	uction .		75
	4.2	The sr	oherical E	DecGMCA algorithm	77
		4.2.1	Toward	a DBSS method for spherical data	77
		4.2.2	Update	of A	79
		4.2.3	Update	of \mathbf{S}	79
		1.2.0	4 2 3 1	Regularization strategies	80
			4232	Numerical comparisons of the regularization strategies	82
	4.3	Impler	nentation	of the SDecGMCA algorithm	84
	1.0	4.3.1	Choice of	of the thresholding parameters	86
		432	A two-st	age minimization approach	88
		433	Converg	ence proprieties	89
	4.4	Nume	rical expe	riments	90
	1.1	4.4.1	Toy more	lel	90
			4.4.1.1	Impact of the Tikhonov regularization	91
			4.4.1.2	Comparisons with other blind source separation meth-	01
				ods	93
			4.4.1.3	Varving observation parameters	94
		4.4.2	Applicat	ion to realistic astrophysical data	95
			4.4.2.1	Example of result	95
			4.4.2.2	Comparison with other blind source separation method	s100
			4.4.2.3	Impact of the noise	100
	4.5	Conclu	usion	I	103
-	т	4 1	1.4.		
9	JOII		nivolutio	te	r 107
	E 1	Introd	letric da	ta	107
	5.1 5.9	Drinoi	nlo		107
	0.2	5 9 1	Contort		100
		5.2.1	Operato	rs on tonsors	100
		5.2.2	Forward	observation model	109
		0.2.0	5021	Assumptions on the poise	109
		594	0.2.0.1	Assumptions on the hoise	111
	52	J.2.4 Towar	de non co	nlanar DRSS	110
	J.J	10war 5 3 1	Undata	ר מפשע האוקים	112 119
		0.0.1	5 2 1 1	Subproblem I: deconvolve and unmiv the visibilities	112
			5319	Subproblem I: acconvolve and unnix the visibilities .	11/
			5.3.1.2 5.3.1.2	Begularization of the sources	115
			0.0.1.0		119

		5.3.2	Update of \mathbf{A}	15
	5.4	The w	·GMCA algorithm	16
		5.4.1	Regularization strategies	16
		5.4.2	Improving the estimation of the mixing matrix	16
		5.4.3	Thresholding parameters	17
		5.4.4	A two-step minimization	117
		5.4.5	Complexity and convergence proprieties	22
	5.5	Nume	rical experiments	122
		5.5.1	Characterization of wGMCA on synthetic data 1	123
			5.5.1.1 Data generation $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots 1$	123
			5.5.1.2 Result example $\ldots \ldots 1$	126
			5.5.1.3 Impact of the non-coplanarity $\ldots \ldots \ldots \ldots \ldots 1$	126
			5.5.1.4 Impact of the regularization hyperparameters 1	127
			5.5.1.5 Impact of the observation parameters $\ldots \ldots \ldots \ldots 1$	129
		5.5.2	Comparisons with other methods on realistic data 1	129
			5.5.2.1 Description of the data $\ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots 1$	129
			5.5.2.2 Sequential deconvolution and separation methods 1	133
			$5.5.2.3$ Results \ldots 1	133
	5.6	Conclu	usion	136
Co	onclu	ision a	nd perspectives 1	39
\mathbf{A}	Eleı	ments	of signal processing for random processes 1	43
в	\mathbf{Sph}	erical	signal analysis 1	47
	B.1	Spheri	ical harmonics	47
	B.2	Convo	lution products	48
	B.3	HEAL	Pix discretization	48
\mathbf{C}	Ada	ptatio	on of SDecGMCA for coplanar data 1	51
D	Open source codes 1			

Synthèse

Contexte et problèmes

Les algorithmes de séparation aveugle de sources (SAS) sont des méthodes nonsupervisées qui permettent d'analyser des données multivaluées en les décomposant automatiquement en composantes élémentaires. Ils sont utilisés dans de nombreux domaines, par exemple en biologie pour l'analyse des données électroencéphalographiques, en télédétection dans le cadre d'observations multispectrales terrestres ou bien en sciences des données pour la fouille de textes [Gillis, 2020].

Les travaux de cette thèse visent notamment l'analyse d'images multispectrales en astrophysique. Dans ce domaine, les algorithmes standard de SAS seront confrontés par la complexité des données multispectrales des projets clefs à venir, comme le radio-interféromètre continental Square Kilometer Array¹ ou l'observatoire à rayons X Athena². A cet égard, plusieurs problèmes dans le domaine du traitement du signal doivent être résolus :

- La résolution spatiale accrue atteinte par certains télescopes devrait permettre de distinguer des émissions très faibles, jusqu'à quelques ordres de grandeur en dessous des émissions les plus puissantes. En outre, le bruit instrumental peut dominer le signal d'intérêt. Par conséquent, le premier point concerne le développement de méthodes de séparation très précises.
- Les observations multispectrales sont déformées par des réponses instrumentales dépendant de la longueur d'onde. Dans le cas de mesures large bande, les méthodes de SAS doivent faire face à des données particulièrement hétérogènes, avec des résolutions spatiales variables.
- Les données radio-interférométriques sont obtenues par corrélation de paires d'antennes. En première approximation, chaque paire d'antennes produit une mesure ponctuelle dans le domaine harmonique. Le nombre de paires étant fini, les données interférométriques sont incomplètes. Par conséquent, les méthodes

¹https://www.skatelescope.org/ ²https://sci.esa.int/web/athena

SAS appliquées aux données interférométriques doivent inclure des régularisations spécifiques.

- Dans le cas de la radio-interférométrie à longue ligne de base (c'est-à-dire grande distance entre les antennes les plus éloignées), l'hypothèse classique d'onde incidente plane devient obsolète. Le modèle d'observation prenant en compte cet effet non coplanaire est beaucoup plus complexe. Un phénomène similaire est à considérer dans le cas d'observations à très grand champ. Les méthodes de séparation précises doivent ainsi tenir compte de ces effets non coplanaires.

Plan

Ce manuscrit est constitué de cinq chapitres. Le chapitre 1 présente le problème standard de SAS linéaire et passe brièvement en revue les algorithmes qui permettent sa résolution. Dans le chapitre 2, nous proposons un algorithme de séparation de sources semi-supervisé qui exploite la potentielle connaissance de certaines des composantes recherchées. Le chapitre 3 présente les bases de la radioastronomie dans l'objectif d'en tirer des modèles d'observation multispectraux afin d'aborder ensuite la SAS à partir de données radio. Dans le chapitre 4, nous étudions un algorithme de SAS et de déconvolution conjointe pour l'analyse des données multispectrales très grand champ acquises sur la sphère. Finalement, dans le chapitre 5, nous nous concentrons sur le développement d'un algorithme de SAS et de déconvolution conjointe à partir de données radio-interférométriques non coplanaires.

Chapitre 1 : une brève revue des algorithmes de séparation aveugle de sources

Le problème de SAS linéaire est communément écrit sous forme matricielle :

$$\mathbf{Y} = \mathbf{AS} + \mathbf{N},\tag{1}$$

où $\mathbf{Y} \in \mathbb{R}^{J \times P}$ contient les observations multispectrales, $\mathbf{A} \in \mathbb{R}^{J \times I}$ est la matrice de mélange, $\mathbf{S} \in \mathbb{R}^{I \times P}$ est la matrice des sources, $\mathbf{N} \in \mathbb{R}^{J \times P}$ le terme de bruit, avec I le nombre de sources, J le nombre de canaux d'observation/bandes spectrales et P le nombre de pixels par observation/source. Sous cette formulation, la matrice \mathbf{A} contient dans ses colonnes les I spectres associés aux composantes recherchées. Le but de la SAS est d'estimer conjointement \mathbf{A} et \mathbf{S} à partir de \mathbf{Y} ; pour ce faire, on distingue trois familles de méthodes de résolution.

La première famille repose sur l'analyse en composantes indépendantes [Comon and Jutten, 2010]. Néanmoins, en astrophysique, l'hypothèse sous-jacente d'indépendance statistique des sources est peu vérifiée en pratique, car les sources présentent couramment des corrélations partielles.

La deuxième famille d'algorithmes de SAS s'appuie sur l'optimisation d'une fonction coût. D'une manière générale, on écrira celle-ci de la façon suivante :

$$\underset{\mathbf{A},\mathbf{S}}{\operatorname{argmin}} \mathcal{D}(\mathbf{Y}, \mathbf{AS}) + h_{\mathcal{S}}(\mathbf{S}) + h_{\mathcal{A}}(\mathbf{A}), \qquad (2)$$

avec \mathcal{D} le terme d'attache aux données, qui dépend à la fois du modèle de mélange et des statistiques du bruit, et $h_{\mathcal{S}}$ et $h_{\mathcal{A}}$ des fonctions de régularisation qui tiennent compte respectivement des a priori sur les sources et sur la matrice de mélange.

Le problème (2) peut être résolu avec les méthodes de factorisation en matrices positives [Gillis, 2020], qui supposent la positivité des coefficients de \mathbf{A} et \mathbf{S} ($h_{\mathcal{S}}$ et $h_{\mathcal{A}}$ sont alors des fonctions caractéristiques sur des orthants positifs). Ces méthodes sont populaires pour leurs garanties de convergence, voire d'identifiabilité. Cependant, pour le traitement des données astrophysiques, elles ne sont pas adaptées car (i) leurs schémas de minimisation offrent peu de flexibilité pour implémenter des régularisations et des modèles de mélange plus élaborés, et (ii) elles manquent parfois de robustesse par rapport à l'initialisation.

On préférera donc les méthodes de SAS parcimonieuse pour résoudre le problème (2). Celles-ci reposent sur la parcimonie des sources dans un domaine bien choisi (d'où $h_{\mathcal{S}}(\mathbf{S}) = \|\mathbf{\Lambda} \odot (\mathbf{S}\mathbf{W}^{\top})\|_1$ avec $\mathbf{W} \in \mathbb{R}^{P' \times P}$ une transformée et $\mathbf{\Lambda} \in \mathbb{R}^{I \times P'}$ des hyperparamètres de régularisation). En particulier, l'algorithme GMCA [Bobin et al., 2015] se distingue par son cadre flexible et frugal basé sur un moindres carrés alterné projeté, ses heuristiques et sa robustesse établie en SAS.

La troisième famille de méthodes de SAS, la moins explorée jusqu'à présent, repose sur l'apprentissage automatique. Cette famille est assez disparate, que ce soit par le type de problème considéré (modèle de mélange, hypothèses et a priori sur les sources ou les processus de mélange) ou par la méthodologie (séparation bout-enbout, apprentissage de régularisation, etc.). Les avantages de ces méthodes sont de potentiellement mieux exploiter les connaissances a priori des solutions et d'améliorer l'interprétabilité des résultats.

Chapitre 2 : introduction de contraintes apprises en séparation aveugle de sources

L'un des défis de la SAS appliquée aux données réelles est de fournir des résultats de séparation physiquement interprétables. En particulier, il est primordial que les fuites entre composantes restent limitées, que ce soit dans le domaine des sources ou dans

le domaine spectral.

Les jeux de données multispectraux analysés concernent souvent des processus qui sont au moins partiellement connus, et à cet égard, le recours à des approches totalement non supervisées ne semble pas le plus adéquat. Par conséquent, dans ce chapitre, nous proposons un algorithme de séparation semi-supervisé qui exploite la potentielle connaissance a priori de certains des spectres recherchés.

La méthode proposée s'inscrit dans le cadre général de l'apprentissage de régularisation pour approche variationnelle en problèmes inverses [Adler and Öktem, 2017]. Cette approche hybride permet d'associer apprentissage automatique et information présente dans une fonction coût (modèle d'observation, modèle de bruit et fonctions de régularisation déjà en place). Plus particulièrement, la fonction coût considérée est³ :

$$\underset{\mathbf{A} \ge \mathbf{0}, \mathbf{S} \ge \mathbf{0}}{\operatorname{argmin}} \frac{1}{2} \left\| \mathbf{Y} - \mathbf{A} \mathbf{S} \right\|_{2}^{2} + \left\| \mathbf{\Lambda} \odot \left(\mathbf{S} \mathbf{W}^{\top} \right) \right\|_{1} + \sum_{i \in \mathbb{I}} \iota_{\mathcal{M}^{(m_{i})}} \left(\mathbf{A}_{:i} \right).$$
(3)

Nous faisons le choix de contraindre certains spectres, pour rappel présents dans \mathbf{A} , à appartenir à des variétés $\{\mathcal{M}^{(m_i)}\}_{i\in\mathbb{I}}$ préalablement apprises. Pour ce faire, nous employons l'autoencodeur interpolateur, qui est un outil que nous avons développé pour régulariser des problèmes de démélange tels que la SAS.

L'algorithme proposé, sGMCA (GMCA semi-aveugle), s'appuie sur un moindres carrés alterné projeté et réemploie les heuristiques de GMCA. Des expériences numériques approfondies sur données astrophysiques réalistes montrent que la régularisation introduite permet de rejeter efficacement les fuites entre composantes, améliorant ainsi de manière significative l'estimation des sources et des spectres comparé aux méthodes totalement aveugles, y compris dans des contextes difficiles (observations bruitées, spectres corrélés, émissions déséquilibrées).

Chapitre 3 : mesure et imagerie du ciel radio en vue de traiter la séparation aveugle de sources

La radioastronomie est une branche de l'astronomie qui concerne l'observation du ciel dans le domaine des fréquences radio. La mesure précise des sources radio, qui sont généralement des signaux de très faible puissance, nécessite l'utilisation d'instruments massifs, les radiotélescopes, qui mobilisent des ressources importantes tant en termes matériel que de calcul.

Dans ce chapitre, les principes fondamentaux des mesures en radioastronomie sont présentés. Les modèles d'acquisition des émissions radio sont ensuite développés, d'une part avec les télescopes à antenne unique et d'autre part avec les interféromètres.

 $^{^{3}\}mathrm{La}$ fonction coût ici présentée est volontairement simplifiée pour ne pas devoir aborder des détails techniques d'implémentation.

Enfin, les modèles d'acquisition sont étendus au cas multispectral en vue d'aborder ultérieurement le problème de SAS. On différencie trois cas :

1. Les cas du télescope à antenne unique et de l'interféromètre coplanaire peuvent tous deux être réunis dans le même cadre pour les petits champs de vue. Le modèle, exprimé directement dans l'espace de Fourier ou des visibilités, s'écrit pour chaque bande fréquentielle ν :

$$\tilde{y}^{(\nu)}(u,v) = \tilde{H}^{(\nu)}(u,v) \ \tilde{x}^{(\nu)}(u,v), \tag{4}$$

avec $y^{(\nu)}$ les observations multispectrales, $H^{(\nu)}$ l'opérateur d'observation (dans le cas d'un télescope à antenne unique, sa réponse instrumentale, et dans le cas d'un interféromètre, le masque issu de la distribution des antennes) et $x^{(\nu)}$ les mélanges recherchés, et où le tilde symbolise la transformée de Fourier bidimensionnelle.

2. Dans le cas d'observations très grand champ, le modèle d'acquisition d'un télescope à antenne unique s'écrit de manière similaire dans le domaine des harmoniques sphériques :

$$\tilde{y}^{(\nu)}(\mathfrak{l},\mathfrak{m}) = \tilde{H}^{(\nu)}(\mathfrak{l}) \ \tilde{x}^{(\nu)}(\mathfrak{l},\mathfrak{m}),\tag{5}$$

avec les mêmes correspondances que ci-dessus hormis le tilde qui désigne ici la transformée en harmoniques sphériques.

3. Finalement, le modèle d'acquisition d'un interféromètre non coplanaire s'écrit :

$$\tilde{y}^{(\nu)}(u,v,w) = \tilde{H}^{(\nu)}(u,v,w) \ \left[\mathcal{F} \circ \left(G^{(w)} \odot x^{(\nu)} \right) \right](u,v,w), \tag{6}$$

avec \mathcal{F} la transformée de Fourier bidimensionnelle, $G^{(w)}(l,m) = e^{-2\pi i w \left(\sqrt{1-l^2-m^2}-1\right)}$ qui rend compte des effets non coplanaires et les mêmes correspondances que ci-dessus. On se restreint aux petits champs de vue, l'interférométrie à très grand champ de vue qui exploite le formalisme sphérique [Carozzi, 2015] étant encore un champ relativement inexploré.

Le chapitre 4 aborde la SAS dans le deuxième cas (et dans le premier par extension, comme nous allons le voir). Le chapitre 5 s'inscrit dans le cadre du troisième cas.

Chapitre 4 : séparation aveugle de sources et déconvolution conjointe d'observations grand champ

Les méthodes standard de SAS ne peuvent être employées sur les données multispectrales radio qu'après un prétraitement qui permet de tenir compte de la réponse du télescope (due aux effets instrumentaux voire à un échantillonnage incomplet dans le cas d'un interféromètre). Toutefois, inclure la réponse du télescopes directement au sein du processus de séparation devrait être plus performant, car cela permettrait de prendre en compte précisément les modèles d'observation, de mélange et de bruit en une seule opération.

Dans ce chapitre, nous étudions ainsi un algorithme de séparation aveugle de sources et de déconvolution (SASD) conjointe. Nous traitons le cas des données grand champ échantillonnées sur la sphère, qui sont maintenant courantes dans des domaines scientifiques tels que la radioastronomie.

La fonction coût suivante est considérée :

$$\underset{\mathbf{A} \ge \mathbf{0}, \mathbf{S} \ge \mathbf{0}}{\operatorname{argmin}} \sum_{(\mathfrak{l}, \mathfrak{m}) \in \mathcal{P}} \frac{1}{2} \left(\left\| \tilde{\mathbf{Y}}_{:(\mathfrak{l}, \mathfrak{m})} - \operatorname{diag} \left(\tilde{\mathbf{H}}_{:\mathfrak{l}} \right) \mathbf{A} \tilde{\mathbf{S}}_{:(\mathfrak{l}, \mathfrak{m})} \right\|_{2}^{2} + \sum_{i=1}^{I} \varepsilon^{(i, \mathfrak{l}, \mathfrak{m})} \left\| \tilde{\mathbf{S}}_{i(\mathfrak{l}, \mathfrak{m})} \right\|^{2} \right) + \left\| \mathbf{\Lambda} \odot \left(\mathbf{S} \mathbf{W}^{\top} \right) \right\|_{1}$$

$$\tag{7}$$

avec $\mathbf{H} \in \mathbb{R}^{J \times P}$ la matrice constituée des réponses du télescope (une réponse par ligne, c'est-à-dire par canal), et où le tilde rend compte d'une transformation en harmoniques sphériques. La résolution est basée sur un moindres carrés projeté alterné afin de combiner rapidité et précision. Par rapport au problème standard de SAS, la procédure fait appel à une régularisation supplémentaire, introduite par les $\{\varepsilon^{(i,\mathfrak{l},\mathfrak{m})}\}_{i,\mathfrak{l},\mathfrak{m}}$, pour traiter un problème naturellement mal posé. Nous étudions de manière approfondie l'impact de la régularisation dans la minimisation des moindres carrés et introduisons des schémas de régularisation adaptés aux statistiques des sources à estimer.

Sur la base de ces techniques de régularisation, nous présentons l'algorithme SDecGMCA (GMCA déconvolution sur la sphère) qui s'avère être une procédure de minimisation robuste et efficace. Des expériences numériques sur des exemples jouets et sur des données issues de simulations astrophysiques évaluent les performances de l'algorithme proposé dans un large éventail de scénarios de mélange. Des comparaisons avec des méthodes standard de SAS sont également effectuées, elles montrent que SDecGMCA est compétitif.

Nous notons finalement que la méthode proposée s'adapte aisément au cas plan, en remplaçant l'analyse en harmoniques sphériques par celle de Fourier.

Chapitre 5 : séparation aveugle de sources et déconvolution conjointe à partir de données interférométriques non coplanaires

L'avènement des radio-interféromètres à grande échelle impose le développement de nouvelles méthodes d'analyse capables de traiter les effets non coplanaires, dits effets w. Comme nous l'avons vu au chapitre 3, ces effets proviennent des longues lignes de base, qui annulent généralement l'hypothèse de coplanarité des lignes de base, ainsi que des observations grand champ, car les ondes incidentes ne sont plus planes. De façon similaire au chapitre 4, le modèle standard de mélange linéaire devient obsolète, ce qui nécessite la conception d'une méthode de SASD dédiée.

Dans ce chapitre, nous proposons un algorithme pour aborder la SASD avec des données interférométriques non coplanaires. Pour ce faire, nous reprenons le cadre de w-stacking [Offringa et al., 2014] qui offre un bon compromis entre précision et coût calculatoire pour les problèmes de déconvolution de données non coplanaires. Les données sont supposées être échantillonnées sur un grille régulière et placées dans un tenseur $\boldsymbol{\mathcal{Y}} \in \mathbb{R}^{J \times W \times P}$, avec W le nombre de plans w considérés. Le modèle d'observation s'écrit⁴ :

$$\boldsymbol{\mathcal{Y}} = \boldsymbol{\mathcal{H}} \odot \left[\boldsymbol{\mathcal{F}} \circ \boldsymbol{\mathcal{G}} \right] (\mathbf{AS}) + \boldsymbol{\mathcal{N}}, \tag{8}$$

avec $\mathcal{H} \in \mathbb{R}^{J \times W \times P}$ le masque de l'interféromètre (dans l'espace des visibilités), \mathcal{F} la transformée de Fourier bidimensionnelle, \mathcal{G} l'opérateur rendant compte de l'effet non coplanaire et $\mathcal{N} \in \mathbb{R}^{J \times W \times P}$ un bruit gaussien.

La fonction coût considérée est adaptée de la façon suivante :

$$\underset{\mathbf{A} \ge \mathbf{0}, \mathbf{S} \ge \mathbf{0}}{\operatorname{argmin}} \frac{1}{2} q(\boldsymbol{\mathcal{Y}} - \boldsymbol{\mathcal{H}} \odot [\boldsymbol{\mathcal{F}} \circ \boldsymbol{\mathcal{G}}] (\mathbf{AS})) + \frac{1}{2} \sum_{i, \omega, p} \varepsilon^{(i, \omega, p)} \left| [\boldsymbol{\mathcal{F}} \circ \boldsymbol{\mathcal{G}}] (\mathbf{S})_{i \omega p} \right|^2 + \left\| \boldsymbol{\Lambda} \odot \left(\mathbf{SW}^\top \right) \right\|_1,$$
(9)

avec q une forme quadratique dépendant de la covariance de \mathcal{N} . Le problème formulé étant particulièrement complexe à traiter, nous choisissons de le résoudre avec un moindres carrés projeté alterné pour son faible coût calculatoire. Comme pour SDecGMCA, l'ajout d'une régularisation est cependant rendue nécessaire pour permettre l'inversion par les moindres carrés. Nous adaptons les schémas de régularisations introduits dans le chapitre précédent au cas non coplanaire.

L'algorithme proposé est appelé wGMCA. Des expériences numériques sur des données synthétiques et réalistes sont réalisées dans différents contextes expérimentaux (degré de non-coplanarité, densité d'échantillonnage de l'espace des visibilités, distribution et nombre de canaux, niveau de bruit, conditionnement du mélange). La méthode proposée présente une certaine robustesse à l'initialisation et au choix des hyperparamètres. Comparé aux méthodes standard de déconvolution et de séparation séquentielles, l'algorithme proposé s'avère performant, étant capable d'estimer les sources avec une meilleure résolution et précision, ce qui augmente la qualité globale de la séparation.

 $^{^4 \}mathrm{On}$ donne ici une forme simplifiée du modèle d'observation, à titre de résumé.

Notations and acronyms

Notations

Matrices and tensors

- **x**: a vector.
- X: a matrix.
- \mathcal{X} : a three-dimensional tensor.
- \mathbf{x}_i , \mathbf{X}_{ij} and $\mathbf{\mathcal{X}}_{ijk}$: respectively, the *i*th, (i, j)th and (i, j, k)th entries of vector \mathbf{x} , matrix \mathbf{X} and tensor $\mathbf{\mathcal{X}}$.
- $\mathbf{X}_{i:}, \mathbf{X}_{:j}, \boldsymbol{\mathcal{X}}_{:jk}, \boldsymbol{\mathcal{X}}_{i:k}$ and $\boldsymbol{\mathcal{X}}_{ij:}$: column vectors formed by fixing the indicated indices of matrix \mathbf{X} or tensor $\boldsymbol{\mathcal{X}}$, depending on the case.
- $\mathcal{X}_{i::}, \mathcal{X}_{:j:}$ and $\mathcal{X}_{::k}$: matrices formed by fixing the indicated index of tensor \mathcal{X} .
- x^* : the complex conjugate of x, where x can be a scalar, a vector, a matrix or a tensor.
- \mathbf{X}^{\top} : the transpose of matrix \mathbf{X} .
- **X**[†]: the conjugate-transpose of matrix **X**.

Operators

- $\|\cdot\|_p$, $p \in \mathbb{N}$: the ℓ_p -norm for vectors, and the entrywise ℓ_p -norm for matrices.
- $\|\cdot\|_{\ell_p}$, $p \in \mathbb{N}$: the matrix ℓ_p -norm induced by the vector ℓ_p -norm.
- diag(·): the diagonal operator, which produces a diagonal matrix from the input parameter (*e.g.*, if **x** is a vector of size P, then diag(**x**) is a diagonal matrix of size $P \times P$ such that diag(**x**)_{pp} = **x**_p).

- $\mathcal{F}(\cdot)$: the 2-dimensional continuous Fourier transform or the spherical harmonic transform⁵, depending if the input function is defined on a plane or on the sphere. The tilde is used as a shortcut: $\tilde{x} := \mathcal{F}(x)$.
- **F**: the 2-dimensional discrete Fourier matrix or the discrete harmonic transform matrix, depending on the input. The tilde is likewise used as a shortcut:
 - $-\tilde{\mathbf{x}} \coloneqq \mathbf{F}\mathbf{x}$, noting furthermore that in the spherical case, $\tilde{\mathbf{x}}_{(\mathfrak{l},\mathfrak{m})}$ designates the spherical harmonic coefficient associated to degree \mathfrak{l} and mode \mathfrak{m} ,
 - \mathbf{X} : the stack of the discrete Fourier transforms or spherical harmonic transforms of the rows of \mathbf{X} ($\mathbf{\tilde{X}}_{i:} \coloneqq \mathbf{F}\mathbf{X}_{i:}$, which is equivalently written $\mathbf{\tilde{X}} \coloneqq \mathbf{X}\mathbf{F}^{\top}$).
- *: the convolution product between two functions, in the continuous or discrete case depending on the situation. For functions defined on the sphere, the isotropic convolution product is used.
- \odot : the entrywise multiplication, also known as the Hadamard product. When written in a superscript of a matrix, it means a entrywise exponent (*e.g.*, $[\mathbf{X}^{\odot 2}]_{ij} = (\mathbf{X}_{ij})^2 \neq [\mathbf{X}^2]_{ij}$).
- \oslash : the entrywise division.
- $\bullet\,$ o: the function composition operator.
- $\iota_{(\cdot)}$: the characteristic function of the set in subscript.
- $\Pi_{(\cdot)}$: the orthogonal projection onto the set in subscript.
- $\mathcal{T}_{\lambda}(\cdot) \coloneqq \operatorname{sign}(\cdot) \max(|\cdot| \lambda, 0)$: the soft-thresholding operator with threshold λ where the sign, max and absolute-value functions are intended entrywise.

Dummy indices

The following dummy indices are reserved for one usage only, to facilitate readability:

- p: pixel or sample,
- i: source,
- j: observation channel,
- *n*: iteration,

⁵Elements of spherical signal analysis are provided in Appendix B.

- s: scale (of a multiscale representation),
- *m*: spectrum model,
- \mathfrak{l} and \mathfrak{m} : degree and mode (in a spherical harmonic decomposition),
- ω : w-plane (in the context of noncoplanar interferometry).

Moreover, dummy indices usually iterate from 1 to their variable in upper case (*e.g.* $p \in [1 \dots P], i \in [1 \dots I]$).

Specific terms

The following variables refer to the same quantities throughout the manuscript:

- $\mathbf{Y} \in \mathbb{R}^{J \times P}$ and $\boldsymbol{\mathcal{Y}} \in \mathbb{R}^{J \times W \times P}$: noisy data.
- $\mathbf{A} \in \mathbb{R}^{J \times I}$: mixing matrix.
- $\mathbf{S} \in \mathbb{R}^{I \times P}$: source matrix.
- $\mathbf{N} \in \mathbb{R}^{J \times P}$ and $\mathcal{N} \in \mathbb{R}^{J \times W \times P}$: noise term.
- $\mathbf{H} \in \mathbb{R}^{J \times P}$ and $\mathcal{H} \in \mathbb{R}^{J \times W \times P}$: measurement operator (*e.g.*, antenna power pattern, interferometric mask).
- $\mathbf{W} \in \mathbb{R}^{P' \times P}$: transform towards a sparsyfing domain (that is, where the signals are sparse).
- $\Lambda \in \mathbb{R}^{I \times P'}$: sparse hyperparameters,
- κ and K: parameters associated to the sparse hyperparameters (see Section 1.3.3.4),
- $\mathcal{D}(\mathbf{Y}, \mathbf{AS})$ or $\mathcal{D}(\mathbf{\mathcal{Y}}, \mathbf{AS})$: data-fidelity term,
- $h_{\mathcal{A}}(\mathbf{A})$: mixing matrix regularization term,
- $h_{\mathcal{S}}(\mathbf{S})$: source matrix regularization term,
- $\mathcal{K}_{\mathcal{A}} \coloneqq \{\mathbf{A} \in \mathbb{R}^{J \times I}, \forall (j, i), \mathbf{A}_{ji} \geq 0, \forall i, \|\mathbf{A}_{:i}\|_{2} \leq 1\}$: intersection of the non-negative orthant for mixing matrices and the unit balls of the columns of the mixing matrices.
- $\mathcal{O}_{\mathcal{A}} := \{ \mathbf{A} \in \mathbb{R}^{J \times I}, \forall (j, i), \mathbf{A}_{ji} \geq 0 \}$: non-negative orthant for mixing matrices.
- $\mathcal{O}_{\mathcal{S}} := \{ \mathbf{S} \in \mathbb{R}^{I \times P}, \forall (i, p), \mathbf{S}_{ip} \ge 0 \}$: non-negative orthant for source matrices.

- I: identity matrix.
- 0 and 1: vector or matrix of zeros and ones, respectively.
- c: speed of light in vacuum.
- ν : electromagnetic frequency.
- λ : electromagnetic wavelength.

Miscellaneous

- := : is defined by.
- $(\cdot)^{(i)}$: *i*th element of a sequence of variables (*e.g.*, of an iterative scheme, of a set).
- \hat{x} : the estimate of x, where x can be a scalar, a vector, a matrix or a tensor.
- x^{\diamond} : the ground truth of x, where x can be a scalar, a vector, a matrix or a tensor.

Attention is drawn to the fact that there is no link between two variables that have the same letter but different cases or styles (*e.g.*, $\Lambda_{i:} \neq \lambda$, $l \neq \mathfrak{l}$). An attempt has been made to avoid using variables with the same letters in a single equation.

Acronyms

- (p)ALS: (projected) alternating least squares.
- BCD: block coordinate descent.
- BSS: blind source separation.
- C_A: mixing matrix criterion.
- EoR: epoch of reionization.
- FBS: forward-backward splitting.
- FFT: fast Fourier transform.
- GMCA: generalized morphological component analysis.
 - DecGMCA: deconvolution GMCA.

- SDecGMCA: spherical deconvolution GMCA.
- sGMCA: semi-blind GMCA.
- wGMCA.
- HALS: hierarchical alternating least squares.
- HPBW: half power beam width.
- IAE: interpolatory autoencoder.
- ICA: independent component analysis.
- MAD: median absolute deviation estimator.
- MU: multiplicative update.
- NMF: nonnegative matrix factorization.
- NMSE: normalized mean square error.
- PALM: proximal alternating linearized minimization.
- SAD: spectral angular distance.
- SBSS: sparse blind source separation.
- SDR, SIR, SNR, SAR: signal-to-distortion, interference, noise and artifact ratio, respectively.

Introduction

Context

The latest technological developments allow for a continuous increase in data acquisition capabilities in scientific experiments. This requires the development of new and more accurate data analysis methods to extract valuable information.

This thesis focuses on the unsupervised analysis of multivalued data. Multivalued data are composed of several measurements of the same scene or process under different but coherent observation conditions. They are found in various applications, for instance in biology with electroencephalographic data (see Fig. 1), in chemistry with liquid chromatography-mass spectrometry data (see Fig. 2) or in physics with multispectral imaging (see Fig. 3). In the given examples, the observed scenes or processes are constituted of distinct elementary components (respectively neurons, chemical compounds, atoms or molecules such as iron or water); each component contributes differently in each observation, thus paving the way for their unmixing.

Blind source separation (BSS) algorithms are unsupervised methods that are the cornerstone of multivalued data analysis by allowing for physically meaningful data decompositions. The unsupervised aspect allows to separate the components automatically, without any a priori knowledge of their presence or their properties (depending on the case, spatial distribution, spectral or temporal response, *etc.*).

The developments of this thesis are especially aimed at the analysis of multispectral⁶ image data in astrophysics. Two applications are particularly identified:

• 21-cm intensity mapping. The measurement of the hydrogen 21-cm line emission is expected to provide key information on the history of our Universe. As illustrated in Figure 4, this signal is not observed directly but mixed with other components from our galaxy (*e.g.* synchrotron and free-free emissions) or

⁶Sometimes the term multispectral is reserved to the case where there are less than ten spectral bands, the term hyperspectral being used otherwise. Hereafter the term multispectral will be used interchangeably for both cases.



Figure 1: Scalp electroencephalographic data. Electrodes placed on a person's scalp measure the resulting electrical activity of the underlying neurons, and potentially non-cerebral artifacts such as muscles. The data have a temporal variable and a spatial variable, that is related to the positioning of the electrodes. Image from [Hu et al., 2007].



Figure 2: Liquid chromatography–mass spectrometry acquisition and data. A chemical fluid is introduced into a chromatography column; the compounds that constitute the fluid flow out at different times and are then analyzed with a mass spectrometer. The data have a temporal variable and a mass-to-charge variable. Image from [Rapin, 2014, Chenot, 2017].



Figure 3: Hyperspectral image data from the OMEGA spectrometer which was aboard the Mars Express Orbiter (https://mars-express.cnes.fr/en/MEX/index. htm). A two-dimensional image is acquired by the spectrometer OMEGA along 128 spectral bands, at wavelengths ranging from 0.98 µm to 2.73 µm. The data have a wavelength/electromagnetic-frequency variable and a two-dimensional spatial variable. Image from [Chenot, 2017].

beyond (*e.g.*, radio galaxies and clusters, present in the form of point sources). The resulting problem can therefore be stated in the form of a BSS problem. Two sub-applications are distinguished, depending on the redshift/frequency range considered:

- at a redshift z from 6 to 30, which requires multichannel radio observations in the range [50, 200] MHz, the 21-cm line emission can only be measured with precision by continental radio-interferometers, such as the forthcoming Square Kilometer Array⁷ (SKA). This signal should provide more information on the epoch of reionization (EoR), which is an early period of the Universe that is poorly known.
- at later times (post-reionization), for a redshift z from 0.1 to 1, which corresponds to the radio frequency range [700, 1300] MHz, the 21-cm line emission can also be measured by single-dish radio-telescopes. Accessing this signal will help understanding the formation of the large-scale structures in the Universe.
- X-ray spectro-imaging. BSS algorithms can be employed to separate the elementary components constituting astronomical objects (*e.g.*, supernova remnants), in the aim of better understanding the physical processes that occur. The forthcoming Athena⁸ telescope will provide measurements in a particularly

⁷https://www.skatelescope.org/ ⁸https://sci.esa.int/web/athena



Figure 4: Galactic and extra-galactic foregrounds contaminating the EoR signal. Image from [Jelić, 2010].

large bandwidth with a fine spectral resolution. Undetected sources so far may be extracted, as long as BSS algorithms are improved.

Problems and challenges

In the context of astrophysics, standard BSS algorithms will be challenged by the deluge of data from key upcoming projects. For BSS methods to provide valuable information, several problems in the field of signal processing must be addressed:

- The increased spatial resolution achieved by some telescopes should allow distinguishing very weak emissions – up to a few orders of magnitude below the most powerful ones. Therefore, the first issue concerns developing very precise separation methods. In addition, instrumental noise can dominate the signal of interest.
- Multispectral observations are distorted by wavelength-dependent instrumental responses. In the case of large-band measurements, BSS methods must cope with particularly heterogeneous data, with varying spatial resolutions.
- Interferometric data are obtained by correlation of pairs of antennas. As a first approximation, each pair of antennas produces a point measurement in the harmonic domain. The number of pairs being finite, the interferometer sampling

is incomplete. Therefore, BSS methods applied to interferometric measurements must include specific regularizations.

• In the case of long baseline interferometry (*i.e.*, a large distance between the most distant antennas), the classical assumption of incident plane waves becomes obsolete. The instrumental response taking into account the non-coplanar effect is much more complex; as a first approximation, it becomes non-stationary on the observed field. A similar phenomenon is to be considered in the case of single-dish wide-field observations. Accurate separation methods must take into account these non-coplanar effects.

Contributions

To address these issues, the thesis focuses on three objectives:

- **I** Develop a separation method that allows taking into account prior physical knowledge of expected components. Resorting to such semi-supervised or semi-blind approach is expected to better discriminate between sources, especially the less powerful ones, and provide more physically relevant information.
- **II** Design a BSS scheme that can process multispectral data distorted by potentially channel-dependent measurement operators. For instance, this is the case of large-band observations which have varying instrumental responses along the channels or interferometric data whose measurement operators in Fourier space are masks. Coping with such data requires tackling an extra deconvolution step, thus leading to a joint deconvolution and blind source separation (DBSS) problem. This becomes highly challenging when applied to datasets sampled on the sphere, such as those provided by wide-field observations.
- **III** Develop a BSS method for non-coplanar interferometric data. This is a novel topic in several respects. Indeed, the single-channel deconvolution of non-coplanar interferometric data is already a research topic in itself; adding blind source separation from multi-channel data is an entirely open problem.

Outline

This manuscript builds upon five chapters. Chapter 1 briefly reviews blind source separation algorithms for processing multispectral data in astrophysics. Chapter 2 proposes a semi-supervised separation method that allows leveraging the potential prior knowledge of expected components, thus tackling objective \mathbf{I} . Chapter 3

presents the basics of radio astronomy, with the primary goal of obtaining observational models to tackle BSS from radio data afterward. Chapter 4 investigates a joint deconvolution and blind source separation algorithm to analyze multispectral data acquired on the sphere, answering objective Π . Finally, Chapter 5 focuses on developing a joint deconvolution and blind source separation algorithm for non-coplanar radio-interferometric data, thus meeting objective Π .

Chapter 1

A quick review of blind source separation

This chapter introduces blind source separation (BSS) from a mathematical perspective. The main resolution methods in the literature are presented. In doing so, the mathematical optimization tools that we will need later on are introduced. In addition, performance metrics for BSS are provided.

1.1 Mixture model

Astrophysical multispectral data are generically constituted of observations of the same region of the sky at different wavelengths across the electromagnetic spectrum. In this respect, let $\mathbf{Y}_{j:} \in \mathbb{R}^{P}$ be a measurement made of P samples or pixels at a channel $j \in [1 \dots J]$, which corresponds to a given wavelength or spectral band. $\mathbf{Y}_{j:}$ is typically a two-dimensional image, which is flattened into a vector.

The linear mixture model provides a good approximation to decompose multispectral data in astrophysics (and even more generally in the other application fields of BSS [Comon and Jutten, 2010]). In this regard, $\mathbf{Y}_{j:}$ can be expressed as the weighted sum of I elementary sources $\{\mathbf{S}_{i:} \in \mathbb{R}^{P}\}_{i \in [1...I]}$:

$$\mathbf{Y}_{j:} = \sum_{i=1}^{I} \mathbf{A}_{ji} \mathbf{S}_{i:} + \mathbf{N}_{j:}, \qquad (1.1)$$

where $\mathbf{A}_{ji} \in \mathbb{R}$ is the contribution of component *i* at channel *j* and $\mathbf{N}_{j:}$ is a noise term¹. Equation (1.1) can be rewritten with matrices (see Fig. 1.1), yielding:

$$\mathbf{Y} = \mathbf{AS} + \mathbf{N},\tag{1.2}$$

¹The noise is written as additive to simplify the notation; however, we do not make this assumption yet.



Figure 1.1: Illustration of the (noiseless) matrix product $\mathbf{Y} = \mathbf{AS}$ in BSS. The rows of \mathbf{S} include sources and the columns of \mathbf{A} contain spectra. \mathbf{Y} includes in its rows the observed mixtures, per channel.

where $\mathbf{Y} \in \mathbb{R}^{J \times P}$ contains the multispectral mixtures, $\mathbf{A} \in \mathbb{R}^{J \times I}$ is the so-called mixing matrix and $\mathbf{S} \in \mathbb{R}^{I \times P}$ is the source matrix. With such notations, the columns of \mathbf{A} include the electromagnetic spectra of the elementary components. In this thesis, the over-determined case is considered, *i.e.*, the number of channels is assumed to be larger than the number of sources $(J \ge I)$.

In the scope of BSS, the objective is to jointly estimate the matrices **A** and **S** from their noise-corrupted product **Y**. It should be noted that **A** and **S** can be recovered up to two indeterminacies at best, namely of scale and permutation, which stem from the matrix product model. As an illustration, if \mathbf{A}^{\diamond} and \mathbf{S}^{\diamond} are the ground-truth matrices, then $\alpha \mathbf{A}^{\diamond}$ and $\alpha^{-1} \mathbf{S}^{\diamond}$ have the same product (scale indeterminacy), as well as $[\mathbf{A}_{:I}^{\diamond} \dots \mathbf{A}_{:1}^{\diamond}]$ and $[\mathbf{S}_{I:}^{\diamond} \dots \mathbf{S}_{1:}^{\diamond}]^{\top}$ (permutation indeterminacy). The performance metrics that will be employed will account for these degeneracies (see Subsection 1.5).

BSS is a very challenging, ill-posed problem that requires the use of prior information. Different kinds of priors have led to diverse classes of component separation algorithms, which are reviewed hereinafter.

1.2 Independent component analysis

Independent component analysis (ICA) is a statistical method that assumes that the sources are mutually independent random signals, of which at most one follows a Gaussian distribution, the others being non-Gaussian.

The mixing of random sources in the context of BSS tends to make the observations Gaussian. This is a consequence of the central limit theorem, which establishes the convergence in law of the sum of a sequence of random variables to the normal distribution. The gist of ICA is to use the "non-Gaussianity" of the sources as a proxy for their independence. ICA aims to find the linear transformation that maximizes the sources' non-Gaussianity to separate them.

Numerous ICA-based methods exist, they mainly differ in the kind of measure of independence. Among the most renowned are [Comon and Jutten, 2010]:

- Infomax, which minimizes the mutual information between the sources.
- JADE, that diagonalizes the fourth-order cumulant tensor of the sources, which amounts to minimizing the fourth-order cross-cumulants thus favoring the independence of the sources.
- FastICA, which maximizes the negentropy of the sources, that is the difference between the statistical entropy of a Gaussian distribution and the entropy of the sources.

Regarding BSS applications in astrophysics, three limitations of ICA-based methods are identified. First, sources often exhibit spatial correlations, for example when they originate from the same physical processes; therefore, the mutual independence assumption is not fulfilled. Second, ICA-based methods are generally not robust to noise contamination – that said, we note the existence of derivations such as spectral matching ICA (SMICA) [Cardoso et al., 2002] that address this specific point. Third, they offer a rather inflexible framework, making it difficult to consider more specific mixture models or add particular priors.

1.3 Cost-function-based methods

1.3.1 General framework

The source separation problem can be formulated as a cost function to be minimized. In this section, the following setting will generically be considered:

$$\underset{\mathbf{A},\mathbf{S}}{\operatorname{argmin}} \mathcal{D}(\mathbf{Y}, \mathbf{AS}) + h_{\mathcal{S}}(\mathbf{S}) + h_{\mathcal{A}}(\mathbf{A}), \qquad (1.3)$$

where $\mathcal{D}(\cdot, \cdot)$ is the data-fidelity term, which depends on both the mixing model and the noise statistics, and $h_{\mathcal{S}}(\cdot)$ and $h_{\mathcal{A}}(\cdot)$ are regularization terms that account for the priors on the sources and the mixing matrix, respectively.

Because of the product AS, Eq. (1.3) is not convex. In consequence, no algorithm can guarantee convergence to the global minimum. At best, some methods can ensure the convergence towards a stationary point, which will depend on the starting point.

1.3.2 Non-negative matrix factorization

Non-negative matrix factorization (NMF) is an extensive class of source separation algorithms which assume that both the mixing matrix and the sources are constituted of non-negative elements. In this regard, the regularization terms in Eq. (1.3) are given by $h_{\mathcal{S}}(\mathbf{S}) = \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S})$ and $h_{\mathcal{A}}(\mathbf{A}) = \iota_{\mathcal{O}_{\mathcal{A}}}(\mathbf{A})$, with $\mathcal{O}_{\mathcal{S}}$ and $\mathcal{O}_{\mathcal{A}}$ non-negative orthants².

NMF methods generally assume that the data-fidelity term is based on a β -divergence. β -divergences are estimators that can be deduced from several noise distributions encountered in physical applications. If x and y are two non-negative scalars, the β -divergence between x and y is given by:

$$d_{\beta}(x,y) \coloneqq \begin{cases} \frac{x}{y} - \log \frac{x}{y} - 1 & \text{if } \beta = 0, \\ x \log \frac{x}{y} - x + y & \text{if } \beta = 1, \\ \frac{1}{\beta(\beta-1)} \left(x^{\beta} + (\beta-1)y^{\beta} - \beta x y^{\beta-1} \right) & \text{otherwise.} \end{cases}$$
(1.4)

In the case of NMF, the derived data-fidelity is the sum of the β -divergences applied to all entries of the data matrix and the mixture model:

$$\mathcal{D}(\mathbf{Y}, \mathbf{AS}) \coloneqq \sum_{j, p} d_{\beta} \left(\mathbf{Y}_{jp}, (\mathbf{AS})_{jp} \right).$$
(1.5)

Three particular cases are worth being mentioned:

- $\beta = 2$: the Frobenius norm, stemming from a Gaussian noise,
- $\beta = 1$: the Kullback-Leibler divergence, derived from a Poisson contamination (*e.g.*, in low count observations, as in X-ray imaging),
- $\beta = 0$: the Itakura–Saito divergence, which comes from a multiplicative noise (and which is preferably used in audio separation for the resulting perceptual proprieties).

²
$$\mathcal{O}_{\mathcal{S}} \coloneqq \{ \mathbf{S} \in \mathbb{R}^{I \times P}, \forall (i, p), \mathbf{S}_{ip} \ge 0 \} \text{ and } \mathcal{O}_{\mathcal{A}} \coloneqq \{ \mathbf{A} \in \mathbb{R}^{J \times I}, \forall (j, i), \mathbf{A}_{ji} \ge 0 \}.$$

Multiplicative update A standard minimization scheme for solving Eq. (1.3) in the case of a β -divergence-based data-fidelity term is the Multiplicative update algorithm (MU, see Alg. 1) [Lee and Seung, 2000, Févotte and Idier, 2011]. It is an iterative procedure that decreases the cost function by alternative updates of **A** and **S** using element-wise multiplications and divisions; it can be interpreted as a particular gradient descent method [Gillis, 2020]. The MU framework presents several advantages, among which:

- it ensures the convergence towards a stationary point,
- it is easy to implement, with no hyperparameter to tune,
- it scales well with large datasets and is among the fastest NMF algorithms for $\beta < 2$.

Algorithm 1 Multiplicative update (MU) Inputs: data Y, starting points A and S while convergence not reached do

(1) Update **S** with **A** fixed

$$\mathbf{S} \leftarrow \mathbf{S} \odot \left(\mathbf{A}^{\top} \left[(\mathbf{A}\mathbf{S})^{\odot(\beta-2)} \odot \mathbf{Y} \right] \right) \oslash \left(\mathbf{A}^{\top} [\mathbf{A}\mathbf{S}]^{\odot(\beta-1)}$$

(2) Update **A** with **S** fixed
 $\mathbf{A} \leftarrow \mathbf{A} \odot \left(\left[(\mathbf{A}\mathbf{S})^{\odot(\beta-2)} \odot \mathbf{Y} \right] \mathbf{S}^{\top} \right) \oslash \left([\mathbf{A}\mathbf{S}]^{\odot(\beta-1)} \mathbf{S}^{\top} \right)$

Outputs: mixing matrix A, sources S

Hierarchical alternating least squares Another algorithm that has gained interest is Hierarchical alternating least squares (HALS, see Alg. 2) [Gillis and Glineur, 2012]. It allows minimizing Eq. (1.3) in the case of additive Gaussian noise ($\beta = 2$, therefore leading to $\mathcal{D}(\mathbf{Y}, \mathbf{AS}) = \|\mathbf{Y} - \mathbf{AS}\|_2^2/2$). Built upon the block-projected gradient descent scheme, it significantly speeds up the minimization compared to MU by sequentially updating the rows of **S** and the columns of **A**. The convergence towards a stationary point is also guaranteed.

Extensions The MU and in a lesser extent the HALS minimization schemes have been adapted to include extra constraints or regularizations (*e.g.*, the orthogonality of the mixing matrix to promote spectra that do not overlap, or the smoothness of the sources [Gillis, 2020]).
Algorithm 2 Hierarchical alternating least squares (HALS)

Inputs: data Y, starting points A and S while convergence not reached do

(1) Update **S** with **A** fixed
for
$$i = 1, 2, ..., I$$
 do
 $\mathbf{S}_{i:} \leftarrow \max\left(\mathbf{0}, \frac{1}{\|\mathbf{A}_{:i}\|_{2}^{2}}\left(\mathbf{Y}^{\top}\mathbf{A}_{:i} - \sum_{k \neq i} \left(\mathbf{A}_{:i}^{\top}\mathbf{A}_{:k}\right)\mathbf{S}_{k:}\right)\right)$
(2) Update **A** with **S** fixed
for $i = 1, 2, ..., I$ do
 $\mathbf{A}_{:i} \leftarrow \max\left(\mathbf{0}, \frac{1}{\|\mathbf{S}_{i:}\|_{2}^{2}}\left(\mathbf{Y}\mathbf{S}_{i:} - \sum_{k \neq i} \left(\mathbf{S}_{i:}^{\top}\mathbf{S}_{k:}\right)\mathbf{A}_{:k}\right)\right)$

Outputs: mixing matrix \mathbf{A} , sources \mathbf{S}

In this thesis, the developed methods will be compared to sparse NMF (SNMF), which enforces the sparsity of the sources by considering the following source regularization:

$$h_{\mathcal{S}}(\mathbf{S}) = \lambda \|\mathbf{S}\|_1 + \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}), \tag{1.6}$$

with λ a sparsity regularization hyperparameter.

In practice, the sparsity regularization associated with the scale indeterminacy inherent to BSS gives degenerated solutions such as $\|\mathbf{S}_{i:}\|_2 \to 0$ and $\|\mathbf{A}_{:i}\|_2 \to \infty$. To prevent this phenomenon, a constraint on the norm of the columns of the mixing matrix \mathbf{A} is added, yielding the following mixing matrix constraint:

$$h_{\mathcal{A}}(\mathbf{A}) = \iota_{\mathcal{B}_{\mathcal{A}}}(\mathbf{A}) + \iota_{\mathcal{O}_{\mathcal{A}}}(\mathbf{A}) = \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}),$$
(1.7)

where $\mathcal{B}_{\mathcal{A}} \coloneqq {\mathbf{A} \in \mathbb{R}^{J \times I}, \forall i, \|\mathbf{A}_{:i}\|_{2} \leq 1}$ and thus $\mathcal{K}_{\mathcal{A}} \coloneqq \mathcal{B}_{\mathcal{A}} \cap \mathcal{O}_{\mathcal{A}}$.

[Le Roux et al., 2015] proposed an adaptation of the MU framework to tackle SNMF (see Alg. 3). The algorithm includes renormalizations of the columns of the mixing matrix so as to fulfill the norm constraint $\iota_{\mathcal{B}_{\mathcal{A}}}(\cdot)$. It distinguishes itself from other SNMF algorithms by processing the normalization more rigorously.

In the context of NMF, the pure-pixel assumption states that each component is active alone in at least one sample of the observed data. This strong assumption offers theoretical guarantees of mixture separability [Donoho and Stodden, 2003].

Regarding applications, NMF algorithms are very popular in remote sensing (e.g., for multispectral terrestrial observations), in part since the data generally satisfy the pure-pixel assumption.

Yet, with a view to tackling multispectral data in astrophysics, NMF presents several

Algorithm 3 Sparse non-negative matrix factorization (SNMF) by [Le Roux et al., 2015]

Inputs: data **Y**, regularization hyperparameter λ , starting points **A** and **S**

 $\mathbf{A} \leftarrow \left[\mathbf{A}_{:1} / \left\|\mathbf{A}_{:1}\right\|_{2}, \mathbf{A}_{:2} / \left\|\mathbf{A}_{:2}\right\|_{2}, \dots, \mathbf{A}_{:I} / \left\|\mathbf{A}_{:I}\right\|_{2}\right] \text{ \# column-wise normalization}$

while convergence not reached do

(1) Update **S** with **A** fixed $\mathbf{S} \leftarrow \mathbf{S} \odot \left(\mathbf{A}^{\top} \left[(\mathbf{A}\mathbf{S})^{\odot(\beta-2)} \odot \mathbf{Y} \right] \right) \oslash \left(\mathbf{A}^{\top} [\mathbf{A}\mathbf{S}]^{\odot(\beta-1)} + \lambda \right)$ (2) Update **A** with **S** fixed $\mathbf{U} \leftarrow \mathbf{A} \odot \left(\mathbf{1}\mathbf{1}^{\top} \left[\mathbf{A} \odot \left(\left[(\mathbf{A}\mathbf{S})^{\odot(\beta-1)} \right] \mathbf{S}^{\top} \right) \right] \right) \quad \text{# normalization purpose}$ $\mathbf{V} \leftarrow \mathbf{A} \odot \left(\mathbf{1}\mathbf{1}^{\top} \left[\mathbf{A} \odot \left(\left[(\mathbf{A}\mathbf{S})^{\odot(\beta-2)} \odot \mathbf{Y} \right] \mathbf{S}^{\top} \right) \right] \right) \quad \text{# normalization purpose}$ $\mathbf{A} \leftarrow \mathbf{A} \odot \left(\left[(\mathbf{A}\mathbf{S})^{\odot(\beta-2)} \odot \mathbf{Y} \right] \mathbf{S}^{\top} + \mathbf{U} \right) \oslash \left([\mathbf{A}\mathbf{S}]^{\odot(\beta-1)} \mathbf{S}^{\top} + \mathbf{V} \right)$ $\mathbf{A} \leftarrow [\mathbf{A}_{:1}/ \|\mathbf{A}_{:1}\|_{2}, \mathbf{A}_{:2}/ \|\mathbf{A}_{:2}\|_{2}, \dots, \mathbf{A}_{:I}/ \|\mathbf{A}_{:I}\|_{2}]$

Outputs: mixing matrix A, sources S

limits. Firstly, the minimization schemes do not allow the implementation of more complex regularizations (*e.g.*, sparsity in a transformed domain), which are nevertheless crucial in practice for separation quality, nor do they allow for the consideration of more complex mixture models. Secondly, NMF is sometimes shown to lack robustness with respect to the initial point; in other words, the algorithm can stabilize on a poor local minimum if the initialization is inadequate.

1.3.3 Sparse blind source separation

The methods developed in this thesis are part of the sparse BSS (SBSS) framework. Indeed, it allows a finer modeling of both priors and mixing model. First, we specify the assumptions we will commonly make on the cost function. Then we review algorithms in the literature that tackle the minimization problem thereby formed.

1.3.3.1 Assumptions on the cost function

Data-fidelity term Thereafter, the noise contamination is assumed additive Gaussian. The data-fidelity term is therefore chosen as an Euclidean distance:

$$\mathcal{D}(\mathbf{Y}, \mathbf{AS}) = \frac{1}{2} \|\mathbf{Y} - \mathbf{AS}\|_2^2.$$
(1.8)

Sources Sparsity is a powerful prior for solving inverse problems [Mallat, 2008]. As we have seen, SNMF promotes the sparsity of the sources in the direct domain. However, astrophysical sources are generally not point-like; they instead present extended and multiscale structures. Therefore, the sparsity of the sources is enforced in a transformed domain – typically based on two-dimensional wavelets.

Let $\mathbf{W} \in \mathbb{R}^{P' \times P}$ be a dictionary of a chosen sparsifying representation. Following the adopted notations, the source matrix expressed in the transformed domain writes $\mathbf{SW}^{\top} \in \mathbb{R}^{I \times P'}$. We choose to enforce the source sparsity with an analysis formulation and not a synthesis formulation, as the former is more robust in the context of SBSS when sources are approximately sparse [Rapin, 2014]. Therefore, the regularization term writes:

$$h_{\mathcal{S}}(\mathbf{S}) = \left\| \mathbf{\Lambda} \odot \left(\mathbf{S} \mathbf{W}^{\top} \right) \right\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}).$$
(1.9)

The regularization hyperparameter $\Lambda \in \mathbb{R}^{I \times P'}$ is written in matrix form to allow it to depend on the samples. This general framework will enable using the reweighted- ℓ_1 scheme later on in order to reduce regularization bias [Candès et al., 2008]. Moreover, as in NMF, the non-negativity constraint on the sources (in the direct domain) is added. Indeed, this assumption is often, if not always, verified in astrophysical applications. However, it is highlighted that it can easily be removed in all developments of this manuscript if desired.

Mixing matrix As for SNMF, in order to prevent the scale degeneracy triggered by the sparse regularization of the sources, a norm constraint is added to the non-negativity constraint of the mixing matrix, yielding:

$$h_{\mathcal{A}}(\mathbf{A}) = \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}),\tag{1.10}$$

where it is reminded that $\mathcal{K}_{\mathcal{A}} = \{ \mathbf{A} \in \mathbb{R}^{J \times I}, \forall (j, i), \mathbf{A}_{ji} \ge 0, \forall i, \|\mathbf{A}_{:i}\|_2 \le 1 \}.$

Altogether, the SBSS problem writes:

$$\operatorname{argmin}_{\mathbf{A},\mathbf{S}} \frac{1}{2} \|\mathbf{Y} - \mathbf{A}\mathbf{S}\|_{2}^{2} + \|\mathbf{\Lambda} \odot \left(\mathbf{S}\mathbf{W}^{\top}\right)\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}) + \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}).$$
(1.11)

Equation (1.11) is multi-convex, *i.e.*, convex with respect to both variables when the other is fixed, which advocates the use of alternating minimization schemes. However, it is composed of non-differentiable terms, which prevent resorting to smooth optimization techniques and rather call for the use of proximal algorithms, which are based on proximal operators [Parikh and Boyd, 2014].

1.3.3.2 Proximal operators of the regularization terms

In this section, we first define proximal operators and then provide the analytical expressions of the proximal operators associated with the two non-smooth regularization terms, $h_{\mathcal{A}}(\cdot)$ and $h_{\mathcal{S}}(\cdot)$, which we will need in the sequel.

Definition Let f be a proper³, lower semi-continuous⁴, convex function from a Hilbert space to $\mathbb{R} \cup \{+\infty\}$. The proximal operator of f, denoted prox_f , is defined by:

$$\operatorname{prox}_{f}(y) = \operatorname*{argmin}_{x} f(x) + \frac{1}{2} \|x - y\|_{2}^{2}.$$
(1.12)

The assumptions on f guarantee the existence and uniqueness of the minimum, and thus that prox_{f} is well-defined.

In some respects, proximal operators can be interpreted as generalizing projections. The simplest case is when f is the indicator function of a convex set, as prox_f is the orthogonal projection on this set. This interpretation can help appreciate why proximal operators constrain or regularize the solutions in proximal algorithms.

Proximal operator of the mixing matrix constraint term $\mathcal{K}_{\mathcal{A}}$ being a convex set, the proximal operator of the mixing matrix constraint $h_{\mathcal{A}}$ is the orthogonal projection on $\mathcal{K}_{\mathcal{A}}$. The latter is shown to be the composition of the orthogonal projections on $\mathcal{B}_{\mathcal{A}}$ and $\mathcal{O}_{\mathcal{A}}$:

$$\operatorname{prox}_{h_{\mathcal{A}}}(\mathbf{A}) = \Pi_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}) = \Pi_{\mathcal{B}_{\mathcal{A}}} \circ \Pi_{\mathcal{O}_{\mathcal{A}}}(\mathbf{A}),$$
(1.13)

where $[\Pi_{\mathcal{O}_{\mathcal{A}}}(\mathbf{A})]_{ji} = \max(0, \mathbf{A}_{ji})$ and $[\Pi_{\mathcal{B}_{\mathcal{A}}}(\mathbf{A})]_{:i} = \mathbf{A}_{:i} / \max(1, \|\mathbf{A}_{:i}\|_{2}).$

Proximal operator of the source regularization term The proximal operator of the source regularization term h_S has unfortunately no analytical form. It can be estimated numerically using iterative schemes, for instance with Generalized forward-backward splitting [Raguet et al., 2013]. However, this solution is likely too slow when included in alternating minimization schemes. Instead, we prefer to resort to approximations to derive an analytical formulation of the proximal operator.

In certain cases, the proximal operator of a sum of functions is the composition of the proximal operators of these functions [Yu, 2013]. In our case, the conditions are not entirely verified. Still, we will approximate the proximal operator of $h_{\mathcal{S}}$ as the composition of the proximal operators of the two terms it is composed of, namely the non-negativity constraint and the ℓ_1 -penalization.

³A proper function has value in $\mathbb{R} \cup \{+\infty\}$ and is not identically worth $+\infty$.

⁴A lower semi-continuous function verifies for every point x_0 of its domain $\lim_{x \to \infty} f(x) \ge f(x_0)$.

Second, the proximal operator associated to the sparse regularization term $\|\Lambda \odot (\mathbf{SW}^{\top})\|_1$ has no analytical form, unless **W** is orthogonal. In this case, it amounts to a soft-thresholding in the transformed domain, backprojected in the direct domain:

$$\operatorname{prox}_{\left\|\mathbf{\Lambda}\odot\left(\cdot\,\mathbf{W}^{\top}\right)\right\|_{1}}(\mathbf{S}) = \mathcal{T}_{\mathbf{\Lambda}}\left(\mathbf{S}\mathbf{W}^{\top}\right)\mathbf{W},\tag{1.14}$$

where $[\mathcal{T}_{\Lambda}(\mathbf{Y})]_{ip} = \operatorname{sign}(\mathbf{Y}_{ip}) \max(0, |\mathbf{Y}_{ip}| - \Lambda_{ip})$ is the soft-threshold operator. It is noted that this expression can be employed in proximal algorithms when \mathbf{W} is a tight frame [Elad, 2006]. Hereafter, we will keep this result as an approximation for transformations that are not tight frames.

Thus, the proximal operator of the source regularization term is approximated as:

$$\operatorname{prox}_{h_{\mathcal{S}}}(\mathbf{S}) \approx \Pi_{\mathcal{O}_{\mathcal{S}}}\left(\mathcal{T}_{\mathbf{\Lambda}}\left(\mathbf{S}\mathbf{W}^{\top}\right)\mathbf{M}^{\top}\right),\tag{1.15}$$

where **M** verifies $\mathbf{MW} = \mathbf{I}$ and $[\Pi_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S})]_{ip} = \max(0, \mathbf{S}_{ip}).$

1.3.3.3 Algorithms for sparse blind source separation

Several methods can tackle multi-convex non-smooth optimization problems such as SBSS. In this section, we propose to review some proximal algorithms for solving the SBSS problem. We do not aim for exhaustivity but instead choose to focus on the essential elements for the rest of the manuscript.

We refer the interested reader to [Chenot, 2017], [Kervazo, 2019] and the references hereinbelow for the theoretical details and the general formulation of the presented algorithms (*i.e.*, outside the scope of SBSS).

Block coordinate descent Block coordinate descent (BCD) is a simple iterative method to tackle multi-convex minimization problems [Tseng, 2001]. The gist of BCD is to minimize the cost function sequentially along each block coordinate, holding the others fixed.

In the case of SBSS, a natural choice of block coordinate is \mathbf{A} and \mathbf{S} . This leads to the following minimization subproblems, which are sequentially solved at each BCD iteration:

$$\begin{cases} \operatorname{argmin}_{\mathbf{S}} \frac{1}{2} \|\mathbf{Y} - \mathbf{AS}\|_{2}^{2} + \|\mathbf{\Lambda} \odot (\mathbf{SW}^{\top})\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}), \qquad (1.16) \end{cases}$$

$$\underset{\mathbf{A}}{\operatorname{argmin}} \frac{1}{2} \|\mathbf{Y} - \mathbf{AS}\|_{2}^{2} + \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}).$$
 (1.17)

Both equations being convex and composed of smooth and non-smooth terms, they can commonly be solved using the Forward-backward splitting (FBS) scheme [Combettes and Wajs, 2005]. The BCD algorithm derived to solve Eq. (1.11) is reported

in Alg. 4.

BCD is demonstrated to converge to a stationary point for the SBSS problem. Nevertheless, it is not used in practice for SBSS because of its high computational cost, as it nests two iterative schemes (each BCD iteration includes two FBS schemes).

Algorithm 4 Block coordinate descent (BCD) for SBSS

Inputs: data **Y**, starting points **A** and **S**, sparsifying transform **W**, thresholds Λ , parameters $\mu \in]0, 1[$ and $\nu \in]0, 1[$

while convergence not reached do

(1) Update S with A fixed: solve Eq. (1.16) using FBS $\gamma \leftarrow 2\mu \|\mathbf{A}\mathbf{A}^{\top}\|_{\ell_{2}}^{-1}$ while convergence not reached do $\mathbf{S} \leftarrow \mathbf{S} - \gamma \nabla_{\mathbf{S}} \mathcal{D}(\mathbf{Y}, \mathbf{A}\mathbf{S}) = \mathbf{S} - \gamma \mathbf{A}^{\top} (\mathbf{A}\mathbf{S} - \mathbf{Y})$ # gradient update $\mathbf{S} \leftarrow \operatorname{prox}_{\gamma h_{\mathcal{S}}}(\mathbf{S})$ # proximal update (2) Update A with S fixed: solve Eq. (1.17) using FBS $\gamma \leftarrow 2\nu \|\mathbf{S}^{\top}\mathbf{S}\|_{\ell_{2}}^{-1}$ while convergence not reached do $\mathbf{A} \leftarrow \mathbf{A} - \gamma \nabla_{\mathbf{A}} \mathcal{D}(\mathbf{Y}, \mathbf{A}\mathbf{S}) = \mathbf{A} - \gamma (\mathbf{A}\mathbf{S} - \mathbf{Y}) \mathbf{S}^{\top}$ # gradient update $\mathbf{A} \leftarrow \operatorname{prox}_{\gamma h_{\mathcal{A}}}(\mathbf{A})$ # proximal update **Outputs:** mixing matrix A, sources S

Proximal alternating linearized minimization The core idea of Proximal alternating linearized minimization (PALM) [Bolte et al., 2014] is to distribute the FBS iterative updates in the alternating scheme. Where BCD minimizes exactly subproblems (1.16) and (1.17), PALM minimizes proximal linearizations of these equations. Therefore, the update of each block coordinate is only a gradient update of the smooth data-fidelity term followed by a proximal update of the non-smooth regularization (see Alg. 5 for its application to SBSS).

The advantage of PALM is that the computational cost is low if explicit expressions of the proximal operators are available. In addition, PALM offers convergence guarantees in the SBSS context. The main drawbacks of PALM are the choice of the sparse regularization hyperparameters Λ , which can be tedious to fix, and the poor robustness to initial points – the starting point must be chosen carefully in order to have meaningful solutions [Kervazo et al., 2020a].

Projected alternating least squares As its name suggests, Projected alternating least squares (pALS) is an alternating algorithm, each iteration of which consists of an exact minimization of the data-fidelity term followed by the application of the

Algorithm 5 Proximal alternating linearized minimization (PALM) for SBSS

Inputs: data **Y**, starting points **A** and **S**, sparsifying transform **W**, thresholds Λ , parameters $\mu \in]0, 1[$ and $\nu \in]0, 1[$

while convergence not reached ${\bf do}$

Outputs: mixing matrix A, sources S

proximal operator of the related regularization term (see Alg. 6 for its application in SBSS).

The pALS algorithm has a very low computational cost when both the minimization of the data-fidelity term and the proximal operators have explicit formulations, like in SBSS. As such, it was among the first algorithms employed for source separation [Paatero and Tapper, 1994]. The main drawback of pALS is that it has no proof of convergence; in fact, it is not even shown that the stabilization point, when it exists, is a local minimizer of Eq. (1.11).

Algorithm 6 Projected alternating least squares (pALS) for SBSS	
--	--

Inputs: data \mathbf{Y} , starting point \mathbf{A} , sparsifying transform \mathbf{W} , thresholds $\mathbf{\Lambda}$

while convergence not reached do

(1) Update \mathbf{S} with \mathbf{A} fixed	
$\mathbf{S} \leftarrow \operatorname{argmin}_{\mathbf{S}} \mathcal{D}(\mathbf{Y}, \mathbf{AS}) = \mathbf{A}^{+} \mathbf{Y}$	<pre># least-square update</pre>
$\mathbf{S} \leftarrow \mathrm{prox}_{h_{\mathcal{S}}}(\mathbf{S})$	<pre># proximal update</pre>
(2) Update \mathbf{A} with \mathbf{S} fixed	
$\mathbf{A} \leftarrow \operatorname{argmin}_{\mathbf{A}} \mathcal{D}(\mathbf{Y}, \mathbf{AS}) = \mathbf{YS}^+$	<pre># least-square update</pre>
$\mathbf{A} \leftarrow \mathrm{prox}_{h_{\mathcal{A}}}(\mathbf{A})$	<pre># proximal update</pre>

Outputs: mixing matrix A, sources S

1.3.3.4 A particular pALS algorithm: Generalized morphological component analysis

This section presents the Generalized morphological component analysis (GMCA) framework [Bobin et al., 2007, Bobin et al., 2015, Kervazo, 2019]. Built upon the pALS scheme, GMCA proposes automatic parameter tuning strategies and heuristics, making it particularly reliable and easy to use in practice. Incidentally, the source separation methods developed in this thesis are based on GMCA.

Morphological diversity The principle of morphological diversity was first introduced in the context of Morphological component analysis (MCA) [Starck et al., 2005]. The sources were assumed to be sparse in different domains, allowing for their unmixing.

The concept of morphological diversity was extended for sources that are sparse in the same representation. Recall that due to the sparsity assumption, each source is encoded in a limited number of active samples when expressed in the transformed domain. This time, morphological diversity assumes that these few active samples are *diverse* from one source to another, *i.e.*, they are located in different positions or have sufficiently disjoint support. We note that real physical sources often verify this property (at least those from different physical processes).

The gist of GMCA is to use these samples to disentangle the mixture; as they originate from only one source, they are particularly discriminating for the unmixing.

Adaptive hyperparameter setting The source identification and unmixing is made possible thanks to the source thresholding induced by the ℓ_1 -penalization. The thresholding parameters in Λ play a central role for the robustness of the pALS minimization scheme:

- if they are too low:
 - the algorithm is burdened with uninformative samples for the unmixing (in general, the low amplitude samples are not discriminating as they share the same support between the sources),
 - the noise from the data is not sufficiently removed, which is also likely to interfere with the unmixing.
- if they are too high, too few samples are kept, which creates a lack of statistics and makes the separation problem more ill-posed.

The strength of GMCA lies in proposing an automatic thresholding strategy. For the sake of clarity, let us assume in this paragraph that the sources are sparse in the direct

domain (*i.e.*, $\mathbf{W} = \mathbf{I}$), noting that the principle is easily extendable to the general case. Consider that the GMCA algorithm has stabilized to a point $(\mathbf{A}^{(\infty)}, \mathbf{S}^{(\infty)})$; denoting $(\mathbf{A}^{\diamond}, \mathbf{S}^{\diamond})$ the ground-truth solution, the estimate of the sources after the least-square update reads:

$$\hat{\mathbf{S}} = \mathbf{A}^{(\infty)^{+}} \mathbf{Y}$$

$$= \mathbf{A}^{(\infty)^{+}} \mathbf{A}^{\diamond} \mathbf{S}^{\diamond} + \mathbf{A}^{(\infty)^{+}} \mathbf{N}$$

$$\approx \mathbf{S}^{\diamond} + \mathbf{A}^{\diamond +} \mathbf{N},$$
(1.18)

assuming that the algorithm has converged to the ground truth mixing matrix ($\mathbf{A}^{(\infty)} \approx \mathbf{A}^{\diamond}$). $\hat{\mathbf{S}}$ is therefore composed of a sparse term, \mathbf{S}^{\diamond} , and a Gaussian noise term, $\mathbf{A}^{\diamond+}\mathbf{N}$, whose noise level depends on the matrix row. In this regard, let $\{\sigma^{(i)}\}_{i\in[1...I]}$ denote the noise term's standard deviations per row.

The thresholding allows cutting the noise term while having a limited impact on the (sparse) term of interest \mathbf{S}^{\diamond} . A standard rule for the choice of the thresholding parameter in such context – denoising a sparse signal – is to fix it proportional to the noise level, leading to:

$$\mathbf{\Lambda}_{i:} = \kappa \, \sigma^{(i)} \, \mathbf{1},\tag{1.19}$$

with κ a hyperparameter. For example with $\kappa = 3$, a sample of $\hat{\mathbf{S}}$ whose amplitude is lower than $3\sigma^{(i)}$ has more than 99% chance of being noise.

The $\{\sigma^{(i)}\}_i$ can be estimated numerically using the Median absolute deviation (MAD) estimator. The MAD is a metric of the variability of a dataset; applied to a vector **n**, it is defined by:

$$MAD(\mathbf{n}) \coloneqq median(|\mathbf{n} - median(\mathbf{n})|).$$
(1.20)

If **n** is normally distributed with standard deviation σ , the MAD empirically relates to σ by $\sigma \approx 1.48 \text{ MAD}(\mathbf{n})$. The MAD is robust to sparse outliers, it can therefore be used to estimate the standard deviations $\{\sigma^{(i)}\}_i$ directly from $\hat{\mathbf{S}}$:

$$\mathrm{MAD}\left(\hat{\mathbf{S}}_{i:}\right) = \mathrm{MAD}\left(\mathbf{S}_{i:}^{*} + \left[\mathbf{A}^{\diamond +}\mathbf{N}\right]_{i:}\right) \approx \mathrm{MAD}\left(\left[\mathbf{A}^{\diamond +}\mathbf{N}\right]_{i:}\right) \approx \frac{1}{1.48} \,\sigma^{(i)}. \quad (1.21)$$

Yet, using Eq. (1.21) to derive the noise levels and ultimately the thresholds requires knowledge of $\hat{\mathbf{S}}$ and, therefore, the ground-truth mixing matrix \mathbf{A}^{\diamond} , which is precisely unknown. The solution proposed by GMCA to avoid this obstacle is to set the thresholds at each iteration with respect to the noise level in the current leastsquare estimate of the sources (*i.e.*, $\mathbf{A}^{(n-1)^+}\mathbf{Y}$ at iteration *n*). Altogether, this results in the following thresholds at iteration *n*:

$$\mathbf{\Lambda}_{i:}^{(n)} = \frac{\kappa}{1.48} \operatorname{MAD}\left(\left[\mathbf{A}^{(n-1)^{+}}\mathbf{Y}\right]_{i:}\right) \mathbf{1}.$$
(1.22)

Decrease of the thresholds The argument of the MAD in Eq. (1.22) writes $\mathbf{A}^{(n-1)^+}\mathbf{Y} = \mathbf{A}^{(n-1)^+}\mathbf{A}^{\diamond}\mathbf{S}^{\diamond} + \mathbf{A}^{(n-1)^+}\mathbf{N}$. During the first iterations, the mixing matrix estimate is likely far from the solution; thus, $\mathbf{A}^{(n-1)^+}\mathbf{A}^{\diamond} \neq \mathbf{I}$, such that $\mathbf{A}^{(n-1)^+}\mathbf{A}^{\diamond}\mathbf{S}^{\diamond}$ may not be sparse. Consequently, the MAD overestimates the noise levels $\{\sigma^{(i)}\}_i$, resulting in high thresholds. This phenomenon is actually beneficial, as it leads to keeping the highest amplitude samples only, which are precisely the most discriminating for determining the mixing directions of \mathbf{A} . From iteration to iteration, as the estimate of \mathbf{A} is refined, $\mathbf{A}^{(n-1)^+}\mathbf{A}^{\diamond}\mathbf{S}^{\diamond}$ becomes sparser, and the thresholds decrease.

In practice, this decay of the thresholds dramatically improves the robustness of the separation process with respect to the spurious local minima. Beyond the MAD, a specific threshold decreasing process is also implemented. At each iteration, the thresholds are further increased to keep a specific percentage K of the most significant amplitude samples which are above the thresholds determined by the noise level (Eq. 1.21). The percentage K is increased linearly at each iteration from 0 (or, more precisely, a very small percentage so as not to cut off the entire signal) to a user-defined K_{max} . This procedure provides an additional implicit regularization by selecting only the most meaningful samples. In this regard, K_{max} sets the intensity of the regularization; more specifically, the smaller the K_{max} , the slower the decrease of the thresholds and the greater the regularization. Assuming that the coefficients of $\mathbf{A}^{(n-1)^+}\mathbf{Y}$ are sorted in order of descending modulus, this procedure is implemented by:

$$\mathbf{\Lambda}_{i:}^{(n)} = \left| \left[\mathbf{A}^{(n-1)^{+}} \mathbf{Y} \right]_{ip_{0}} \right| \mathbf{1}, \qquad (1.23)$$
$$\mathbf{A}^{(n-1)^{+}} \mathbf{Y}_{in} \ge \frac{\kappa}{m} \operatorname{MAD}([\mathbf{A}^{(n-1)^{+}} \mathbf{Y}]_{in}) \}) \right|.$$

with $p_0 = \lfloor K \operatorname{card}(\{p, |[\mathbf{A}^{(n-1)^+}\mathbf{Y}]_{ip}| \ge \frac{\kappa}{1.48} \operatorname{MAD}([\mathbf{A}^{(n-1)^+}\mathbf{Y}]_{i:})\}) \rfloor.$

 ℓ_1 -reweighting The soft-thresholding induced by the ℓ_1 -penalization introduces a bias in the estimation of the sources. This phenomenon can be reduced by resorting to a ℓ_1 -reweighting scheme [Candès et al., 2008]. Basically, it amounts to replacing the ℓ_1 -penalty on the sources in the cost function by a log-sum penalty. In practice, the implementation of ℓ_1 -reweighting consists in deriving sample-wise thresholds based on the samples' value; the greater the expected sample's amplitude, the smaller the threshold and thus the smaller the bias. Denoting $\lambda_i^{(n)}$ the initial non-sample-dependent threshold of source *i* at iteration *n*, the ℓ_1 -reweighted threshold for source *i* at sample *p* is given by:

$$\Lambda_{ip}^{(n)} = \frac{\lambda_i^{(n)}}{1 + \frac{|\mathbf{s}_{ip}^{(n-1)}|}{\lambda_i^{(n)}}}.$$
(1.24)

Since the starting point is likely to be quite far from the sought-after sources, the ℓ_1 -reweighting scheme is only applied when the algorithm is stabilized, in order not

to favor spurious solutions.

Dealing with a multiresolution representations Multiresolution representations (*e.g.*, wavelets, curvelets, to only name two) are versatile and yet effective analytic signal representations that are well adapted to provide sparse representations for a wide range of natural data. As detailed in [Starck et al., 2010], such signal representations decompose the data into one coarse scale and several detail scales. The latter bears invaluable information to disentangle the sources while the former is a mere low frequency approximation of the data that is useless for the separation process. Therefore, the coarse scales of the data are generally removed for the estimation of the mixing matrix. In addition, the thresholding is not performed on the coarse scale, since it not sparse by nature.

The GMCA algorithm is summarized in Alg. 7 in the case of a multiresolution representation. The starting mixing matrix has a very limited impact on the finale solution, it can either be chosen randomly or from a PCA estimation. The choice of the stopping criterion is also of little importance, one can for example choose a relative variation on the estimated sources $(e.g., \|\mathbf{S}^{(i)} - \mathbf{S}^{(i-1)}\|_2 / \|\mathbf{S}^{(i)}\|_2 \leq \epsilon$, with $\epsilon = 1e - 6$) coupled with a maximum number of iterations in the unlikely event where the algorithm cycles $(e.g., n_{max} = 200)$.

Thanks to the algorithm's adaptive strategies, the choice of the few hyperparameters is relatively straightforward. The threshold parameter κ is generally set ranging from 1 to 3, depending on the sparsity of the expected sources ($\kappa = 3$ for very sparse sources, $\kappa = 1$ for approximately sparse sources). Concerning the maximum support of the sources, $K_{max} = 0.5$ is suitable in practice for a large class of signals. Lastly, the minimum number of iterations, which affects the decrease of the thresholds, must be large enough to avoid local minima; it can typically be set to $n_{min} = 100$.

1.4 Learning-based techniques

We identify two limitations of BSS algorithms when applied to the analysis of multispectral data in physical applications.

The first limitation is that BSS algorithms do not guarantee the interpretability of separation results, owing to their completely unsupervised nature. For example, this manifests in approximately separated sources, *i.e.*, contaminated by residuals in the form of leakage from other sources or in components that are not identified at all.

The second limitation concerns the adaptability of BSS algorithms. The regularizations employed, namely independence, non-negativity and sparsity, are not very flexible and difficult to adapt to the nature of the signals to unmix. However, a regularization that increases the contrast between the components is essential to promote Algorithm 7 Generalized morphological component analysis (GMCA) for sources which are sparse in a multiresolution representation

Inputs: data **Y**, number of sources *I*, multiresolution representation (with *S* detail scales of analysis operator $\mathbf{W}_{\mathbf{d}}$, one coarse scale of analysis operator $\mathbf{W}_{\mathbf{c}}$ & a synthesis operator **M**), thresholding parameters $\kappa \& K_{max}$, min. number of iterations n_{min}

 $n \leftarrow 0$ $\mathbf{Y_d} \leftarrow \mathbf{Y} \mathbf{W_d}^\top$ # detail scales of the data $\mathbf{A}^{(0)} \leftarrow \mathrm{PCA}(\mathbf{Y}, I)$ while convergence not reached do $n \leftarrow n+1$ (1) Update S with A fixed (1-i) Least-square update $\overline{\mathbf{S}^{(n)} \leftarrow \operatorname{argmin}_{\mathbf{S}} \mathcal{D}(\mathbf{Y}, \mathbf{A}^{(n-1)}\mathbf{S})} = \mathbf{A}^{(n-1)^{+}}\mathbf{Y}$ (1-ii) Proximal update (1-ii-a) Determination of the thresholds $\mathbf{S}_{\mathbf{d}}^{(n)} \leftarrow \mathbf{S}^{(n)} \mathbf{W}_{\mathbf{d}}^{\top}$ # detail scales of the sources Estimate the noise levels per source and per detail scale $\{\sigma^{(i,s)}\}_{i \in [1...I], s \in [1...S]}$ using the MAD $K^{(n)} \leftarrow K_{max} \min(n/n_{min}, 1)$ # % of significant samples kept Derive the thresholds $\{\lambda^{(i,s)}\}_{i\in[1,\dots,I],s\in[1,\dots,S]}$ w.r.t the noise levels so as to keep only $K^{(n)}\%$ of the most significant samples per scale, and write them in $\Lambda^{(n)}$ if $n > n_{min}$ then $\mathbf{\Lambda}^{(n)} \leftarrow \mathbf{\Lambda}^{(n)} \oslash (1 + |\mathbf{S}_{\mathbf{d}}^{(n-1)}| \oslash \mathbf{\Lambda}^{(n)})$ # perform ℓ_1 -reweighting (1-ii-b) Soft-thresholding $\mathbf{S_d}^{(n)} \leftarrow \mathcal{T}_{\mathbf{\Lambda}^{(n)}}(\mathbf{S_d}^{(n)})$ (1-ii-c) Non-negativity constraint $\mathbf{S}_{\mathbf{c}}^{(n)} \leftarrow \mathbf{S}^{(n)} \mathbf{W}_{\mathbf{c}}^{\top}$ # coarse scales of the sources $\mathbf{S}^{(n)} \leftarrow [\mathbf{S_d}^{(n)} \mathbf{S_c}^{(n)}] \mathbf{M}^\top$ # synthesis of the sources

(2) Update A with S fixed $\frac{(2\cdot i) \text{ Least-square update}}{\mathbf{S}_{\mathbf{d}}^{(n)} \leftarrow \mathbf{S}^{(n)} \mathbf{W}_{\mathbf{d}}^{\top}} \qquad \texttt{# detail scales of the sources}$ $\mathbf{A}^{(n)} \leftarrow \operatorname{argmin}_{\mathbf{A}} \mathcal{D}(\mathbf{Y}_{\mathbf{d}}, \mathbf{A} \mathbf{S}_{\mathbf{d}}^{(n)}) = \mathbf{Y}_{\mathbf{d}} \mathbf{S}_{\mathbf{d}}^{(n)^{+}}$ $\frac{(2\cdot ii) \operatorname{Proximal update}}{\mathbf{A}^{(n)} \leftarrow \operatorname{prox}_{h, 4}(\mathbf{A}^{(n)})}$

Outputs: mixing matrix A, sources S

 $\mathbf{S}^{(n)} \leftarrow \Pi_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}^{(n)})$

their separation.

Thanks to its fine modeling capabilities, machine learning has recently proven to be a particularly powerful technique for solving specific inverse problems (e.g., denoising, deconvolution) and has been the subject of numerous studies [McCann et al., 2017].

In this respect, learning-based approaches could address the limitations identified above in the context of BSS. However, little research has been conducted in this area to this day. As we have seen, BSS is a particularly challenging problem with interdependent variables; learning-based approaches are likely more difficult to value. In this section, methods that incorporate machine learning in BSS are reviewed.

End-to-end separation methods The first class of methods focuses on separating the sources directly from their mixture using neural networks (*e.g.*, with generative adversarial networks [Subakan and Smaragdis, 2018], with deep generative priors [Jayaram and Thickstun, 2020]). Most of the techniques proposed so far deal with the single-channel case, which is severely underdetermined and whose stakes differ substantially from the multichannel case we wish to address. Among the few multi-channel source separation methods, [Kameoka et al., 2019] introduces a multichannel variational autoencoder (VAE), which builds upon a set of conditional VAEs (one per expected source) to identify and retrieve the sources.

It is worth noting that autoencoders have also been used to perform nonlinear independent component analysis [Brakel and Bengio, 2017, Hyvärinen et al., 2019]. The gist is to learn a non-linear mapping (a neural network) towards a latent space with independent features. The encoder can then be used to separate mixtures composed of independent sources – the non-linear aspect comes from the neural network that the encoder constitutes. This is probably a promising approach to tackle non-linear BSS problems, for which methods remain scarce today (and that are out of the scope of this thesis).

The main drawback of these end-to-end methods is that they do not consider the mixture model nor necessarily the noise statistics, which is crucial in scientific applications. In addition, they only work with identified mixture classes that must be included in the learning stage.

Learning the priors to unmix Another way of incorporating machine learning in BSS is to combine a standard variational approach, *i.e.*, based on minimizing a cost-function, with a learned prior. This makes it possible to leverage the knowledge of the mixing model and the noise statistics, as well as the modeling capabilities of machine learning.

In this respect, research was primarily conducted in the specific framework of audio

source separation; for example in [Nugraha et al., 2016] the sources are iteratively refined by a neural-network-based denoiser within a standard unsupervised source separation scheme.

Unrolling blind source separation Deep unrolling or deep unfolding was initially proposed to accelerate classical iterative algorithms. In brief, it consists in replacing the updates of an iterative scheme (typically derived from a variational approach) with layers of a neural network [Monga et al., 2021]. For similar performance, the number of layers needed in the network is significantly reduced compared to the number of iterations of the iterative algorithm. Indeed, it can be thought that the updates are optimized during the network's learning. An advantage that gains importance concerns the network's interpretability, as its architecture stems from a variational framework.

Several algorithms have been successfully unrolled in the context of BSS, such as MU [Nasser et al., 2022] and PALM [Fahes et al., 2022].

1.5 Performance metrics for BSS

Facing a bivariate problem, two performance metrics will be used each time to evaluate the estimation quality of BSS algorithms, one for the mixing matrix and the other for the sources. The metrics will be expressed on an inverse logarithmic scale for the sake of precision (so that the higher the metrics, the more accurate the estimates).

1.5.1 Mixing matrix

Spectral angular distance The spectra that constitute the mixing matrix can be evaluated individually with the spectral angular distance (SAD). This metric stems from an angular interpretation, in which the estimated and ground truth spectra are considered as geometrical vectors. The SAD associated to the *i*th spectrum in a mixing matrix \mathbf{A} is defined by:

$$\operatorname{SAD}^{(i)} \coloneqq -10 \log_{10} \left(\operatorname{arccos} \left(\frac{\mathbf{A}_{:i}^{\top} \mathbf{A}^{\diamond}_{:i}}{\|\mathbf{A}_{:i}\|_{2} \|\mathbf{A}^{\diamond}_{:i}\|_{2}} \right) \right), \qquad (1.25)$$

with \mathbf{A}^{\diamond} the ground truth mixing matrix. We define the overall SAD as the geometric mean over all spectra in \mathbf{A} :

$$SAD \coloneqq -10 \log_{10} \left(\left(\prod_{i=1}^{I} \arccos \left(\frac{\mathbf{A}_{:i}^{\mathsf{T}} \mathbf{A}^{\diamond}_{:i}}{\|\mathbf{A}_{:i}\|_{2} \|\mathbf{A}^{\diamond}_{:i}\|_{2}} \right) \right)^{\frac{1}{I}} \right)$$

$$\coloneqq \frac{1}{I} \sum_{i=1}^{I} SAD^{(i)}$$

$$(1.26)$$

Mixing matrix criterion The mixing matrix criterion, denoted C_A , is a global criterion over all spectra that is given by:

$$C_{\mathbf{A}} \coloneqq -10 \log_{10} \left(\operatorname{mean} \left(|\mathbf{A}^{+}\mathbf{A}^{\diamond} - \mathbf{I}| \right) \right)$$
$$\coloneqq -10 \log_{10} \left(\frac{1}{I^{2}} \sum_{i_{1}=1}^{I} \sum_{i_{2}=1}^{J} |\mathbf{A}^{+}\mathbf{A}^{\diamond} - \mathbf{I}|_{i_{1}i_{2}} \right).$$
(1.27)

The SAD has the advantage of quantifying the estimation quality of each spectrum, which may be crucial in some cases (*e.g.*, to assess the impact of disparately emitting components further), but the derived global criterion is somewhat ad hoc. The C_A is, on the contrary, a global metric; based on a pseudo-inverse calculation, it is very sensitive to poorly separated spectra, which therefore correlate.

1.5.2 Sources

Signal-to-distortion ratio Estimation errors in the source domain can have several origins that may be of interest to quantify. In this regard, [Vincent et al., 2006] proposes to decompose the *i*th estimate of a source $\mathbf{S}_{i:}$ as follows:

$$\mathbf{S}_{i:} = \mathbf{s}_{\mathbf{t}}^{(i)} + \mathbf{e}_{\mathbf{i}}^{(i)} + \mathbf{e}_{\mathbf{n}}^{(i)} + \mathbf{e}_{\mathbf{a}}^{(i)}, \qquad (1.28)$$

where $\mathbf{s_t}^{(i)}$ is the term that actually corresponds to the sought-after source, and $\mathbf{e_i}^{(i)}$, $\mathbf{e_n}^{(i)}$ and $\mathbf{e_a}^{(i)}$ are error terms originating respectively from the other sources, the noise, and from other causes (*e.g.*, thresholding artifacts). More specifically, denoting $P_{(\cdot)}$ the orthogonal projector on the subspace generated by the elements given in subscript, each term is given by:

- $\mathbf{s}_{\mathbf{t}}^{(i)} \coloneqq P_{\mathbf{S}_{i:}^*}(\mathbf{S}_{i:}),$
- $\mathbf{e_i}^{(i)} \coloneqq P_{\mathbf{S}^\diamond}(\mathbf{S}_{i:}) \mathbf{s_t}^{(i)},$

- $\mathbf{e}_{\mathbf{n}}^{(i)} \coloneqq P_{\mathbf{S}^\diamond, \mathbf{N}}(\mathbf{S}_{i:}) \mathbf{e}_{\mathbf{i}}^{(i)},$
- $\mathbf{e}_{\mathbf{a}}^{(i)} \coloneqq \mathbf{S}_{i:} P_{\mathbf{S}^\diamond, \mathbf{N}}(\mathbf{S}_{i:}).$

Four performance metrics are derived from this decomposition, namely the signal-todistortion ratio (SDR), that accounts for all contaminations in $\mathbf{S}_{i:}$

$$SDR^{(i)} \coloneqq -10 \log_{10} \left(\frac{\|\mathbf{e_i}^{(i)} + \mathbf{e_n}^{(i)} + \mathbf{e_a}^{(i)}\|_2^2}{\|\mathbf{s_t}^{(i)}\|_2^2} \right).$$
(1.29)

the signal-to-interference ratio (SIR), which assesses in particular the contaminations coming from the other sources:

$$\operatorname{SIR}^{(i)} \coloneqq -10 \log_{10} \left(\frac{\|\mathbf{e_i}^{(i)}\|_2^2}{\|\mathbf{s_t}^{(i)}\|_2^2} \right), \qquad (1.30)$$

the signal-to-noise ratio (SNR), that evaluates specifically noise leakages:

$$\mathrm{SNR}^{(i)} \coloneqq -10 \log_{10} \left(\frac{\|\mathbf{e_n}^{(i)}\|_2^2}{\|\mathbf{s_t}^{(i)} + \mathbf{e_i}^{(i)}\|_2^2} \right), \tag{1.31}$$

and the signal-to-artifact ratio (SAR), which accounts for the remaining contaminations:

$$SAR^{(i)} \coloneqq -10 \log_{10} \left(\frac{\|\mathbf{e}_{\mathbf{a}}^{(i)}\|_{2}^{2}}{\|\mathbf{s}_{\mathbf{t}}^{(i)} + \mathbf{e}_{\mathbf{i}}^{(i)} + \mathbf{e}_{\mathbf{n}}^{(i)}\|_{2}^{2}} \right).$$
(1.32)

Similarly to the SAD, we define a global SDR, SIR, SNR and SAR for all sources, using geometric means.

Normalized mean square error While being particularly informative, the source metrics defined above are not adapted for mixture models that include instrumental responses. In that case, we will resort to a more basic normalized mean square error (NMSE). We define both a source-dependent NMSE:

NMSE⁽ⁱ⁾ = -10 log₁₀
$$\left(\frac{\|\mathbf{S}_{i:} - \mathbf{S}_{i:}^*\|_2^2}{\|\mathbf{S}_{i:}^{\diamond}\|_2^2} \right)$$
, (1.33)

as well as an overall NMSE, which is the geometric mean of the source-dependant NMSEs.

When working with heterogeneous data in terms of resolution, it is common in physical applications to reduce the resolution of the formed images to that of the lowresolution channel to attenuate potential deconvolution artifacts. Consider a low-pass kernel \mathbf{h} that brings the resolution to a user-defined one; the NMSE is adapted as follows:

NMSE⁽ⁱ⁾ = -10 log₁₀
$$\left(\frac{\|\mathbf{h} * (\mathbf{S}_{i:} - \mathbf{S}_{i:}^*) \|_2^2}{\|\mathbf{h} * \mathbf{S}_{i:}^{\diamond} \|_2^2} \right),$$
 (1.34)

and likewise for the global metric.

Chapter 2

Introducing learned constraints in blind source separation

2.1 Introduction

When applied to real data, one of the challenges of BSS is to provide physically interpretable separation results. In particular, it is of paramount importance that leakages between components remain limited, whether in the source or spectral domain.

The multispectral datasets which are analyzed often concern processes that are at least partially known, and in this respect, resorting to totally unsupervised approaches may not be the most adequate. Therefore, it makes sense to develop semi-supervised separation methods that account for prior knowledge of sought-after components; in this chapter, we propose a way of introducing learning-based priors in source separation.

First, we discuss how to include learned priors that efficiently discriminate between components in BSS. We then introduce the interpolatory autoencoder (IAE), which is a learning-based approach that allows building regularizations for unmixing problems, including BSS. This leads us afterward to propose a *semi-blind* source separation framework that associates the SBSS variational approach with priors learned by the IAE. The resulting algorithm, named semi-blind GMCA (sGMCA), is finally tested on realistic astrophysical datasets. It is shown to reject leakages efficiently, allowing for improved separation.

2.2 Towards a semi-blind source separation framework

2.2.1 Constraining the spectra

In BSS, the role of regularization is to gain contrast between components to allow their unmixing. So far, we have considered the sparsity of the sources in a transformed domain (and the non-negativity of the sources and the mixing matrix, but this constraint is not particularly restrictive). However, we identify two limitations with the aim of tackling astrophysical applications.

First, sources may correlate spatially; the morphological diversity assumption being not strictly respected, SBSS algorithms struggle to separate them clearly [Picquenot et al., 2019]. Second, the instruments may have a low spatial resolution, either natively (*e.g.*, Athena, which will have an excellent spectral resolution but will produce images of few pixels) or due to measurement settings (*e.g.*, the interferometers in low radio frequency bands). SBSS algorithms can no longer rely solely on the sparsity of the sources, which requires high spatial frequency information. A solution could be to learn new representations for the sources, but this seems complicated as they generally present a large variability.

Therefore, we resort to constraining the spectra that constitute the mixing matrix. Indeed, contrary to sources, spectra in multispectral data generally present a lower variability and exhibit strong structures (see examples of realistic spectra in Fig. 2.1), which makes them particularly suitable for their modeling with machine learning.

To do so, we choose to employ autoencoders to constrain the sought-after spectra to belong to learned, low-dimensional manifolds. The resulting regularization is expected to reject significantly leakages, allowing for an improved disentanglement between the components. In addition, enforcing spectra to belong to learned and physically interpretable manifolds is of prime interest with a view to tackling physical applications.

2.2.2 Machine learning and blind source separation

End-to-end methods The first way to integrate machine learning into BSS is with end-to-end separation, which directly estimates the sources and mixing process from the observed mixtures [Kameoka et al., 2019, Jayaram and Thickstun, 2020]. Nonetheless, such methods are not well adapted to account for the exact mixture model and the noise statistics, which is key in scientific applications. Moreover, in the context of BSS, all possible combinations of **A** and **S** must be considered during learning, which can quickly become voluminous.



Figure 2.1: Ensemble of three emission models, illustrating the low variability of spectra.

Consequently, hybrid techniques that combine standard variational approaches [Scherzer et al., 2008], *i.e.*, based on the minimization of a cost-function, and learned priors are better suited. In this framework, designing an effective BSS algorithm with interpretable solutions requires building regularizations for the mixing matrix and/or the sources that precisely account for the properties of the sought-after signals.

Associating a variational approach with a learning-based regularization In the literature, combining variational approaches for inverse problems with regularization-learning procedures comes in different flavors. In [Adler and Oktem, 2017, Adler and Öktem, 2018, Gilton et al., 2020, Yang et al., 2016, the proposed learning architecture builds upon a residual network to mimic standard proximal algorithms, in which regularization learning is one element of the learning scheme. So far, it is unclear whether this approach is well-suited for multiconvex problems (i.e.,convex according to each variable when the other variables are fixed) such as BSS. A different way of solving inverse problems with learned regularizations consists in first inverting the observation operator (e.q.), the mixing matrix, when the sources alone are estimated) with a fast and simple procedure (e.q., least-square solution)and then clean inversion artifacts thanks to a learned denoiser. This technique has been investigated both with standard convolutional networks denoisers [Hammernik et al., 2018, Jin et al., 2017, Romano et al., 2017, Sureau et al., 2020] or with generative models with adversarial training [Lunz et al., 2018]. The underlying denoiser depends on the linear operator to be inverted, and in the case of BSS, this means that it should depend on the sources and the mixing matrix that are precisely the variables to be estimated.

Prior work in BSS Incorporating learning-based priors in classical BSS methods is not novel. Research was mostly conducted in the specific framework of audio source separation. A neural-network-based denoiser has been used in [Nugraha et al., 2016] to refine the source estimate within a standard source separation scheme. Generative models built upon variational autoencoders have also been exploited in the context of single channel source separation [Narayanaswamy et al., 2020]. However, these methods are not easily transposable for the multichannel BSS problem that we propose to tackle. Moreover, as said above, regularizing the sources might not be constraining enough to efficiently discriminate between the components and limit leakages, especially in the presence of unknown morphologies or statistical distributions.

2.3 The interpolatory autoencoder

In this subsection, we present the interpolatory autoencoder (IAE), which is an approach that we first proposed to build learning-based regularizations for unmixing problems, and that we will use thereafter for semi-blind source separation.

2.3.1 Principle

2.3.1.1 Learning to travel on a manifold

From a general viewpoint, the manifold hypothesis [Fefferman et al., 2013] states that natural signals, such as spectra in the context of BSS, lie on low-dimensional manifolds embedded in larger spaces. In contrast to learning the manifold structure straight from the data, an interesting approach proposed in [Culpepper and Olshausen, 2009] consists in estimating how to travel on it. The gist of this method is to learn a transport operator from pairs of samples. The approach has recently been combined with a variational autoencoder architecture to build a generative model allowing to learn a latent space with a topology of increased interpretability [Connor and Rozell, 2020].

Inspired by this seminal work, we instead build upon the availability of examples, whether derived from simulations or measurements, which can be exploited to learn an efficient generative model from a reduced number of training samples. In this context, the first ingredient of the proposed method is to build some regularization by *exploring the manifold from these examples*, which we call "anchor points", rather than trying to capture the overall manifold structure. This will be conveniently carried out by building new samples on a manifold as an interpolation or barycenter of the anchor points.

2.3.1.2 Preliminary: affine hulls

We provide a brief mathematical background on affine spaces and, more particularly, affine hulls, which allows us to define barycenters and barycentric projections.

The affine hull $\operatorname{aff}(\mathcal{K})$ of a set $\mathcal{K} \subset \mathbb{R}^J$ is the set of all affine combinations, also referred to as barycentric combinations, of elements of \mathcal{K} , *i.e.*:

aff
$$(\mathcal{K}) \coloneqq \left\{ \sum_{n=1}^{N} \lambda^{(n)} \boldsymbol{\alpha}^{(n)}, N > 0, \boldsymbol{\alpha}^{(n)} \in \mathcal{K}, \lambda^{(n)} \in \mathbb{R}, \sum_{n=1}^{N} \lambda^{(n)} = 1 \right\}.$$
 (2.1)

Illustrations of affine hulls are shown in Fig. 2.2. Thereafter, the elements of \mathcal{K} will be considered affinely independent, such that each element of $\operatorname{aff}(\mathcal{K})$ is uniquely decomposed by affine combination of elements of \mathcal{K} ; in other words, \mathcal{K} constitutes a basis of $\operatorname{aff}(\mathcal{K})$. Assuming that $\operatorname{card}(\mathcal{K}) = N \geq 2$, let us form the matrix $\Phi := [\boldsymbol{\alpha}^{(1)} \dots \boldsymbol{\alpha}^{(N)}] \in \mathbb{R}^{N \times J}$ constituted by the elements of \mathcal{K} . The definition of the affine hull with the matrix-vector framework reads as follows:

aff
$$(\mathcal{K}) = \{ \Phi \lambda, \ \lambda \in \mathcal{S} \},$$
 (2.2)

where $\mathcal{S} := \{ \boldsymbol{\lambda} \in \mathbb{R}^N, \mathbf{1}^\top \boldsymbol{\lambda} = 1 \}$. The elements of $\boldsymbol{\lambda}$ are called barycentric coordinates.



Figure 2.2: Illustrations of affine hulls and barycentric projections. (a) $\Pi_{\text{aff}(\mathcal{K})}(\mathbf{x}) = -0.87 \boldsymbol{\alpha}^{(1)} + 1.87 \boldsymbol{\alpha}^{(2)}$, thus the barycentric coordinates of the projection are $\boldsymbol{\lambda} = (-0.87, 1.87)^{\top}$. (b) $\Pi_{\text{aff}(\mathcal{K})}(\mathbf{x}) = 1.10 \boldsymbol{\alpha}^{(1)} - 0.47 \boldsymbol{\alpha}^{(2)} + 0.38 \boldsymbol{\alpha}^{(3)}$, thus the barycentric coordinates of the projection are $\boldsymbol{\lambda} = (1.10, -0.47, 0.38)^{\top}$.

Affine hulls being convex sets, it is possible to define orthogonal projections onto them. They are given by:

$$\Pi_{\mathrm{aff}(\mathcal{K})}(\mathbf{x}) = \Phi \boldsymbol{\lambda} \quad \text{s.t.} \quad \boldsymbol{\lambda} = \operatorname*{argmin}_{\boldsymbol{\lambda} \in \mathcal{S}} \frac{1}{2} \| \boldsymbol{\Phi} \boldsymbol{\lambda} - \mathbf{x} \|^{2}$$
$$= \Phi^{+} \mathbf{x} + \frac{1 - \mathbf{1}^{\top} \Phi^{+} \mathbf{x}}{\mathbf{1}^{\top} (\Phi^{\top} \Phi)^{-1} \mathbf{1}} \left(\Phi^{\top} \Phi \right)^{-1} \mathbf{1}.$$
(2.3)

Hereafter, the projection onto an affine hull will be called barycentric projection, as it amounts to finding the closest barycenter to the input vector.

Alternatively, the barycentric coordinates in the projection can be approximated by:

$$\boldsymbol{\lambda} \approx \Pi_{\mathcal{S}} \left(\boldsymbol{\Phi}^{\top} \mathbf{x} \right), \tag{2.4}$$

where the projection on \mathcal{S} is further approached by a rescaling $(\Pi_{\mathcal{S}}(\boldsymbol{\lambda}) \approx \boldsymbol{\lambda}/(\mathbf{1}^{\top}\boldsymbol{\lambda}))$.

2.3.1.3 Architecture of the interpolatory autoencoder

Let $\mathcal{A} = {\{\boldsymbol{\alpha}^{(n)}\}}_{n \in [1...N]} \subset \mathcal{M}$ be a set of N anchor points of a manifold \mathcal{M} that one seeks to model. In order to learn how to travel on \mathcal{M} by computing barycenters of the anchor points, a neural-network-based autoencoder is introduced; let $\phi : \mathbb{R}^J \to \mathbb{R}^{J'}$ and $\psi : \mathbb{R}^{J'} \to \mathbb{R}^J$ be the forward encoder and the backward decoder, respectively.

The architecture of the IAE is represented in Figure 2.3. First, ϕ encodes the input sample and the anchor points in the latent space, then the code of the input sample is projected on the affine hull of the codes of the chosen anchor points, and finally, ψ decodes the sample in the direct space. The encoder and decoder are optimized so to minimize the reconstruction error of a given training dataset $\mathcal{T} \subset \mathcal{M}$, that is:

$$\underset{\phi,\psi}{\operatorname{argmin}} \sum_{\mathbf{t}\in\mathcal{T}} \left\| \mathbf{t} - \psi \circ \Pi_{\operatorname{aff}(\phi(\mathcal{A}))} \circ \phi(\mathbf{t}) \right\|_{2}^{2},$$
(2.5)

where $\phi(\mathcal{A})$ is a shortcut for $\{\phi(\boldsymbol{\alpha}^{(n)})\}_n$. In doing so, the manifold tends to be linearized in the latent space (see Fig. 2.4).

Once the autoencoder is learned, the manifold can be approximated as the decoding in the direct space of the affine hull of the encoded anchor points, that is:

$$\mathcal{M} \approx \left\{ \mathbf{x} \in \mathbb{R}^{J}, \exists \boldsymbol{\lambda} \in \mathcal{S}, \mathbf{x} = \psi \left(\boldsymbol{\Phi} \boldsymbol{\lambda} \right) \right\},$$
(2.6)

where now $\mathbf{\Phi} \coloneqq [\phi(\boldsymbol{\alpha}^{(1)}) \dots \phi(\boldsymbol{\alpha}^{(N)})] \in \mathbb{R}^{N \times J'}$ to match the IAE's architecture.

2.3.2 Application

The IAE can be used in unmixing problems to constrain the sought-after signals to belong to learned manifolds. To do so, two projection operators on manifold are defined.

Fast interpolation A first projection, named fast interpolation, is defined as a forward pass in the autoencoder:

$$\check{\Pi}_{\mathcal{M}}(\mathbf{x}) \coloneqq \psi \circ \Pi_{\mathrm{aff}(\phi(\mathcal{A}))} \circ \phi(\mathbf{x}) \,. \tag{2.7}$$



Figure 2.3: Diagram of the interpolatory autoencoder (IAE). It is recalled that λ is the vector with the barycentric coordinates associated to \mathbf{x} , and $\mathbf{\Phi} = [\phi(\boldsymbol{\alpha}^{(1)}) \dots \phi(\boldsymbol{\alpha}^{(N)})]$.



Figure 2.4: Illustration of the IAE data modeling and the barycentric span projection. In this example, the manifold has one dimension; two anchor points are therefore required.

The corresponding barycentric coordinates are $\check{\boldsymbol{\lambda}} \coloneqq \operatorname{argmin}_{\boldsymbol{\lambda} \in \mathcal{S}} \frac{1}{2} \| \boldsymbol{\Phi} \boldsymbol{\lambda} - \boldsymbol{\phi}(\mathbf{x}) \|^2$, whose solution is given in Eq. (2.3). The advantage of this projection is that it is fast, as it is based on analytical formulations only. However, it does not guarantee the accuracy of the reconstruction, as the encoder $\boldsymbol{\phi}$ is not aimed to be robust to contaminations, either from measurement noise or external interference. This is why a second, more precise projection is implemented.

Generative modeling Before this, we note that the IAE allows us to derive a generative model. It is obtained by exploring the latent domain:

$$g(\boldsymbol{\lambda}) \coloneqq \psi(\boldsymbol{\Phi}\boldsymbol{\lambda}), \ \boldsymbol{\lambda} \in \mathcal{S},$$
 (2.8)

which amounts to decoding barycenters of the encoded anchor points. Unlike classical generative models, for example derived from variational autoencoders, the proposed generative model has no statistical basis, but it is not needed for the considered application.

Barycentric span projection The second projection on manifold, called barycentric span projection, is defined by:

$$\Pi_{\mathcal{M}}(\mathbf{x}) \coloneqq g(\boldsymbol{\lambda}) \quad \text{s.t.} \quad \boldsymbol{\lambda} = \operatorname*{argmin}_{\boldsymbol{\lambda} \in \mathcal{S}} \frac{1}{2} \|g(\boldsymbol{\lambda}) - \mathbf{x}\|_{2}^{2}.$$
(2.9)

The projection boils down to finding the latent parameters which minimize the Euclidean distance between the input \mathbf{x} and the image space of the generative model, *i.e.*, the decoding of the affine hull of the encoded anchor points. The optimization problem is a constrained non-linear least-square problem, which does not admit a closed-form solution. It can nonetheless be estimated by a gradient-descent-based algorithm (*e.g.*, the Adam optimizer) since g is differentiable; the sum-to-one constraint can be addressed by defining an element of $\boldsymbol{\lambda}$ as one minus the sum of the other elements.

In physical applications, \mathcal{M} is generally not convex, and thus neither is the projection. Therefore, the choice of the initialization of the descent algorithm is crucial, as the quality of the solution depends on it. The barycentric coordinates $\check{\lambda}$ of the fast interpolation seen above are generally accurate; therefore, when they are used to initialize the descent algorithm of the barycentric span projection, the optimization tends to be locally convex. This is particularly appealing when integrating the barycentric span projection as a regularization scheme for unmixing problems.

2.4 The sGMCA algorithm

We now return to the BSS problem. As we concluded, we will keep the variational approach but constrain the spectra of the mixing matrix to belong to learned manifolds;

for this purpose, we will use the IAE.

2.4.1 Principle

The proposed source separation algorithm is based on a the SBSS framework. In accordance with Chapter 1, the considered mixture model is $\mathbf{Y} = \mathbf{AS} + \mathbf{N}$, and the cost function writes:

$$\underset{\mathbf{A},\mathbf{S}}{\operatorname{argmin}} \mathcal{D}\left(\mathbf{Y},\mathbf{AS}\right) + h_{\mathcal{S}}\left(\mathbf{S}\right) + h_{\mathcal{A}}\left(\mathbf{A}\right), \qquad (2.10)$$

where, in accordance with Chapter 1, $\mathcal{D}(\mathbf{Y}, \mathbf{AS}) = \|\mathbf{Y} - \mathbf{AS}\|_2^2/2$ and $h_{\mathcal{S}}(\mathbf{S}) = \|\mathbf{\Lambda} \odot (\mathbf{SW}^{\top})\|_1 + \iota_{O_S}(\mathbf{S})$ with $\mathbf{\Lambda}$ the sparse hyperparameters, \mathbf{W} is a sparsifying transform and $\mathcal{O}_{\mathcal{S}}$ is the non-negative orthant for source matrices. The difference concerns the regularization term on the mixing matrix $h_{\mathcal{A}}(\cdot)$ that accounts for the IAE modeling of the spectra, as discussed below.

Application of the generative models A first approach would amount to optimizing the spectra directly in the latent space of the IAE models, leading to the following cost function:

$$\min_{\mathbf{S}, \left\{\boldsymbol{\lambda}^{(i)}\right\}_{i \in [1...I]}} \frac{1}{2} \left\| \mathbf{Y} - \mathbf{A}\mathbf{S} \right\|_{2}^{2} + \left\| \mathbf{\Lambda} \odot \left(\mathbf{S}\mathbf{W}^{\top}\right) \right\|_{1} \text{ s.t. } \forall i \in [1...I], \ \mathbf{A}_{:i} = g^{(m_{i})} \left(\boldsymbol{\lambda}^{(i)}\right),$$
(2.11)

where m_i is the index of the IAE spectrum model which is associated to component i, and thus $g^{(m_i)}(\cdot)$ is the spectrum generative model associated to IAE model m_i . However, the optimization according to the $\{\boldsymbol{\lambda}^{(i)}\}_{i\in[1...I]}$ is problematic in two respects. First, the sources intervene in the application of the models, which may pose issues of convergence as the impact of an estimation bias of \mathbf{S} on the $\{\boldsymbol{\lambda}^{(i)}\}_{i\in[1...I]}$ is not necessarily well controlled. Second, the computational cost is high because the optimization requires inverting the sources.

Rather, we take the IAE models into consideration by constraining the spectra to belong to the image spaces of the IAE models (which are supposed to approximate the underlying spectrum manifolds $\{\mathcal{M}^{(m_i)}\}_{i\in[1...I]}$). To that end, the following constraint term on the mixing matrix can be considered: $\sum_{i\in[1...I]} \iota_{\mathcal{M}^{(m_i)}}(\mathbf{A}_{:i})$. This makes it possible to decouple the application of the models from the sources, which facilitates the problem resolution.

The semi-blind case Until now, we supposed that all the spectra of the mixing matrix were known and modeled. However, in the more general *semi-blind* approach, we suppose that among the I elementary components, M have a spectrum modeled

by an IAE model and I - M are fully unknown. Let $\mathbb{I} \subset [1 \dots I]$ be the indices of the modeled components.

If a spectrum is not modeled (*i.e.*, $i \notin \mathbb{I}$), it is constrained to belong to $\mathcal{K}_a \coloneqq \{\mathbf{a} \in \mathbb{R}^J, \|\mathbf{a}\|_2 \leq 1, \forall j, \mathbf{a}_j \geq 0\}$, as above in SBSS (see Section 1.3.3). This constraint is reminded to ensure the spectra' non-negativity and prevent the scale degeneracy inherent to SBSS.

The constraint term on the mixing matrix finally reads as:

$$h_{\mathcal{A}}(\mathbf{A}) \coloneqq \sum_{i \in \mathbb{I}} \iota_{\mathcal{M}^{(m_i)}}(\mathbf{A}_{:i}) + \sum_{i \notin \mathbb{I}} \iota_{\mathcal{K}_a}(\mathbf{A}_{:i}).$$
(2.12)

To summarize, sGMCA seeks the solution to the following problem:

$$\min_{\mathbf{A},\mathbf{S}} \frac{1}{2} \|\mathbf{Y} - \mathbf{A}\mathbf{S}\|_{2}^{2} + \|\mathbf{\Lambda} \odot (\mathbf{S}\mathbf{W}^{\top})\|_{1} + \iota_{O_{S}}(\mathbf{S}) + \sum_{i \in \mathbb{I}} \iota_{\mathcal{M}^{(m_{i})}}(\mathbf{A}_{:i}) + \sum_{i \notin \mathbb{I}} \iota_{\mathcal{K}_{a}}(\mathbf{A}_{:i}). \quad (2.13)$$

2.4.2 Minimization scheme

The sGMCA method is described in Algorithm 8; it is based on GMCA, which, as we have seen, offers a flexible framework and provides robust and adaptive heuristics. The sources and the mixing matrix are initialized with GMCA. At this point the solution is approximate; the mixing matrix and the sources are likely to be contaminated by residuals of other components (referred to as "leakages"). **S** and **A** are updated alternatively and iteratively until convergence is reached. According to the projected alternating least-square scheme, each update comprises a least-square estimate, so as to minimize the data-fidelity term \mathcal{D} , followed by the application of the proximal operator of the corresponding regularization term. The procedure is stopped when either the estimated sources have stabilized $(\|\mathbf{S}^{(k)} - \mathbf{S}^{(k-1)}\|_2 / \|\mathbf{S}^{(k)}\|_2 \le \epsilon \text{ with } \epsilon = 10^{-6}$ in practice) or a maximal number of iterations have been reached (50 in practice).

Equation (2.13) is non-convex. It is nevertheless convex with respect to **S** when **A** is fixed. This is however not the case for **A** when **S** is fixed, since the manifolds $\{\mathcal{M}^{(m_i)}\}_{i\in\mathbb{I}}$ are unlikely to be convex. Fortunately, if the first guess obtained by the least-square estimate is decent, and if the IAE generative models are accurate enough, the projections tend to be locally convex (see results in Section 2.5.1). Being a non-convex problem, convergence to a critical point can at best be guaranteed. As far as we know, it is not proven that pALS converges, and in this respect, neither does sGMCA. However, in all tests performed, the algorithm always stabilized. We note that the stability is sometimes not strict, as minor cycles between the update of **S** and **A** can appear. This is most likely due to the convexity assumption of the spectrum update, which remains approximate.

Algorithm 8 sGMCA

Inputs: data **Y**, number of sources I, set of M (already trained) IAE generative models of spectra, sparsifying representation (analysis operator **W** and synthesis operator **M**), thresholding parameters κ and K_{max}

 $\mathbf{A}, \mathbf{S} \leftarrow \text{GMCA}(\mathbf{Y}, I, \mathbf{W}, \mathbf{M}, \kappa, K_{max})$ while convergence not reached do

(1) Update S with A fixed
(1-i) Least squares: S ← A⁺Y
(1-ii) Determination of the thresholding parameters Λ
(1-iii) Regularization: S ← max (T_Λ (SW^T) M^T, 0)
(2) Update A with S fixed
(2-i) Least squares: A ← YS⁺
(2-ii) Spectrum identification & model-to-spectrum mapping: determine I & {m_i}_{i∈I} (see Alg. 9)
(2-iii) Constraint application:
for i ∉ I do
(2-iii-a) Projection on K_a: A_{:i} ← Π_{Ka} (A_{:i})
for i ∈ I do
(2-iii-b) Manifold projection: A_{:i} ← Π_{M^(mi)} (A_{:i})

Outputs: mixing matrix A, sources S

2.4.3 Source update

As seen previously, following the pALS minimization scheme, the source update reads as:

$$\mathbf{S} \leftarrow \operatorname{prox}_{h_{\mathcal{S}}}(\mathbf{A}^{+}\mathbf{Y}) \approx \max\left(\mathcal{T}_{\mathbf{\Lambda}}\left(\left(\mathbf{A}^{+}\mathbf{Y}\right)\mathbf{W}^{\top}\right)\mathbf{M}^{\top}, \mathbf{0}\right),$$
 (2.14)

with $\mathcal{T}_{\Lambda}(\cdot)$ the soft-thresholding operator and **M** an analysis operator such that $\mathbf{MW} = \mathbf{I}$. The thresholds Λ are set automatically according to the strategies implemented in GMCA (*i.e.*, the thresholds are set with respect to the estimated noise level according to a " κ - σ " rule and then adapted with a ℓ_1 -reweighting scheme, see Section 1.3.3.4).

2.4.4 Mixing matrix update

2.4.4.1 Adaptation of the barycentric span projection for spectra

The projection of a spectrum $\mathbf{a} \in \mathbb{R}^J$ on a manifold \mathcal{M} parameterized by an IAE model is adapted as follows:

$$\Pi_{\mathcal{M}}(\mathbf{a}) \coloneqq g(\boldsymbol{\lambda}) \quad \text{s.t.} \quad \boldsymbol{\lambda} = \operatorname*{argmin}_{\boldsymbol{\lambda} \in \mathcal{S}, \, \rho \in \mathbb{R}^+} \frac{1}{2} \left\| \rho \, g(\boldsymbol{\lambda}) - \mathbf{a} \right\|_2^2.$$
(2.15)

Compared to the earlier defined barycentric span projection (Eq. 2.9), a scale coefficient ρ has been added to account for the scale indeterminacy on **a**, which is inherent to BSS.

Regarding the choice of the starting points of the underlying descent algorithm, the choice of the barycentric coordinate vector is unchanged $(\check{\boldsymbol{\lambda}})$ and ρ is initialized so as to minimize $\|\mathbf{a} - \rho \,\check{\Pi}_{\mathcal{M}}(\mathbf{a})\|_2$, yielding $\check{\rho} \coloneqq \mathbf{a}^{\top} \check{\Pi}_{\mathcal{M}}(\mathbf{a}) / / \|\check{\Pi}_{\mathcal{M}}(\mathbf{a})\|_2^2$.

2.4.4.2 Spectrum identification and model-to-spectrum mapping

The least-square estimate, which constitutes the first step of the mixing matrix update, is equal to $\bar{\mathbf{A}} := \mathbf{YS}^+$. Before applying the mixing matrix constraint, the following questions arise: (i) among the I spectra in $\bar{\mathbf{A}}$, which can be modeled by the provided generative models, and (ii) more specifically, which model to associate with which identified spectrum? In other words, it raises the question of determining I and $\{m_i\}_{i\in\mathbb{I}}$ of Eq. (2.13).

In physical applications where only one set of data is being worked on, identification can usually be made by hand (on the first iteration only, as the order does not change in the subsequent iterations). Otherwise, we propose the procedure for identifying spectra in Algorithm 9. The gist is to iteratively associate the spectrummodel pair (m, i) that minimizes the Euclidean distance between the *interference-free* spectra and their projections on manifolds, that is:

$$\epsilon^{(m,i)} \coloneqq \left\| \bar{\mathbf{A}}_{:i} - \mathbf{M}\boldsymbol{\mu} - \check{\rho}^{(m,i)} \,\check{\Pi}_{\mathcal{M}^{(m)}} \big(\bar{\mathbf{A}}_{:i} - \mathbf{M}\boldsymbol{\mu} \big) \right\|_{2}.$$
(2.16)

Indeed, for robustness purposes, it is essential to account for the interferences that are likely to contaminate the input spectra in $\bar{\mathbf{A}}$. This is done by subtracting the identified spectra up to the current iteration in \mathbf{M} , weighted by coefficients $\boldsymbol{\mu}$. These can be determined with a coarse grid-search, which is sufficient for the sake of identification. Likewise, the fast projection is employed as its accuracy is enough to identify spectra.

Algorithm 9 Spectrum identification and model-to-spectrum mapping

Input: mixing matrix $\bar{\mathbf{A}}$, set of M IAE generative models of spectra

Initialize empty matrix $\mathbf{M} \in \mathbb{R}^{J \times 0}$ $\mathbb{I} \leftarrow [], \mathbb{I}^C \leftarrow [1 \dots I], \mathbb{M} \leftarrow [1 \dots M]$

for k in $1 \dots M$ do

 $i, m, \boldsymbol{\mu} \leftarrow \operatorname*{argmin}_{i \in \mathbb{I}^{C}, m \in \mathbb{M}, \boldsymbol{\mu} \in \mathbb{R}^{k-1}} \left\| \bar{\mathbf{A}}_{:i} - \mathbf{M}\boldsymbol{\mu} - \check{\rho}^{(m,i)} \,\check{\mathbf{\Pi}}_{\mathcal{M}^{(m)}} \left(\bar{\mathbf{A}}_{:i} - \mathbf{M}\boldsymbol{\mu} \right) \right\|^{2}$ Append $\check{\rho} \,\check{\mathbf{\Pi}}_{\mathcal{M}^{(m)}} \left(\bar{\mathbf{A}}_{:i} - \mathbf{M}\boldsymbol{\mu} \right)$ to **M** Append *i* to \mathbb{I} , remove *i* from \mathbb{I}^{C} , remove *m* from \mathbb{M} $m_{i} \leftarrow m$

Output: indices of the modeled spectra I, model-to-spectrum map $\{m_i\}_{i\in\mathbb{I}}$

2.4.4.3 Constraint application

It is recalled that the proximal operator of the characteristic function of a convex set is the orthogonal projection on the aforementioned set. Depending on whether a spectrum $\mathbf{A}_{:i}$ is modeled or not, the applied constraint differs:

• Unknown spectrum: projection on \mathcal{K}_a . If a spectrum is not constrained (*i.e.*, $i \notin \mathbb{I}$), the update reads as:

$$\mathbf{A}_{:i} \leftarrow \Pi_{\mathcal{K}_a} \left(\bar{\mathbf{A}}_{:i} \right) = \frac{\max\left(\mathbf{A}_{:i}, \mathbf{0} \right)}{\max\left(1, \left\| \max\left(\bar{\mathbf{A}}_{:i}, \mathbf{0} \right) \right\|_2 \right)}.$$
(2.17)

• Modeled spectrum: projection on manifold. If a spectrum is constrained (*i.e.*, $i \in \mathbb{I}$), it is projected on its associated manifold:

$$\mathbf{A}_{:i} \leftarrow \Pi_{\mathcal{M}^{(m_i)}} \left(\bar{\mathbf{A}}_{:i} \right). \tag{2.18}$$



Figure 2.5: Spatial templates obtained from Chandra X-ray observations of the Cassiopeia A supernova remnant (logarithmic scale).

2.5 Numerical experiments

In this section, the proposed method is evaluated on a realistic toy model of the Cassiopeia A supernova remnant observed by the X-ray space telescope Chandra¹ (see [Picquenot et al., 2019] for more details about these data). The data are composed of I = 4 sources of size $P = 346 \times 346$ (see Fig. 2.5), specifically one synchrotron source (radiation of the energetic charged particles in the supernova remnant by the synchrotron process), one thermal source (emission from a 10^7 K plasma including continuum and lines emissions) and two Gaussian line emissions. The mixtures are observed over J = 75 channels. An absorbed power-law model and a hot plasma emission model produced using the Astrophysical Plasma Emission Code (APEC [Foster and Heuer, 2020]) convolved with the spectral response of the Chandra telescope are used to generate the sets of synchrotron and thermal spectra (see Fig. 2.6a and 2.6b). The emission line spectra are modeled as Gaussian kernels, with widths proportional to the center as shown in Fig. 2.6c.

We define three experimental parameters, namely:

- the signal-to-noise ratio SNR, which is defined as the ratio of signal energy $\|\mathbf{AS}\|_2^2$ to the noise energy $\|\mathbf{N}\|_2^2$,
- the distance between the center of the two Gaussian line spectra δ (expressed in terms of spectral sample), to test how the proposed method separates two components whose spectra correlate this happens for example when the underlying physical processes are similar,
- the amplitude ratio k, which is a scalar by which the thermal and Gaussian sources are multiplied, so as to unbalance the sources and test how the algorithm

 $^{^1}$ chandra.harvard.edu



Figure 2.6: Ensemble of three emission models. The colored thick lines are the chosen anchor points in the context of the IAE modeling.

recovers the least energetic components, which is a common situation in physical applications.

The sGMCA code that is used is open source (see Appendix D). To the best of our knowledge, sGMCA is the first BSS method that makes use of a learned prior on the mixing matrix. We compare sGMCA with two benchmark experiments:

- an oracle version of GMCA, where \mathbf{S} (respectively \mathbf{A}) is estimated with the ground-truth \mathbf{A} (respectively \mathbf{S}); it provides an upper-bound for the reconstruction performances,
- an alternate version of sGMCA in which the spectra are regularized with a nearest-neighbor search among the spectra of the training sets; this experiment allows to better highlight the benefit of the IAE regularization.

The sGMCA algorithm is also compared to three standard BSS algorithms, which have no data-driven prior on the spectra, namely GMCA, HALS [Gillis and Glineur, 2012] and SNMF [Le Roux et al., 2015] (see Chapter 1). For the GMCA-based algorithms, the sparsity of the sources is enforced in the starlet (*i.e.*, an isotropic undecimated wavelet) representation with two details scales [Starck et al., 2015].

The estimated spectra and mixing matrices are assessed with the spectral angular distance (SAD, see Chapter 1). Concerning the sources, they are evaluated with the signal-to-distortion, interference, noise and artifacts ratios (respectively SDR, SIR, SNR² and SAR, see definitions in Chapter 1). The metrics are expressed in an inverse logarithmic scale, so the greater they are, the more accurate the estimates.

 $^{^2 \}rm Note that SNR in italics refers to the signal-to-noise ratio of the observed data, while SNR in roman is the criterion.$

Value		
Sync.	Therm.	Gauss.
3	3	2
541	601	600
178	198	200
10000	10000	100000
	0.1	
	4	
	0.1	
$1\mathrm{e}-4$		
Adam		
	25	
	Sync. 3 541 178 10 000	Value Sync. Therm. 3 3 541 601 178 198 10000 10000 0.1 4 0.1 1e - 4 Adam 25

Table 2.1: Parameters of the learning stage.

2.5.1 Learning of the generative models

Before tackling the source separation problem, the generative models of the three families of spectra need to be learned. The first step is to choose the anchor points and their numbers. The minimum number of anchor points required by the IAE is the dimension of the underlying manifold plus one. The thermal and synchrotron spectra depend on two physical parameters, therefore we will select three anchor points. The Gaussian spectra depend on one parameter only, thus two anchor points are needed. The choice of the anchor points is in practice not particularly critical; as long as they are not colinear, the reconstruction performances are satisfying. The anchor points are selected by hand, so that they tend to maximize their contrasts (see Fig. 2.6). The three sets of spectra are decomposed into training, validation and test sets. The reconstruction performances are found to stabilize as of four layers, which is therefore the number of layers chosen for all three models. The values of the hyper-parameters of the IAE models are summarized in Table 2.1.

Figure 2.7 shows examples of projections of spectra from the test sets on manifolds modeled by IAE, alongside with the projection errors. The spectra are well reproduced on a reasonably wide dynamic, except for the low-amplitude samples. To correct this, the spectra could be learned in a logarithmic scale, but it is out of the scope of the present work. The median SAD of each type of spectra and in overall are reported in Table 2.2. These results give the IAE modeling error bounds, that sGMCA cannot exceed.

Figure 2.8 shows two optimization landscapes of projections on manifold, that is Eq. (2.15) as a function of the latent parameter λ . As stated earlier, the optimization landscape is quite convex near the solution, which allows using proximal minimization



Figure 2.7: Examples of projections of spectra from the test sets on manifolds modeled by IAE. Top row: spectra (solid lines: results of projection, dashed lines: ground truth). Bottom row: projected spectra over ground-truth spectra (plotted if ground truth greater than 1e - 4). The figures on a same row share the same ordinate range.

Spectra	Sync.	Therm.	Gauss.	Overall
SAD (dB)	25.26	25.92	24.81	25.18

Table 2.2: Median reconstruction SAD over the test sets by the IAE.



Figure 2.8: Optimization landscapes of the manifold projections, that is $\|\mathbf{a} - \rho^* g(\boldsymbol{\lambda})\|_2^2$ as a function of $\boldsymbol{\lambda}$ with the ground truth ρ^* , in an inverse logarithmic scale. Blue cross: result of the manifold projection, red circle: initialization (by the fast projection). Top row: overview. Bottom row: zoom on the solution. (a) Synchrotron model, with $\boldsymbol{\lambda}_3 = 1 - \boldsymbol{\lambda}_1 - \boldsymbol{\lambda}_2$ to fulfill the sum-to-one constraint of the latent parameter inherent to the IAE. (b) Gaussian model, the dashed line is the sum-to-one constraint $\boldsymbol{\lambda}_1 + \boldsymbol{\lambda}_2 = 1$.

schemes as long as the initialization is decent.

2.5.2 Results

2.5.2.1 Overall results

Firstly, let us compare qualitatively sGMCA to the BSS algorithms on a typical run. Figure 2.9 shows the estimated spectra along with the estimation errors. The three BSS methods are particularly prone to interferences, which result in either portions of spectra misestimated at zero or with leakages of other components, which is problematic for astrophysical interpretations. On the contrary, sGMCA manages to remove most interferences and recovers satisfactory spectra. Figures 2.10 and 2.11 show the estimates of respectively the synchrotron source and the Gaussian II source, as well as the associated estimation errors. Residuals of other sources and/or reconstruction artifacts are clearly visible in the estimation errors of the three blind methods. The sGMCA algorithm provides more accurate sources, whose error equally originates from interference, noise contamination and artifacts.

Concerning computation times, sGMCA is underiably slower than the blind methods due to the manifold projections. For example, in the tests performed, the projection of the four spectra takes approximately three seconds and 50 iterations are typically necessary. In contrast, the blind methods run in a few seconds.



Figure 2.9: Example of estimated spectra, with SNR = 40 dB, $\delta = 20$ and k = 1. Top row: spectra (solid lines: estimation, dashed lines: ground truth). Bottom row: estimated spectra over ground truth spectra (plotted if ground truth greater than 1e - 4). The figures on a same row share the same ordinate range.

2.5.2.2 Impact of the experimental parameters

In order to confirm quantitatively these results, three experiments are performed, in each of which an experimental parameter is varied. They are constituted of Monte-Carlo trials with varying spectra (taken from the test sets) and noise realizations. The overall results are reported in Fig. 2.12 and are commented on in the paragraphs below.

The sparsity hyperparameter of SNMF is found to be very sensitive to the noise level, noise realization and mixing matrix. In the reported results, a fixed mean "good" hyperparameter $\lambda = 10^{-7}$ is taken.

Noise level The overall tendencies are consistent; the higher the SNR, the better the estimations. As we had observed previously on the example, sGMCA estimates particularly well the spectra; the gain in SAD is from a few to 10 dB compared to


Figure 2.10: Example of the estimated synchrotron source, with SNR = 40 dB, $\delta = 20$ and k = 1. Top: estimations (logarithmic scale), bottom: absolute error (logarithmic scale). The figures on a same row share the same color scale.

GMCA, depending on the noise level. Since the sources are better disentangled, this results in a significant gain of the SIR and thus the SDR.

The additional prior information to which the nearest-neighbor benchmark algorithm has access leads to improved performances compared to GMCA, but not as much as sGMCA. This emphasizes the advantage of the regularization provided by the IAE; by modeling the manifolds on which the spectra evolve, the latter are reconstructed much more precisely.

It is noted that the performance metrics of sGMCA and GMCA can reach a plateau at a high *SNR*. According to the hyperparameter tuning strategy of GMCA and sGMCA, the thresholds applied to the sources are low at a low noise level, inducing an underregularization. Interestingly, the sGMCA is less sensitive to this effect, most likely because the components are better separated.

Collinearity of the spectra Similarly to the previous experiment, sGMCA outperforms the BSS algorithms, with appreciable gains in SAD by 10 to 20 dB and in SDR by 5 to 10 dB compared to GMCA. When the two Gaussian spectra tend to coincide ($\delta \rightarrow 0$) the source estimates of the GMCA-based algorithms decline because the least-square update of the sources in the minimization scheme becomes



Figure 2.11: Example of the estimated Gaussian II source, with SNR = 40 dB, $\delta = 20$ and k = 1. Top: estimations (logarithmic scale), bottom: absolute error (logarithmic scale). The figures on a same row share the same color scale.

ill-conditioned. Moreover, the GMCA metrics are relatively insensitive to δ ; a more detailed analysis shows that the GMCA errors are dominated by the synchrotron component, on which δ has indeed a negligible effect.

Unbalanced sources The reported source metrics concern only the thermal and Gaussian components, since we want to assess the impact on the hidden sources. Again, sGMCA allows to recover more precisely the spectra, with a considerable gain in SAD ranging from 8 dB to 12 dB compared to GMCA, and up to 5 dB for the SDR.

When the sources are unbalanced, GMCA estimates precisely the synchrotron spectrum but very poorly the three other spectra (they are contaminated by the synchrotron component). The accurate estimation of the synchrotron spectrum makes it possible to separate the synchrotron source from the data in the least-square update of the sources, which allows for the retrieval of the hidden sources, hence the acceptable SDR (see examples of estimated sources when the synchrotron source is a hundred times brighter in Fig. 2.13 and 2.14).

On the contrary, the nearest-neighbor benchmark algorithm performs particularly poorly; it generally fails to identify the spectra of the three hidden sources, because



Figure 2.12: Median metrics over 100 realizations, and first and third quartile for sGMCA, when varying the (a) noise level, (b) collinearity of the Gaussian line spectra, (c) unbalance of the sources. The top row (SAD) concerns only the spectra, the other rows (SDR, SIR, SNR and SAR) concern only the sources. (c) The reported SDR, SIR, SNR and SAR are calculated over the thermal and Gaussian sources only.

at the first iteration (which we recall to be the output of GMCA) the nearest neighbors of all four spectra are synchrotron spectra.



Figure 2.13: Example of the estimated synchrotron source, with SNR = 40 dB, $\delta = 20$ and k = 0.01. Top: estimations (logarithmic scale), bottom: absolute error (logarithmic scale). The figures on a same row share the same color scale.

2.5.2.3 Impact of the generative modeling

In the previous experiments, all four components were constrained by generative models. In this subsection, we focus on the impact of having one or more components fully unknown. To that end, Monte Carlo experiments are performed where the IAE models of the synchrotron, thermal and/or Gaussian spectra are removed, both in the balanced and an unbalanced case. The results are reported in Table 2.3.

In the balanced case, the performance metrics are particularly sensitive to the presence or not of the thermal model. The thermal source strongly correlates with the two Gaussian line sources, it is therefore not surprising that constraining the thermal spectrum notably improves the unmixing.

In the unbalanced case, where the synchrotron source has a hundred times larger norm than the other sources, it is particularly advantageous to include the synchrotron model in addition to the thermal model. This reduces the leakage of the synchrotron component into the estimates of the other hidden components.



Figure 2.14: Example of the estimated Gaussian II source, with SNR = 40 dB, $\delta = 20$ and k = 0.01. Top: estimations (logarithmic scale), bottom: absolute error (logarithmic scale). The figures on a same row share the same color scale.

2.6 Conclusion

We introduce a novel source separation approach to tackle physical multispectral data. Compared to standard blind source separation methods, the objective is twofold: to better discriminate between sources and to ensure the provision of physically relevant information. For this purpose, we make use of learned priors, which are based on generative models, on the spectra of the sought-after components in a standard variational framework. Extensive numerical experiments on realistic astrophysical data show that the introduced regularization efficiently rejects inter-component leakages, thus improving significantly the estimation of both the sources and the spectra, including in challenging settings.

		Components modeled with a generative model				
k	Metric (dB)	All	Therm. & Gauss.	& Sync. & Gauss.	Gauss.	None
1	SAD SDR SIR SNR SAR	$21.67 \\ 43.93 \\ 44.16 \\ 31.98 \\ 30.48$	$20.71 \\ 42.41 \\ 42.16 \\ 31.97 \\ 29.42$	15.59 39.06 33.35 31.96 22.81	15.32 37.74 33.33 31.91 22.57	$12.54 \\ 34.24 \\ 17.49 \\ 31.91 \\ 16.30$
0.01	SAD SDR SIR SNR SAR	$17.73 \\ 36.64 \\ 42.54 \\ 31.87 \\ 29.24$	12.27 34.39 35.97 31.68 27.00	$8.83 \\ 36.96 \\ 36.70 \\ 31.67 \\ 25.05$	8.27 37.02 37.05 31.62 25.00	$5.78 \\ 34.56 \\ 18.22 \\ 31.47 \\ 17.31$

Table 2.3: Median metrics over 100 realizations. The sGMCA algorithm is run in various semi-blind configurations, with different known and unknown components. k is recalled to be the ratio between the norm of the thermal or a Gaussian source (they have the same norm) and the norm of the synchrotron source. For k = 0.01, the source metrics are calculated over the thermal and Gaussian sources.

Chapter 3

Measuring and imaging the radio sky with a view to performing blind source separation

Radio astronomy is a branch of astronomy dealing with the observation of the sky in the radio frequency domain. The precise measurement of radio sources, which are generally low-power signals, requires the use of massive instruments, radio telescopes, which mobilize significant resources both in material and computational terms.

In this chapter, the fundamentals of radio astronomy measurements are first introduced. The acquisition model of cosmic radio emissions is then developed, on the one hand with single-dish telescopes and on the other hand with interferometers. Lastly, the acquisition models are extended to the multispectral case, with a view to addressing the BSS problem later on.

The content of this chapter builds upon the reference books in radio-astronomy [Taylor et al., 1999] and [Thompson et al., 2017], as well as the lecture notes [Klein, 2011] and [Sasao and Fletcher, 2005].

3.1 Measuring astronomical radio emissions

3.1.1 Characterization of radio sources

In astronomy, the physical processes at the origin of radio emissions are chaotic. The main examples are:

• thermal radiation due to the random movement of heated particles, which includes cosmological signals such as the cosmic microwave background or the EoR signals, but also astrophysical objects such as interstellar clouds,

- synchrotron emission due to the random acceleration of electrons in magnetic fields,
- spectral line radiation produced by the random transition between different energy states of atoms and molecules.

It follows that the received signals are random, and that a statistical processing stage is required to derive any valuable information. For this purpose, the random process framework is employed; some elements are recalled in Appendix A. Hereafter, we make the standard assumptions of radio astronomy concerning radio emissions. Firstly, radio sources are approximated as wide-sense stationary ergodic random processes. Moreover, they are supposed to be spatially incoherent – that is the cross correlation of the electric field incoming of two different directions is null.

In order to characterize radio sources, astronomers rely on spectral densities (which are constant quantities in time, thanks to the stationarity hypothesis). The most employed is the brightness $B^{(\nu)}$, also called monochromatic or specific intensity – ν corresponds to a radio frequency. It is defined as the electromagnetic radiation power coming from a unit solid angle in the sky, through a perpendicular unit area, per unit of frequency bandwidth around frequency ν , and thus expressed in terms of Wm⁻²Hz⁻¹sr⁻¹. The brightness has the advantage of being conserved in empty space (such that it is the same at the source and at the receiver), and of not depending on the measurement/reception conditions (namely the field of view, the bandwidth and the effective area).

Another spectral density quantity of interest is the flux $S^{(\nu)}$ (also referred to as spectral flux density). It is the amount of radiation energy through a unit-area cross section, per unit frequency bandwidth around ν and per unit of time, and is expressed in terms of Wm⁻²Hz⁻¹ or Jansky (1 Jy = 1e - 26 Wm⁻²Hz⁻¹). It relates to the intensity by: $S^{(\nu)} = \int B^{(\nu)}(\mathbf{s}) d\Omega$.

3.1.2 Radio antennas

An antenna is a device that can radiate or receive an electromagnetic field (see Fig. 3.1). Let us define some of its main characteristics.

Power pattern The power pattern describes how well the antenna radiates or receives the electromagnetic power according to the angular direction. It depends mostly on the antenna geometry. The fundamental reciprocity theorem in electromagnetic shows that the pattern is the same for radiation and reception.

Figure 3.2 illustrates an example of a normalized power pattern. It typically features a primary lobe, centered on the so-called aperture axis, and side lobes. The



Figure 3.1: Illustration of a single-dish telescope scanning the sky brightness.

half power beam width (HPBW) is defined as the angular width in which the power pattern amplitude is greater than half of the maximum. It affects the directivity of the antenna, and consequently the resolution when forming images afterwards. In rule of thumb, the HPBW relates to the radiation wavelength λ and the antenna diameter D by the relation HPBW $\approx \lambda/D = c/(\nu D)$. In this regard, Fig. 3.2 shows the effect of a frequency change on the normalized power pattern. In addition, side lobes deteriorate the antenna directivity and thus the resolution, they are therefore attenuated as much as possible during design.

Consider a diffuse source in the sky with a brightness $B^{(\nu)}(\mathbf{s})$. The flux measured by an antenna targeting the direction of $\mathbf{s_0}$ is the brightness weighted by the antenna normalized power pattern $P_n^{(\nu)}$: $S^{(\nu)}(\mathbf{s_0}) = \int P_n^{(\nu)}(\boldsymbol{\sigma}) B^{(\nu)}(\mathbf{s_0} + \boldsymbol{\sigma}) d\Omega$ with $d\Omega$ an infinitesimal solid angle element of direction $\boldsymbol{\sigma}$.

Antenna aperture The antenna aperture $A_e^{(\nu)}$, which is expressed in terms of m^2 , relates the received flux to the measured power spectral density $W^{(\nu)}$ by $W^{(\nu)} = A_e^{(\nu)} S^{(\nu)}/2$ (the term 1/2 implies that a single component polarization is measured by the antenna).

Thereafter, the following model between the brightness of an astronomical radio source and the power spectral density at the output of an antenna targeting a point s_0 in the sky will be more simply considered:

$$W^{(\nu)}(\mathbf{s_0}) = \int A^{(\nu)}(\boldsymbol{\sigma}) B^{(\nu)}(\mathbf{s_0} - \boldsymbol{\sigma}) \,\mathrm{d}\Omega, \qquad (3.1)$$

where $A^{(\nu)}(\boldsymbol{\sigma}) = A_e^{(\nu)} P_n^{(\nu)}(-\boldsymbol{\sigma})/2$ is the antenna instrumental response.

Measuring and imaging the radio sky with a view to performing blind source separation



Figure 3.2: Example of normalized power patterns $P_n^{(\nu)}$ for different frequencies ν and illustration of the HPBW.

Receiving systems In practice, an antenna measures a voltage. A signal processing stage is needed at the output of the antenna to derive the power spectral density, so to determine the brightness (see Fig. 3.1). Several techniques can be distinguished:

- correlator: it consists in the combination of a bandpass filter around a frequency of interest, a square law detector and an integrator. Thanks to the ergodicity hypothesis, the output approximates the power spectrum on the considered band.
- filterbank spectrometer: it is the generalization of the correlator to several frequency bands. The voltage is sent on a series of bandpass filters (a filterbank), which are connected to square law detectors and integrators. The combination of the outputs gives the whole power spectrum.
- autocorrelation spectrometer: this is a hardware implementation of the Wiener-Khinchin theorem (see Appendix A). Each sample is multiplied with the other successively delayed samples, and the ensemble is summed. This allows to calculate the autocorrelation of the signal, from which the power spectral density is deduced by Fourier transform.
- Fourier spectrometer: this is the most recent technique. The whole signal is Fourier transformed, and the square modulus of the result is calculated to derive the power spectrum. The main burden of this technique is the computation of the Fourier transform, which can be particularly costly for big datasets.

3.2 Single-dish imaging

An extended source can be imaged by scanning the sky at several locations and joining the measurements [Mangum, J. G. et al., 2007]. When doing so, several detrimental effects must be addressed, for instance the ones due to the aliasing or the gridding. The aim of this section being to develop eventually a model for BSS with radio data, we will consider them beyond the scope of the thesis and ignore them.

It is worth reminding that when tackling multispectral measurements, the antenna responses $\{A^{(\nu)}\}_{\nu}$ depend on the channel frequencies (as seen earlier, their resolution varies roughly in $\lambda/D = c/(\nu D)$).

Two cases can further be distinguished, depending on whether dealing with a small or large field of view.

3.2.1 Small fields of view

When imaging a small portion of the sphere, the problem can be reformulated on a plane. To do so, let us consider the plane tangent to the celestial sphere at the center of the field of view, parameterized by the system coordinates (l, m). Equation (3.1) may be rewritten as follows:

$$W^{(\nu)}(l,m) = \iint A^{(\nu)}(l',m') B^{(\nu)}(l-l',m-m') dl' dm'$$

= $[A^{(\nu)} * B^{(\nu)}] (l,m),$ (3.2)

where * is the usual two-dimensional convolution product. In the perspective of performing BSS, it is furthermore possible to simplify the model by expressing it in the (spatial) Fourier domain; indeed, denoting (u, v) the coordinates in Fourier space, the convolution product simplifies to an element-wise product:

$$\tilde{W}^{(\nu)}(u,v) = \tilde{A}^{(\nu)}(u,v) \,\tilde{B}^{(\nu)}(u,v). \tag{3.3}$$

3.2.2 Large fields of view

When imaging larger field of views, potentially the entire sphere, the spherical formulation is inevitable. As in the planar case, the model can be expressed with a spherical convolution product¹[Driscoll and Healy, 1994], that we will denote likewise *:

$$W^{(\nu)}(\mathbf{s}) = \left[A^{(\nu)} * B^{(\nu)}\right](\mathbf{s}). \tag{3.4}$$

¹Several convolution products are defined on the sphere [Roddy and McEwen, 2021], we consider here the so-called isotropic convolution product, which requires the beams to be isotropic. See Appendix B for more details.

Similarly to the planar case, the spherical convolution product is diagonalized by the spherical harmonic transform. The model therefore becomes:

$$\tilde{W}^{(\nu)}(\mathfrak{l},\mathfrak{m}) = \tilde{A}^{(\nu)}(\mathfrak{l})\tilde{B}^{(\nu)}(\mathfrak{l},\mathfrak{m}).$$
(3.5)

Interferometric imaging 3.3

As seen in Section 3.1.2, the resolution of an antenna is roughly equal to λ/D . In order to increase the resolution at a fixed wavelength, it is therefore necessary to increase the antenna diameter D. However, there is a limitation in terms of mechanical feasibility, typically at $D \sim 100$ m. At a low radio frequency of 100 MHz, for the study of the EoR signal for instance, this leads to a maximum resolution of 1.7° , which is far too coarse for scientific applications. Here lies one of the bottlenecks of single-dish imaging: In order to image the sky at low radio frequencies with a sufficiently fine resolution, antennas with an aperture diameter of potentially several kilometers would be required.

This is one of the main motivations for the use of radio-interferometers. By combining several antennas, it is possible on paper to achieve resolutions as fine as desired. However, the price to pay is both a more complex signal processing chain and more complex algorithms to form images.

3.3.1A basic twin interferometer

Let us first consider a two-dimensional two-element interferometer plugged to a correlator (see Fig. 3.3). The antennas are separated by a distance d (the so-called baseline) and supposed identical, such that they have the same instrumental response. In addition, let us consider a monochromatic point source of frequency ν in the sky, that forms an angle θ with the zenith.

The two antennas measure the incoming wavefront with a geometric delay $\tau(\theta) =$ $d\sin(\theta)/c$. The voltages at the output of the antennas can therefore be written $v_1(t,\theta) = V \cos(2\pi\nu t)$ and $v_2(t,\theta) = V \cos(2\pi\nu(t-\tau(\theta)))$. These are multiplied and then averaged on a sufficiently long time T by the correlator, yielding:

$$r(\theta) = \lim_{T \to \infty} \frac{1}{2T} \int_{-T}^{T} v_1(t,\theta) v_2(t,\theta) dt$$

=
$$\lim_{T \to \infty} \frac{1}{2T} \frac{V^2}{2} \int_{-T}^{T} \left(\cos\left(4\pi\nu t - \tau(\theta)\right) + \cos\left(2\pi\nu\tau(\theta)\right) \right) dt \qquad (3.6)$$

=
$$\frac{V^2}{2} \cos(2\pi\nu\tau(\theta)).$$



Figure 3.3: A basic twin interferometer with a point source monochromatic source.



Figure 3.4: Illustration of the gain in resolution by the interferometer. $\Delta \theta_{0,min}$ and $\Delta \theta_0$ are the angular resolutions of a single antenna and the interferometer, respectively.

The term $V^2/2$ is proportional to the power $P_0(\theta)$ measured by each antenna (the θ dependency comes from the antennas' non-isotropic response). The following power can be derived at the output of the correlator (see Fig. 3.4):

$$P(\theta) = P_0(\theta) \cos(2\pi\nu\tau(\theta)). \tag{3.7}$$

It is the power received by each antenna modulated by a so-called fringe pattern $\cos(2\pi\nu\tau(\theta))$.

This expression allows to highlight the gain in resolution provided by an interferometer. While an antenna can only distinguish sources with an angular distance $\Delta\theta$ greater than its HPBW ($\Delta\theta_{0,min} \approx$ HPBW, see Fig. 3.4), an interferometer can discern closer sources thanks to the fringe pattern (so long as $\Delta\theta_{min} \approx \nu \tau(\theta) \approx d/\lambda$, see Fig. 3.4). Thus, simply by moving antennas apart, images with a desired resolution can be produced.

The above reasoning neglects the motion of Earth, which causes θ to vary in time. It follows that the fringe pattern is not purely sinusoidal after integration over a long time. Hereafter, we will continue neglecting this phenomenon, supposing that the signal processing stage tracks the sources in the sky.

3.3.2 The visibility function

We now move to a more general three-dimensional setting with a diffuse multispectral radio source (Fig. 3.5). Let $\mathbf{s} = \mathbf{s}_0 + \boldsymbol{\sigma}$ be a unit-vector aiming towards the radio



Figure 3.5: Two elements of an interferometer in the general case.

source from the interferometer, with s_0 a reference unit-vector which will become the center of the imaged area.

In the case where the receiving system is the same correlator as before, the output power in a small frequency band $\Delta \nu$ around a frequency ν relates to the brightness $B^{(\nu)}(\boldsymbol{\sigma})$ by:

$$P^{(\nu)}(\mathbf{s_0}) = \Delta \nu \int A^{(\nu)}(\mathbf{s_0} - \boldsymbol{\sigma}) B^{(\nu)}(\boldsymbol{\sigma}) \cos(2\pi\nu\tau(\boldsymbol{\sigma})) \,\mathrm{d}\Omega, \qquad (3.8)$$

with $\tau(\boldsymbol{\sigma}) = \mathbf{d}^{\top} \mathbf{s}/c$ the geometric delay and $A^{(\nu)}(\boldsymbol{\sigma})$ the antennas' instrumental response (which are still supposed identical). Note the similarity of this expression with the output power of a single-dish telescope (Eq. 3.1). Developing the cosinus term yields:

$$P^{(\nu)}(\mathbf{s_0}) = \Delta\nu\cos(2\pi\nu\mathbf{d}^{\top}\mathbf{s_0}/c) \int A^{(\nu)}(\mathbf{s_0} - \boldsymbol{\sigma}) B^{(\nu)}(\boldsymbol{\sigma}) \cos(2\pi\nu\mathbf{d}^{\top}\boldsymbol{\sigma}/c) d\Omega - \Delta\nu\sin(2\pi\nu\mathbf{d}^{\top}\mathbf{s_0}/c) \int A^{(\nu)}(\mathbf{s_0} - \boldsymbol{\sigma}) B^{(\nu)}(\boldsymbol{\sigma}) \sin(2\pi\nu\mathbf{d}^{\top}\boldsymbol{\sigma}/c) d\Omega.$$
(3.9)

Let us now define the complex-valued visibility function $V^{(\nu)}$ as:

$$V^{(\nu)}(\mathbf{d}, \mathbf{s_0}) = |V^{(\nu)}| e^{i\varphi_{V^{(\nu)}}} \coloneqq \int A_n^{(\nu)}(\mathbf{s_0} - \boldsymbol{\sigma}) B^{(\nu)}(\boldsymbol{\sigma}) e^{-2\pi i\nu \mathbf{d}^{\top} \boldsymbol{\sigma}/c} \,\mathrm{d}\Omega, \qquad (3.10)$$

where $A_n^{(\nu)}(\mathbf{s}) \coloneqq A^{(\nu)}(\mathbf{s}) / A^{(\nu)}(\mathbf{0})$ is the normalized antenna response. Separating the real and imaginary parts yields:

$$\operatorname{Re}\left(V^{(\nu)}\right) = \int A_{n}^{(\nu)}\left(\mathbf{s_{0}}-\boldsymbol{\sigma}\right) B^{(\nu)}\left(\boldsymbol{\sigma}\right) \cos\left(2\pi i\nu \mathbf{d}^{\top}\boldsymbol{\sigma}/c\right) d\Omega = \left|V^{(\nu)}\right| \cos\left(\varphi_{V^{(\nu)}}\right),$$
$$\operatorname{Im}\left(V^{(\nu)}\right) = -\int A_{n}^{(\nu)}(\mathbf{s_{0}}-\boldsymbol{\sigma}) B^{(\nu)}\left(\boldsymbol{\sigma}\right) \sin\left(2\pi i\nu \mathbf{d}^{\top}\boldsymbol{\sigma}/c\right) d\Omega = \left|V^{(\nu)}\right| \sin\left(\varphi_{V^{(\nu)}}\right).$$
$$(3.11)$$

The power then relates to the visibility by:

$$P^{(\nu)}(\mathbf{s_0}) = \Delta \nu A^{(\nu)}(\mathbf{0}) \left| V^{(\nu)} \right| \cos \left(2\pi \nu \mathbf{d}^\top \mathbf{s_0} / c - \varphi_{V^{(\nu)}} \right).$$
(3.12)

The interferometer output is a fringe pattern whose period is that of a fictitious source at the reference position \mathbf{s}_0 . Moreover, the amplitude of the fringe pattern is proportional to that of the visibility, and the phases of the fringe pattern and the visibility are the same (noting that $\varphi_{V(\nu)}$ is measured relative to the reference position \mathbf{s}_0).

In practice, the receiving systems are more sophisticated than the basic correlator considered above. Firstly, special structures allow obtaining the complex visibility directly (complex correlators). Secondly, as with single-dish imaging, they can be associated with spectrometers in order to derive the visibility function on a frequency band or on several frequency channels at once. Hereafter, it will be considered that each antenna pair gives directly a sample of the visibility function. In particular, the calibration aspects, which amount for instance to determine exactly the antennas' response, are set aside.

In summary, in single-dish imaging, the antennas are particularly directive so that the sky brightness can be measured directly point by point. In interferometric imaging, the antennas are usually small and so have large beams. Thus, each pair of antennas acquires the brightness over a large area of the sky, but with a modulation (by a fringe pattern (Eq. 3.8) or equivalently by a complex exponential (Eq. 3.10)). In order to recover the brightness, algorithmic processing is required.

3.3.3 Coordinate systems for imaging

In this thesis, we will limit ourselves to interferometric data acquired on a small field of view, allowing us to synthesize images faithfully on a plane. Indeed, wide-field



Figure 3.6: Coordinate systems for interferometric imaging. Image from [Thompson et al., 2017].

interferometry, which uses the spherical formalism, is still a relatively unexplored domain that requires a high computational cost [Carozzi, 2015]. Let us introduce two Euclidean coordinate systems dedicated for interferometric imaging problems (see Fig. 3.6).

Let (u, v, w) represent the interferometer baseline coordinates, where u, v and w respectively point towards East, North and the center of the region of interest in the sky \mathbf{s}_0 . The terms u, v and w are expressed in terms of wavelength.

Positions in the sky are indicated with the direction cosines with respect to u and v, and denoted respectively l and m. The (l, m)-plane upon which the images are formed is the plane tangent to the celestial sphere at the reference position s_0 .

Before going any further, some remarks are made. Firstly, the field of view of the antennas being large, it is assumed that a single pointing is sufficient to image an area of the sky, and that the antennas' response in this area is constant $(A_n^{(\nu)}(\mathbf{s_0} - \boldsymbol{\sigma}) \approx 1)$. Thus, to simplify the notations, the dependency in $\mathbf{s_0}$ will not be indicated, nor will the antennas' responses².

Secondly, we will work with a single fixed frequency channel for the time being – multi-frequency aspects will be addressed in Section 3.4. Thus, the dependency in ν will be temporarily omitted for the sake of readability.

With the latter assumptions and the coordinate system presented above, the following results are obtained in the definition of the visibility function (3.10):

$$2\pi\nu \frac{\mathbf{d}^{\top}\boldsymbol{\sigma}}{c} = ul + vm + w\left(\sqrt{1 - l^2 - m^2}\right),$$

$$2\pi\nu \frac{\mathbf{d}^{\top}\mathbf{s_0}}{c} = w,$$

$$d\Omega = \frac{dldm}{\sqrt{1 - l^2 - m^2}},$$
(3.13)

The visibility function then writes:

$$V(u, v, w) = \iint \frac{B(l, m)}{\sqrt{1 - l^2 - m^2}} e^{-2\pi i w \left(\sqrt{1 - l^2 - m^2} - 1\right)} e^{-2\pi i (ul + vm)} dl dm.$$
(3.14)

The equation is further simplified. Firstly, we include the denominator term in the brightness: $B(l,m) \leftarrow B(l,m)/\sqrt{1-l^2-m^2}$ (noting that the actual brightness can trivially be deduced by multiplying the estimated brightness by $\sqrt{1-l^2-m^2}$). Secondly, we define the function $G(l,m,w) \coloneqq e^{-2\pi i w (\sqrt{1-l^2-m^2}-1)}$, that we will refer to

²It is noted that the antennas' responses could always be taken into account by dividing the estimated brightness by $A_n^{(\nu)}$.

as the w-term. This leads to:

$$V(u, v, w) = \iint B(l, m) G(l, m, w) e^{-2\pi i (ul + vm)} dl dm.$$
(3.15)

3.3.4 Towards the forward observation model

An interferometer samples the visibility function in the (u, v, w)-space based on the positions of the antennas. In this regard, let $\{(u_k, v_k, w_k)\}_{k \in [1...K]}$ denote the baseline coordinates of an N-element interferometer. The number of measurements K is the number of pairs of antennas, which relates to the number of antennas N by K = N(N-1)/2. It is possible to define a sampling function:

$$S(u, v, w) = \sum_{k=0}^{K} \delta(u - u_k) \delta(v - v_k) \delta(w - w_k), \qquad (3.16)$$

with δ the Dirac distribution. The sampling function allows to conveniently write the forward model of the interferometer:

$$V_S(u, v, w) = S(u, v, w) \iint B(l, m) G(l, m, w) e^{-2\pi i (ul + vm)} dl dm, \qquad (3.17)$$

where V_S corresponds to the data acquired by the interferometer.

3.3.5 Deconvolution algorithms

In radio interferometry, deconvolution consists in estimating the brightness from the measured visibility samples. In this section, we review common deconvolution algorithms in the field. First, the coplanar case is addressed, which constitutes the main case of study so far, both from a methodological and an application point of view. Then, the non-coplanar case is tackled, which is a much more novel subject.

Until now, no mention has been made of the gridding of the data, which consists in placing the interferometric measurements on a regular grid, particularly with a view to using the fast Fourier transform (FFT), which is significantly faster than the discrete Fourier transform. The gridding is generally performed by convolution with a so-called gridding function. The most advanced deconvolution methods include the gridding in a concern of precision. The choice is made below not to mention the gridding, because the BSS algorithms developed thereafter will not include it, and the developments would be unnecessarily complex.

3.3.5.1 The coplanar case

The coplanarity assumption is satisfied when the argument of the *w*-term *G* always remains negligible: $2\pi w(1 - \sqrt{1 - l^2 - m^2}) \ll 1$. This occurs when both:

- the antennas of the interferometer all lie on a plane (which is mostly the case if the baselines are small), implying that $w \ll 1$,
- the size of view is narrow, which leads to $l \ll 1$ and $m \ll 1$, and thus $1 \sqrt{1 l^2 m^2} \ll 1$.

Consequently, the *w*-term disappears $(G(l, m, w) \approx 1)$. The visibility becomes a function of *u* and *v* only, and relates to the brightness by:

$$V(u,v) = \iint B(l,m) e^{-2\pi i (ul+vm)} \,\mathrm{d}l \,\mathrm{d}m, \qquad (3.18)$$

which is none other than the two-dimensional Fourier transform: $V(u, v) = [\mathcal{F} \circ B](u, v) = \tilde{B}(u, v).$

Therefore, the forward observation model of the interferometer writes:

$$V_S(u,v) = S(u,v) \ B(u,v)$$
 (3.19)

Again, note the similarity to the single-dish case (Eq. 3.3). The main difference is that the convolution operator is here a mask, with zero values, and not a smoother power pattern.

Two major classes of deconvolution algorithms can be distinguished in the literature: the methods built upon CLEAN and those based on an optimization framework [Thompson et al., 2017].

CLEAN-based algorithms Since the brightness and the visibility are Fourier pairs, one may try to inverse Fourier transform the data in order to retrieve the brightness:

$$B_d(l,m) \coloneqq [\mathcal{F}^{-1} \circ V_S](l,m). \tag{3.20}$$

The resulting image B_d generally has strong artifacts, due to the partial sampling of the visibility space, and is therefore called "dirty" image. Indeed,

$$B_d(l,m) = [S_d * B](l,m), \tag{3.21}$$

which shows that the dirty image is the true image convolved by the so-called "dirty beam" $S_d := \mathcal{F}^{-1}(S)$, that is the inverse Fourier transform of the sampling function.

The CLEAN algorithm [Högbom, 1974] has been the reference deconvolution method in radio astronomy for many years already. It is a greedy algorithm that has similarities with the matching pursuit algorithm [Mallat and Zhang, 1993]. It supposes that the observed sky is constituted of point sources. The gist of CLEAN is to iteratively:

- (i) backward step: find the brightest point in the dirty image (location and amplitude) and add it in a list,
- (ii) forward step: remove a portion of the brightest point from the dirty image, using the dirty beam S_d .

The reconstructed ("cleaned") image is formed by adding the identified point sources and convolving it with a predefined beam in order to match the interferometer resolution.

A first improvement of CLEAN was proposed in [Clark, 1980]; it accelerates the algorithm by considering approximate dirty beams and by using minor and major cycles. This has been incorporated in a second upgrade [Schwab, 1984], which allows a gain in accuracy by working directly with the ungridded visibilities.

Several extensions of CLEAN have since been proposed to deal with more specific cases. In particular there exists multi-resolution versions of CLEAN to tackle extended sources [Wakker and Schwarz, 1988, Starck et al., 2002, Cornwell, 2008], some of which include wavelet decompositions.

CLEAN-based methods are still the reference today in physical applications. Indeed, they are quite modular and simple to implement. Above all, they are compatible with large data sets such as can be found in radio-interferometry.

Optimization-based methods The second family of deconvolution algorithms, more recent, relies on sparsity-based convex optimization [Carrillo et al., 2012, Girard et al., 2015]. The problem is formulated in the form of a cost function which classically includes the physical model of the acquisition by the interferometer and some priors on the solution (among which sparsity in one or more transformed domains). The main obstacle to the use of these methods in scientific applications is their computational cost, which can quickly become prohibitive. The latest developments are therefore focused on the acceleration/parallelization of the algorithms [Carrillo et al., 2014, Thouvenin et al., 2021]

3.3.5.2 The non-coplanar case

The non-coplanar effects come from the *w*-term *G* which modulates the sky brightness by introducing a phase difference. It is noticeable as soon as large fields of view (*i.e.*, $1-l^2-m^2 \ll 1$) and/or long antenna baselines (which generally leads to $w \gg 1$) are involved. It causes a spread spectrum in the visibility space [Wiaux et al., 2009]. To illustrate this, consider the case where *w* is non-zero and constant; then the visibility writes $V(u, v, w) = [\tilde{G}^{(w)} * \tilde{B}](u, v)$ with $G^{(w)}(l, m) \coloneqq G(l, m, w)$. The consequence is that the handy Fourier transform can no longer be used to switch from the visibility to the image space, as in the coplanar case. Measuring and imaging the radio sky with a view to performing blind 70 source separation

The main methods for dealing with non-coplanar effects are reviewed below. Some deal with specific cases, e.g., a wide-field but a coplanar interferometer, or a small w-effect. We note that all the methods were first proposed to be included in the CLEAN framework.

Combination of snapshots This method concerns the particular case of coplanar radio-interferometers (*i.e.*, all antennas lie on a plane). When long observations are performed, because of Earth's rotation, the baselines of the interferometer move, which has the effect of rotating the plane and thus breaking the coplanar nature of the measurements. A first technique amounts to decompose the measurements in a set of so-called snapshots, in each of which the planar baseline assumption is valid, and to correct the rotation so as to apply the usual planar inverting techniques.

More precisely, when dealing with coplanar baselines, the w coordinate terms are shown to be linear combinations of the u and v terms: $w = u \tan Z \sin \chi - v \tan Z \cos \chi$ with Z the zenith angle and χ the parallactic angle of the interferometer [Perley, 1999]. Introducing the new coordinates system (l', m'):

$$\begin{cases} l' = l + \tan Z \sin \chi \left(\sqrt{1 - l^2 - m^2} - 1 \right) \\ m' = m - \tan Z \cos \chi \left(\sqrt{1 - l^2 - m^2} - 1 \right), \end{cases}$$
(3.22)

the *w*-term disappears. While this method is shown to suffer a lack of precision in practice, the idea of grouping coplanar measurements can be borrowed in other techniques for speed considerations (*e.g.*, [Cornwell et al., 2012] with *w*-projection, which is presented below).

Three-dimensional Fourier transform The visibility function can be embedded in a triple integral [Perley, 1999]:

$$V(u, v, w) = \iiint B(l, m) \delta\left(n - \sqrt{1 - l^2 - m^2}\right) e^{-2\pi i (ul + vm + w(n-1))} dl \, dm \, dn, \quad (3.23)$$

which is a particular three-dimensional Fourier transform of the function $(l, m, n) \mapsto B(l, m)\delta(n - \sqrt{1 - l^2 - m^2})$. The latter can be inverted by:

$$B(l,m) = \iiint V(u,v,w) e^{2\pi i \left(ul + vm + w(\sqrt{1-l^2 - m^2} - 1)\right)} \mathrm{d}u \, \mathrm{d}v \, \mathrm{d}w, \qquad (3.24)$$

either with a three-dimensional FFT, or with a two-dimensional FFT along the (u, v) axes combined with a Discrete Fourier Transform along the w axis.

This technique has not been used in practice for both memory and computational reasons. [Smith et al., 2017] suggests it could still be of interest in some forthcoming scenarios, with particularly non-coplanar baselines combined with very large field of views.

w-projection Recalling that the Fourier transform of a pointwise product is the convolution of their Fourier transforms, the visibility function equivalently writes:

$$V(u, v, w) = \left[V(\cdot, \cdot, 0) * \tilde{G}^{(w)}\right](u, v), \qquad (3.25)$$

with $G^{(w)}(l,m) \coloneqq G(l,m,w)$ (therefore, $\tilde{G}^{(w)}$ is the two-dimensional Fourier transform of G with respect to (l,m) with a fixed w). Under the small angle approximation, the following analytical form of $\tilde{G}^{(w)}$ can be derived: $\tilde{G}^{(w)}(u,v) = i/w e^{-\pi i (u^2 + v^2)/w}$ [Cornwell et al., 2008].

Consequently, the visibilities for any w can be deduced from those at w = 0 (which are the Fourier transform of the sought-after brightness) by a convolution with a known kernel \tilde{G} . This is very advantageous in iterative schemes where the data must be estimated with the observation model, as is the case with CLEAN. Indeed, rather than performing one Fourier transform per sampled w-value, which is time consuming, a single Fourier transform is performed, and then the w-modulations are applied.

Facet-based The field of view is decomposed into small facets, tangent to the sphere, in which the *w*-term is approximately constant [Cornwell and Perley, 1992]. Each facet has a locale dirty beam, which can be used for deconvolution. The main issue of this method concerns the association of the facets for the estimation of the residual in the CLEAN procedure, which requires a costly reprojection.

To address this, [Kogan and Greisen, 2009] proposes a coplanar faceting, which exploits the w-projection technique. However, corrections of overlapping areas between facets still must be dealt with.

w-stacking The gist of *w*-stacking is to group the sampled visibilities by close *w*-values and affect them the same *w*-value. The brightness can then be derived by (i) calculating the inverse Fourier transform of the visibilities for each group of *w*-values, (ii) applying the inverse *w*-effect $(1/G^{(w)})$ and (iii) associating the obtained images [Humphreys and Cornwell, 2011, Offringa et al., 2014]. This technique, associated to CLEAN and coined WSCLEAN for *w*-stacking CLEAN, is the one mainly used in scientific applications, as it is particularly fast compared to the aforementioned methods.

3.4 Tackling multi-frequency data

3.4.1 Multi-frequency interferometric deconvolution

Before investigating BSS with multi-wavelength radio-interferometric measurements, it is worthwhile exploring the literature on multi-channel deconvolution.

Measuring and imaging the radio sky with a view to performing blind 72 source separation

An interferometer samples the visibility space as a function of antenna distribution and wavelength. In the case of wideband measurements, a pair of antennas with baseline coordinates (d_u, d_v, d_w) will produce interferometric samples at $(u, v, w) = (d_u/\lambda, d_v/\lambda, d_w/\lambda)$ for each wavelength λ . There is thus a dilation effect in the sampling, which leads to an improvement of the brightness spatial-frequency coverage. The difficulty is that the brightness depends on the wavelength. However, some radio sources have particular spectral structures, for instance smooth (*e.g.*, power law) or known a priori. This can be advantageously used to synthesize brightness images from multi-wavelength data.

CLEAN-based extensions for multi-channel data are numerous, they are generally referred to as MFCLEAN, for multi-frequency CLEAN [Conway et al., 1990]. Extensions to tackle extended sources have been proposed [Rau, U. and Cornwell, T. J., 2011], as well as non-coplanar data [Offringa and Smirnov, 2017].

Research has also been conducted on optimization-based approaches. For instance [Abdulaziz et al., 2019] proposes a low-rank regularization to favor spatial correlations between channels.

3.4.2 Towards blind source separation

In the aim of performing BSS on radio data, we propose to summarize the acquisition models seen above. We take the opportunity to rename the variables to make them consistent with the specific notations of this manuscript. Three cases are distinguished as detailed below and summarized in Table 3.1.

Euclidean/coplanar case The cases of the single-dish telescope and the coplanar interferometer can both be combined in the same framework for small fields of view. The model, expressed directly in the Fourier/visibility space, writes:

$$\tilde{y}^{(\nu)}(u,v) = \tilde{H}^{(\nu)}(u,v) \ \tilde{x}^{(\nu)}(u,v), \tag{3.26}$$

with $\tilde{y}^{(\nu)}$ the multispectral data ($\tilde{y}^{(\nu)} = \tilde{W}^{(\nu)}$ for a single-dish telescope, $\tilde{y}^{(\nu)} = V_S^{(\nu)}$ for a coplanar interferometer), $\tilde{H}^{(\nu)}$ the measurement operator ($\tilde{H}^{(\nu)} = \tilde{A}^{(\nu)}$ for a single-dish telescope, $\tilde{H}^{(\nu)} = S^{(\nu)}$ for a coplanar interferometer) and $\tilde{x}^{(\nu)}$ the sought-after mixtures ($\tilde{x}^{(\nu)} = \tilde{B}^{(\nu)}$).

Spherical case This concerns the case of the single-dish telescope for wide-field observations. Similarly to the Euclidean case, the model is:

$$\tilde{y}^{(\nu)}(\mathfrak{l},\mathfrak{m}) = \tilde{H}^{(\nu)}(\mathfrak{l}) \ \tilde{x}^{(\nu)}(\mathfrak{l},\mathfrak{m}), \tag{3.27}$$

with the same correspondences as the above single-dish telescope case.

		Wide field of view	Small field of view	
Single-dish telecope		$\Big[\tilde{H}^{(\nu)}\odot\tilde{x}^{(\nu)}\Big](\mathfrak{l},\mathfrak{m})$	$\left[ilde{H}^{(u)}\odot ilde{x}^{(u)} ight](u,v)$	
Interferometer	Coplanar	not considered		
	Non-coplanar	not considered	$\left[\tilde{H}^{(\nu)} \odot \left(\mathcal{F} \circ \left[G^{(w)} \odot x^{(\nu)}\right]\right)\right](u, v, w)$	

Table 3.1: Summary of the considered radio telescope acquisition models. Wide-field interferometry, which uses the spherical formalism [Carozzi, 2015], is a relatively new field and will not be considered hereafter in the context of joint deconvolution and BSS.

Non-coplanar case For non-coplanar interferometric measurements, the considered model is:

$$\tilde{y}^{(\nu)}(u,v,w) = \tilde{H}^{(\nu)}(u,v,w) \ \left[\mathcal{F} \circ \left(G^{(w)} \odot x^{(\nu)} \right) \right](u,v,w), \tag{3.28}$$

with $G^{(w)}(l,m) \coloneqq e^{-2\pi i w \left(\sqrt{1-l^2-m^2}-1\right)}$ and the same correspondences as the coplanar interferometer case seen above.

Chapter 4 addresses the second case (and the first one by extension, as we will see). Chapter 5 tackles the third case.

Measuring and imaging the radio sky with a view to performing blind 74 source separation

Chapter 4

Joint deconvolution and blind source separation for wide-field observations

4.1 Introduction

Data acquired in radio astronomy are distorted by channel-dependent instrumental responses, whether they come from antenna beams in the case of single-dish telescopes or from incomplete samplings in the case of interferometers, as we saw in the previous chapter, Eq. (3.26). If the instrumental responses are assembled in a matrix $\mathbf{H} \in \mathbb{R}^{J \times P}$ (one operator per row, *i.e.*, channel), the linear mixture model becomes $\mathbf{Y}_{j:} = \mathbf{H}_{j:} * [\mathbf{AS}]_{j:} + \mathbf{N}_{j:}$ or more simply in the Fourier domain $\tilde{\mathbf{Y}}_{j:} = \tilde{\mathbf{H}}_{j:} \odot [\mathbf{AS}]_{j:} + \tilde{\mathbf{N}}_{j:}$. In this case, coping with now heterogeneous or incomplete data requires tackling an extra deconvolution step, thus leading to a joint deconvolution and blind source separation (DBSS) problem. A mathematically similar problem arises with other kind of astrophysical measurements, *e.g.*, the X-ray observations from the future European mission Athena¹, or compressive hyperspectral imaging [Golbabaee et al., 2013, Kobarg et al., 2014].

Being an ill-posed matrix factorization problem, BSS alone is already a challenging inverse problem. This is all-the-more complex when a channel-dependent operator further needs to be inverted as it can be ill-conditioned or even not invertible. Standard BSS methods cannot be employed directly unless the data \mathbf{Y} are pre-processed so as to obtain new data with a common resolution. However, jointly solving both deconvolution and separation is expected to yield much better results, allowing to more precisely account for the forward observation model and noise in a single pass.

¹https://sci.esa.int/web/athena

To the best of our knowledge, joint DBSS has been seldom investigated. The closest work known so far has been introduced by [Kleinsteuber and Shen, 2012], who proposed a BSS algorithm to analyze incomplete data in the framework of compressed sensing. This can be regarded as a special case of DBSS where the measurement operator is defined as a projection onto a low-dimensional measurement subspace. However, the proposed method is not compatible in our case since it only applies to compressively sensed measurements.

More recently, the first joint DBSS method was introduced [Jiang et al., 2017]. The proposed DecGMCA algorithm enforces the sparsity of the sources in a transformed domain \mathbf{W} by seeking a stationary point of the following cost function:

$$\operatorname{argmin}_{\mathbf{A},\mathbf{S}} \sum_{j=1}^{J} \frac{1}{2} \left\| \tilde{\mathbf{Y}}_{j:} - \tilde{\mathbf{H}}_{j:} \odot [\mathbf{A}\tilde{\mathbf{S}}]_{j:} \right\|_{2}^{2} + \left\| \mathbf{\Lambda} \odot \left(\mathbf{S} \mathbf{W}^{\top} \right) \right\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}) + \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}), \quad (4.1)$$

where the same notations as Chapter 1 are employed. DecGMCA builds upon the GMCA framework. However, in the setting of DBSS, resorting to a projected ALS optimization scheme also raises a major difficulty: The least-square problem with respect to \mathbf{S} is generally ill-conditioned, if not ill-posed, and needs to be regularized, which has a significant impact on the quality of the separation results.

Contributions In this chapter, we investigate a new joint DBSS algorithm to analyze data that are sampled on the sphere. This is essential to cope with the kind of wide-field spherical data, which are now common in scientific fields such as astronomy. In contrast to the standard case, analyzing spherical data raises extra difficulties due to the high computational cost of their manipulation, which makes the design of a computationally efficient and reliable algorithm essential. Therefore, we first aim to extend the algorithm DecGMCA [Jiang et al., 2017] to tackle joint deconvolution and separation problems from spherical data. As described in Section 4.2, the method is based on a projected alternating least-square minimization in order to combine rapidity and precision. Compared to the BSS problem, the procedure calls for an extra regularization to deal with a naturally ill-conditioned, if not ill-posed, problem. Yet, the regularization strategy and its impact on the solution have not been examined. For that purpose, we introduce several regularization strategies in Subsection 4.2.3.1 that significantly improve the separation quality. Based on these results, we introduce in Section 4.3 a new algorithm, coined SDecGMCA, to tackle efficiently joint deconvolution and blind source separation problems. Finally, in Section 4.4, numerical experiments, which involve both toy examples and realistic astrophysical simulations, are presented.

We also provide in Appendix C the equations of SDecGMCA adapted for planar data,

which is the case studied in [Jiang et al., 2017]. Indeed, there are noticeable parallels with the spherical case that can help for interpretation.

4.2 The spherical DecGMCA algorithm

In this section, we first adapt the DBSS method proposed in [Jiang et al., 2017] to tackle spherical data. Furthermore, we investigate in depth the ad hoc, yet necessary, regularization procedure in the update of the sources used in the DecGMCA algorithm. We show that it has a significant impact on the reconstruction quality of the sources. We then introduce two new regularization strategies, which noticeably outperform the one of DecGMCA.

4.2.1 Toward a DBSS method for spherical data

In the following, the multichannel data $\mathbf{Y} \in \mathbb{R}^{J \times P}$ are supposed to be sampled on the sphere with HEALPix [Górski et al., 2005] (see Appendix B for mathematical derivations related to signal analysis on the sphere). Moreover, the measurement operators are assumed linear and isotropic. According to observation model in the spherical case of the previous chapter, if the measurement operators are stacked in a matrix $\mathbf{H} \in \mathbb{R}^{J \times P}$, the mixture model writes:

$$\mathbf{Y}_{j:} = \left[\mathbf{AS}\right]_{j:} * \mathbf{H}_{j:} + \mathbf{N}_{j:}, \tag{4.2}$$

where * denotes the isotropic convolution operator on the sphere [Driscoll and Healy, 1994].

Quite similarly to the Euclidean case, the isotropic convolution product is shown to simplify in the spherical harmonic domain:

$$\tilde{\mathbf{Y}}_{j(\mathfrak{l},\mathfrak{m})} = \tilde{\mathbf{H}}_{j(\mathfrak{l},\mathfrak{m})}[\mathbf{A}\tilde{\mathbf{S}}]_{j(\mathfrak{l},\mathfrak{m})} + \tilde{\mathbf{N}}_{j(\mathfrak{l},\mathfrak{m})}, \qquad (4.3)$$

where the subscript $(\mathfrak{l}, \mathfrak{m})$ refers to the column of the corresponding matrix with degree \mathfrak{l} and mode \mathfrak{m} . Since the convolution kernels are isotropic, $\tilde{\mathbf{H}}$ does not depend on the mode \mathfrak{m} ; we will remove it hereafter, slightly abusing the notation. Moreover, the harmonic projections are performed up to a degree \mathfrak{l}_{max}^2 ; consequently, $\tilde{\mathbf{Y}}$, $\tilde{\mathbf{H}}$ and $\tilde{\mathbf{N}}$ are of size $J \times (\mathfrak{l}_{max} + 1)^2$, and $\tilde{\mathbf{S}}$ is of size $I \times (\mathfrak{l}_{max} + 1)^2$. The mixing model is rewritten by grouping the channels, yielding:

$$\tilde{\mathbf{Y}}_{:(\mathfrak{l},\mathfrak{m})} = \operatorname{diag}\left(\tilde{\mathbf{H}}_{:\mathfrak{l}}\right) \mathbf{A}\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} + \tilde{\mathbf{N}}_{:(\mathfrak{l},\mathfrak{m})}, \qquad (4.4)$$

²According to the usage of HEALPix, $l_{max} \leq 3N_{side} - 1$, where N_{side} is the resolution parameter of the spherical sampling.

where diag $\left(\tilde{\mathbf{H}}_{:\mathfrak{l}}\right) \in \mathbb{C}^{J \times J}$, in agreement with the notations of the manuscript.

To build an estimator for the sources S and the mixing matrix A, the same assumptions are made as in Chapter 1 regarding SBSS. These are briefly recalled and adapted to the new mixing model:

• *Data-fidelity term:* The noise is considered as Gaussian; therefore, the data-fidelity term writes:

$$\mathcal{D}(\mathbf{Y}, \mathbf{AS}) = \sum_{(\mathfrak{l}, \mathfrak{m}) \in \mathcal{P}} \frac{1}{2} \left\| \tilde{\mathbf{Y}}_{:(\mathfrak{l}, \mathfrak{m})} - \operatorname{diag}\left(\tilde{\mathbf{H}}_{:\mathfrak{l}} \right) \mathbf{A} \tilde{\mathbf{S}}_{:(\mathfrak{l}, \mathfrak{m})} \right\|_{2}^{2}, \quad (4.5)$$

with $\mathcal{P} \coloneqq \{(\mathfrak{l}, \mathfrak{m}) \in \mathbb{N} \times \mathbb{Z}, \mathfrak{l} \leq \mathfrak{l}_{max}, |\mathfrak{m}| \leq \mathfrak{l}\}$ the set of multipoles.

• Source regularization: The sources are assumed to be both sparse in a given dictionary $\mathbf{W} \in \mathbb{R}^{P' \times P}$ and nonnegative, hence the following regularization term on \mathbf{S} :

$$h_{\mathcal{S}}(\mathbf{S}) = \left\| \mathbf{\Lambda} \odot \left(\mathbf{S} \mathbf{W}^{\top} \right) \right\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}), \qquad (4.6)$$

where $\mathbf{\Lambda} \in \mathbb{R}^{I \times P'}$ is a matrix that stores the thresholding parameters and $\mathcal{O}_{\mathcal{S}} = \{ \mathbf{S} \in \mathbb{R}^{I \times P}, \forall (i, p), \mathbf{S}_{ip} \geq 0 \}.$

• Mixing matrix constraint: To mitigate the scale indeterminacy of the product \tilde{AS} , the columns of \tilde{A} are enforced to be in the unit ball. Imposing in addition the nonnegativity of the coefficients of the matrix yields the following constraint term:

$$h_{\mathcal{A}}(\mathbf{A}) = \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}), \tag{4.7}$$

where it is reminded that $\mathcal{K}_{\mathcal{A}} \coloneqq \{\mathbf{A} \in \mathbb{R}^{J \times I}, \forall (j, i), \mathbf{A}_{ji} \ge 0, \forall i, \|\mathbf{A}_{:i}\|_{2} \le 1\}.$

To sum up, the mixing matrix and the sources will be estimated by looking for a stationary point of the following cost function:

$$\operatorname{argmin}_{\mathbf{A},\tilde{\mathbf{S}}} \sum_{(\mathfrak{l},\mathfrak{m})\in\mathcal{P}} \frac{1}{2} \left\| \tilde{\mathbf{Y}}_{:(\mathfrak{l},\mathfrak{m})} - \operatorname{diag}\left(\tilde{\mathbf{H}}_{:\mathfrak{l}}\right) \mathbf{A}\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} \right\|_{2}^{2} + \left\| \mathbf{\Lambda} \odot \left(\tilde{\mathbf{S}}\mathbf{F}^{*}\mathbf{W}^{\top}\right) \right\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}}\left(\tilde{\mathbf{S}}\mathbf{F}^{*}\right) + \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}), \quad (4.8)$$

where \mathbf{F} is the spherical harmonic transform, in accordance with the notations of the manuscript.

BSS is already a challenging non-convex problem. When it comes to DBSS, a significant challenge is a need for a robust yet effective algorithm with a reasonable

computational burden. The above cost-function combines non-differentiable terms, such as the sparsity-enforcing ℓ_1 -regularization term or the characteristic functions. Therefore, designing a minimizer based on plain proximal algorithms would make perfect sense [Parikh and Boyd, 2014]. However, this would lead to algorithms with very high computational costs, especially since DBSS problems are generally ill-conditioned. Moreover, when BSS only is considered, tuning the regularization parameters, which has a major impact on the separation performances, is already particularly complex with standard proximal algorithms [Kervazo et al., 2020a].

In contrast, we have seen that projected alternating least squares (pALS) are particularly effective methods for providing fast SBSS algorithms. Furthermore, they come with almost automatic strategies to fix the regularization parameters robustly. We thus propose to use this framework for spherical DBSS.

4.2.2 Update of A

Updating the mixing matrix **A** when the sources **S** are fixed requires solving the least-square problem $\underset{\mathbf{A}}{\operatorname{argmin}} \mathcal{D}(\mathbf{Y}, \mathbf{AS})$, which yields for each channel $j \in [1, J]$:

$$\left(\mathbf{A}_{j:}\right)^{\top} \leftarrow \left(\sum_{(\mathfrak{l},\mathfrak{m})\in\mathcal{P}} \tilde{\mathbf{Y}}_{j(\mathfrak{l},\mathfrak{m})} \tilde{\mathbf{H}}_{j\mathfrak{l}}^{*} \left(\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})}\right)^{\dagger}\right) \left(\sum_{(\mathfrak{l},\mathfrak{m})\in\mathcal{P}} \left|\tilde{\mathbf{H}}_{j(\mathfrak{l},\mathfrak{m})}\right|^{2} \tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} \left(\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})}\right)^{\dagger}\right)^{-1}.$$
 (4.9)

In this equation, the estimate is defined with sums over all multipoles, which are significantly numerous than the sources. Consequently, the matrix on the right is generally invertible if not well conditioned.

In a second phase, the proximal operator of $h_{\mathcal{A}}$ is applied on the solution. As seen in Chapter 1, the update reads as follows, for each spectrum *i*:

$$\mathbf{A}_{:i} \leftarrow \frac{\max\left(\mathbf{A}_{:i}, \mathbf{0}\right)}{\max\left(1, \left\|\max\left(\mathbf{A}_{:i}, \mathbf{0}\right)\right\|_{2}\right)}.$$
(4.10)

4.2.3 Update of S

Updating \mathbf{S} assuming \mathbf{A} is fixed leads to the following optimization problem:

$$\operatorname{argmin}_{\tilde{\mathbf{S}}} \sum_{(\mathfrak{l},\mathfrak{m})\in\mathcal{P}} \frac{1}{2} \left\| \tilde{\mathbf{Y}}_{:(\mathfrak{l},\mathfrak{m})} - \operatorname{diag}\left(\tilde{\mathbf{H}}_{:\mathfrak{l}}\right) \mathbf{A}\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} \right\|_{2}^{2} + \left\| \mathbf{\Lambda} \odot \left(\tilde{\mathbf{S}}\mathbf{F}^{*}\mathbf{W}^{\top}\right) \right\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}}\left(\tilde{\mathbf{S}}\mathbf{F}^{*}\right).$$

$$(4.11)$$

According to the pALS minimization scheme, the solution of the above minimization problem is first approximated with a least-square estimate by finding a solution to

$$\underset{\tilde{\mathbf{S}}}{\operatorname{argmin}} \sum_{(\mathfrak{l},\mathfrak{m})\in\mathcal{P}} \frac{1}{2} \left\| \tilde{\mathbf{Y}}_{:(\mathfrak{l},\mathfrak{m})} - \operatorname{diag}\left(\tilde{\mathbf{H}}_{:\mathfrak{l}} \right) \mathbf{A} \tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} \right\|_{2}^{2},$$
(4.12)

which is given for each harmonic coefficient $(\mathfrak{l}, \mathfrak{m}) \in \mathcal{P}$ by:

$$\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} \leftarrow \mathbf{M}^{(\mathfrak{l})^{-1}} \mathbf{A}^{\top} \operatorname{diag}\left(\tilde{\mathbf{H}}_{:\mathfrak{l}}^{*}\right) \tilde{\mathbf{Y}}_{:(\mathfrak{l},\mathfrak{m})},\tag{4.13}$$

with $\mathbf{M}^{(l)} \coloneqq \mathbf{A}^{\top} \operatorname{diag} \left(|\tilde{\mathbf{H}}_{:l}|^2 \right) \mathbf{A}.$

In contrast to standard BSS problems, it is essential to notice that the above solution is not necessarily stable with respect to noise as the matrices $\mathbf{M}^{(l)}$ may be ill-conditioned. Additionally, it is not always unique as $\mathbf{M}^{(l)}$ may not be invertible; this might occur when the convolution kernels vanish for some spherical harmonics multipoles.

Borrowing ideas from sparsity enforcing deconvolution methods such as ForWard [Neelamani et al., 2004], [Jiang et al., 2017] proposes adding an extra Tikhonov regularization [Bertero et al., 2021] to the least-square problem:

$$\underset{\tilde{\mathbf{S}}}{\operatorname{argmin}} \sum_{(\mathfrak{l},\mathfrak{m})\in\mathcal{P}} \frac{1}{2} \left(\left\| \tilde{\mathbf{Y}}_{:(\mathfrak{l},\mathfrak{m})} - \operatorname{diag}\left(\tilde{\mathbf{H}}_{:\mathfrak{l}}\right) \mathbf{A} \tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} \right\|_{2}^{2} + \sum_{i=1}^{I} \varepsilon^{(i,\mathfrak{l})} \left\| \tilde{\mathbf{S}}_{i(\mathfrak{l},\mathfrak{m})} \right\|^{2} \right), \quad (4.14)$$

The set $\{\varepsilon^{(i,\mathfrak{l})} \geq 0, i \in [1...I], \mathfrak{l} \in \mathbb{N}, \mathfrak{l} \leq \mathfrak{l}_{max}\}$ are the Tikhonov regularization coefficients. They depend on the degree \mathfrak{l} and on the source *i*. For all $(\mathfrak{l}, \mathfrak{m}) \in \mathcal{P}$, updating the sources can now be recast as:

$$\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} \leftarrow \left(\mathbf{M}^{(\mathfrak{l})} + \underset{i \in [1...I]}{\operatorname{diag}} \left(\varepsilon^{(i,\mathfrak{l})}\right)\right)^{-1} \mathbf{A}^{\top} \operatorname{diag} \left(\tilde{\mathbf{H}}_{:\mathfrak{l}}^{*}\right) \tilde{\mathbf{Y}}_{:(\mathfrak{l},\mathfrak{m})}.$$
(4.15)

As with standard deconvolution problems, how the Tikhonov regularization coefficients are fixed dramatically impacts the quality of the regularized least-square solution and, eventually, the whole separation process. This is discussed in depth in the following section.

In a second step, the proximal operator of $h_{\mathcal{S}}$ is applied to the above least-square solution. According to Chapter 1, it can be approximated by:

$$\tilde{\mathbf{S}} \leftarrow \Pi_{\mathcal{O}_{\mathcal{S}}} \left(\mathcal{T}_{\mathbf{\Lambda}} \left(\tilde{\mathbf{S}} \mathbf{F}^* \mathbf{W}^\top \right) \mathbf{M}^\top \right) \mathbf{F}^\top$$
 (4.16)

where we remind $[\mathcal{T}_{\Lambda}(\mathbf{X})]_{ip} = \operatorname{sign}(\mathbf{X}_{ip}) \max(0, |\mathbf{X}_{ip}| - \Lambda_{ip})$ is the soft-threshold operator, **M** verifies $\mathbf{MW} = \mathbf{I}$ and $[\Pi_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S})]_{ip} = \max(0, \mathbf{S}_{ip}).$

4.2.3.1 Regularization strategies

In [Jiang et al., 2017], the regularization strategy is defined somewhat arbitrarily with hyperparameters fixed to an ad hoc small value (e.g., 1e - 3). In this section, we will

highlight that this regularization has a significant impact on the estimation accuracy. We then study different strategies to tune these critical parameters in a more efficient and adaptive way.

Let c be a positive number, that will be called regularization hyperparameter. We will further investigate four different regularization strategies:

• Strategy #1 Let us first consider the naive strategy where the regularization parameters are chosen independently of the source i and the frequency l:

$$\varepsilon^{(i,\mathfrak{l})} = c. \tag{4.17}$$

• Strategy #2 The strategy presented in [Jiang et al., 2017] is also considered:

$$\varepsilon^{(i,\mathfrak{l})} = c\,\lambda_{\max}(\mathbf{M}^{(\mathfrak{l})}),\tag{4.18}$$

where $\lambda_{\max}(\cdot)$ returns the greatest eigenvalue. In [Jiang et al., 2017], the main motivation of this choice was to set regularization parameters that scale with the data.

• Strategy #3 In Eq. (4.13), the errors that contaminate the observations are amplified by the inverse of the smallest eigenvalue of $\mathbf{M}^{(l)}$, denoted $\lambda_{\min}(\mathbf{M}^{(l)})^{-1}$. Limiting the noise amplification to c amounts to choosing $\varepsilon^{(i,l)}$ such that $\varepsilon^{(i,l)} + \lambda_{\min}(\mathbf{M}^{(l)}) \leq c \iff \varepsilon^{(i,l)} \leq c - \lambda_{\min}(\mathbf{M}^{(l)})$. Bearing in mind that $\varepsilon^{(i,l)} \geq 0$, it is possible to set $\varepsilon^{(i,l)} = \max(0, c - \lambda_{\min}(\mathbf{M}^{(l)}))$. A change of variable is finally operated to facilitate the interpretation of the hyperparameter c, yielding for each source i and frequency l:

$$\varepsilon^{(i,\mathfrak{l})} = \max\left(0, c - \frac{\lambda_{\min}(\mathbf{M}^{(\mathfrak{l})})}{\lambda_{\min}(\mathbf{A}^{\top}\mathbf{A}) + \epsilon}\right),\tag{4.19}$$

with $\epsilon = 1e - 2$, to prevent numerical issues. Since the sequence $(\lambda_{\min}(\mathbf{M}^{(l)}))_{l \in \mathbb{N}}$ decreases, the sequence $(\varepsilon^{(i,l)})_{l \in \mathbb{N}}$ increases. Consequently, the higher frequencies are more penalized, while the lower ones are preserved. This is advantageous because most of the information from physical sources is generally in low spatial frequencies. Moreover, the maximum operator in (4.19) allows to have $\varepsilon^{(i,l)} = 0$ for smaller frequencies, provided that c is small enough; this allows to keep the smaller frequencies unbiased.

Strategy #4 As in the Euclidean case, it is possible to derive an angular power spectrum from a harmonic decomposition, which describes the power distribution along the degrees l (see Appendix B). If the angular power spectra of the sources {S_i:S_i} and the noise S_{nn} (which is assumed observation independent, hence the notation with n) are known, the strategy which

minimizes the mean square error of the unpenalized least-square problem is $\varepsilon^{(i,l)} = \mathfrak{S}_{\mathbf{nnl}} / \mathfrak{S}_{\mathbf{S}_{i:}\mathbf{S}_{i:l}}$. This is reminiscent of a Wiener deconvolution filter. Yet, it has to be reminded that the update of **S** is further followed by a thresholding step, which likely alters the properties of the regularization strategy. Therefore, the regularization strategy is adapted by adding a hyperparameter c, yielding for each source i and degree \mathfrak{l} :

$$\varepsilon^{(i,\mathfrak{l})} = c \, \frac{\mathfrak{S}_{\mathbf{n}\mathfrak{l}}}{\mathfrak{S}_{\mathbf{s}_i:\mathbf{S}_{i:l}}}.\tag{4.20}$$

The hyperparameter c tips the balance between the Tikhonov regularization and the sparsity regularization.

4.2.3.2 Numerical comparisons of the regularization strategies

In this paragraph, we propose illustrating the impact of the above regularization strategies with numerical experiments. For this purpose, we consider the non-blind source separation case, *i.e.*, using the ground truth mixing matrix \mathbf{A}^{\diamond} .

Description of the data The data are generated as follows:

- the sources **S** are sampled on the sphere using the HEALPix pixelization, with parameters $N_{side} = 128$ and $\mathfrak{l}_{max} = 384$. They are random nonnegative signals that are sparse in the spherical starlet (isotropic undecimated wavelets) domain [Starck et al., 2005], and further band-limited to a cut-off frequency $\mathfrak{l}_{max}/6 = 64$. See example in Fig. 4.5a.
- the mixing matrices **A** are random nonnegative matrices with a given condition number.
- the convolution kernels $\tilde{\mathbf{H}}_{j\mathfrak{l}} = \exp\left(-\frac{\mathfrak{l}(\mathfrak{l}+1)}{r(j)(r(j)+1)}\log 2\right)$ are Gaussian-shaped, with resolutions r evenly spread between the minimum resolution r_{min} , which is a parameter to be set, and \mathfrak{l}_{max} . The resolution is defined as the full width at half maximum (FWHM) in the spherical harmonic domain of the convolution kernel (see example in Fig. 4.1).

Throughout this section, we consider I = 4 sources. Moreover, we define four parameters that characterize the observations:

- number of observations J (J = 8 by default),
- mixing matrix condition number $cond(\mathbf{A})$ ($cond(\mathbf{A}) = 2$ by default),
- minimum resolution r_{min} ($\mathfrak{l}_{max}/8 = 48$ by default),
- overall signal-to-noise ratio (10 dB by default).



Figure 4.1: Example of a Gaussian filter $\mathbf{H}_{i:}$, with a resolution r = 144.

Method The joint deconvolution and non-blind separation is performed by applying the source update proposed in Subsection 4.2.3 with the ground truth mixing matrix (*i.e.*, a least-square update and a proximal update). It is noted that the observations are unmixed and deconvolved at the resolution of the best-resolved channel, which amounts to replacing $\tilde{\mathbf{H}}_{jl}$ by $\tilde{\mathbf{H}}_{jl}/\tilde{\mathbf{H}}_{jbl}$ with j_b the index of the best-resolved channel. Moreover, for strategy #4, the regularization parameters are calculated with the angular power spectra of the ground truth sources. The thresholds are set in the same way than during the finale iteration of SDecGMCA, which is discussed in the following section.

We test how the four regularization strategies we defined in the previous section behave with respect to the number of observations, the mixing matrix condition number, the minimum resolution of the convolution kernel and the SNR. Each experiment is composed of 100 trials, with varying random sources, mixing matrices and noise realizations. For each point, the optimal regularization hyperparameter is fine-tuned using a grid-search. The performances of each strategy is evaluated based on the NMSE of the estimated sources.

Results Figure 4.2 shows the mean NMSE as a function of each of the four aforementioned observation parameters. Strategy #4 provides the best reconstruction qualities in all observation scenarios. Among the other strategies, that do not assume the sources to be known, strategy #3 achieves better results. It is mostly thanks to the non-linear maximum operator, which allows keeping the lower frequencies unbiased, where most of the sources energy is located (see example Fig. 4.3, where $\varepsilon^{(i,1)} = 0$ for $\mathfrak{l} \leq 44$). Strategy #2 gives poor results; indeed, it biases more significantly the lower frequencies than the higher ones.

Table 4.1 shows the range of variation of the mean optimal regularization hyperpa-




Figure 4.2: Mean NMSE over 100 realizations as a function of observation parameters.

rameters when the observation parameters vary. Contrary to the three first strategies, the optimal regularization hyperparameter for strategy #4 is rather insensitive to the observation parameters (typically $c_{opt} \sim 0.5$). The noticeable exception is regarding the SNR. When the SNR is low, the denoising by the sparsity regularization is particularly promoted, hence a marked decrease of c_{opt} .

4.3 Implementation of the SDecGMCA algorithm

We shed light on how critical the choice of the Tikhonov regularization strategy is. In this section, the details of implementation of SDecGMCA are described. In particular, a novel two-stage algorithm which exploits the proposed regularization strategies is introduced.

SDecGMCA is summarized in Algorithm 10. The mixing matrix \mathbf{A} is initialized with a Principal Component Analysis (PCA) performed on the observations. For this step only, the data are first reconvolved so as they share a common resolution; to



Figure 4.3: Example of regularization parameters $\varepsilon^{(i,l)}$ as a function of the frequency \mathfrak{l} for the default observation parameters and with the mean optimal hyperparameters.

Parameter	Range	Regularization strategy			
1 difuiiletei	1001150	#1	#2	#3	#4
$cond(\mathbf{A})$	1.5 to 14	1.04	1.84	0.31	0.10
r_{min}	2 to 350	> 3.30	> 2.60	0.75	0.11
J	4 to 24	0.49	2.02	1.11	0.25
SNR (dB)	-10 to 40	1.18	1.08	0.41	0.76

Table 4.1: Range of variation of the mean optimal regularization hyperparameter c_{opt} , in terms of order of magnitude, when the observation parameters vary. In some cases, c_{opt} is smaller than the lower bound of the search interval (1e - 5), hence the lower limit of the range, which is indicated by the sign ">".

avoid noise amplification, this resolution is that of the worse-resolved channel. SDecGMCA is comprised of two stages. The first stage estimates a first guess of the mixing matrix and the sources (**warm-up**); it is required to provide robustness with respect to the initial point. The second stage refines the separation by employing a more precise regularization strategy (**refinement**). The values of the parameters of SDecGMCA according to the stage are summarized in Table 4.2.

Algorithm 10 SDecGMCA

Inputs: data \mathbf{Y} , masks $\tilde{\mathbf{H}}$, number of sources I, regularization hyperparameters c_{wu} and c_{ref} , variance of the noise σ^2 , sparsifying representation (analysis operator \mathbf{W} and synthesis operator \mathbf{M}), thresholding parameters κ and K_{max}

 $\mathbf{A} \leftarrow \mathrm{PCA}(\mathbf{Y}, I)$

(1) Warm-up stage stage \leftarrow warm-up while convergence not reached do Update c and K according to Table 4.2 $\tilde{\mathbf{S}} \leftarrow$ UpdateS (stage, $\tilde{\mathbf{Y}}$, $\tilde{\mathbf{H}}$, \mathbf{A} , c, σ^2 , \mathbf{W} , \mathbf{M} , κ , K) $\mathbf{A} \leftarrow$ UpdateA ($\tilde{\mathbf{Y}}$, $\tilde{\mathbf{H}}$, $\tilde{\mathbf{S}}$) (2) Refinement stage stage \leftarrow refinement while convergence not reached do $\tilde{\mathbf{A}} = \mathbf{W} + \tilde{\mathbf{S}}$ ($\mathbf{A} = \tilde{\mathbf{A}} = \tilde{\mathbf{X}} = \tilde{\mathbf{W}} + \mathbf{A} = \mathbf{A}$

$$\tilde{\mathbf{S}} \leftarrow \text{UpdateS}\left(\text{stage, } \tilde{\mathbf{Y}}, \ \tilde{\mathbf{H}}, \ \mathbf{A}, \ c_{ref}, \ \sigma^2, \ \mathbf{W}, \ \mathbf{M}, \ \kappa, \ K_{max} \right)$$

 $\mathbf{A} \leftarrow \text{UpdateA}\left(\tilde{\mathbf{Y}}, \ \tilde{\mathbf{H}}, \ \tilde{\mathbf{S}}\right)$

return A, S

4.3.1 Choice of the thresholding parameters

As emphasized in Chapter 1, the thresholding parameters play a central role for the robustness of the minimization scheme with respect to the spurious local minima and the precision of the estimates. In this respect, the same threshold selection strategies as GMCA are implemented in SDecGMCA, namely:

• the thresholds are adjusted with respect to the noise level contaminating the current estimation of the sources (" $\kappa - \sigma$ " rule, see Eq. 1.22),

86

Algorithm 11 UpdateS

Inputs: stage, $\tilde{\mathbf{Y}}$, $\tilde{\mathbf{H}}$, \mathbf{A} , c, σ^2 , \mathbf{W} , \mathbf{M} , κ , K

- (1) Least-square update
- for $\mathfrak{l} = 0, 1, \dots, \mathfrak{l}_{max}$ and $m = -\mathfrak{l}, -\mathfrak{l} + 1, \dots, \mathfrak{l}$ do

Determine the Tikhonov regularization parameters $\{\varepsilon^{(i,\mathfrak{l})}\}_i$ according to the stage and c

$$\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} \leftarrow \left(\mathbf{A}^{\top} \operatorname{diag}\left(|\tilde{\mathbf{H}}_{:\mathfrak{l}}|^{2}\right) \mathbf{A} + \operatorname{diag}_{i \in [1...I]}\left(\varepsilon^{(i,\mathfrak{l})}\right)\right)^{-1} \mathbf{A}^{\top} \operatorname{diag}\left(\tilde{\mathbf{H}}_{:\mathfrak{l}}^{*}\right) \tilde{\mathbf{Y}}_{:(\mathfrak{l},\mathfrak{m})}$$

(2) Regularization

Calculate the thresholds Λ with a support-based strategy and a ℓ_1 -reweighting if need be, according to the stage

if stage is warm-up then

$$\begin{split} \tilde{\mathbf{S}} &\leftarrow \mathcal{T}_{\mathbf{\Lambda}} \Big(\tilde{\mathbf{S}} \mathbf{F}^* \mathbf{W}^\top \Big) \, \mathbf{M}^\top \mathbf{F}^\top \\ \mathbf{else} \\ \tilde{\mathbf{S}} &\leftarrow \max \left(\mathcal{T}_{\mathbf{\Lambda}} \Big(\tilde{\mathbf{S}} \mathbf{F}^* \mathbf{W}^\top \Big) \, \mathbf{M}^\top, \mathbf{0} \Big) \, \mathbf{F}^\top \end{split}$$

return $\tilde{\mathbf{S}}$

(1) Least-square update
for
$$j = 1, 2, ..., J$$
 do
$$\mathbf{A}_{j:} \leftarrow \left(\sum_{(\mathfrak{l},\mathfrak{m})\in\mathcal{P}} \tilde{\mathbf{Y}}_{j(\mathfrak{l},\mathfrak{m})} \tilde{\mathbf{H}}_{j\mathfrak{l}}^{*} \left(\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})}\right)^{\dagger}\right) \left(\sum_{(\mathfrak{l},\mathfrak{m})\in\mathcal{P}} \left|\tilde{\mathbf{H}}_{j(\mathfrak{l},\mathfrak{m})}\right|^{2} \tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})} \left(\tilde{\mathbf{S}}_{:(\mathfrak{l},\mathfrak{m})}\right)^{\dagger}\right)^{-1}$$

(2) Constraint application for i = 1, ..., I do $\mathbf{A}_{:i} \leftarrow \max(\mathbf{A}_{:i}, \mathbf{0})$ $\mathbf{A}_{:i} \leftarrow \mathbf{A}_{:i} / \max(\|\mathbf{A}_{:i}\|_{2}, 1)$

return A

Stage	1: warm-up	2: refinement
Regularization strategy	#3	#4
Regularization hyperparameter c	$10 c_{wu} \rightarrow c_{wu}$	c_{ref}
Active source support K	$0 \to K_{max}$	K_{max}
ℓ_1 -reweighting	No	Yes
Nonnegativity constraint on ${f S}$	No	Yes

Joint deconvolution and blind source separation for wide-field observations

Table 4.2: Parameters	of SDecGMCA	for each	stage.
-----------------------	-------------	----------	--------

- the thresholds are decreased along the iterations to avoid inaccurate local minima (see Eq. 1.23),
- the thresholds are adapted sample-wise with a ℓ_1 -reweighting scheme so as to reduce thresholding artifacts [Candès et al., 2008] (see Eq. 1.24).

4.3.2 A two-stage minimization approach

We highlighted in Section 4.2.3.2 that regularization strategy #4 provides significantly better results in the non-blind source separation case. Since, this strategy is defined based on some estimate of the sources, we propose proceeding with a two-step approach: (i) a warm-up step, whose goal is to provide a quick rough estimate of **A** and **S**, with increased robustness with respect to the initialization and spurious local minima, and (ii) a refinement step that makes use of the more precise regularization strategy #4.

Warm-up stage During the warm-up, regularization strategy #3 is employed. Indeed, as seen in Section 4.2.3.2, it provides the best results among the three first regularization strategies that require no additional information about the sources. During this stage, the estimates of the sources are likely to be dominated by the artifacts due to the noise amplification, especially if the regularization hyperparameter at warm-up c_{wu} is too small. The estimation of **A** is then erroneous, and the algorithm might get stuck in a local minimum. This phenomenon can be alleviated by over-regularizing the sources during the first iterations, *i.e.*, by taking a higher c_{wu} [Jiang et al., 2017]. In doing so, the noise contamination is reduced, and only the major features of the sources are kept. In the spirit of [Jiang et al., 2017], the regularization hyperparameter is then progressively decreased along the iterations to the input value, in order to refine the estimations of **S** and consequently **A**. The decrease of the warm-up regularization hyperparameter markedly improves the robustness of the separation in terms of convergence. Moreover, the earlier mentioned support-based strategy is achieved during the warmup. Since the starting point of this step is likely to be quite far from the sought-after sources, no reweighted ℓ_1 is applied as it would tend to favor spurious solutions. For the same reason, the nonnegativity constraint is likewise not applied to the sources during this stage (but it is for **A**).

The warm-up ends when the decrease of c_{wu} and the increase of K are completed, and when the estimations of the sources have converged, that is when $||\tilde{\mathbf{S}}^{(n)} - \tilde{\mathbf{S}}^{(n-1)}||_2/||\tilde{\mathbf{S}}^{(n)}||_2 \leq \epsilon_{wu}$. In practice, ϵ_{wu} does not require to be very small $(e.g., \epsilon_{wu} = 1e - 2)$.

Refinement stage During the second stage, the estimations are refined by using the more precise regularization strategy #4 introduced in 4.2.3.1. The regularization parameters are calculated with the angular power spectra of the sources estimated at the previous iteration, which are assumed to be close enough to the ground truth ones:

$$\varepsilon^{(i,\mathfrak{l})^{(n)}} = c_{ref} \frac{\mathfrak{S}_{\mathbf{nn}\mathfrak{l}}}{\mathfrak{S}_{i:}^{(n-1)} \mathbf{S}_{i:}^{(n-1)}}, \qquad (4.21)$$

where *n* is an iteration and c_{ref} is the regularization hyperparameter at refinement. \mathfrak{S}_{nn} is deduced from the SNR of the observations. Concerning the choice of the thresholds, the active support is kept constant at K_{max} . The ℓ_1 -reweighting of the sources we described above is applied during this stage, as well as the nonnegativity constraint on the sources. The refinement ends when the estimations of the sources have converged, that is when $||\tilde{\mathbf{S}}^{(n)} - \tilde{\mathbf{S}}^{(n-1)}||_2/||\tilde{\mathbf{S}}^{(n)}||_2 \leq \epsilon_{ref}$ (for instance $\epsilon_{ref} = 1e - 5$).

Setting K_{max} to a value different from 1 generally improves the separation, but it biases the estimation of the sources. Therefore, after the refinement step, a final estimation of the sources with **A** fixed and K = 1 is performed. This final step does not appear in Algorithm 10 for clarity.

4.3.3 Convergence proprieties

DBSS requires solving a multi-convex optimization problem. No method can guarantee finding in general the global minimum. At best, it is possible to guarantee a convergence toward a local minimum (*e.g.*, BCD [Tseng, 2001], PALM [Bolte et al., 2014]). SDecGMCA is built upon a projected ALS scheme. To the best of our knowledge, is has not been theoretically proved that projected ALS algorithms converge. However, we empirically show that the regularization parameters tend to stabilize along the iteration, as well as the estimates of **A** and **S** (see Fig. 4.4).



Figure 4.4: Illustration of the empirical stabilization. The relative variation is equal to $\|\mathbf{X}^{(n)} - \mathbf{X}^{(n-1)}\|_2 / \|\mathbf{X}^{(n)}\|_2$, with $\mathbf{X} = \mathbf{A}$ or \mathbf{S} depending on the case. The peak at iteration #101 corresponds to the switch from the warm-up stage to the refinement stage.

4.4 Numerical experiments

In this section, we investigate the performances of the SDecGMCA algorithm on synthetic toy examples that allow performing Monte-Carlo trials, as well as realistic simulations with partial sky coverage. The code that is used is open source (see Appendix D).

For all these experiments, we make use of the HEALPix pixelisation on the sphere.

Comparison criteria are first based on the NMSE, which measures the reconstruction quality of the sources. As well, we make use of the mixing matrix criterion C_A to assess the quality of the estimated mixing matrices. The mixing matrix criterion is more appropriate to compare DBSS and BSS methods with different source regularization, since it only depends on A.

We define the oracle as the solution of the non-blind problem (*i.e.*, knowing the ground truth mixing matrix) with regularization strategy #4 and the optimal regularization hyperparameter. It provides an upper bound of the NMSE that SDecGMCA can reach.

4.4.1 Toy model

The same toy examples as in Section 4.2.3.2 are employed, with a similar parameterization. The input parameters of SDecGMCA are summarized in Table 4.3.

Parameter	Notation	Value
Minimum number of iterations at warm-up	N_{wu}	100
Number of detail scales	S	3
Thresholding parameter	κ	3
Maximum source support	K_{max}	0.5

Table 4.3: Input parameters of SDecGMCA for the toy example

4.4.1.1 Impact of the Tikhonov regularization

In standard BSS, the common sources of error are interference (*i.e.*, leakage between sources due to a poor estimation of the mixing matrix), noise contamination (*i.e.*, back-projected noise from the observations which is not filtered) and artifacts (*i.e.*, remaining errors such as thresholding artifacts) [Vincent et al., 2006]. Within the scope of joint deconvolution and BSS, this straightforward decomposition is not suitable anymore. For instance, when the mixing matrix is known, the estimation of the sources alone can generate interference in addition to the deconvolution artifacts due to the biasing Tikhonov regularization. Generally speaking, we can consider that a separation is successful when deconvolution artifacts dominate the reconstruction errors.

In this subsection, the data have the default observation parameters (see Section 4.2.3.2). Figure 4.5 shows an example of solution by SDecGMCA where the reconstructions errors are dominated by deconvolution artifacts (see Fig. 4.5f).

We first assess the impact of the regularization hyperparameters on the performances of SDecGMCA. To this end, we execute SDecGMCA with different warm-up and refinement regularization hyperparameters. Since the sources are updated similarly in the non-blind problem, the mean optimal regularization hyperparameters found in Section 4.2.3.2 are taken as reference regularization hyperparameters. Each experiment is performed 100 times, with varying sources, mixing matrices and noise realizations.

The mean NMSE and the mean C_A are reported in Table 4.4. The choice of c_{wu} for the warm-up stage has little impact on the performance metrics; indeed, the loss of NMSE and C_A due to a poor choice of c_{wu} is, respectively, at most 0.42 dB and 0.77 dB in the tested range. On the contrary, the selection of c_{ref} for the refinement stage may be more critical. It is, however, very interesting to highlight that, in a range of one order of magnitude around the optimal hyperparameter, the loss of NMSE and C_A remains quite limited (*i.e.*, respectively about -2.21 dB and -1.05 dB at most). It is noted that the oracle mean NMSE is 25.74 dB; thus, the extra estimation of the mixing matrix is only the origin of a 0.95 dB loss in NMSE.



Figure 4.5: DBSS example by SDecGMCA on the toy model. The graphics of sources, observations and error are in logarithmic scale.

			c_{ref}	$(\times c_{reform})$	_{pt})	
		10^{-1}	$10^{-0.5}$	10 ⁰	$10^{0.5}$	10^{1}
c_{wu}	$\begin{array}{c} 10^{0} \rightarrow 10^{-1} \\ 10^{0.5} \rightarrow 10^{-0.5} \\ \mathbf{10^{1}} \rightarrow \mathbf{10^{0}} \\ 10^{1.5} \rightarrow 10^{0.5} \\ 10^{2} \rightarrow 10^{1} \end{array}$	22.86 22.99 23.00 23.06 23.06	24.43 24.58 24.58 24.59 24.23	24.61 24.79 24.79 24.65 24.4	22.6 22.83 22.82 22.44 22.43	18.11 18.35 18.34 18.10 <i>17.93</i>
$(\times c_{wuopt})$	$ \begin{array}{c} \hline 10^0 \rightarrow 10^{-1} \\ 10^{0.5} \rightarrow 10^{-0.5} \\ \hline 10^1 \rightarrow 10^0 \\ 10^{1.5} \rightarrow 10^{0.5} \\ 10^2 \rightarrow 10^1 \end{array} $	25.62 25.60 25.57 25.28 25.28	$25.20 \\ 25.19 \\ 25.15 \\ 24.80 \\ 24.56$	24.86 24.86 24.81 24.37 24.25	$24.41 \\ 24.46 \\ 24.39 \\ 23.75 \\ 23.84$	22.66 22.81 22.80 22.07 <i>22.04</i>

Table 4.4: Mean NMSE (top) and C_A (bottom) in dB, over 100 realizations, performed by SDecGMCA on the toy model as a function of c_{wu} and c_{ref} . These are given as multiples of c_{wuopt} and c_{refopt} , which are the mean optimal hyperparameters for the non-blind problem. It is noted that the oracle mean NMSE is 25.74 dB.

4.4.1.2 Comparisons with other blind source separation methods

In this paragraph, comparisons with other blind source separation methods are performed. Since few DBSS methods have been investigated so far, a natural comparison would be with DecGMCA. However, as it has not directly been designed for data sampled on the sphere, a direct comparison cannot be performed. We rather propose to substitute DecGCMA's regularization strategy #2 within SDecGMCA to quantify the impact of the regularization strategy. In contrast to [Jiang et al., 2017], where the regularization parameter is chosen in an ad hoc manner, we employ the optimal regularization hyperparameters; this method is therefore referred to as oDecGMCA (for optimized DecGMCA). Moreover, in order to highlight the benefit of combining deconvolution and BSS, we propose to compare SDecGMCA to three standard BSS methods:

- GMCA (with the decreasing thresholding strategy and the ℓ_1 -reweighting),
- Hierarchical alternating least squares (HALS) [Cichocki et al., 2007, Gillis and Glineur, 2012], which is recalled to be a nonnegative matrix factorization algorithm solving $\operatorname{argmin}_{\mathbf{A} \ge 0, \mathbf{S} \ge 0} \|\mathbf{X} \mathbf{AS}\|_2^2$ using a block coordinate descent with multiplicative updates,
- Beta sparse nonnegative matrix factorization (β -SNMF) [Cherni et al., 2020], which is a nonnegative matrix factorization algorithm promoting the sparsity

	C_A	$\mathrm{NMSE}_{\mathrm{w}}$	NMSE
SDecGMCA	24.81	27.08	24.79
oDecGMCA	23.01	20.94	15.03
GMCA	21.98	19.35	N/A
HALS	8.17	5.83	N/A
β -SNMF	8.43	7.01	N/A

Table 4.5: Mean performance metrics in dB, over 100 realizations, achieved by different algorithms on the toy model.

of **S** in the direct domain, which solves $\operatorname{argmin}_{\mathbf{A} \ge 0, \mathbf{S} \ge 0} \|\mathbf{X} - \mathbf{AS}\|_2^2 + \lambda \|\mathbf{S}\|_1$ by resorting also to block coordinate descent with multiplicative updates.

As these methods can only process observations with the same resolution, the observations are deconvolved beforehand to the resolution of the worse-resolved observation so as to avoid noise amplification. The NMSE is adapted to take account of the resolution loss:

$$\text{NMSE}_{w} = -10 \log_{10} \left(\frac{\|\mathbf{H}_{j_{w}:} * \mathbf{S}^{\diamond} - \mathbf{S}\|_{2}^{2}}{\|\mathbf{H}_{j_{w}:} * \mathbf{S}^{\diamond}\|_{2}^{2}} \right),$$
(4.22)

where j_w is the index of the worse-resolved observation channel. The NMSE_w of the DBSS methods can be calculated by deteriorating the estimated sources.

The performance metrics achieved by the different DBSS and BSS algorithms are reported in Table 4.5. Compared to oDecGMCA, SDecGMCA performs a significant gain in NMSE and a moderate increase in C_A . This result confirms that the choice of the regularization strategy is crucial for the estimation of the sources. Moreover, the BSS algorithms achieve poor results; indeed, high-frequency information, which is essential for the separation process, is lost during deconvolution.

4.4.1.3 Varying observation parameters

Let us evaluate the sensitivity of the different source separation algorithms to the observation parameters, that is the mixing matrix condition number, the minimum resolution of the convolution kernels, the number of observations and the SNR. For SDecGMCA and oDecGMCA, the mean optimal regularization hyperparameter found above with the non-blind problem are used. At each point, the algorithms are executed 30 times with varying sources, mixing matrices and noise realizations. The mean performance metrics are reported in Figure 4.6.

The NMSE achieved by SDecGMCA is close to the oracle; the loss is typically of 1 dB. In every scenario, SDecGMCA clearly outperforms oDecGMCA in terms of NMSE and C_A . Overall, the tendencies are consistent; the performance metrics

increase with increasing minimum resolution, number of observations and SNR, while decrease with increasing mixing matrix condition number. The first notable exception is the NMSE_w against the minimum resolution. It is due to the fact that the reference $\mathbf{H}_{jw} * \mathbf{S}$ in the definition of the NMSE_w varies from one point to the other. The second exception concerns the SNR; the performance metrics stabilize or decrease when there is little noise. This is an effect of the implicit regularization provided by the noise. According to the proposed threshold tuning strategy, when the noise level is low, the thresholds are low and the sparsity constraint is loosened. On the contrary, a higher noise level yields higher thresholds that tend to select high amplitude coefficients, which better discriminate between the sources.

Finally, let us assess the sensitivity of SDecGMCA to the regularization hyperparameter at refinement as a function of the observation parameters (see Fig. 4.7). Both NMSE and C_A losses are limited in a range of one order of magnitude around the optimal regularization hyperparameter (typically -2 dB). The noticeable exception is when **A** is ill-conditioned. The higher sensibility to the regularization hyperparameter may come from the induced ill-condition of the ($\mathbf{M}^{(l)}$)_{$l \in \mathbb{N}$} (matrices which are inverted in Eq. 4.15).

4.4.2 Application to realistic astrophysical data

Joint deconvolution and blind source separation is now performed on realistic astrophysical simulations. These data are composed of J = 25 observations, which are built as mixtures of three sources. The latter are associated with a synchrotron, a thermal and an iron line emissions as displayed in Fig. 4.8, with spectra displayed in panel (d). It is commonplace in such an application that only a partial sky coverage is observed, which is simulated by projecting the sources on a limited portion of the sphere (see Fig. 4.8). Similarly to the synthetic data, we set the pixelization parameters to $N_{side} = 128$ and $\mathfrak{l}_{max} = 384$. The 25 observations have resolutions evenly spread between $\mathfrak{l}_{max}/8$ and \mathfrak{l}_{max} .

4.4.2.1 Example of result

We firstly consider that the observations are corrupted by a noise of 10 dB (see two observations on Fig. 4.8e and 4.8f). Figure 4.9 shows a solution given by SDecGMCA (with $K_{max} = 0.2$ to overcome the correlations between the sources). The errors are dominated by deconvolution artifacts. Actually, the estimated sources are very close to the oracle estimates that use the ground truth mixing matrix (NMSE = 21.17 dB vs. oracle NMSE = 22.58 dB). Moreover, putting aside a small leakage in the lower frequencies of the emission line spectrum in the synchrotron spectrum, all three spectra are well reconstructed (C_A = 18.02 dB).



Figure 4.6: Mean performance metrics over more than 30 realizations as a function of observation parameters on the toy model. The default values of the observation parameters are summarized in Section 4.2.3.2.



Figure 4.7: Mean performance metrics over more 30 realizations performed by SDecGMCA on the toy model with the optimal regularization hyperparameter at warm-up c_{wuopt} and with different regularization hyperparameters at refinement c_{ref} , which are indicated as multiples of c_{refopt} in the legends.



chrotron and thermal emissions

(e) Low resolution observation $\mathbf{Y}_{3:}$, with syn- (f) High resolution observation $\mathbf{Y}_{20:}$, with mostly emission line source

Figure 4.8: Realistic data and example of observations with a SNR of 10 dB. The plots of sources and observations are in logarithmic scale.



Figure 4.9: DBSS example by SDecGMCA on realistic data. The plots of sources and errors are in logarithmic scale.

	C_A	$\mathrm{NMSE}_{\mathrm{w}}$	NMSE
SDecGMCA	18.02	26.92	21.17
oDecGMCA	15.78	25.89	18.16
GMCA	16.60	25.84	N/A
HALS	7.78	10.29	N/A
β -SNMF	7.93	10.38	N/A

Table 4.6: Mean performance metrics in dB, over 50 realizations, achieved by different algorithms. The oracle mean NMSE is 22.58 dB.

4.4.2.2 Comparison with other blind source separation methods

Let us compare the results returned by the previously considered DBSS and BSS algorithms. The mean performance metrics over 50 noise realizations are reported in Table 4.6. Equivalently to the synthetic data, SDecGMCA markedly overcomes oDecGMCA in terms of estimation error on both \mathbf{S} and \mathbf{A} .

The estimates of the thermal source by the five algorithms are reported in Fig. 4.10 and 4.11, along with the ground truth and the oracle estimation. As said earlier, the SDecGMCA and oracle estimates are remarkably similar. The oDecGMCA algorithm reconstructs the source with a slightly lower resolution. More importantly, it is contaminated by outlier pixels, which are likely due to the regularization favoring the higher frequencies combined with too low threshold. All the finer details are lost in the GMCA estimate because the observations are degraded beforehand. Both HALS and β -SNMF do not correctly denoise the sources; this highlights the advantage of the sparsity constraint in a transformed domain.

4.4.2.3 Impact of the noise

We then focus on the behavior of the different algorithms with respect to the SNR. The SNR is, in fact, the observation parameter that generally has the most significant impact on the results and robustness of the algorithms.

In the first place, the optimal regularization parameters are estimated (by resolving the non-blind problem). The results are plotted in Fig. 4.12. As observed with the synthetic data, the optimal regularization hyperparameter of strategy #4 (refinement strategy) is relatively insensitive to the SNR and is worth approximately 0.5. The optimal regularization hyperparameter at warm-up (strategy #3) is more sensitive to the SNR. However, as with the synthetic data, it has limited influence on the result of the separation (thereafter, we set $c_{wu} = 1e - 3$, independently of the SNR).

The performance metrics of SDecGMCA as a function of the SNR and for different regularization hyperparameters at refinement are plotted in Fig. 4.13. Similarly to



Figure 4.10: Estimated source $\mathbf{S}_{2:}$ by different algorithms. The plots are in logarithmic scale.



(a) Ground truth



Figure 4.11: Zoom on the estimated source $S_{2:}$ by different algorithms. The plots share the same logarithmic scale.



Figure 4.12: Mean optimal regularization hyperparameter of the non-blind problem, over 20 noise realizations, as a function of the SNR and the regularization strategy on the realistic data.

the synthetic data, the NMSE and C_A stabilize when there is little noise. In overall, the choice of c_{ref} has little impact on the finale result (approximately -2 dB at most for both metrics). Overestimating c_{ref} improves the estimation of **A** to the detriment of the estimation of **S**, and vice versa.

The different DBSS and BSS algorithms are finally compared for different noise levels. The results are reported in Fig. 4.14. Compared to the synthetic sources, the realistic sources are more correlated and less sparse in the starlet domain. This can result in a decrease in robustness. In order to compare the different methods, we calculated the mean performance metrics only over the successful realizations (that is with a C_A close enough to the maximum C_A across all realizations). SDecGMCA demonstrates a satisfactory robustness to noise in terms of convergence. As with the synthetic data, the performance metrics tend to stabilize when there is little noise. These experiments highlight that oDecGMCA is not robust to low noise level; the few cases that converge return good metrics, hence an apparent improvement in C_A and NMSE at high SNR compared to SDecGMCA. The performances of GMCA notably decrease at a low noise level. Indeed, the degree of sparsity of the sources is markedly decreased at low resolution; when combined with a low noise level (and thus a smaller threshold regularization), the separation process is deteriorated.

4.5 Conclusion

This chapter investigates a new joint deconvolution and sparse blind source separation algorithm to analyze multichannel spherical data. We first thoroughly investigate



Figure 4.13: Mean performance metrics over 20 noise realizations performed by SDecGMCA on the realistic data as a function of the SNR, with $c_{wu} = 1e - 3$ at warm-up and with different regularization hyperparameters c_{ref} at refinement, which are indicated as multiples of $c_{ref_{opt}}$ in the legends.

the impact of the regularization scheme in the least-square minimization, which is proven to impact the quality of the separation process broadly. We further introduce dedicated regularization schemes that better adapt to the statistics of the sources to be estimated. Based on these regularization techniques, we present a two-step minimization algorithm coined SDecGMCA, which is shown to provide a robust and effective minimization procedure. Numerical experiments on both toy examples and realistic astronomical simulations evaluate how the proposed algorithm performs in a wide range of mixing scenarios. Comparisons with standard BSS methods are further carried out, which shows that the proposed SDecGMCA algorithm is competitive.



Figure 4.14: Mean performance metrics and convergence rate over 20 noise realizations performed by different DBSS and BSS algorithms on the realistic data as a function of the SNR.

Chapter 5

Joint deconvolution and blind source separation with non-coplanar interferometric data

5.1 Introduction

The advent of large-scale radio-interferometers mandates the development of new analysis methods that can tackle non-coplanar effects. As seen in Chapter 3, these effects arise from long baselines (*i.e.*, large distance between the most distant antennas), which generally break the assumption that the antennas are on the same plane, as well as wide-field observations, because the incident waves are no longer plane. Similarly to Chapter 4, the standard linear mixture model becomes obsolete, which calls for the design of a dedicated DBSS algorithm.

In the present chapter, we propose a method to address DBSS with non-coplanar interferometric data. As with DecGMCA, jointly tackling deconvolution and separation is expected to yield better results, allowing for a more accurate accounting of the forward model and noise in a single pass. The algorithm, called wGMCA, is tested on synthetic and realistic data, and compared to sequential standard non-coplanar deconvolution and separation methods. To the best of our knowledge, there exists no BSS method that can tackle measurements with non-coplanar effects.

5.2 Principle

5.2.1 Context

Recall the imaging model of non-coplanar radio interferometers on which we concluded Chapter 3 (Eq. 3.28):

$$\tilde{y}^{(\nu)}(u,v,w) = \tilde{H}^{(\nu)}(u,v,w) \left[\mathcal{F} \circ \left(G^{(w)} \odot x^{(\nu)} \right) \right](u,v,w) + \tilde{n}^{(\nu)}(u,v,w), \quad (5.1)$$

with:

- ν a spectral band, which is hereafter confused with its central frequency,
- (u, v, w) the baseline coordinate system and (l, m) the sky coordinate system,
- $\tilde{y}^{(\nu)}(u, v, w)$ the observed data at frequency ν ,
- $\tilde{H}^{(\nu)}(u, v, w)$ the response of the interferometer, which is a mask that depends on the position of the antennas and ν ,
- \mathcal{F} the two-dimensional Fourier transform operator, which is applied in this case on the u and v axes,
- $G^{(w)}(l,m) \coloneqq e^{-2\pi i w \left(\sqrt{1-l^2-m^2}-1\right)}$ the *w*-modulation term arising from the non-coplanar effect,
- $x^{(\nu)}(l,m)$ the sought-after image at frequency ν ,
- $\tilde{n}^{(\nu)}(u, v, w)$ the noise term at frequency ν .

In Chapter 3, we reviewed the main deconvolution algorithms for non-coplanar interferometric data, and we highlighted that the w-stacking method [Offringa et al., 2014] offered a decent compromise between accuracy and computational cost. Therefore, we choose to build upon the w-stacking framework to derive a non-coplanar DBSS algorithm.

In a multispectral context, consider an interferometer that probes the sky on J channels between a low frequency ν_{min} and a high frequency ν_{max} . Thereafter, we define the normalized baseline coordinate system (u, v, w) according to the high-frequency channel; in other words, if (d_u, d_v, d_w) denotes the baseline coordinate system, $u \coloneqq d_u \nu_{max}/c$, $v \coloneqq d_v \nu_{max}/c$ and $w \coloneqq d_w \nu_{max}/c$. Consequently, a pair of antennas with baseline (d_u, d_v, d_w) will produce interferometric samples at $(u\nu/\nu_{max}, v\nu/\nu_{max}, w\nu/\nu_{max})$ for each frequency ν .

Following the *w*-stacking framework, the *w*-axis is uniformly discretized into W values and, for each channel, the interferometric samples are assigned to their nearest *w*-plane. The number of planes W should be great enough so that the maximal

angle difference between two samples of the same *w*-plane is small, which yields the criterion $W \gg 2\pi(\max(w) - \min(w)) \max_{l,m}(1 - \sqrt{1 - l^2 - m^2})$ [Offringa et al., 2014]. For each channel and for each *w*-plane, the interferometric samples are then gridded on a uniform (u, v) grid of size $\sqrt{P} \times \sqrt{P}$ and flattened in a vector of size P. This leads to the obtaining of a three-dimensional tensor $\boldsymbol{\mathcal{Y}} \in \mathbb{C}^{J \times W \times P}$.

5.2.2 Operators on tensors

Before going further, let us define three operators on three-dimensional tensors:

- As previously, \odot represents the element-wise product. If \mathcal{T} and \mathcal{U} are both in $\mathbb{C}^{I \times J \times K}$, then so is $\mathcal{T} \odot \mathcal{U}$, and $(\mathcal{T} \odot \mathcal{U})_{ijk} \coloneqq \mathcal{T}_{ijk} \mathcal{U}_{ijk}$.
- \boxtimes_n denotes the *n*-mode product, which can be viewed as the multiplication of the *n*th axis of a tensor by a matrix. For instance, if $\mathcal{T} \in \mathbb{C}^{I \times J \times K}$ is a tensor and $\mathbf{M} \in \mathbb{C}^{N \times I}$ a matrix, then $\mathbf{M} \boxtimes_1 \mathcal{T} \in \mathbb{C}^{N \times J \times K}$ is a tensor such that $(\mathbf{M} \boxtimes_1 \mathcal{T})_{:ik} \coloneqq \mathbf{M} \mathcal{T}_{:jk}$.
- \square a binary operator, which we define as combining two matrices to produce a tensor that is the element-wise product of the two matrices along their second axis. More specifically, if $\mathbf{M} \in \mathbb{C}^{J \times K}$ and $\mathbf{P} \in \mathbb{C}^{I \times K}$, then $\mathbf{M} \boxdot \mathbf{P} \in \mathbb{C}^{I \times J \times K}$ such that $(\mathbf{M} \boxdot \mathbf{P})_{ij:} = \mathbf{M}_{j:} \odot \mathbf{P}_{i:}$ (note the index swap, which is introduced for practical reasons with respect to the following observation model).

5.2.3 Forward observation model

The forward model which makes the link between the observed data $\boldsymbol{\mathcal{Y}}$ and the mixture \mathbf{AS} is:

$$\boldsymbol{\mathcal{Y}} = \boldsymbol{\mathcal{H}} \odot (\mathbf{F} \boxtimes_3 (\mathbf{G} \boxdot (\mathbf{AS}))) + \boldsymbol{\mathcal{N}}, \tag{5.2}$$

where $\mathcal{H} \in \mathbb{C}^{J \times W \times P}$ is the mask of the interferometer at each channel and each w-plane, $\mathbf{F} \in \mathbb{C}^{P \times P}$ is the 2D Fourier transform matrix (reordered so as to tackle flattened images), $\mathbf{G} \in \mathbb{C}^{W \times P}$ accounts for the *w*-terms due to the non-coplanar effect, $\mathbf{A} \in \mathbb{R}^{J \times I}$ and $\mathbf{S} \in \mathbb{R}^{I \times P}$ are the usual mixing and source matrices, and finally $\mathcal{N} \in \mathbb{C}^{J \times W \times P}$ is a zero-mean complex Gaussian noise. An illustration of the model is shown in Fig. 5.1.

Two specifications on the model are made. Firstly, in agreement with the response of an interferometer in a non-coplanar setting (Eq. 3.17), the elements of **G** are given by:

$$\mathbf{G}_{\omega p} \coloneqq \exp\left(-2\pi i w(\omega) \left(\sqrt{1 - l(p)^2 - m(p)^2} - 1\right)\right),\tag{5.3}$$



Figure 5.1: Illustration of the (noiseless) forward observation model. The pictures show absolute values, except those concerning the w-term matrix \mathbf{G} , which show real parts.

where $w(\omega)$ is the value of w associated to the plane indexed by ω , and l(p) and m(p) are the direction cosines associated to pixel p.

Secondly, as in the planar case, the mask verifies a conjugate-symmetry propriety: $\mathcal{H}_{j\omega p} = \mathcal{H}^*_{j\omega' p'}$ where ω' is the plane index such that $w(\omega') = -w(\omega)$ and p' is the pixel such that l(p') = -l(p) and m(p') = -m(p). Since the observed mixtures are real-valued, the same conjugate-symmetry applies for the data \mathcal{Y} and the noise \mathcal{N} .

In Eq. (5.2), the mixing matrix \mathbf{A} is applied on the source axis of \mathbf{S} , while operators \mathbf{G} and \mathbf{F} are applied on the pixel axis of \mathbf{S} . It follows that the model can be rewritten equivalently:

$$\boldsymbol{\mathcal{Y}} = \boldsymbol{\mathcal{H}} \odot \left(\mathbf{A} \boxtimes_1 \left(\mathbf{F} \boxtimes_3 \left(\mathbf{G} \boxdot \mathbf{S} \right) \right) \right) + \boldsymbol{\mathcal{N}}.$$
(5.4)

The latter formulation is numerically advantageous because the costly Fourier transform is evaluated fewer times, since the overdetermined case is considered $(J \ge I)$.

5.2.3.1 Assumptions on the noise

The independence of the noise between two channels and two *w*-planes is clear. In this regard, let $\{\Sigma^{(j,\omega)} \in \mathbb{R}^{P \times P}\}_{j,\omega}$ denote the covariance matrices of \mathcal{N} along the pixel axis.

The gridding performed on the interferometric measurements of a single channel and a single *w*-plane mixes the data and introduces noise correlation. However, for numerical and memory reasons, we will neglect this effect and thus consider that the $\{\Sigma^{(j,\omega)}\}_{j,\omega}$ are diagonal. This assumption is all the more true if the convolution kernel used for the gridding is small, incidentally. We note that the diagonal terms of the covariance matrices are not constant since the visibility space is not sampled uniformly by the interferometer.

To summarize, the considered noise is zero-mean complex Gaussian, independent but not identically distributed. The variances per sample are stored in diagonal matrices $\{\Sigma^{(j,\omega)}\}_{j,\omega}$.

5.2.4 Cost-function

The previously seen SBSS framework is adapted to tackle non-coplanar interferometric data.

Data-fidelity term The noise is reminded to be additive and Gaussian, with a known covariance. The data-fidelity term is chosen as the noise neg-log-likelihood, which leads to:

$$\mathcal{D}(\mathbf{A}, \mathbf{S}) \coloneqq \frac{1}{2} q \left(\mathbf{\mathcal{Y}} - \mathbf{\mathcal{H}} \odot \left(\mathbf{A} \boxtimes_1 \left(\mathbf{F} \boxtimes_3 \left(\mathbf{G} \boxdot \mathbf{S} \right) \right) \right) \right), \tag{5.5}$$

where q is a quadratic form, which depends on the noise covariance.

Source regularization The sources are assumed to be both nonnegative and sparse in a transformed domain (*e.g.*, wavelets); a nonnegative constraint and a ℓ_1 -penalization term on the transformed coefficients of the sources are therefore considered. Let $\mathbf{W} \in \mathbb{R}^{P' \times P}$ be a dictionary of the considered sparsifying transform; the regularization term is then given by:

$$h_{\mathcal{S}}(\mathbf{S}) \coloneqq \left\| \mathbf{\Lambda} \odot \left(\mathbf{S} \mathbf{W}^{\top} \right) \right\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}), \qquad (5.6)$$

where $\Lambda \in \mathbb{R}^{I \times P'}$ is the regularization hyperparameter, written in matrix form to allow it to be sample and source dependent, and $\mathcal{O}_{\mathcal{S}} := \{\mathbf{X} \in \mathbb{R}^{I \times P}, \forall (i, p), \mathbf{X}_{ip} > 0\}.$

Mixing-matrix constraint In order to prevent the scale degeneracy $\|\mathbf{S}_{i:}\|_2 \to 0$ and $\|\mathbf{A}_{:i}\|_2 \to \infty$ that is inherent to SBSS, the spectra constituting the mixing matrix are enforced to belong to the unit ball. Assuming furthermore the nonnegativity of the spectra leads to the following constraint:

$$h_{\mathcal{A}}(\mathbf{A}) \coloneqq \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}), \tag{5.7}$$

where $\mathcal{K}_{\mathcal{A}} = \left\{ \mathbf{A} \in \mathbb{R}^{J \times I}, \forall (i, j), \mathbf{A}_{ji} \ge 0, \forall i, \|\mathbf{A}_{:i}\| \le 1 \right\}.$

To summarize, the following cost-function is considered:

$$\min_{\mathbf{A},\mathbf{S}} \frac{1}{2} q \left(\boldsymbol{\mathcal{Y}} - \boldsymbol{\mathcal{H}} \odot \left(\mathbf{A} \boxtimes_1 \left(\mathbf{F} \boxtimes_3 \left(\mathbf{G} \boxdot \mathbf{S} \right) \right) \right) \right) + \left\| \boldsymbol{\Lambda} \odot \left(\mathbf{S} \mathbf{W}^\top \right) \right\|_1 + \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}) + \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}).$$
(5.8)

5.3 Towards non-coplanar DBSS

We choose to tackle the above non-coplanar DBSS problem with a pALS scheme, since it already proved to be efficient and reliable in the context of BSS [Kervazo, 2019]. Moreover, we propose to keep the general architecture of the planar version of SDecGMCA, noting that the latter algorithm already addresses a particular case of wGMCA (with a single *w*-plane at w = 0, that is W = 1 and **G** a row vector of ones).

5.3.1 Update of S

Following the pALS scheme, the aim is to solve approximately:

$$\min_{\mathbf{S}} \frac{1}{2} q \left(\boldsymbol{\mathcal{Y}} - \boldsymbol{\mathcal{H}} \odot \left(\mathbf{A} \boxtimes_{1} \left(\mathbf{F} \boxtimes_{3} \left(\mathbf{G} \boxdot \mathbf{S} \right) \right) \right) \right) + \left\| \boldsymbol{\Lambda} \odot \left(\mathbf{S} \mathbf{W}^{\top} \right) \right\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}}(\mathbf{S}), \quad (5.9)$$

by first minimizing the data-fidelity term and then applying the proximal operator of the two regularization terms.

However, the direct minimization of the data-fidelity term with a fixed **A** is intractable, as it requires inverting W matrices of size $IP \times IP$ (in the simpler cases, $P = 128^2$ typically). The minimization is therefore decomposed into two subproblems.

5.3.1.1 Subproblem I: deconvolve and unmix the visibilities

We propose to firstly estimate $\mathcal{X} := \mathbf{F} \boxtimes_3 (\mathbf{G} \boxdot \mathbf{S}) \in \mathbb{C}^{I \times W \times P}$, that is the sources modulated by the *w*-terms and projected in Fourier space. Indeed, the problem becomes separable along the *w* and pixel axes, which reduces the computation costs.

For each plane ω and pixel p, let $\mathbf{y}^{(\omega,p)} \eqqcolon \mathbf{\mathcal{Y}}_{:\omega p} \in \mathbb{C}^{J}$ be the multichannel visibilities, $\mathbf{x}^{(\omega,p)} \rightleftharpoons \mathbf{\mathcal{X}}_{:\omega p} \in \mathbb{C}^{I}$ be the Fourier transform of the sources modulated by the *w*-terms and $\mathbf{H}^{(\omega,p)} \coloneqq \operatorname{diag}(\mathbf{\mathcal{H}}_{:\omega p}) \in \mathbb{C}^{J \times J}$ be the mask of the interferometer recast in a diagonal matrix. Since the noise is independent along the channel axis, the covariance matrix of the noise which contaminates $\mathbf{y}^{(\omega,p)}$ is diagonal and given by $\mathbf{\Psi}^{(\omega,p)} \coloneqq \operatorname{diag}_{i}(\mathbf{\Sigma}^{(j,\omega)}{}_{pp})$.

Equipped with these handy matrix-vector notations, and temporarily omitting the dependencies in ω and p for clarity, the quadratic form q takes a simple form such that the first subproblem can be written for each w-plane and each pixel:

$$\min_{\mathbf{x}} \frac{1}{2} \left\| \boldsymbol{\Psi}^{-\frac{1}{2}} \left(\mathbf{y} - \mathbf{H} \mathbf{A} \mathbf{x} \right) \right\|_{2}^{2}.$$
 (5.10)

The above problem is very similar to the least squares tackled in the source update of SDecGMCA (see Eq. 4.12). The operator \mathbf{HA} is likely to be ill-posed, since the mask can be mostly null. We therefore add an extra Tikhonov regularization so that Eq. (5.10) can be inverted directly:

$$\min_{\mathbf{x}} \frac{1}{2} \left\| \Psi^{-\frac{1}{2}} \left(\mathbf{y} - \mathbf{H} \mathbf{A} \mathbf{x} \right) \right\|_{2}^{2} + \frac{1}{2} \sum_{i=1}^{I} \varepsilon^{(i)} \left| \mathbf{x}_{i} \right|_{2}^{2},$$
(5.11)

where $\{\varepsilon^{(i)}\}_i$ are the Tikhonov regularization coefficients, which generally depend on the sources, the pixels and the *w*-planes. The solution is then given by:

$$\hat{\mathbf{x}} \coloneqq \left(\mathbf{A}^{\top} \mathbf{H}^{\dagger} \mathbf{\Psi}^{-1} \mathbf{H} \mathbf{A} + \mathbf{D}\right)^{-1} \mathbf{A}^{\top} \mathbf{H}^{\dagger} \mathbf{\Psi}^{-1} \mathbf{y}, \qquad (5.12)$$

where $\mathbf{D} = \operatorname{diag}_{i}(\varepsilon^{(i)})$ is the diagonal matrix constituted of the Tikhonov regularization coefficients for all sources.

5.3.1.2 Subproblem II: correct the non-coplanar effect

114

The aim is now to recover **S** from the previous estimation of \mathcal{X} . Let $\mathbf{s}^{(i)} \coloneqq \mathbf{S}_{i:} \in \mathbb{R}^{P}$ be source $i, \mathbf{\Gamma}^{(\omega)} \coloneqq \text{diag}(\mathbf{G}_{\omega:}) \in \mathbb{C}^{P \times P}$ be the *w*-terms of plane ω recast in a diagonal matrix, $\boldsymbol{\xi}^{(i,\omega)} \coloneqq \mathcal{X}_{i\omega:} \in \mathbb{C}^{P}$ be the Fourier transform of source i modulated by the aforementioned *w*-term, and $\hat{\boldsymbol{\xi}}^{(i,\omega)}$ be the estimation of $\boldsymbol{\xi}^{(i,\omega)}$ formed by Eq. (5.12) $(i.e., \hat{\boldsymbol{\xi}}_{p}^{(i,\omega)} \coloneqq \hat{\mathbf{x}}_{i}^{(\omega,p)})$.

The second subproblem is separable in the source axis; it can be written for each source:

$$\min_{\mathbf{s}^{(i)}} \sum_{\omega=1}^{W} \frac{1}{2} \left\| \mathbf{\Omega}^{(i,\omega)^{-\frac{1}{2}}} \left(\hat{\boldsymbol{\xi}}^{(i,\omega)} - \mathbf{F} \boldsymbol{\Gamma}^{(\omega)} \, \mathbf{s}^{(i)} \right) \right\|_{2}^{2}.$$
(5.13)

The diagonal matrices $\{\Omega^{(i,\omega)}\}_{i,\omega}$ allow for a weighting to account for the uncertainty on $\{\hat{\boldsymbol{\xi}}^{(i,\omega)}\}_{i,\omega}$; the calculations are detailed in the paragraph below. The solution of Eq. (5.13) is given by:

$$\hat{\mathbf{s}}^{(i)} = \left(\sum_{\omega=1}^{W} {\mathbf{\Gamma}^{(\omega)}}^{\dagger} \mathbf{F}^{\dagger} {\mathbf{\Omega}^{(i,\omega)}}^{-1} \mathbf{F} {\mathbf{\Gamma}^{(\omega)}}\right)^{-1} \left(\sum_{\omega=1}^{W} {\mathbf{\Gamma}^{(\omega)}}^{\dagger} \mathbf{F}^{\dagger} {\mathbf{\Omega}^{(i,\omega)}}^{-1} \hat{\boldsymbol{\xi}}^{(i,\omega)}\right).$$
(5.14)

The analytical evaluation of Eq. (5.14) requires inverting a matrix of size $P \times P$, which is numerically intractable. As a first approximation, if the non-coplanar effects are small (*i.e.*, $\Gamma^{(\omega)} \approx \mathbf{I}$), Eq. (5.14) can be approximated as follows:

$$\hat{\mathbf{s}} \approx \sum_{\omega=1}^{W} {\mathbf{\Gamma}^{(\omega)}}^{\dagger} \mathbf{F}^{\dagger} \mathbf{O}^{(i,\omega)} \hat{\boldsymbol{\xi}}^{(i,\omega)}, \qquad (5.15)$$

where $\mathbf{O}^{(i,\omega)} \coloneqq \left(\sum_{\omega'=1}^{W} \mathbf{\Omega}^{(i,\omega')^{-1}}\right)^{-1} \mathbf{\Omega}^{(i,\omega)^{-1}}$. In any case, a finer estimation of $\hat{\mathbf{s}}^{(i)}$ can be obtained by the conjugate gradient algorithm [Hestenes and Stiefel, 1952], taking as a starting point the sources obtained by Eq. (5.15).

Calculation of the weights $\{\Omega^{(i,\omega)}\}_{i,\omega}$ Recall that Subproblem I aims at estimating $\mathcal{X} := \mathbf{F} \boxtimes_3 (\mathbf{G} \boxdot \mathbf{S})$, that is the sources modulated by the *w*-terms and projected in Fourier space, from the data \mathcal{Y} . In order to be able to compute numerically the solution, a Tikhonov regularization is introduced. However, it induces an estimation bias:

Bias
$$\left(\hat{\boldsymbol{\mathcal{X}}}_{:\omega p}\right) = \left(\left(\mathbf{M}^{(\omega,p)} + \operatorname{diag}_{i}\left(\varepsilon^{(i,\omega,p)}\right)\right)^{-1}\mathbf{M}^{(\omega,p)} - \mathbf{I}\right)\boldsymbol{\mathcal{X}}_{:\omega p},$$
 (5.16)

where $\mathbf{M}^{(\omega,p)} = \mathbf{A}^{\top} \operatorname{diag}_{j} \left(|\mathcal{H}_{j\omega p}|^{2} / \Sigma^{(j,\omega)}{}_{pp} \right) \mathbf{A}$ and \mathbf{I} is the identity matrix. Moreover, the noise which contaminates the data propagates in the estimate of \mathcal{X} . The following covariance matrices are derived:

$$\operatorname{Cov}\left(\hat{\boldsymbol{\mathcal{X}}}_{:\omega p}\right) = \left(\mathbf{M}^{(\omega, p)} + \operatorname{diag}_{i}\left(\varepsilon^{(i, \omega, p)}\right)\right)^{-1} \mathbf{M}^{(\omega, p)}\left(\mathbf{M}^{(\omega, p)} + \operatorname{diag}_{i}\left(\varepsilon^{(i, \omega, p)}\right)\right)^{-1}.$$
(5.17)

Subproblem II focuses on estimating the sources **S** from the previous estimation of \mathcal{X} . As we have just seen, the bias on $\hat{\mathcal{X}}$ is not null and its variance is not identically distributed. Therefore, we consider the following weighting:

$$\mathbf{\Omega}^{(i,\omega)} \coloneqq \operatorname{Cov}\left(\hat{\boldsymbol{\mathcal{X}}}_{i\omega:}\right) + \operatorname{Bias}\left(\hat{\boldsymbol{\mathcal{X}}}_{i\omega:}\right) \operatorname{Bias}\left(\hat{\boldsymbol{\mathcal{X}}}_{i\omega:}\right)^{\dagger}.$$
(5.18)

Concerning the first term of the sum, the covariance matrix is diagonal by construction and the diagonal elements are given by $\operatorname{Cov}(\hat{\boldsymbol{X}}_{i\omega:})_{pp} = \operatorname{Cov}(\hat{\boldsymbol{X}}_{:\omega p})_{ii}$ (by independence of the input noise). Regarding the second term, the following equality holds: $\operatorname{Bias}(\hat{\boldsymbol{X}}_{i\omega:})_p = \operatorname{Bias}(\hat{\boldsymbol{X}}_{:\omega p})_i$; for numerical reasons, only the diagonal elements of $\operatorname{Bias}(\hat{\boldsymbol{X}}_{i\omega:})$ are calculated.

5.3.1.3 Regularization of the sources

Following the pALS minimization procedure, the proximal operator of $h_{\mathcal{S}}$ is then applied on the above estimation of **S**. It is approximated as the composition of the projection on the nonnegative orthant and \mathcal{T}_{Λ} , that is the soft-threshold operator with thresholds Λ , in the sparsifying domain:

$$\hat{\mathbf{S}} \leftarrow \operatorname{prox}_{h_{\mathcal{S}}}(\hat{\mathbf{S}}) \approx \max\left(\mathcal{T}_{\mathbf{\Lambda}}(\mathbf{S}\mathbf{W}^{\top}) \mathbf{M}^{\top}, \mathbf{0}\right),$$
 (5.19)

with $\mathcal{T}_{\Lambda}(\cdot) \coloneqq \operatorname{sign}(\cdot) \odot \max(|\cdot| - \Lambda, \mathbf{0})$ where the sign, max and absolute value operators are intended element-wise, and \mathbf{M} verifies $\mathbf{MW} = \mathbf{I}$.

5.3.2 Update of A

Minimizing the data-fidelity term with **S** fixed can be performed separately for each channel of **A**. Recalling that $\mathcal{X} \coloneqq \mathbf{F} \boxtimes_3 (\mathbf{G} \boxdot \mathbf{S})$, it leads to:

$$\min_{\mathbf{A}_{j:}} \sum_{\omega=1}^{W} \frac{1}{2} \left\| \boldsymbol{\Sigma}^{(j,\omega)^{-1/2}} \left(\boldsymbol{\mathcal{Y}}_{j\omega:} - \boldsymbol{\mathcal{H}}_{j\omega:} \odot \left(\left(\boldsymbol{\mathcal{X}}_{:\omega:} \right)^{\top} \mathbf{A}_{j:} \right) \right) \right\|_{2}^{2}.$$
(5.20)

whose solution is given by:

$$\left(\hat{\mathbf{A}}_{j:}\right)^{\top} = \left(\sum_{\omega=1}^{W} \sum_{p=1}^{P} \Sigma^{(j,\omega)}{}_{pp}{}^{-1} \left(\mathcal{H}_{j\omega p}\right)^{\dagger} \mathcal{Y}_{j\omega p} \left(\mathcal{X}_{:\omega p}\right)^{\dagger}\right)$$
$$\left(\sum_{\omega=1}^{W} \sum_{p=1}^{P} \Sigma^{(j,\omega)}{}_{pp}{}^{-1} \left|\mathcal{H}_{j\omega p}\right|^{2} \mathcal{X}_{:\omega p} \left(\mathcal{X}_{:\omega p}\right)^{\dagger}\right)^{-1}.$$
 (5.21)

In a second step, the proximal operator of the mixing-matrix constraint $h_{\mathcal{A}}$ is applied. More precisely, for each spectrum *i*:

$$\hat{\mathbf{A}}_{:i} \leftarrow \operatorname{prox}_{h_{\mathcal{A}}} \left(\hat{\mathbf{A}}_{:i} \right) = \frac{\max\left(\mathbf{A}_{:i}, \mathbf{0} \right)}{\max\left(\left\| \max\left(\mathbf{A}_{:i}, \mathbf{0} \right) \right\|_{2}, 1 \right)}.$$
(5.22)

5.4 The wGMCA algorithm

In the previous section, we presented the pALS framework upon which the wGMCA algorithm is built. In this section, the implementation details of wGMCA are described.

5.4.1 Regularization strategies

In the source update, the minimization of Subproblem I relies on a Tikhonov regularization. Drawing inspiration of SDecGMCA, we consider two strategies to determine the regularization coefficients $\{\varepsilon^{(i,\omega,p)}\}_{i,\omega,p}$, so that they depend on a single regularization hyperparameter, denoted c:

• strategy #1: $\varepsilon^{(i,\omega,p)} \coloneqq \max\left(0, c - \frac{\lambda_{min}\left((\mathbf{A}_{j:})^{\top} \operatorname{diag}_{j}\left(|\mathcal{H}_{j\omega p}|^{2}/\boldsymbol{\Sigma}^{(j,\omega)}_{pp}\right)\mathbf{A}_{j:}\right)}{\lambda_{min}\left(\mathbf{A}^{\top}\mathbf{A}\right)}\right)$ where

 $\lambda_{min}(\cdot)$ returns the smallest eigenvalue of the input matrix. As we showed in Chapter 4, the gist of this strategy is to limit the noise amplification.

• strategy #2: $\varepsilon^{(i,\omega,p)} := \frac{c}{|\boldsymbol{\chi}_{i\omega p}|^2}$. This strategy is reminiscent of a Wiener deconvolution filter.

As in SDecGMCA, empirical tests show that strategy #2 provides better source estimates. However, unlike strategy #1, it requires a prior knowledge of the sought-after sources.

5.4.2 Improving the estimation of the mixing matrix

The least-square update of \mathbf{A} (Eq. 5.21) is likely to lack robustness in practice, when working with realistic or real data. In this section, we analyze the reasons for this and deduce an improved update for the mixing matrix.

Firstly, in the estimation of the sources, the convolution kernel is inverted totally. The sparse regularization (and, in a lesser extent, the Tikhonov regularization) helps recovering the missing high-frequency information of the sources that is not sampled by the interferometer. This information may not be reliable to estimate the mixing matrix. Therefore, a first improvement would consist in re-convolving partially the sources, so as to remove high-frequency deconvolution artifacts. Secondly, as seen in Chapter 1, the classical GMCA algorithm is preferably run directly in the transformed sparsifying domain \mathbf{W} , because astrophysical sources correlate less there. It leads in practice to a better estimation of the mixing matrix, the sources being synthesized in the direct domain in a second step after the separation. Consequently, a second way of improvement would involve updating the mixing matrix with the sources and the data expressed in the transformed sparsifying domain.

With this in mind, let $\mathbf{K} \in \mathbb{R}^{P \times P}$ be a user-defined low-pass filter, which attenuates the higher frequencies that are not sampled by the interferometer. We propose to substitute the least squares on \mathbf{A} of Eq. (5.20) by:

$$\min_{\mathbf{A}_{j:}} \sum_{\omega=1}^{W} \frac{1}{2} \left\| \boldsymbol{\Upsilon}^{(j,\omega)^{-1/2}} \left(\boldsymbol{\mathcal{Y}}_{j\omega:} - \boldsymbol{\mathcal{H}}_{j\omega:} \odot \left(\left(\boldsymbol{\mathcal{X}}_{:\omega:}^{\prime} \right)^{\top} \mathbf{A}_{j:} \right) \right) \right\|_{2}^{2},$$
(5.23)

where $\mathcal{X}' \coloneqq \mathbf{F} \boxtimes_3 (\mathbf{G} \boxdot (\mathbf{S} \mathbf{K}^\top \mathbf{W}^\top))$ is constituted of the sources filtered by \mathbf{K} (so as to remove high-frequency deconvolution artifacts), projected in the sparsifying domain \mathbf{W} (in order to attenuate possible correlations between the sources), modulated by the *w*-terms and projected in Fourier space, and $\Upsilon^{(j,\omega)} \coloneqq \Sigma^{(j,\omega)} \odot |\mathbf{F} \mathbf{K}|^2 \odot |\mathbf{F} \mathbf{W}|^2$ accounts for the effects of \mathbf{K} and \mathbf{W} (in doing so, it is assumed that \mathbf{W} is a linear filter, which still includes many transforms like wavelets).

Rigorously, the data \mathcal{Y} should likewise be adapted to take into account **K** and **W**. However, the matrix **G**, which renders the *w*-effect, prevents from having an analytical expression of the transformed data from \mathcal{Y} , as in the standard BSS case or even the coplanar DBSS case, so the approximation is left as is.

The new update reads similarly to Eq. (5.21), replacing $\boldsymbol{\mathcal{X}}$ by $\boldsymbol{\mathcal{X}}'$ and $\boldsymbol{\Sigma}^{(j,\omega)}$ by $\boldsymbol{\Upsilon}^{(j,\omega)}$.

5.4.3 Thresholding parameters

The thresholding procedure of GMCA, which has been detailed in Section 1.3.3.4, is implemented likewise in the wGMCA algorithm. It is reminded to include: an automatic setting of the thresholds based on the level of the propagated noise (" κ - σ "), a support-based strategy that decreases progressively the thresholds (" K_{max} ") and ℓ_1 -reweighting scheme.

5.4.4 A two-step minimization

As for SDecGMCA, the wGMCA algorithm relies on a two-stage minimization procedure. The first stage provides an initial estimate of the mixing matrix and the sources, it ensures some robustness with respect to the initial point (**warm-up**). The second stage refines the estimates by using more accurate techniques (**refinement**).

Stage	1: warm-up	2: refinement
Regularization strategy	#1	#2
Use conjugate gradient	No	If needed, yes
Active source support K	$0 \to K_{max}$	K_{max}
ℓ_1 -reweighting	No	Yes
Enforce non-negativity on ${\bf S}$	No	Yes

Joint deconvolution and blind source separation with non-coplanar interferometric data

Table 5.1: Parameters of wGMCA in the two stages.

The wGMCA algorithm is described in Algorithm 13, with the parameter values of the algorithm according to the step summarized in Table 5.1.

Warm-up The mixing matrix is initialized running Principle Component Analysis (PCA) on the *w*-plane with the most visibility samples.

In the first warm-up phase, regularization strategy #1 is used since it requires no knowledge on the sources. Moreover, the refinement of the source estimate by the conjugate gradient is not performed, even if the non-coplanar effects are large, as the small non-coplanar effect approximation is sufficient to have a first estimate of the sources and the mixing matrix.

Regarding the choice of the thresholds, because the starting point is generally poor, (i) the earlier mentioned support-based strategy is performed to facilitate the identification of the main directions of the mixture, (ii) no reweighted- ℓ_1 is achieved in order to avoid spurious critical points and (iii) the nonnegativity constraint is not applied as it tends to destabilize the algorithm when far from the solution.

The warm-up ends when the sources have converged $(\|\mathbf{S}^{(n)} - \mathbf{S}^{(n-1)}\|_2 / \|\mathbf{S}^{(n)}\|_2 \le \epsilon_{wu}$ where *n* is the iteration and ϵ_{wu} does not need to be particularly small, *e.g.*, $\epsilon_{wu} = 1e - 2$).

Refinement At the beginning of the second refinement phase, the estimates of the sources and the mixing matrix are generally acceptable. This allows the more precise regularization strategy #2 to be employed, using the latest estimate of the sources: $\varepsilon^{(i,\omega,p)(n)} = c_{ref}/|\mathcal{X}_{i\omega p}^{(n-1)}|^2$. In addition, if non-coplanar effects are important, the conjugate gradient algorithm may be used to improve the source estimates.

Concerning the regularization of the sources, (i) the source support is kept constant, (ii) the thresholds are adapted with the ℓ_1 -reweighting procedure and (iii) the non-negativity constraint of the sources is applied.

The refinement ends likewise when the sources have stabilized $(||\mathbf{S}^{(n)} - \mathbf{S}^{(n-1)}||_2 / ||\mathbf{S}^{(n)}||_2 \le \epsilon_{ref}$, with $\epsilon_{ref} = 1e - 4$ in practice).

118

Algorithm 13 wGMCA

Inputs: data \mathcal{Y} , masks \mathcal{H} , w-term matrix \mathbf{G} , number of sources I, diagonal of the noise covariance matrices $\{\Sigma^{(j,\omega)}_{pp}\}_{j,\omega,p}$, regularization hyperparameters c_{wu} and c_{ref} , sparsifying transform \mathbf{W} , thresholding parameters κ and K_{max} , low-pass filter \mathbf{K}

 $\mathbf{A} \leftarrow \text{PCA}(\boldsymbol{\mathcal{Y}}_{:\omega:})$, with ω the *w*-plane with the most visibility samples

(1) Warm-up stage
stage ← warm-up
while convergence not reached do
Update K according to Table 5.1
S ← UpdateS(stage, A, Y, H, G, {Σ^(j,ω)_{pp}}_{j,ω,p}, c_{wu}, W, κ, K)
A ← UpdateA(S, Y, H, G, {Σ^(j,ω)_{pp}}_{j,ω,p}, K)
(2) Refinement stage
stage ← refinement
while convergence not reached do
S ← UpdateS(stage, A, Y, H, G, {Σ^(j,ω)_{pp}}_{j,ω,p}, c_{ref}, W, κ, K_{max})
A ← UpdateA(S, Y, H, G, {Σ^(j,ω)_{pp}}_{j,ω,p}, K)

Outputs: mixing matrix A, sources S
Algorithm 14 UpdateS

Inputs: stage, A, \mathcal{Y} , \mathcal{H} , G, $\{\Sigma^{(j,\omega)}{}_{pp}\}_{j,\omega,p}$, c, W, κ , K

(1) Subproblem I for $\omega = 1, ..., W$ and p = 1, ... P do Determine the Tikhonov regularization coefficients $\{\varepsilon^{(i,\omega,p)}\}_i$ according to the stage and c $\boldsymbol{\mathcal{X}}_{:\omega p} \leftarrow \left(\mathbf{A}^{\top} \operatorname{diag}_j\left(|\boldsymbol{\mathcal{H}}_{j\omega p}|^2 / \boldsymbol{\Sigma}^{(j,\omega)}_{pp}\right) \mathbf{A} + \operatorname{diag}_i\left(\varepsilon^{(i,\omega,p)}\right)\right)^{-1}$ $\mathbf{A}^{\top} \operatorname{diag}_j\left(\boldsymbol{\mathcal{H}}_{j\omega p}^* / \boldsymbol{\Sigma}^{(j,\omega)}_{pp}\right) \boldsymbol{\mathcal{Y}}_{:\omega p}$

(2) Subproblem II for i = 1, ..., I do Determine the weighting matrices $\{\Omega^{(i,\omega)}\}_{\omega}$ (Eq. 5.18) for $\omega = 1, ..., W$ do $\mathbf{O}^{(i,\omega)} \leftarrow \left(\sum_{\omega'=1}^{W} \Omega^{(i,\omega')^{-1}}\right)^{-1} \Omega^{(i,\omega)^{-1}}$ $\mathbf{S}_{i:} \leftarrow \sum_{\omega=1}^{W} \text{diag} (\mathbf{G}_{\omega:})^{\dagger} \mathbf{F}^{\dagger} \mathbf{O}^{(i,\omega)} \boldsymbol{\mathcal{X}}_{i\omega:}$

if stage is refinement and conjugate gradient requested then Refine estimation **S** with conjugate gradient algorithm

(3) Regularization

Calculate the thresholds Λ with a support-based strategy and a $\ell_1\text{-reweighting}$ if need be, according to the stage

if stage is warm-up then

 $\mathbf{S} \leftarrow \mathcal{T}_{\mathbf{\Lambda}} ig(\mathbf{S} \mathbf{W}^{ op} ig) \mathbf{W}$ else

 $\mathbf{S} \leftarrow \max\left(\mathcal{T}_{\boldsymbol{\Lambda}}\!\left(\mathbf{S}\mathbf{W}^{\top}\right)\mathbf{W}, \mathbf{0}\right)$

Output: S

Algorithm 15 UpdateA

Inputs: S, $\boldsymbol{\mathcal{Y}}, \, \boldsymbol{\mathcal{H}}, \, \mathrm{G}, \, \{\boldsymbol{\Sigma}^{(j,\omega)}{}_{pp}\}_{j,\omega,p}, \mathrm{K}$

for
$$i = 1, ..., I$$
 and $\omega = 1, ..., W$ do
 $\mathcal{X}'_{iw:} \leftarrow \mathbf{F} \operatorname{diag}(\mathbf{G}_{\omega:}) \mathbf{W} \mathbf{K}^{\top} \mathbf{S}_{i:}$
for $j = 1, ..., J$ and $\omega = 1, ..., W$ do
 $\Upsilon^{(j,\omega)} \leftarrow \Sigma^{(j,\omega)} \odot |\mathbf{F} \mathbf{K}|^2 \odot |\mathbf{F} \mathbf{W}|^2$

(1) Least squares
for
$$j = 1, ..., J$$
 do
 $(\mathbf{A}_{j:})^{\top} \leftarrow \left(\sum_{\omega=1}^{W} \sum_{p=1}^{P} \mathbf{\Upsilon}^{(j,\omega)}{}_{pp}^{-1} (\mathbf{\mathcal{H}}_{jwp})^{\dagger} \mathbf{\mathcal{Y}}_{jwp} (\mathbf{\mathcal{X}}'_{:\omega p})^{\dagger}\right)^{-1}$
 $\left(\sum_{\omega=1}^{W} \sum_{p=1}^{P} \mathbf{\Upsilon}^{(j,\omega)}{}_{pp}^{-1} |\mathbf{\mathcal{H}}_{jwp}|^{2} \mathbf{\mathcal{X}}'_{:\omega p} (\mathbf{\mathcal{X}}'_{:\omega p})^{\dagger}\right)^{-1}$

(2) Constraints for i = 1, ..., I do $\mathbf{A}_{:i} \leftarrow \max(\mathbf{A}_{:i}, \mathbf{0})$ $\mathbf{A}_{:i} \leftarrow \mathbf{A}_{:i} / \max(\|\mathbf{A}_{:i}\|_{2}, 1)$

Output: A

5.4.5 Complexity and convergence proprieties

122

The computational complexity of deconvolution algorithms, both in terms of time and memory, is one of the main limitations of radio-interferometry. Optimizing the methods is therefore a hot research topic, especially for processing massive datasets of future large-scale radio-interferometers.

The goal of this work being to propose a first DBSS algorithm for non-coplanar interferometric data, we have left aside the computational optimization aspects. As such, the wGMCA algorithm is limited in time by the Fourier transforms in the source update (Subproblem II and proximal update), which scale in $\mathcal{O}(P \log P)$. Concerning memory, the current implementation stores in volatile memory the data tensor \mathcal{Y} , the mask tensor \mathcal{H} and the diagonal elements of the covariance matrices $\{\Sigma^{(j,\omega)}_{pp}\}_{j,\omega,p}$, each of which is of size JWP.

Convergence towards a local minimum is not guaranteed by wGMCA. However, as with the other GMCA-based algorithms, an empirical stabilization is observed.

5.5 Numerical experiments

Two sets of data are considered for the numerical experiments. The first one consists of synthetic sources and mixing matrices; it allows easy modification of the experimental parameters, making it adequate for characterizing the wGMCA algorithm. The second dataset is composed of realistic astrophysical sources and mixing matrices, and is used to compare wGMCA with standard methods that perform the non-coplanar deconvolution and separation sequentially.

The intensity of the non-coplanar effect depends on the size of the field of view, the non-coplanarity of the interferometer and the channel frequency. Rather than providing the three values, we will indicate the maximal phase difference of the *w*-terms $\Delta \varphi_{max} \coloneqq \max_{\omega,p}(|\arg(\mathbf{G}_{\omega p})|) \approx 2\pi \max(w) \max_{l,m}(1 - \sqrt{1 - l^2 - m^2})$, which is a more global metric to account for the non-coplanar effect intensity. Typically, the non-coplanar effect must be addressed as soon as $\Delta \varphi_{max} \gtrsim 0.1$ rad.

In digital signal processing, the signal-to-noise ratio is typically defined as the ratio of the absolute square of the noiseless data on the expectancy of the absolute square of the noise: $\sum_{j\omega p} |(\mathcal{H} \odot \mathbf{F} \boxtimes_3 \mathbf{G} \boxdot \mathbf{AS})_{j\omega p}|^2 / \sum_{j\omega p} (\Sigma_{pp}^{(j,\omega)})$. In our case, the noise is not identically distributed, and consequently neither is the signal-to-noise ratio. As an illustration, Figure 5.2 shows an example of signal-to-noise ratio distribution by an actual interferometer. Such a definition of SNR would lead to high values that poorly reflect the level of data corruption. Therefore, we employ the SNR defined as



Figure 5.2: Example of signal-to-noise ratio distribution $|(\mathcal{H} \odot \mathbf{F} \boxtimes_3 \mathbf{G} \boxdot \mathbf{AS})_{j\omega}|^2/\operatorname{diag}(\Sigma^{(j,\omega)})$ by MeerKAT (see Section 5.5.2.1 for more details) for particular values of j and ω (which are related to ν and w, respectively), in dB. The images all share the same colorbar. The signal-to-noise ratio does not exist in the white areas, where there is no data. The SNR defined according to (5.24) is 41 dB, the median values being here (a) 39 dB, (b) 15 dB and (a) 14 dB. In comparison, the SNR according to the usual definition is 100 dB.

follows:

$$\operatorname{SNR} \coloneqq 10 \log_{10} \left(\max_{j,\omega} \left(\operatorname{median}_p \left(\frac{|(\mathcal{H} \odot \mathbf{F} \boxtimes_3 \mathbf{G} \boxdot \mathbf{AS})_{j\omega p}|^2}{\boldsymbol{\Sigma}_{pp}^{(j,\omega)}} \right) \right) \right).$$
(5.24)

In other words, the median SNR is calculated for each channel and each w-plane, and the highest value is kept. It is found to reflect the degree of data corruption more adequately. In any case, dirty maps will be given to provide an intuition of the level of noise contamination in the data considered.

5.5.1 Characterization of wGMCA on synthetic data

The wGMCA algorithm is first characterized using toy datasets.

5.5.1.1 Data generation

Sources and mixing matrix The normalized baseline coordinate space (u, v, w) is discretized on a grid of size $\sqrt{P} \times \sqrt{P} \times W$. The sources are random nonnegative images of size $\sqrt{P} \times \sqrt{P} = 128 \times 128$ that are sparse in the starlet (an isotropic undecimated wavelet [Starck et al., 2015]) domain, see an example in Fig. 5.4a. The mixing matrices are random nonnegative matrices with a given condition number.



Figure 5.3: Example of masks $\mathcal{H}_{j\omega}$ used in the toy model for some particular values of j and ω (which are related to ν and w, respectively), with a sampling density of $\alpha = 10\%$. All the images are in a logarithmic scale and share the same colorbar.

Masks The channels are set at frequencies evenly distributed between a low frequency ν_{min} and a high frequency ν_{max} . Concerning the masks of the high frequency channel, pairs of points which respect the earlier mentioned symmetry of the interferometric sampling are drawn randomly from the grid according to a centered Gaussian distribution – in this manner, the points near the origin (0,0,0) are favored, as for an actual interferometer. The masks of the other channels are then deduced; for each channel, the normalized baseline coordinates of the interferometric samples are recalculated to take into account the dilation effect (as seen previously, with a factor ν/ν_{max}) and a gridding is performed, firstly along the w-axis by associating each sample to the nearest w-plane, and then along the (u, v) plane using bilinear interpolation. Finally, the masks are all multiplied by a Gaussian-shaped kernel to emulate the angular response of an interferometer. Examples of masks are shown in Figure 5.3.

Parameter	Symbol	Default value
Maximal phase difference of the w -terms	$\Delta \varphi_{max}$	$0.1 \mathrm{rad}$
Sampling density	α	10%
Resolution ratio	r	0.5
Signal-to-noise ratio	SNR	30 dB
Number of channels	J	8
Mixing matrix condition number	$\operatorname{cond}(\mathbf{A})$	2

Table 5.2: Default observation parameters of the numerical experiments on the synthetic data.

Noise The data are generated according to the forward model (5.2) using the gridded sources. The noise variance of each sample is chosen proportional to the mask, *i.e.*, $\Sigma_{pp}^{(j,\omega)} \propto \mathcal{H}_{j\omega p}$. This choice mimics the noise distribution after the gridding process, because the more a zone is sampled, the higher the noise level but the lower the signal-to-noise ratio.

About the gridding The data gridding, which allows placing the interferometric samples on a regular (u, v, w)-grid, introduces errors (mainly aliasing and discretization); in the non-coplanar framework, it causes more particularly smearing and decorrelation effects [Offringa, A. R. et al., 2019]). Different methods exist to limit these errors, the simplest being to increase the number of samples on the grid.

In the experiments conducted, the sources are generated directly on a regular (u, v, w)grid; therefore, the impact of the gridding on the separation cannot be evaluated (it
is anyway outside the scope of this work). Hereafter, the number of w-planes is set
to W = 11; this choice is purely arbitrary as the data are generated directly from the
gridded sources.

Observation parameters Six observation parameters are considered, namely:

- the above-mentioned maximal phase difference of the w-terms $\Delta \varphi_{max}$,
- the interferometer sampling density $\alpha \in]0, 1]$, which is defined as the proportion of nonzero elements in the mask of the high-frequency channel,
- the resolution ratio $r \coloneqq \nu_{min}/\nu_{max}$,
- the SNR as defined in (5.24),
- the number of channels J,
- the mixing matrix condition number cond(A).

Joint deconvolution and blind source separation with non-coplanar interferometric data



(a) One (of four) ground-truth source.



(d) Estimated source.



(b) Dirty mixture at channel #0.



(e) Source absolute estimation error.



(c) Dirty mixture at channel #7.



(f) Mixing matrix. Solid line: estimation, thick transparent line: ground-truth.

Figure 5.4: Example of separation result with the synthetic data. The images are in a logarithmic scale.

5.5.1.2 Result example

A representative example of result is shown in Fig. 5.4 with the default observation parameters. Two dirty mixtures are also shown (see Fig. 5.4b and 5.4c) to illustrate the effects of the interferometer instrumental response and of the noise.

Both sources and mixing matrix are recovered very precisely (see Fig. 5.4d and Fig. 5.4f, respectively). Concerning the sources, we note that the residuals are dominated by deconvolution artifacts (Fig. 5.4e), the interferences between sources being very limited.

5.5.1.3 Impact of the non-coplanarity

Let us first evaluate how the non-coplanarity of the data affects the results of the proposed algorithm. To this end, wGMCA is tested with data subject to a non-coplanar effect of increasing intensity.

As seen earlier in the update of the sources, particularly in Subproblem II (5.13), we resort to a small non-coplanar effect approximation allowing us to write the ana-



Figure 5.5: Median performance metrics of wGMCA with and without the conjugate gradient refinement, over 10 realizations (with first and third quartile), as a function of the maximal phase difference of the *w*-terms $\Delta \varphi_{max}$.

lytical solution (5.15); we specify that the conjugate gradient method can still be used as a refinement when the non-coplanar effect is strong. In this regard, the wGMCA algorithm is executed with and without the conjugate gradient refinement to highlight the range where the use of the conjugate gradient is required.

The results of the experiments, which are generalized with Monte-Carlo trials, are reported in Fig. 5.5. When the conjugate gradient method is not used, the performance unsurprisingly decreases as soon as the non-coplanar effect becomes important, as of 0.1 rad. The effect is particularly pronounced for the source estimate, less for the mixing matrix one.

Nevertheless, the use of the conjugate gradient algorithm allows to maintain the separation performance when the non-coplanar effect is important. We note that the NMSE increases with the non-coplanar effect intensity, which may seem counterintuitive, while the SAD decreases. It is actually a common behavior in BSS, that the degradation of the estimation of one parameter can lead to the improvement of the estimation of the other.

The number of conjugate gradient iterations is also shown in the right panel of Fig. 5.5 (averaged over the iterations of wGMCA). The stopping criterion is based on the maximum number of iterations (1000) or on the ratio of the norms of the residual and the data. We note that the number of iterations remains bounded.

In the following experiments on synthetic data, the non-coplanar effect will be kept constant at $\Delta \varphi_{max} = 0.1$ rad. This will avoid the use of the conjugate gradient algorithm and decrease the execution time.

5.5.1.4 Impact of the regularization hyperparameters

The wGMCA algorithm requires setting two regularization hyperparameters, namely c_{wu} for the warm-up stage and c_{ref} for the refinement step. In this subsection, we assess their impacts on the algorithm.



Figure 5.6: Median optimal hyperparameters at warm-up c_{wu} and refinement c_{ref} over 100 realizations (with first and third quartile) as a function of an observation parameter, on the synthetic data.

As in the previous chapters, let the oracle estimator designate the version of wGMCA in which the sources and the mixing matrix are estimated with the groundtruth mixing matrix and sources, respectively. The oracle can be used to provide optimal values of the two hyperparameters, that is which provide the best reconstructions of the sources based on the NMSE.

We first evaluate the optimal values of c_{wu} and c_{ref} with respect to the observation parameters. To do so, Monte-Carlo experiments are performed; for each realization, the regularization hyperparameters which provide the best reconstructions are searched. The results are reported in Fig. 5.6. Both hyperparameters at warm-up and refinement tend to be stable with respect to the observation parameters, which is valuable for physical applications. This is certainly due to the fact that the deconvolution scheme takes into account the noise statistics.

When working with real data, the optimal hyperparameters are unknown. If the previous experiment shows that these are rather insensitive to the experimental conditions, it does not inform about the consequence of a poor choice of hyperparameters. Therefore, for a selection of observation scenarios, the optimal hyperparameters of the previous experiments c_{wuopt} and c_{refopt} are taken as reference, and the wGMCA algorithm is run with varying values of hyperparameters around these optimal values. The results are reported in Fig. 5.7. In overall, the separation quality is notably insensitive to the choice of the hyperparameters. In particular, the warm-up hyperpa-

rameter c_{wu} has almost no effect on the finale result. The refinement step is the most determining, but the window of acceptable values for the refinement hyperparameter c_{ref} is large, of several orders of magnitude.

5.5.1.5 Impact of the observation parameters

Finally, we test how the previously defined observation parameters influence the separation performance of the wGMCA algorithm. The results of the Monte-Carlo experiments are reported in Fig. 5.8. The performance metrics of the oracle estimator are also plotted; it provides the upper-bound that wGMCA can reach, the difference being due to the blind nature of the unmixing.

In overall, the wGMCA algorithm achieves good results, mostly close to the oracle. The evolutions of both source NMSEs and mixing matrix SADs are consistent, since the better the observation conditions, the better the estimates. Concerning the SNR, the metrics stagnate over 40 dB; this is likely because of the automatic thresholding strategy, as we already observed with sGMCA and SDecGMCA. When the sampling density reaches 100%, both metrics stabilize on a plateau. This is due to the mask generation strategy that favors low spatial frequencies: At some point, the new samples are predominantly high frequency and provide little information about the sources, which are relatively low frequency. We note that the mixing matrix condition number does not affect the SAD, only the NMSE. In fact, it impacts the conditioning of the estimation problem of \mathbf{S} , not particularly that of \mathbf{A} .

5.5.2 Comparisons with other methods on realistic data

In a second stage, the wGMCA algorithm is tested on realistic astrophysical data, and compared to other standard methods.

5.5.2.1 Description of the data

In the following experiments, I = 4 sources of size $P = 128 \times 128$ representing realistic diffuse astrophysical sources are considered (see first column in Fig. 5.10). The same set of realistic emission spectra as in the experiments of sGMCA in Chapter 2 is used, the sole difference being that the spectra are resampled along J = 20 channels (see example in Fig. 5.11).

The masks are based on the MeerKAT telescope [Jonas, 2009]. It is a set of 64 antennas in South Africa, operating on a wide frequency radio-band from approximately 700 MHz up to 14.5 GHz. 75% of the antennas are located within a 500-meter radius, the remaining 25% being distributed over a larger area with a maximum baseline of 8 km. MeerKAT is one of the precursors of SKA and should be integrated to it in



130



(a) Sampling density α . Top row: $\alpha = 1\%$, bottom row: $\alpha = 100\%$.



(b) Resolution ratio r. Top row: r = 0.1, bottom row: r = 0.9.



(c) Signal-to-noise ratio. Top row: SNR = 10 dB, bottom row: SNR = 80 dB.

(d) Number of channels J. Top row: J = 4, bottom row: J = 25.



Figure 5.7: Median performance metrics of wGMCA over 100 realizations (with first and third quartile) as a function of the regularization hyperparameter at warm-up c_{wu} or at refinement c_{ref} , which are both expressed in terms of their optimal value, in various experimental scenarios on the synthetic data. When testing c_{wu} (resp. c_{ref}), the optimal value of c_{ref} (resp. c_{wu}) is taken.



Figure 5.8: Median performance metrics of wGMCA over 100 realizations (with first and third quartile) as a function of an observation parameter, on the synthetic data.



Figure 5.9: Example of realistic masks $\mathcal{H}_{j\omega}$: for some particular values of j and ω (which are related to ν and w, respectively), with an observation time of $\Delta t = 8$ hours. All the images are in a logarithmic scale and share the same colorbar.

the long term.

The masks are generated as follows:

- the (u, v, w)-coordinates of the interferometric samples are calculated based on the MeerKAT array configuration with a random sky pointing, during an userdefined observation time Δt with a 30-minute time step,
- the *w*-axis is discretized into W = 11 planes,
- for each channel, the samples are grouped by w-values, and for each w-plane, the samples are gridded on a uniform $P \times P$ grid. To that end, a Gaussian-sinc kernel of support 11 is used, with an oversampling factor of 63.

Examples of masks are shown in Fig. 5.9. The noise tensor is created in parallel, as well as the covariance matrices. The interferometric data are finally generated using the forward model (5.2).

5.5.2.2 Sequential deconvolution and separation methods

Four methods that perform deconvolution and separation sequentially are tested. To this end, two deconvolution algorithms are considered, namely:

- a Tikhonov-based deconvolution, quite similar to the least squares on **S** by replacing **A** by the identity matrix,
- a WSCLEAN-based algorithm [Offringa et al., 2014] that is adapted to enforce the sparsity of the images in the starlet domain.

These two deconvolution algorithms are then combined with both GMCA and HALS [Gillis and Glineur, 2012], bringing the number of comparison algorithms to four.

It is important to note that the separation step generally fails if the sources are fully deconvolved, likely because of deconvolution artifacts; the mixtures are therefore deconvolved at the resolution of the least-resolved low-frequency channel.

5.5.2.3 Results

A typical example of separation result by wGMCA and a sequential method is shown in Fig. 5.10 and 5.11 for the sources and for the mixing matrix, respectively. Example of dirty mixtures are also provided (see two last rows of Fig. 5.10) to illustrate the effects of the interferometer instrumental response and of the noise.

Regarding the sources, the wGMCA algorithm recovers them quite precisely (see third column of Fig. 5.10). Interestingly, the source on the first line is estimated with the same high resolution. Indeed, it emits mostly in the best-resolved high-frequency channels (blue spectrum in Fig. 5.11), which explains that the dirty image is well resolved. On the contrary, the resolution of the three other source is lower - but still better than the dirty images - because they emit predominantly in the leastresolved low-frequency channels. The sources estimated by Tikhonov+HALS (see fourth column of Fig. 5.10) present a low resolution, as the other sequential method in fact. Indeed, as stated above for the sequential methods, the mixtures can only be deconvolved partially for the separation to succeed.

Concerning the spectra, the estimations of wGMCA are in overall more precise than the sequential methods, with less interference.

Quantitative results derived from Monte Carlo experiments are finally reported in Tables 5.3 and 5.4. Eight configurations are tested, combining two regimes of noise level, two resolution ratios and two observation times.

The gain of wGMCA lies mainly in the estimation of the sources, which is more accurate thanks to the joint deconvolution. Indeed, the high-frequency information present in the data is better rendered in the estimated sources. This results in an improvement in the estimation of the mixing matrix, which is however less significant.

Joint deconvolution and blind source separation with non-coplanar interferometric data



Figure 5.10: Example of separation result in the source domain. Above the line: four sources, in a logarithmic scale, sharing the same colorbar. Below the line: selection of two zooms, in a logarithmic scale, but with different colorbars for the sake of clearness. The data have been generated with the masks of Fig. 5.9, a SNR of 60 dB, and the mixing matrix of Fig. 5.11. The dirty mixtures are at the channels where each source is the most visible (that is, from left to right, channels #15, #6, #12 and #9).



Figure 5.11: Example of estimated mixing matrices. Solid line: estimation, thick transparent solid line: ground-truth.

The principal limitation of the proposed algorithm concerns the best conditioned cases, for instance when the measurements are homogeneous in resolution (which corresponds to a resolution ratio that tends to one) or when the mask has few holes (which corresponds to long observation times). Indeed, addressing the deconvolution jointly does not appear to improve the separation particularly.

5.6 Conclusion

We investigate a joint deconvolution and sparse blind source separation algorithm coined wGMCA to analyze multichannel non-coplanar interferometric data. Numerical experiments on both synthetic and realistic data are achieved in various experimental settings. In overall, it is advantageous to process the non-coplanar deconvolution and the blind source separation in a single pass rather than sequentially, as current standard method would. The proposed algorithm proves to be competitive, estimating sources with improved resolution, which increases the overall separation quality.

	$r = 0.1 \&$ $\Delta t = 8 h$		$r = 0.5 \&$ $\Delta t = 8 h$		$r = 0.1 \&$ $\Delta t = 30 mn$		$r = 0.5 \&$ $\Delta t = 30 mn$	
Method	SAD	NMSE	SAD	NMSE	SAD	NMSE	SAD	NMSE
Oracle	13.08	25.80	15.32	26.27	12.83	28.81	15.05	28.85
wGMCA	8.08	15.30	7.91	13.46	7.70	14.00	7.77	13.29
Tikhonov+GMCA	7.64	11.73	7.60	11.53	7.01	11.68	7.01	11.52
Tikhonov+HALS	7.95	13.26	7.82	13.52	7.17	12.88	6.84	12.28
WSClean+GMCA	6.79	7.59	7.71	9.72	5.97	6.09	6.67	5.47
WSClean+HALS	7.77	8.85	7.98	10.73	5.61	6.54	6.40	6.05

Table 5.3: Median metrics in dB over 100 realizations, with varying masks and mixing matrices, for different resolution ratios r and observation times Δt , in a low noise regime (SNR = 60 dB when $\Delta t = 8$ h, and SNR = 80 dB when $\Delta t = 30$ mn).

Method	$r = 0.1 \&$ $\Delta t = 8 h$		$r = 0.5 \&$ $\Delta t = 8 h$		$r = 0.1 \&$ $\Delta t = 30 mn$		$r = 0.5 \&$ $\Delta t = 30 mn$	
	SAD	NMSE	SAD	NMSE	SAD	NMSE	SAD	NMSE
<i>Oracle</i> wGMCA	12.65 7.70	19.47 12.48	<i>14.69</i> 7.41	23.59 1 2.26	<i>10.93</i> 6.05	27.90 12.27	14.96 7.58	28.75 13.01
Tikhonov+GMCA Tikhonov+HALS WSClean+GMCA	7.32 7.64 6.81	$10.94 \\ 10.78 \\ 7.55$	7.42 7.64 7.65	11.74 12.64 0.58	7.04 7.05 5.98	11.97 12.30 6 10	7.04 7.20 6.67	$11.68 \\ 12.57 \\ 5.47$
WSClean+HALS	7.66	8.85	8.27	10.74	5.98 5.71	6.55	6.67	6.13

Table 5.4: Median metrics in dB over 100 realizations, with varying masks and mixing matrices, for different resolution ratios r and observation times Δt , in a high noise regime (SNR = 10 dB when $\Delta t = 8$ h, and SNR = 30 dB when $\Delta t = 30$ mm).

Conclusion and perspectives

BSS algorithms are the cornerstone of multivalued data analysis in several applications. However, as seen in the introduction, current methods will face several bottlenecks when processing data of forthcoming telescopes in astrophysics. In this regard, we study and propose three algorithms in this manuscript that tackle identified issues.

First, we introduce in Chapter 2 the semi-blind source separation algorithm sGMCA. Built upon a variational framework that allows accounting for the mixing model, sGMCA makes use of learned priors on the sought-after spectra. Compared to fully blind methods, numerical experiments emphasize that the components are better separated, with a noticeable decrease in leakage, especially in challenging contexts (low emissions, high noise, *etc.*). This is particularly beneficial for scientific applications, as the physical interpretability of the separation results is thereby improved.

We then describe in Chapter 4 the SDecGMCA algorithm, which allows performing joint deconvolution and blind source separation from wide-field observations on the sphere. In order to keep a minimization procedure based on pALS, it is necessary to regularize the least squares on the sources, which is otherwise naturally ill-posed. This extra regularization is thoroughly analyzed and optimized, and a strategy for automatic choice of the regularization parameters is derived. Numerous numerical experiments demonstrate that performing the data deconvolution and separation jointly allows a clear improvement of the separation quality.

Finally, we address in Chapter 5 joint deconvolution and source separation in the case of non-coplanar interferometric data, and introduce the wGMCA algorithm. We keep the pALS minimization scheme but have to adapt it to the new mixing model that is computationally more complex. As far as we know, wGMCA is the first BSS method for non-coplanar interferometric data to date. Extensive tests on synthetic and realistic data show the advantage of tackling deconvolution and BSS in a single pass rather than sequentially, as standard methods would.

The advantage of the variational formulation and the projected alternating least squares on which the proposed algorithms are based is their flexibility. In this respect, although the developments are done separately, it would be straightforward to include the learned constraints of sGMCA in joint deconvolution and blind source separation problems (DecGMCA, SDecGMCA, or wGMCA), as the former impact the mixing matrix \mathbf{A} and the latter the sources \mathbf{S} .

Several avenues for improvement can be considered within the framework of BSS for the analysis of astrophysical signals; they are classified into three main categories.

Modeling First, the observation model, and therefore the data-fidelity term of the cost function, may be refined. In radio-interferometry, standard deconvolution methods take the raw (at least calibrated) data as input and perform the gridding jointly with the deconvolution, which improves the reconstruction quality. In this respect, the question could be raised of including the gridding in DecGMCA and wGMCA alongside the deconvolution and the separation.

The linear mixing model considered in this manuscript is only a first-order approximation of the observed processes. A more realistic model should account for spectral variability, that is, when spectra vary from one pixel to another. The underlying unmixing problem is much more complex and ill-posed; still, it can be dealt with by introducing specific regularizations, as the spectra change little from one pixel to another. A solution currently being explored in the laboratory is to constrain the spectra in the latent spaces of IAEs, by regularizing the variation of the associated latent parameters.

Interpretability Ensuring model and result interpretability is crucial for processing real data; two tracks are particularly identified.

The first relates to the robustness of the learned spectrum regularization in sGMCA. If the manifolds are modeled too approximately, there is a risk of systematically biasing the solutions. An improvement would be to relax the current constraint by allowing the spectra to belong to the neighborhoods of the manifolds.

The second track concerns uncertainty quantification. In general, results in the field of physics are always associated with confidence intervals. In this respect, research has recently been conducted to develop numerical methods that also estimate their output uncertainty (see [Arras, Philipp et al., 2019] for example in radio-interferometry). To our knowledge, this has not yet been explored in the context of BSS.

Acceleration The proposed joint deconvolution and blind source separation methods have a computational cost that can quickly become prohibitive with real observations, especially radio interferometric data. Thus, future developments will have to focus on code optimization and parallelization. A way could be to resort to patches on the sources (*e.g.*, see [Kervazo et al., 2020b] in the context of SBSS and [Tasse, C. et al., 2018] in the context of non-coplanar interferometric deconvolution).

Deep unrolling is also a way to speed up the proposed algorithms. This may not necessarily be more difficult to implement than BSS unrolling [Fahes et al., 2022] since it would only be a matter of adapting already learned networks for given types of mixtures **AS** to observation operators **H** or \mathcal{H} under consideration (a single instrumental response in the case of a single-dish telescope, or a class of masks in the case of an interferometer).

Appendix A

Elements of signal processing for random processes

Cosmic signals in radio astronomy are inherently random; therefore, a statistical processing stage is required to derive any valuable information (see Chapter 3). This appendix provides the main definitions and theoretical results of random processes, on which statistical methods in radio astronomy are based.

Let X(t) be a complex random process, that is a function of time whose value at any time t is a random complex variable [Charbit, 1996]. Rigorously, it should be referred to as X or $X(\cdot)$; for the sake of clarity, we will however abuse the notations thereafter to make the nature of variables explicit (*e.g.* time, frequency) and/or distinguish several variables of the same nature in a single expression. A realization of X(t) is denoted x(t).

Statistical moments We firstly define various statistical moments:

• the expectation $\mu_X(t)$:

$$\mu_X(t) := \mathbb{E} \left[X(t) \right],$$

$$:= \int_{\mathbb{R}} \int_{\mathbb{R}} (x_r + ix_i) f_X(x_r, x_i, t) \mathrm{d}x_r \mathrm{d}x_i,$$
 (A.1)

where $f_X(x_r, x_i, t)$ is the probability density function of X(t),

• the autocorrelation $R_{XX}(t_1, t_2)$:

$$R_{XX}(t_1, t_2) \coloneqq \mathbf{E} \left[X(t_1) X(t_2)^* \right], \tag{A.2}$$

• the crosscorrelation between X(t) and another complex random process Y(t), denoted $R_{XY}(t_1, t_2)$:

$$R_{XY}(t_1, t_2) := \mathbb{E} \left[X(t_1) Y(t_2)^* \right].$$
 (A.3)

Stationarity A random process is first-order stationary if its expectation is independent of time:

$$X(t)$$
 is first-order stationary $\iff E[X(t)] = \mu_X \in \mathbb{C}.$ (A.4)

Moreover, it is wide-sense stationary, or second-order stationary, if its expectation is independent of time and its auto-correlation depends on the time difference only:

$$X(t) \text{ is wide-sense stationary} \iff \begin{cases} \mathbb{E}\left[X(t)\right] = \mu_X \in \mathbb{C}, \\ \mathbb{E}\left[X(t_1)X(t_2)^*\right] = R_{XX}(\tau), \ \tau = t_1 - t_2. \end{cases}$$
(A.5)

In this manuscript, "stationarity" actually refers to wide-sense stationarity.

Two random processes are jointly wide-sense stationary if they are both wide-sense stationary and if their cross-correlation depends on the time difference:

X(t) and Y(t) are jointly wide-sense stationary

$$\iff \begin{cases} X(t) \text{ and } Y(t) \text{ are both wide sense-stationary,} \\ E[X(t_1)Y(t_2)^*] = R_{XY}(\tau), \ \tau = t_1 - t_2. \end{cases}$$
(A.6)

"Joint stationarity" likewise refers to joint wide-sense stationarity in the manuscript.

Ergodicity A first-order stationary random process is first-order ergodic, or mean ergodic, if the time average of a single realization of it is equal to the expectation:

$$X(t)$$
 is first-order ergodic $\iff \lim_{T \to \infty} \frac{1}{T} \int_{-T/2}^{T/2} x(t) dt = \mu_X.$ (A.7)

A wide-sense stationary random process is wide-sense ergodic, also said second-order or correlation ergodic, if both expectation and autocorrelation can be retrieved from a single realization as follows:

$$X(t) \text{ is wide-sense ergodic} \iff \begin{cases} \lim_{T \to \infty} \frac{1}{T} \int_{-T/2}^{T/2} x(t) dt = \mu_X, \\ \lim_{T \to \infty} \frac{1}{T} \int_{-T/2}^{T/2} x(\tau) x(\tau - t)^* dt = R_{XX}(\tau). \end{cases}$$
(A.8)

In this manuscript, "ergodicity" actually refers to wide-sense ergodicity.

Power spectra A measure of interest in signal processing is the power spectral density, also referred to as power spectrum. It describes how the power of a time

function distributes along the frequencies, allowing for harmonic analysis.

Consider a function of time x(t) (which can be a realization of a random process) and a rectangular window function $w_T(t)$, such that $w_T(t) = 1$ if $t \in [-T/2, T/2]$ and 0 elsewhere; let $x_T(t) \coloneqq x(t)w_T(t)$ be a truncated version of x(t) and $\tilde{x}_T(\omega)$ its Fourier transform.

The average power of x(t) is worth $P = \lim_{T\to\infty} \frac{1}{T} \int_{\mathbb{R}} |x_T(t)|^2 dt$. According to Parseval's theorem, it is also equal to $P = \lim_{T\to\infty} \frac{1}{T} \int_{\mathbb{R}} |\tilde{x}_T(\omega)|^2 d\omega$. The power spectral density is defined as the integrand of the previous equation:

$$S_{XX}(\omega) \coloneqq \lim_{T \to \infty} \frac{1}{T} \left| \tilde{x}_T(\omega) \right|^2.$$
(A.9)

The cross power spectral density, or simply cross power spectrum, of two random processes x(t) and y(t) is defined similarly:

$$S_{XY}(\omega) \coloneqq \lim_{T \to \infty} \frac{1}{T} \tilde{x}_T(\omega) \tilde{y}_T(\omega)^*.$$
 (A.10)

Proprieties Consider X(t) a stationary random process.

The Wiener-Khinchin theorem relates its power spectrum $S_{XX}(\omega)$ to the Fourier transform of its auto-correlation $\tilde{R}_{XX}(\omega)$:

$$S_{XX}(\omega) = R_{XX}(\omega). \tag{A.11}$$

Let Y(t) be another stationary random process, which is jointly stationary with X(t). Quite similarly to the Wiener-Khinchin theorem, the cross-correlation theorem states that:

$$S_{XY}(\omega) = \hat{R}_{XY}(\omega). \tag{A.12}$$

Appendix B Spherical signal analysis

Astrophysical signals as measured from Earth are by nature defined on the sphere. Often the field of view is narrow and the corresponding portion of the sphere is approximated by a plane, which then facilitates the various numerical calculations. Nevertheless, in some applications such as cosmology, the valuable information lies in large scales; large portions of the sphere if not the whole sphere must therefore be surveyed, and the spherical formalism becomes indispensable. In this appendix, the main results of spherical signal analysis are given.

Let S^2 be the two-dimensional sphere, and f and g two real valued squareintegrable functions defined on the sphere. The inner product is defined by $\langle f, g \rangle := \frac{1}{4\pi} \int_{S^2} fg$.

B.1 Spherical harmonics

Spherical harmonics are the solution functions of Laplace's equation on the sphere, *i.e.*, the spherical functions whose Laplacian is zero. They are shown to depend on two whole parameters, known as multipoles, namely the degree $l \in \mathbb{N}$ and the order $\mathfrak{m} \in [-l, l]$; l essentially corresponds to the angular spatial frequency and \mathfrak{m} to the direction of the oscillations. Thereafter, the (l, \mathfrak{m}) th normalized spherical harmonic will be denoted $y^{(l,\mathfrak{m})}$.

Spherical harmonics are shown to form an orthonormal basis of the squareintegrable spherical functions. Therefore, if $\mathcal{D} := \{(\mathfrak{l}, \mathfrak{m}) \in \mathbb{N} \times \mathbb{Z}, |\mathfrak{m}| \leq \mathfrak{l}\}$ represents the set of multipoles, f can be decomposed as follows:

$$f = \sum_{(\mathfrak{l},\mathfrak{m})\in\mathcal{D}} \tilde{f}^{(\mathfrak{l},\mathfrak{m})} y^{(\mathfrak{l},\mathfrak{m})}, \tag{B.1}$$

where \tilde{f} is the spherical harmonic transform of f, which is worth $\tilde{f}^{(\mathfrak{l},\mathfrak{m})} \coloneqq \langle f, y^{(\mathfrak{l},\mathfrak{m})} \rangle$ at multipole $(\mathfrak{l},\mathfrak{m})$. Since f is assumed real-valued, the harmonic coefficients verify a complex-conjugate symmetry propriety: $\tilde{f}^{(\mathfrak{l},\mathfrak{m})} = \tilde{f}^{(\mathfrak{l},-\mathfrak{m})^*}$.

The spherical harmonic transform is the counterpart of the discrete-time Fourier transform for functions defined on the sphere. An angular power spectrum $\{S_{ff}^{(l)}\}_{l \in \mathbb{N}}$ can even be derived from the spherical harmonic transform of f. It is defined by:

$$S_{ff}^{(\mathfrak{l})} \coloneqq \frac{1}{2\mathfrak{l}+1} \sum_{\mathfrak{m}=-\mathfrak{l}}^{\mathfrak{l}} \left| \tilde{f}^{(\mathfrak{l},\mathfrak{m})} \right|^2.$$
(B.2)

B.2 Convolution products

Contrarily to the Euclidean case, the convolution product on the sphere is not uniquely defined. Spherical convolution products differ in particular by the nature of the convolution kernel (isotropic or not) and the output space of the product (the sphere or another space, in general the three-dimensional rotation group SO(3)) [Roddy and McEwen, 2021].

In the thesis, we employ the isotropic convolution product [Driscoll and Healy, 1994], as we consider isotropic kernels only. Let h be an isotropic real valued square-integrable spherical function. The convolution of f by h is defined by:

$$[f * h](\theta, \varphi) \coloneqq \langle f, \mathcal{R}_{\theta, \varphi}(h) \rangle, \tag{B.3}$$

where $\mathcal{R}_{\theta,\varphi}$ denotes the rotation of angles θ and φ^1 .

The main property of the isotropic convolution product is that it simplifies into a point-wise product in the spherical harmonic space:

$$\widetilde{f * h}^{(\mathfrak{l},\mathfrak{m})} = \frac{4\pi}{2\mathfrak{l} + 1} \widetilde{f}^{(\mathfrak{l},\mathfrak{m})} \widetilde{h}^{(\mathfrak{l},0)^{\dagger}}.$$
(B.4)

It is noted that the harmonic coefficients of h are null for non-null modes, *i.e.*, $\tilde{h}^{(\mathfrak{l},\mathfrak{m})} = 0$ for all $\mathfrak{m} \neq 0$, since h is isotropic.

B.3 Healpix discretization

In order to perform numerical calculations, the data must be sampled on a grid. In this thesis, we make use of the HEALPix discretization of the sphere [Górski et al., 2005]. As showed in Fig. B.1, the sphere is divided in 12 quadrilateral sections of equal area, which are hierarchically subdivided in N_{side}^2 pixels. The resolution parameter N_{side} is the number of pixels along the sides of the 12 initial sections. HEALPix is the most commonly used pixelization of the sphere in astrophysics and geophysics, since it is well adapted for hierarchical analysis and spherical harmonic projections.

¹The rotation depends on two parameters as it is applied on an isotropic function.



Figure B.1: HEALPix discretization of the sphere. Moving clockwise from top left, $N_{side} = 1, 2, 4, 8$. Image from [Górski et al., 2005].

Let $\mathbf{y}^{(\mathfrak{l},\mathfrak{m})} \in \mathbb{C}^{P}$ denote the $(\mathfrak{l},\mathfrak{m})$ th spherical harmonic sampled with HEALPix². The spherical harmonic projection of a vector \mathbf{f} , denoted $\tilde{\mathbf{f}}$, can be approximated at 0th order by:

$$\tilde{\mathbf{f}}_{(\mathfrak{l},\mathfrak{m})} \coloneqq \frac{4\pi}{P} \mathbf{y}^{(\mathfrak{l},\mathfrak{m})^{\dagger}} \mathbf{f}, \qquad (B.5)$$

and the angular power spectrum by:

$$\mathfrak{S}_{\mathbf{ff}\mathfrak{l}} = \frac{1}{2\mathfrak{l}+1} \sum_{\mathfrak{m}=-\mathfrak{l}}^{\mathfrak{l}} \left| \tilde{\mathbf{f}}_{(\mathfrak{l},\mathfrak{m})} \right|^2.$$
(B.6)

In theory, the number of coefficients of $\mathbf{\hat{f}}$ and $\mathfrak{S}_{\mathbf{ff}}$ should be infinite. In practice, as the spherical harmonics sampled with HEALPix form a linearly independent set up to $\mathfrak{l} = 3N_{side} - 1$, the signals are band-limited to a degree $\mathfrak{l}_{max} \leq 3N_{side} - 1$, and the spherical harmonic coefficients are calculated up to \mathfrak{l}_{max} .

HEALPix provides fast routines to calculate spherical harmonic decompositions. The notable difference with the discrete Fourier transform in the Euclidean case is that the projection is not exact.

²According to the HEALPix pixelization scheme, the number of pixels is $P = 12N_{side}^2$.

Appendix C

Adaptation of SDecGMCA for coplanar data

In this appendix, we provide the equations adapted from SDecGMCA for planar data (which is the case addressed by DecGMCA [Jiang et al., 2017]). The similarities with the spherical case are remarkable.

First, the mixture model writes in the Fourier domain for all sample *p*:

$$\tilde{\mathbf{Y}}_{:p} = \operatorname{diag}\left(\tilde{\mathbf{H}}_{:p}\right)\mathbf{A}\tilde{\mathbf{S}}_{:p} + \tilde{\mathbf{N}}_{:p},$$
 (C.1)

and the derived cost-function is:

$$\underset{\mathbf{A},\tilde{\mathbf{S}}}{\operatorname{argmin}} \frac{1}{2} \sum_{p} \left\| \tilde{\mathbf{Y}}_{:p} - \operatorname{diag} \left(\tilde{\mathbf{H}}_{:p} \right) \mathbf{A} \tilde{\mathbf{S}}_{:p} \right\|_{2}^{2} + \left\| \mathbf{\Lambda} \odot \left(\tilde{\mathbf{S}} \mathbf{F}^{*} \mathbf{W}^{\top} \right) \right\|_{1} + \iota_{\mathcal{O}_{\mathcal{S}}} \left(\tilde{\mathbf{S}} \mathbf{F}^{*} \right) + \iota_{\mathcal{K}_{\mathcal{A}}}(\mathbf{A}),$$
(C.2)

where this time \mathbf{F} denotes the discrete Fourier transform or FFT operator. The update of the mixing matrix reads as follows:

$$\left(\hat{\mathbf{A}}_{j:}\right)^{\top} \leftarrow \left(\sum_{p} \tilde{\mathbf{Y}}_{jp} \tilde{\mathbf{H}}_{jp}^{*} \left(\tilde{\mathbf{S}}_{:p}\right)^{\dagger}\right) \left(\sum_{p} \left|\tilde{\mathbf{H}}_{jp}\right|^{2} \tilde{\mathbf{S}}_{:p} \left(\tilde{\mathbf{S}}_{:p}\right)^{\dagger}\right)^{-1}, \quad (C.3)$$

and the proximal step is unchanged.

Concerning the sources, the least-square update reads:

$$\hat{\tilde{\mathbf{S}}}_{:p} \leftarrow \left(\mathbf{M}^{(p)} + \operatorname{diag}_{i \in [1...I]} \left(\varepsilon^{(i,p)} \right) \right)^{-1} \mathbf{A}^{\top} \operatorname{diag} \left(\tilde{\mathbf{H}}_{:p}^{*} \right) \tilde{\mathbf{Y}}_{:p}.$$
(C.4)

with $\mathbf{M}^{(p)} \coloneqq \mathbf{A}^{\top} \operatorname{diag} \left(|\tilde{\mathbf{H}}_{:p}|^2 \right) \mathbf{A}$. At warm-up, the Tikhonov regularization coefficients are given by:

$$\varepsilon^{(i,p)} = \max\left(0, c - \frac{\lambda_{\min}\left(\mathbf{M}^{(p)}\right)}{\lambda_{\min}\left(\mathbf{A}^{\top}\mathbf{A}\right) + \epsilon}\right),\tag{C.5}$$

and at refinement:

$$\varepsilon^{(i,p)} = c \, \frac{\mathfrak{S}_{\mathbf{nn}p}}{\mathfrak{S}_{\mathbf{s}_i;\mathbf{S}_{i:p}}}.\tag{C.6}$$

where \mathfrak{S}_{nn} and $\mathfrak{S}_{\mathbf{s}_i:\mathbf{s}_i:}$ denote the power spectral density of the noise and source i, respectively. The proximal step on the source matrix is unchanged.

Appendix D

Open source codes

The codes presented in the manuscript are open source and can be found online on version 3 of the LGPL:

- sGMCA: https://github.com/RCarloniGertosio/sGMCA,
- SDecGMCA: https://github.com/RCarloniGertosio/SDecGMCA,
- wGMCA: https://github.com/RCarloniGertosio/wGMCA.

Bibliography

- [Abdulaziz et al., 2019] Abdulaziz, A., Dabbech, A., and Wiaux, Y. (2019). Wideband super-resolution imaging in Radio Interferometry via low rankness and joint average sparsity models (HyperSARA). Monthly Notices of the Royal Astronomical Society, 489(1):1230–1248.
- [Adler and Öktem, 2017] Adler, J. and Öktem, O. (2017). Solving ill-posed inverse problems using iterative deep neural networks. *Inverse Problems*, 33(12):124007.
- [Adler and Öktem, 2018] Adler, J. and Öktem, O. (2018). Learned primal-dual reconstruction. *IEEE Transactions on Medical Imaging*, 37(6):1322–1332.
- [Arras, Philipp et al., 2019] Arras, Philipp, Frank, Philipp, Leike, Reimar, Westermann, Rüdiger, and Enßlin, Torsten A. (2019). Unified radio interferometric calibration and imaging with joint uncertainty quantification. A&A, 627:A134.
- [Bertero et al., 2021] Bertero, M., Boccacci, P., and De Mol, C. (2021). Introduction to inverse problems in imaging. CRC press.
- [Bobin et al., 2015] Bobin, J., Rapin, J., Larue, A., and Starck, J.-L. (2015). Sparsity and adaptivity for the blind separation of partially correlated sources. *IEEE Transactions on Signal Processing*, 63(5):1199–1213.
- [Bobin et al., 2007] Bobin, J., Starck, J.-L., Fadili, J., and Moudden, Y. (2007). Sparsity and morphological diversity in blind source separation. *IEEE Transactions on Image Processing*, 16(11):2662–2674.
- [Bolte et al., 2014] Bolte, J., Sabach, S., and Teboulle, M. (2014). Proximal alternating linearized minimization for nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1):459–494.
- [Brakel and Bengio, 2017] Brakel, P. and Bengio, Y. (2017). Learning independent features with adversarial nets for non-linear ICA. arXiv preprint arXiv:1710.05050.
- [Candès et al., 2008] Candès, E. J., Wakin, M. B., and Boyd, S. P. (2008). Enhancing sparsity by reweighted l1 minimization. Journal of Fourier Analysis and Applications, 14(5):877–905.
- [Cardoso et al., 2002] Cardoso, J.-F., Snoussi, H., and Delabrouille, J. (2002). Blind separation of noisy gaussian stationary sources. application to cosmic microwave background imaging. In 2002 11th European Signal Processing Conference, pages 1–4. IEEE.
- [Carozzi, 2015] Carozzi, T. D. (2015). Imaging on a sphere with interferometers: the spherical wave harmonic transform. Monthly Notices of the Royal Astronomical Society: Letters, 451(1):L6–L10.
- [Carrillo et al., 2012] Carrillo, R. E., McEwen, J. D., and Wiaux, Y. (2012). Sparsity Averaging Reweighted Analysis (SARA): a novel algorithm for radiointerferometric imaging. *Monthly Notices of the Royal Astronomical Society*, 426(2):1223–1234.
- [Carrillo et al., 2014] Carrillo, R. E., McEwen, J. D., and Wiaux, Y. (2014). PU-RIFY: a new approach to radio-interferometric imaging. *Monthly Notices of the Royal Astronomical Society*, 439(4):3591–3604.
- [Charbit, 1996] Charbit, M. (1996). Éléments de théorie du signal : aspects aléatoires. Ellipses.
- [Chenot, 2017] Chenot, C. (2017). Parcimonie, diversité morphologique et séparation robuste de sources. PhD thesis, Université Paris-Saclay.
- [Cherni et al., 2020] Cherni, A., Anthoine, S., and Chaux, C. (2020). β-nmf and sparsity promoting regularizations for complex mixture unmixing. application to 2d hsqc nmr. In ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 5795–5799. IEEE.
- [Cichocki et al., 2007] Cichocki, A., Zdunek, R., and Amari, S.-i. (2007). Hierarchical ALS Algorithms for Nonnegative Matrix and 3D Tensor Factorization. In *Proceedings of ICA*, pages 169–176. Springer.
- [Clark, 1980] Clark, B. G. (1980). An efficient implementation of the algorithm 'CLEAN'. Astronomy and Astrophysics, 89(3):377.
- [Combettes and Wajs, 2005] Combettes, P. L. and Wajs, V. R. (2005). Signal recovery by proximal forward-backward splitting. *Multiscale modeling & simulation*, 4(4):1168–1200.

- [Comon and Jutten, 2010] Comon, P. and Jutten, C. (2010). Handbook of Blind Source Separation: Independent Component Analysis and Applications. Academic Press, 1st edition.
- [Connor and Rozell, 2020] Connor, M. and Rozell, C. (2020). Representing closed transformation paths in encoded network latent space. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):3666–3675.
- [Conway et al., 1990] Conway, J., Cornwell, T., and Wilkinson, P. (1990). Multifrequency synthesis-a new technique in radio interferometric imaging. *Monthly Notices of the Royal Astronomical Society*, 246:490.
- [Cornwell, 2008] Cornwell, T. J. (2008). Multiscale clean deconvolution of radio synthesis images. *IEEE Journal of Selected Topics in Signal Processing*, 2(5):793–801.
- [Cornwell et al., 2008] Cornwell, T. J., Golap, K., and Bhatnagar, S. (2008). The noncoplanar baselines effect in radio interferometry: The w-projection algorithm. *IEEE Journal of Selected Topics in Signal Processing*, 2(5):647–657.
- [Cornwell and Perley, 1992] Cornwell, T. J. and Perley, R. A. (1992). Radiointerferometric imaging of very large fields. The problem of non-coplanar arrays. *Astronomy and Astrophysics*, 261:353–364.
- [Cornwell et al., 2012] Cornwell, T. J., Voronkov, M. A., and Humphreys, B. (2012). Wide field imaging for the square kilometre array. In Bones, P. J., Fiddy, M. A., and Millane, R. P., editors, *Image Reconstruction from Incomplete Data VII*, volume 8500, pages 186 – 197. International Society for Optics and Photonics, SPIE.
- [Culpepper and Olshausen, 2009] Culpepper, B. and Olshausen, B. (2009). Learning transport operators for image manifolds. In Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C., and Culotta, A., editors, Advances in Neural Information Processing Systems, volume 22. Curran Associates, Inc.
- [Donoho and Stodden, 2003] Donoho, D. and Stodden, V. (2003). When does nonnegative matrix factorization give a correct decomposition into parts? In Thrun, S., Saul, L., and Schölkopf, B., editors, Advances in Neural Information Processing Systems, volume 16. MIT Press.
- [Driscoll and Healy, 1994] Driscoll, J. R. and Healy, D. M. (1994). Computing fourier transforms and convolutions on the 2-sphere. Adv. Appl. Math., 15(2):202–250.
- [Elad, 2006] Elad, M. (2006). Why simple shrinkage is still relevant for redundant representations? *IEEE Transactions on Information Theory*, 52(12):5559–5569.

- [Fahes et al., 2022] Fahes, M., Kervazo, C., Bobin, J., and Tupin, F. (2022). Unrolling PALM for sparse semi-blind source separation. In *International Conference on Learning Representations*.
- [Fefferman et al., 2013] Fefferman, C., Mitter, S., and Narayanan, H. (2013). Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29:983– 1049.
- [Foster and Heuer, 2020] Foster, A. R. and Heuer, K. (2020). PyAtomDB: Extending the AtomDB Atomic Database to Model New Plasma Processes and Uncertainties. *Atoms*, 8(3):49.
- [Févotte and Idier, 2011] Févotte, C. and Idier, J. (2011). Algorithms for Nonnegative Matrix Factorization with the β -Divergence. Neural Computation, 23(9):2421–2456.
- [Gillis, 2020] Gillis, N. (2020). Nonnegative Matrix Factorization. SIAM Society for Industrial and Applied Mathematics.
- [Gillis and Glineur, 2012] Gillis, N. and Glineur, F. (2012). Accelerated Multiplicative Updates and Hierarchical ALS Algorithms for Nonnegative Matrix Factorization. *Neural Computation*, 24(4):1085–1105.
- [Gilton et al., 2020] Gilton, D., Ongie, G., and Willett, R. (2020). Neumann networks for linear inverse problems in imaging. *IEEE Transactions on Computational Imaging*, 6:328–343.
- [Girard et al., 2015] Girard, J., Garsden, H., Starck, J., Corbel, S., Woiselle, A., Tasse, C., McKean, J., and Bobin, J. (2015). Sparse representations and convex optimization as tools for LOFAR radio interferometric imaging. *Journal of Instrumentation*, 10(08):C08013–C08013.
- [Golbabaee et al., 2013] Golbabaee, M., Arberet, S., and Vandergheynst, P. (2013). Compressive source separation: Theory and methods for hyperspectral imaging. *IEEE Transactions on Image Processing*, 22(12).
- [Górski et al., 2005] Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., and Bartelmann, M. (2005). Healpix: A framework for highresolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759.
- [Hammernik et al., 2018] Hammernik, K., Klatzer, T., Kobler, E., Recht, M. P., Sodickson, D. K., Pock, T., and Knoll, F. (2018). Learning a variational network for reconstruction of accelerated MRI data. *Magnetic Resonance in Medicine*, 79(6):3055–3071.

- [Hestenes and Stiefel, 1952] Hestenes, M. R. and Stiefel, E. (1952). Methods of conjugate gradients for solving. Journal of research of the National Bureau of Standards, 49(6):409.
- [Högbom, 1974] Högbom, J. A. (1974). Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines. Astronomy and Astrophysics Supplement, 15:417.
- [Hu et al., 2007] Hu, S., Stead, M., and Worrell, G. A. (2007). Automatic identification and removal of scalp reference signal for intracranial eegs based on independent component analysis. *IEEE Transactions on Biomedical Engineering*, 54(9):1560– 1572.
- [Humphreys and Cornwell, 2011] Humphreys, B. and Cornwell, T. (2011). Memo 132, analysis of convolutional resampling algorithm performance. Technical report, Square Kilometre Array.
- [Hyvärinen et al., 2019] Hyvärinen, A., Sasaki, H., and Turner, R. (2019). Nonlinear ica using auxiliary variables and generalized contrastive learning. In Chaudhuri, K. and Sugiyama, M., editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 859–868. PMLR.
- [Jayaram and Thickstun, 2020] Jayaram, V. and Thickstun, J. (2020). Source separation with deep generative priors. In III, H. D. and Singh, A., editors, *Proceedings* of the 37th International Conference on Machine Learning, volume 119 of Proceedings of Machine Learning Research, pages 4724–4735. PMLR.
- [Jelić, 2010] Jelić, V. (2010). Cosmological 21cm experiments: Searching for a needle in a haystack. PhD thesis, Rijksuniversiteit Groningen.
- [Jiang et al., 2017] Jiang, M., Bobin, J., and Starck, J.-L. (2017). Joint multichannel deconvolution and blind source separation. SIAM Journal on Imaging Sciences, 10(4):1997–2021.
- [Jin et al., 2017] Jin, K. H., McCann, M. T., Froustey, E., and Unser, M. (2017). Deep convolutional neural network for inverse problems in imaging. *IEEE Trans*actions on Image Processing, 26.
- [Jonas, 2009] Jonas, J. L. (2009). Meerkat—the south african array with composite dishes and wide-band single pixel feeds. *Proceedings of the IEEE*, 97(8):1522–1530.
- [Kameoka et al., 2019] Kameoka, H., Li, L., Inoue, S., and Makino, S. (2019). Supervised determined source separation with multichannel variational autoencoder. *Neural Computation*, 31(9):1891–1914.

- [Kervazo, 2019] Kervazo, C. (2019). Optimization framework for large-scale sparse blind source separation. PhD thesis, Université Paris-Saclay.
- [Kervazo et al., 2020a] Kervazo, C., Bobin, J., Chenot, C., and Sureau, F. (2020a). Use of palm for ℓ1 sparse matrix factorization: Difficulty and rationalization of a two-step approach. *Digital Signal Processing*, 97:102611.
- [Kervazo et al., 2020b] Kervazo, C., Liaudat, T., and Bobin, J. (2020b). Faster and better sparse blind source separation through mini-batch optimization. *Digital Signal Processing*, 106:102827.
- [Klein, 2011] Klein, U. (2011). Radio astronomy: tools, applications and impacts.
- [Kleinsteuber and Shen, 2012] Kleinsteuber, M. and Shen, H. (2012). Blind source separation with compressively sensed linear mixtures. *IEEE Signal Processing Letters*, 19(2).
- [Kobarg et al., 2014] Kobarg, J., Maass, P. qnd Oetje, J., Tropp, O., Hirsch, E., Sagiv, C., Golbabaee, M., and Vandergheynst, P. (2014). Numerical experiments with maldi imaging data. *Computational Mathematics*, 40(3).
- [Kogan and Greisen, 2009] Kogan, L. and Greisen, E. W. (2009). Faceted imaging in aips. AIPS memo.
- [Le Roux et al., 2015] Le Roux, J., Weninger, F., and Hershey, J. (2015). Sparse NMF – half-baked or well done? Technical Report TR2015-023, Mitsubishi Electric Research Laboratories, Cambridge MA, USA.
- [Lee and Seung, 2000] Lee, D. and Seung, H. S. (2000). Algorithms for non-negative matrix factorization. In Leen, T., Dietterich, T., and Tresp, V., editors, Advances in Neural Information Processing Systems, volume 13. MIT Press.
- [Lunz et al., 2018] Lunz, S., Öktem, O., and Schönlieb, C.-B. (2018). Adversarial regularizers in inverse problems. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31, pages 8507–8516. Curran Associates, Inc.
- [Mallat, 2008] Mallat, S. (2008). A Wavelet Tour of Signal Processing: The Sparse Way. Academic Press, 3 edition.
- [Mallat and Zhang, 1993] Mallat, S. and Zhang, Z. (1993). Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415.

- [Mangum, J. G. et al., 2007] Mangum, J. G., Emerson, D. T., and Greisen, E. W. (2007). The on the fly imaging technique. A&A, 474(2):679–687.
- [McCann et al., 2017] McCann, M. T., Jin, K. H., and Unser, M. (2017). Convolutional neural networks for inverse problems in imaging: A review. *IEEE Signal Processing Magazine*, 34(6):85–95.
- [Monga et al., 2021] Monga, V., Li, Y., and Eldar, Y. C. (2021). Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44.
- [Narayanaswamy et al., 2020] Narayanaswamy, V., Thiagarajan, J. J., Anirudh, R., and Spanias, A. (2020). Unsupervised Audio Source Separation Using Generative Priors. In Proc. Interspeech 2020, pages 2657–2661.
- [Nasser et al., 2022] Nasser, R., Eldar, Y. C., and Sharan, R. (2022). Deep unfolding for non-negative matrix factorization with application to mutational signature analysis. *Journal of Computational Biology*, 29(1):45–55.
- [Neelamani et al., 2004] Neelamani, R., Choi, H., and Baraniuk, R. (2004). Forward: Fourier-wavelet regularized deconvolution for ill-conditioned systems. *IEEE Trans*actions on Signal Processing, 52(2):418–433.
- [Nugraha et al., 2016] Nugraha, A. A., Liutkus, A., and Vincent, E. (2016). Multichannel audio source separation with deep neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(9):1652–1664.
- [Offringa et al., 2014] Offringa, A. R. et al. (2014). WSCLEAN: an implementation of a fast, generic wide-field imager for radio astronomy. *Monthly Notices of the Royal Astronomical Society*, 444(1):606–619.
- [Offringa and Smirnov, 2017] Offringa, A. R. and Smirnov, O. (2017). An optimized algorithm for multiscale wideband deconvolution of radio astronomical images. *Monthly Notices of the Royal Astronomical Society*, 471(1):301–316.
- [Offringa, A. R. et al., 2019] Offringa, A. R., Mertens, F., van der Tol, S., Veenboer, B., Gehlot, B. K., Koopmans, L. V. E., and Mevius, M. (2019). Precision requirements for interferometric gridding in the analysis of a 21 cm power spectrum. A&A, 631:A12.
- [Paatero and Tapper, 1994] Paatero, P. and Tapper, U. (1994). Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2):111–126.

- [Parikh and Boyd, 2014] Parikh, N. and Boyd, S. (2014). Proximal algorithms. Foundations and Trends in optimization, 1(3):127–239.
- [Perley, 1999] Perley, R. A. (1999). Imaging with Non-Coplanar Arrays. In Taylor, G. B., Carilli, C. L., and Perley, R. A., editors, Synthesis Imaging in Radio Astronomy II, volume 180 of Astronomical Society of the Pacific Conference Series, page 383.
- [Picquenot et al., 2019] Picquenot, A., Acero, F., Bobin, J., Maggi, P., Ballet, J., and Pratt, G. (2019). Novel method for component separation of extended sources in x-ray astronomy. A&A, A139(627).
- [Raguet et al., 2013] Raguet, H., Fadili, J., and Peyré, G. (2013). A generalized forward-backward splitting. SIAM Journal on Imaging Sciences, 6(3):1199–1226.
- [Rapin, 2014] Rapin, J. (2014). Décompositions parcimonieuses pour l'analyse avancée de données en spectrométrie pour la santé. PhD thesis, Université Paris-Sud.
- [Rau, U. and Cornwell, T. J., 2011] Rau, U. and Cornwell, T. J. (2011). A multiscale multi-frequency deconvolution algorithm for synthesis imaging in radio interferometry. A&A, 532:A71.
- [Roddy and McEwen, 2021] Roddy, P. J. and McEwen, J. D. (2021). Sifting convolution on the sphere. *IEEE Signal Processing Letters*, 28:304–308.
- [Romano et al., 2017] Romano, Y., Elad, M., and Milanfar, P. (2017). The little engine that could: Regularization by denoising (RED). SIAM J. Imaging Sci., 10(4):1804–1844.
- [Sasao and Fletcher, 2005] Sasao, T. and Fletcher, A. B. (2005). Introduction to VLBI systems.
- [Scherzer et al., 2008] Scherzer, O., Grasmair, M., Grossauer, H., Haltmeier, M., and Lenzen, F. (2008). Variational Methods in Imaging. Springer Publishing Company, Incorporated, 1 edition.
- [Schwab, 1984] Schwab, F. R. (1984). Relaxing the isoplanatism assumption in selfcalibration; applications to low-frequency radio interferometry. *The Astronomical Journal*, 89:1076–1081.
- [Smith et al., 2017] Smith, D. M. P., Young, A., and Davidson, D. B. (2017). Reconsidering the advantages of the three-dimensional representation of the interferometric transform for imaging with non-coplanar baselines and wide fields of view. A&A, 603:A40.

- [Starck et al., 2002] Starck, J., Pantin, E., and Murtagh, F. (2002). Deconvolution in astronomy: A review. Publications of the Astronomical Society of the Pacific, 114(800):1051–1069.
- [Starck et al., 2005] Starck, J.-L., Elad, M., and Donoho, D. L. (2005). Image decomposition via the combination of sparse representations and a variational approach. *IEEE transactions on image processing*, 14(10):1570–1582.
- [Starck et al., 2015] Starck, J.-L., Murtagh, F., and Bertero, M. (2015). Starlet transform in astronomical data processing. In Scherzer, O., editor, *Handbook of Mathematical Methods in Imaging*, pages 2053–2098. Springer New York, New York, NY.
- [Starck et al., 2010] Starck, J.-L., Murtagh, F., and Fadili, J. (2010). Sparse Image and Signal Processing - Wavelets, Curvelets, Morphological Diversity. Cambridge University Press.
- [Subakan and Smaragdis, 2018] Subakan, Y. C. and Smaragdis, P. (2018). Generative adversarial source separation. In 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 26–30.
- [Sureau et al., 2020] Sureau, F., Lechat, A., and Starck, J.-L. (2020). Deep learning for a space-variant deconvolution in galaxy surveys. *Astronomy and astrophysics*, 11(A67).
- [Tasse, C. et al., 2018] Tasse, C., Hugo, B., Mirmont, M., Smirnov, O., Atemkeng, M., Bester, L., Hardcastle, M. J., Lakhoo, R., Perkins, S., and Shimwell, T. (2018). Faceting for direction-dependent spectral deconvolution. A&A, 611:A87.
- [Taylor et al., 1999] Taylor, G. B., Carilli, C. L., and Perley, R. A. (1999). Synthesis Imaging in Radio Astronomy II, volume 180 of Astronomical Society of the Pacific Conference Series.
- [Thompson et al., 2017] Thompson, A. R., Moran, J. M., and Swenson, George W., J. (2017). Interferometry and Synthesis in Radio Astronomy, 3rd Edition. Astronomy and Astrophysics Library. Springer.
- [Thouvenin et al., 2021] Thouvenin, P.-A., Abdulaziz, A., Dabbech, A., Repetti, A., and Wiaux, Y. (2021). Parallel faceted imaging in radio interferometry via proximal splitting (Faceted HyperSARA): when precision meets scalability. https://arxiv. org/abs/2003.07358.
- [Tseng, 2001] Tseng, P. (2001). Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of optimization theory and applications*, 109(3):475–494.

- [Vincent et al., 2006] Vincent, E., Gribonval, R., and Fevotte, C. (2006). Performance measurement in blind audio source separation. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(4):1462–1469.
- [Wakker and Schwarz, 1988] Wakker, B. P. and Schwarz, U. J. (1988). The Multi-Resolution CLEAN and its application to the short-spacing problem in interferometry. *Astronomy and Astrophysics*, 200(1-2):312–322.
- [Wiaux et al., 2009] Wiaux, Y., Puy, G., Boursier, Y., and Vandergheynst, P. (2009). Spread spectrum for imaging techniques in radio interferometry. *Monthly Notices* of the Royal Astronomical Society, 400(2):1029–1038.
- [Yang et al., 2016] Yang, Y., Sun, J., Li, H., and Xu, Z. (2016). Deep ADMM-net for compressive sensing MRI. In *Proceedings of the 30th International Conference* on Neural Information Processing Systems, NIPS'16, page 10–18, Red Hook, NY, USA. Curran Associates Inc.
- [Yu, 2013] Yu, Y.-L. (2013). On decomposing the proximal map. In Burges, C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K., editors, Advances in Neural Information Processing Systems, volume 26. Curran Associates, Inc.