Sparse reconstruction of the dark matter mass map from weak gravitational lensing

Francois Lanusse

NNT : 2015SACLS102

These de doctorat
de
L'Universite de Paris-Saclay
Preparee a
L'Universite Paris-Sud

Ecole Doctorale N° 127 :
Astronomie et Astrophysique d'Île-de-France

Spécialité de doctorat : Sciences de l'Univers

Par

François Lanusse

Reconstruction parcimonieuse de la carte de masse de matière noire par effet de lentille gravitationnelle

Thèse présentée et soutenue à Gif-Sur-Yvette, le 20 Novembre 2015 :

Composition du jury :

M. Yannick Mellier	Directeur de Recherche, Institut d'Astrophysique de Paris	Président
M. James Bartlett	Professeur, Université Paris Diderot	Rapporteur
M. Alan Heavens	Professeur, Imperial College London	Rapporteur
Mme. Nabila Aghanim	Directeur de Recherche CNRS, Université Paris Saclay	Examinatrice
Mme. Sophie Maurogordato	Directeur de Recherche CNRS, Observatoire de Nice	Examinatrice
M. Jean-Luc Starck	Directeur de Recherche, CEA	Directeur de Thèse
M. Benjamin Wandelt	Professeur, Université Pierre et Marie Curie	Invité

ABSTRACT

Keywords: weak gravitational lensing, cosmology, sparsity

Gravitational lensing, that is the distortion of the images of distant galaxies by intervening massive objects, has been identified as one of the most promising probes to help answer questions relative to the nature of dark matter and dark energy. As the lensing effect is caused by the total matter content, it can directly probe the distribution of the otherwise invisible dark matter. By measuring the shapes of distant galaxies and statistically estimating the deformations caused by gravitational lensing, it is possible to reconstruct the distribution of the intervening mass. This mass-mapping process can be seen as an instance of a linear inverse problem, which can be ill-posed in many situations of interest, especially when mapping the dark matter on small angular scales or in three dimensions. As a result, recovering a meaningful mass-map in these situations is not possible without prior information.

In recent years, a class of methods based on a so-called sparse prior has proven remarkably successful at solving similar linear inverse problems in a wide range of fields such as medical imaging or geophysics. The primary goal of this thesis is to apply these sparse regularisation techniques to the gravitational lensing problem in order to build next-generation dark matter mass-mapping tools. We propose in particular new algorithms for the reconstruction of high-resolution 2D mass-maps and 3D mass-maps and demonstrate in both cases the effectiveness of the sparse prior. We also apply the same sparse methodologies to the reconstruction the primordial density fluctuation power spectrum from measurements of the Cosmic Microwave Background which constitutes another notoriously difficult inverse problem. We apply the resulting algorithm to reconstruct the primordial power spectrum using data from the Planck satellite.

Finally, we investigate new methodologies for the analysis of cosmological surveys in spherical coordinates. We develop a new wavelet transform for the analysis of scalar fields on the 3D ball. We also conduct a comparison of methods for the 3D analysis of spectroscopic galaxy survey.

RESUME

Mots clés : Effet de lentille gravitationnelle faible, cosmologie, parcimonie

L'effet de lentille gravitationnelle, qui se traduit par une déformation des images nous parvenant de galaxies lointaines, constitue l'une des techniques les plus prometteuse pour répondre aux nombreuses questions portant sur la nature de l'énergie sombre et de la matière noire. Cet effet de lentille étant sensible à la masse totale, il permet de sonder directement la distribution de matière noire, qui resterait autrement invisible. En mesurant la forme d'un grand nombre de galaxies lointaines, il est possible d'estimer statistiquement les déformations causées par l'effet de lentille gravitationnelle puis d'en inférer la distribution de masse de la lentille. La reconstruction de ces cartes de masse constitue un problème inverse qui se trouve être mal posé dans un certain nombre de situations d'intérêt, en particulier lors de la reconstruction de la carte de masse aux petites échelles ou en trois dimensions. Dans ces situations, il devient impossible de reconstruire une carte sans l'ajout d'information a priori.

Une classe particulière de méthodes, basées sur un a priori de parcimonie, s'est révélée remarquablement efficace pour résoudre des problèmes inverses similaires pour un large champ d'applications tels que la géophysique et l'imagerie médicale. Le but principal de cette thèse est donc d'adapter ces techniques de régularisation parcimonieuses au problème de la cartographie de la matière noire afin de developper une nouvelle generation de méthodes. Nous développons en particulier de nouveaux algorithmes permettant la reconstruction de cartes de masse bi-dimensionnelles de haute résolution ainsi que de cartes de masse tri-dimensionnelles.

Nous appliquons de plus les mêmes méthodes de régularisation parcimonieuse au problème de la reconstruction du spectre de puissance des fluctuations primordiales de densités à partir de mesures du fond diffus cosmologique, ce qui constitue un problème inverse particulièrement difficile à résoudre. Nous développons un nouvel algorithme pour résoudre ce problème, que nous appliquons aux données du satellite Planck.

Enfin, nous investiguons de nouvelles méthodes pour l'analyse de relevés cosmologiques exprimés en coordonnées sphériques. Nous développons une nouvelle transformée en ondelettes pour champs scalaires exprimés sur la boule 3D et nous comparons différentes méthodes pour l'analyse cosmologique de relevés de galaxies spectroscopiques.

RÉSUMÉ DE LA THÈSE EN FRANÇAIS

RÉGULARISATION PARCIMONIEUSE DE PROBLÈMES INVERSES

Parcimonie

La parcimonie est un concept mathématique permettant de décrire une classe de signaux dont la représentation dans une base appropriée est essentiellement composée de coefficients nuls à l'exception d'un faible nombre de coefficients non nuls. De tels signaux sont dit parcimonieux. Cette notion en apparence très simple constitue en réalité un outil très puissant pour adresser une multitude de problèmes pratiques couvrant une large gamme de champs d'application.

Tout particulièrement, la parcimonie fournit un outil remarquable pour résoudre la classe particulière des problèmes inverses linéaires. Dans ces problèmes, le but est de reconstituer un signal inconnu à partir d'un jeu d'observations dégradées par l'effet d'un opérateur linéaire, généralement dit mal-posé i.e. non directement inversible ou possédant une inverse très instable. Dans ce contexte, il existe en général une infinité de solutions potentielles au problème, c'est à dire permettant de reproduire les observations, sans pour autant être proche du signal recherché. Il est généralement impossible d'estimer une solution proche de la vérité sans l'ajout d'information a priori sur le signal à estimer. Le concept de parcimonie fournit justement un cadre mathématique pour résoudre ce type de problèmes lorsque le signal à reconstruire peut être considéré comme parcimonieux. Dans ce cas, la solution du problème inverse peut être vue comme la solution d'un problème d'optimisation où l'on recherche le signal le plus parcimonieux compatible avec les observations.

Résoudre en pratique ce problème d'optimisation n'est pas chose facile mais grâce à de récentes avancées dans le domaine de l'optimisation convexe (théorie proximale) nous avons à présent à notre disposition une série d'algorithmes permettant de résoudre efficacement ces problèmes pouvant aussi intégrer différentes contraintes additionnelles. Ces algorithmes sont la clé de l'application des méthodes parcimonieuses à des problèmes de grande taille tels que ceux qui seront adressés dans cette thèse.

Ondelettes

Bien souvent, un signal n'est pas directement parcimonieux dans le domaine dans lequel il est observé, mais sa représentation dans un domaine approprié peut tout de même être parcimonieuse. C'est par exemple le cas d'une sinusoïde qui n'est pas parcimonieuse dans le domaine temporel, mais qui le devient dans le domaine fréquentiel, après transformée de Fourier. Cet exemple met en valeur la notion de représentation parcimonieuse.

De même que la transformée de Fourier est une représentation parcimonieuse pour des signaux stationnaires harmoniques, la transformée en ondelettes est une représentation parcimonieuse pour une vaste classe de signaux dits lisses par morceaux, incluant un grand nombre de signaux naturels. De par leur flexibilité et leur faible coût algorithmique, les ondelettes sont un outil de

base pour les techniques parcimonieuses. En particulier, dans toutes les applications présentées dans cette thèse, l'a priori de parcimonie est basé sur une représentation en ondelettes.

Application à la reconstruction du spectre de puissance primordial

La physique gouvernant les touts premiers instants de l'Univers est encore largement inconnue. Le modèle cosmologique standard nécessite une courte période d'accélération exponentielle appelée inflation, juste après la singularité initiale. Ce mécanisme d'inflation est en particulier essentiel pour expliquer l'origine des fluctuations de densité dans l'Univers primordial qui au fil du temps ont évolué principalement sous l'effet de la gravité pour former la structure à grande échelle visible aujourd'hui. Cependant, la physique de l'inflation n'est à ce jour pas élucidée et toute information permettant de discriminer entre divers modèles est d'un intérêt crucial pour la cosmologie moderne. En particulier, l'étude du spectre de puissance de ces perturbations primordiales peut permettre de discerner des traces particulières à certains modèles d'inflation, comme des oscillations ou d'autres sortes de traits localisés.

Malgré que le spectre de puissance primordial ne soit plus directement observable aujourd'hui, il est encore accessible indirectement en étudiant le spectre de puissance des fluctuations de densité observées plus tard dans l'age de l'Univers. C'est par exemple le cas du fond diffus cosmologique qui présente une image de ces fluctuations de densité après quelques 380 000 ans d'évolution sous l'influence de phénomènes physiques bien compris. Cette évolution pouvant être décrite comme un processus linéaire il est possible d'écrire le spectre de puissance du fond diffus cosmologique comme le résultat de l'action d'un opérateur linéaire sur le spectre de puissance primordial: estimer le spectre de puissance primordial devient un cas typique de problème inverse linéaire.

Malheureusement, résoudre ce problème inverse est non trivial en raison de plusieurs facteurs dont le fait que l'opérateur linéaire en question n'est pas inversible et surtout du fait de la variance cosmique (l'incertitude sur le spectre de puissance du fond diffus cosmologique venant du fait que nous devons l'estimer à partir d'une seule réalisation).

Nous proposons une nouvelle méthode pour résoudre ce problème inverse basée sur la régularisation parcimonieuse. Cette méthode nous permet de reconstruire le spectre de puissance primordial avec suffisamment de qualité pour discerner entre divers types de déviations par rapport au modèle le plus simple d'une loi de puissance. Nous démontrons l'efficacité de la méthode sur divers types de spectres en simulant des observations de la mission WMAP et nous l'appliquons aux données du satellite Planck. Malgré la sensibilité et la robustesse de notre méthode, nous ne détectons aucune déviation d'une simple loi de puissance. Ces résultats sont en parfait accord avec la dernière étude sur le sujet menée par la collaboration Planck.

CARTOGRAPHIE DE LA MATIÈRE NOIRE PAR EFFET DE LENTILLE GRAVITATIONNELLE FAIBLE

L'effet de lentille gravitationnelle faible

L'effet de lentille gravitationnelle se produit lorsque la lumière nous arrivant de galaxies lointaines passe à proximité de structure massives, agissant comme des lentilles en déformant l'espace temps

et causant au passage une déformation de l'image de ces objets lointains. Lorsque ces déformations restent faibles (on parle alors d'effet de lentille gravitationnelle faible), l'image qui nous arrive de ces galaxies se trouve très légèrement étirée selon une direction particulière, c'est ce que l'on appelle le cisaillement gravitationnel. En conséquence, la forme mesurée d'une galaxie (caractérisée par son ellipticité) peut être vue comme la somme de la forme intrinsèque de la galaxie et du champ de cisaillement à la position de cette galaxie. Si l'on suppose que les formes intrinsèques des galaxies sont distribuées de manière aléatoire on s'attend à ce que la moyenne de leurs ellipticités tende vers zero alors qu'au contraire les déformations causées par le cisaillement sont cohérentes et ne se moyennent pas à zero. Ainsi, mesurer la forme d'un grand nombre de galaxies nous permet d'avoir un estimateur bruité mais non baisé (du moins en première approximation) du champ de cisaillement.

L'intérêt cosmologique fondamental du cisaillement gravitationnel réside dans le fait qu'il est sensible à la masse totale de la lentille et donc en particulier à la matière noire qui représente la grande majorité du contenu total en matière de l'Univers (autour de 85%) mais qui reste complètement invisible autrement que par ses effets gravitationnels. Une des applications de cet effet de cisaillement est la cartographie de la distribution de masse d'une lentille (et donc majoritairement de la matière noire qui la constitue) en essayant d'inverser l'opérateur linéaire reliant cette distribution de masse au champ de cisaillement observé. Là encore, ce problème constitue un exemple typique de problème inverse linéaire.

Dans le cas où l'on cherche à reconstruire une carte bidimensionnelle d'une lentille situé à une distance bien connue, le problème n'est pas nécessairement mal posé. En effet, l'opérateur de cisaillement possède un inverse explicite et stable dans le cas où le problème est posé sur une grille régulière. Cependant, comme discuté plus en détails dans la prochaine section, en pratique le cisaillement n'est mesuré que là où l'on observe des galaxies en arrière plan, qui sont distribuées de façon aléatoire, ce qui complique l'inversion, particulièrement aux petites échelles.

Dans le cas où l'on cherche à reconstruire une carte tridimensionnelle de la distribution de matière à l'origine du champ de cisaillement, le problème devient extrêmement mal posé. Plusieurs méthodes ont été proposées dans la littérature pour essayer de reconstruire cette carte tridimensionnelle mais sans grand succès, ce qui a largement limité jusqu'à ce jour l'application de ce genre de techniques. En particulier, les cartes estimées par ces méthodes échouent à reconstruire correctement les structures le long de la ligne de visée, ainsi une lentille située à une distance spécifique se retrouve complètement étalée le long de la ligne de visée dans la carte reconstruite, si bien qu'il est très difficile d'en déduire les propriétés de la lentille, tels sa masse et sa distance.

Cartographie de haute résolution combinant cisaillement et flexion

Lorsque l'on cherche à reconstruire une carte bidimensionnelle d'une lentille à une distance connue, on se trouve confronté à deux difficultés. Premièrement, le champ de cisaillement n'est accessible qu'à la position particulière des galaxies en arrière plan, ce qui rend l'inversion de l'opérateur de cisaillement mal posé si l'on ne lisse les données sur une grille régulière au préalable. Deuxièmement, les mesures d'ellipticité sont très largement dominées par la forme intrinsèque des galaxies, ce qui rend l'estimation du champ de cisaillement extrêmement bruitée si l'on n'applique pas un lissage visant à réduire le niveau de bruit. Dans tous les cas, un simple lissage permet de reconditionner le problème et de réduire le niveau de bruit mais au prix de la perte de l'information aux petites échelles.

La méthode que nous proposons a pour but de reconstruire une carte de masse de haute résolution, c'est à dire en préservant les détails aux petites échelles lorsque ceux-ci sont discernables du bruit, en appliquant une régularisation parcimonieuse au problème inverse nous évitant d'avoir recours à tout lissage des données. De plus, afin de gagner en sensibilité, en particulier aux petites échelles, nous proposons de combiner l'effet de cisaillement gravitationnel avec l'effet de flexion gravitationnel (causant des déformations d'ordre supérieur) particulièrement sensible aux détails. Nous testons l'algorithme résultant sur un jeux de simulations correspondant à un relevé typique effectué avec le satellite Hubble. Nous montrons en particulier l'apport de la flexion qui nous permet de résoudre certaines sous-structures de petites tailles d'amas de galaxies qui ne peuvent pas être reconstruites uniquement à partir du cisaillement.

Cartographie tridimensionnelle

Contrairement au problème de la cartographie bidimensionnelle, la reconstruction tridimensionnelle de la carte de matière noire implique l'inversion d'un opérateur linéaire supplémentaire agissant selon la ligne de visée qui se trouve être particulièrement mal posé. Il est donc nécessaire de faire appel à une régularisation particulièrement efficace pour tenter de résoudre ce problème. Les méthodes de l'état de l'art reposent soit sur un re-conditionnement de cet opérateur en utilisant une décomposition en valeurs singulières, soit sur un filtrage de Wiener basé sur un a priori du spectre de puissance du signal recherché. Ces deux approchent se sont révélées particulièrement inefficaces. La première du fait que les valeurs singulières de faible amplitude restent fondamentales pour reconstruire une distribution de masse correctement localisée le long de la ligne de visée. La seconde du fait qu'un a priori sur le spectre de puissance n'est absolument pas adapté à la distribution de masse que l'on cherche à reconstruire, très difficilement caractérisable par son spectre de puissance.

Nous proposons donc une méthode basée sur l'a priori que la distribution de matière doit être parcimonieuse le long de la ligne de visée. En effet, le signal que l'on cherche à reconstruire est constitué d'halos de matière noire (autour des amas de galaxies) situés à des distances bien précises tandis qu'en dehors de ces zones particulières, la densité de matière est en moyenne très faible et peut être correctement approximée à zéro étant donné le niveau de bruit. Nous démontrons sur une simulation particulière que notre algorithme est capable de reconstruire correctement la densité de matière d'un halo massif (avec la bonne amplitude et correctement localisé en distance) alors que les méthodes précédentes produisent une distribution non localisée, à partir de laquelle il est impossible d'inférer directement la masse ou la distance de l'objet.

Nous caractérisons notre méthode sur un vaste jeux de simulations basées sur les caractéristiques de la mission Euclid et nous montrons que notre algorithme est capable de produire des cartes de densité en trois dimensions à partir desquels nous pouvons directement mesurer la distance des amas de galaxies ainsi qu'estimer leurs masses. Cette nouvelle approche ouvre donc la voie vers de nouvelles applications de la cartographie tridimensionnelle par cisaillement gravitationnel jusqu'alors impossibles en raison des limitations des méthodologies existantes.

COSMOLOGIE SUR LA BOULE

Ondelettes isotropes sur la boule 3D

Les relevés cosmologiques, et en particulier les relevés de galaxies, sont observés en coordonnées sphériques (position sur le ciel combinée à la mesure de la distance des objets par rapport à nous) et vivent donc naturellement sur la boule tridimensionnelle. Afin d'extraire un maximum d'information cosmologique il est nécessaire de développer des outils d'analyse appropriés à la géométrie particulière de ces relevés. Dans ce contexte, nous proposons une nouvelle construction en ondelettes directement définie sur la boule et donc particulièrement bien adaptée à des fins d'études cosmologiques.

Cette ondelette est la première ondelette sur la boule permettant à la fois analyse et synthèse. Reposant sur la décomposition de Fourier-Bessel sphérique, elle est construite par extension sur la boule des concepts développés sur la sphère utilisant la décomposition en harmoniques sphériques. Contrairement à d'autres constructions proposées ultérieurement dans la littérature, cette ondelette est non séparable, ce qui lui confère des propriétés telles que l'isotropie ou l'invariance par translation qui sont extrêmement désirables dans nombre d'applications.

Analyse tridimensionnelle de relevés de galaxies spectroscopiques

Les relevés de galaxies spectroscopiques fournissent d'excellentes sondes pour contraindre en particulier les paramètres de l'énergie sombre. Cependant, cette information est principalement renfermée dans l'évolution radiale de la distribution de galaxie. Il est donc primordial de correctement extraire l'information radiale pour exploiter tout le potentiel de cette sonde. L'approche classique à ce problème, dite tomographique, consiste à séparer le relevé en couches de rayons différents puis de calculer des spectres de puissances angulaires croisés entre les couches. Une autre approche, jusqu'ici beaucoup moins explorée, consiste à procéder à une décomposition harmonique sur la boule (transformée de Fourier-Bessel sphérique) du relevé puis de calculer le spectre de puissance des coefficients obtenus.

En théorie, ces deux approches devraient permettre de capturer la même information. Cependant, en pratique, elles ne sont pas sensibles de la même manière à un certain nombre d'effets et de systématiques et ne sont donc pas nécessairement équivalentes. Nous procédons à une analyse comparative à l'aide de matrices de Fisher afin d'investiguer l'impact de l'exclusion des échelles non-linéaires (difficiles à interpréter) ainsi que l'incertitude sur le biais des galaxies (source de systématiques) pour ces deux approches différentes.

Nous trouvons que l'analyse basée sur la décomposition de Fourier-Bessel est plus à même de distinguer l'incertitude sur le biais des galaxies de l'impact de l'énergie sombre et permet une réjection plus nette des échelles indésirables. Ceci nous amène à conclure que cette approche devrait être favorisée par rapport à une analyse tomographique classique pour les prochains relevés tel celui de la mission Euclid.

PUBLICATIONS

BOOK CHAPTER

- F. Lanusse, J.-L. Starck, A. Woiselle, and M.J. Fadili. 3D Sparse Representations. In *Advances in Imaging and Electron Physics*, Academic Press: Elsevier Inc, pages 99–204. 2014.

REFEREED ARTICLES

- F. Lanusse, J.-L. Starck, A. Leonard, and S. Pires. High resolution weak lensing mass-mapping combining shear and flexion. *Submitted to Astronomy & Astrophysics*.
- F. Lanusse, A. Rassat, and J.-L. Starck. 3D galaxy clustering with future wide-field surveys: Advantages of a spherical Fourier-Bessel analysis. *Astronomy & Astrophysics*, 578:A10, June 2015.
- A. Leonard, F. Lanusse, and J.-L. Starck. Weak lensing reconstructions in 2D and 3D: implications for cluster studies. *Monthly Notices of the Royal Astronomical Society*, 449(1):1146–1157, May 2015.
- A. Möller, V. Ruhlmann-Kleider, F. Lanusse, J. Neveu, N. Palanque-Delabrouille, and J.-L. Starck. SNIa detection in the SNLS photometric analysis using Morphological Component Analysis. *Journal of Cosmology and Astroparticle Physics*, 2015(04):041–041, April 2015.
- F. Lanusse, P. Paykari, J. L. Starck, F. Sureau, J. Bobin, and A. Rassat. PRISM: Recovery of the primordial spectrum from Planck data. *Astronomy and Astrophysics*, 571:L1, November 2014.
- P. Paykari, F. Lanusse, J. L. Starck, F. Sureau, and J. Bobin. PRISM: Sparse recovery of the primordial power spectrum. *Astronomy and Astrophysics*, 566:A77, June 2014.
- A. Leonard, F. Lanusse, and J.-L. Starck. GLIMPSE: accurate 3D weak lensing reconstructions using sparsity. *Monthly Notices of the Royal Astronomical Society*, 440(2):1281, May 2014.
- F. Lanusse, A. Rassat, and J. L. Starck. Spherical 3D isotropic wavelets. *Astronomy and Astrophysics*, 540(3):A92, April 2012.

CONFERENCE PROCEEDINGS

- F. Lanusse, A. Leonard, and J.-L. Starck. Density reconstruction from 3D lensing: Application to galaxy clusters. In *Proceedings of the International Astronomical Union*, volume 10, pages 104–106, May 2014.
- P. Paykari, F. Lanusse, J.-L. Starck, F. Sureau, and J. Bobin. PRISM: Sparse recovery of the primordial spectrum from WMAP9 and Planck datasets. In *Proceedings of the International Astronomical Union*, volume 10, pages 60–63, May 2014.

- A. Rassat, F. Lanusse, D. Kirk, O. Host, and S. Bridle. Combining Probes. In *Proceedings of the International Astronomical Union*, volume 10, pages 192–201, May 2014.
- F. Lanusse, A. Leonard, and J.-L. Starck. Imaging dark matter using sparsity. In Dimitri Van De Ville, Vivek K. Goyal, and Manos Papadakis, editors, *SPIE 8858, Wavelets and Sparsity XV*, page 885824, September 2013.
- F. Lanusse and J.-L. Starck. 3D sparse representations on the sphere and applications in astronomy. In Dimitri Van De Ville, Vivek K. Goyal, and Manos Papadakis, editors, *SPIE 8858, Wavelets and Sparsity XV*, page 88580K, September 2013.

SOFTWARE

- MRS3D: Implementation of the 3D Spherical Wavelet Transform on the Sphere in C++ with an IDL interface.
  http://jstarck.free.fr/mrs3d.html
- PRISM: Sparse recovery of the Primordial Power Spectrum implemented in C++ with an IDL interface, part of the iSAP package.
  http://www.cosmostat.org/research/cmb/prism
- CosmicPy: Python package for interactive cosmology, with an embedded fast C++ library for efficient computation of Spherical Fourier-Bessel power spectra.
  https://cosmicpy.github.io
- GLIMPSE: Implementation in C++ of the sparse reconstruction of 3D dark matter density from weak gravitational lensing.
  http://www.cosmostat.org/research/wl/glimpse

ACKNOWLEDGMENTS

I would like first of all to extend a very warm thank you to Y. Mellier, A. Heavens, J. Bartlett, N. Aghanim, S. Maurogordato, and B. Wandelt for doing me the honour of being part of my defense committee. I am particularly grateful to Alan Heavens and Jim Bartlett for taking the time to thoroughly review this manuscript.

I would like to acknowledge support from the CNES who co-financed, with the CEA, this thesis through a doctoral research grant.

Then of course, my next thought is for my PhD advisor, Jean-Luc, who taught me so much over the last 5 years, about wavelets and sparsity of course, but more importantly about scientific research in general and about the particular qualities necessary to become a successful researcher. I am also immensely grateful to Anais Rassat, Adrienne Leonard, and Paniez Paykari for all their help, support and guidance not only as we worked together on the various projects presented in this manuscript but also far beyond that.

Ces trois années passées à CosmoStat ont été formidables et je tiens à remercier tous les membres de l'équipe, Sandrine, Jérôme, Martin et Florent qui ont toujours été pret à répondre à mes questions et à apporter leur aide et leur expertise. Bien sûr, je dois aussi remercier tous mes co-thésards, en particulier pour m'avoir supporté pendant la rédaction, en commençant par Chieh-An avec qui j'ai partagé un bureau et de nombreux délires, mais aussi Fureddo-san, Ming et Cécile (ça va me manquer de ne plus entendre parler de A et S). Je pense aussi aux anciens, Daniel, Antoine, Simon (et ses wings), Rémy, Gabriel, Julien, Jérémy (qui m'a montré la voie de l'analyse) ainsi qu'aux nouveaux, Austin, Sam et Joana qui ont rendu ces derniers mois particulièrement sympathiques. Je remercie aussi tous les autres thésards et membres du Service d'Astrophysique dont Maelle, Tugdual, Corentin, Andrea, Bertrand, René et Pierre, pour n'en citer que quelques uns.

C'est de plus l'occasion de remercier les amis rencontrés en Master NPAC avec qui j'ai partagé beaucoup des joies et des peines de la thèse et sans qui ces quatre dernières années auraient été beaucoup plus ternes. Je pense tout spécialement à Anais qui ne se doutait pas de ce dans quoi j'allais l'entrainer mais qui a toujours conservé sa joie de vivre communicative. Une mention spéciale à Cyril et Camille qui m'ont entrainé un peu partout en France et au delà des océans. Et puis un grand merci à Alexis pour ses soirées GoT and Pizza ainsi qu'à Vivien, Clémentine, au grand JF et à tous les autres.

Ce manuscrit n'aurait peut-être pas vu le jour sans le groupe de rédaction du LAL avec qui j'ai partagé un été très intense en émotions. Un grand merci à Marta (Praticamente perfetta sotto ogni aspetto), Laura (Ti perdono per Word), Alexis, Marjia, Roman, Cyril et Camille (encore eux)... et puis Choppy aussi, même si je n'ai toujours pas bien compris ce qu'un première année fait au labo les dimanches après-midi en plein été... Et puis n'oublions pas ces figures qui ont marqué la rédaction comme Philippe et ses parties de cache-cache, Vianney dont on ne sait toujours pas où il était ou Shia Labeouf et sa motivation débordante.

J'en profite aussi pour remercier mes amis de ma vie précédente qui m'ont permis d'échapper à la thèse de temps en temps, en particulier merci à Vincent et surtout à Kevin et Marie !

Enfin, ça va presque sans dire mais je remercie bien évidemment toute ma famille qui m'a toujours soutenu et qui m'a permis d'en arriver là aujourd'hui.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ALGORITHMS

ACRONYMS

ACS Advanced Camera for Surveys

AIM Analytic Image Models

AKDE Adaptive Kernel Density Estimate

BAO Baryon Acoustic Oscillations

BPDN Basis Pursuit DeNoise

CDM Cold Dark Matter

CMB Cosmic Microwave Background

CWT Continuous Wavelet Transform

FB Forward-Backward

DCT Discrete Cosine Transform

DETF Dark Energy Task Force

DFT Discrete Fourier Transform

DL Dictionary Learning

DSBT Discrete Spherical Fourier-Bessel Transform

DWT Discrete Wavelet Transform

EOS Equation Of State

FFT Fast Fourier Transform

FIR Finite Impulse Response

FISTA Fast Iterative Shrinkage-Thresholding Algorithm

FLRW Friedmann-Lemaître-Robertson-Walker

FoM Figure of Merit

GFB Generalised Forward-Backward

GR General Relativity

GREAT3 third GRavitational lEnsing Accuracy Testing

HEALPix Hierarchical Equal Area isoLatitude Pixelization of a sphere

HST Hubble Space Telescope

ISTA Iterative Soft-Thresholding Algorithm

IUWT Isotropic Undecimated Wavelet Transform

LGMCA Local-Generalized Morphological Component Analysis

LSS Large Scale Structure

MAD Median Absolute Deviation

MCA Morphological Component Analysis

MCMC Monte Carlo Markov Chain

MRA MutliResolution Analysis

NDFT Non-equispaced Discrete Fourier Transform

NFFT Non-equispaced Fast Fourier Transform

PDF Probability Density Function

PPS Primordial Power Spectrum

PSF Point Spread Function

RSD Redshift Space Distortions

SBT Spherical Bessel Transform

SFB Spherical Fourier-Bessel

SFBT Spherical Fourier-Bessel Transform

SHT Spherical Harmonics Transform

SIH Statistically Isotropic and Homogeneous

SNLS SuperNova Legacy Survey

SNR Signal to Noise Ratio

SPT South Pole Telescope

STEP Shear TEsting Programme

STFT Short-Term Fourier Transform

SVD Singular Value Decomposition

UWT Undecimated Wavelet Transform

VST Variance Stabilisation Transform

WMAP Wilkinson Microwave Anisotropy Probe

1 INTRODUCTION

Cosmology has entered in the last decade a new era of precision. With the success of the WMAP and Planck missions, the standard cosmological model is now well established and its main parameters are measured at the percent level. Yet, in this model, conventional baryonic matter only accounts for a mere 5 % of the total energy content today, while an overwhelming 95 % of the Universe is composed of 27 % of dark matter and 68 % of dark energy. If these two components are necessary for the model to fit the observations, they remain largely mysterious and understanding the physical nature of this dark universe constitutes one of the main challenges of modern cosmology.

Gravitational lensing, that is the distortion of the images of distant galaxies by intervening massive objects, has been identified as one of the most promising probes to help answer questions relative to the nature of dark matter and dark energy. Indeed, as the lensing effect is caused by the total matter content, it can directly probe the distribution of dark matter. In the weak lensing regime, where those deformations are small, there is a direct linear mapping between the dark matter distribution and the measured lensing effect. Reconstructing the dark matter map from weak lensing therefore constitutes a linear inverse problem which can be ill-posed in practice when mapping the dark matter in three dimensions and when part of the data is masked. In these situations, the information is degraded in a way that makes it impossible to recover a meaningful dark matter map without additional prior information.

Ill-posed linear inverse problems such as dark matter mass-mapping are in fact ubiquitous and well studied in the signal processing literature. In recent years, a class of methods based on a so called sparse prior has proven remarkably successful at solving inverse problems in a wide range of fields such as medical imaging or geophysics. These methods rely on the assumption that the signals to recover are sparse, i.e. can be represented with a small number of non-zero coefficients, and use this information to constrain the solution of the inverse problem. As the sparsity of a signal depends on the basis used to represent it, a key aspect of these methods is the choice of an appropriate sparse representation for the signals of interests.

The primary goal of this thesis is therefore to apply these sparse regularisation techniques to the weak lensing problem in order to build next-generation dark matter mass-mapping tools which can eventually help us address the challenges of modern cosmology.

The aim of the first part of this thesis is to introduce the concepts and tools of sparse regularisation. In Chapter 3 I introduce the notion of sparsity, how it can be applied in practice to solve inverse problem, and the sparse optimisation algorithms used in the rest of this thesis. In complement, I provide in Chapter 4 an introduction to wavelets, which are a particularly successful and versatile class of sparse representations, at the core of all the applications presented in the following chapters. Equipped with these tools I present in Chapter 5 a first application of these methods to the problem of the recovery of the power spectrum of primordial perturbations. In this canonical example of a linear inverse problem, the aim is to use measurements of the anisotropies of the

Cosmic Microwave Background we observe today to reconstruct this primordial power spectrum, which carries valuable information about the physics at play in the primordial universe.

In the second part of this thesis I focus on the particular problem of weak lensing mass-mapping. After introducing the weak lensing framework in Chapter 6, I first address in Chapter 7 the problem of the high-resolution 2D mass-mapping of galaxy clusters from weak lensing alone. I develop a new approach, based on sparse regularisation which makes use of higher order lensing information in order to improve the resolution of the mass map and allow the detection of small substructure. The ability to reconstruct mass maps at these resolutions from weak lensing alone can be of considerable interest for investigating the nature of dark matter. In Chapter 8 I consider the more difficult problem of reconstructing the dark matter density in three dimensions, using tomographic weak lensing information. The sparse reconstruction method I develop in this chapter is capable of recovering both the masses and distances of dark matter halos and represents a significant improvement over previous state-of-the-art techniques, opening the door to new applications for this kind of analysis.

The last part of this thesis focuses on the 3D analysis of cosmological surveys in spherical coordinates. By nature, cosmological surveys which include distance estimates, such as spectroscopic galaxy surveys, live on the 3D ball. Preserving this spherical geometry in the subsequent analysis is crucial to avoid unnecessary mixing of purely radial or angular effects. In Chapter 9, I develop the first practical wavelet transform on the 3D ball which allows analysis and synthesis in both continuous and discrete settings. This isotropic, shift invariant wavelet is ideally suited for the analysis of galaxy surveys. Finally, I present in Chapter 10 a comparison of methods for the 3D analysis of spectroscopic galaxy survey. The first method investigated is a partial, tomographic expansion of the data while the second is a full harmonic expansion on the 3D ball. As demonstrated, the full 3D expansion is more robust to unavoidable systematic effects than a more standard tomographic analysis, especially for constraining dark energy parameters.

2 COSMOLOGICAL CONTEXT

CONTENTS

The purpose of this first chapter is to provide the general cosmological background for the rest of this thesis. We introduce in particular the standard cosmological model and present some of the current challenges of modern cosmology. This chapter is complemented by Chapter 6 which will be dedicated to gravitational lensing.

2.1 THE ΛCDM MODEL

In this section, we briefly describe the cosmological concordance model, which has been so far extremely successful at describing cosmological observations. This model is based on two primary components, *Cold Dark Matter* (CDM) and a non-vanishing cosmological constant Λ, and is consequently known as ΛCDM.

In the currently commonly accepted picture of the Universe, shortly after an initial singularity, the Universe underwent a very brief period of exponential expansion, known as *inflation* (see Section 2.3.1), which imprinted small fluctuations of quantum origin on an otherwise extremely smooth background. This short period of inflation ended around 10^{-32}s after the initial singularity, marking the beginning of an era of much slower expansion and progressive cooling. After this initial phase, the Universe was filled with a hot plasma of baryons, electrons, and photons. The mean free path of photons remained extremely short during a long time and this plasma was therefore opaque. As the Universe expanded it slowly cooled down, until the temperature reached the 3000 K mark, about 380,000 years after the Big Bang, at which point electrons and protons combined to form neutral atoms. During this event, called *recombination*, the mean free path of photons suddenly increased to reach the order of the Hubble scale and the Universe became transparent. The first light emitted at this epoch is still visible to this day as the *Cosmic Microwave*

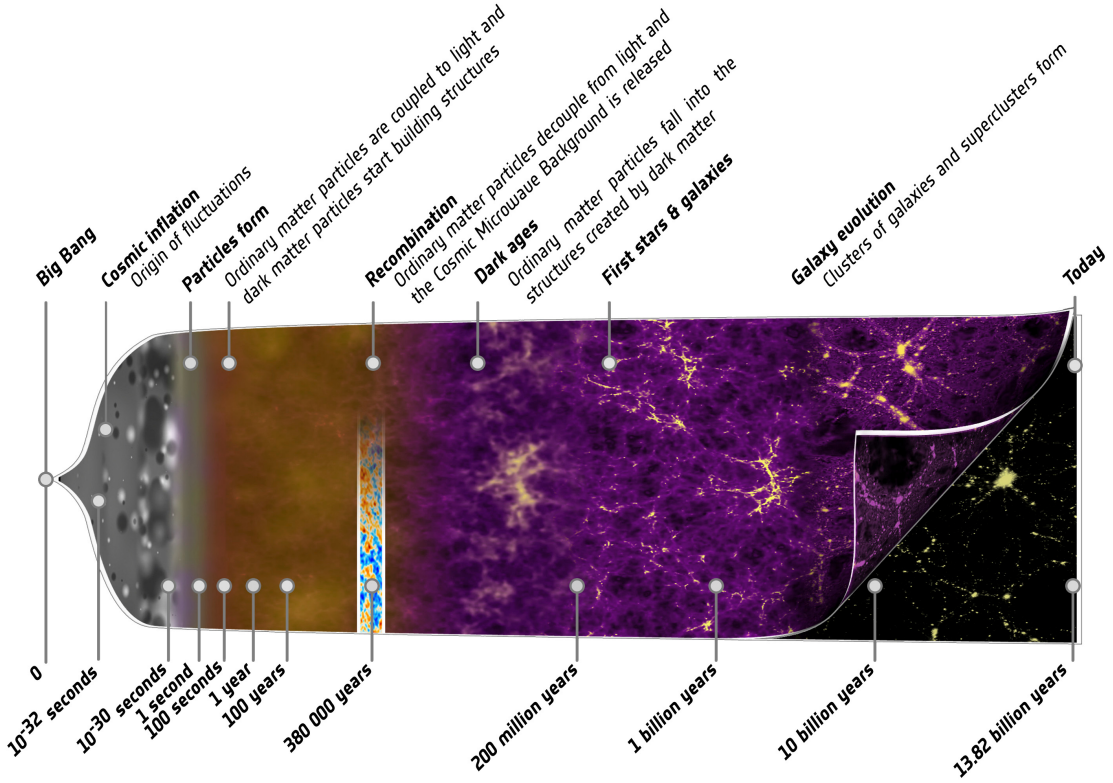

Figure 2.1: Cosmic history of the Universe from the Big Bang to present day. Credit: ESA - C. Carreau

Background (CMB), now observed in the microwave domain at a temperature of 2.725 K and identified for the first time by Penzias and Wilson (1965). The evolution of the Universe then remained dominated by its matter content, allowing the growth of the initial density perturbations into the large scale structure we observe today, through gravitational collapse. About 4 billion years ago, the Universe entered a new phase of accelerated expansion, believed to be driven by dark energy, which will eventually stop the growth of structures. Our understanding of this history of the Universe is summarised on Figure 2.1.

In order to describe this evolution, the ΛCDM model assumes that the Universe is composed of the following fluids, each behaving and influencing the dynamics of the Universe in a distinct way.

DARK ENERGY The total energy content of the Universe today is dominated at 69.11 % by a mysterious dark energy. As we will see in the following section, this energy has negative pressure and is driving a new phase of accelerated expansion (Riess et al., 1998; Perlmutter et al., 1999). So far, this energy remains undistinguishable from a non-vanishing cosmological constant Λ in Einstein's equations. One the main goals of the next generation of cosmological surveys will be to determine whether dark energy is indeed due to a cosmological constant or if it behaves as a dynamical field.

COLD DARK MATTER The second largest contribution to the total energy content of the Universe, accounting for 25.84 %, is due to *dark matter*. This non-baryonic matter does not interact through the electromagnetic interaction, and is therefore invisible, hence its *dark* denomination. If the true

Parameter name	Symbol	Value
Hubble constant	H_0	67.74 ± 0.46
Dark energy density	Ω_Λ	0.6911 ± 0.0062
Total matter density	Ω_m	0.3089 ± 0.0062
Physical Baryon density	$\Omega_b h^2$	0.02230 ± 0.00014
Power spectrum normalisation	σ_8	0.8159 ± 0.0086
Spectral index	n_s	0.9667 ± 0.0040
Reionisation optical depth	τ	0.066 ± 0.012

Table 1: Main cosmological parameters 68 % confidence limits for the ΛCDM from Planck Collaboration et al. (2015b). These parameters are derived from a combination of Planck CMB power spectra with the lensing reconstruction and external data (BAO+JLA+H_0).

nature of dark matter remains unknown, it is an essential component of the standard cosmological model and although it has not be directly detected or produced to this day, its existence is necessary to explain a number of observations, in particular, the famous weak lensing mass-mapping of the Bullet cluster (Clowe et al., 2004, 2006). Several flavours of dark matter have been proposed, but the currently favoured model is *Cold Dark Matter* (CDM), where the *cold* denomination implies that dark matter particles are non-relativistic.

BARYONIC MATTER The remaining 4.86 % are composed of ordinary baryonic matter, essentially in the form of hydrogen and helium gas in the intergalactic medium.

RELATIVISTIC SPECIES Finally, the Universe today contains trace amounts of electromagnetic radiation, mainly due to the CMB photons. Cosmological neutrinos can also contribute to this component.

One of the success of the ΛCDM model is its ability to describe this evolution of the Universe and reproduce all cosmological observations from a small set of parameters, the most relevant of which are summarised on Table 1.

2.2 THE HOMOGENEOUS ISOTROPIC UNIVERSE

2.2.1 Cosmological General Relativity and the FLRW metric

Our comprehension of the Universe and its evolution relies on Einstein's theory of *General Relativity* (GR). In this geometrical theory of gravity, what appears like a gravitational force is only the consequence of the curvature of the 4 dimensional spacetime. In turn the curvature of spacetime will be generated by its energy-momentum content. Applied to cosmological scales, GR makes the link between the geometry of the Universe and its matter and energy content and describes their dynamical evolution with time.

The local geometry of spacetime is entirely described by the metric tensor $g_{\mu\nu}$ which allows the computation of distances between two points by defining a line element $\mathrm{d}s^2$:

$$\mathrm{d}s^2 = g_{\mu\nu}\,\mathrm{d}x^\mu\,\mathrm{d}x^\nu \tag{2.1}$$

where x^μ is some set of coordinates with indices ranging from 0 to 3, with by convention 0 being the time dimension and 1-3 the spatial dimensions. As a special example, let us consider a flat spacetime. In this case, the metric reduces to the Minkowski metric $\eta_{\mu\nu}$ of special relativity defined in Cartesian coordinates by:

$$\eta = \begin{pmatrix} -c^2 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{2.2}$$

In GR, in the absence of external forces, the inertial motion of particles (including massless photons) in a curved spacetime is determined by the *geodesic equation*:

$$\frac{\mathrm{d}^2x^\mu}{\mathrm{d}\lambda^2} + \Gamma^\mu_{\alpha\beta}\frac{\mathrm{d}x^\alpha}{\mathrm{d}\lambda}\frac{\mathrm{d}y^\beta}{\mathrm{d}\lambda} = 0 \tag{2.3}$$

where the affine parameter λ is a scalar parametrising the particle along its trajectory and $\Gamma^\mu_{\alpha\beta}$ are Christoffel symbols, function of the metric $g_{\mu\nu}$ and its partial derivatives:

$$\Gamma^\mu_{\alpha\beta} = \frac{1}{2}g^{\mu\nu}\left(\frac{\partial g_{\nu\alpha}}{\partial x^\beta} + \frac{\partial g_{\nu\beta}}{\partial x^\alpha} - \frac{\partial g_{\alpha\beta}}{\partial x^\nu}\right) \tag{2.4}$$

Trajectories which verify Equation (2.3) are known as *spacetime geodesics*. This equation will be especially important to present the gravitational lensing effect in Chapter 6 as it describes how light propagates in a curved spacetime and in particular how images from distant galaxies are lensed by the intervening large scale structure of the Universe. In the particular case of a flat spacetime, described by the metric $\eta_{\mu\nu}$, the Christoffel symbols simply vanish, as the components of the metric are constants, and Newton's first law of motion can be recovered:

$$\frac{\mathrm{d}^2x^\mu}{\mathrm{d}t^2} = 0 \tag{2.5}$$

In this case, the geodesic is simply a straight line as one would expect in the absence of external forces.

If the motion of free falling objects can be explained by the curvature of spacetime, the question is now what causes spacetime to acquire curvature in the first place. This is described in GR by the famous Einstein field equations which relate the curvature to the local energy and momentum content of spacetime through the following expression:

$$G_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} \tag{2.6}$$

where $T_{\mu\nu}$ is the energy-momentum tensor and $G_{\mu\nu}$ is the Einstein tensor which measures the curvature of spacetime and can be expressed in terms of the Christoffel symbols and their derivatives

but remains ultimately a function of the metric $g_{\mu\nu}$. The energy-momentum tensor describes the local density and flux of energy and momentum of all the content of spacetime, including matter and radiation. Solving analytically the Einstein field equations without any additional considerations such as symmetries of the system is generally not possible. In particular, applying these equations to describe the evolution of the entire Universe will require some simplifying assumptions, as we will see now.

One of the cornerstones of cosmology is the *Cosmological Principle* which implies that the Universe is both *isotropic* and *homogeneous* on large scales. Although this principle can seem at odds with our own direct experience of the Universe, which is extremely inhomogeneous at our scale, it is remarkably well verified on large scales by CMB and Large Scale Structure (LSS) observations.

Thanks to the Cosmological Principle, the Einstein field equations can be greatly simplified. First and foremost, as was shown independently by Robertson and Walker in the mid 1930s, the metric describing an homogeneous and isotropic Universe is unique and is known as the Friedmann-Lemaître-Robertson-Walker (FLRW) metric:

$$\mathrm{d}s^2 = -c^2\,\mathrm{d}t^2 + a^2(t)\left(\mathrm{d}r^2 + f_K(r)^2\,\mathrm{d}\Omega^2\right) \tag{2.7}$$

where r is the time independent *comoving distance*, $a(t)$ is the *scale factor* describing an isotropic scaling of the cosmological comoving coordinate system with time and the function $f_K(r)$ is the transverse comoving distance which depends on the curvature K of the Universe:

$$f_K(r) = \begin{cases} K^{-1/2}\sin\left(K^{-1/2}r\right) & \text{for } K > 0 \text{ (spherical)} \\ r & \text{for } K = 0 \text{ (flat)} \\ |K|^{-1/2}\sinh\left(|K|^{-1/2}r\right) & \text{for } K < 0 \text{ (hyperbolic)} \end{cases} \tag{2.8}$$

The second simplification implied by the Cosmological Principle is that the matter-energy content of the Universe can be described as a perfect fluid which is completely characterised in a comoving inertial frame by a simple energy-momentum tensor of the form:

$$T^{\mu\nu} = \begin{pmatrix} \rho & 0 & 0 & 0 \\ 0 & p & 0 & 0 \\ 0 & 0 & p & 0 \\ 0 & 0 & 0 & p \end{pmatrix} \tag{2.9}$$

where ρ is the *energy density* and p the *pressure* of the fluid. As we will see in the next section, the total energy density and pressure can be separated into different contributions for the various components of the Universe, each with their own equations of state.

Thanks to the simplification of both the metric and the energy-momentum made possible by the Cosmological Principle, the Einstein field equations can now be used to describe the dynamical evolution of the Universe.

2.2.2 The dynamics of the universe

Combining the FLRW metric and the energy-momentum tensor of a perfect fluid in Einstein field equations yields two independent equations relating the dynamical evolution of the Universe to its matter and energy content. These are known as the Friedmann equations:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{Kc^2}{a^2} \tag{2.10}$$

$$\frac{\ddot{a}}{a} = -\frac{4}{3}\pi G\left(\rho + \frac{3p}{c^2}\right) \tag{2.11}$$

Let us introduce the dimensionless Hubble parameter $H(t) = \dot{a}(t)/a(t)$ whose value at present time $t = t_0$, i.e. for $a = 1$, is the *Hubble constant* $H_0 = \dot{a}(t_0)$. The Hubble constant therefore corresponds to the speed of the expansion of the Universe today. H_0 is usually defined in terms of the dimensionless reduced Hubble constant h according to:

$$H_0 = 100\ h\ \mathrm{km\ s^{-1}\ Mpc^{-1}} \tag{2.12}$$

Given the first Friedmann equation, the Hubble parameter at any given time or equivalently, any given scale factor can simply be deduced from the total energy density ρ. To describe the evolution of ρ with the scale factor, both Friedmann equations can be combined to yield the following conservation relation:

$$\frac{\mathrm{d}\rho}{\mathrm{d}a} + \frac{3}{a}\left(\rho + \frac{p}{c^2}\right) = 0 \tag{2.13}$$

For the various components of the Universe, one can define an *Equation Of State* (EOS) relating the pressure of the fluid to its energy density through a parameter w:

$$p = c^2 w\rho\ . \tag{2.14}$$

In particular, for non-relativistic matter (cold dust), $w_m = 0$ while for relativistic species $w_r = 1/3$. For these two particular cases, the EOS is constant and Equation (2.13) is readily solved:

$$\rho_r(a) = \rho_{r,0}a^{-4} \quad ; \quad \rho_m(a) = \rho_{m,0}a^{-3} \tag{2.15}$$

As one can see from the acceleration equation Equation (2.11), a fluid with an EOS with $w < -1/3$ would drive an acceleration of the expansion. In particular, a cosmological constant can be interpreted as a fluid with a fixed EOS $w_\Lambda = -1$, in which case its energy density remains constant with time. However, in order to investigate the nature of dark energy, one can let its EOS free to vary with time, which leads to the following expression:

$$\rho_{de} = \rho_{de,0}a^{f(a)} \qquad \text{with} \qquad f(a) = -\frac{3}{\ln a}\int_0^{\ln a}(1 + w_{de}(a'))\,\mathrm{d}\ln a' \tag{2.16}$$

To restrict the functional space of $w_{de}(a)$, several parametrisation have been used in the literature. Throughout this work we will be using the common Linder (Linder, 2003) parametrisation $w_{de}(a) = w_0 + w_a(1-a)$. This parametrisation makes explicit a possible dependence of the EOS with time through the w_a parameter and would reduce to a simple cosmological constant for $w_0 = -1$ and $w_a = 0$. For this specific form of the EOS, the function f becomes $f(a) = -3(1+w_0) - 3w_a\left(1 + \frac{1-a}{\ln(a)}\right)$. The total energy density $\rho(t)$ can be expressed as the sum of

$$\rho(a) = \rho_m(a) + \rho_r(a) + \rho_{de}(a) \tag{2.17}$$

$$= \rho_{r,0}a^{-4} + \rho_{m,0}a^{-3} + \rho_{de,0}a^{f(a)} \tag{2.18}$$

One can see from Equation (2.10) that for a specific value ρ_c of the density ρ, the curvature of the Universe K vanishes. This value is known as the *critical density* and can be expressed as:

$$\rho_c(a) = \frac{3H^2(a)}{8\pi G} \tag{2.19}$$

It is common to rescale the density ρ by the critical ρ_c and define a dimensionless density parameter Ω:

$$\Omega(a) = \frac{\rho(a)}{\rho_c(a)} \tag{2.20}$$

which leads to

$$\Omega(a) = \Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_{de} a^{f(a)} \tag{2.21}$$

By definition of the critical density, $\Omega = 1$ would correspond to a flat universe otherwise, for non-zero curvature $\Omega \neq 1$. The effect of curvature can conveniently be interpreted as an additional energy density by defining a curvature density parameter $\Omega_k(a) = 1 - \Omega(a)$.

$$H^2(a) = H_0^2 \left[\Omega_r a^{-4} + \Omega_m a^{-3} + \Omega_k a^{-2} + \Omega_{de} a^{f(a)}\right] \tag{2.22}$$

2.2.3 Distances

Redshift

In an expanding Universe described by the FLRW metric it is easily shown that the light coming to us from distant sources undergo a shift in frequency which can be expressed as the ratio of the scales factor at the emission and observation:

$$\frac{\lambda_{\text{obsv}}}{\lambda_{\text{emit}}} = \frac{a_{\text{obsv}}}{a_{\text{emit}}} \tag{2.23}$$

In the absence of expansion $a_{\text{obsv}} = a_{\text{emit}}$ and the wavelength of the light is not affected as it propagates through the Universe. However, if the Universe expands, the scale factor at the time of observation is larger than it was when the light was emitted by the distant source, in which case $\lambda_{\text{obsv}} > \lambda_{\text{emit}}$: the spectrum is shifted towards the red. This shift in wavelength is quantified by the dimensionless *cosmological redshift* z defined as:

$$z = \frac{\lambda_{\text{obsv}} - \lambda_{\text{emit}}}{\lambda_{\text{obsv}}} \tag{2.24}$$

Given that the scale factor is normalised at $a = 1$ at present time, the measured redshift z of a distant source is directly linked to the scale factor at the time of emission by:

$$a = \frac{1}{1+z} \tag{2.25}$$

The redshift of distant sources can directly be linked to their distance through the cosmological model and therefore this effect provides us with an invaluable tool to probe the depth of the Universe.

This link between redshift and distance was first established by Edwin Hubble in 1929 who observed a roughly linear relation between the velocity and the distance of "extra-galactic nebulae" (Hubble, 1929) which is known as Hubble's law : $v = H_0 r$. This linear law is only true for the local Universe and can be recovered from the definition of the Hubble parameter. On cosmological scales, the evolution of the Hubble parameter with time needs to be taken into account.

Comoving distance

The comoving distance r has been introduced with the formulation of the FLRW metric and corresponds to the distance between two points in the cosmological comoving coordinate system. As the Universe is expanding, the physical proper distance between objects is always fluctuating (increasing) but the comoving distance remains unaffected by the expansion by factoring out the scale factor a.

As was mentioned in the previous paragraph, the cosmological model allows us to relate the observed redshift of a distant source to its comoving distance with respect to the observer. The expression of this relation can be worked out from the definition of the metric $\mathrm{d}r = a^{-1}c\,\mathrm{d}t$ and the definition of the Hubble parameter $\mathrm{d}t = a^{-1}H(a)^{-1}\,\mathrm{d}a$:

$$r(a) = \int_a^1 \frac{c\,\mathrm{d}a'}{a'^2 H(a')}\,. \tag{2.26}$$

This expression gives us the comoving distance of an object given that we observe now light that it emitted when the scale factor of the Universe was equal to a. Equivalently this relation can be expressed in term of the observable redshift, given Equation (2.25) :

$$r(z) = \int_0^z \frac{c\,\mathrm{d}z'}{H(z')}\,. \tag{2.27}$$

Angular diameter distance

The angular diameter distance, as its name implies, is linked to another fundamental approach to measuring distances based on geometrical considerations. In Euclidian geometry, the distance to an object of a given size can simply be related to its apparent angular diameter. In complete analogy, the angular diameter distance D_A is defined so that an object of physical size Δ observed at a redshift z will have an apparent angular size $\delta\theta$ according to:

$$D_A(z) = \frac{\Delta}{\delta\theta} \tag{2.28}$$

This distance can directly be expressed in terms of the transverse comoving distance $f_K(r(a))$ defined in Equation (2.8):

$$D_A(a) = a f_K(r(a)) = \frac{f_K(r(z))}{z} \tag{2.29}$$

Luminosity distance

A last essential measure of distances is the *luminosity distance*, which is based on the scaling of the flux of a distant source with its distance. It is defined so that the observed flux F of a source at distance D_L with an intrinsic luminosity L correspond to:

$$F = \frac{L}{4\pi D_L^2} \tag{2.30}$$

This distance can be expressed in terms of the transverse comoving distance f_K or the equivalently in terms of the angular diameter distance as:

$$D_L(a) = a^{-1} f_K(r(a)) = a^{-2} D_A(a) \tag{2.31}$$

$$= (1+z) f_K(r(z)) = (1+z)^2 D_A(z) \tag{2.32}$$

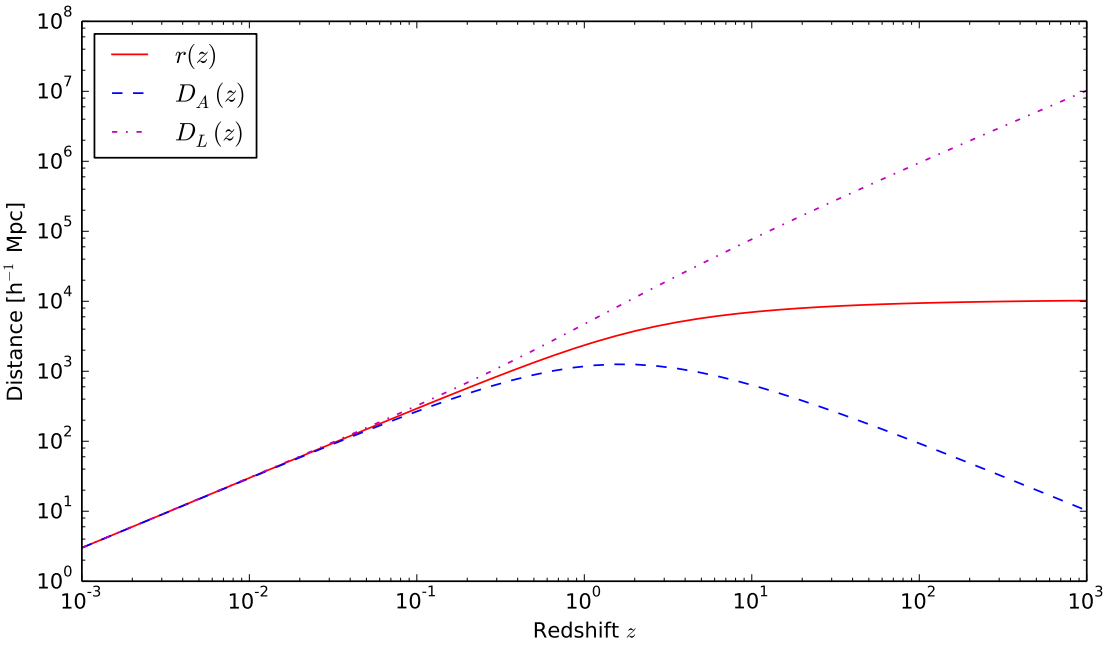

Figure 2.2: Cosmological distances as a function of redshift in a flat universe with $\Omega_m = 0.25$ and $\Omega_{de} = 0.75$.

The three cosmological distances presented in this section are plotted as a function of redshift on Figure 2.2.

2.3 LINEAR STRUCTURE FORMATION

The Universe described in the previous section is smooth and homogeneous, which if course is at odds with observations of the LSS. It is however possible to describe the formation and evolution of these structures by treating them as small fluctuations on top of the smooth FLRW background which leads to the linear theory outlined in this section.

2.3.1 Primordial perturbations

The structures we observe today can be described by the gravitational collapse of small inhomogeneities present in the early Universe. The commonly accepted mechanism to explain the origin of these perturbations is known as inflation (Guth, 1981; Linde, 1982) which corresponds to a short period of exponential expansion underwent by the Universe, ending approximately 10^{-30}s after the initial singularity. In simple inflationary models, this expansion is driven by a scalar field, the *inflaton*, which can exert a negative pressure (much like dark energy) for a limited amounted of time as the field slowly rolls towards a ground state, at which point inflation stops. In this process, quantum fluctuations of the inflaton are imprinted in the metric and reach macroscopic scales under this extremely fast period of expansion. As inflation ends, these perturbations of quantum origin remain and will act as seeds for the eventual formation of large scale structures.

The scale dependence of these perturbations are encoded by the *Primordial Power Spectrum* (PPS), generally expressed in the following parametric form:

$$P_{\mathrm{p}}(k) = A_s \left(\frac{k}{k_p}\right)^{n_s - 1 + \frac{1}{2}\alpha_s \ln(k/k_p)} , \tag{2.33}$$

where k_p is a given pivot scale, A_s is the overall amplitude of the power spectrum, n_s is known as the *spectral index* and α_s is an optional *running* parameter. An exactly scale-invariant PPS, known as the Harrison-Zeldovich model, which sets $n_s = 1$ and $\alpha_s = 0$ (Harrison, 1970; Zeldovich, 1972), is now significantly disfavoured by CMB constraints (Planck Collaboration et al., 2015b), and a near scale-invariant spectrum is instead preferred by current observations.

To this day, there is no direct evidence to confirm the validity of the inflationary paradigm, but inflation is nonetheless a remarkably successful phenomenological model. In addition to providing an explanation for the origin of primordial perturbations, it was realised by Guth that a period of exponential expansion of the Universe could explain some of the major issues in cosmology not directly addressed by the standard model, in particular the horizon and flatness problems described below.

THE HORIZON PROBLEM As mentioned in the previous section isotropy and homogeneity are part of the fundamental assumptions of cosmology. Yet on the largest scales the Universe appears simply too isotropic. In particular, the level of isotropy of the CMB is baffling (temperature anisotropies are at the 10^{-5} level) as it would have required regions of the universe at opposite position on the sky to have been in causal contact in the past, in order to reach a common thermal equilibrium, which should not have been possible under matter or radiation domination. The period of exponential expansion associated with inflation conveniently solves this problem by allowing the entire observable Universe to have originally occupied a very small volume. Within this small volume, even the largest scales observable today would have been in causal contact and reached thermal equilibrium before the onset of inflation.

THE FLATNESS PROBLEM Current observations seem to suggest that the Universe today is in fact flat, with $\Omega_K \approx 0$. In itself this result is not paradoxical, however the curvature parameter steadily increases with the scale factor during matter and radiation domination, which means that in order for the curvature to be small today, the Universe would have needed to be exceptionally flat in the past (Ω_K of the order of 10^{-16} at the epoch of nucleosynthesis). Such small values seem oddly specific and pose an acute fine-tuning problem. Inflation provides a mechanism to explain why the Universe was so flat at early times as curvature is naturally diluted during a period of exponential expansion. As a result, any non-zero initial curvature can be sufficiently suppressed by the end of inflation to fall within the current constraints, thus eliminating the fine-tuning problem.

2.3.2 Linear growth

As discussed above, inflation can be invoked to produce small density perturbations in an otherwise homogeneous universe. These primordial perturbations will act as seeds which will grow, driven by gravitational collapse, to ultimately form the structures present in our Universe today. In the limit where these perturbations remain small, their growth can be described by a linear evolution as we will describe in this section.

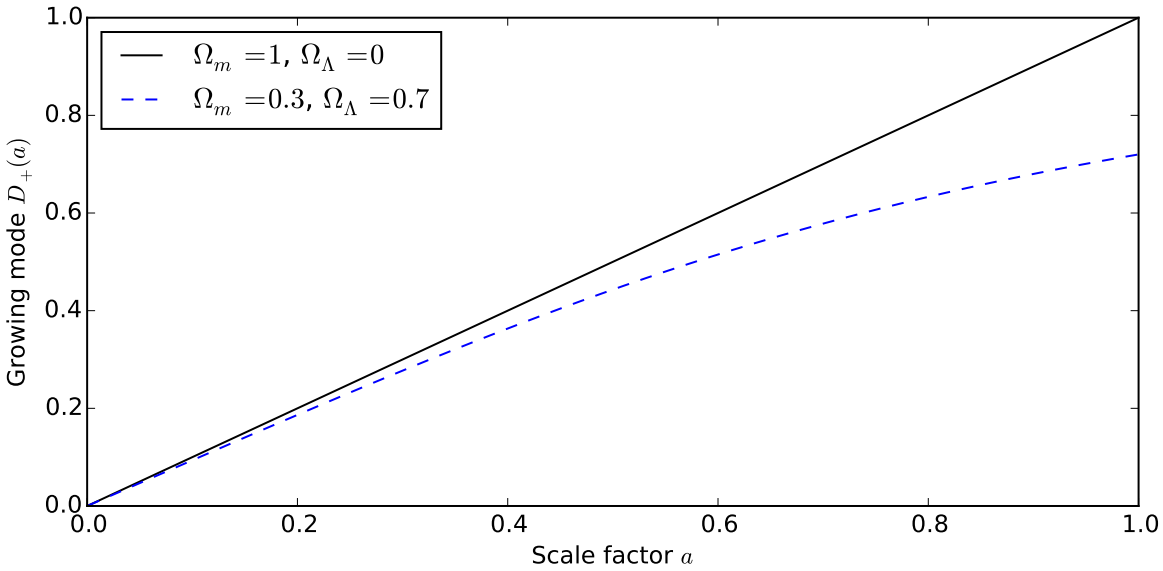

Figure 2.3: Linear growth factor as a function of scale factor for an Einstein-de Sitter universe ($\Omega_m = 1$) and a flat universe with a cosmological constant. The presence of a non vanishing cosmological constant suppresses the growth of structure when the universe leaves matter domination.

Let us introduce the *matter density contrast* $\delta(\boldsymbol{x}, a)$ at comoving position $\boldsymbol{x}$ and at scale factor a, defined in terms of the matter density ρ_m as:

$$\delta(\boldsymbol{x}, a) = \frac{\rho_m(\boldsymbol{x}, a) - \bar{\rho}_m(a)}{\bar{\rho}_m(a)} \,, \tag{2.34}$$

where $\bar{\rho}_m(a)$ is the mean matter density at scale factor a. In the limit of small perturbations $\delta \ll 1$ (verified on large scales and at early times), the evolution of the density contrast can be described by linear perturbation theory. It can be shown (e.g. Dodelson, 2003; Peebles, 1980) that in the limit of small peculiar velocities and for collisionless dark matter, the density contrast follows a differential equation of the form:

$$\ddot{\delta} + 2H\dot{\delta} - 4\pi G\bar{\rho}_m\delta = 0 \,. \tag{2.35}$$

In this very simple differential equation, the matter density $\bar{\rho}_m$ can be interpreted as a source term while the Hubble parameter H behaves as a damping term. The growth of structure is therefore suppressed by the expansion of the Universe. Another important point is that this equation allows the decoupling of spatial and temporal coordinates and accepts general solutions of the form:

$$\delta(\boldsymbol{x}, a) = D_+(a)f_1(\boldsymbol{x}) + D_-(a)f_2(\boldsymbol{x}) \,, \tag{2.36}$$

where f_1 and f_2 are particular initial conditions while D_+ and D_- are respectively a growing and decaying mode. We are only interested in the growing mode D_+, as D_- eventually decays and does not affect structure formation. The general solution of this growing mode is:

$$D_+(a) = \frac{5\Omega_m}{2}\frac{H(a)}{H_0}\int_0^a \frac{\mathrm{d}a'}{(a'H(a')/H_0)^3} \,. \tag{2.37}$$

This expression is normalised so that $D_+(a) = a$ during matter domination. Figure 2.3 illustrates the impact of dark energy on this growth factor, which clearly suppresses growth compared to a

matter dominated Universe. For the rest of this thesis, we will define the normalised linear growth factor D as:

$$D(a) = \frac{D_+(a)}{D_+(a=1)} \,. \tag{2.38}$$

This normalisation ensures that the growth factor today is equal to 1, so that the density contrast $\delta(\boldsymbol{x}, a)$ at a given time a can simply be described in terms of the density contrast today $\delta(\boldsymbol{x})$ as:

$$\delta(\boldsymbol{x}, a) = \delta(\boldsymbol{x}) D(a) \,. \tag{2.39}$$

2.3.3 Matter power spectrum

As can be demonstrated using stationarity and isotropy arguments, the Fourier modes of a *Statistically Isotropic and Homogeneous* (SIH) field are uncorrelated. In particular, as the density contrast $\delta(\boldsymbol{x})$ can be expected to verify this SIH condition as a consequence of the Cosmological Principle, this implies that the two-point correlation function of the 3D Fourier transform of the density field can be expressed as:

$$< \delta(\boldsymbol{k})^* \delta(\boldsymbol{k'}) >= (2\pi)^3 \, P(k) \, \delta_D(\boldsymbol{k} - \boldsymbol{k'}) \,. \tag{2.40}$$

where δ_D is the Dirac delta function and $P(k)$ is defined as the 3D matter power spectrum. One of the main cosmological parameters, σ_8, quantifies the normalisation of this power spectrum and is defined by the variance of the density contrast within a window of $8h^{-1}$Mpc.

Under linear evolution, the Fourier modes of the density contrast involve independently, and this evolution can be described in terms of a linear transfer function $T(k)$, defined in Fourier space as:

$$T(k) = \frac{\delta(k, a=1)}{\delta(k, a=0)} \frac{\delta(k=0, a=0)}{\delta(k=0, a=1)} \tag{2.41}$$

This function therefore accounts for the modulation of each Fourier mode between their original value at the end of inflation ($a = 0$) and their current value ($a = 1$). This includes not only the total linear growth described previously but also accounts for the different times of horizon crossing of different scales and most importantly for *Baryon Acoustic Oscillations* (BAO) effects, which are due to the propagation of acoustic waves in the baryon-photon plasma before recombination. The proper computation of this transfer function requires the use of Boltzmann codes such as CLASS (Lesgourgues, 2011) or CAMB (Lewis et al., 2000) but a number of analytical fitting formulae are also available when an exact computation is not essential. In particular, we use in Chapter 10 the common fitting formula from Eisenstein and Hu (1998). This transfer function allows us to relate the power spectrum today $P(k)$ to the PPS, according to:

$$P(k) = T^2(k) \, P_{\rm p}(k) \,. \tag{2.42}$$

This linear transfer function only accounts for the evolution of density fluctuations in the linear regime. On small scales, and at late times, the evolution of the density contrast becomes non-linear and this simple description of the growth of structure breaks down. Non linear corrections can be added to the linear power spectrum using for instance the Halofit code corresponding to Smith et al. (2003). The impact of these non linearities is an enhancement of the power spectrum on small scales. A comparison of linear and non-linear power spectra is shown in Figure 2.4.

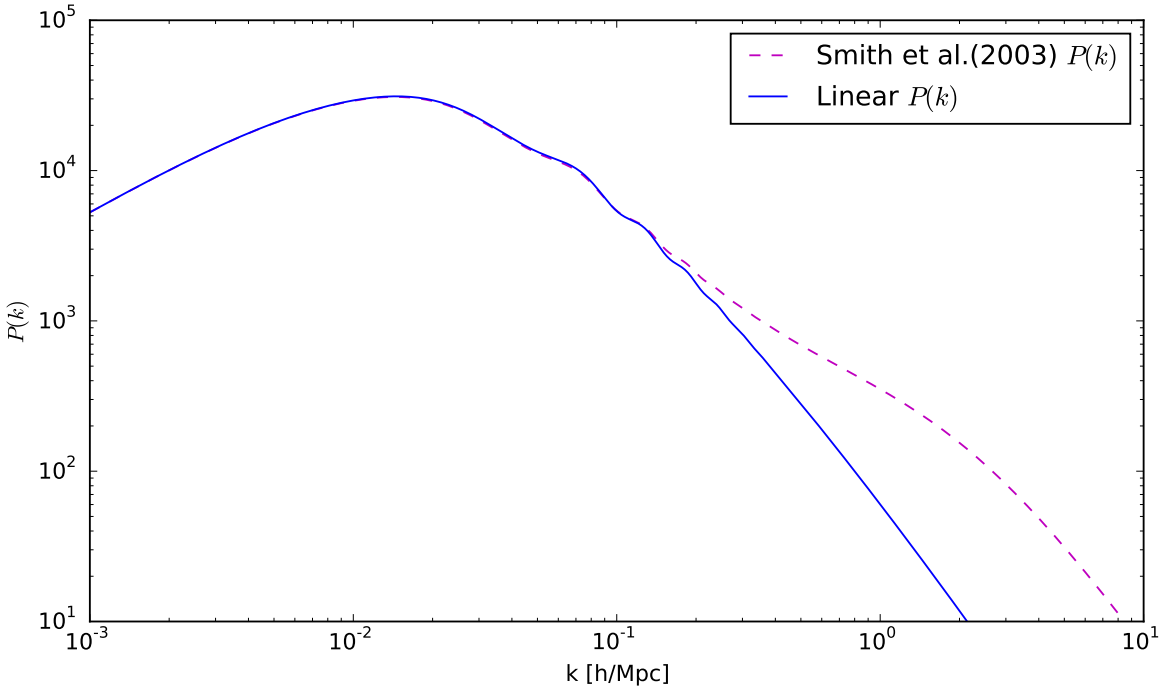

Figure 2.4: Matter power spectrum as a function of scale factor. The linear power spectrum computed with the fitting formula from Eisenstein and Hu (1998) is compared to the non-linear power spectrum computed following Smith et al. (2003). In both cases, BAOs are causing characteristic wiggles on the power spectrum.

2.4 CONCLUSION

If the parameters of ΛCDM are now well constrained (see Table 1), understanding the nature of dark matter, dark energy and the physics of inflation are the main challenges of modern cosmology.

In the rest of this thesis, we will rely on this standard cosmological model and develop new tools that can be used to help address the pressing questions of modern cosmology. In Chapter 5 we will propose a new tool for probing inflationary models by reconstructing the PPS. In Chapter 7 and Chapter 8 we will develop new tools for mapping the dark matter distribution, which can eventually be used to investigate different dark matter models and constrain cosmological parameters. Finally in Chapter 10 we present a comparison of methods for the 3D analysis of galaxy surveys, with a particular focus on constraints on the dark energy EOS.

Part I

SPARSE REGULARISATION OF INVERSE PROBLEMS

3 SPARSITY

CONTENTS

Sparsity is a concept used to describe signals which, when expressed in an appropriate basis, can be represented with a small number of coefficients. More fundamentally, this apparently simple notion implies that the intrinsic number of degrees of freedom of a sparse signal is low, which makes in fact a very strong statement about the nature of the signal. Sparsity provides a mathematical framework to characterise this class of signals and leverage their properties to address a wide range of practical problems.

A particular class of problems efficiently solved using sparsity are linear inverse problems. In such problems, one aims at recovering an unknown signal from a set of observations, degraded by a linear operator. As we will see in this chapter, if the signal to recover is known to be sparse, it can be robustly estimated from the measurements, even if the linear operator involved is not formally invertible.

The aim of this chapter is to introduce the mathematical concept of sparsity, how it can be applied to the regularisation of inverse problems, and most importantly, provide the algorithmic tools allowing us to efficiently solve these problems in practice. These tools, combined with the sparse representations introduced in Chapter 4, will form the core of the methods developed in the rest of this thesis to address various problems in the cosmological context.

3.1 INTRODUCTION TO SPARSITY

The concept of sparsity is very akin to Occam's razor: if two hypotheses can explain equally well the observations, the one with the fewest assumptions (i.e. the simplest one) should be preferred. Sparsity provides a mathematical framework to quantify the complexity of potential solutions to an inverse problem and therefore identify the simplest one.

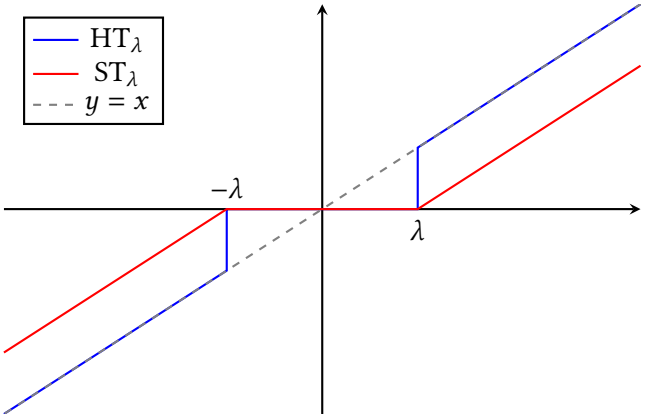

Figure 3.1: Response of the Hard Thresholding (blue) and Soft Thresholding (red) operators. Hard Thresholding only sets to zero coefficients with an amplitude lower than λ. Soft Thresholding shrinks the amplitude of coefficients by λ and sets to zero coefficients of amplitude lower than λ.

A signal is considered to be sparse when most of its coefficients are zero. More precisely, if we consider a discrete signal $\boldsymbol{x} \in \mathbb{R}^N$ then $\boldsymbol{x}$ is said to be exactly k-sparse if only k of its coefficients $x_i|_{i\in[\![1\dots N]\!]}$ are non zero. Behind this definition is the notion that for two signals living in the same space, a sparse signal has a lower number of degrees of freedom than a non-sparse signal.

Quantifying the sparsity of a given signal $\boldsymbol{x}$ is simply a matter of counting the number of active (i.e. non-zero) coefficients of $\boldsymbol{x}$. To describe this measure mathematically, we introduce the ℓ_0 pseudo-norm, noted $\|\cdot\|_0$, which counts the number of active coefficients of a signal. The ℓ_0 norm can formally be defined by considering the active support of a signal $\text{Supp}(\boldsymbol{x})$, the set of indices of non-zero coefficients, defined as

$$\text{Supp}(\boldsymbol{x}) = \left\{ i \in [\![1, N]\!] \quad | \quad x_i \neq 0 \right\} . \tag{3.1}$$

Therefore $\| \boldsymbol{x} \|_0$ is simply the number of elements of the set $\text{Supp}(\boldsymbol{x})$ which leads to the following definition

$$\forall \boldsymbol{x} \in \mathbb{R}^N, \quad \| \boldsymbol{x} \|_0 = \text{Card}(\text{Supp}(\boldsymbol{x})) , \tag{3.2}$$

where $\text{Card}(X)$ is the cardinality of set X.

In practice however, this mathematical definition of exactly k-sparse signals is not always adapted to natural signals. Indeed, a signal may have only a few coefficients with a high amplitude, yet small coefficients are seldom identically zero, in which case its ℓ_0 norm is potentially high. Nevertheless, such a signal can be considered as *compressible*, or weakly sparse, in the sense that it can be well approximated by a strictly k-sparse signal. Indeed, a good approximation of such a signal but with a lower ℓ_0 norm can be built by setting negligible coefficients to zero and keeping only the k most significant coefficients.

To mathematically describe this operation, one can introduce the Hard Thresholding operator HT_λ which sets to zero coefficients of amplitude lower than a given threshold $\lambda \in \mathbb{R}$:

$$\forall i \in [\![1, N]\!], \quad \text{HT}_\lambda(\boldsymbol{x})_i = \begin{cases} x_i & \text{if } |x_i| \geq \lambda \\ 0 & \text{otherwise} \end{cases} . \tag{3.3}$$

The response of the Hard Thresholding operator is illustrated in Figure 3.1. Using this operator, one can build a strictly k-sparse non-linear approximation $\tilde{\boldsymbol{x}}_{|_k} = \text{HT}_{\lambda_k}(\boldsymbol{x})$ of a signal $\boldsymbol{x}$ by adjusting

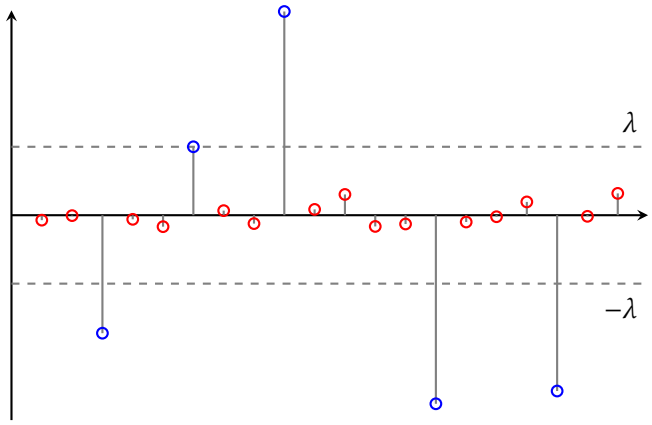

Figure 3.2: Example of sparse approximation for a compressible signal. By setting the threshold λ to the amplitude of the fifth largest coefficient, only significant (blue) coefficients are preserved after applying HT_λ while non-significant coefficients (red) are set to zero.

the level of the threshold λ_k to the amplitude of the k-th largest coefficient of $\boldsymbol{x}$. Provided that the original signal $\boldsymbol{x}$ had only k significant coefficients (i.e. coefficients of non negligible amplitude) the error induced by the non-linear approximation will be small. Figure 3.2 illustrates this concept of sparse approximation of a compressible signal. Only five coefficients (marked in blue) are significant while the rest of the coefficients (marked in red) have negligible amplitude. By setting the threshold value λ to the amplitude of the fifth largest coefficient, after Hard Thresholding all but the five most significant coefficients will be set to zero.

The quality of this non-linear approximation can be quantified by introducing a distance d between two signals and measuring the distance between the original signal and its approximation. Depending on the specific application various distances can be used but the most common is based on the quadratic ℓ_2 norm which quantifies the energy of the approximation error. Given two signals $\boldsymbol{x}$ and $\boldsymbol{y}$ in $\mathbb{R}^N$, the distance d between these two signals is

$$d(\boldsymbol{x}, \boldsymbol{y}) = \| \boldsymbol{x} - \boldsymbol{y} \|_2^2 \ . \tag{3.4}$$

Therefore, one can quantify the degradation of the signal caused by a sparse approximation as $d(\tilde{\boldsymbol{x}}_{|_k}, \boldsymbol{x})$.

Using these concepts, we can formally define a class of signals said to be *compressible* if the sorted magnitude of their coefficients decays quickly according to a power law. If $i(n)$ is the index of the n-th largest coefficient of $\boldsymbol{x}$, such that $|x_{i(0)}| \geq |x_{i(1)}| \geq \ldots \geq |x_{i(N)}|$, then a signal is compressible if there exists $C \in \mathbb{R}^+$ and $s > 1/2$ such that:

$$\forall n \in [\![1, N]\!], \quad |x_{i(n)}| \leq C n^{-s} \ . \tag{3.5}$$

With this definition, the larger the value of s, the faster the decay of the coefficients. For compressible signals, the k-sparse non-linear approximation error can be bounded:

$$d(\tilde{\boldsymbol{x}}_{|_k}, \boldsymbol{x}) \leq \frac{C^2}{2s-1} k^{-2s+1} \ . \tag{3.6}$$

This upper bound guarantees that the faster the decay of the coefficients, the smaller the non-linear approximation error is when approximating a compressible signal by an exactly k-sparse signal.

As a result, most of the developments presented in the following for exactly sparse signals still hold to a very good approximation if the signals are in fact compressible.

Nevertheless, although the notion of strict sparsity can be extended to compressibility, most natural signals are neither sparse nor compressible in the space in which they are measured. For instance, a sinusoidal signal sampled in the time domain is clearly not sparse as most of its coefficients are of significant amplitude. However, through a Fourier transform, such a signal can equivalently be represented in the frequency domain, in which it is extremely sparse (a sine is 1-sparse in the frequency domain).

This leads to the fundamental idea that the sparsity of a signal will be intrinsically linked to the domain used to represent it. Therefore finding appropriate signal representations that maximise the sparsity of certain classes of signals is an active and important research field for the application of sparse methodologies.

3.1.1 Sparse signal representation

Any given signal can be represented in a variety of domains without loss of information, but depending on the properties of the signal and on the application, a given representation may be preferred. In particular, finding signal representations that maximise the sparsity of the coefficients is of special interest in the perspective of applying sparse methodologies.

A signal can be modelled as a linear combination of elementary templates called *atoms*. A family of atoms which can span the functional space in which the signal to represent lives is called a *dictionary*. More formally, given a signal $\boldsymbol{x} \in \mathbb{R}^N$, $\boldsymbol{x}$ can be represented in a dictionary $\Phi \in \mathbb{R}^{N \times P}$ with $P \geq N$ as a linear combination of P atoms $\boldsymbol{\phi}_i \in \mathbb{R}^N$:

$$\boldsymbol{x} = \Phi \boldsymbol{\alpha} = \sum_{i=1}^{P} \boldsymbol{\phi}_i \alpha_i \ , \tag{3.7}$$

where $\boldsymbol{\alpha} \in \mathbb{R}^P$ are the coefficients of $\boldsymbol{x}$ in dictionary Φ. An important distinction can be made at this point between redundant and non-redundant dictionaries. If the atoms of Φ form a basis of $\mathbb{R}^N$, the size of the dictionary is $N \times N$ and the decomposition $\boldsymbol{\alpha}$ is unique and non redundant as $\boldsymbol{\alpha}$ has the same size as $\boldsymbol{x}$. On the contrary, when $P \geq N$, the atoms of the dictionary are not linearly independent and the decomposition $\boldsymbol{\alpha}$ is no longer unique. On one hand, non-redundant representations will lead to simpler sparse optimisation problems but, on the other hand, over-complete dictionaries will offer more flexibility in their design to better suit specific applications.

A fundamental example of non-redundant representation is the *Discrete Fourier Transform* (DFT):

$$\boldsymbol{x} = \mathbf{F}\hat{\boldsymbol{x}} \ , \tag{3.8}$$

where $\hat{\boldsymbol{x}}$ denotes the DFT of vector $\boldsymbol{x}$ and $\mathbf{F}$ is the Fourier matrix defined as :

$$F_{n,k} = \frac{1}{\sqrt{N}} e^{i2\pi kn/N} \ . \tag{3.9}$$

The Fourier transform is the classical example of sparse representation for stationary periodic signals. Indeed, a sine, which is not sparse in the time domain, becomes exactly 2-sparse in the

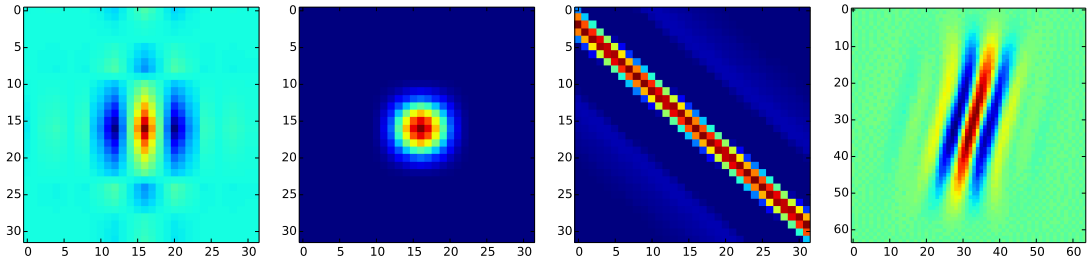

(a) Orthogonal wavelet **(b)** Isotropic wavelet **(c)** Ridgelet **(d)** curvelet

Figure 3.3: Examples of atoms from various dictionaries. As a general rule, to achieve a sparse representation, one should choose a dictionary whose atoms correspond to the morphology of the signal.

Fourier domain. Furthermore, the DFT benefits from efficient *Fast Fourier Transform* (FFT) algorithms with a complexity in $O(N \log N)$. We will also mention the very closely related *Discrete Cosine Transform* (DCT) defined by:

$$\alpha_k = \sum_{n=0}^{N-1} x_n \cos\left(\frac{\pi}{N}(n+\frac{1}{2})k\right) \tag{3.10}$$

and its inverse formula:

$$x_k = \frac{2}{N}\left(\frac{1}{2}\alpha_0 + \sum_{n=1}^{N-1} \alpha_n \cos\left(\frac{\pi}{N}(k+\frac{1}{2})n\right)\right). \tag{3.11}$$

The DCT dictionary can be extended to 2D discrete signals as a separable product of the DCTs along each direction. This dictionary is particularly well suited to represent textures and can be efficiently implemented using fast cosine transform algorithms with a complexity in $O(N \log N)$. This makes the DCT a good candidate for image compression applications.

A major drawback of these dictionaries is that their atoms are non local. As a result, they are not efficient at representing local features or more generally non stationary signals. These limitations can be mitigated by analysing the signal in windows small enough so that the properties of the signal can be assumed to remain stationary, which leads to the *Short-Term Fourier Transform* (STFT) or the block-DCT, used in the popular JPEG format.

Yet an even more efficient solution to overcome these limitations is to resort to wavelets. Wavelets are by construction designed to probe a signal in both time and frequency and provide sparse representations for piecewise smooth signals. Again, efficient algorithms are available for their computation, with complexities in $O(N)$ compared to the $O(N \log N)$ of the FFT or the DCT. In general, wavelets lead to very sparse representations of natural images and have replaced the DCT in the more efficient JPEG 2000 format. Orthogonal and redundant wavelet transforms will be introduced in detail in Chapter 4 and most of the applications presented in this thesis will be based on wavelet dictionaries.

A plethora of more complex sparse representations have been developed over the years since the first wavelets[1] with the aim of designing dictionaries adapted to specific applications or classes of signals. For instance, one of the shortcomings of wavelets is their lack of directionality which makes them rather inefficient at representing highly anisotropic features such as edges on 2D images. This has led to the development of ridgelets(Candes and Donoho, 1999) and curvelets (Starck

1 See http://tinyurl.com/wits-wavelets-starlet for a great compilation.

et al., 2002; Candès and Donoho, 1999) which are respectively extremely effective representations for lines and edges. Examples of typical atoms of these dictionaries are shown on Figure 3.3.

An alternative to the traditional design of sparse representations was proposed in Aharon et al. (2006) and consists in directly learning sparse representations from examples of the signal. This approach, known as *Dictionary Learning* (DL), involves solving an optimisation problem to find a dictionary which leads to the sparsest representation of the signals in the training set. Because these dictionaries are specifically built to optimise the sparsity of a given signal, they usually outperform more generic dictionaries such as wavelets (see for instance Beckouche et al. (2013) in the context of astronomical image denoising).

Figure 3.4 illustrates the sparse representation of a natural image using an orthogonal wavelet dictionary. The sorted amplitudes of both the pixels of the image and its wavelet coefficients are plotted in Figure 3.4b. As can be seen, the amplitude of the coefficients decays much faster in the wavelet domain than in the direct domain. In particular, because of the power law decay of the wavelet coefficients, this image can be considered compressible in this wavelet dictionary. The bottom panels show the non-linear approximation of the input image setting 99% of the coefficients to 0 in the direct and wavelet domains. Whereas, the remaining 1% most significant wavelet coefficients retain most of the information, it is not the case of the 1% most significant pixels.

3.2 SPARSE REGULARISATION OF INVERSE PROBLEMS

As mentioned at the beginning of this chapter, a key application of the sparsity concept is the regularisation of linear inverse problems. Such problems cover a very wide range of applications and typically involve recovering an unknown signal from a series of linear measurements. Most of the time, either because of noise or degeneracies in the linear operator involved in the measurements, the solution of the inverse problem is not unique and additional information is required to recover the signal of interest. In this context, sparsity offers a framework to use the morphology of the signal as a powerful regularising prior.

3.2.1 Basis Pursuit Denoise

In this section, we consider general linear problems of the form:

$$\boldsymbol{y} = \mathbf{A}\boldsymbol{x} + \boldsymbol{n} \,, \tag{3.12}$$

where $\boldsymbol{x} \in \mathbb{R}^N$ is an unknown signal to recover, $\boldsymbol{y} \in \mathbb{R}^M$ contains the measurements and $\boldsymbol{n} \in \mathbb{R}^M$ is an additive noise, assumed to be Gaussian with diagonal covariance $\sigma^2\mathbf{I_d}$. Finally $\mathbf{A}$ is a bounded linear operator which degrades the signal $\boldsymbol{x}$ and which is typically ill-behaved so that $\mathbf{A}^{-1}$ does not exist or is extremely unstable with respect to the noise. For instance, in the case of a deconvolution problem, $\mathbf{A}$ is a convolution by a blurring kernel which removes the high-frequency details of $\boldsymbol{x}$. Another example is the inpainting problem where $\mathbf{A}$ is a binary mask which sets to zero portions of the signal $\boldsymbol{x}$. In both instances, the operator is not invertible and without additional prior information, an infinite number of solutions $\tilde{\boldsymbol{x}}$ are possible.

To constrain the space of possible solutions, we consider as additional prior information that the signal $\boldsymbol{x}$ is sparse in an adapted dictionary $\boldsymbol{\Phi}$ i.e. there exists a set of coefficients $\boldsymbol{\alpha}$ sparse such

(a) Input image

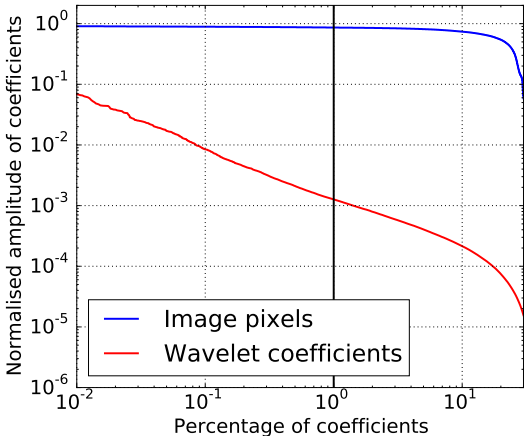

(b) Sorted amplitude of coefficients

(c) Non-linear approximation using 1 % of wavelet coefficients

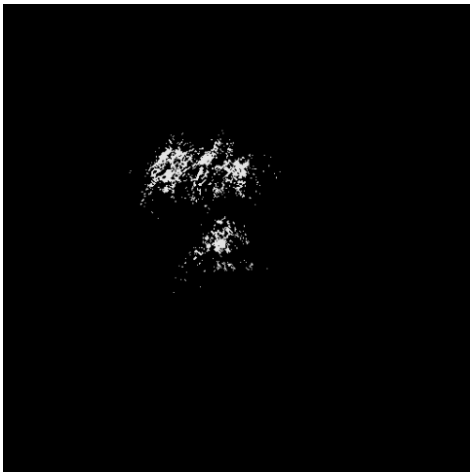

(d) Non-linear approximation using 1 % of image pixels

Figure 3.4: Non-linear approximation of a natural image using an orthogonal wavelet dictionary. The sorted amplitudes of the image pixels and corresponding wavelet coefficients are shown in (b). This image is not sparse in the direct domain but the fast decay of its wavelet coefficients indicate that it is compressible in this wavelet dictionary. The vertical black line indicates the first percentile of the coefficients. The non-linear approximations obtained by setting to zero all coefficients except for the first percentile are shown in (c) and (d). Credit: Input image from NASA/APL/SwRI

that $\boldsymbol{x} = \boldsymbol{\Phi\alpha}$. Thus the inverse problem can be regularised by imposing that the solution not only has to fit the data but also needs to be sparse in the dictionary $\boldsymbol{\Phi}$.

Finding the sparsest solution that fits the observations can be stated as an optimisation problem of the form:

$$\arg\min_{\boldsymbol{\alpha}} \| \boldsymbol{\alpha} \|_0 \quad s.t. \quad \| \boldsymbol{y} - \mathbf{A\Phi\alpha} \|_2^2 \leq \epsilon \,, \tag{3.13}$$

where $\epsilon > 0$ is a parameter based on the level of noise which relaxes the data fidelity term to avoid over-fitting the noise. As introduced at the beginning of this chapter, the ℓ_0 norm in the left-hand side of this expression is a measure of the sparsity of the solution and will promote the sparsest solution.

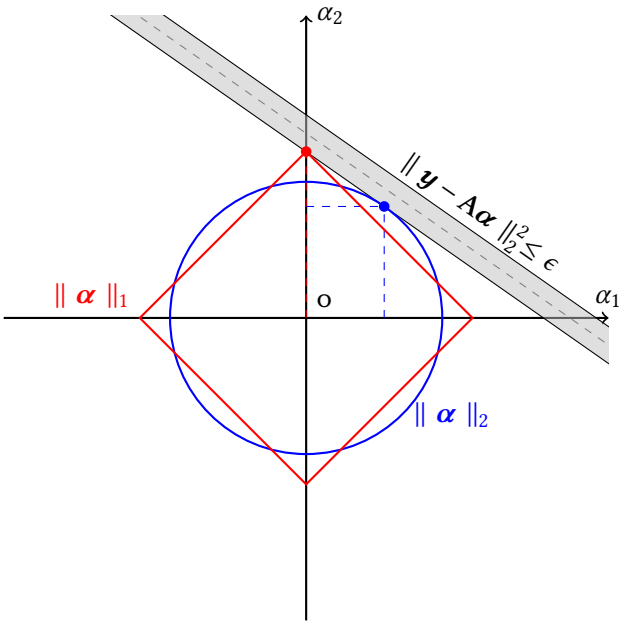

Figure 3.5: Geometrical representation of the inverse problem. The shaded gray area represents the feasible set allowed by the data fidelity constraint. The smallest ℓ_1 (red) and ℓ_2 (blue) balls intersecting the feasible set each define a different solution of the optimisation problem (red and blue dots). The ℓ_1 solution is sparse, with only one non zero coefficient, contrary to the ℓ_2 solution.

Although this problem leads to the sparsest solution and is what we would aim to solve, it is computationally hard to solve in practice (it is actually NP-hard Natarajan (1995)). As a result, we will rather consider a convex relaxation of this problem obtained by replacing the ℓ_0 norm by the ℓ_1 norm:

$$\arg\min_{\boldsymbol{\alpha}} \| \boldsymbol{\alpha} \|_1 \quad s.t. \quad \| \boldsymbol{y} - \mathbf{A}\boldsymbol{\Phi}\boldsymbol{\alpha} \|_2^2 \leq \epsilon \ . \tag{3.14}$$

This problem is known as *Basis Pursuit DeNoise* (BPDN). Both terms are now convex which will enable the development of fast algorithms to perform the minimisation.

Although not equivalent to the original ℓ_0 formulation, this relaxation of the problem still leads to sparse solutions. For sufficiently sparse signals, the relaxed ℓ_1 problem has even been shown to recover the exact solution of the ℓ_0 problem in the absence of noise (Donoho and Huo, 2001). Figure 3.5 illustrates the behaviour of the ℓ_1 norm on a simple two dimensional problem. Solving the BPDN amounts to finding the smallest ℓ_1 ball (red) intersecting the feasible set (grey shaded area). The point of intersection (red dot) is the solution of the problem. As can be seen, the solution is 1-sparse (only one active coefficient). In contrast, if one tries to address the same problem using an ℓ_2 norm, the ℓ_2 ball (blue) will intersect the feasible set at a point which is not sparse (blue dot), it has active coefficients in all dimensions. More generally, all ℓ_p norms with $p \in [0, 1]$ tend to promote sparse solutions but only the ℓ_1 norm is convex and thus leads to tractable optimisation problems.

The BPDN problem can also be equivalently recast either with explicit constraint on the ℓ_1 norm:

$$\arg\min_{\boldsymbol{\alpha}} \| \boldsymbol{y} - \mathbf{A}\boldsymbol{\Phi}\boldsymbol{\alpha} \|_2^2 \quad s.t. \quad \| \boldsymbol{\alpha} \|_1 \leq \tau \ , \tag{3.15}$$

which is commonly known as the Lasso problem, or in the form of an augmented Lagrangian:

$$\arg\min_{\boldsymbol{\alpha}} \frac{1}{2} \| \boldsymbol{y} - \mathbf{A}\Phi\boldsymbol{\alpha} \|_2^2 + \lambda \| \boldsymbol{\alpha} \|_1 \quad . \tag{3.16}$$

These three formulations are equivalent in the sense that for an appropriate choice of parameters ϵ, τ and λ they yield the same solution. However, the correspondence between these parameters is not trivial and not known a priori. Therefore, depending on the specific application and on algorithmic considerations one formulation or the other can be preferred.

3.2.2 Analysis sparsity prior

The formulation of the sparse recovery problem presented in the previous paragraph was under a so-called *synthesis* prior: the minimisation problem aims at finding a set of coefficients $\boldsymbol{\alpha}$ which are sparse and from which the solution can be synthesised as $\boldsymbol{x} = \Phi\boldsymbol{\alpha}$. However, one can also impose an *analysis* sparsity prior by requiring the analysis coefficients $\Phi^t\boldsymbol{x}$ to be sparse. In this case, the sparse recovery problem can be formulated as:

$$\arg\min_{\boldsymbol{x}} \frac{1}{2} \| \boldsymbol{y} - \mathbf{A}\boldsymbol{x} \|_2^2 + \lambda \| \Phi^t\boldsymbol{x} \|_1 \quad . \tag{3.17}$$

The two formulations are completely equivalent in the case where Φ is an orthogonal dictionary. Indeed, in this case, $\Phi^t = \Phi^{-1}$ and the analysis formulation only amounts to a change of variables. However, in the case of redundant dictionaries the two problems are no longer equivalent.

The synthesis approach has received considerable attention over the last decade regarding theoretical guarantees and efficient algorithms. In contrast, the analysis model has been much less investigated. Yet early works have advocated the use of the analysis prior Starck et al. (2004, 2005). The fundamental distinctions between these two priors was made explicit in Elad et al. (2007) and more recent results have shown superior recovery properties under this prior compared to the synthesis approach. For instance, Selesnick and Figueiredo (2009) compare both priors on denoising and deconvolution applications using undecimated wavelets and report superior results in both instances under the analysis model.

Several factors can explain this difference of performance between the two priors. First and foremost, under the synthesis prior the recovered signal is constrained to the space spanned by a few atoms of the dictionary whereas under the analysis model the solution can be any arbitrary signal of $\mathbb{R}^N$. Therefore, the atoms of an optimal synthesis dictionary should contain examples of the signal (for instance Diracs for the recovery of spikes). However, even when a dictionary is specifically built for an application (e.g. using dictionary learning Aharon et al. (2006)), several atoms are usually needed to reproduce examples of the signal. For instance synthesising positive structures using wavelets requires a significant number of coefficients to compensate the oscillatory nature of wavelet atoms. On the contrary, under the analysis prior, the recovered signal is not constrained to belong to a specific subspace spanned by a few columns of the dictionary.

Another aspect in favour of the analysis prior is the reduced number of unknowns involved in the optimisation problem. Indeed, when using redundant dictionaries, the dimension of the space of coefficients in the dictionary Φ can be much larger than the dimension of the signal to recover. In this case, the analysis formulation has much fewer unknowns thus leading to a simpler and more stable optimisation problem.

Recovery guarantees have recently been shown for the analysis prior, in particular in the context of Compressed Sensing (Candès et al., 2011) and in the presence of noise (Vaiter et al., 2013). A new framework, known as the *cosparse analysis model* (Nam et al., 2011), as also been proposed to explain the behaviour of the analysis prior. This framework adopts a different view point on the problem by defining the notion of cosparsity which focuses on the number of zero coefficients instead of the number of non-zero coefficients (see Nam et al. (2013) for a review).

3.2.3 Enhanced sparsity through ℓ_1-reweighting

The ℓ_1-minimisation problem provides a robust and efficient framework to enforce sparse regularisation on inverse problems. However, compared to the ℓ_0 regularisation, it has one major drawback, namely its dependence on the magnitude of the coefficients. As a result, in the presence of noise, the solution of the ℓ_1 minimisation can often be of biased amplitude. Indeed, without sufficient evidence from the data, the prior will prefer solutions with smaller amplitude. Consequently, it is often necessary to perform some additional debiasing step after convergence of the ℓ_1 recovery to correct for this effect.

An elegant approach to this problem has been proposed in Candès et al. (2008). The proposed method consists in iteratively solving an ℓ_1 minimisation, each time with a specific weighting of the ℓ_1 norm based on the previous estimate of the solution. The idea is to tend to a solution where non-zero coefficients are equitably penalized independently of their amplitude, thus getting closer to the solution of the ℓ_0 minimisation problem. The strength of this approach is that it only involves solving a series of convex problems which can efficiently be addressed by the proximal algorithms presented in the next section.

The weighted ℓ_1 minimisation problem is defined as:

$$\arg\min_{\boldsymbol{\alpha}} \frac{1}{2} \parallel \boldsymbol{y} - \mathbf{A}\boldsymbol{\Phi}\boldsymbol{\alpha} \parallel_2^2 + \lambda \parallel \mathbf{W}\boldsymbol{\alpha} \parallel_1 , \tag{3.18}$$

where $\mathbf{W}$ is a diagonal matrix of positive weights $w_1, w_2, \ldots, w_n$. The aim is to find a set of weights which tune the shape of the ℓ_1 border in order to promote a sparser solution and attenuate the dependence of the penalisation on the amplitude of the coefficients.

The question is now how to set adequate weights. In Candès et al. (2008), the authors propose the following iterative scheme to set the weights and converge to a solution:

1. Set the iteration count $\ell = 0$ and initialise the weights to $1 : \forall i, \quad w_i^{(0)} = 1$.
2. Solve the weighted ℓ_1 minimisation problem:

$$\boldsymbol{\alpha}^{(\ell)} = \arg\min_{\boldsymbol{\alpha}} \frac{1}{2} \parallel \boldsymbol{y} - \mathbf{A}\boldsymbol{\Phi}\boldsymbol{\alpha} \parallel_2^2 + \lambda \parallel \mathbf{W}^{(\ell)}\boldsymbol{\alpha} \parallel_1 .$$

3. Update the weights based on the solution $\boldsymbol{\alpha}^{(\ell)}$:

$$w_i^{(\ell+1)} = \frac{1}{|\alpha_i^{(\ell)}| + \epsilon} , \tag{3.19}$$

   where ϵ is a small positive parameter preventing infinite weights.

4. Terminate on convergence. Otherwise, increment ℓ and go to step 2.

As shown by Candés, this procedure has some analytical justification as it can be linked to solving the inverse problem under a log-sum penalty of the form $g(\boldsymbol{\alpha}) = \sum_i \log(|\alpha_i| + \epsilon)$ which is a better approximation of the ℓ_0 penalty but which is no longer convex. Unfortunately, this means that this procedure is sensitive to the choice of initialisation and more importantly to the choice of parameter ϵ. Candés recommends empirically setting ϵ to a small value of the order of the smallest expected non-zero coefficient of the signal.

In practice, we find that this reweighting procedure does improve the results of the ℓ_1 minimisation and we have implemented it to address the problems presented in Chapter 5 and Chapter 7. However, we advocate using a different update rule for the weights $\mathbf{W}$ based on the significance of the coefficients recovered at each iteration with respect to the level of noise:

$$w_i^{(\ell+1)} = \begin{cases} \frac{1}{|\alpha_i^{(\ell)}|/K\sigma_i} & \text{if } |\alpha_i^{(\ell)}| \geq K\sigma_i \\ 1 & \text{if } |\alpha_i^{(\ell)}| < K\sigma_i \end{cases} , \tag{3.20}$$

where σ_i is an estimate of the standard deviation of coefficient α_i due to noise and K is a given significance level. The idea behind this rule is to reduce the penalty on coefficients which have already been identified as significant at previous iterations. However, for coefficients below the threshold $K\sigma$ we keep the weights to 1 as we do not want to increase the penalty which would prevent new features from being detected at subsequent iterations.

Finally note that we have introduced the reweighted-ℓ_1 procedure in a synthesis setting. It can also be applied to improve the results of the analysis ℓ_1 problem as was demonstrated in Candès et al. (2008, 2011). The iterative scheme is identical except for the optimisation problem in step 2 which becomes:

$$\boldsymbol{x}^{(\ell)} = \arg\min_{\boldsymbol{x}} \frac{1}{2} \parallel \boldsymbol{y} - \mathbf{A}\boldsymbol{x} \parallel_2^2 + \lambda \parallel \mathbf{W}^{(\ell)}\boldsymbol{\Phi}^t\boldsymbol{x} \parallel_1 .$$

For the applications considered in the rest of this work, we have found the reweighted-ℓ_1 analysis to yield significantly better results compared to the synthesis approach.

3.3 PROXIMAL ALGORITHMS FOR SPARSE RECOVERY

The previous section has introduced the notion of sparse regularisation to address inverse problems which involves finding a solution by solving an optimisation problem. As was mentioned, solving the ℓ_0 minimisation problem, which is NP-hard, is not tractable, however its ℓ_1 relaxation is convex and can efficiently be solved by algorithms derived from proximal calculus. The aim of this section is to provide these algorithmic tools which make solving large scale optimisation problems possible in practice.

3.3.1 Elements of proximal calculus

We begin by introducing some notions of convex analysis. For a reference on this subject, we point the interested reader to Zalinescu (2002); Bauschke and Combettes (2011). Let us consider a function $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ and define the *domain of* f, noted dom f, as the subset of $\mathbb{R}^n$ where f does not reach $+\infty$:

$$\text{dom } f = \{\boldsymbol{x} \in \mathbb{R}^n \quad | \quad f(\mathbf{x}) < +\infty\} . \tag{3.21}$$

Property	$f(\boldsymbol{x})$	$\text{prox}_f(\boldsymbol{x})$
Translation	$f(\boldsymbol{x}-\boldsymbol{z}),\ \boldsymbol{z}\in\mathbb{R}^N$	$\boldsymbol{z}+\text{prox}_f(\boldsymbol{x}-\boldsymbol{z})$
Scaling	$f(\boldsymbol{x}/a),\ a\in\mathbb{R}^*$	$a\,\text{prox}_{f/a^2}(\boldsymbol{x}/a)$
Reflection	$f(-\boldsymbol{x})$	$-\text{prox}_f(-\boldsymbol{x})$
Conjugation	$f^*(\boldsymbol{x})$	$\boldsymbol{x}-\text{prox}_f(\boldsymbol{x})$
Separability	$f(\boldsymbol{x})=\sum_{i=1}^n f_i(x_i)$	$(\text{prox}_{f_1}(x_1),\dots,\text{prox}_{f_n}(x_n))$

Table 2: Useful properties of the proximity operator. See Combettes and Pesquet (2011) for a more extensive compilation. In this table, the conjugation property is relative to the *convex conjugation* defined for a function f as $f^*:\boldsymbol{x}\longmapsto \sup_{\boldsymbol{y}} <\boldsymbol{x},\boldsymbol{y}> -f(\boldsymbol{y})$.

A function f will be said *proper* if its domain $\text{dom}\, f$ is nonempty. We also define the *epigraph of* f, noted $\text{epi}\, f$:

$$\text{epi}\, f = \{(\boldsymbol{x},\lambda)\in\mathbb{R}^n\times\mathbb{R} \quad | \quad f(\boldsymbol{x})\leq\lambda\}\subset\mathbb{R}^{n+1}\ . \tag{3.22}$$

The epigraph is useful to characterise several important properties, in particular the *convexity* and *lower semicontinuity* of f:

$$\text{epi}\, f \text{ is closed} \Leftrightarrow f \text{ is lower semicontinuous} \tag{3.23}$$

$$\text{epi}\, f \text{ is convex} \Leftrightarrow f \text{ is convex} \tag{3.24}$$

We will note Γ_0 the class of proper lower semicontinuous convex functions of $\mathbb{R}^n$. Functions in Γ_0 are therefore characterised by a non empty closed convex epigraph.

The proximal operator, introduced by Moreau (1962) is at the center of the different algorithms introduced in the following sections. This operator, which can be seen as an extension of the convex projection operator is defined as:

Definition 3.1. *Let $F\in\Gamma_0$. For every $\boldsymbol{x}$ the function $\boldsymbol{y}\mapsto \frac{1}{2}\parallel\boldsymbol{x}-\boldsymbol{y}\parallel^2+F(\boldsymbol{y})$ achieves its infimum at a unique point defined as* $\text{prox}_F(\boldsymbol{x})$.

Therefore, given a proper lower semicontinuous convex function F, the proximity operator of F is uniquely defined by:

$$\text{prox}_F(\boldsymbol{x}) = \arg\min_{\boldsymbol{y}} \frac{1}{2}\parallel\boldsymbol{y}-\boldsymbol{x}\parallel_2^2+F(\boldsymbol{y})\ . \tag{3.25}$$

To manipulate this operator, some useful properties and calculus rules are listed in Table 2.

In order to illustrate this operator in practice in a simple case, consider $F=i_C$ the indicator function of a closed convex set C. Then the proximity operator reduces to the orthogonal projector onto C, noted proj_C:

$$\text{prox}_{i_C}(\boldsymbol{x}) = \arg\min_{\boldsymbol{y}\in C} \frac{1}{2}\parallel\boldsymbol{y}-\boldsymbol{x}\parallel_2^2 = \text{proj}_C(\boldsymbol{x})\ . \tag{3.26}$$

Similarly to the simple case of the indicator function, explicit expressions for the proximity operator exist for a number of different simple functions. Some useful examples are listed in Table 3.

In the context of sparse optimisation, we will be particularly interested in the proximity operator of the ℓ_1 norm which appears in the various problems stated in the previous section. Thankfully, the proximity operator of $F(\boldsymbol{x})=\lambda\parallel\boldsymbol{x}\parallel_1$ is explicit and corresponds to *Soft Thresholding*:

$$\text{prox}_{\lambda\parallel\cdot\parallel_1}(\boldsymbol{x}) = \text{ST}_\lambda(\boldsymbol{x})\ , \tag{3.27}$$

$f(\boldsymbol{x})$	$\text{prox}_f(\boldsymbol{x})$
$i_{\cdot \geq 0}(\boldsymbol{x})$	$[\boldsymbol{x}]_+$
$i_{\lvert\cdot_i\rvert \leq \lambda}(\boldsymbol{x})$	$x - \text{ST}_\lambda(\boldsymbol{x})$
$\lambda \parallel x \parallel_1$	$\text{ST}_\lambda(x)$

Table 3: Examples of explicit proximity operators. See Combettes and Pesquet (2011) for a more extensive compilation.

where the Soft Thresholding operator ST_λ is defined for $\boldsymbol{x} \in \mathbb{R}^N$ as:

$$\forall i \in [\![1, N]\!], \quad \text{ST}_\lambda(\boldsymbol{x})_i = \begin{cases} x_i - \lambda & \text{if } x_i \geq \lambda \\ 0 & \text{if } |x_i| \leq \lambda \\ x_i + \lambda & \text{if } x_i \leq \lambda \end{cases}. \tag{3.28}$$

The Soft Thresholding operator shrinks the amplitude of the coefficients by an amount λ and sets to 0 coefficients of amplitude smaller than λ. The response of this operator is compared to the Hard Thresholding in Figure 3.1.

Another situation of special interest for sparse optimisation is the precomposition of the ℓ_1 norm with a bounded linear operator. Indeed, under the analysis prior (see Section 3.2.2), the regularisation term is not directly the ℓ_1 norm of the variable but $F(\boldsymbol{x}) = \lambda \parallel \boldsymbol{\Phi}^t \boldsymbol{x} \parallel_1$. Unfortunately, in the general case, this function no longer admits an explicit proximity operator which makes solving the ℓ_1 analysis problem significantly more challenging than its synthesis counterpart. In the absence of an explicit formula, the proximity operator can still be evaluated by going back to its definition and directly solving the optimisation problem involved. It can be shown (e.g. Appendix of Rapin et al., 2014) that for an arbitrary dictionary $\boldsymbol{\Phi}$, the proximity operator of F can be evaluated as the solution of:

$$\text{prox}_{\lambda \parallel \boldsymbol{\Phi}^t \cdot \parallel_1}(\boldsymbol{x}) = \boldsymbol{x} - \boldsymbol{\Phi} \left(\arg\min_{\boldsymbol{u}} \frac{1}{2} \parallel \boldsymbol{x} - \boldsymbol{\Phi u} \parallel_2^2 + i_{\lvert\cdot_i\rvert \leq \lambda}(\boldsymbol{u}) \right) \tag{3.29}$$

Note that in the optimisation problem involved, the second term is the indicator function of the ℓ_1 ball whose proximity operator is explicit (see Table 3). Thus solving this problem can be efficiently addressed with the proximal algorithms introduced in the next section. An explicit algorithm to evaluate this proximity operator using Forward-Backward splitting (see Section 3.3.2.1) is provided in Algorithm 3.1.

Algorithm 3.1 Evaluation of $\text{prox}_{\lambda \parallel \boldsymbol{\Phi}^t \cdot \parallel_1}$

Require: Gradient step $0 < \mu < \frac{1}{\|\boldsymbol{\Phi}\|^2}$.
1: **for** $n = 0$ to $N_{\text{max}} - 1$ **do**
2: $\quad \tilde{\boldsymbol{u}}^{(n+1)} = \boldsymbol{u}^{(n)} + \mu \boldsymbol{\Phi}^t \left(\boldsymbol{x} - \boldsymbol{\Phi u}^{(n)} \right)$
3: $\quad \boldsymbol{u}^{(n+1)} = \tilde{\boldsymbol{u}}^{(n+1)} - \text{ST}_\lambda(\tilde{\boldsymbol{u}}^{(n+1)})$
4: **end for**
5: **return** $\text{prox}_{\lambda \parallel \boldsymbol{\Phi}^t \cdot \parallel_1}(\boldsymbol{x}) = \boldsymbol{x} - \boldsymbol{\Phi u}^{(N_{\text{max}})}$

3.3.2 Proximal algorithms

Based on proximal calculus a number of minimisation algorithms have been developed in recent years for the purpose of solving sparse optimisation problems. In this section, we gather a few proximal algorithms which have been useful over the course of this thesis.

3.3.2.1 Forward-Backward

The simplest and most well-known algorithm is the *Forward-Backward* (FB) algorithm (Combettes and Wajs, 2005) which aims at solving problems of the form:

$$\arg\min_{\boldsymbol{x}} F(\boldsymbol{x}) + G(\boldsymbol{x}) \,, \tag{3.30}$$

where F and G are two proper lower semicontinuous convex functions but in the case where F is differentiable with β the Lipschitz constant of ∇F, the gradient of F. Under these conditions, the following algorithm converges to the minimum of Equation (3.30):

Algorithm 3.2 Forward-Backward Splitting algorithm

Require: Gradient step $0 < \mu < \frac{1}{\beta}$.
1: **for** $n = 0$ to $N_{\max} - 1$ **do**
2: $\quad \tilde{\boldsymbol{x}}^{(n+1)} = \boldsymbol{x}^{(n)} + \mu \nabla F\left(\boldsymbol{x}^{(n)}\right)$
3: $\quad \boldsymbol{x}^{(n+1)} = \text{prox}_{\mu G}\left(\tilde{\boldsymbol{x}}^{(n+1)}\right)$
4: **end for**

This algorithm applies in particular to the ℓ_1 recovery problem stated in Equation (3.16) with $F(\boldsymbol{\alpha}) = \frac{1}{2} \parallel \boldsymbol{y} - \mathbf{A}\boldsymbol{\Phi}\boldsymbol{\alpha} \parallel_2^2$ and $G(\boldsymbol{\alpha}) = \lambda \parallel \boldsymbol{\alpha} \parallel_1$. Indeed, the quadratic data fidelity term is differentiable, with $\nabla F(\alpha) = \boldsymbol{\Phi}^t \mathbf{A}^* (\boldsymbol{y} - \mathbf{A}\boldsymbol{\Phi}\boldsymbol{\alpha})$ and $\beta = \parallel \mathbf{A}\boldsymbol{\Phi} \parallel^2$ where $\parallel . \parallel$ is the operator norm. As was presented in the previous section, the proximal operator of the G is in this case a simple soft thresholding $\text{prox}_G(\boldsymbol{x}) = \text{ST}_\lambda(\boldsymbol{x})$. Therefore, for the ℓ_1 recovery problem, forward-backward splitting yields a simple *Iterative Soft-Thresholding Algorithm* (ISTA):

$$\boldsymbol{\alpha}^{(n+1)} = ST_{\mu\lambda}\left(\boldsymbol{\alpha}^{(n)} + \mu \boldsymbol{\Phi}^t \mathbf{A}^* \left(\boldsymbol{y} - \mathbf{A}\boldsymbol{\Phi}\boldsymbol{\alpha}\right)\right) \tag{3.31}$$

Although this algorithm is proven to converge to the solution, the convergence rate of this simple iteration is slow. A simple variant of ISTA was proposed in Beck and Teboulle (2009) which converges in $O(1/n^2)$ compared to $O(1/n)$ for ISTA. This variant, called *Fast Iterative Shrinkage-Thresholding Algorithm* (FISTA), is based on a specific update rule which combines the current and previous estimates at each iteration. This algorithm is presented in Algorithm 3.3.

3.3.2.2 Generalised Forward-Backward

The Forward-Backward algorithm is very useful in a number of situations but it is limited to the minimization of only 2 terms. For some applications, additional constraints are desirable/required but cannot be included in a single regularisation term G without losing the explicit formulation of the proximity operator. This is for instance the case if one wants to include an additional positivity constraint on the solution of Equation (3.16). An extension of the FB algorithm was proposed

Algorithm 3.3 Fast Iterative Shrinkage-Thresholding (FISTA)

Require: Gradient step $0 < \mu < \frac{1}{\beta}$.

$t_0 = 1$
for $n = 0$ to $N_{\max} - 1$ **do**
  $\tilde{z}^{(n)} = x^{(n)} + \mu \nabla F\left(x^{(n)}\right)$
  $z^{(n)} = \text{prox}_{\mu G}\left(\tilde{z}^{(n)}\right)$
  $t_{n+1} = \frac{1+\sqrt{1+4t_n^2}}{2}$
  $x^{(n+1)} = z^{(n)} + \left(\frac{t_n - 1}{t_{n+1}}\right)\left(z^{(n)} - z^{(n-1)}\right)$
end for

in Raguet et al. (2013) to handle an arbitrary number of additional proximable constraints. This Generalised Forward-Backward (GFB) algorithm aims at solving problems of the form:

$$\arg\min_{x} F(x) + \sum_{i=1}^{n} G_i(x) \,, \tag{3.32}$$

where F and G_i are proper lower semicontinuous convex functions and only F needs to be differentiable. Compared to the simple FB algorithm, the regularisation G can be split into as many terms G_i as necessary so that each G_i has a simple form with an explicit proximal operator. The algorithm solving this problem is given in Algorithm 3.4.

Algorithm 3.4 Generalized Forward-Backward Algorithm

Require: Gradient step $0 < \mu < \frac{2}{\beta}$. Weights ω_i such that $\sum_i \omega_i = 1$.

for $k = 0$ to $k_{\max} - 1$ **do**
  for $i \in [\![1, n]\!]$ **do**
    $z_i^{(k+1)} = z_i^{(k)} + \text{prox}_{\frac{\mu}{\omega_i} G_i}\left(2x^{(k)} - z_i^{(k)} - \mu \nabla F(x^{(k)})\right) - x^{(k)}$
  end for
  $x^{(k+1)} = \sum_{i=1}^{n} \omega_i z_i^{(k+1)}$
end for

3.3.2.3 *Chambolle-Pock primal-dual algorithm*

The algorithms presented so far rely on the evaluation of at least one proximity operator at each iteration. Although in simple cases, a closed form expression of the operator can be used (for instance soft thresholding for ℓ_1 regularisation), in more complex cases the proximity operator is not explicit and these algorithms lose some of their appeal as a nested optimisation problem needs to be solved at each iteration. This is for instance the case of the ℓ_1 analysis problem in which the regularisation term $G(x) = \| \Phi^t x \|_1$ does not have an explicit proximity operator for an arbitrary linear operator Φ.

Chambolle and Pock proposed in Chambolle and Pock (2011) an elegant way to address problems of the form:

$$\arg\min_{x} F(x) + G(\mathbf{W}x) \,, \tag{3.33}$$

where F and G are two functions in Γ_0 and $\mathbf{W}$ is a linear operator. Note that the first term F does not need to be differentiable, and even if it is, this property is not used by the algorithm. These authors recast this optimisation problem in the following primal-dual form:

$$\arg\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} \langle \mathbf{W}\boldsymbol{x}, \boldsymbol{y} \rangle + F(\boldsymbol{x}) - G^*(\boldsymbol{y}) \,, \tag{3.34}$$

where G^* is the *convex conjugate* of the original function G, defined by $G^*(\boldsymbol{y}) = \max_{\boldsymbol{y}'}\{\langle \boldsymbol{y}, \boldsymbol{y}' \rangle - G(\boldsymbol{y}')\}$. Note that in the above primal-dual formulation, the term $G^*(\boldsymbol{y})$ no longer involves the precomposition by the linear operator $\mathbf{W}$ and as long as the G possesses an explicit proximity operator, so does G^*, thanks to the conjugation formula in Table 2. The algorithm proposed by Chambolle and Pock is detailed in Algorithm 3.5.

Algorithm 3.5 Chambolle-Pock primal-dual algorithm

Require: $L = \| \mathbf{W} \|$, $\tau > 0$, $\sigma > 0$, $L^2\sigma\tau < 1$.

for $k = 0$ to $k_{\max} - 1$ **do**
  $\boldsymbol{y}^{(k+1)} = \text{prox}_{\sigma G^*}(\boldsymbol{y}^{(k)} + \sigma \mathbf{W}\boldsymbol{x}^{(k)})$
  $\tilde{\boldsymbol{x}}^{(k+1)} = \text{prox}_{\tau F}(\boldsymbol{x}^{(k)} - \tau \mathbf{W}^t \boldsymbol{y}^{(k+1)})$
  $\boldsymbol{x}^{(k+1)} = 2\tilde{\boldsymbol{x}}^{(k+1)} - \tilde{\boldsymbol{x}}^{(k)}$
end for

3.3.2.4 *Vu primal-dual algorithm*

The drawback of the primal-dual Chambolle-Pock algorithm is that it does not exploit the differentiability of F when this function F is actually differentiable. A wider framework was proposed in Vu (2013) which can accommodate at the same time the precomposition by a linear operator and a differentiable term. In particular, it can be used to solve problems of the form:

$$\arg\min_{\boldsymbol{x}} F(\boldsymbol{x}) + G(\mathbf{W}\boldsymbol{x}) + H(\boldsymbol{x}) \,, \tag{3.35}$$

where F is convex and differentiable with a Lipschitzian gradient of constant β, $(H, G) \in \Gamma_0^2$ and $\mathbf{W}$ is a non-zero linear operator. The algorithm proposed by Vu to solve a problem of this form is given in Algorithm 3.6. As an example of application of this algorithm, consider the following ℓ_1

Algorithm 3.6 Vu primal-dual algorithm

Require: $\tau > 0$, $\sigma > 0$ with $1 - \tau\sigma \| \mathbf{W} \|^2 > \tau\beta/2$.

for $k = 0$ to $k_{\max} - 1$ **do**
  $\boldsymbol{x}^{(k+1)} = \text{prox}_{\tau H}\left(\boldsymbol{x}^{(k)} + \tau(\nabla F(\boldsymbol{x}^{(k)}) - \mathbf{W}^t \boldsymbol{u}^{(k)}\right)$
  $\boldsymbol{u}^{(k+1)} = \text{prox}_{G^*}(\boldsymbol{u}^{(k)} + \sigma \mathbf{W}(2\boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)}))$
end for

analysis problem, with positivity constraint:

$$\arg\min_{\boldsymbol{x}} \frac{1}{2} \| \boldsymbol{y} - \mathbf{A}\boldsymbol{x} \|_2^2 + \lambda \| \boldsymbol{\Phi}^t \boldsymbol{x} \|_1 + i_{\cdot \geq 0}(\boldsymbol{x}) \,. \tag{3.36}$$

This problem is not efficiently addressed by any of the previous algorithms as it contains three different terms and one of these ($\lambda \parallel \Phi^t \boldsymbol{x} \parallel_1$) does not have an explicit proximal operator. Yet, it can be easily solved using Algorithm 3.6, where steps 2 and 3 are specialised to:

$$\boldsymbol{x}^{(k+1)} = \left[\boldsymbol{x}^{(k)} + \tau(\nabla F(\boldsymbol{x}^{(k)}) - \Phi \boldsymbol{u}^{(k)}\right]_+ \tag{3.37}$$

$$\boldsymbol{u}^{(k+1)} = (\mathrm{Id} - \mathrm{ST}_\lambda)(\boldsymbol{u}^{(k)} + \sigma \Phi^t (2\boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)})) \tag{3.38}$$

3.4 CONCLUSION

In this chapter, we introduced the notion of sparsity and how it can be applied to the regularisation of linear inverse problems. As we have seen, under a sparsity prior on the solution, these problems can be recast as convex optimisation problems and we introduced the algorithmic tools necessary to efficiently solve such problems.

However, the success of sparse recovery not only relies on efficient algorithms but more crucially it depends on the quality of the dictionaries used to sparsely represent the signal of interest. This question has barely been discussed in this chapter but will be developed in more details in the next chapter where we introduce wavelets, a celebrated family of sparse representations, extensively used throughout this thesis. Combined with the algorithms introduced in this chapter, wavelets will be at the core of all the applications presented in Chapter 5, Chapter 7 and Chapter 8.

4 WAVELETS

CONTENTS

Wavelets are a wide class of functions, localised in time and frequency, which can be used to efficiently represent non stationary signals. The analysis of a signal through its wavelet coefficients makes it possible to identify features of different scales and at different positions, contrary to, for instance, a Fourier analysis which is limited to scale.

In the context of the sparse regularisation of inverse problems, our interest in wavelets stems from their ability to provide sparse representations for most natural signals. In fact, wavelets lead to sparse representations for the very general class of piecewise smooth signals. However, in sparse recovery applications, depending on the specific type of wavelet, the quality of the end result can vary greatly and therefore much consideration should go into the choice of an appropriate wavelet for a given application.

In this chapter, we begin by introducing the fundamentals of the continuous and discrete wavelet transforms, the latter being the one used in practice for digital signal processing. We review a few specific wavelet constructions in 1D and 2D used in the rest of this work and outline the principles of the fast transforms which make wavelets a particularly attractive option for building sparse representations. This chapter also serves as an introduction to the Spherical 3D wavelets presented in Chapter 9 which will extend to the 3D ball some of the concepts presented here.

4.1 CONTINUOUS WAVELETS

4.1.1 Definition

The *Continuous Wavelet Transform* (CWT) was defined by Morlet and Grossmann (Grossmann and Morlet, 1984) for functions in $\mathbf{L}^2(\mathbb{R})$, the space of square integrable functions. To ensure an invertible wavelet transform, a wavelet is defined as a real function ψ which verifies the following *admissibility condition*:

$$C_\psi = \int_0^\infty |\hat{\psi}(\nu)|^2 \nu^{-1} \mathrm{d}\nu < +\infty \,, \tag{4.1}$$

where $\hat{\psi}$ is the Fourier transform of ψ. In order to verify this condition, wavelets must at least have vanishing mean $\int \psi(t)\,\mathrm{d}t = 0$. The function ψ is called a *mother wavelet*, from which a family of *daughter wavelets* can be defined through scaling and translation. Given a scaling parameter $a \in \mathbb{R}^+$ and a translation parameter $b \in \mathbb{R}$, the daughter wavelet $\psi_{a,b}$ is defined as

$$\forall x \in \mathbb{R}, \quad \psi_{a,b}(x) = \frac{1}{\sqrt{a}} \psi\left(\frac{x-b}{a}\right) \,. \tag{4.2}$$

The CWT is defined as the projection of a function $f \in L_2(\mathbb{R})$ onto the family of daughter wavelets. The coefficients of this projection are called wavelet coefficients and are obtained by taking the inner product of f and $(\psi_{a,b})_{a>0,b\in\mathbb{R}}$:

$$\begin{aligned} \forall a \in \mathbb{R}^+, b \in \mathbb{R}, \quad W_f(a,b) &= < f, \psi_{a,b} > & (4.3)\\ &= \int_{\mathbb{R}} f(x) \psi^*_{a,b}(x) \mathrm{d}x & (4.4)\\ &= \frac{1}{\sqrt{a}} \int_{\mathbb{R}} f(x) \psi^*\left(\frac{x-b}{a}\right) \mathrm{d}x & (4.5) \end{aligned}$$

Interestingly, these wavelet coefficients can also be expressed in terms of a convolution product by defining the function $\bar{\psi}_a(x) = \frac{1}{\sqrt{a}} \psi^*(\frac{-x}{a})$, in which case:

$$\forall a \in \mathbb{R}^+, b \in \mathbb{R}, \quad W_f(a,b) = f * \bar{\psi}_a(b) \tag{4.6}$$

This expression hints at an implementation of the wavelet transform using linear filters which will be an essential aspect of the discrete wavelets introduced in the next section.

This wavelet decomposition is invertible (Grossmann and Morlet, 1984) and the function f can be recovered from the wavelet coefficients $(W_f(a,b))_{a>0,b\in\mathbb{R}}$ with the following formula:

$$f(x) = \frac{1}{C_\psi} \int_{\mathbb{R}^+} \int_{\mathbb{R}} W_f(a,b) \psi_{a,b}(x) \frac{\mathrm{d}a}{a^2} \mathrm{d}b \tag{4.7}$$

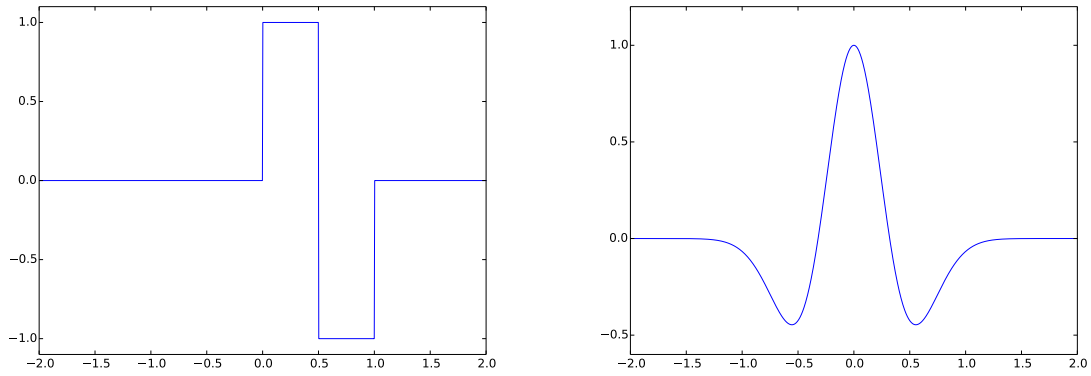

(a) Haar wavelet $\psi(t)$

(b) Mexican hat wavelet $\psi(t)$

Figure 4.1: Mother wavelet function for the Haar and Mexican hat wavelets.

4.1.2 Examples

Haar wavelet

The simplest and first example of wavelet function is due to Haar (Haar, 1910) who defined the following piecewise constant function:

$$\psi(t) = \begin{cases} 1 & \text{if } t \in [0,\ \frac{1}{2}[\\ -1 & \text{if } t \in [\frac{1}{2},\ 1[\\ 0 & \text{otherwise} \end{cases} \tag{4.8}$$

One of the particular properties of this wavelet, already recognised by Haar, is that the discrete set $\{\psi_{j,n}\}_{(j,n)\in\mathbb{Z}^2}$ of shifted and scaled versions of the mother wavelet defined by the following equation is an orthonormal basis of $\mathbf{L}^2(\mathbb{R})$:

$$\forall (j,n) \in \mathbb{Z}^2, \quad \psi_{j,n}(t) = \frac{1}{\sqrt{2^j}}\psi\left(\frac{t - 2^j n}{2^j}\right) \tag{4.9}$$

Therefore, this wavelet is also the first example of a discrete wavelet decomposition where the scaling and translation parameters can be discretised. In practice this wavelet is extremely simple to compute, as it only involves evaluating finite differences. However, its main drawback is its lack of regularity which can cause severe artefacts in a number of applications.

Mexican hat wavelet

Another very common example of continuous wavelet is the Mexican hat wavelet (or Ricker wavelet) which is built from the second derivative of a Gaussian. The expression of the wavelet is:

$$\psi(t) = \frac{2}{\pi^{1/4}\sqrt{3\sigma}}\left(\frac{t^2}{\sigma^2} - 1\right)\exp\left(\frac{-t^2}{2\sigma^2}\right) \tag{4.10}$$

Contrary to the Haar wavelet, the Mexican hat is extremely regular. Its 2D generalisation is well known in the field of computer vision under the name of Laplacian of Gaussian and is often used as a blob detector.

An example of continuous wavelet decomposition using the Mexican hat wavelet is presented in Figure 4.2. The top panel shows a non stationary 1D signal which is analysed to produce the *scalogram* showed in the bottom panel. The x-axis of the scalogram represents the time while the y-axis represents the scale. As can be seen, the signal contains two contributions which are clearly separated on the scalogram in time and frequency.

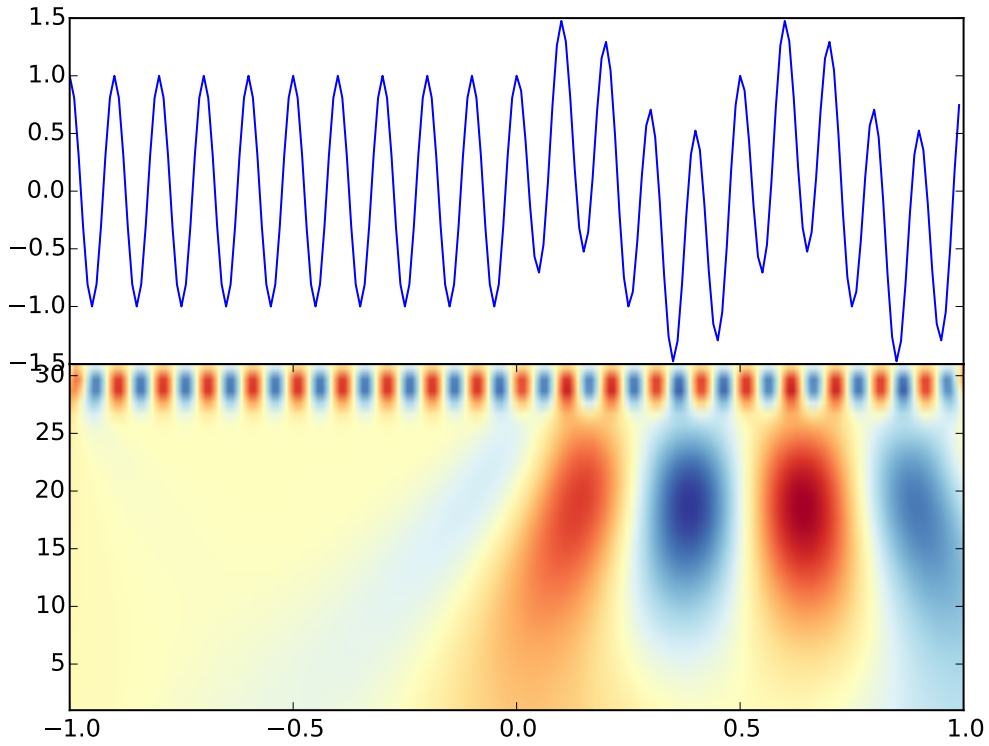

Figure 4.2: CWT of the signal in the top panel using the Mexican hat wavelet. The input signal contains two contributions, a high frequency cosine that spans the entire sequence and a low frequency sine that starts at $t = 0$. These two contributions are clearly identifiable in time and frequency on the scalogram.

4.2 ORTHOGONAL AND BI-ORTHOGONAL WAVELETS

The wavelet transform introduced above is set in a continuous framework. However, for signal processing purposes, a discrete wavelet transform is required. There are several ways to discretise this transform but to be useful in practice such a transform must be fast and admit an exact reconstruction formula. This has lead to the development of the Multiresolution Analysis framework (Mallat, 1989; Meyer, 1992) which can be used to build orthogonal wavelet basis.

4.2.1 MultiResolution Analysis

Mallat introduced in Mallat (1989) the concept of *MutliResolution Analysis* (MRA) in order to provide a framework for building orthogonal wavelet bases. The strength of this approach is that it can be used to define a discrete wavelet transform, with exact reconstruction from a discrete set of coefficients, which can be efficiently implemented using simple linear filters.

The idea of multiresolution analysis is to build a sequence of approximations f_j of a function $f \in \mathbf{L}^2(\mathbb{R})$ by smoothing f with a kernel of increasing width proportional to 2^j. Each of these approximations f_j will belong to a subspace $V_j \subset \mathbf{L}^2(\mathbb{R})$ which will regroup all possible approximations at scale 2^j. Then, f_j can be seen as the orthogonal projection of f on V_j. For each approximation subspace V_j one can build an orthogonal subspace W_j which will contain all the details lost between two consecutive approximations, such that $f_{j-1} = f_j + w_j$ with $w_j \in W_j$:

$$V_{j-1} = V_j \oplus W_j \ . \tag{4.11}$$

By recursively using this decomposition of approximations subspaces V_j for all $j \in \mathbb{Z}$ it will directly follow that the entire space $\mathbf{L}^2(\mathbb{R})$ can be decomposed into a direct sum of detail subspaces W_j:

$$\mathbf{L}^2(\mathbb{R}) = \bigoplus_{j=-\infty}^{\infty} W_j \ . \tag{4.12}$$

Mallat shows that one can build a wavelet ψ such that for each $j \in \mathbb{Z}$ the family of functions

$$\left\{ \psi_{j,n} = \frac{1}{\sqrt{2^j}} \psi \left(\frac{t - 2^j n}{2^j} \right) \right\}_{n \in \mathbb{Z}} \ , \tag{4.13}$$

forms an orthonormal basis of W_j. Finally, the whole family of wavelets $\{\psi_{j,n}\}_{(j,n)\in\mathbb{Z}^2}$ forms an orthonormal basis of $\mathbf{L}^2(\mathbb{R})$.

To formally develop these ideas, Mallat and Meyer (Mallat, 1989; Meyer, 1992) introduce the following definition of multiresolutions.

Definition 4.1 (Multiresolution approximation)**.** *Let f be a function in $L_2(\mathbb{R})$. A multiresolution approximation is a sequence $\{V_i\}_{j\in\mathbb{Z}}$ of closed embedded subspaces $\{0\} \subset \ldots \subset V_1 \subset V_0 \subset V_{-1} \subset \ldots \subset L_2(\mathbb{R})$ which verify the following properties:*

- *Translational invariance: $\forall j,k \in \mathbb{Z}, f(t) \in V_j \Leftrightarrow f(t - 2^j k) \in V_{j+1}$*
- *Causality: $f(t) \in V_j \Leftrightarrow f(t/2) \in V_{j+1}$*
- *Limit conditions: $\bigcap_{j\in\mathbb{Z}} V_j = \{0\}$ and $\overline{\bigcup_{j\in\mathbb{Z}} V_j} = L_2(\mathbb{R})$*
- *There must exist a function θ such that $\{\theta(t-n)\}_{n\mathbb{Z}}$ is a Riesz basis of V_0.*

The purpose of the Riesz basis $\{\theta(t-n)\}_{n\mathbb{Z}}$ of V_0 is to provide a sampling theorem: any function of V_0 can be represented using a discrete set of coefficients. Thanks to the properties of the multiresolution approximation, Mallat shows that the function θ can be used to construct a function ϕ, called *scaling function*, such that the family $\{\phi_{j,n}\}_{n\in\mathbb{Z}}$, defined as follows, forms an orthonormal basis of V_j for all $j \in \mathbb{Z}$:

$$\phi_{j,n}(t) = \frac{1}{\sqrt{2^j}} \phi \left(\frac{t-n}{2^j} \right) \ . \tag{4.14}$$

Therefore, it becomes possible to compute f_j the approximation of a function f at scale 2^j by taking its inner product with the basis functions $\{\phi_{j,n}\}_{n\in\mathbb{Z}}$ of the approximation space V_j:

$$f_j = P_{V_j} f = \sum_{n\in\mathbb{Z}} < f, \phi_{j,n} > \phi_{j,n} = \sum_{n\in\mathbb{Z}} a_j[n]\phi_{j,n} \ , \tag{4.15}$$

where P_{V_j} is the orthogonal projector on V_j and $a_j[n] =< f, \phi_{j,n} >$ are the discrete *approximation coefficients* of f at the scale 2^j. This extends the sampling theorem for V_0 provided by the Riesz basis to all spaces V_j: for all $j \in \mathbb{Z}$ the approximation f_j is uniquely defined by the set of approximation coefficients $\{a_j[n]\}_{n\in\mathbb{Z}}$.

A key aspect of multiresolution analysis is that approximation coefficients at one resolution can be computed from the coefficients of the previous resolution by convolution with a discrete filter entirely defined by the scaling function ϕ.
Indeed, since ϕ belongs to V_0, the scaled function $\frac{1}{\sqrt{2}}\phi(\frac{t}{2})$ is in V_1 by causality and therefore in V_0 since $V_1 \subset V_0$, which means that it can be decomposed on the basis $\{\phi_{0,n}\}_{n\in\mathbb{Z}}$ of V_0 leading to the following *scaling equation*:

$$\frac{1}{\sqrt{2}}\phi(\frac{t}{2}) = \sum_{n\in\mathbb{Z}} h[n]\phi(t-n) \ , \tag{4.16}$$

where $h[n] =< \frac{1}{\sqrt{2}}\phi(\frac{t}{2}), \phi(t-n) >$. More generally, any basis function $\phi_{j+1,n}$ at scale 2^{j+1} can be decomposed on the basis $\{\phi_{j,p}\}_{p\in\mathbb{Z}}$ of V_j using the same filter $h[n]$:

$$\phi_{j+1,p} = \sum_{n\in\mathbb{Z}} h[n-2p]\phi_{j,n} \ . \tag{4.17}$$

Taking the inner product of f with both terms of this equation leads to the following expression:

$$a_{j+1}[p] = \sum_{n\in\mathbb{Z}} h[n-2p]a_j[n] \ . \tag{4.18}$$

This fundamental relation between approximation coefficients of two consecutive resolutions means that these coefficients can *recursively* be computed using a simple discrete filter, without the need of actually computing the (continuous) inner products between f and the basis functions $\phi_{j,n}$. This relation is at the heart of the fast wavelet transform algorithms detailed in the next section.

The filter $h[n]$ derived from the scaling equation in Equation (4.16) exhibits interesting properties in relation to the scaling function ϕ in Fourier space. In particular Mallat (1989); Meyer (1992) show the following theorem:

Theorem 4.1 (Conjugate Mirror Filter). *If $\phi \in \mathbf{L}^2(\mathbb{R})$ is a scaling function, then the Fourier series of $h[n]$ verifies $\hat{h}(0) = \sqrt{2}$ and is a Conjugate Mirror Filter (CMF) i.e. satisfies*

$$\forall \omega \in \mathbb{R}, \quad |\hat{h}(\omega)|^2 + |\hat{h}(\omega+\pi)|^2 = 2. \tag{4.19}$$

Reciprocally, if $\hat{h}$ is a CMF (i.e. verifies Equation (4.19)), 2π periodic, continuously differentiable around $\omega = 0$ and satisfies $\hat{h}(0) = \sqrt{2}$ and $\inf_{\omega\in[-\pi/2,\pi/2]} |\hat{h}(\omega)| > 0$ then

$$\hat{\phi}(\omega) = \prod_{p=1}^{+\infty} \frac{\hat{h}(2^{-p}\omega)}{\sqrt{2}} \tag{4.20}$$

is the Fourier transform of a scaling function $\phi \in \mathbf{L}^2(\mathbb{R})$.

So far, the multiresolution analysis has allowed us to build orthonormal bases of the approximation subspaces V_j and to recursively compute approximation coefficients. However, it is also possible to capture the details lost between two consecutive approximations as a set of discrete wavelet coefficients. Indeed, since for a given scale 2^j, V_j is included in V_{j-1}, one can define W_j the orthogonal supplement of V_j in V_{j-1} so that:

$$V_{j-1} = V_j \oplus W_j \ . \tag{4.21}$$

Finally Mallat (1989); Meyer (1992) show that one can build an orthonormal basis $\{\psi_{j,n}\}_{n\in\mathbb{Z}}$ of W_j from a wavelet function ψ associated to the scaling function ϕ. This wavelet function is defined Fourier space as:

$$\hat{\psi}(\omega) = \frac{1}{\sqrt{2}}\hat{g}(\frac{\omega}{2})\hat{\phi}(\frac{\omega}{2}) \ , \tag{4.22}$$

where we introduce the discrete filter $g[n]$, defined by its Fourier series:

$$\hat{g}(\omega) = e^{-i\omega}\hat{h}^*(\omega + \pi) \ . \tag{4.23}$$

If one defines $\psi_{j,n} = \frac{1}{\sqrt{2}}\psi\left(\frac{t-2^j n}{2^j}\right)$, then for any scale 2^j, $\{\psi_{j,n}\}_{n\in\mathbb{Z}}$ is an orthonormal basis of W_j and for all scales $\{\psi_{j,n}\}_{(n,j)\in\mathbb{Z}^2}$ is an orthonormal basis of $\mathbf{L}^2(\mathbb{R})$.

This result has many important consequences. First and foremost it implies that $\mathbf{L}^2(\mathbb{R})$ can be decomposed as a direct sum of orthogonal subspaces W_j:

$$\mathbf{L}^2(\mathbb{R}) = \bigoplus_{j=-\infty}^{\infty} W_j \ . \tag{4.24}$$

which means that any function of $\mathbf{L}^2(\mathbb{R})$ can uniquely be decomposed into a set of discrete *wavelet coefficients* by an orthogonal projection:

$$f = \sum_{j\in\mathbb{Z}}\sum_{n\in\mathbb{Z}} < f, \psi_{j,n} > \psi_{j,n} = \sum_{j\in\mathbb{Z}}\sum_{n\in\mathbb{Z}} d_j[n]\psi_{j,n} \tag{4.25}$$

where $d_j[n] =< f, \psi_{j,n} >$ are the wavelet coefficients of f at scale 2^j and position $2^j n$.

Furthermore, similarly to the scaling equation in Equation (4.16), the wavelet function $\psi_{j+1,n}$ is a function of V_j, since $V_j = V_{j+1} \oplus W_{j+1}$, and can therefore be decomposed on the basis $\{\phi_{j,n}\}_{n\in\mathbb{Z}}$ of V_j. This decomposition can be written in terms of the filter g and leads to:

$$\psi_{j+1,p} = \sum_{n\in bZ} g[n-2p]\phi_{j,n} \ . \tag{4.26}$$

Taking the inner product of f with both terms of this equation yields:

$$d_{j+1}[p] = \sum_{n\in\mathbb{Z}} g[n-2p]a_j[n] \tag{4.27}$$

Again, this relation is fundamental to derive fast wavelet transform algorithms as it shows that wavelet coefficients at scale 2^{j+1} can simply be computed from the approximation coefficients at scale 2^j using a simple discrete filter and eliminates the need to actually evaluate the inner product of f with the wavelet $\psi_{j,n}$.

4.2.2 Vanishing moments and size of the support

As was mentioned in the introduction of this chapter, our particular interest in wavelets stems from their ability to provide sparse representations for a wide variety of natural signals. In fact, as we will discuss now, wavelets bases are sparse representations for piecewise smooth functions and can be optimised to maximise this sparsity.

A wavelet ψ is said to possess p vanishing moments if

$$\forall k \in [\![0, p[\![, \quad \int_{-\infty}^{+\infty} x^k \psi(x)\,\mathrm{d}x = 0\ , \tag{4.28}$$

which means that ψ is orthogonal to any polynomial of degree less than or equal to $p-1$. Now consider a function f which is assumed to be locally of class C^k then it can be locally approximated by a Taylor expansion of order k. If $k < p$, the wavelet will be orthogonal to the Taylor polynomial and therefore the coefficients $| < f, \psi_{j,n} > |$ will be small on fine scales at the vicinity of the expansion.

Thus, in order to provide a sparse representation for a wider class of signals, a wavelet with a large number of vanishing moments would be preferable from this point of view. However, this conclusion is only valid where the function f is locally smooth. In practice, signals are often only piecewise smooth and exhibit a number of singularities (for instance sharp edges in an image) which will lead to significant coefficients even on small scales as they cannot be approximated by low order polynomials. Therefore, a large number of vanishing moments is not enough, it is also important to limit the impact of singularities in the signal to a small number of wavelet coefficients. This can be achieved by reducing the size of the support of the wavelet.

To limit the impact of these singularities the support of the wavelet needs to be of minimal size so that only coefficients at the close vicinity will be affected. If the wavelet has compact support of size N, at each scale 2^j only N wavelet coefficients $< f, \psi_{j,n} >$ will be affected by an isolated singularity. The amplitude of these coefficients is potentially high and neglecting them usually greatly impacts the quality of the representation. Thus, using a wavelet with large support can lead to a very inefficient representation for signals with a significant number of singularities. To maximise the sparsity of the representation in this case a minimal support size should be preferred.

In general, the size of the support and the number of vanishing moments are not necessarily correlated. However, in the specific case of wavelet bases the size of the support scales with the number of vanishing moments as described by the following property due to Daubechies (Daubechies, 1988):

Proposition 4.1. *If ψ is a wavelet with p vanishing moments that generates an orthonormal basis of $\mathbf{L}^2(\mathbb{R})$, then it has a support of size larger than or equal to $2p-1$. In particular, Daubechies wavelets reach this minimum with a compact support equal to $[-p+1, p]$.*

Therefore, vanishing moments and support size are two competing properties and Daubechies wavelets (presented in the following section) are optimal in that they have minimal support for a given number of vanishing moments. Depending on the nature of the signal to represent, one may therefore wish to favour a large number of vanishing moments for smooth signals with only few singularities or on the contrary a small support if the singularities are predominant.

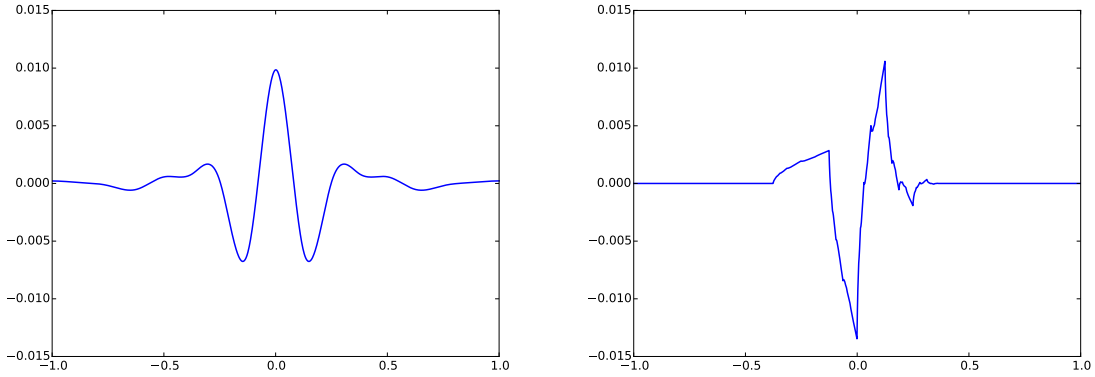

(a) Battle-Lemarié, $p = 4$

(b) Daubechies, $p = 2$

Figure 4.3: Battle-Lemarié and Daubechies wavelets with respectively 4 and 2 vanishing moments. By construction, the support of the Daubechies wavelet is compact whereas the Battle-Lemarié wavelet has infinite support.

4.2.3 A few wavelet bases

Battle-Lemarié

Battle-Lemarié wavelets (Battle, 1987; Lemarié, 1988) are based on spline multiresolution approximations. The scaling function is defined in Fourier space as:

$$\hat{\phi}(\omega) = \frac{\exp(-i\epsilon\omega/2)}{\omega^{m+1}\sqrt{S_{2m+2}(\omega)}}, \tag{4.29}$$

where $S_n(\omega) = \sum_{k=-\infty}^{+\infty} \frac{1}{(\omega+2k\pi)^n}$, and $\epsilon = 1$ if m is even or $\epsilon = 0$ if m is odd. And the corresponding wavelet is:

$$\hat{\psi}(\omega) = \frac{\exp(-i\omega/2)}{\omega^{m+1}} \sqrt{\frac{S_{2m+2}(\omega/2+\pi)}{S_{2m+2}(\omega)S_{2m+2}(\omega/2)}}. \tag{4.30}$$

For splines of degrees m, the wavelet ψ has $p = m + 1$ vanishing moments. Although Battle-Lemarié wavelets are not compactly supported, they have an exponential time decay. The Battle-Lemarié wavelet of order 3 (i.e 4 vanishing moments) is illustrated on Figure 4.3.

Daubechies wavelets

The wavelets proposed by Daubechies (Daubechies, 1988) have the orthogonal wavelets with a support of minimum size for a given number of vanishing moments p. When $p = 1$ the Haar wavelet is recovered. The regularity of the scaling and wavelet functions increase with the number of vanishing moments p. The Daubechies wavelet for $p = 2$ is illustrated on Figure 4.3.

4.2.4 Filter banks and the Fast Pyramidal decomposition

As mentioned in the previous section, the key to the fast wavelet transform is the possibility to recursively compute wavelet and approximation coefficients from one resolution to the next using

a couple of linear filters (h, g) as demonstrated by Equation (4.18) and Equation (4.27). A second essential feature is the ability to reconstruct the signal from its wavelet coefficients. For orthogonal wavelets, Mallat shows the same filters (h, g) form a couple of perfect reconstruction filters.

This defines the fast orthogonal wavelet transform, usually just called *Discrete Wavelet Transform* (DWT), which, given a function $f \in \mathbf{L}^2(\mathbb{R})$, recursively computes its approximation $a_j[n] =< f, \phi_{j,n} >$ and detail (or wavelet) coefficients $d_j[n] =< f, \psi_{j,n} >$ according to:

$$a_{j+1}[p] = \sum_{n\in\mathbb{Z}} h[n-2p]a_j[n] = [a_j * \bar{h}]_{\downarrow 2}[p] \,, \tag{4.31}$$

$$d_{j+1}[p] = \sum_{n\in\mathbb{Z}} g[n-2p]a_j[n] = [a_j * \bar{g}]_{\downarrow 2}[p] \,, \tag{4.32}$$

where $\bar{h}[n] = h[-n]$ and $\bar{g}[n] = [-n]$ and $[\cdot]_{\downarrow 2}$ stands for the decimation by a factor 2.

Conversely, approximation coefficients at scale 2^j can be reconstructed from approximation and detail coefficients at scale 2^{j+1} with:

$$a_j[p] = \sum_{n\in\mathbb{Z}} h[p-2n]a_{j+1}[n] + \sum_{n\in\mathbb{Z}} g[p-2n]d_{j+1}[n] \,, \tag{4.33}$$

$$= [a_{j+1}]_{\uparrow 2} * h[p] + [d_{j+1}]_{\uparrow 2} * g[p] \,. \tag{4.34}$$

The DWT can be implemented as a simple cascade of linear filters as illustrated on Figure 4.4. This filter bank implementation of the discrete wavelet is extremely, with a complexity in $O(N)$, in contrast to $O(N \log(N))$ of the FFT.

With orthogonal wavelets, the analysis and synthesis wavelets are identical and the same filters are used for the decomposition and reconstruction. However, perfect reconstruction can also be achieved, under certain conditions, if a different wavelet is used for the reconstruction. Let (h, g) and $(\tilde{h}, \tilde{g})$ be the filters associated with the analysis and synthesis wavelets. Then the two wavelets are said to be *bi-orthogonal* and perfect reconstruction is possible if the following bi-orthogonal conditions, proposed by Vetterli (Vetterli, 1986), are verified by the filter bank $(h, g, \tilde{h}, \tilde{g})$:

$$\hat{h}^*(\omega + \pi)\widehat{\tilde{h}}(\omega) + \hat{g}^*(\omega + \pi)\widehat{\tilde{g}} = 0 \,, \tag{4.35}$$

and

$$\hat{h}^*(\omega)\widehat{\tilde{h}}(\omega) + \hat{g}^*(\omega)\widehat{\tilde{g}}(\omega) = 2 \,. \tag{4.36}$$

Thanks to the relaxation of the strict requirements of orthogonality, wavelets pairs can be created that are symmetric, regular and compactly supported (Cohen et al., 1992). Given a bi-orthogonal filter bank $(h, g, \tilde{h}, \tilde{g})$, the same Fast wavelet transform algorithm holds is illustrated on Figure 4.4. Compared to orthogonal wavelets, only the reconstruction formula is modified and becomes:

$$a_j[p] = [a_{j+1}]_{\uparrow 2} * \tilde{h}[p] + [d_{j+1}]_{\uparrow 2} * \tilde{g}[p] \,. \tag{4.37}$$

4.2.5 2D Bi-Orthogonal Wavelets

The 1D wavelet introduced in the previous section can be extended to higher dimensions, and in particular to 2D by considering separable products of 1D wavelets. However, such a product can

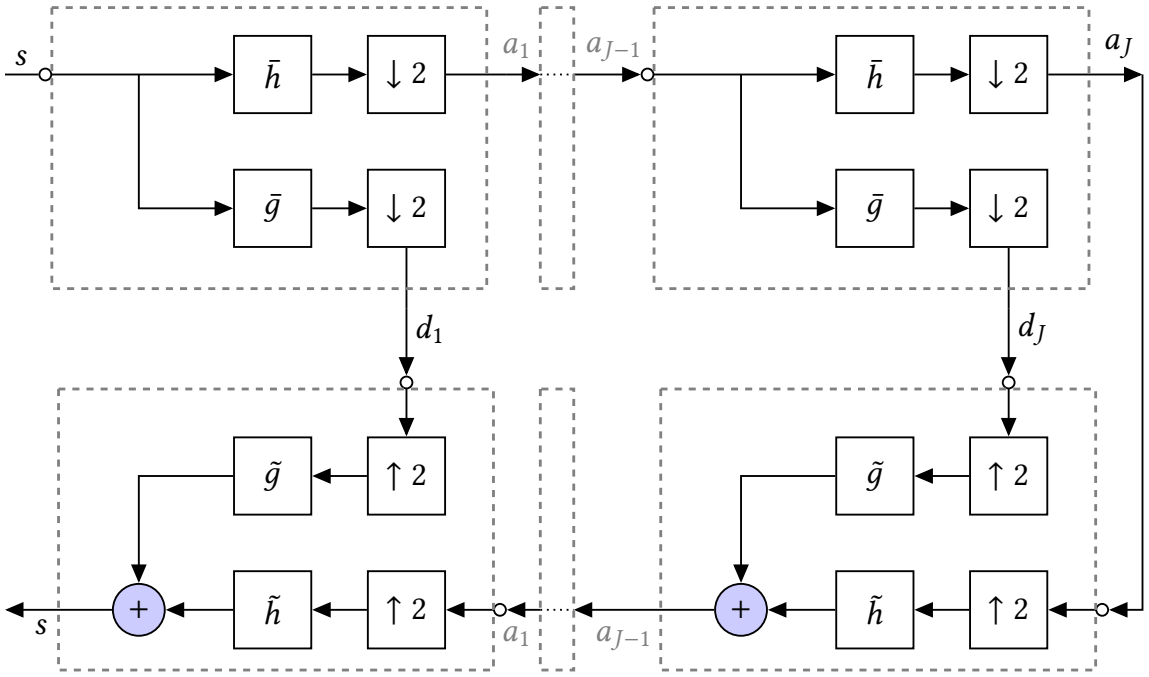

Figure 4.4: Fast pyramidal wavelet transform algorithm for a bi-orthogonal filter bank $(h, g, \tilde{h}, \tilde{g})$, note that in the case of orthogonal wavelets the same algorithm applies with $\tilde{h} = h$ and $\tilde{g} = g$. The top part of the diagram represents the fast decomposition of an input signal s into a set of coefficients $\{d_1, d_2, \ldots, a_J\}$. The bottom part of the diagram represents the fast reconstruction of s from its decomposition coefficients.

be built in several ways depending on the nature of the dimensions. If the two dimensions are homogeneous, it is generally useful to probe details along each direction using the same physical scale. In the contrary, in the case where the dimensions are not homogeneous (for instance, distance and time or distance and energy) then it can useful to probe each direction using independent scales, for instance to isolate slow moving large scale objects.

For the last case, a wavelet basis of $\mathbf{L}^2(\mathbb{R}^2)$ can easily be constructed using the basis functions $\psi_{j_1,j_2,n_1,n_2}(x_1, x_2) = \psi_{j_1,n_1}(x_1)\psi_{j_2,n_2}(x_2)$. Note that the scale j_1 and j_2 are independent. Such a decomposition leads to 3 kinds of coefficients:

- detail-detail coefficients: $d^1_{j_1} d^2_{j_2}$
- approximation-detail coefficients: $a^1_{j_1} d^2_{j_2}$ and $d^1_{j_1} a^2_{j_2}$
- approximation-approximation coefficients: $a^1_{j_1} a^2_{j_2}$

This construction is not optimal when the dimensions of the signal are homogeneous, like for 2D images for instance. Indeed, in this case, we want to probe the signal at a the same given scale for all dimensions. Instead, a decomposition of $\mathbf{L}^2(\mathbb{R}^2)$ can be built by defining the concept of separable multiresolutions such that the approximation of a function $f(x_1, x_2)$ at scale 2^j will be the projection of f on the subspace $V^2_j = V_j \otimes V_j$. If one defines the detail space W^2_j the orthogonal supplement of V^2_j in V^2_{j-1}, then, as demonstrated for instance in Mallat (1999), a wavelet basis for W^2_j

can be built from separable products of 1D scaling functions ϕ and wavelet functions ψ. Consider the three functions:

$$\psi^1(x) = \phi(x_1)\psi(x_2), \quad \psi^2(x) = \psi(x_1)\phi(x_2), \quad \psi^3(x) = \psi(x_1)\psi(x_2)\,, \tag{4.38}$$

and denote for $1 \leq k \leq 3$

$$\psi^k_{j,n}(x) = \frac{1}{2^j}\psi^k\left(\frac{x_1 - 2^j n_1}{2^j}, \frac{x_2 - 2^j n_2}{2^j}\right). \tag{4.39}$$

Then for $j \in \mathbb{Z}$ the family of wavelets $\{\psi^1_{j,n}, \psi^2_{j,n}, \psi^3_{j,n}\}_{n\in\mathbb{Z}^2}$ is an orthonormal basis of W^2_j. And the family $\{\psi^1_{j,n}, \psi^2_{j,n}, \psi^3_{j,n}\}_{(j,n)\in\mathbb{Z}\times\mathbb{Z}^2}$ is an orthonormal basis of $\mathbf{L}^2(\mathbb{R}^2)$.

From this result, one can build a 2D DWT to compute the following approximation and detail coefficients:

$$a_{j+1} = [a_j * \bar{h}\bar{h}]_{\downarrow 2}\,, \tag{4.40}$$
$$d^1_{j+1} = [a_j * \bar{h}\bar{g}]_{\downarrow 2}\,, \tag{4.41}$$
$$d^2_{j+1} = [a_j * \bar{g}\bar{h}]_{\downarrow 2}\,, \tag{4.42}$$
$$d^3_{j+1} = [a_j * \bar{g}\bar{g}]_{\downarrow 2}\,, \tag{4.43}$$

where we use the notation $hg[x_1, x_2] = h[x_1]g[x_2]$. Note that for d^1 and d^2, the high-pass filter g is only applied in one direction, this will lead to purely horizontal and vertical wavelet coefficients. For d^3, the details are computed by applying the high-pass filter g in both directions, the resulting wavelet coefficients are diagonal.

Conversely, the reconstruction can recursively be implemented using the following formula:

$$a_j = [a_{j+1}]_{\uparrow 2} * hh + [d^1_{j+1}]_{\uparrow 2} * hg + [d^2_{j+1}]_{\uparrow 2} * gh + [d^3_{j+1}]_{\uparrow 2} * gg\,. \tag{4.44}$$

Just as in the 1D case, this reconstruction formula can also use a different set of filters $(\tilde{h}, \tilde{g})$ while preserving the perfect reconstruction property if the filter bank $(h, g, \tilde{h}, \tilde{g})$ verifies the bi-orthogonal conditions Equation (4.35) and Equation (4.36).

Since the family of wavelets $\{\psi^1_{j,n}, \psi^2_{j,n}, \psi^3_{j,n}\}_{(j,n)\in\mathbb{Z}^2}$ forms a basis of $\mathbf{L}^2(\mathbb{R}^2)$, the wavelet decomposition of an image is non redundant and the wavelet coefficients can be arranged to form an image of the same size as the input image. This representation of the wavelet decomposition, due to Mallat, is illustrated on Figure 4.5.

4.3 THE STARLET TRANSFORM

In the previous section we introduced the framework for building non-redundant discrete wavelet transforms. Although this construction has many advantages, it has one major drawback: the lack of translational invariance of the wavelets. Indeed, for orthogonal wavelets, at scale 2^j, the wavelet atoms $\psi_{j,n}$ are positioned using a grid of step $2^j n$. Although this is enough to capture the signal, for restoration applications it is also useful to compute wavelet coefficients at scale 2^j at each position. This results in redundant wavelet transforms where more coefficients than necessary to reconstruct the signal are computed. These wavelet decompositions offer more flexibility, at the cost of an increased redundancy.

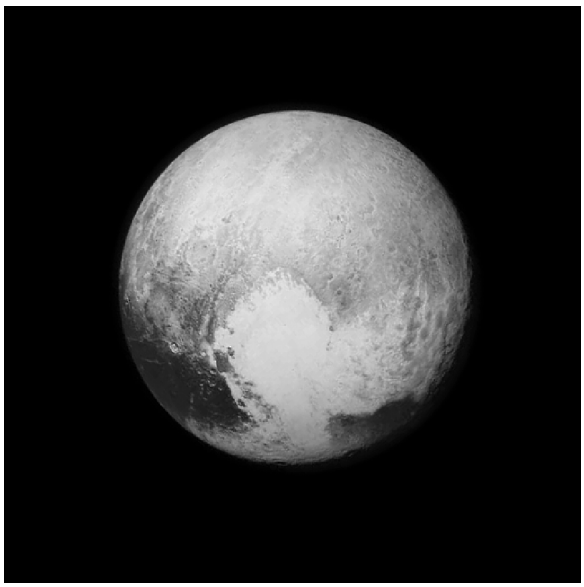

(a) Input image

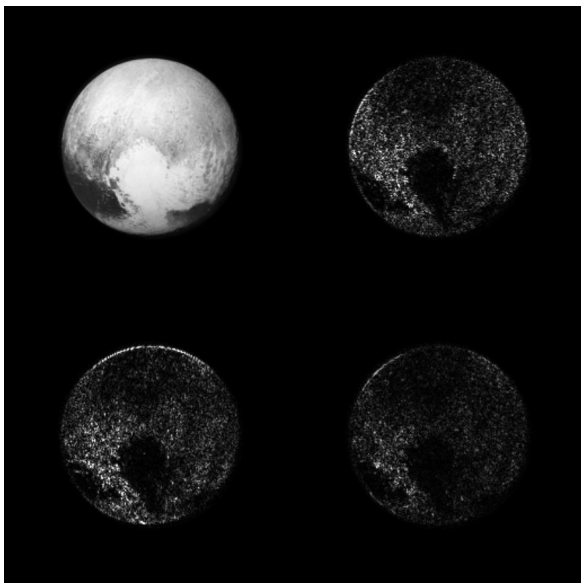

(b) DWT with one scale

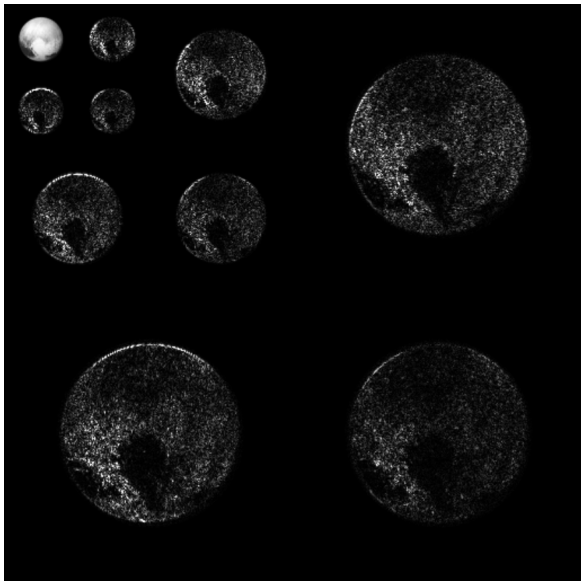

(c) DWT with 3 scales

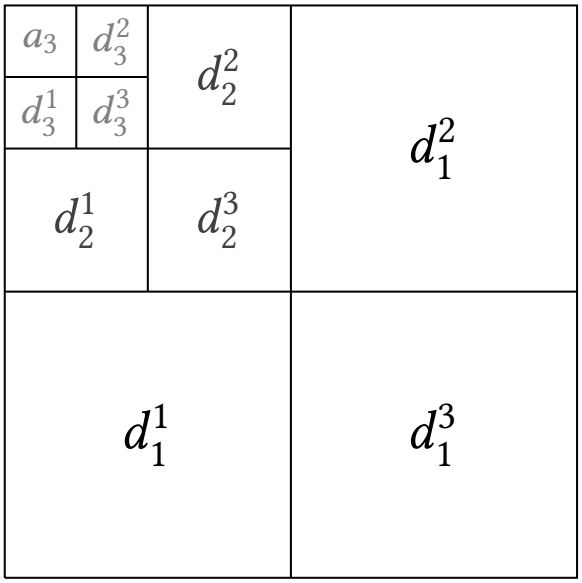

(d) Mallat representation

Figure 4.5: 2D Discrete Wavelet Transform of Pluto image.
2D DWT of New Horizons' image of Pluto (a) using an orthogonal Battle-Lemarié wavelet. The image is first decomposed into 3 wavelets sub-bands (for horizontal, vertical and diagonal directions) and a smooth approximation at the finest scale (b). Then the decomposition is recursively applied on the smooth approximation to compute the next 2 scales (c). The coefficients of the decomposition are arranged according to Mallat's representation. Credit: Input image from NASA/APL/SwRI.

4.3.1 À trous wavelet transform

The basic idea to build an *Undecimated Wavelet Transform* (UWT) from the DWT introduced in the previous section is to remove the decimation step. The resulting transform can efficiently be

performed with the "à trous" algorithm (with holes in french) (Holschneider et al., 1989; Shensa, 1992) which computes the wavelet decomposition as:

$$a_{j+1}[k] = a_j * \bar{h}^{(j)}[k] = \sum_{l\in\mathbb{Z}} h[l]a_j[l-2^j k]\,, \tag{4.45}$$

$$d_{j+1}[k] = a_j * \bar{g}^{(j)}[k] = \sum_{l\in\mathbb{Z}} g[l]a_j[l-2^j k]\,, \tag{4.46}$$

Compared to the DWT algorithm, different filters h^j and g^j are used at every scale and are obtained from the original filters h and g by inserting $2^j - 1$ zeros between every sample, thus the name "à trous". For instance the filters $h^{(1)}$ and $h^{(2)}$ can be written as:

$$h^{(1)} = [\ldots, h[-2], \quad 0, \quad h[-1], \quad 0, \quad h[0], \quad 0, \quad h[1], \quad 0, \quad h[2], \ldots] \tag{4.47}$$

$$h^{(2)} = [\ldots, h[-2], \quad 0,0,0, \quad h[-1], \quad 0,0,0, \quad h[0], \quad 0,0,0, \quad h[1], \quad 0,0,0, \quad h[2], \ldots] \tag{4.48}$$

It can be shown (Shensa, 1992) that the coefficients computed at position $2^j\ n$ using the à trous transform still correspond to the coefficients of the original DWT decomposition. The signal can still be reconstructed from these coefficients using the filter bank $(\tilde{h}, \tilde{g})$ with the following formula:

$$a_j[k] = \tilde{h}^{(j)} * a_{j+1}[k] + \tilde{g}^{(j)} * d_{j+1}[k]. \tag{4.49}$$

One important feature of the UWT is that the filter bank $(h, g, \tilde{h}, \tilde{g})$ no longer needs to verify the de-aliasing condition Equation (4.35) for exact reconstruction as the decimation step has been removed. It only needs to verify the exact reconstruction formula:

$$\hat{h}^*(\omega)\widehat{\tilde{h}}(\omega) + \hat{g}^*(\omega)\widehat{\tilde{g}}(\omega) = 2 \tag{4.50}$$

Therefore, this framework leaves much more flexibility for the design of filter banks which are no longer restricted by the bi-orthogonal conditions. This will be exploited in the next section to build the Starlet, a wavelet adapted to astronomical images.

Just like for the DWT, the UWT can be extended to 2D images using the same concept of separable multiresolutions. The decomposition of an image using J scales will lead to $3J$ wavelet sub bands (3 for each scale: horizontal, vertical and diagonal) and one smooth approximation, each with the same size as the input image. The redundancy of this transform in 2D is therefore $3J + 1$.

4.3.2 Isotropic undecimated wavelet: the Starlet transform

As mentioned in the previous paragraph, the UWT leaves a lot of freedom in the design of analysis and synthesis filter banks. In particular, this flexibility can be exploited to build wavelet decompositions adapted to astronomical data where objects are most of the time more or less isotropic. This has lead to the development of the *Isotropic Undecimated Wavelet Transform* (IUWT) (Starck and Murtagh, 2006). This transform is built around two criteria:

- The filters h, g must be symmetric.
- In 2-D or higher dimension, the wavelet and scaling function must be isotropic.

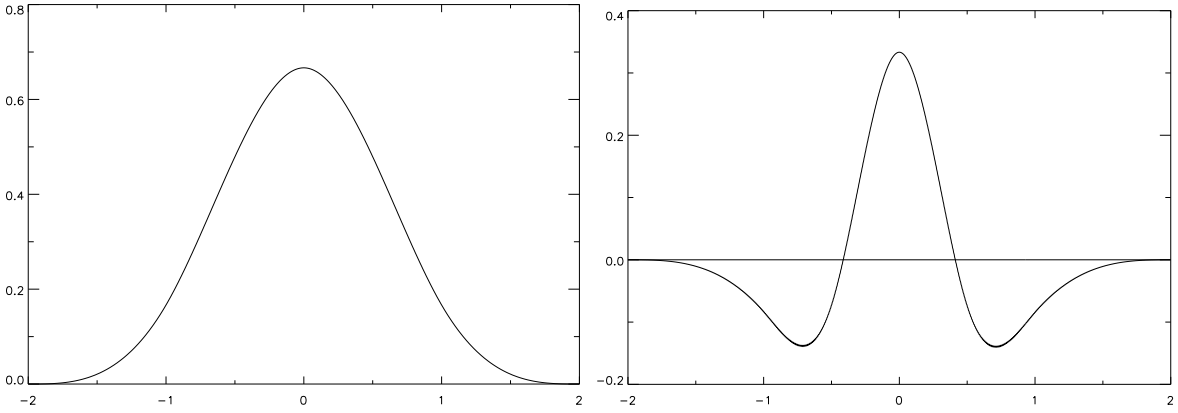

(a) Scaling function $\phi_{(1D)}$ **(b)** Wavelet function $\psi_{(1D)}$

Figure 4.6: B_3-spline scaling function and associated wavelet function corresponding to the "astro"-filter.

One simple way to construct such a decomposition is to start from an isotropic scaling function ϕ and to define the wavelet function as the difference between two successive approximations:

$$\frac{1}{4}\psi(\frac{x_1}{2}, \frac{x_2}{2}) = \phi(x_1, x_2) - \frac{1}{4}\phi(\frac{x_1}{2}, \frac{x_2}{2}) \tag{4.51}$$

Therefore, if ϕ is isotropic, so is ψ. This relationship between wavelet and scaling function can equivalently be written in terms of the associated filters h and g:

$$g[k, l] = \delta[k, l] - h[k, l] \tag{4.52}$$

where $\delta[k, l] = 1$ is $(k, l) = (0, 0)$ and $\delta[k, l] = 0$ otherwise. It can be shown Starck et al. (2007) that for any pair of even-symmetric analysis Finite Impulse Response filters $(h, g = \delta - h)$, this filter bank implements a frame decomposition for which perfect reconstruction is possible using Finite Impulse Response (FIR) filters. Indeed, based on the simple structure of the filter g, it can easily be seen that exact reconstruction can be achieved by simple summation of the wavelet coefficients:

$$a_0[k, l] = a_J[k, l] + \sum_{j=1}^{J} d_j[k, l] \tag{4.53}$$

Therefore, the IUWTde is entirely defined by the choice of the scaling function ϕ and associated filter h. The "astro"-filter, a specific choice of scaling function motivated by astronomical applications, was proposed in (Starck and Murtagh, 2006) to define the *Starlet* transform. This scaling function, a B-spline of order 3, is defined as:

$$\phi_{(1D)}(x) = \frac{1}{12}\left(|x-2|^3 - 4|x-1|^3 + 6|x|^3 - 4|x+1|^3 + |x+2|^3\right) \tag{4.54}$$

from which the N-dimensional scaling function can be built as a separable product of N ϕ_{1D} : $\phi(x_1, x_2) = \phi_{1D}(x_1)\phi_{1D}(x_2)$. These scaling and wavelet functions are illustrated on Figure 4.6. As can be seen, the wavelet function is very regular, with a small compact support and minimum oscillations. These properties make the Starlet an excellent choice for many restoration applications.

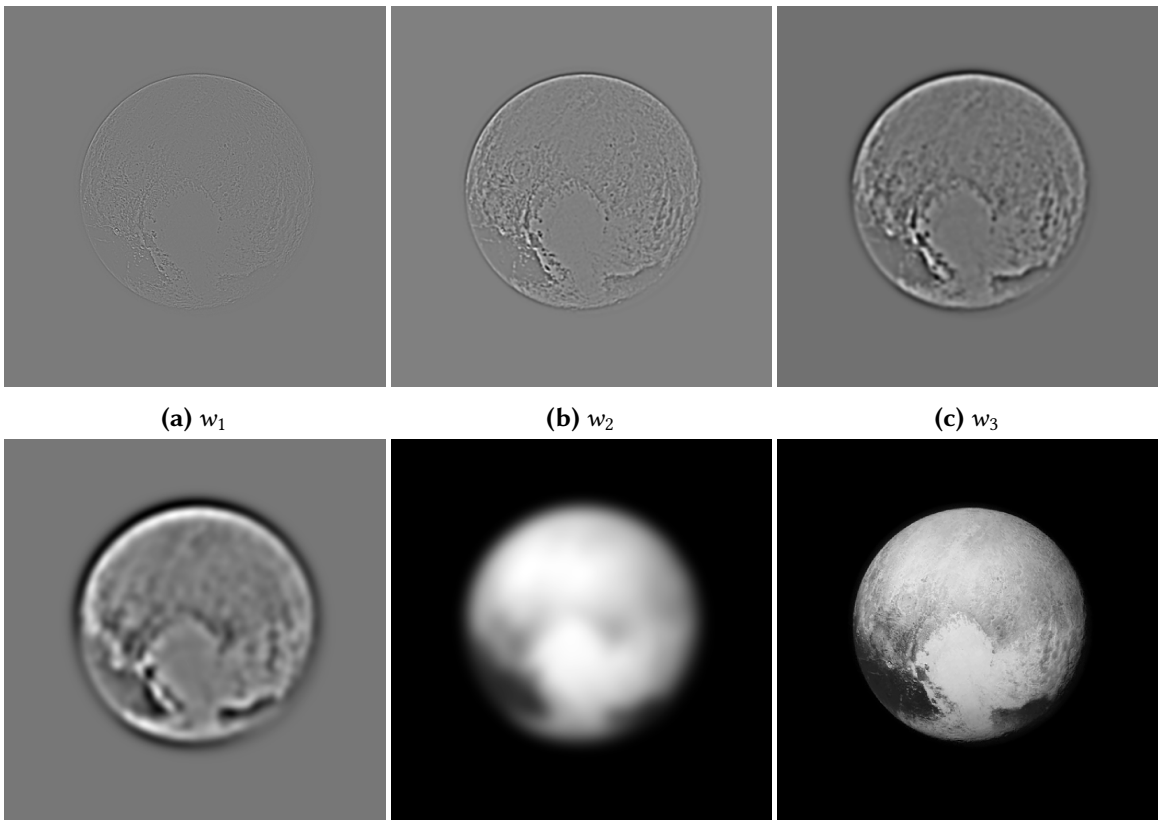

(a) w_1 **(b)** w_2 **(c)** w_3

(d) w_4 **(e)** a_4 **(f)** Input

Figure 4.7: Starlet transform of the Pluto image. Images (a)-(d) are the wavelet coefficients w_j for increasing scale j and (e) is the smooth approximation a_4. (f) is the input image which can be recovered by simple summation of the wavelet coefficients: (f) = (a) + (b) + (c) + (d) + (e). Credit: Input image from NASA/APL/SwRI.

The associated FIR filters h and g are:

$$h_{(1D)}[k] = \frac{1}{16}[1, 4, 6, 4, 1] \tag{4.55}$$

$$h[k, l] = h_{(1D)}[k]h_{(1D)}[l] \tag{4.56}$$

$$g[k, l] = \delta[k, l] - h[k, l] \tag{4.57}$$

The Starlet has the advantage of using a separable 2D scaling function which makes the computation of the 2D convolution products faster as they can be implemented as consecutive 1D convolutions along rows and columns.

Because of its isotropy, the Starlet only has one wavelet band per scale in contrast to the UWT which has three different directions for each scale. The redundancy of the Starlet is thus lower and is equal to $J + 1$ where J is the number of scales. An illustration of a Starlet decomposition is given in Figure 4.7.

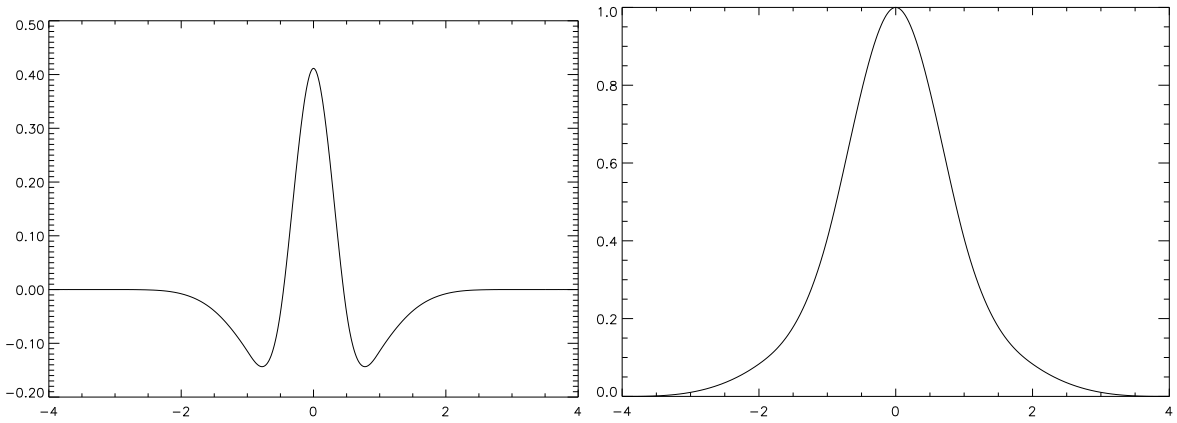

(a) Analysing wavelet ψ_{1D}^{gen2} **(b)** Synthesis function $\tilde{\psi}_{1D}^{gen2}$

Figure 4.8: Analysis wavelet $\psi_{(1D)}^{gen2}$ and positive synthesis function $\tilde{\psi}_{(1D)}^{gen2}$ for the second generation Starlet transform.

4.3.3 Second generation positive reconstruction Starlet

The Starlet transform introduced in the previous paragraph was defined based on considerations on the analysis (mainly that the analysis wavelet must be isotropic). However, no constraints were put on the synthesis (other than allowing exact reconstruction). Yet, for restoration applications, where a signal is reconstructed from modified wavelet coefficients, the quality of the results depends strongly on the properties of the synthesis operation. In particular, in the astronomical domain, most signals are positive, therefore a positive reconstruction formula (where positive wavelet coefficients lead to a positive reconstructed image) is a very desirable property to avoid ringing artefacts around positive structures.

With the freedom of wavelet design made possible by the UWT. Starck et al. (2007) have proposed a *second generation Starlet* which admits a positive reconstruction filter. This new transform uses the following filter bank:

$$h[k,l] = h_{(1D)}[k]h_{(1D)}[l] \tag{4.58}$$
$$g[k,l] = \delta[k,l] - h * h[k,l] \tag{4.59}$$
$$\tilde{h}[k,l] = h[k,l] \tag{4.60}$$
$$\tilde{g}[k,l] = \delta[k,l] \tag{4.61}$$

where h is the same low pass filter as for the first generation wavelet. One can check that this filter bank verifies the perfect reconstruction condition defined in Equation (4.50). The analysis filter g still corresponds to a wavelet as it still has zero mean but it is no longer the case of the reconstruction filter $\tilde{g}$ which is now positive. The analysis wavelet and synthesis function associated to these two filters are illustrated on Figure 4.8.

4.4 CONCLUSION

In this chapter, we introduced the theoretical background of wavelets as well as the practical fast algorithms available to compute the wavelet transform. In the context of sparse regularisation of inverse problems, wavelets are of special interest as they provide a family of sparse representations.

The applications presented in the rest of this thesis will rely heavily on the wavelet constructions introduced in this chapter. In particular, Starlets will be used in the context of weak lensing mass-mapping in Chapter 7 and Chapter 8 while the application presented in Chapter 5 will be based on undecimated bi-orthogonal wavelets.

Finally, the framework presented in this chapter in 1D and 2D will be extended to the 3D ball in Chapter 9 where we introduce a new isotropic 3D wavelet.

5 APPLICATION: SPARSE RECONSTRUCTION OF THE PRIMORDIAL POWER SPECTRUM

CONTENTS

This chapter presents, as a direct application of the sparse methodology introduced in Chapter 3 and Chapter 4, a new approach to the recovery of the power spectrum of primordial perturbations from measurements of the CMB. As will be described in the first section, this problem is a typical instance of an ill-posed linear inverse problem, with, as an additional difficulty, measurements contaminated by a multiplicative noise. We address this problem, using sparse regularisation, assuming that the PPS is sparse in a wavelet dictionary.

The resulting algorithm, coined PRISM, is tested on an extensive set of Wilkinson Microwave Anisotropy Probe (WMAP) simulations and applied to both the WMAP nine-year and Planck 2013 data, processed with the Local-Generalized Morphological Component Analysis (LGMCA) component separation pipeline (Bobin et al., 2013). We demonstrate how small features on the PPS can be accurately recovered using PRISM. When applied to WMAP and Planck data, we do not detect any significant deviations from the currently preferred near scale-invariant model, in accordance with the latest results from Planck (Planck Collaboration et al., 2015a).

The results presented in this chapter were published in Paykari et al. (2014) and Lanusse et al. (2014). This work has been conducted in collaboration with Paniez Paykari, Jean-Luc Starck, and the CosmoStat team at CEA/Saclay.

5.1 THE PRIMORDIAL POWER SPECTRUM RECONSTRUCTION PROBLEM

The PPS describes the initial curvature perturbations that over time evolved to form the large-scale structure we observe today. Because the physics of the early Universe are encoded in the PPS, it represents an invaluable probe of primordial cosmology, and measuring it is a crucial research area in modern cosmology. The currently favoured model describing the physics of the early Universe, inflation (Guth, 1981; Linde, 1982), produces initial perturbations from quantum fluctuations during an epoch of accelerated exponential expansion (see Section 2.3.1). This inflation process produces a power spectrum of specific shape and can leave characteristic features. For the simplest inflation models, the power spectrum, generated by almost purely adiabatic perturbations, is predicted to be nearly scale invariant. Hence, it is often expressed in terms of an amplitude A_s and a spectral index n_s with an optional 'running' α_s,

$$P(k) = A_s \left(\frac{k}{k_p}\right)^{n_s - 1 + \frac{1}{2}\alpha_s \ln\left(k/k_p\right)} , \tag{5.1}$$

where k_p is a pivot scale. We consider here only the first-order expansion of the spectral index, although higher orders can be considered (e.g., Debono et al., 2010). Exact scale invariance, known as the Harrison-Zeldovich model, which sets $n_s = 1$ (and $\alpha_s = 0$) (Harrison, 1970; Zeldovich, 1972), has been ruled out by different datasets. Instead, the near scale-invariant spectrum with $n_s < 1$ fits the current observations very well (e.g., Planck Collaboration et al., 2015b). More complex models generating deviations from scale invariance include those with features on the potential (Starobinsky, 1992; Adams et al., 2001; Wang et al., 2005; Hunt and Sarkar, 2004; Joy et al., 2008; Hunt and Sarkar, 2007; Pahud et al., 2009; Lerner and McDonald, 2009; Kumazaki et al., 2011; Meerburg et al., 2012; Ashoorioon and Krause, 2006; Ashoorioon et al., 2009), a small number of e-folds (Powell and Kinney, 2007; Nicholson and Contaldi, 2008), or other exotic inflationary models (Lesgourgues, 2000; Feng and Zhang, 2003; Mathews et al., 2004; Jain et al., 2009; Romano and Sasaki, 2008; Piao et al., 2004; Choudhury et al., 2013; Choudhury and Mazumdar, 2014). Therefore, determining the shape of the PPS will allow us to evaluate how well these models of the early Universe compare to the observations, possibly rule out some of the proposed models, and thus provide some much needed insight into the conditions of the primordial Universe.

The difficulty of course is that the PPS is no longer directly observable today. Nevertheless, the initial curvature perturbations present in the very early Universe and described by the PPS have evolved with time to form the structures we observe today. In particular, the CMB still bears a relatively clean imprint of the initial perturbations. Indeed, thanks to the linearity of the physical processes at play at the time of the CMB emission, the initial perturbations can be linearly mapped to the CMB temperature anisotropies through a transfer function, so that the theoretical CMB power spectrum $C_\ell^{\rm th}$ can be linked to the unknown $P(k)$ through:

$$C_\ell^{\rm th} = 4\pi \int_0^\infty d\ln k \Delta_\ell^2(k) P(k) , \tag{5.2}$$

where ℓ is the angular multipole and $\Delta_\ell(k)$ is the angular transfer function of the radiation anisotropies, which depends on the cosmological parameters responsible for the evolution of the Universe. This transfer function exhibits a high level of degeneracy due to projection effects which makes recovering the PPS $P(k)$ an instance of an ill-posed linear inverse problem. This problem proves to be

particularly difficult as the statistics of the measured CMB power spectrum are not Gaussian and an accurate estimation of this power spectrum is impeded by instrumental noise and various masks applied to the data. Finally, as the CMB spectrum is jointly sensitive to the primordial spectrum and the cosmological parameters in the transfer function, there is an induced degeneracy between them. The impact and level of this degeneracy have been investigated in (Paykari and Jaffe, 2010). A joint estimation of the cosmological parameters and a free form PPS would be prohibitively expensive to perform (as the parameter space can become very large). As a result, a parametric form of the PPS is assumed when jointly estimating this spectrum along with the other cosmological parameters. This potentially hides degeneracies between the cosmological parameters in the transfer function and the form of $P(k)$. One way to break this degeneracy is by adding extra information, such as polarisation or LSS data (Hu and Okamoto, 2004; Nicholson and Contaldi, 2009; Mortonson et al., 2009).

There are generally two approaches to determine the shape of the PPS, one by parametrisation and the second by reconstruction. Numerous parametric approaches that search for features with a similar form to those in complex inflationary models have been performed along with a simple binning of $P(k)$ (Bridle et al., 2003; Parkinson et al., 2005; Sinha and Souradeep, 2006; Sealfon et al., 2005; Mukherjee and Wang, 2005; Bridges et al., 2006, 2007; Covi et al., 2006; Hazra et al., 2010; Joy et al., 2009; Verde and Peiris, 2008; Paykari and Jaffe, 2010; Guo et al., 2011; Goswami and Prasad, 2013). Non-parametric methods, which make no assumptions about the model of the early Universe, have also been probed (Hannestad, 2001; Wang and Mathews, 2002; Matsumiya et al., 2002; Shafieloo and Souradeep, 2004; Bridle et al., 2003; Kogo et al., 2004a; Mukherjee and Wang, 2003; Hannestad, 2003; Kogo et al., 2004b; Tocchini-Valentini et al., 2005; Leach, 2006; Shafieloo et al., 2007; Shafieloo and Souradeep, 2008; Nagata and Yokoyama, 2008, 2009; Nicholson and Contaldi, 2009; Nicholson et al., 2010; Hazra et al., 2013). For an extensive review on how to search for features in the PPS using a wide range of methods, refer to the following papers and the references therein, which provide a sample on non-parametric reconstruction: deconvolution (Tocchini-Valentini et al., 2006; Ichiki and Nagata, 2009; Ichiki et al., 2010), Richardson-Lucy deconvolution (Hamann et al., 2010; Shafieloo et al., 2007), smoothing splines (Verde and Peiris, 2008; Peiris and Verde, 2010; Sealfon et al., 2005; Gauthier and Bucher, 2012), linear interpolation (Hannestad, 2003; Bridle et al., 2003), and Bayesian model selection (Bridges et al., 2009; Vazquez et al., 2012).

The recent Planck mission CMB temperature anisotropy data constrain the spectral index to $n_s = 0.9655 \pm 0.0062$ (Planck Collaboration et al., 2015b), ruling out exact scale invariance at over 5σ. Planck also failed to find a statistically significant running of the scalar spectral index, obtaining $\alpha_s = -0.0134 \pm 0.0090$. On the other hand, high-resolution CMB experiments, such as the South Pole Telescope (SPT)[1], report a small running of the spectral index; $-0.046 < \alpha_s < -0.003$ at 95% confidence (Hou et al., 2014). However, in general, any such detections have been weak and were consistent with zero.

Furthermore, the Planck collaboration extensively investigated features in the PPS. Initially, in Planck Collaboration et al. (2014a) a penalised likelihood approach indicated that there might be a feature near the highest wavenumbers probed by Planck at an estimated significance of $\sim 3\sigma$. This nominally statistically significant feature was detected around $k \sim 0.13$ Mpc^{-1}. However, it has

1 http://pole.uchicago.edu/spt/index.php

been confirmed since that the large dip at $\ell \sim 1800$ in the CMB power spectrum, which is associated with residual electromagnetic interference generated by the drive electronics of the 4 K cooler, is in fact responsible for the features detected at these high wavenumbers. The detection of this feature was subsequently retracted in Planck Collaboration et al. (2015a) and the main conclusion of this analysis is now that Planck does not find any significant departure from a simple power law model.

5.2 MODELLING

5.2.1 Empirical power spectrum

A CMB experiment, such as Planck, measures the CMB temperature anisotropy $\Theta(\vec{p})$ in direction $\vec{p}$, which is described as $T(\vec{p}) = T_{\rm CMB}[1 + \Theta(\vec{p})]$. This anisotropy field can be expanded in terms of spherical harmonic functions $Y_{\ell m}$ as

$$\Theta(\vec{p}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\vec{p}) \,, \tag{5.3}$$

with $a_{\ell m}$ being the spherical harmonic coefficients. The CMB anisotropy $\Theta(\vec{p})$ is assumed to be Gaussian distributed, which makes the $a_{\ell m}$ independent and identically distributed (i.i.d.) Gaussian variables with zero mean, $\langle a_{\ell m} \rangle = 0$, and variance

$$\langle a_{\ell m} a^*_{\ell' m'} \rangle = \delta_{\ell \ell'} \delta_{m m'} C_\ell^{\rm th} \,, \tag{5.4}$$

where $C_\ell^{\rm th}$ is the CMB temperature angular power spectrum introduced in Equation (5.2). However, we only observe a realisation of this underlying power spectrum on our sky, which we can estimate using the *empirical power spectrum estimator* defined as

$$\widehat{C}_\ell^{\rm th} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |a_{\ell m}|^2 \,, \tag{5.5}$$

where $\widehat{C}_\ell^{\rm th}$ is an unbiased estimator of the true underlying power spectrum; this becomes $\langle \widehat{C}_\ell^{\rm th} \rangle = C_\ell^{\rm th}$ in the case of noiseless CMB data over full sky.

For a given ℓ, the empirical power spectrum follows a χ^2 distribution with $2\ell + 1$ degrees of freedom, as it is a sum of the squares of independent Gaussian random variables. To account for this variability, we recast the relation between $\widehat{C}_\ell^{\rm th}$ and $C_\ell^{\rm th}$ as

$$\widehat{C}_\ell^{\rm th} = C_\ell^{\rm th} \, Z_\ell \,, \tag{5.6}$$

where $Z_\ell = \sum_m |a_{\ell m}|^2 / L C_\ell^{\rm th}$, which is a random variable representing a multiplicative noise distributed according to

$$L Z_\ell \sim \chi^2_L \,, \qquad \text{where } L = 2\ell + 1 \,. \tag{5.7}$$

In particular, the standard deviation of the empirical power spectrum estimator for a given ℓ is $\sqrt{(2/L)}\, C_\ell^{\rm th}$.

5.2.2 Accounting for instrumental noise and partial sky coverage

So far, we have considered that the CMB anisotropy data was available on the full sky which is not possible in practice because of the different Galactic foregrounds. Applying a mask on the sky results in the a modification of the spherical harmonic coefficients of the CMB temperature anisotropy,

$$\tilde{a}_{\ell m} = \int \Theta(\vec{p}) W(\vec{p}) Y^*_{\ell m}(\vec{p}) d\vec{p} \,, \tag{5.8}$$

where $W(\vec{p})$ is the window function applied to the data. The presence of the window function induces correlations between the $a_{\ell m}$ coefficients at different ℓ and different m and hence Equation (5.4) is no longer true.

One can define the *pseudo power spectrum* $\widetilde{C}_\ell$ as the application of the empirical power spectrum estimator on the spherical harmonic coefficients of the masked sky. When data is contaminated with additive Gaussian stationary noise, the pseudo power spectrum is

$$\widetilde{C}_\ell = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |\tilde{a}_{\ell m} + \tilde{n}_{\ell m}|^2 \,, \tag{5.9}$$

where $\tilde{n}_{\ell m}$ are the spherical harmonic coefficients of the masked instrumental noise.

Following the MASTER method from Hivon et al. (2002), the pseudo power spectrum $\widetilde{C}_\ell$ and the empirical power spectrum $\widehat{C}^{\rm th}_\ell$ can be related through their ensemble averages,

$$\langle \widetilde{C}_\ell \rangle = \sum_{\ell'} M_{\ell\ell'} \langle \widehat{C}^{\rm th}_{\ell'} \rangle + \langle \widetilde{N}_\ell \rangle \,, \tag{5.10}$$

where $M_{\ell\ell'}$ describes the mode-mode coupling between modes ℓ and ℓ' resulting from computing the transform on the masked sky. We note that in this expression $\langle \widehat{C}^{\rm th}_{\ell'} \rangle = C^{\rm th}_{\ell'}$ and we introduce the notations

$$C_\ell = \langle \widetilde{C}_\ell \rangle \qquad \text{and} \qquad N_\ell = \langle \widetilde{N}_\ell \rangle \,, \tag{5.11}$$

where C_ℓ and N_ℓ refer to the CMB and the noise power spectra of the masked maps, respectively.

We will also work under the approximation that the pseudo power spectrum $\widetilde{C}_\ell$ still follows a χ^2 distribution with $2\ell+1$ degrees of freedom and can be modelled as

$$\widetilde{C}_\ell = C_\ell Z_\ell \,, \tag{5.12}$$

$$= \left(\sum_{\ell'} M_{\ell\ell'} C^{\rm th}_{\ell'} + N_\ell \right) Z_\ell \,, \tag{5.13}$$

where Z_ℓ is defined in Equation (5.7).

5.2.3 Transfer function of the radiation anisotropies

Equation (5.13) relates the observables $\widetilde{C}_\ell$ to the theoretical CMB anisotropy power spectrum $C^{\rm th}_\ell$, taking into account instrumental noise, sample variance, and masking. The theoretical power spectrum $C^{\rm th}_\ell$ is itself related to the PPS through the convolution operation defined in Equation (5.2).

For a finite sampling of the wavenumber k, this convolution can be recast as a matrix operator $\mathbf{T}$ acting on the discretely sampled primordial spectrum, now referred to as P_k,

$$C_\ell^{\rm th} \simeq \sum_k T_{\ell k} P_k \ , \tag{5.14}$$

with matrix elements $T_{\ell k} = 4\pi \Delta \ln k \, \Delta_{\ell k}^2$, where $\Delta \ln k$ is the logarithmic k interval for the discrete sampling chosen in the integration of the system of equations. Because of the non-invertibility of the $\mathbf{T}$ operator, recovering the PPS P_k from the true CMB power spectrum $C_\ell^{\rm th}$ constitutes an ill-posed inverse problem. Finally, the complete problem we aim to solve can be condensed in the following form:

$$\widetilde{C}_\ell = \left(\sum_{\ell' k} M_{\ell \ell'} T_{\ell' k} P_k + N_\ell \right) Z_\ell \ . \tag{5.15}$$

We assume that the masked instrumental noise power spectrum N_ℓ is known for a given experiment. It can be computed from a JackKnife data map or from realistic instrumental noise simulations. Therefore, in the power spectrum of the data $\widetilde{C}_\ell$, only the PPS P_k remains unknown. Here we assume that the cosmology is known and hence operator $\mathbf{T}$ is known.

As already mentioned, the degeneracies of the transfer function make the recovery of the PPS an ill-posed inverse problem. The presence of the multiplicative noise Z_ℓ further complicates the problem and requires a specific treatment.

5.2.4 Variance stabilisation

One of the major difficulties encountered while solving this problem comes from the multiplicative nature of the noise Z_ℓ in Equation (5.13). Indeed, all the methods introduced in Chapter 3 to solve linear inverse problems assume additive Gaussian noise. One way to address this problem is to apply a variance stabilisation scheme to the data to turn the multiplicative noise into an additive Gaussian noise of unit variance. This approach has been applied to the CMB power spectrum in Paykari et al. (2012) to provide a proper treatment of the non-Gaussian noise on $\widehat{C}_\ell^{\rm th}$ based on the Wahba Variance Stabilisation Transform (VST). After the variance stabilisation is applied, the noise on $\widehat{C}_\ell^{\rm th}$ can be treated as an additive Gaussian noise with zero mean and unit variance. The VST operator $\mathcal{T}$ is defined as

$$\mathcal{T} : x \in \mathbb{R}^+ \mapsto \frac{\ln x - \mu_L}{\sigma_L} \ , \tag{5.16}$$

where $\mu_L = \psi_0(L/2) - \ln(L/2)$ and $\sigma_L^2 = \psi_1(L/2)$, where ψ_m is the polygamma function $\psi_m(t) = \frac{d^{m+1}}{dt^{m+1}} \ln \Gamma(t)$. We denote C_ℓ^s as the stabilised empirical power spectrum after applying the VST and get

$$C_\ell^s = \mathcal{T}(\widehat{C}_\ell^{\rm th}) = \frac{\ln C_\ell^{\rm th}}{\sigma_L} + \epsilon_\ell \ , \tag{5.17}$$

where $\epsilon_\ell = (\ln Z_\ell - \mu_L)/\sigma_L \sim \mathcal{N}(0,1)$. The inverse operator of $\mathcal{T}$ can be defined as

$$\mathcal{R} : x \in \mathbb{R} \mapsto \exp(\sigma_L x) \ . \tag{5.18}$$

Equipped with this VST tool, the question is now how to incorporate it within the reconstruction problem to make it tractable with the methodologies and algorithm at our disposal. The VST cannot simply be applied to the input data $\tilde{C}_\ell$ to solve the problem in the transformed space as the logarithmic VST operator would make the problem non-linear.

However, it can be used to build an estimator of the quantity we are interested i.e. the residuals $R_\ell(X)$ between the true noiseless pseudo power spectrum C_ℓ and the reconstructed power spectrum $C_\ell(X) = \boldsymbol{MT}X + N_\ell$ given a PPS X:

$$R_\ell(X) = C_\ell - C_\ell(X) \,. \tag{5.19}$$

Note that the true pseudo power spectrum C_ℓ is unknown and so is $R_\ell(x)$. The aim is to find an estimator $\overline{R}_\ell(X)$ of $R_\ell(X)$ based on the measured $\tilde{C}_\ell$ so that:

$$\overline{R}_\ell(X) = R_\ell(X) + n_\ell \tag{5.20}$$

where n_ℓ is a Gaussian noise with zero mean and covariance Σ. If such an estimator can be found then the inverse problem can be stated as a minimisation problem with a χ^2 data fidelity term of the form $\frac{1}{2} \parallel \Sigma^{-1/2}\overline{R}_\ell(X) \parallel_2^2$.

Consider the following difference:

$$\mathcal{T}(\widetilde{C}_\ell) - \frac{\ln(C_\ell(X))}{\sigma_L} = \frac{\ln(C_\ell) - \ln(C_\ell(X))}{\sigma_L} + \epsilon_\ell \,, \tag{5.21}$$

$$= \frac{1}{\sigma_L} \ln\left(\frac{C_\ell}{C_\ell(X)}\right) + \epsilon_\ell \,, \tag{5.22}$$

$$= \frac{1}{\sigma_L} \ln\left(1 + \frac{R_\ell(X)}{C_\ell(X)}\right) + \epsilon_\ell \,, \tag{5.23}$$

where ϵ_ℓ is the Gaussian noise with zero mean introduced in Equation (5.17). Assuming that the residual $R_\ell(X)$ is small compared to $C_\ell(X)$, one can linearise the above equation, to a good approximation, as

$$\mathcal{T}(\widetilde{C}_\ell) - \frac{\ln(C_\ell(X))}{\sigma_L} \simeq \frac{1}{\sigma_L C_\ell(X)} R_\ell(X) + \epsilon_\ell \,, \tag{5.24}$$

and

$$R_\ell(X) \simeq C_\ell(X)\sigma_L \left(\mathcal{T}(\widetilde{C}_\ell) - \frac{\ln(C_\ell(X))}{\sigma_L}\right) - C_\ell(X)\sigma_L\epsilon_\ell \,. \tag{5.25}$$

In this expression, the variance of the noise, i.e. the second term in the above equation, depends on the current estimate $C_\ell(X)$. As we will need to estimate the variance of the noise propagated to the wavelet coefficients using Monte Carlo simulations, it would be too expensive to estimate this every time $C_\ell(X)$ changes. Therefore, we opted for an additional approximation and replace the term $C_\ell(X)\sigma_L$ by $C_\ell(X^0)\sigma_L$, where X^0 is now a fixed fiducial power spectrum which can be the initial guess of the solution. We can now define the estimator $\overline{R}_\ell(X)$ for $R_\ell(X)$ as

$$\overline{R}_\ell(X) \equiv C_\ell(X^0)\sigma_L \left(\mathcal{T}(\widetilde{C}_\ell) - \frac{\ln(C_\ell(X))}{\sigma_L}\right) \,. \tag{5.26}$$

We can now verify that this estimator behaves as expected:

$$\overline{R}_\ell(X) \simeq \frac{C_\ell(X^0)}{C_\ell(X)} R_\ell(X) + C_\ell(X^0)\sigma_L\epsilon_\ell \,, \tag{5.27}$$

$$\simeq \alpha\, R_\ell(X) + n_\ell \,, \tag{5.28}$$

where we have introduced for convenience $\alpha(X) = \frac{C_\ell(X^0)}{C_\ell(X)}$ and $n_\ell = C_\ell(X^0)\sigma_L\epsilon_\ell$. Note that in this expression, the approximation comes from the development to first order of the logarithm

and should remain small as long as the residuals are small. In particular, the quality of the approximation improves as the estimate of the PPS X converges to the true $P(k)$. However, unless $C_\ell(X^0) = C_\ell(X)$, the factor α is not equal to 1 thus this estimator yields a biased estimate of the amplitude of $R_\ell(X)$. Nevertheless, it still verifies the fixed-point property $\overline{R}_\ell(P_k^{\rm th}) = 0$ necessary to converge towards the solution and unless the estimated solution X deviates significantly from X^0, the ratio $C_\ell(X^0)/C_\ell(X)$ remains limited to within a few percents. Finally, in the iterative algorithm proposed in the next section, the fiducial PPS X^0 can be reset several times as the current estimated X as the algorithm converges towards a solution, thereby removing any potential bias on the residuals once the algorithm has converged. Besides these small caveats, the noise n_ℓ on the estimator $\overline{R}_\ell(X)$ is now Gaussian, additive, with zero mean and a fixed covariance Σ independent of the current estimate of the solution X, assumed to be diagonal with $\Sigma_\ell = (C_\ell(X^0)\sigma_L)^2$.

5.3 THE PRISM ALGORITHM

5.3.1 Formulation of the inverse problem

The problem of reconstructing the PPS is stated in Equation (5.15). Solving this problem has three inherent difficulties:

1. the singularity of the convolution operator $T_{\ell k}$, which makes the inverse problem ill-posed even in the absence of noise;
2. the multiplicative noise on the power spectrum;
3. the mask applied to the maps, inducing correlations on the power spectrum.

To address this complex inverse problem, we adopt the sparse regularisation framework introduced in Chapter 3 to recast the reconstruction of the PPS as a convex optimisation problem of the form:

$$\arg\min_X \frac{1}{2} \parallel C_\ell - (\mathbf{M}\mathbf{T}X + N_\ell) \parallel_2^2 + \lambda \parallel \mathbf{\Phi}^t X \parallel_1 , \tag{5.29}$$

where X is the reconstructed estimate for the PPS P_k. The first term in Equation (5.29) imposes a ℓ_2 fidelity constraint to the data while the second term promotes the sparsity of the solution in dictionary $\mathbf{\Phi}$. The parameter λ tunes the sparsity constraint. This problem naturally incorporates the inversion of both the MASTER mixing matrix and the radiative transfer function and thus addresses the first and third difficulties stated earlier.

However this problem is what we would ideally like to solve but because of sample variance, the true pseudo-power spectrum C_ℓ is unknown, not the actual measurements $\widetilde{C}_\ell$. This is linked to the second difficulty; the measurements are contaminated with a multiplicative noise which cannot be handled with the formulation of Equation (5.29). To overcome this issue, we use the variance stabilisation scheme introduced in Section 5.2.4 to state an alternative optimisation problem which is now tractable by the convex optimisation algorithms presented in Chapter 3:

$$\arg\min_X \frac{1}{2} \parallel \Sigma^{-1/2}\overline{R}_\ell(X) \parallel_2^2 + \lambda \parallel \mathbf{\Phi}^t X \parallel_1 , \tag{5.30}$$

where the estimator $\overline{R}_\ell(X)$ has been introduced in Equation (5.26). The following section details the algorithm used to solve this problem and addresses the problem of the choice of the regularisation parameter λ.

5.3.2 Sparse reconstruction algorithm

As was mentioned before, the solution of the ℓ_1 regularised problem tends to be biased. To avoid this problem, in practice the PRISM algorithm uses the re-weighted ℓ_1 approach of Candès et al. (2008) (see Section 3.2.3) which tends to give results closer to the ℓ_0 problem. This technique amounts to solving a sequence of weighted ℓ_1 problems of the form

$$\min_X \frac{1}{2} \parallel \frac{1}{C_\ell(X^0)\sigma_L}\overline{R}_\ell(X) \parallel_2^2 + \lambda \parallel \mathbf{W}\mathbf{\Phi}^t X \parallel_1 , \tag{5.31}$$

where $\mathbf{W}$ is a diagonal matrix applying a different weight for each wavelet coefficient, effectively modifying the shape of the ℓ_1 ball. By iteratively solving this problem and using the the previous estimate of X to define weights $\boldsymbol{W}$ for the next iteration, the sparsity of the solution can be reinforced. The steps of this reweighted analysis-based ℓ_1 recovery are summarised below:

1. Set $j = 0$, for each diagonal element of the weighting matrix $\mathbf{W}$ set $w_i^j = 1$. Set the first guess X^0 by fitting a pure scale-invariant PPS to the data $\widetilde{C}_\ell$.
2. Solve the weighted ℓ_1 problem (5.31) yielding a solution X^j.
3. Update the fiducial PPS with the current estimate: $X^0 = X^j$
4. Estimate the threshold levels λ_i by Monte-Carlo simulation.
5. Compute $\alpha_i^j = \mathbf{\Phi} X^j$ and update the weights according to:

$$w_i^{j+1} = \begin{cases} \frac{1}{|\alpha_i^j|/K\lambda_i} & \text{if } |\alpha_i^j| \geq K\lambda_i \\ 1 & \text{if } |\alpha_i^j| < K\lambda_i \end{cases} , \tag{5.32}$$

   where λ_i is the standard deviation propagated to the wavelet coefficients (see Section 5.3.3) and K is a given significance level.
6. Terminate on convergence or when reaching the maximum number of iterations, otherwise go to step 2.

In practice, we find that three iterations of this procedure are enough to reach satisfying convergence and de-biasing of our results and we see no further improvements by performing additional re-weightings.

To solve the relaxed problem (5.31) given a weighting matrix $\mathbf{W}$, one can use the popular algorithm ISTA, introduced in Section 3.3.2.1. This proximal forward-backward iterative scheme relies on the iteration

$$\widetilde{X}^{n+1} = X^n + \mu \mathbf{T}^t \mathbf{M}^t \frac{1}{(C_\ell(X^0)\sigma_L)^2}\overline{R}_\ell(X^n) , \tag{5.33}$$

$$X^{n+1} = \text{prox}_{K\mu\|\lambda \odot W\Phi^t \cdot\|_1}\left(\widetilde{X}^{n+1}\right) , \tag{5.34}$$

where μ is an adapted step size and $\mathrm{prox}_{K\mu\|\lambda\odot W\Phi^t\cdot\|_1}$ is the proximal operator corresponding to the sparsity constraint. The gradient descent step μ has to verify

$$0 < \mu \leq \frac{2}{\| \mathbf{T}^t\mathbf{M}^t(C_\ell(X^0)\sigma_L)^{-2}\mathbf{M}\mathbf{T} \|} , \tag{5.35}$$

where $\| \cdot \|$ is the spectral norm of the operator.

As was explained in Section 3.3, in the absence of a closed-form expression for the proximal operator, its value can be estimated by solving a nested optimisation problem:

$$\begin{cases} \hat{u} = \arg\min_{|u_i|\leq K\mu\lambda_i w_i} \frac{1}{2} \| \mathbf{\Phi} u - x \|_2^2 \\ \mathrm{prox}_{K\mu\|\lambda\odot W\mathbf{\Phi}^t\cdot\|_1}(x) = x - \mathbf{\Phi}\hat{u} \end{cases} . \tag{5.36}$$

We solve this optimisation problem at each iteration of the algorithm, using FISTA (Beck and Teboulle, 2009), a fast variant of ISTA.

The details of the algorithm solving this weighted problem are provided in Algorithm 5.1.

5.3.3 Choice of wavelet dictionary and regularisation parameter

As mentioned in the previous section, the regularisation parameter K can be set according to a desired significance level. In Equation (5.36), it can be seen that the wavelet coefficients u_i are constrained within a weighted ℓ_1 ball and correspond to the non-significant part of the signal. In order to place the radius of this ℓ_1 ball according to the expected level of noise for each wavelet coefficient, we propagate the noise on the estimator $\overline{R}_\ell$ from Equation (5.34) through the operator $\mathbf{\Phi}\mathbf{T}^t\mathbf{M}^t(C_\ell(X^0)\sigma_L)^{-2}$ and estimate its variance at each pixel and each wavelet scale. In practice, we estimate this noise level using Monte Carlo simulations of the noise on $\overline{R}_\ell$. We set each λ_i to the resulting variance for each wavelet coefficient. As a result, coefficients below $K\lambda_i$ will be considered as part of the noise and one only need to set a global parameter K to tune the sparsity constraint according to the noise level. In Paykari et al. (2014), we used a very conservative threshold at $K = 5$ which robustly suppressed the noise while still allowing us to recover all of our test features. For the analysis of the Planck data in Lanusse et al. (2014), we relaxed this parameter to $K = 4$ for better sensitivity.

The choice of wavelet $\mathbf{\Phi}$ will have an impact on the performance of feature recovery. In the following study, we use bi-orthogonal Battle-Lemarié wavelets of order 1 (see Section 4.2.3). These wavelets are very regular, have limited oscillations, are exponentially localised and have two vanishing moments, which makes them well suited to recover a near scale-invariant power spectrum in logarithmic scale. We have also tried various other orthogonal and bi-orthogonal wavelets but the Battle-Lemarié provided the best results on the features we tested. More physically motivated dictionaries could be used to reconstruct a specific type of feature predicted by a given theory.

5.4 VALIDATION AND RESULTS FOR WMAP NINE-YEAR DATA

5.4.1 Numerical simulations

To assess the performance of our non-linear algorithm we perform a series of reconstructions for three different types of PPS: a near scale-invariant spectrum with $n_s = 0.972$ (Hinshaw et al.,

Algorithm 5.1 Weighted analysis-based P_k sparse recovery

Require:
Pseudo power spectrum of the data: $\widetilde{C}_\ell$,
Instrumental noise power spectrum N_ℓ,
First guess PPS X^0,
Sparsity constraint parameter K,
Weights w_i for each wavelet coefficients.

Initialise $C_\ell^0 = \mathbf{MT}X^0$.
Compute variance σ_i of noise $\sim \mathcal{N}(0,1)$ propagated to wavelet coefficients through $\mathbf{\Phi T}^t\mathbf{M}^t(C_\ell(X^0)\sigma_L)^{-2}$ from Monte Carlo simulations.

for $n = 0$ to $N_{\max} - 1$ **do**
  $\overline{R}_\ell^n = C_\ell^0\sigma_L\left(\mathcal{T}(\widetilde{C}_\ell) - \frac{\ln(\mathbf{MT}X^n + N_\ell)}{\sigma_L}\right)$
  $\widetilde{X}^{n+1} = X^n + \mu\mathbf{T}^t\mathbf{M}^t(C_\ell(X^0)\sigma_L)^{-2}\overline{R}_\ell^n$

  Computing $\text{prox}_{\lambda\mu\|\mathbf{W}\mathbf{\Phi}^t\cdot\|_1}$:
  Initialise $u_1 = y_0 = \mathbf{\Phi}^t\overline{X}^{n+1}$, $t_1 = 1$.
  for $k = 1$ to $K_{\max} - 1$ **do**
    $\overline{u}_k = u_k + \mu'\mathbf{\Phi}\left(\overline{X}^{n+1} - \mathbf{\Phi}^t u_k\right)$
    $y_k = \overline{u}_k - \text{ST}_{\mu w_i K\sigma_i}(\overline{u}_k)$
    $t_{k+1} = (1 + \sqrt{1 + 4t_k^2})/2$
    $u_{k+1} = y_k + \frac{t_k - 1}{t_{k+1}}(y_k - y_{k-1})$
  end for

  Update of the reconstruction:
  $X^{n+1} = \widetilde{X}^{n+1} - \mathbf{\Phi}u_{K_{\max}}$
end for
Return: The reconstructed PPS $P_k = X^{N_{max}}$.

2013), a spectrum with a small running of the spectral index with $n_s = 0.972$ and $\alpha_s = -0.017$ (Hou et al., 2014), and a spectrum with $n_s = 0.972$ with a compensated feature around $k = 0.03$ Mpc^{-1}. The first two simple models are the most favoured by the current data and the spectrum with the feature (investigated in other works, see Nicholson and Contaldi (2009)) is only used to demonstrate the ability of the algorithm to detect and reconstruct isolated features. In all cases, the cosmological parameters responsible for the evolution of the Universe in the radiation transfer function are kept the same and according to the WMAP nine-year parameters; (Hinshaw et al., 2013), $\Omega_b h^2 = 0.02264$, $\Omega_c h^2 = 0.1138$, $\Omega_\Lambda = 0.721$, and $\tau = 0.089$.

For a thorough comparison of our simulations to the WMAP nine-year data we perform the Monte Carlo simulations at the level of the five WMAP frequency channels, taking into account the propagation of the instrumental noise through the component separation and masking steps. For each of the three test primordial spectra we produce a set of 2000 pseudo power spectra $\widetilde{C}_\ell$ by processing the simulated channels through the LGMCA component separation pipeline (Bobin

et al., 2013) before computing the empirical power spectrum of the masked maps. In detail, the simulations are produced using the following steps.

- **Frequency channels**: We simulate CMB maps at the five WMAP channels at frequencies 23, 33, 41, 61, and 94 GHz. The frequency dependant beams are perfectly isotropic Point Spread Functions (PSFs) and their profiles have been obtained as the mean value of the beam transfer functions at each frequency as provided by the WMAP consortium (nine year version).
- **Instrumental noise**: Noise maps for each channel have been generated as Gaussian realisations of pixel variance maps obtained by combining the nine one-year full-resolution hit maps as provided by the WMAP consortium.
- **Cosmic microwave background**: Gaussian realisations of the CMB are computed from the three power spectra $C_\ell^{\rm th}$, which were obtained by applying the radiation transfer function $\mathbf{T}$ to each of the three test PPS. The transfer function is computed using CLASS[2] (Lesgourgues, 2011) according to the best-fit WMAP nine-year cosmology. The CMB signal for each channel is then obtained by applying the corresponding beam to the simulated CMB map as well as the HEALPix window for *nside* of 1024.
- **LGMCA Component Separation**: Full sky 15 arcmin resolution maps are obtained by applying LGMCA, with the precomputed set of parameters (Bobin et al., 2013), to the five simulated channels for CMB and noise. Noisy full sky maps are obtained by adding the resulting signal and noise maps.
- **Masking**: Final maps are obtained by applying the WMAP mask kq85 mask with $f_{sky} = 0.75$.

The pseudo power spectra are obtained by applying the empirical power spectrum estimator to the simulated maps. The noise power spectrum N_ℓ is estimated by averaging the 2000 pseudo spectra of masked noise maps. Figure 5.1 shows an example of a masked noisy CMB map obtained from our simulation process. Figure 5.2 shows the pseudo power spectra for the three test primordial spectra as well as the instrumental noise power spectrum estimated from the simulations. The light blue crosses show one realisation of the pseudo power spectrum for the near scale-invariant PPS and the pink crosses show the one with a small running. As can be seen, the three different CMB spectra lie well within each other's noise band and on large and small scales they become almost indistinguishable. Hence to accurately reconstruct the three underlying PPS from these CMB spectra, a very good handle on both the instrumental noise and the sample variance is required.

5.4.2 Reconstructions of primordial power spectra

To apply PRISM to the simulated data, we build a transfer function $\mathbf{T}'$ adapted to the simulations so that it includes the effects of the 15 arcmin beam from LGMCA and the HEALPix window of $nside = 1024$. Using the same radiation transfer function $\mathbf{T}$ as computed for the simulations, the resulting transfer matrix $\mathbf{T}'$ can be written as

$$\mathbf{T}' = b_\ell^2 h_\ell^2 \mathbf{T}\mathbf{Q} \,, \tag{5.37}$$

2 http://class-code.net/

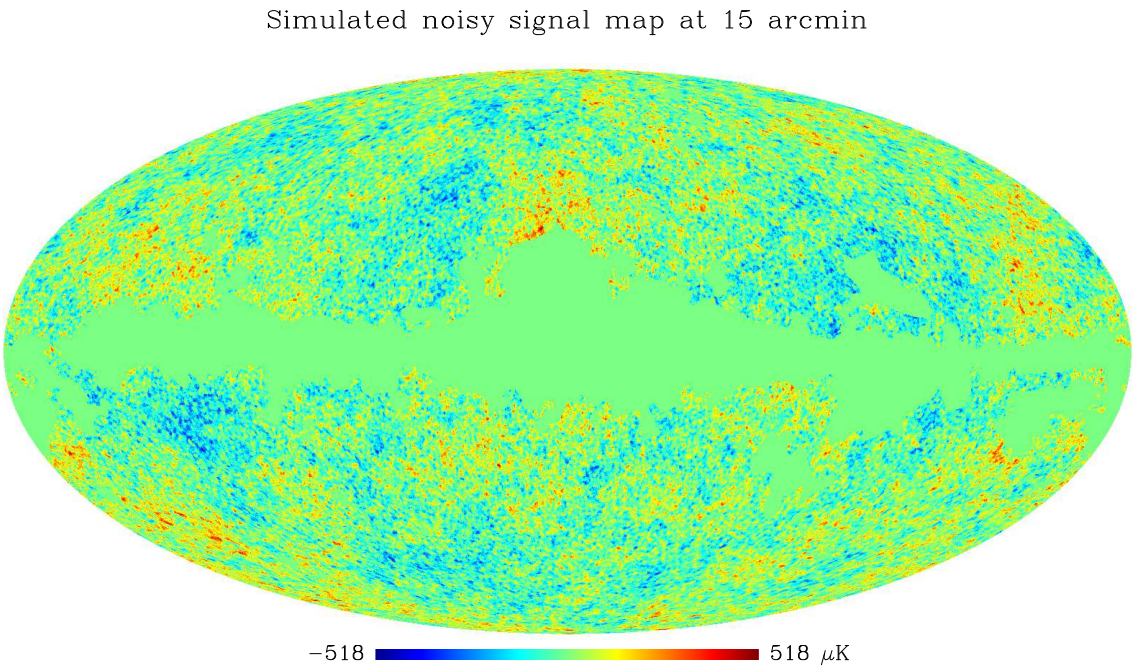

Figure 5.1: A simulated noisy CMB map at 15 arcmin resolution obtained from LGMCA and masked with the WMAP kq85 mask. The noise level corresponds to the WMAP nine-year data. This map was generated from a CMB power spectrum for a primordial spectrum with $n_s = 0.972$ and $\alpha_s = 0$.

where b_ℓ^2 and h_ℓ^2 are the beam and the HEALPix window, respectively, and $\mathbf{Q}$ is an operator performing a linear interpolation from the linear sampling in k of the CLASS transfer function $\mathbf{T}$ to a logarithmic scale using 838 points in the range $k \sim 10^{-4} - 0.15$ Mpc^{-1}. We also compute the MASTER coupling matrix $\mathbf{M}_{\mathrm{kq85}}$ corresponding to the kq85 high-resolution temperature analysis mask used in the simulations.

We now have all the ingredients necessary in our algorithm: $\mathbf{M}_{\mathrm{kq85}}$, $\mathbf{T}'$, and $\mathbf{\Phi}$, which we use to construct our algorithm and apply it to the 3×2000 simulated pseudo power spectra. We use the same set of parameters in PRISM for three types of primordial spectra: a $K\sigma$ significance level for the sparsity constraint with $K = 5$, three reweightings, and $N_{max} = 400$ iterations per reweighting.

In Figure 5.3a we show the reconstructed primordial spectra in the range $k \sim 0.001 - 0.10$ Mpc^{-1}. The blue lines show the 2000 reconstructed spectra for the spectrum with $n_s = 0.972$ and $\alpha_s = 0.0$ and the cyan lines show the reconstructions for the spectrum with $n_s = 0.972$ and $\alpha_s = -0.017$. In each case, the orange line is the mean of the reconstructions and the red line is the fiducial one.

The reconstruction of the PPS is limited by different effects on different scales. On very large scales, there are fundamental physical limitations placed on the recovery of the PPS by both the cosmic variance and the more severe geometrical projection of the modes. The physical limitations in the radiation transfer function places an inherent limitation at large scales meaning the PPS cannot be fully recovered on these scales, even in a perfect CMB measurement. On the other hand, on small scales we are limited by the instrumental noise. This leaves us with a window through which we can recover the PPS with a good accuracy. Nevertheless, as can be seen, for $k > 0.015$ Mpc^{-1} the PRISM algorithm can reconstruct the PPS to a great accuracy and easily distinguishes between the two types of spectra.

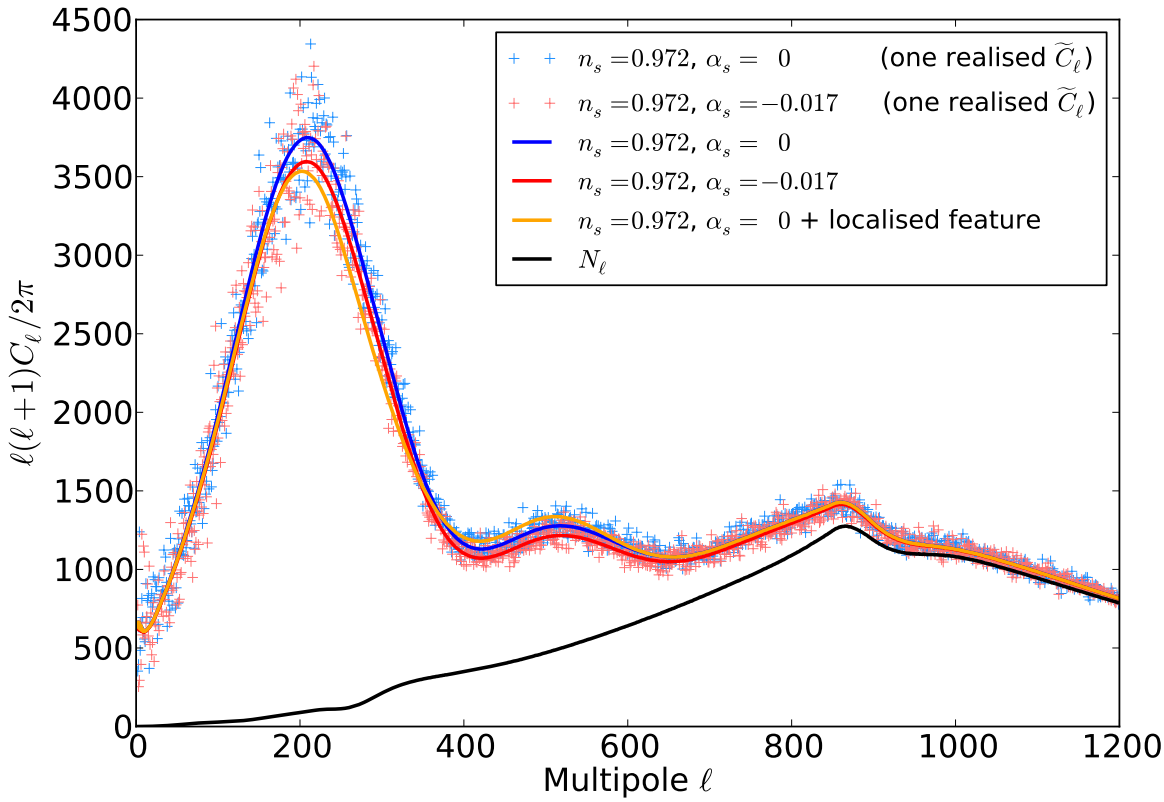

Figure 5.2: CMB pseudo power spectra for the three types of PPS. The blue solid line shows the pseudo spectrum based on a primordial spectrum with $n_s = 0.972$ and $\alpha_s = 0$. The light blue crosses show one simulation of this spectrum, computed from the map in Figure 5.1. The red line shows the pseudo spectrum for a primordial spectrum with $n_s = 0.972$ and $\alpha_s = -0.017$ and the orange line corresponds to a power spectrum with a localised feature at $k = 0.03$ Mpc^{-1}. These spectra include the effects of the mask, the 15 arcmin beam, the HEALPix window for *nside* of 1024, and the instrumental noise power spectrum, which is shown by a solid black line.

Figure 5.3b shows the 2000 CMB spectra obtained from the reconstructed primordial power spectra of each type. The blue lines show the CMB power spectra obtained from the near scale-invariant primordial spectra and the cyan lines show the ones for the primordial spectrum with a running. In each case, the orange line shows the mean of the reconstructions and the red line shows the fiducial spectrum. Comparing these CMB spectra to the input simulated ones, shown in Figure 5.2, illustrates the performance of the PRISM algorithm.

Figure 5.4 shows the performance of PRISM in reconstructing a localised feature in the PPS. The green lines show the 2000 individual reconstructions, the orange solid line shows the mean of the reconstructions, and the fiducial spectrum is shown in red. As can be seen, both the position and the amplitude of the feature can be recovered with good accuracy.

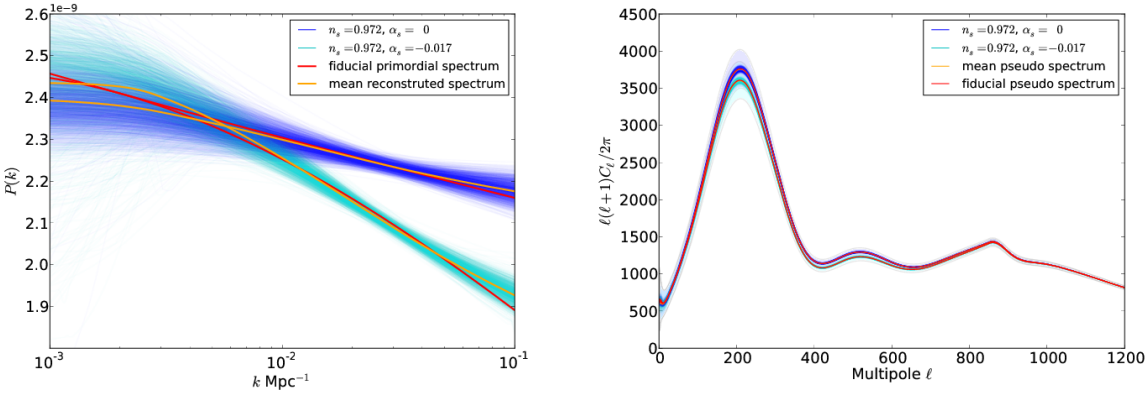

(a) Reconstructed primordial power spectra

(b) Corresponding CMB pseudo power spectra

Figure 5.3: Reconstructions for the PPS and their corresponding CMB pseudo spectra are shown. In blue we show the 2000 reconstructed spectra with $n_s = 0.972$ and $\alpha_s = 0$ and in cyan the reconstruction for $n_s = 0.972$ and $\alpha_s = -0.017$. In both cases the mean of the reconstructions is shown in orange and the fiducial input spectrum is shown in red. As can be seen, for $k > 0.015$ Mpc^{-1} PRISM can reconstruct the primordial power spectra with such accuracy that the two are easily distinguishable, despite their very similar forms in C_ℓ space. The shaded regions in the right-hand plot correspond to the 1σ sample (cosmic) variance, which demonstrates the similarity of the two types of CMB spectra.

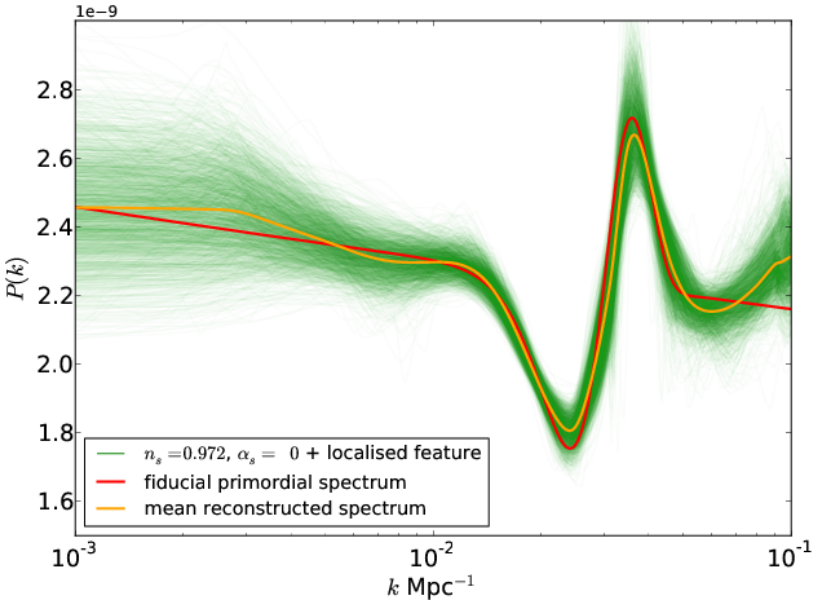

Figure 5.4: Reconstruction of the PPS with $n_s = 0.972$, $\alpha_s = 0.0$, and an additional feature around $k = 0.03$ Mpc^{-1} shown in green. The 2000 reconstructions are superimposed with their mean shown in orange. The fiducial input spectrum is shown in red.

5.4.3 Reconstruction from WMAP nine-year CMB spectrum

In the WMAP nine-year analysis (Hinshaw et al., 2013), the cosmological parameters in the radiation transfer function are fitted along with n_s and A_s, hence a power law form for the PPS is

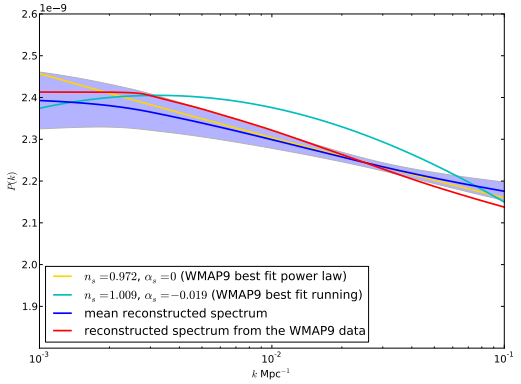

(a) Reconstructed primordial spectrum from WMAP nine-year data

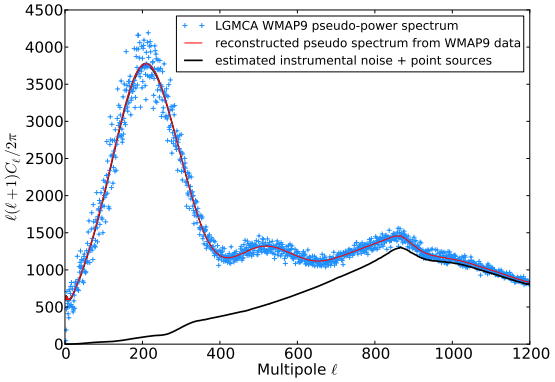

(b) Corresponding CMB pseudo power spectrum

Figure 5.5: Reconstruction of the PPS from the LGMCA WMAP nine-year data and its corresponding pseudo spectrum are shown in red. For comparison, we also show the mean of the reconstruction for $n_s = 0.972$ and $\alpha_s = 0$ with a solid dark blue line with the 1σ interval around the mean shown as a shaded blue region. The WMAP nine-year fiducial PPS with $n_s = 0.972$ and $\alpha_s = 0$ is shown in yellow and in cyan we show the best-fit PPS with a running from WMAP nine-year data with $n_s = 1.009$ and $\alpha_s = -0.019$. On the right, we plot the LGMCA WMAP nine-year pseudo power spectrum (blue crosses) and the estimated instrumental noise power spectrum including the point sources power spectrum is shown (solid black line). The very small blue region corresponds to the 1σ interval around the mean reconstructed spectrum (i.e. blue region on the left plot).

assumed. This means the transfer function computed using these best-fit parameters will always allow a power-law PPS to fit the observed data. However, reconstructing a free form PPS from the data, assuming the fiducial transfer function, allows us to test this null hypothesis by looking for significant deviations between the reconstructed spectrum from data and the simulations.

The WMAP nine-year data is processed using LGMCA as described in Bobin et al. (2013), which is the same pipeline as the one used to produce the simulations. As mentioned previously, a good handle on the noise power spectrum is critical in order to yield an unbiased reconstruction of the PPS. We estimate the noise power spectrum from the WMAP nine-year data by subtracting the cross-power spectrum from the auto-power spectrum and applying a denoising, using the TOUSI algorithm (Paykari et al., 2012). To account for the effect of point sources, which were not accounted for in the simulations, we add an estimate of the point sources power spectrum, computed from 100 simulations, to the estimated noise power spectrum. Figure 5.5b shows the pseudo-power spectrum computed from the LGMCA WMAP nine-year map (blue crosses) and the estimated instrumental noise power spectrum (black solid line). We note that in theory, the noise power spectrum could be computed from simulations. However, after comparing our estimated noise power spectrum from the 2000 simulations to the actual noise power spectrum in the WMAP nine-year data we found a small bias that we could not account for in the simulations. Hence we opted to use the data itself to estimate the noise power spectrum.

We apply PRISM, with the same parameters as in the simulations, to the WMAP nine-year LGMCA CMB pseudo power spectrum. The reconstructed PPS is shown in red in Figure 5.5a. In this figure, we overlay the 1σ interval around the mean of reconstructed primordial near scale-

invariant spectrum, obtained from the simulations. The best-fit power-law power spectrum from WMAP nine-year data with $n_s = 0.972$ and $\alpha_s = 0$ is shown in yellow, while the best-fit power spectrum with a running from WMAP nine-year data with $n_s = 1.009$ and $\alpha_s = -0.019$ is shown in cyan (Hinshaw et al., 2013). As can be seen, the reconstructed power spectrum from the data does not exhibit a significant deviation from the best-fit near scale-invariant spectrum. The small departure from the 1σ interval at small scales is not significant, especially since our simulations did not thoroughly take into account additional effects such as a beam uncertainty and point sources. To conclude, we find no significant departure from the WMAP nine-year best-fit near scale-invariant spectrum.

5.5 RESULTS FOR PLANCK 2013 DATA

Compared to WMAP, the Planck satellite is able to probe the CMB up to much higher multipoles and thus can constrain the PPS over a wider range of scales. In this section, we apply the PRISM algorithm on the first public release of the Planck data (hereafter Planck PR1) processed with the LGMCA[3] pipeline.

To estimate the mean and variance of the PRISM reconstruction, we set up a simulation pipeline for Planck data, mimicking the complete LGMCA pipeline applied to the actual data. The steps are very similar to the one described in Section 5.4.1.

We adopt as a base line the Planck PR1 best fit cosmology, which we used for the radiation transfer function $\mathbf{T}$ (computed using CAMB[4] instead of CLASS) and to define a fiducial near scale invariant PPS with $A_s = 2.215 \times 10^{-9}$ and $n_s = 0.9626$. The lensing contribution to the CMB temperature power spectrum, also computed with CAMB for the PR1 fiducial cosmology, was taken into account as an additional contribution to N_ℓ.

Using this model, we generate CMB realisations for the nine Planck frequency channels. On each of these realisations, we added simulated instrumental noise maps provided by the Planck team[5]. Unfortunately, only 100 out of their 1000 independent noise realisations were made publicly available, which dramatically reduced the number of simulations we were able to perform for the Planck data compared to the 2000 realisations used in the previous section for WMAP data.

For each simulation, we processed the nine frequency channels through LGMCA, with the pre-computed set of parameters (Bobin et al., 2014). Full-sky noisy maps with a 5 arcmin resolution were obtained, which were then masked using a Galactic and point sources mask with $f_{\rm sky} = 0.76$.

The pseudo power spectra were obtained by applying the empirical power spectrum estimator to the masked maps. We also built an estimate of the instrumental noise power spectrum N_ℓ by processing the noise maps through the same pipeline in LGMCA, masking the resulting noise maps and applying the pseudo spectrum estimator to these masked maps. We set our estimate of N_ℓ to the average of the 100 noise pseudo spectra.

In contrast to the approach taken in the previous section, where the algorithm was initialised to a scale-invariant power spectrum, we modified our choice of initialisation and used the best-fit Planck PR1 PPS to initialise the algorithm, as we search for small deviations from the best-fit power

3 LGMCA codes and Planck PR1 data are available at http://www.cosmostat.org/planck_pr1.html
4 http://camb.info
5 http://wiki.cosmos.esa.int/planckpla

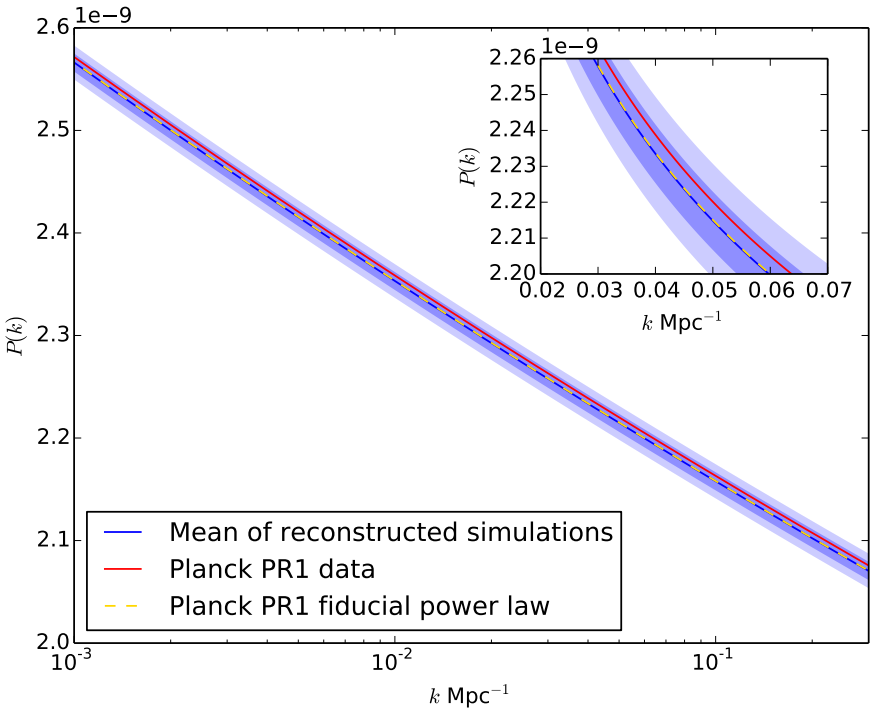

Figure 5.6: Reconstructions for the primordial power spectra from 100 simulations and Planck 2013 data. The 1σ and 2σ dispersion of the reconstructed spectra from the simulations are shown as blue bands around the mean of the reconstructions (blue line). We note that these bands do not include the errors due to point sources and beam uncertainties. The Planck fiducial power spectrum used for the simulations, with $A_s = 2.215 \times 10^{-9}$ and $n_s = 0.9626$, is shown as a yellow dashed line. The inset shows a close-up of the main figure in linear scale.

law that already fits the data. With this choice of initialisation, the reconstruction will not depart from the best-fit power law in the absence of evidence from the data.

As mentioned previously, we lowered the regularisation parameter K to $K = 4$ in this analysis of the Planck data, compared to $K = 5$ in the previous section. This lower level of regularisation increases the sensitivity of the algorithm while still robustly rejecting the noise due to sampling variance.

For Planck data, we expect to be able to effectively constrain the PPS in the range $k \sim 0.005 - 0.20$ Mpc^{-1}.

In Figure 5.6 we show the reconstructed spectra from the simulations and the data. The mean reconstructed power spectrum perfectly fits the input PR1 best-fit power law in the entire reconstructed range. Of course this does not mean that the algorithm is able to perfectly reconstruct an unknown power spectrum over this entire range, but that with the regularisation level used for these reconstructions, no significant departures from the best-fit power law have been detected. The reconstructed spectrum from the LGMCA PR1 power spectrum remains within the 1σ bar of the reconstructed spectra from the PR1 best-fit power law. Thus, we find no significant departure from the PR1 best-fit near scale-invariant spectrum.

As a complementary test of PRISM on Planck-like data we assessed the algorithm's ability to recover a small local departure from the best-fit PR1 power law. We created a set of CMB simulations from a fiducial PPS with a small localised test feature causing a dip in the angular power spectrum

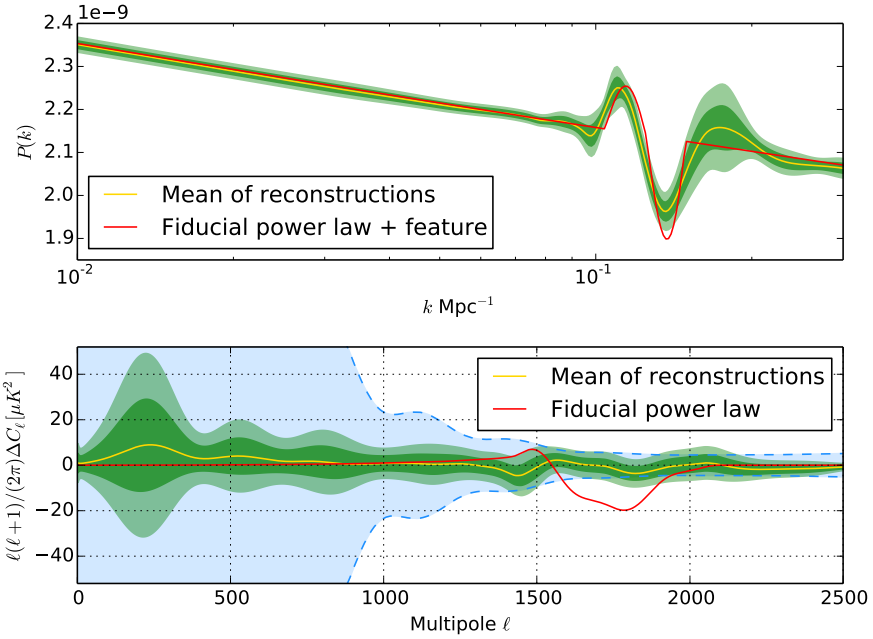

Figure 5.7: Top panel shows a fiducial PPS with a feature around $k = 0.125$ Mpc^{-1} in red and in green contours the 1σ and 2σ dispersion of 100 reconstructions from simulated CMB spectra. The mean of reconstructions is shown in yellow. The bottom panel shows the residuals ΔC_ℓ between the C_ℓ for the fiducial PPS with a feature and the C_ℓ for the best-fit Planck power law in red and for the mean reconstructed PPS in yellow. The green bands indicate the 1σ and 2σ bands for the ΔC_ℓ from the simulations, the dashed blue lines show the 1σ region due to cosmic variance.

around $\ell \sim 1800$. The aim of this set of simulations was to mimic the feature that Planck Collaboration et al. (2014b) proposed to be accountable for the large dip in the angular power spectrum, which was later confirmed as being caused by residual electromagnetic interferences. Our test PPS was built from the best-fit PR1 power law with an added feature around $k = 0.125$ which causes a dip in the angular power spectrum around $\ell \sim 1800$. This feature and the residuals ΔC_ℓ between the fiducial angular power spectrum and the PR1 best fit C_ℓ are shown in Figure 5.7.

From this test PPS, we generated a set of 100 CMB simulations using the exact same procedure as previously mentioned, and we applied PRISM to the measured angular pseudo-power spectra with the exact same parameters. As can be seen in Figure 5.7, the feature is successfully detected, and the reconstruction shows little bias in position and amplitude. Using the PPS reconstructed with PRISM enables a much better fit to the data than a power law, and the reconstructed angular power spectra fall inside the 1σ region due to cosmic variance. If such a feature existed in the LGMCA processed Planck PR1 data, PRISM would therefore have been able to detect it.

5.6 CONCLUSION

In this chapter, we presented a direct application of the sparse regularisation framework introduced in Chapter 3 and Chapter 4. Based on this framework, we proposed a new non-parametric method for the recovery of the Primordial Power Spectrum from the measured CMB power spectrum. This inverse problem is particularly important in the current cosmological context as the PPS is a

fundamental probe of inflation and an accurate reconstruction of this power spectrum can provide some much needed insight into the physics of the early Universe.

As we demonstrated on simulations, our method is able to robustly recover small isolated features as well as the general shape of the PPS. We applied our reconstruction algorithm on WMAP nine-year and Planck data but despite the sensitivity of the method we do not detect any significant deviations from the best-fit near scale-invariant power spectrum.

To reconstruct the PPS over a wider range of scales and with more accuracy, this method can easily be extended to include polarisation and LSS information.

Part II

WEAK LENSING MASS-MAPPING

6 WEAK GRAVITATIONAL LENSING

CONTENTS

Under General Relativity, gravitation can be understood as a consequence of the curvature of space-time, caused by the presence of massive objects. One of the predictions of this theory is that light propagating in a curved space-time will also be deflected with a specific angle. This prediction led to one of the first observational confirmations of Einstein theory in 1919 by Eddington's expedition to observe the deflection of starlight around the sun during a total solar eclipse Dyson et al. (1920). Almost a century latter, gravitational lensing is now a well established cosmological tool.

At the scale of the Universe, the light coming to us from distant galaxies propagates through the large scale structure and gets deflected by the gravitational potentials of these structures as illustrated by Figure 6.1. Because of differential deflections of the light bundles coming from these galaxies, their apparent images will undergo some deformations ranging from the most subtle shearing to extreme stretching and the formation of arcs. All these effects are referred to as gravitational lensing but two main categories can be identified, depending on the strength of the effect: *strong* and *weak* gravitational lensing.

Strong lensing mostly occurs when the light from distant galaxies encounters a massive galaxy cluster along the way or passes within the vicinity of another galaxy. In such cases, extreme effects can be observed such as the formation of multiple images, arcs and even Einstein rings. An example

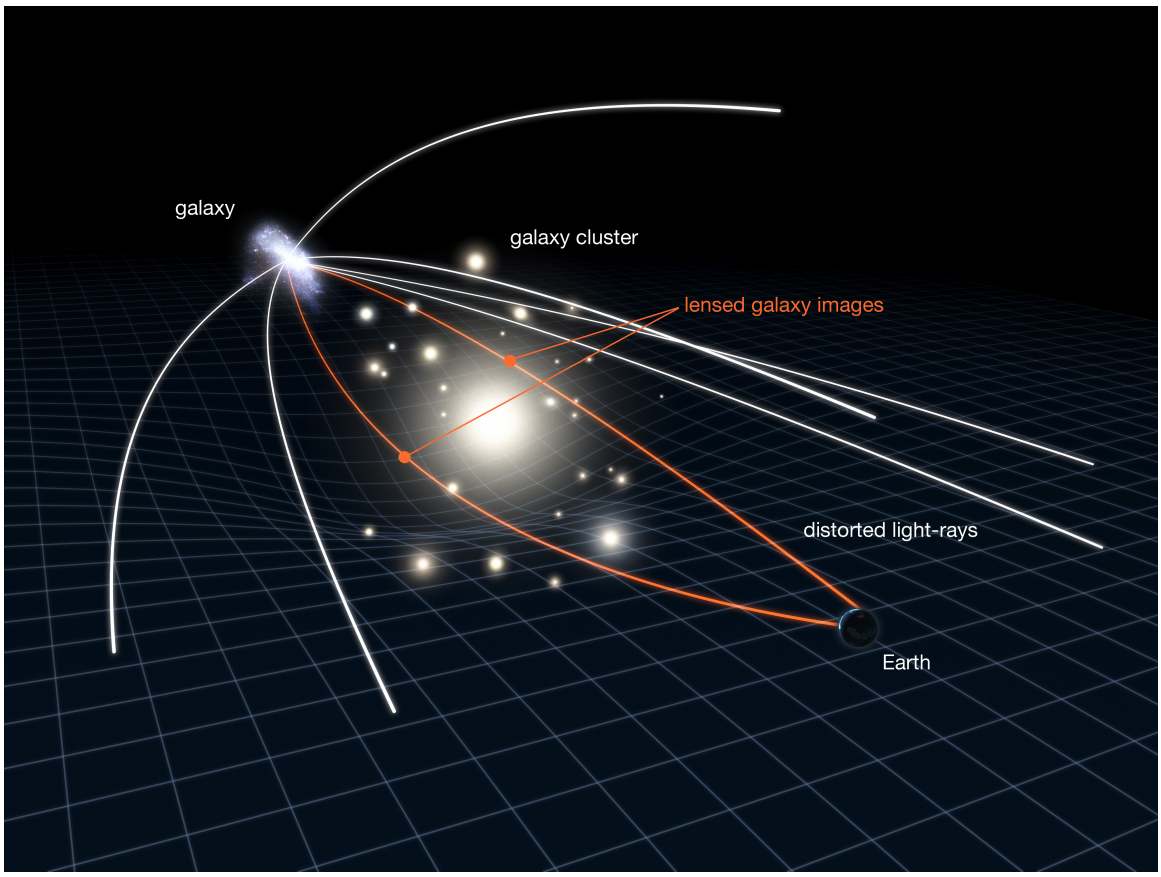

Figure 6.1: Illustration of the lensing effect caused by a massive cluster of galaxies. Credit: NASA, ESA & L. Calçada

of a strong lensing system is shown on Figure 6.2. In this image of the Abell 1689 galaxy cluster, the presence of a massive lens can be seen by eye with the presence of a large number of arcs.

Although the spectacular deformations induced by strong lensing provide a wealth of information about galaxy clusters and cosmology, in the rest of this thesis we will more particularly focus on the less extreme weak lensing regime. Contrary to strong lensing, the deformations in the weak regime are very subtle and cannot be identified by eye. Although much less extreme, weak lensing is ubiquitous. Virtually all galaxies in the sky are lensed to some extent by the large scale structure and the statistical analysis of this effect allows us to directly probe the matter content of the Universe, and thus constrain cosmology.

The goal of this chapter is to introduce the various notions of weak gravitational lensing which will be required for the following chapters. In particular, we review existing mass-mapping techniques for the reconstruction of the dark matter distribution in 2D and 3D.

Figure 6.2: Abell 1689 Galaxy cluster, as seen by the ACS camera of the Hubble Space Telescope. Several strongly lensed galaxies are visible. Credit: NASA, N. Benitez (JHU), T. Broadhurst (Racah Institute of Physics/The Hebrew University), H. Ford (JHU), M. Clampin (STScI), G. Hartig (STScI), G. Illingworth (UCO/Lick Observatory), the ACS Science Team and ESA

6.1 LENSING FORMALISM

6.1.1 Propagation of light bundles in weak gravitational fields

This section outlines the derivation of the lensing equations within the context of GR, assuming weak gravitational fields described by a Newtonian potential $\Phi \ll c^2$ and slowly moving gravitational lenses $v \ll c$ with respect to the cosmological flow. These basic assumptions are very well verified in the context of lensing by the large scale structure of the Universe and will always be verified within the context of this thesis.

Under GR, light propagates along null-geodesics of the space-time metric which are influenced by the spatial curvature (see Section 2.2.1) . If the curvature exhibits local variation then light-rays within a light bundle will be affected by slight differences in the local curvature which will result in a deformation of the bundle.

A thorough derivation of the lens equation describing this effect can be found in Seitz et al. (1994); Bartelmann (2010) and we outline here some of the main results leading to this equation. Let us

consider a fiducial light ray at the center of an infinitesimally thin light beam propagating, for now, in an arbitrary space-time. This light ray propagates along a null geodesic denoted γ_0^μ which can be parametrised by an affine parameter λ chosen so that $\lambda = 0$ at the observer and increases with the distance to the observer. Consider now the separation vector between the fiducial ray and a neighbouring light ray defined by $\xi^\mu(\lambda, \boldsymbol{\theta}) = \gamma^\mu(\lambda, \boldsymbol{\theta}) - \gamma_0^\mu(\lambda)$ where $\boldsymbol{\theta}$ is the angular separation between the two rays. To study the deformation of the light bundle, we are more particularly interested in the transverse components of this separation vector which can be described by the projection of the 4-vector on a 2-dimensional screen tangent to the sphere of directions seen by the observer and perpendicular to the light ray. If one introduces ξ_1 and ξ_2, the components of $\xi^\mu(\lambda, \boldsymbol{\theta})$ on an orthogonal basis of this screen, then the vector $\boldsymbol{\xi}(\lambda, \boldsymbol{\theta}) = (\xi_1, \xi_2)$ follows the following equation of geodesic deviation:

$$\frac{\mathrm{d}^2\boldsymbol{\xi}(\lambda, \boldsymbol{\theta})}{\mathrm{d}\lambda^2} = \mathcal{T}(\lambda)\boldsymbol{\xi}(\lambda, \boldsymbol{\theta}) \,, \tag{6.1}$$

which describes the variation of the separation vector, and thus the deformation of the light bundle, with respect to the affine parameter λ in terms of the *optical tidal matrix* $\mathcal{T}(\lambda)$. This matrix can be expressed in terms of the Riemann curvature tensor and encodes the effect of the local space-time curvature on the propagation of light. Furthermore, based on the definition of λ, the separation vector follows the initial conditions $\boldsymbol{\xi}(0) = 0$ and $\frac{\mathrm{d}\boldsymbol{\xi}}{\mathrm{d}\lambda} = \boldsymbol{\theta}$.

The question is now how to relate the tidal matrix $\mathcal{T}$, which depends on the curvature, to the matter content of the Universe. When considering local perturbations to the homogeneous Universe under a weak, quasi-static, Newtonian gravitational potential Φ, the tidal matrix can be explicitly computed by considering the following metric:

$$\mathrm{d}s^2 = a^2(\eta)\left(-\left(1 + \frac{2\Phi}{c^2}\right)\mathrm{d}\eta^2 + \left(1 - \frac{2\Phi}{c^2}\right)\left(\mathrm{d}r^2 + f_K^2(r)\,\mathrm{d}\Omega^2\right)\right) \,. \tag{6.2}$$

This metric is the superposition of the FLRW metric describing the isotropic and homogeneous Universe and a comoving Newtonian metric characterising local density fluctuations of scales much smaller than the Hubble length. These two contributions can be worked out independently and, as a result, for weak perturbations, the tidal matrix can be expressed as the sum of two terms:

$$\mathcal{T} = \mathcal{T}_{\mathrm{bg}} + \mathcal{T}_{\mathrm{cl}} \,, \tag{6.3}$$

where $\mathcal{T}_{\mathrm{bg}}$ captures the lensing effect caused by the smooth homogeneous Universe while $\mathcal{T}_{\mathrm{cl}}$ describes the tidal effects resulting from the inhomogeneities of the density distribution. As derived in Seitz et al. (1994), the tidal matrix takes the form:

$$\mathcal{T}_{ij}(\lambda) = \underbrace{-\frac{4\pi G}{c^2}\rho_0(1+z)^5\,\delta_{ij}}_{\text{Background contribution}} \underbrace{-\frac{(1+z)^2}{c^2}\left(2\partial_i\partial_j\Phi + \delta_{ij}\,\partial_{ij}^2\Phi\right)}_{\text{Clump contribution}} \,, \tag{6.4}$$

where $\partial_i = \frac{\partial}{\partial\xi_i}$. The background term remains proportional to the identity matrix which means that in a homogeneous isotropic Universe, light bundles only undergo an isotropic scaling as a function of distance which corresponds to the angular diameter distance. On the contrary, the clump contribution has non vanishing off-diagonal terms which will induce anisotropic deformations of the light bundle and will be responsible for the shearing effect that we will introduce latter.

Injecting the expression of the optical tidal matrix in the equation of geodesic deviations leads to the following equation describing the evolution of the comoving separation vector $\boldsymbol{x} = \frac{1}{a}\boldsymbol{\xi}$ with respect to the comoving distance :

$$\frac{\mathrm{d}^2\boldsymbol{x}}{\mathrm{d}r^2} + K\boldsymbol{x} = -\frac{2}{c^2}\left(\nabla_\perp\Phi(\boldsymbol{x}, r) - \nabla_\perp\Phi(\boldsymbol{0}, r)\right) \ , \tag{6.5}$$

where $\nabla_\perp = (\frac{\partial}{\partial x_1}, \frac{\partial}{\partial x_2})$ is the comoving gradient in the 2D screen transverse to the fiducial light ray and K is the spatial curvature. This equation can be solved by considering the following boundary conditions:

$$\boldsymbol{x}(0) = 0 \quad ; \quad \frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}r}(0) = \boldsymbol{\theta} \ , \tag{6.6}$$

where $\boldsymbol{\theta}$ is the apparent separation angle between the fiducial ray and its neighbour at the position of the observer. In this case, the solution of Equation (6.5) is given by:

$$\boldsymbol{x}(r) = f_K(r)\boldsymbol{\theta} - \frac{2}{c^2}\int_0^r \mathrm{d}r' f_K(r - r')\left(\nabla_\perp\Phi(\boldsymbol{x}(r'), r') - \nabla_\perp\Phi(\boldsymbol{0}, r')\right) \ , \tag{6.7}$$

In the absence of density fluctuations, the second term vanishes and this equation reduces to $\boldsymbol{x}(r) = f_K(r)\boldsymbol{\theta}$ which corresponds to the angular diameter distance. Note that in this expression, the Newtonian potential Φ is formally integrated along the perturbed light path $\boldsymbol{x}(r')$. However, this expression can be simplified under the assumption of small deviations by integrating the potential along the unperturbed light path $f_K(r')\boldsymbol{\theta}$ instead, leading to the so called *Born approximation*. This approximation is generally very well justified by numerical simulations and will be assumed throughout this work. Under the Born approximation, the relative deflections of the fiducial ray and its close neighbour are negligible when computing the difference of the transverse potential gradients, so that $\nabla_\perp\Phi(f_K(r')\boldsymbol{\theta}, r') - \nabla_\perp\Phi(\boldsymbol{0}, r') \simeq \nabla_\perp\left(\Phi(f_K(r')\boldsymbol{\theta}, r') - \Phi(\boldsymbol{0}, r')\right)$ where the potential Φ can be redefined for convenience to $\Phi(f_K(r')\boldsymbol{\theta}, r') - \Phi(\boldsymbol{0}, r')$. Therefore, the previous equation can be simplified as:

$$\boldsymbol{x}(r) = f_K(r)\boldsymbol{\theta} - \frac{2}{c^2}\int_0^r \mathrm{d}r' f_K(r - r')\nabla_\perp\Phi(f_K(r')\boldsymbol{\theta}, r') \ . \tag{6.8}$$

We can now derive the lens equation for a source at distance r_s by introducing the angular separation vector $\boldsymbol{\beta} = \boldsymbol{x}/f_K(r_s)$ defining the unlensed angular position in the source plane. In the absence of density fluctuations, the observed angular separation $\boldsymbol{\theta}$ would correspond to $\boldsymbol{\beta}$. However, in presence of lensing by the large scale structure these two quantities are linked by the following lens equation:

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \boldsymbol{\alpha}(\boldsymbol{\theta}) \ , \tag{6.9}$$

where the deflection angle $\boldsymbol{\alpha}$ is defined as:

$$\boldsymbol{\alpha} = \frac{2}{c^2}\int_0^{r_s} \mathrm{d}r' \frac{f_K(r_s - r')}{f_K(r_s)}\nabla_\perp\Phi(f_K(r')\boldsymbol{\theta}, r') \ . \tag{6.10}$$

In this expression, the gradient of the Newtonian potential is taken with respect of the transverse comoving coordinates in the screen perpendicular to the propagation of the light beam. However, it will be much more convenient for the following to recast this equation in terms of the derivatives with respect to the angular coordinates $\boldsymbol{\theta}$. This leads to the definition of the *effective lensing potential* ψ:

$$\psi(\boldsymbol{\theta}) = \frac{2}{c^2}\int_0^{r_s} \mathrm{d}r' \frac{f_K(r_s - r')}{f_K(r_s)f_K(r')}\Phi(f_K(r')\boldsymbol{\theta}, r') \ . \tag{6.11}$$

With this definition, the deflection angle can simply be expressed as the gradient of the effective lensing potential:

$$\boldsymbol{\alpha} = \nabla\psi \ , \tag{6.12}$$

where the gradient is now taken in terms of the angular coordinates $\boldsymbol{\theta}$. Note however that this expression is only valid as long as the curvature of the sky can be neglected i.e. within the *flat-sky approximation*. We will generally consider this approximation to hold for fields of view smaller than 10×10 degrees.

6.1.2 Weak lensing quantities

Shear and convergence

To describe the lensing effect in terms of a set of local quantities, Equation (6.9) can be approximated using a Taylor expansion to first order which holds to very good accuracy when the lensing is small. The lensing equation becomes:

$$\beta_i \simeq A_{ij}\theta_j \ , \tag{6.13}$$

where the Jacobian $\mathbf{A} = \dfrac{\partial \boldsymbol{\beta}}{\partial \boldsymbol{\theta}}$ is known as the amplification matrix, defined in terms of the derivatives of the lensing potential as:

$$A_{ij}(\boldsymbol{\theta}) = \frac{\partial \beta_i}{\partial \theta_j} = \delta_{ij} - \partial_i\partial_j\psi(\boldsymbol{\theta}) \ . \tag{6.14}$$

The amplification matrix can be separated into a trace and trace-free part:

$$\mathbf{A} = \begin{pmatrix} 1-\kappa & 0 \\ 0 & 1-\kappa \end{pmatrix} - \begin{pmatrix} \gamma_1 & \gamma_2 \\ \gamma_2 & -\gamma_1 \end{pmatrix} \ , \tag{6.15}$$

which is parametrised in terms of the complex *shear* $\gamma = \gamma_1 + i\gamma_2 = |\gamma|e^{2i\phi}$ and the *convergence* κ. The convergence only contributes to the diagonal of the matrix and leads to isotropic deformation of sources. On the other hand, the shear has non-vanishing off-diagonal terms which will cause anisotropic shearing of the source images. Both effects are illustrated on Figure 6.3. As we will see in the next section, it will also be interesting to factorise the term $1-\kappa$ in the amplification matrix:

$$\mathbf{A} = (1-\kappa)\left(\mathrm{Id} - \begin{pmatrix} g_1 & g_2 \\ g_2 & -g_1 \end{pmatrix}\right) \ , \tag{6.16}$$

where we have introduced the *reduced shear* $g = \dfrac{\gamma}{1-\kappa}$. As we will see, shape measurement methods used to estimate the weak lensing effect are only sensitive to the reduced shear g. In the weak lensing regime where $\kappa \ll 1$, the reduced shear can be approximated to γ but this approximation no longer holds when considering the lensing effect of massive galaxy clusters.

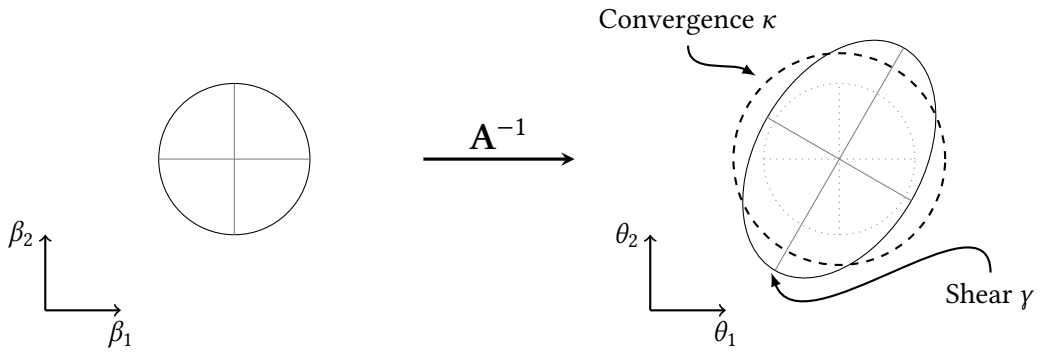

Figure 6.3: Deformation of a circular source under the effect of the amplification matrix $\mathbf{A}$. The effects of the anisotropic shear and isotropic convergence are separated.

Convergence and shear can be explicitly expressed in terms of the derivatives of the lensing potential according to:

$$\kappa = \frac{1}{2}\left(\partial_1^2\psi + \partial_2^2\psi\right) = \frac{1}{2}\Delta\psi \tag{6.17}$$

$$\gamma_1 = \frac{1}{2}\left(\partial_1^2\psi - \partial_2^2\psi\right) \tag{6.18}$$

$$\gamma_2 = \partial_1\partial_2\,\psi\ . \tag{6.19}$$

These differential equations linking shear, convergence and lensing potential will be at the heart of the weak lensing mapping techniques presented in Section 6.3.

The convergence κ will be of particular interest to us. It can indeed be directly interpreted as a line of sight projection of the matter density distribution between the observer and the source. According to Equation (6.17), the convergence can be interpreted as a source term of the 2D Poisson equation:

$$\kappa = \frac{1}{2}\Delta\psi\ , \tag{6.20}$$

$$= \frac{1}{c^2}\int_0^r \frac{f_K(r-r_s)}{f_K(r)f_K(r')}\Delta\Phi(f_k(r')\boldsymbol{\theta}, r')\ . \tag{6.21}$$

The Laplacian in this expression is taken with respect to the angular coordinates. This 2D angular Laplacian can be turned into a 3D Laplacian in comoving coordinates by adding an additional term $\frac{\partial^2\Phi}{\partial x_3^2}$ to the integrand using the argument that this term is expected to vanish under the line of sight integral under homogeneity arguments. Applying the 3D Poisson equation to the Newtonian potential then leads to the following expression of the convergence as a function of the 3D density contrast $\delta = (\rho - \bar{\rho})/\bar{\rho}$:

$$\kappa(\boldsymbol{\theta}, r) = \frac{3H_0^2\Omega_m}{2c^2}\int_0^r \mathrm{d}r' \frac{f_K(r')f_K(r-r')}{f_K(r)}\frac{\delta(f_K(r')\boldsymbol{\theta}, r')}{a(r')}\ . \tag{6.22}$$

Alternatively, this expression can also be written directly in terms of the 3D density ρ:

$$\kappa(\boldsymbol{\theta}, r) = \frac{4\pi G}{c^2}\int_0^r \mathrm{d}r' \frac{f_K(r')f_K(r-r')}{f_K(r)}\rho(f_K(r')\boldsymbol{\theta}, r')\ . \tag{6.23}$$

Magnification

Another consequence of the lensing effect is the modification of the fluxes of lensed sources. Indeed, due to shape distortions, the apparent solid angle of a lensed sources is altered while an intrinsic property of lensing ensures that the surface brightness of the source remains constant. As a result, the fluxes of lensed sources scales with their solid angle.

Conservation of the surface brightness is a consequence of the Liouville theorem which states that the volume in phase space occupied by a swarm of particles is invariant as they travel along their geodesic world lines if no absorption or emission takes place Misner et al. (1973). Consider the flux of photons emitted by a source characterised by a phase-space density $f(\boldsymbol{x}, \boldsymbol{p})$ defined as:

$$f(\boldsymbol{x}, \boldsymbol{p}, t) = \frac{\mathrm{d}N}{\mathrm{d}^3\boldsymbol{x}\mathrm{d}^3\boldsymbol{p}} \,. \tag{6.24}$$

Following Mollerach and Roulet (2002), this expression can be transformed, considering that the beam energy is $\mathrm{d}E = E_\gamma \mathrm{d}N$ where $E_\gamma = cp$ is the energy of individual photons, while $\mathrm{d}\boldsymbol{p} = p^2\mathrm{d}p\mathrm{d}\Omega$ and $\mathrm{d}\boldsymbol{x} = c\mathrm{d}A\mathrm{d}t$ with Ω the solid angle and A is the area perpendicular to the direction of propagation. Then phase space density becomes:

$$f(\boldsymbol{x}, \boldsymbol{p}, t) = \frac{\mathrm{d}E}{E_\gamma c\ \mathrm{d}A\ \mathrm{d}t\ p^2\mathrm{d}p\ \mathrm{d}\Omega} = \frac{\mathrm{d}E}{hcp^3\ \mathrm{d}\nu\ \mathrm{d}\Omega\ \mathrm{d}A\ \mathrm{d}t} = \frac{I(\nu)}{hcp^3} \,, \tag{6.25}$$

where $I\nu$ is the surface brightness of the source. Therefore, the Liouville theorem implies that the quantity $\frac{I(\nu)}{hcp^3}$ is constant along the trajectory of the beam in a curved space-time and therefore is not affected by the presence of gravitational lenses along its path.

The *magnification* μ can then be expressed as the ratio of lensed and unlensed fluxes which reduces to the ratio of solid angles as the surface brightness is preserved by lensing:

$$\mu = \frac{I\ \mathrm{d}\Omega}{I_0\ \mathrm{d}\Omega_0} = \frac{\mathrm{d}\Omega}{\mathrm{d}\Omega_0} \,, \tag{6.26}$$

where I_0 and $\mathrm{d}\Omega_0$ are the unlensed surface brightness and solid angles. Hence, the magnification reduces to a purely geometrical effect which can be expressed in terms of the amplification matrix $\mathbf{A}$ as:

$$\mu = |\det(\mathbf{A})|^{-1} = \frac{1}{(1-\kappa)^2 - \gamma_1^2 - \gamma_2^2} \,. \tag{6.27}$$

Note that this expression can lead to singularities where the determinant of the Jacobian vanishes. The points verifying $\det(\mathbf{A}) = 0$ are called critical points and form critical curves. Their images by the coordinate transform are known as *caustics*. Although the magnification of sources crossing the critical curves is extreme, it remains finite despite the infinite magnification predicted by this formalism which no longer holds in these situations.

Gravitational Flexion

Convergence and shear have been derived from a first order Taylor expansion of the lens equation. However, when the lensing effect becomes strong enough, higher order terms can be included in the expansion to accurately describe the distortions. This leads to the gravitational flexion Goldberg and Bacon (2005) formalism in which the lens equation is expanded to second order:

$$\beta_i \simeq A_{ij}\theta_j + \frac{1}{2}D_{ijk}\theta_j\theta_k \,, \tag{6.28}$$

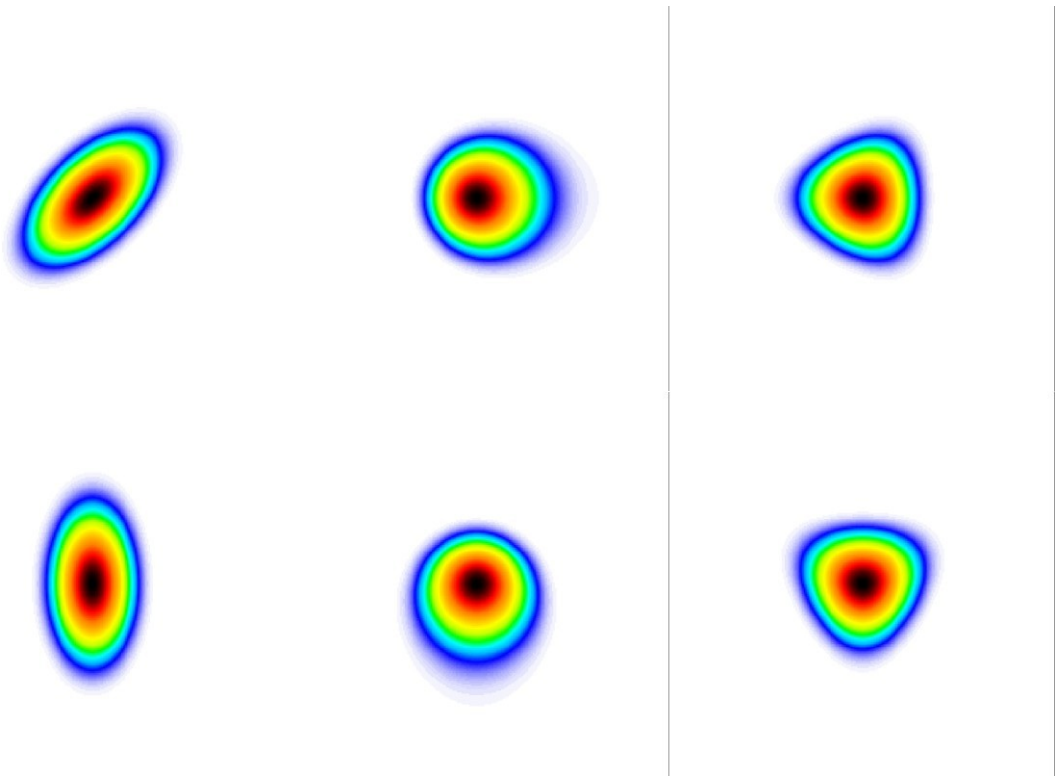

Figure 6.4: Distorsion of a circular image under the effects of shear γ (left), first flexion $\mathcal{F}$ (middle) and second flexion $\mathcal{G}$ (right). Credit: Figure from Bartelmann (2010).

where the third order lensing tensor D_{ijk} can be derived from $D_{ijk} = \partial_k A_{ij}$. This additional term gives rise to third order derivatives of the lensing potential which can be summarised as a spin-1 field $\mathcal{F} = \nabla\kappa$ and a spin-3 field $\mathcal{G} = \nabla\gamma$ respectively called first and second flexion. The tensor D_{ijk} can be separated into two components $D_{ijk} = \mathcal{F}_{ijk} + \mathcal{G}_{ijk}$ which can each be expressed in terms of the first and second flexion:

$$\mathcal{F}_{ij1} = -\frac{1}{2}\begin{pmatrix} 3\mathcal{F}_1 & \mathcal{F}_2 \\ \mathcal{F}_2 & \mathcal{F}_1 \end{pmatrix} \quad ; \quad \mathcal{F}_{ij2} = -\frac{1}{2}\begin{pmatrix} \mathcal{F}_2 & \mathcal{F}_1 \\ \mathcal{F}_1 & 3\mathcal{F}_2 \end{pmatrix}, \tag{6.29}$$

and

$$\mathcal{G}_{ij1} = -\frac{1}{2}\begin{pmatrix} \mathcal{G}_1 & \mathcal{G}_2 \\ \mathcal{G}_2 & -\mathcal{G}_1 \end{pmatrix} \quad ; \quad \mathcal{G}_{ij2} = -\frac{1}{2}\begin{pmatrix} \mathcal{G}_2 & -\mathcal{G}_1 \\ -\mathcal{G}_1 & -\mathcal{G}_2 \end{pmatrix}. \tag{6.30}$$

Being due to derivatives of the shear and convergence, flexion only becomes significant and mesurable when κ and γ exhibit significant variations across the length of a source, typically a galaxy. The effects of both flexion components are illustrated on Figure 6.4.

Similarly to the case of the reduced shear, these two flexion components are not directly observable. Instead, only the reduced flexion fields are measured from galaxy images:

$$F = \frac{\mathcal{F}}{1-\kappa} \quad ; \quad G = \frac{\mathcal{G}}{1-\kappa}. \tag{6.31}$$

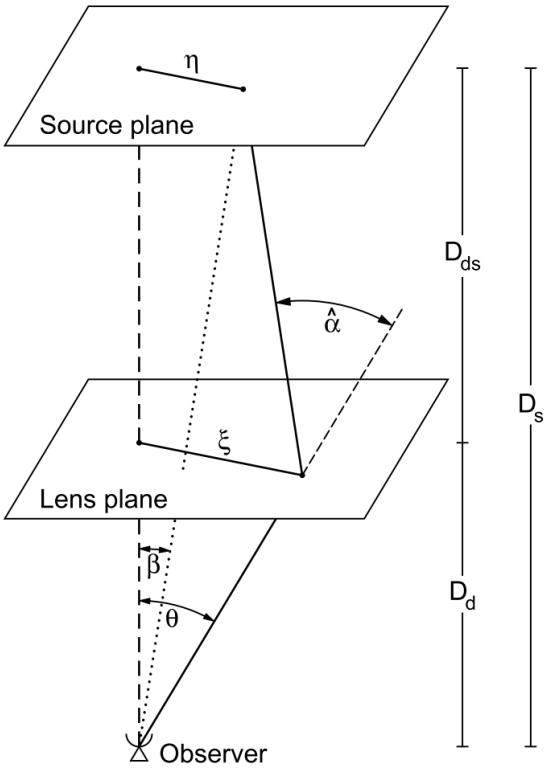

Figure 6.5: Illustration of the thin lens approximation. Credit: Figure from Bartelmann and Schneider (2001).

6.1.3 Thin lens approximation

The lensing equations derived in Section 6.1.1 take into account the full 3D distribution of matter along the line of sight. However, in practice, it is generally unlikely to find several significant structures along the same line of sight. Yang et al. (2011) find on numerical simulation that the vast majority of high convergence peaks are due to a single lens along the line of sight. In such cases, it is much more convenient to model the matter distribution as a single thin lens (the radial extension of structures such as galaxy clusters is always small compared to cosmological distances) and to neglect the rest of the density fluctuations along the line of sight. For strong lenses, this approximation generally holds at the percent level (Schneider, 2014). Considering a single thin lens along the line of sight, Equation (6.23) can be simplified in the following way:

$$\kappa(\theta_i) = \frac{4\pi G}{c^2} \int_0^{r_s} \mathrm{d}r' \frac{f_K(r') f_K(r_s - r')}{f_K(r_s)} \rho(f_K(r')\theta_i, r') \,, \tag{6.32}$$

$$= \frac{4\pi G}{c^2} \frac{f_K(r_l) f_K(r_s - r_l)}{f_K(r_s)} \int \mathrm{d}r' \rho(f_K(r')\theta_i, r') \,, \tag{6.33}$$

where r_l is the distance to the lens and r_s is the distance to the source. This expression can be interpreted as the ratio between two quantities:

$$\kappa(\theta_i) = \frac{\Sigma(\theta_i)}{\Sigma_{crit}} \,, \tag{6.34}$$

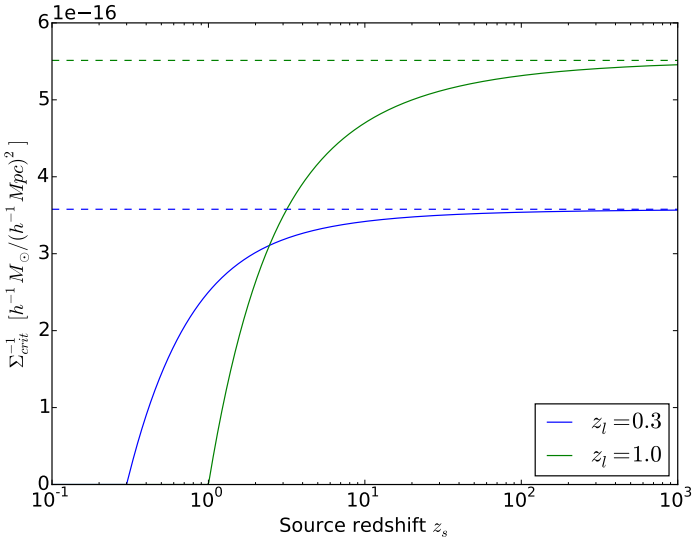

Figure 6.6: Inverse critical surface mass density as a function of source redshift for two lenses at redshifts $z_l = 0.3$ and $z_l = 1.0$. The dashed lines correspond to the asymptotes for $z_s \to \infty$.

where $\Sigma(\theta_i)$ is the surface mass density of the lens, corresponding to the line of sight integration of the 3D mass density ρ:

$$\Sigma(\theta_i) = \int \mathrm{d}r \rho(f_K(r)\theta_i, r) \,, \tag{6.35}$$

and Σ_{crit} is known as the critical surface mass density defined as:

$$\Sigma_{crit}^{-1} = \frac{4\pi G}{c^2} \frac{f_K(r_l) f_K(r_s - r_l)}{f_K(r_s)} \,, \tag{6.36}$$

$$= \frac{4\pi G}{c^2} \frac{D_L D_{LS}}{D_S} \,. \tag{6.37}$$

Therefore, for a given surface mass density $\Sigma(\theta_i)$, the intensity of the lensing effect only depends on the ratio of distances between the observer, the lens and the source, as illustrated on Figure 6.5.

Up to now, we have described the lensing effect for sources located on a single redshift plane. However, as can be seen in Equation (6.34), the intensity of the lensing effect is directly proportional to the ratio of distances $\frac{D_{LS}}{D_S}$ between the lens and the source. As a result, for a given lens, images of galaxies at different redshifts will experience different convergence $\kappa(z)$ and shear $\gamma(z)$ as a function of their individual redshifts z. This evolution in redshift is determined by the critical surface mass density which can be expressed as a function of source redshift as:

$$\Sigma_{crit}^{-1}(z) = \frac{4\pi G}{c^2} \frac{f_K(r_l) f_K(r(z) - r_l)}{f_K(r(z))} \,. \tag{6.38}$$

As the comoving distance convergences to a finite value when the redshift tends to infinity, so does $\Sigma_{crit}(z)$. The behavior of the critical mass density as a function of source redshift is illustrated on Figure 6.6. As the critical surface mass density converges to a finite value $\Sigma_{crit}(z) \xrightarrow[z\to\infty]{} \Sigma_{crit}^{\infty}$, its

limit at infinity can be used as a reference to define a convergence field κ_∞ independent of the redshifts of individual sources:

$$\kappa_\infty(\boldsymbol{\theta}) = \frac{\Sigma(\boldsymbol{\theta})}{\Sigma^\infty_{crit}} \,. \tag{6.39}$$

Given this convergence for sources at infinite redshift, the convergence for a source at finite redshift $z_s > z_l$ can be expressed as:

$$\kappa(\boldsymbol{\theta}, z_s) = \frac{\Sigma(\boldsymbol{\theta})}{\Sigma_{crit}(z_s)} \,, \tag{6.40}$$

$$= \frac{\Sigma^\infty_{crit}}{\Sigma_{crit}(z_s)} \frac{\Sigma(\boldsymbol{\theta})}{\Sigma^\infty_{crit}} \,, \tag{6.41}$$

$$= Z(z_s)\, \kappa_\infty(\boldsymbol{\theta}) \,, \tag{6.42}$$

where Z is a cosmological weight function defined for a lens at redshift z_l as:

$$Z(z_s) = \frac{\Sigma^\infty_{crit}}{\Sigma_{crit}(z_s)} H(z_s - z_l) \,, \tag{6.43}$$

$$= \lim_{u\to\infty} \frac{f_K(r(z_s) - r(z_l))\; f_K(u)}{f_K(r(u) - r(z_l))\; f_K(z_s)} H(z_s - z_l) \,, \tag{6.44}$$

where H is the Heaviside step function which only accounts for the fact that sources located at lower redshift than the lens are not lensed. The following lensing quantities depend linearly on the convergence and can therefore be rescaled in the same way:

$$\gamma(\boldsymbol{\theta}, z_s) = Z(z_s)\gamma_\infty(\boldsymbol{\theta}) \quad ; \quad \mathcal{F}(\boldsymbol{\theta}, z_s) = Z(z_s)\mathcal{F}_\infty(\boldsymbol{\theta}) \quad ; \quad \mathcal{G}(\boldsymbol{\theta}, z_s) = Z(z_s)\mathcal{G}_\infty(\boldsymbol{\theta}) \,. \tag{6.45}$$

6.1.4 Mass-sheet degeneracy

One of the most notorious issue in weak lensing mass mapping is the so called mass-sheet degeneracy. This degeneracy is due to the fact that the shear is left invariant by the addition of a constant mass-sheet to the lens surface density. As a result, the observed reduced shear is invariant under the following *λ-transformation*:

$$\kappa' = \lambda\kappa + (1 - \lambda) \,, \tag{6.46}$$

Indeed, if g' is the reduced shear generated by κ' then:

$$g' = \frac{\lambda\gamma}{1 - \lambda\kappa - 1 + \lambda} = \frac{\lambda\gamma}{\lambda - \lambda\kappa} = g \,, \tag{6.47}$$

By the same mechanism, the reduced flexions F and G are also invariant under the same λ-transformation.

This issue is particularly problematic for measuring the mass of galaxy clusters from weak-lensing. Indeed, these measurements are typically performed on small fields where it cannot simply be assumed that the convergence goes to zero outside the field without biasing the measurement (Bartelmann, 1995). Although it is true that from shear alone the mass-sheet degeneracy can not be lifted, it can nonetheless be mitigated when including additional information about the relative distances between the lens and the different sources. Indeed, as was pointed out in Bradac et al. (2004), knowledge of individual photometric redshifts of background galaxies is enough to constrain the mass-sheet for strong enough lenses.

Let us consider $\gamma_\infty(\theta)$ and $\kappa_\infty(\theta)$ the shear and convergence of a given lens, for sources at infinite redshift. The actual shear and convergence applied to a galaxy at a specific redshift z_s can be expressed as $\kappa(\theta, z_s) = Z(z_s)\kappa_\infty(\theta)$ and $\gamma(\theta, z_s) = Z(z_s)\gamma_\infty(\theta)$ where Z is a cosmological weight introduced in Equation (6.43). If we now consider the reduced shear measured on the galaxy as a function of κ_∞ and γ_∞, it takes the form:

$$g(\theta, z_s) = \frac{Z(z_s)\gamma_\infty(\theta)}{1 - Z(z_s)\kappa_\infty(\theta)} \,. \tag{6.48}$$

This expression alone is just a rewriting of the measured reduced shear, which makes explicit the dependency on the redshift of the source. Therefore, the convergence is still subject to the mass-sheet degeneracy which can now be made explicit in term of z_s:

$$\kappa' = \lambda\kappa + \frac{1-\lambda}{Z(z_s)} \,. \tag{6.49}$$

It is easily shown that such a transformation leaves the measured reduced shear invariant. However, as soon as the reduced shear is measured for galaxies at two different redshifts, the mass-sheet degeneracy is automatically lifted thanks to the different cosmological weights. As a result, simultaneously constraining the convergence to fit the reduced shear at different redshifts formally leads to $\lambda = 1$:

$$\begin{cases} \kappa &= \lambda\kappa_0 + \frac{1-\lambda}{Z} \\ \kappa &= \lambda'\kappa_0 + \frac{1-\lambda'}{Z'} \end{cases} , \tag{6.50}$$

$$\Rightarrow (\lambda - \lambda')\kappa_0 = \frac{Z'(1-\lambda) - Z(1-\lambda')}{ZZ'} \,, \tag{6.51}$$

$$\Rightarrow \lambda = \lambda' \text{ and } (Z' - Z)(1-\lambda) = 0 \,, \tag{6.52}$$

Therefore, if the redshifts of the sources are different, $Z \neq Z'$, which leads to $\lambda = 1$.

Note that the degeneracy is only lifted if the convergence is strong enough. In the linear regime, the transformation $\kappa' = \kappa + \lambda$ has no dependency on the redshift of the source. Therefore, this mechanism is only efficient outside of the weak lensing regime. For the purpose of mapping galaxy clusters this is will generally be the case and using the redshift information of individual galaxies will contribute to mitigate the degeneracy and thus provide more reliable density reconstructions for the purpose of weighting clusters.

6.2 WEAK LENSING MEASUREMENT

The previous section introduced the gravitational lensing formalism which describes how the images of distant galaxies are distorted by the matter overdensities along the line of sight. Now the question is how to measure in practice these distortions on telescope images, especially when the distortions are weak. This section introduces the basic principles of shear measurements along with the main practical difficulties.

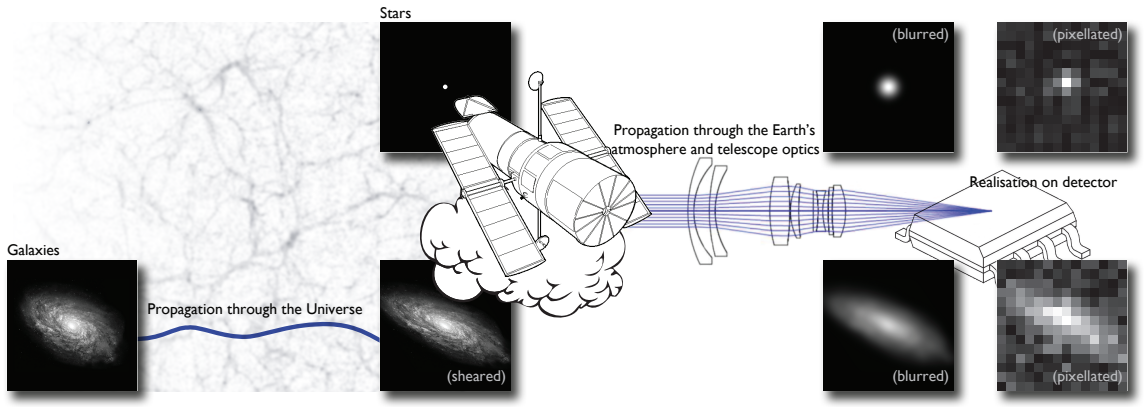

Figure 6.7: Measurement process of weak gravitational lensing, from the original galaxy image to the observed pixelated galaxy image. The lensed galaxy image is first blurred by the PSF of the instrument and then pixelised by the detector. To distinguish between the lensing effect and instrumental PSF, the latter can be measured by imaging stars which can be considered as point sources. Credit: Figure from Mandelbaum et al. (2014), adapted from Kitching et al. (2011).

6.2.1 Imaging process

Measurements of weak lensing in practice with sufficient accuracy to derive reliable cosmological constraints has proven to be a difficult task which requires excellent image quality and advanced image processing techniques. The weak lensing effect we are interested in measuring induces only very faint deformations of the images of background galaxies which are extremely difficult to disentangle from instrumental effects.

The different steps of a typical acquisition process are illustrated on Figure 6.7. The original galaxy image is first slightly deformed as the light propagates through the large scale structure of the Universe. In the case of space based instrument such as Euclid, the light of the distant galaxy then directly reaches the telescope without having to propagate through the atmosphere which is an important source of perturbations. Nevertheless, even without deformations due to atmospheric turbulence, because of fundamental limitations of any optical system, the image of the galaxy formed by the telescope on the focal plane is not perfect and will be slightly blurred. Finally this blurred image is acquired by a detector to produce the noisy and pixelated galaxy image actually measured. These additional effects corrupting the images of the observed galaxies need to be precisely taken into account in order to extract the lensing signal.

The response of the instrument can be characterised by its *Point Spread Function* (PSF) i.e. the impulse response of the entire imaging system from the optics to the detector. As the main effect of weak lensing will be to shear the image of galaxies, it is particularly crucial to account for any anisotropies in the PSF, which creates a spurious shearing orders of magnitudes stronger than the cosmological signal. For space-based telescopes, which are not hindered by the atmosphere, this PSF can be directly estimated using an optical model of the instrument (e.g Rhodes et al., 2007). Another approach consists in estimating the PSF by imaging a clean sample of stars of our own galaxy which can be considered as point sources. As illustrated on the top of Figure 6.7, the image of the star will undergo the same blurring effects and directly produce an image of the PSF. Another difficulty comes from the fact that the PSF is generally not stationary across the field and

is only sampled on a limited number of points using stars. Therefore it must be interpolated at the position of each observed galaxy to properly correct the measured shape.

6.2.2 Galaxy ellipticity as a shear estimator

As was mentioned in the previous section, gravitational shear causes an anisotropic stretching of the galaxy images which directly affects their apparent ellipticity. We now describe how the observed galaxy ellipticities can be used to build shear estimators.

In order to define this ellipticity for arbitrary galaxy images, we first need to introduce descriptors of what is actually measured by the detector: the surface brightness $I(\boldsymbol{\theta})$. One way of characterising the image of an isolated galaxy is to consider the moments of the $I(\boldsymbol{\theta})$, starting with the first brightness moment which defines the centroid of the galaxy:

$$\overline{\boldsymbol{\theta}} = \frac{\int \mathrm{d}\boldsymbol{\theta}\, I(\boldsymbol{\theta}) W(\boldsymbol{\theta}) \boldsymbol{\theta}}{\int \mathrm{d}\boldsymbol{\theta}\, I(\boldsymbol{\theta}) W(\boldsymbol{\theta})}\,, \tag{6.53}$$

where $W(\boldsymbol{\theta})$ is an arbitrary weight function which plays an important role in practice, to ensure the convergence of the integral and to control the impact of instrumental noise. Once the centroid of the image is identified, information about the shape of the galaxy can be captured by considering the second brightness moments $\{Q_{ij}\}_{i,j\in\{1,2\}}$ defined as:

$$Q_{ij} = \frac{\int \mathrm{d}\boldsymbol{\theta}\, I(\boldsymbol{\theta}) W(\boldsymbol{\theta}) (\theta_i - \overline{\theta}_i)(\theta_j - \overline{\theta}_j)}{\int \mathrm{d}\boldsymbol{\theta}\, I(\boldsymbol{\theta}) W(\boldsymbol{\theta})}\,. \tag{6.54}$$

From this tensor, one can build the following *complex ellipticity*:

$$\epsilon = \frac{Q_{11} - Q_{22} + 2iQ_{12}}{Q_{11} + Q_{22} + 2(Q_{11}Q_{22} - Q_{12}^2)^{1/2}}\,. \tag{6.55}$$

This quantity has the interesting property that one can directly relate the complex ellipticity of the unlensed source image $\epsilon^{(s)}$ to that of the observed lensed image ϵ in terms of the reduced shear (Seitz and Schneider, 1997):

$$\epsilon = \begin{cases} \dfrac{\epsilon^s + g}{1 + g^*\epsilon^s} & \text{when } |g| \leq 1 \\ \\ \dfrac{1 - g\epsilon^*}{\epsilon^* - g^*} & \text{when } |g| > 1 \end{cases}\,, \tag{6.56}$$

where ϵ^* designates the complex conjugate of the ellipticity. Although several definitions of the complex ellipticity have been proposed in the literature, this specific choice has the remarkable property that expectation value of the measured ϵ does not depend on the probability distribution of the source ellipticity as long as it is isotropic (see (Seitz and Schneider, 1997), Appendix A1):

$$< \epsilon >_{\epsilon_s} = \begin{cases} g & \text{when } |g| \leq 1 \\ \\ \dfrac{1}{g^*} & \text{when } |g| > 1 \end{cases}\,. \tag{6.57}$$

Therefore, for sub-critical lensing and assuming that the intrinsic orientations of the source galaxies are uniformly distributed, this definition of the ellipticity provides an unbiased estimator of the reduced shear.

In practice, to estimate this expectation value and therefore the reduced shear, it is necessary to average the ellipticities measured for a population of galaxies experiencing a similar shear i.e. in the same small patch of the sky:

$$\bar{\epsilon} = \sum_{k=0}^{N_{gal}} \epsilon_k \ . \tag{6.58}$$

As is generally assumed, if the intrinsic ellipticity of the galaxies is Gaussian distributed with zero mean and variance σ_ϵ^2, then this estimator can be modelled as:

$$\bar{\epsilon} = g + n \qquad \text{with} \qquad n \sim \mathcal{N}(0, \frac{\sigma_\epsilon^2}{N_{gal}}) \ , \tag{6.59}$$

where n is a Gaussian *shape noise*.

6.2.3 Shape measurement methods

Now that we have described how the ellipticity ϵ provides an estimator for the shear, the question is how to measure this ellipticity in practice on blurred, pixelated and noisy images acquired through the steps described at the beginning of this section. There are mainly two families of methods: direct ellipticity measurement using *moments* of the image or ellipticity estimation through *forward-fitting* methods.

In the family of ellipticity measurements based on moments, the most notable example is the *KSB* method (Kaiser et al., 1995) and its improved *KSB+* successor (Luppino and Kaiser, 1997; Hoekstra et al., 1998). This method relies on the weighted second moments of the light distribution and approximates the effect of the PSF by a linear operation, assuming that both shear and PSF anisotropies are small.

The second family of forward fitting approaches is more recent with the first method proposed by Kuijken (1999). Contrary to the previous methods which use moments of the light distribution, in this class of methods a model for the unlensed galaxy is used and then convolved with the PSF and sheared to fit the observed image, yielding an estimate of the shear. A notable example of this class of methods is *lensfit* (Miller et al., 2007; Kitching et al., 2008a) which provides a fully Bayesian estimation of the shear. Another example is the maximum likelihood *gfit* (Gentile et al., 2012) method.

Despite the large number of methods developed so far, accurate shear measurements meeting the requirements of the future Stage-IV dark energy surveys are still difficult to obtain. Staring with the Shear TEsting Programme (STEP) program (Heymans et al., 2006), a series of challenges have been organised to test these algorithms, investigate their limitations and determine the methodoligical improvements necessary to reach the requirements of future surveys. The third GRavitational lEnsing Accuracy Testing (GREAT3) (Mandelbaum et al., 2014) is the latest of these challenges and provides the most recent assessment of the performance of shear measurement methods. The general conclusion of this challenge (Mandelbaum et al., 2015) is that fitting methods generally outperform moments methods. In particular, GREAT3 demonstrated that even when using very

simple galaxy models (either a single or a combination of two Sérsic profiles), fitting methods were still effective when including realistic galaxy morphologies in the simulations (GREAT3 galaxies were based on Hubble Space Telescope (HST) imaging).

6.3 2D MASS MAPPING

6.3.1 Kaiser-Squires inversion

It was originally realised by Kaiser and Squires (Kaiser and Squires, 1993) that weak lensing can be used to directly map the mass distribution. Indeed, weak lensing provides a direct estimate of the shear γ (in the weak regime) which can be related to the convergence κ through a set of differential equations established in Section 6.1.2. In the Fourier domain, these relations can be written as:

$$\tilde{\gamma}_1 = \frac{1}{2}\left(k_1^2 - k_2^2\right)\tilde{\psi} \quad ; \quad \tilde{\gamma}_2 = k_1 k_2 \tilde{\psi} \quad ; \quad \tilde{\kappa} = \frac{1}{2}(k_1^2 + k_2^2)\tilde{\psi} \,. \tag{6.60}$$

As can be seen, the convergence can directly be obtained from the lensing potential. In turn, the lensing potential can be independently estimated from both components of the shear. The idea of the Kaiser-Squires method is to combine this set of equations to build a minimum variance estimator of the convergence, which leads to

$$\hat{\tilde{\kappa}}_\gamma = \frac{k_1^2 - k_2^2}{k^2}\tilde{\gamma}_1 + \frac{2k_1 k_2}{k^2}\tilde{\gamma}_2 \,, \tag{6.61}$$

where $k^2 = k_1^2 + k_2^2$. An important point to notice is that this formula is not defined for $k_1 = k_2 = 0$, which means that the convergence can only be recovered up to a constant term. This is the famous *mass-sheet degeneracy*. Equivalently, this estimator can be expressed in the direct domain as a convolution product of the form:

$$\kappa(\boldsymbol{\theta}) - \kappa_0 = \frac{1}{\pi}\int \mathrm{d}\boldsymbol{\theta}' \mathcal{D}(\boldsymbol{\theta} - \boldsymbol{\theta}')\gamma(\boldsymbol{\theta}') \,, \tag{6.62}$$

where κ_0 is a constant offset due to the mass sheet degeneracy which cannot be determined from the shear alone and $\mathcal{D} = \frac{1}{(\boldsymbol{\theta}^*)^2}$ is just the Fourier transform of the kernel in Equation (6.61).

An important aspect of this estimator is its linearity, which allows us to derive the noise on the estimated convergence given the noise on the input shear. Assuming that both components of the shear are uncorrelated and have variance σ_ϵ^2, the variance of the estimator becomes:

$$< \hat{\tilde{\kappa}}^* \hat{\tilde{\kappa}} > = \frac{(k_1^2 - k_2^2)^2}{k^4}\sigma_\epsilon^2 + \frac{4k_1^2 k_2^2}{k^4}\sigma_\epsilon^2 \,, \tag{6.63}$$

$$= \sigma_\epsilon^2 \,. \tag{6.64}$$

Thus the Kaiser-Squires estimator has the same noise variance as the input shear. As was mentioned in Section 6.1.2 the noise on the shear estimates can be modelled as an additive Gaussian noise of variance $\frac{\sigma_\epsilon^2}{N_g}$ where N_g is the mean number of galaxies per angular bins. Therefore, the convergence maps built from a binned shear map will have a white Gaussian noise of the same variance $\frac{\sigma_\epsilon^2}{N_g}$.

The Kaiser-Squires estimator is defined in terms of the shear γ whereas in practice, only the reduced shear $g = \frac{\gamma}{1-\kappa}$ is accessible. This is not a particular issue when mapping the large scale dark matter distribution, as the lensing signal on large scale is well within the weak regime so that $g \simeq \gamma$. However, on smaller scales, for instance at the close vicinity of galaxy clusters, the convergence can become significant so that the reduced shear can no longer be directly used as an estimator of the shear. Unfortunately, the inversion problem becomes non-linear and can no longer be solved by a simple convolution. Nevertheless, as was noted in Seitz and Schneider (1995), the Kaiser-Squires estimator can still be used in an iterative scheme which applies a correction factor to the reduced shear based on an estimate $\bar{\kappa}$ of the convergence which is refined at each iteration. The estimator becomes:

$$\kappa(\boldsymbol{\theta}) - \kappa_0 = \frac{1}{\pi} \int \mathrm{d}\boldsymbol{\theta}' [1 - \bar{\kappa}(\boldsymbol{\theta}')] \mathcal{D}(\boldsymbol{\theta} - \boldsymbol{\theta}') g(\boldsymbol{\theta}') \,, \tag{6.65}$$

which is generally found to converge to a solution in a few iterations by replacing $\bar{\kappa}$ with the current estimate of the convergence. As was described in Section 6.1.2, the solution of this problem is still subject to the mass-sheet degeneracy. Contrary to the weak regime, this degeneracy is not a simple additive constant term, instead the convergence is recovered up to the λ-transformation introduced in the previous section:

$$\kappa' = \lambda\kappa + (1 - \lambda) \,. \tag{6.66}$$

The measured reduced shear g remains invariant under this transformation.

Although this inversion technique has some very desirable properties including linearity (in the weak regime), minimum variance and the ability to accommodate reduced shear, it has one major shortcoming, namely that is not well defined on bounded domains, even less so on domains with complex geometries and missing data which occurs in actual surveys. Indeed, the convolution in Equation (6.62) is performed over the entire $\mathbb{R}^2$ plane.

This has led to the development of complementary or alternative methods, the most important of which being the Seitz and Schneider inversion (Seitz and Schneider, 1996). This particular method aims at solving the inversion problem using local differential relations between the convergence and the derivatives of the shear and is theoretically well defined on bounded domains for noiseless data. However, in practice, the Seitz and Schneider method is not very robust to noise, unless the data is first smoothed by a large Gaussian, and it was found that in general, even on small field, the Kaiser-Squires inversion generally faired better than the Seitz-Schneider method(Schneider, 2006).

Another alternative to the simple Kaiser-Squires method has also been developed to mitigate the impact of the mask by Pires et al. (2009). This method is an application of sparse inpainting to the weak lensing inversion problem. The solution is required to be sparse in a DCT dictionary, which is shown to dramatically reduce the impact of the mask on second order statistics.

6.3.2 Extension to flexion reconstruction

Similarly to the Kaiser-Squires inversion for the shear, Bacon et al. (2006) proposed a minimum variance filter for estimating the convergence from first flexion $\mathcal{F}$. As the flexion is only the gradient of the convergence (see Section 6.1.2), these two quantities can be related in the Fourier domain as:

$$\tilde{\mathcal{F}_1} = -ik_1\tilde{\kappa} \quad ; \quad \tilde{\mathcal{F}_2} = -ik_2\tilde{\kappa} \,. \tag{6.67}$$

With the same approach as the Kaiser-Squires inversion, the two components of the first flexion can be combined in order to yield a minimum variance estimator of the convergence, which takes the simple expression:

$$\hat{\tilde{\kappa}}_F = \frac{ik_1\tilde{\mathcal{F}}_1 + ik_2\tilde{\mathcal{F}}_2}{k^2} \,. \tag{6.68}$$

Note that just as with the Kaiser-Squires inversion, this expression is only valid for $k_1 \neq 0$ and $k_2 \neq 0$, which means that the flexion reconstruction is still subject to the mass sheet degeneracy. Contrary to the previous Kaiser-Squires inversion, this estimator does not have a flat noise power spectrum. Indeed, assuming uncorrelated flexion measurements with intrinsic variance σ_F^2, the variance of this estimator is:

$$< \hat{\tilde{\kappa}}_F^* \hat{\tilde{\kappa}}_F > = \frac{k_1^2}{k^4}\sigma_F^2 + \frac{k_2^2}{k^4}\sigma_F^2 \,, \tag{6.69}$$

$$= \frac{1}{k^2}\sigma_F^2 \,. \tag{6.70}$$

This makes the reconstruction of mass maps from flexion alone using this estimator very problematic on large scales where the noise will always dominate the signal. This means that applying any low-pass smoothing actually reduces the signal to noise ratio of the reconstruction.

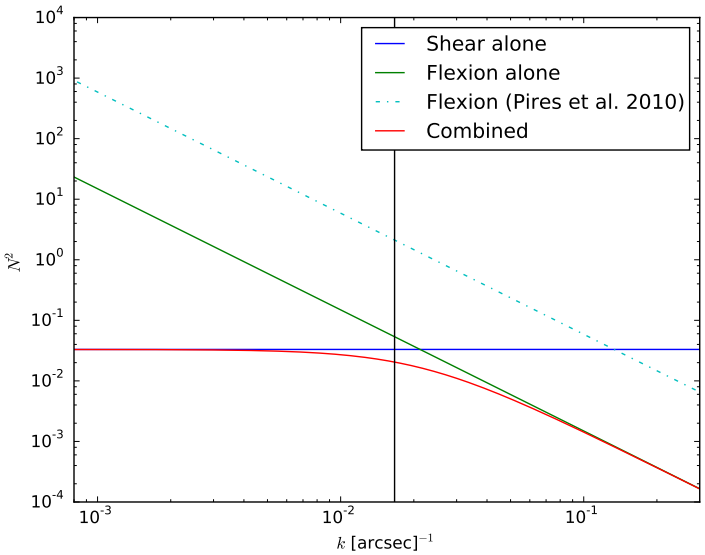

Figure 6.8: Noise power spectra of the minimum variance estimators for shear alone (blue line), flexion alone (green line) and shear and flexion combined (red line). The dashed cyan line shows the flexion noise level wrongly estimated in Pires and Amara (2010). The vertical line indicates the arcminute scale computed as $k = 2\pi/r$. These noise power spectra assume $\sigma_F = 0.04$ arcsec^{-1}, $\sigma_\gamma = 0.3$ and $n_g = 50$ gal/arcmin^2 with 14 arcsecs pixels, as in Pires and Amara (2010).

This has led Pires and Amara (2010) to conclude that flexion was unusable for mass mapping, claiming that the Signal to Noise Ratio (SNR) of the flexion reconstruction could only surpass that of the shear reconstruction at the scale of a few arcsecs. However, this conclusion was based on an erroneous interpretation of Fourier frequencies with respect to angular scales. Using the same setting as Pires and Amara (2010), we find the angular scale where flexion reaches the noise level

of the shear to be around 45 arcsecs. As a result, based solely on noise levels using these Fourier estimators, we argue that flexion can be used to help resolve sub-arcmins structures. Figure 6.8 shows the noise power spectra of a shear inversion alone (blue) and flexion inversion alone (green), the arc-minute scale is marked by the vertical line.

It is clear from Figure 6.8 that although flexion can help constrain small features, it is not competitive with respect to the shear on scales larger than 1 arcmin. To simultaneously benefit from both shear and flexion, Bacon et al. (2006) proposed a minimum variance filter combining both measurements. Although they present reconstructions using this combined filter, they do not explicitly provide its expression, which can easily be derived (see Appendix B) and takes the form:

$$\hat{\tilde{\kappa}}_{\gamma F} = \frac{1}{k^2 + \frac{\sigma_F^2}{\sigma_\gamma^2}} \left(ik_1 \tilde{\mathcal{F}}_1 + ik_2 \tilde{\mathcal{F}}_2 + \frac{\sigma_F^2}{\sigma_\gamma^2} \left(\frac{k_1^2 - k_2^2}{k^2} \tilde{\gamma}_1 + \frac{2k_1 k_2}{k^2} \tilde{\gamma}_2 \right) \right) . \tag{6.71}$$

The noise variance of this estimator is now:

$$< \hat{\tilde{\kappa}}^*_{\gamma F} \hat{\tilde{\kappa}}_{\gamma F} >= \frac{\sigma_F^2}{k^2 + \frac{\sigma_F^2}{\sigma_\gamma^2}} . \tag{6.72}$$

Figure 6.8 illustrates the noise power spectrum of this estimator (red line) in realistic conditions. We see that the combined estimator starts improving over the shear alone around the arcminute scale.

An example of reconstruction using this combined estimator is shown on Figure 6.9. The input convergence map is reconstructed from shear alone, flexion alone and combining shear and flexion. As expected, the shear alone reconstruction is very noisy, especially on small scales because of its flat noise power spectrum. On the contrary, the flexion alone reconstruction is noise dominated on scales larger than 1 arcmin. However, by combining both shear and flexion information, the noise is effectively suppressed on small scales and not amplified on large scales which makes the cluster and some of its substructure clearly identifiable without any additional filtering. Note however that this simple example only illustrates the noise properties of these estimators but is unrealistic in the sense that it does not consider the problem of missing data. A reconstruction strategy combining shear and flexion and properly accounting for missing data will be presented in Chapter 7.

6.3.3 Maximum likelihood methods

The direct inversion methods presented so far only amount to applying a filter to the measured shear and/or flexion. They are not regularised and another filtering step is usually required to beat down the noise and recover the signal. Furthermore, this framework does not allow the combination of additional constraints such as strong lensing or magnification.

Another approach originally suggested in Bartelmann et al. (1996) consists in defining a likelihood function which incorporates both a data fidelity term and a regularisation term to avoid over-fitting the data and thus control the level of noise. This class of methods can seamlessly incorporate all available information (shear, flexion, magnification, strong lensing, etc.) into a joint reconstruction. They have mainly be used in the context of strong lensing.

The general approach in non-parametric maximum likelihood methods is to define a discretised lensing potential ψ on a grid. The choice of reconstructing the lensing potential and not the convergence is due to the more direct relationships between the potential ψ and the observables (for

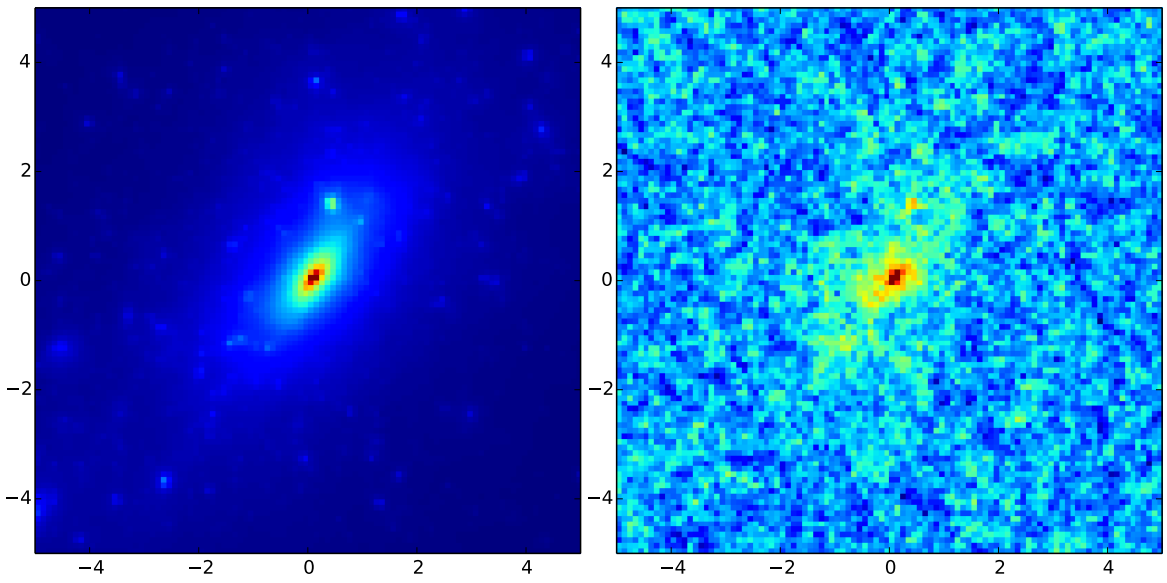

(a) Input convergence map

(b) Combined estimator $\hat{\kappa}_{\gamma F}$

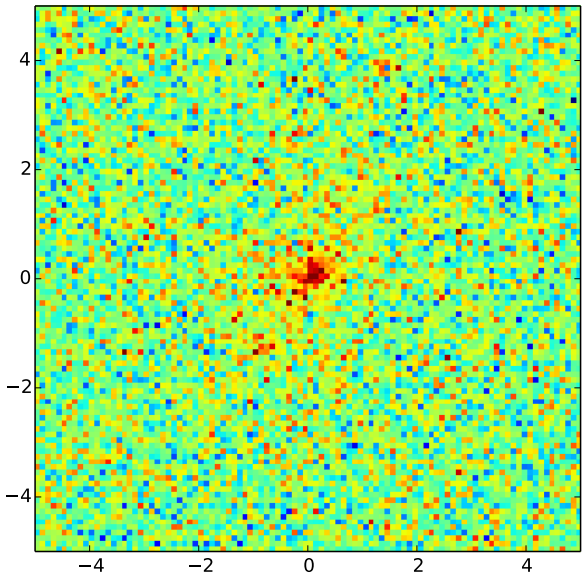

(c) Shear alone $\hat{\kappa}_{\gamma}$

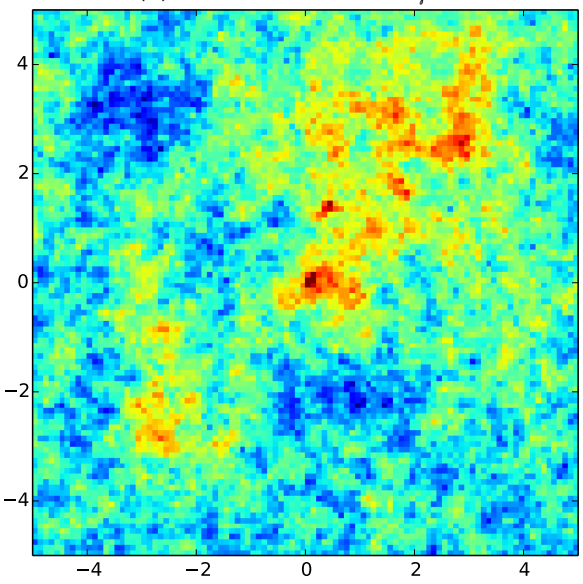

(d) Flexion alone $\hat{\kappa}_{F}$

Figure 6.9: Reconstruction of a small 10×10 arcmin^2 field containing a massive cluster using the shear alone Kaiser-Squires estimator (c), the flexion alone minimum variance estimator (d) and the combined minimum variance estimator (b). The noise level correspond to $\sigma_F = 0.04$ arcsec^{-1}, $\sigma_\gamma = 0.3$ and $n_g = 50$ gal/arcmin^2 with 6 arcsecs pixels.

instance computing the shear from the convergence involves an inverse Laplacian). The potential is fitted to the different constraints available by minimising a multi-component χ^2-function:

$$\chi^2(\psi) = \chi_1^2(\psi) + \chi_2^2(\psi) + \chi_3^2(\psi) + \dots \,, \tag{6.73}$$

where each χ_i^2 corresponds to a different observational constraint. For instance the χ^2 for weak lensing data can be expressed as:

$$\chi_g^2 = \sum_{i=1}^{N_{gal}} \frac{|\epsilon_i - g(\boldsymbol{\theta}_i, \psi)|^2}{\sigma_i^2} \,, \tag{6.74}$$

where $g(\boldsymbol{\theta}_i, \psi)$ is the reduced shear evaluated at the position of the galaxy of index i, given the lensing potential ψ.

Despite combining different constraints, the inversion problem remains ill-posed and minimising the data fidelity χ^2 without additional regularisation leads to an important over-fitting of the noise. To avoid over-fitting, a regularisation term R is added to the total χ^2:

$$\chi^2_{reg}(\psi) = \chi^2(\psi) + \lambda R(\psi) \,, \tag{6.75}$$

where λ tunes the strength of the regularisation. Different kinds of regularisation functions R have been used in the literature such as the maximum-entropy (Seitz et al., 1998) or the quadratic deviation from a model (Bradac et al., 2005).

Recovering the lensing potential ψ now amounts to minimising the regularised $\chi^2_{reg}(\psi)$. Solving this optimisation problem is challenging in practice as the χ^2 is usually a non-linear function of the potential ψ, which requires one to iteratively solve a linearisation of the problem. Furthermore, the computational cost of evaluating and/or differentiating the χ^2-function scales rapidly with the size and resolution of the field which limits the applicability of these methods to relatively small areas.

6.4 3D MASS MAPPING

3D weak lensing was first introduced in Taylor (2001) as an extension of 2D mass mapping. However, since then only a handful of additional methods have been proposed (Hu and Keeton, 2002; Simon et al., 2009; VanderPlas et al., 2011; Leonard et al., 2012) and despite the increasing availability of lensing and photometric redshift data, very few studies have actually applied 3D lensing techniques to lensing surveys. In this section, we review existing 3D lensing methods and highlight some of their drawbacks which can explain why 3D lensing approaches have not been more widespread until now.

In Chapter 8 we will introduce a new 3D lensing technique, based on sparse regularisation, which overcomes most of the limitations of the linear methods presented in this section.

6.4.1 The Taylor direct inversion method

The first 3D weak lensing mass-mapping method was proposed by Taylor (2001) and further developed in Bacon and Taylor (2003). This method aims at mapping the 3D gravitational potential $\boldsymbol{\Phi}$ from the tomographic lensing potential ψ. These two potentials are related through the following line of sight integration:

$$\psi(\boldsymbol{r}) = 2 \int_0^r \mathrm{d}r' \left(\frac{r - r'}{r r'} \right) \boldsymbol{\Phi}(\boldsymbol{r'}) \,, \tag{6.76}$$

assuming a flat Universe and where r is the radial comoving distance. If the lensing potential is known, the 3D gravitational potential can readily be recovered as (Taylor, 2001):

$$\Phi = \frac{1}{2}\partial_r r^2 \partial_r \psi \ . \tag{6.77}$$

The lensing potential is not directly observable but can be estimated from the measured tomographic shear through a Kaiser-Squires inversion:

$$\tilde{\psi} = 2\partial^{-4}\partial_i\partial_j\gamma_{ij} \ , \tag{6.78}$$

where ∂^{-2} is the inverse 2D Laplacian. This inversion is of course subject to the mass-sheet degeneracy, which means that the tomographic lensing potential can only be recovered up to an additive function of the comoving distance r: $\tilde{\psi}(r,\boldsymbol{\theta}) = \psi(r,\boldsymbol{\theta}) + f(r)$. This degeneracy can however be mitigated for large enough surveys since for each tomographic slice the mean potential is expected to vanish due to homogeneity and isotropy properties, which leads to $<\tilde{\psi}>=<\psi>+f=f$. In this case, subtracting the mean of estimated lensing potential per slice removes the bias caused by the mass sheet degeneracy. Based on this lensing potential estimator, the 3D gravitational potential estimator becomes:

$$\tilde{\Phi} = \frac{1}{2}\partial_r r^2 \partial_r \left(\tilde{\psi} - <\tilde{\psi}>\right) \ . \tag{6.79}$$

Although Equation (6.79) formally allows the recovery of the 3D gravitational potential, it requires the evaluation of both angular and radial derivatives which cannot be directly computed from a discretely sampled galaxy catalogue. The first step of the method is to perform angular and radial binning of the catalogue before applying an additional angular and/or radial smoothing in order to compute these derivatives on a 3D grid. It is important to stress that in the radial direction the Newtonian potential is recovered through two consecutive differentiations which strongly amplifies the noise, thus the radial smoothing has a strong impact on the overall variance of the estimator (Bacon and Taylor (2003) find that the noise variance on the gravitational potential is proportional to the inverse fifth power of the width of radial smoothing).

In practice, even with binning and smoothing the galaxy catalogue, the primary map recovered by the estimator in Equation (6.79) is still highly noise dominated and an additional filtering step is required to recover a meaningful mass map.

Wiener filtering of the primary map was proposed in Bacon and Taylor (2003) and Hu and Keeton (2002) to beat down the noise assuming a given covariance matrix for the signal. In Bacon and Taylor (2003), the signal covariance matrix is just proportional to the unit matrix, with an amplitude tuned to match the expected signal for a cluster of a given mass. A more accurate modelling of the signal covariance matrix for recovering the large scale structure was suggested in Hu and Keeton (2002) based on the matter power spectrum, but the principles of the Wiener filter remain the same. The results of this method on simulations is illustatred on Figure 6.10 from Bacon and Taylor (2003). In this example, a space based lensing survey is simulated with $n_g = 100$ galaxies/arcmin^2 and an intrinsic shape noise of $\sigma_\epsilon = 0.2$. The field contains two NFW profile clusters at $z = 0.25$ and $z = 0.4$ of mass $8 \times 10^{13} M_\odot$ within a radius of 2 arcmin. The field reconstructed after Wiener filtering is shown on Figure 6.10b where the white crosses indicate the positions of the clusters. This method was first applied to the COMBO-17 survey (Wolf et al., 2003) in Taylor et al. (2004) to produce a 3D map of the dark matter environment of the A901/2 supercluster. It was also used in Massey et al. (2007) to map the 3D distribution of dark matter in the COSMOS survey (Scoville et al., 2007).

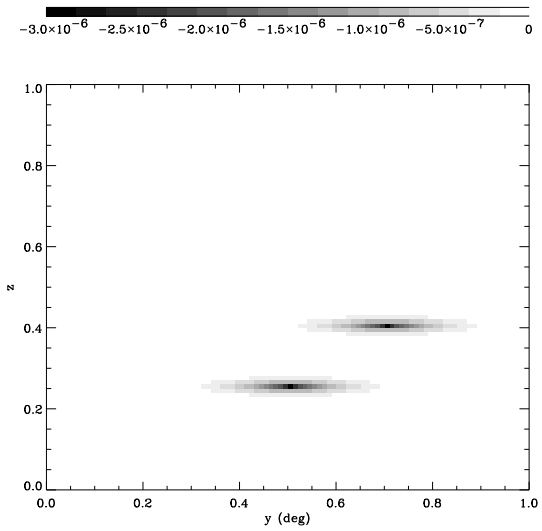

(a) Input gravitational potential

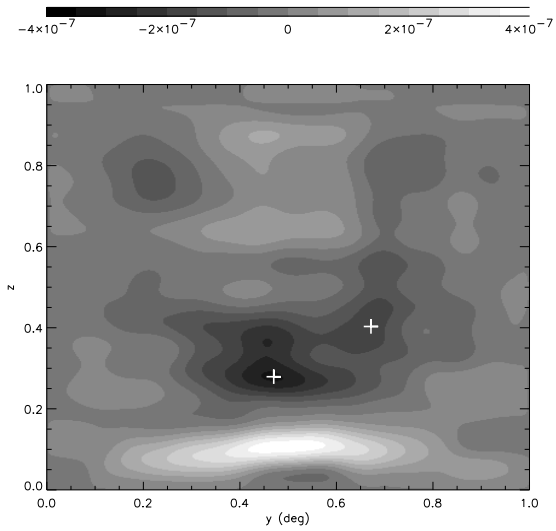

(b) Reconstructed gravitational potential after Wiener filtering

Figure 6.10: Reconstruction of two simulated clusters of masses $8 \times 10^{13} M_{\odot}$ within a radius of 2 arcmin for a space based lensing survey ($n_g = 100$ galaxies/arcmin^2 and $\sigma_\epsilon = 0.2$) using the 3D gravitational potential technique and Wiener filtering. Figures from Bacon and Taylor (2003)

For the recovery of isolated clusters, Hu and Keeton (2002) suggest two alternatives to the Wiener filter: the Maximum Entropy Method (MEM) and the Point Source method. They find that both methods can achieve remarkable localisation in redshift of individual clusters, for a good choice of prior. Interestingly, the Point Source method described by the authors is nothing more than an ℓ_0 sparse recovery problem. They are looking for the sparsest solution (the smallest number of halos) that fits the measured tomographic shear. This approach bears great resemblance to the 3D lensing technique presented in Chapter 8.

6.4.2 Regularised linear inversion

In the previous section the mapping was done in two steps. A first noisy estimation of the 3D gravitational potential using a Taylor inversion. However, this inversion of the lensing kernel is extremely ill-conditioned and tends to amplify the noise. Therefore, a second filtering step is required in order to produce a meaningful map. The limitation of this approach is that the inversion and filtering are considered in two separate steps and in its implementation the covariance matrix of the noise after inversion is computed by randomisation of the shear catalogue, not derived from the covariance matrix of the shear.

Simon et al. (2009) proposed a different approach, combining the inversion of the lensing kernel and the filtering of the map into one single minimum variance estimator. They consider the inversion of the following matrix problem:

$$\boldsymbol{\gamma} = \mathbf{P}\mathbf{Q}\boldsymbol{\delta} + \boldsymbol{n} \tag{6.80}$$

where $\mathbf{Q}$ is a matrix obtained by discretisation of the line of sight lensing efficiency kernel (see Chapter 8 for its expression) so that $\boldsymbol{\kappa} = \mathbf{Q}\boldsymbol{\delta}$ and $\mathbf{P}$ is the 2D Kaiser-Squires operator which com-

putes the shear from the convergence $\boldsymbol{\gamma} = \mathbf{P}\boldsymbol{\kappa}$. The authors propose the following minimum variance estimator to invert this problem:

$$\widehat{\boldsymbol{\delta}} = \left(\alpha \mathrm{Id} + \mathbf{S}\mathbf{R}^t\mathbf{N}^{-1}\mathbf{R}\right)^{-1}\mathbf{S}\mathbf{R}^t\mathbf{N}^{-1}\boldsymbol{\gamma} \tag{6.81}$$

In this expression, $\mathbf{N}$ is the covariance matrix of the noise at the level of the pixelated shear and can therefore be explicitly described in terms of the shear ellipticity variance and individual shear weights. This is in contrast with the previous approach. $\mathbf{S}$ is a model for the signal covariance, $\mathbf{R} = \mathbf{P}\mathbf{Q}$ is the forward 3D lensing operator relating density contrast to shear and α is a tuning parameter. For $\alpha = 0$, this estimator reduces to a simple minimum variance estimator without any prior on the signal. However, since the inversion of $\mathbf{R}$ is ill-posed, in the absence of regularisation the solution is extremely noisy. On the contrary, for $\alpha = 1$ this estimator becomes a Wiener filter and the regularisation of the inversion comes from the assumed model for the signal covariance $\mathbf{S}$. Therefore α makes it possible to tune the strength of the regularisation, to mitigate potential biases induced by the prior. Simon et al. (2009) propose two different priors for their Wiener filter by considering separately the radial and transverse power spectra of the 3D matter density.

This leads on one hand to a purely *transverse Wiener filter* which only puts a prior on the angular correlations of the 3D density field and consider different redshift slices as independent. Although this approach successfully attenuates the noise in the angular domain it does not add information to help constrain the redshift of the structures. In particular, the authors find that transverse filter induces an important and systematic broadening and biasing of the structures in redshift. On the other hand, the purely *radial Wiener filter* treats different lines of sight as independent and only consider the radial correlation of the 3D matter density. This method is actually very similar to the Wiener filter applied in the previous section for the Taylor inversion. However, this filter is much less successful at controlling the noise and the authors find that an additional Gaussian smoothing in the angular domain is necessary to recover a meaningful map.

The transverse Wiener filter method was applied in Simon et al. (2012) to the STAGES survey (Gray et al., 2009) which covers the same Abell 901/902 supercluster as the COMBO-17 survey mentioned in the previous section. Although the Wiener filtered map was able to recover the main structures, the authors conclude with a pessimistic view of the future of 3D lensing arguing that the linear filter they propose remains heavily biased and noisy and that only the most massive structures have a chance of being detected in a 3D map.

VanderPlas et al. (2011) follow a similar approach but instead of using a prior on the power spectrum to regularise the inversion they propose a Singular Value Decomposition (SVD) regularisation of the lensing operator which does not make any assumptions on the signal to recover. In their approach, the authors simply truncate a given percentage of lowest singular values which improves the condition number of the operator and greatly reduces the amplification of the noise during the inversion. This estimator can be written as:

$$\widehat{\boldsymbol{\delta}} = \mathbf{V}\overline{\boldsymbol{\Sigma}^{-1}}\mathbf{U}^t\mathbf{N}^{-1/2}\boldsymbol{\gamma} \tag{6.82}$$

where $\mathbf{U}, \mathbf{V}$ and $\boldsymbol{\Sigma}$ correspond to the SVD of the operator $N^{-1/2}\mathbf{R} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^t$ and $\overline{\Sigma^{-1}}$ is the regularised inverse singular value matrix matrix where a given percentage of the lowest singular values are discarded. Unfortunately, the modes associated with the remaining highest amplitude singular values are very broad in redshift and do not allow a good redshift resolution. Improving the redshift resolution necessarily requires to include modes associated with lower singular values at the price

of increasing the noise in the reconstruction. These authors find very comparable results with the radial Wiener filtering although their SVD method is about 500 times faster than the Wiener filter (thanks to a clever implementation of the SVD for sparse matrices).

These three different methods are presented on Figure 6.11 for a simulated NFW profile with a mass of $10^{15} M_{\odot}$ at redshift $z = 0.55$ under noise conditions similar to that of the Euclid survey ($n_{gal} = 30$ arcmin^{-2} and intrinsic ellipticity standard deviation $\sigma_{\epsilon} = 0.25$). This figure illustrates the main issues with the linear inversion methods. First and foremost, the recovered structures are systematically biased in redshift. This issue is particularly predominant for the transverse Wiener filter and was thoroughly studied in Simon et al. (2009). Another fundamental limitation is the redshift resolution of the reconstruction, all three methods presented here present a trade-off between redshift resolution and noise-reduction. In particular, VanderPlas et al. (2011) reach the conclusion that linear, non-parametric methods are fundamentally limited in redshift resolution.

These important limitations have hindered so far the generalisation of 3D lensing to more data sets. In Chapter 8 we present a new approach to this problem based on non-linear sparse recovery which is not bound by the same limitations as the linear methods presented here. In particular, we demonstrate a much improved redshift resolution and a robust estimation of the density contrast, which significatively broadens the range of applications of 3D lensing.

6.5 CONCLUSION

We introduced in this chapter the framework behind weak gravitational lensing and we also reviewed the fundamentals of mass-mapping techniques using gravitational shear to infer the dark matter distribution in 2D and in 3D. These concepts will be put to use in the following chapters where we develop new mass-mapping techniques based on the sparse optimisation framework introduced in the first part of this thesis.

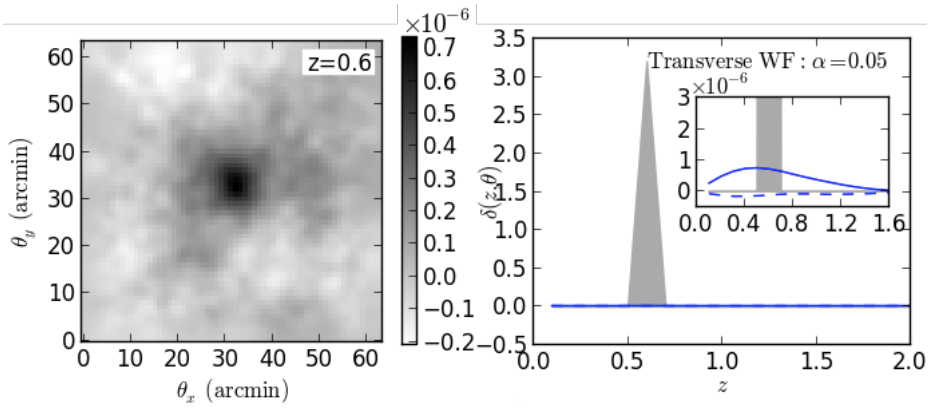

(a) Transverse Wiener filter

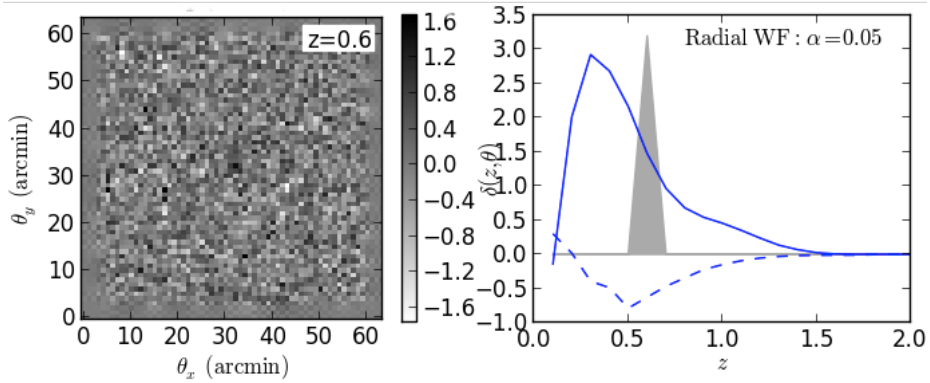

(b) Radial Wiener filter

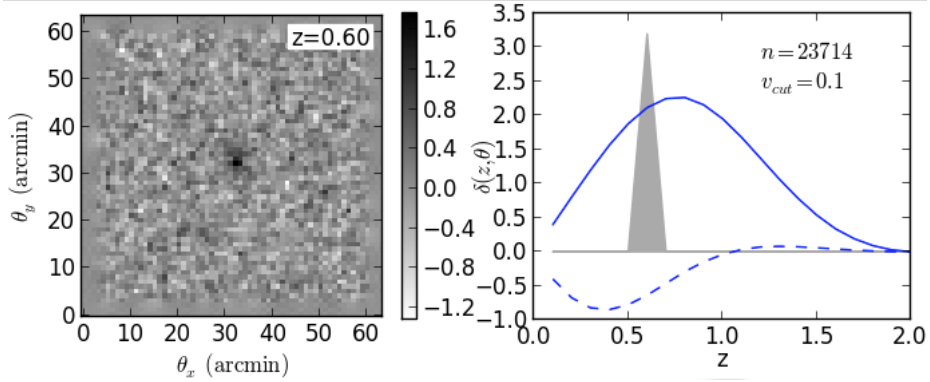

(c) SVD

Figure 6.11: Reconstruction of a $10^{15}\,M_\odot$ cluster at redshift $z = 0.6$ using the 3 linear inversion methods. The noise level corresponds to a survey with $n_{gal} = 30\text{gal/arcmin}^{-2}$ and $\sigma_\epsilon = 0.25$. The left panels show a slice in the 3D reconstructed density contrast at $z = 0.6$. The right panels show the central line of sight in the reconstruction. The gray wedges show the input density constrast and the solid blue lines show the E-mode signal and the dashed line show the B-mode signal.

7 HIGH RESOLUTION MASS MAPPING COMBINING SHEAR AND FLEXION

CONTENTS

In this chapter, we present a new weak lensing mass mapping technique, specifically designed to reconstruct the surface mass density of galaxy clusters with high resolution. As gravitational shear is generally noise dominated on small scales, we combine shear with gravitational flexion which allows us to recover small cluster substructures at the 10 arcsec scale.

While relying on Fourier estimators, our method does not require any binning or smoothing of the input galaxy catalogue, thus preserving small scale information. Instead, we treat the irregular galaxy sampling as part of an inverse problem which we address with the sparse regularisation framework introduced in Chapter 3.

We test our reconstruction method on a set of realistic weak lensing simulations corresponding to typical *Hubble Space Telescope* (HST)/*Advanced Camera for Surveys* (ACS) cluster observations and demonstrate our ability to recover substructures with the inclusion of flexion.

7.1 THE CLUSTER DENSITY MAPPING PROBLEM

Although very powerful, strong lensing is only effective within the Einstein radius of the cluster which represents for most clusters only a fraction of their total area. Outside of this region, constraints from gravitational shear are generally very poor below the arcminute scale which makes

the detection of substructures extremely unlikely. A promising avenue to bridge the gap between shear and strong lensing, and help resolve substructure outside of the innermost regions of galaxy clusters is gravitational flexion, which has already been shown to provide valuable constraints at these intermediate scales (Leonard et al., 2007).

Measuring flexion has proven to be a difficult task in practice but a new promising shape measurement method has recently been proposed by Cain et al. (2011). Whereas previous methods were based on shapelet (Goldberg and Bacon, 2005) or brightness moments (Okura et al., 2007), this new approach, called *Analytic Image Models* (AIM), is the first model fitting technique applied to the measurement of flexion. Among the benefits of this new method are a complete invariance under the mass-sheet degeneracy and a better characterisation of the measurement errors. Cain et al. (2015) demonstrated the possibility to use AIM measurements in combination to strong lensing to reconstruct cluster sub-structures.

Very few methods have been developed for mass-mapping including flexion. The first reconstruction technique using flexion has been proposed in Bacon et al. (2006) as an extension of the Kaiser-Squires Fourier estimator (see Section 6.3.2). Leonard et al. (2009); Leonard and King (2010) proposed an extension of aperture mass filters to the case of flexion. Finally, Cain et al. (2015) introduced a maximum likelihood approach combining shear, flexion and strong lensing for performing a joint reconstruction of the mass map. This last method being the most advanced combined reconstruction technique, we will be particularly interested in comparing our results to theirs.

In this chapter we develop a new approach for mapping the surface mass density of galaxy clusters from shear and flexion, which we specifically design to be able to recover small scale substructures. This mass-mapping problem is complex and involves a number of sub-problems listed below.

IRREGULAR SAMPLING In any lensing survey, the shear field is only sampled at the position of background galaxies which are randomly distributed. When mapping the matter distribution on large scales, these galaxies can be binned into regular pixels, smaller than the scales of interest, thus providing a regular sampling of the field from which the convergence map can be conveniently computed (for instance using a Kaiser-Squires technique). Although this approach is convenient, information on scales smaller than the bin size are irremediably lost. In the method developed in this chapter we aim to avoid any smoothing or binning of the input data, thus preserving all available information. This comes at a cost however, as the mass-mapping problem becomes in this case an ill-posed inverse problem. As we will demonstrate, this problem can be successfully solved using the sparse regularisation framework introduced in Chapter 3.

NOISE REGULARISATION Intrinsic shear and flexion noise becomes a particular problem on small scales as the signal cannot be averaged over a large number of galaxies. In fact, very few details can be expected to reach a significant SNR on the smallest scales (smaller than 10 arcseconds). Yet, very locally the signal can still become significant, for instance at the vicinity of small substructures. While a simple smoothing is an effective way of suppressing these noise dominated scales, it will also suppress any small substructure, even the few that reach significant amplitudes. Therefore, ideally the noise regularisation should be adaptive and allow the detection of features at different scales based on their SNR. This is an issue with most maximum likelihood techniques which generally assume a smoothness prior to regularise the reconstruction when using shear in-

formation. In the method developed in this chapter, our regularisation will be based on a mutliscale sparsity prior, which is a much more powerful denoising technique.

NON-LINEARITY While the mass-mapping problem remains linear in the weak lensing regime ($\kappa \ll 1.$), at the close vicinity of the galaxy clusters we aim to reconstruct in this chapter this condition no longer holds. The inversion method needs to properly take into account the reduced shear $g = \frac{\gamma}{1-\kappa}$ and reduced flexion $F = \frac{\mathcal{F}}{1-\kappa}$ in order to avoid important systematic biases in the reconstructed map. However, in this case, the inversion problem becomes non-linear and the algorithms for sparse regularisation presented in Chapter 3 no longer formally apply. The problem will require some linearisation scheme in order to remain tractable by the same methods.

MASS-SHEET DEGENERACY A well known issue with all weak lensing mass-mapping techniques is the mass-sheet degeneracy. In the non-linear regime we are considering in this chapter, the reconstructed convergence map is invariant under the λ-transform introduced in Section 6.1.4. However, as was explained in that section, when including redshift information for individual galaxies, the degeneracy can be lifted at the close vicinity of the most significant structures. As our method avoids binning the input galaxies, it is possible to incorporate individual reshift estimates for the sources. Even if not very effective in general at breaking the mass-sheet degeneracy, the redshift information is still important to ensure a proper scaling of the reconstructed surface mass density from which physical masses can be measured.

COMBINING FLEXION Small scales details such as cluster substructures are generally lost when using shear information alone as the noise becomes very important on small scales. On the contrary, while flexion becomes noise dominated on large scales, on small scales (below 0.5 arcmin) it becomes far less noisy than the shear and thus becomes sensitive to substructures which would have been otherwise lost. Shear and flexion are therefore complementary but the inclusion of flexion adds to the complexity of the inverse problem. In particular, it is important to ensure that shear and flexion are optimally combined, in a way which is coherent with the noise regularisation.

In order to disentangle these various difficulties, the method presented in the following sections will progressively build up in complexity. We start by addressing the linear weak lensing inversion and the noise regularisation in Section 7.2. In Section 7.3, we incorporate the non-linearity induced by the reduced shear, we add redshift information for individual galaxies and we extend the framework to include flexion. Finally, in Section 7.4 we demonstrate on realistic cluster simulation that our method is successful at reconstructing the surface mass density and we illustrate the importance of flexion for the recovery of small-scale substructures.

7.2 SPARSE REGULARISATION OF THE LINEAR INVERSE PROBLEM

In Fourier based methods, the fundamental motivation behind regular binning is to enable the computation of the DFT of the shear, in order to apply the Kaiser-Squires estimator. Indeed, the DFT is not well defined outside of a regular grid with no gaps. Since in practice the shear is only sampled at the position of randomly distributed background galaxies, this binning is therefore a necessary step to ensure a well defined inversion. However, preventing gaps in the sampling of the

shear can require a potentially coarse binning, thus irremediably losing small scale information. Even then, gaps in the data are unavoidable due to various various masks applied in practice to the survey (bright stars, CCD defects,...).

We propose in this section a new method for reconstructing the convergence map, based on Fourier estimators, which does not require any binning or smoothing of the input shear. Although this approach allows us to preserve small-scale information which would otherwise be lost, it does turn the reconstruction problem into an ill-posed inverse problem, which we address using the sparse regularisation framework introduced in Chapter 3.

We only consider the linear weak lensing problem in this section, reconstructing the convergence map assuming knowledge of the shear. The full treatment of the complete problem is addressed in the next section.

7.2.1 Non-equispaced Discrete Fourier Transform

When either the spatial or frequency space is not regularly sampled, the DFT can be generalised to the *Non-equispaced Discrete Fourier Transform* (NDFT). The NDFT is of course no longer an orthogonal transform and is usually not even invertible. We will consider the case where only the spatial nodes $\boldsymbol{x} = (x_l)_{0 \leq l < M}$ are arbitrary while the frequency nodes $\boldsymbol{k} = [\![0, N[\![$ are N regularly spaced integers. Computing the NDFT from a set of Fourier coefficients $\hat{\boldsymbol{f}} = \left(\hat{f}_k\right)_{0 \leq k < N}$ simply amounts to evaluating the trigonometric polynomials:

$$\forall l \in [\![0, M[\![, \quad f_l = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} \hat{f}_k e^{2\pi i k x_l} \,. \tag{7.1}$$

This operation can more conveniently be expressed using matrix notations as:

$$\boldsymbol{f} = \mathbf{T}\hat{\boldsymbol{f}} \quad \text{with} \quad T_{lk} = \frac{1}{\sqrt{N}} e^{2\pi i k x_l} \,, \tag{7.2}$$

where $\mathbf{T}$ is the NDFT matrix, with matrix elements T_{lk}. Note that in the case of equispaced spatial nodes such that $x_l = \frac{1}{N} l$, this operation corresponds to the conventional DFT and $\mathbf{T}$ reduces to the Fourier matrix $\mathbf{F}$ defined as:

$$F_{lk} = \frac{1}{\sqrt{N}} e^{2\pi i k l / N} \,. \tag{7.3}$$

This Fourier matrix is unitary and its inverse is simply its Hermitian conjugate: $\mathbf{F}^{-1} = \mathbf{F}^*$. On the contrary, for non-equispaced spatial nodes $\boldsymbol{x}$, the operator $\mathbf{T}$ is typically neither orthogonal nor admits an inverse. Still, one can consider the adjoint NDFT operator $\mathbf{T}^*$:

$$T^*_{kl} = \frac{1}{\sqrt{N}} e^{-2\pi i k x_l / N} \,. \tag{7.4}$$

Although this adjoint operation no longer corresponds to the inverse of the transform, it can be used in practice to estimate the inverse through a least squares problem of the form:

$$\hat{\boldsymbol{f}} = \arg\min_{\hat{\boldsymbol{x}}} \| \boldsymbol{f} - \mathbf{T}\hat{\boldsymbol{x}} \|_2^2 \,. \tag{7.5}$$

This problem can efficiently be solved using iterative algorithms (in particular using a conjugate gradient) which involve the computation of both $\mathbf{T}$ and $\mathbf{T}^*$.

It is therefore important to have fast algorithms for the computation of the NDFT and its adjoint. Note that a naive evaluation of the sum in Equation (7.1) would scale as $O(N \times M)$ which is prohibitively large for most applications (including for our mass-mapping problem). In this work, we use a fast approximate algorithm[1] to evaluate the NDFT and its adjoint, called *Non-equispaced Fast Fourier Transform* (NFFT) (Keiner et al., 2009), which only scales as $O(N \log(N) + |\log(\epsilon)|M)$ where ϵ is the desired accuracy. For a given sampling of the frequency space (i.e. for a given N), the NFFT only linearly scales with the number M of spatial nodes x_l, which in our case will correspond to the number of galaxies in the survey.

7.2.2 Sparse regularisation of the inversion problem

Consider a lensing survey with N_g galaxies. We can write the expression of the shear $\boldsymbol{\gamma} = (\gamma_i)_{i \in [\![0, N_g [\![}$ at the position of each galaxy given a convergence map $\boldsymbol{\kappa}$ as

$$\boldsymbol{\gamma} = \mathbf{TPF}^* \boldsymbol{\kappa} \tag{7.6}$$

In this expression, $\mathbf{F}$ is the Fourier matrix and $\mathbf{T}$ is the NDFT matrix defined for arbitrary spatial nodes $\boldsymbol{x}$ placed at the position of each galaxy in the survey. The diagonal operator $\mathbf{P}$ implements the transformation from convergence to shear in Fourier space:

$$\hat{\boldsymbol{\gamma}} = \mathbf{P}\hat{\boldsymbol{\kappa}} = \left(\frac{k_1^2 - k_2^2}{k^2} + i \frac{2k_1 k_2}{k^2} \right) \hat{\boldsymbol{\kappa}} \tag{7.7}$$

Of course, this expression is not defined for $k_1 = k_2 = 0$, which corresponds to the well known mass-sheet degeneracy and by convention we will set the mean to 0. We are using complex notations for both shear and convergence, with in particular $\boldsymbol{\kappa} = \boldsymbol{\kappa}_E + i\boldsymbol{\kappa}_B$ where κ_E and κ_B are respectively E- and B-modes maps.

An important point to stress is that the operator $\mathbf{P}$ is unitary, with $\mathbf{P}^* \mathbf{P} = \mathrm{Id}$, just like the Fourier matrix $\mathbf{F}$, and as such is readily invertible. Solving Equation (7.6) therefore reduces to the inversion of the NDFT operator, which is the only difficult step. As mentioned in the previous section, if the spatial nodes are not regularly spaced, estimating the inverse NDFT is in the general case an ill-posed inverse problem.

Our aim is to apply the sparse regularisation framework introduced in Chapter 3 to the inversion of the NDFT operator, thus yielding an estimate of the convergence. We consider the following sparse optimisation problem:

$$\arg\min_{\boldsymbol{\kappa}} \frac{1}{2} \| \boldsymbol{\gamma} - \mathbf{TPF}^* \boldsymbol{\kappa} \|_2^2 + \lambda \| \boldsymbol{w} \circ \boldsymbol{\Phi}^t \boldsymbol{\kappa} \|_1 + i_{\Im(\cdot)=0}(\boldsymbol{\kappa}) \,. \tag{7.8}$$

Let us detail each term in this expression. The first term is a quadratic data fidelity term, where we assumed a shear covariance matrix proportional to the identity matrix for simplicity. In practice, different inverse variance weights can be applied to each galaxy, but the covariance matrix remains diagonal as we are treating galaxies independently. The second term is an analysis-based sparsity constraint where $\boldsymbol{\Phi}$ is a dictionary providing a sparse representation of the signal we want to recover, $\circ$ is the Hadamard product, and $\boldsymbol{w}$ is a vector of weights allowing us to adaptively adjust

1 The C++ library NFFT 3 is available at https://www-user.tu-chemnitz.de/~potts/nfft

the ℓ_1 ball based on the local level of noise (see next section). Finally, the last term imposes the imaginary part of the solution (i.e. B-modes) to vanish. Remember that the indicator function i_C of a set C is defined by

$$i_C(x) = \begin{cases} 0 & \text{if } x \in C \\ +\infty & \text{otherwise} \end{cases} \tag{7.9}$$

Here, $C = \left\{\boldsymbol{\kappa} \in \mathbb{C}^{N\times N} \quad | \quad \Im(\boldsymbol{\kappa}) = 0\right\}$ where $N \times N$ is the size of the reconstruction grid, and as a result any solution with a non-zero imaginary part is excluded. This additional constraint is crucial as the irregular galaxy sampling leads to important leakage between E- and B-modes. As E-modes are constrained by the sparsity prior, most of the signal would tend to leak towards B-modes without this additional term. Vanishing B-modes is of course a completely physically motivated prior as gravitational lensing only produces E-modes, but this also means that this method will be sensitive to spurious B-modes resulting from uncorrected systematics which will contaminate the recovered signal.

To efficiently solve this problem, we adopt the primal-dual algorithm from Vu (2013), introduced in Section 3.3.2.4. This algorithm is capable of handling the 3 terms of Equation (7.8) and especially the lack of explicit proximal operator for the analysis-based sparsity constraint. Given that the proximity operator of $i_{\Im(\cdot)=0}(\kappa)$ is simply the real part of κ, $\text{prox}_{i_{\Im(\cdot)=0}}(\kappa) = \Re(\boldsymbol{\kappa})$, the specialisation of this algorithm to the problem at hand is straightforward and described in Algorithm 7.1.

Algorithm 7.1 Analysis-based κ sparse recovery from shear

Require:
Shear of each galaxy in the survey $\boldsymbol{\gamma}$.
Sparsity constraint parameter $\lambda > 0$.
Weights $w_i > 0$.
$\tau = 2/(\| \mathbf{\Phi} \|^2 + \| \mathbf{T} \|^2)$.

$\boldsymbol{\kappa}^{(0)} = 0$
$\forall i, \quad \lambda'_i = \lambda w_i$
for $n = 0$ to $N_{\text{max}} - 1$ **do**
  $\boldsymbol{\nabla}^{(n)} = \mathbf{F}\mathbf{P}^*\mathbf{T}^*\left(\boldsymbol{\gamma} - \mathbf{T}\mathbf{P}\mathbf{F}^*\boldsymbol{\kappa}^{(n)}\right)$
  $\boldsymbol{\kappa}^{(n+1)} = \Re\left[\boldsymbol{\kappa}^{(n)} + \tau\left(\boldsymbol{\nabla}^{(n)} - \mathbf{\Phi}\boldsymbol{\alpha}^{(n)}\right)\right]$
  $\boldsymbol{\alpha}^{(n+1)} = (\text{Id} - \text{ST}_{\lambda'})\left(\boldsymbol{\alpha}^{(n)} + \mathbf{\Phi}^t\left(2\boldsymbol{\kappa}^{(n+1)} - \boldsymbol{\kappa}^{(n)}\right)\right)$
end for
return $\kappa^{(N_{\text{max}})}$.

The solution of this optimisation problem tends to be biased as a side effect of the ℓ_1 sparsity constraint. In order to correct for this bias and improve the quality of the solution we implement the reweighted-ℓ_1 strategy of Candès et al. (2008) which was detailed in Section 3.2.3. This method relies on iteratively solving Equation (7.8), adjusting each time the weights $\boldsymbol{w}$ based on the previous estimate of the solution. This procedure is described below:

1. Set the iteration count $\ell = 0$ and initialise the weights $\boldsymbol{w}^{(0)}$ according to the procedure described in Section 7.2.4.

2. Solve the weighted ℓ_1 minimisation problem of Equation (7.8) using Algorithm 7.1, yielding a solution $\boldsymbol{\kappa}^{(\ell)}$.

3. Update the weights based on the wavelet transform of the solution $\boldsymbol{\alpha}^{(\ell)} = \Phi^t \boldsymbol{\kappa}^{(\ell)}$:

$$w_i^{(\ell+1)} = \begin{cases} \dfrac{w_i^{(0)}}{|\alpha_i^{(\ell)}|/\lambda w_i^{(0)}} & \text{if } |\alpha_i^{(\ell)}| \geq \lambda w_i^{(0)} \\ w_i^{(0)} & \text{if } |\alpha_i^{(\ell)}| < \lambda w_i^{(0)} \end{cases}, \tag{7.10}$$

4. Terminate on convergence. Otherwise, increment ℓ and go to step 2.

In practice we find that 3 to 5 re-weightings are generally sufficient to reach a satisfying solution.

7.2.3 Choice of dictionary

For any sparse regularisation method, an appropriate choice of dictionary is important for the quality of the result. This is especially true for noise dominated problems where the prior takes prevalence when the data is not constraining. Previous sparsity based methods developed for weak lensing mass-mapping either employed starlets (see Section 4.3.2) for denoising (Starck et al., 2006) or DCT for inpainting (Pires et al., 2009). Indeed, at small scale, the non-Gaussian convergence signal essentially generated by isolated galaxy clusters is extremely well represented using the starlet dictionary which features isotropic atoms, adapted to the average circular profile of dark matter halos. On the other hand, on large scales, the DCT is more efficient at capturing the Gaussian part of the convergence signal and has proven to be an excellent dictionary to inpaint missing data due to masks without altering the power spectrum of the reconstructed maps.

In this work, we aim at reconstructing the convergence map on small scales and we therefore adopt the starlet dictionary for the reasons stated above. However, we find that when using the starlet alone, details at the finest scale are not sufficiently constrained, in particular spurious isolated pixels tend to contaminate the solution, as shown on Figure 7.1a. To help penalise this unwanted behaviour, we build an hybrid dictionary by concatenating to the starlet dictionary the first scale of an undecimated bi-orthogonal wavelet transform, more specifically a Battle-Lemarié wavelet of order 5 (see Section 4.2.3). Contrary to starlet atoms, Battle-Lemarié wavelets are much more oscillatory and have a larger support (formally infinite but with an exponential decay). This makes them relatively inefficient at sparsely representing singularities such as isolated pixels, which are therefore more strongly penalised by the sparsity prior.

To illustrate the benefits of using this hybrid dictionary, we compare in Figure 7.1 the results of our reconstruction algorithm in a simple noiseless case when using starlets alone or the hybrid dictionary described above, all other parameters being kept fixed. As this simple qualitative comparison demonstrates, starlets alone tend to create small pixel-sized artefacts, even in the absence of noise. These are completely eliminated by the inclusion of the Battle-Lemarié wavelets.

We add that although we find this dictionary to be effective at regularising the mass-mapping inversion, it remains generic and was not specifically designed for an optimal representation of convergence maps. More specific dictionaries could be used just as well, for instance using DL techniques, and would potentially improve further our results.

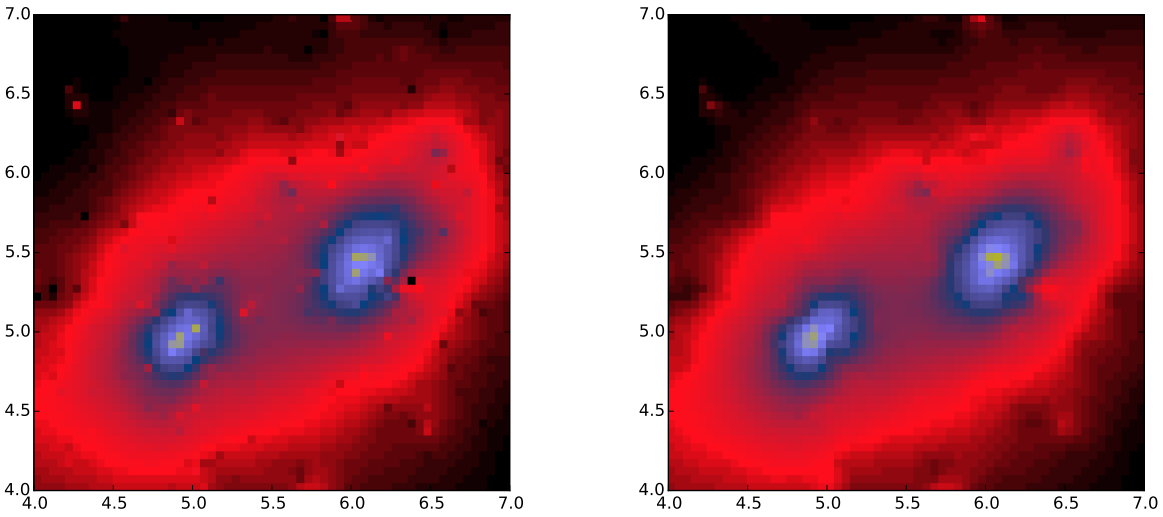

(a) Starlet dictionary

(b) Starlet + Battle-Lemarié dictionary

Figure 7.1: Comparison of reconstruction using starlets or a combination of starlets and Battle-Lemarié wavelets as the dictionary $\mathbf{\Phi}$. The Battle-Lemarié atoms are very effective at penalising isolated pixels.

7.2.4 Adjusting the sparsity constraint

A recurring issue with sparse recovery problems such as the one stated in Equation (7.8) is the choice of the regularisation parameter λ. There is unfortunately no general rule indicating how to set this parameter in practice. For this application, we adopt the approach that was proposed in Chapter 5, which consists in defining this parameter with respect to the noise level.

Formally, the parameter λ scales the ℓ_1 ball used in the sparsity constraint. In practice, it defines the level of Soft Thresholding applied to the dual variable $\boldsymbol{\alpha}$ (see line 6 in Algorithm 7.1) and therefore discriminates between significant and non-significant coefficients. While this threshold can be set according to a given sparsity model of the signal to recover, for noise dominated problems, it is much more crucial to define this threshold with respect to the noise level. As the noise statistics vary across the field, depending on the specific galaxy distribution, we introduce a vector of weights $\boldsymbol{w}$ (found in the ℓ_1 term in Equation (7.8)) with the purpose of locally scaling the sparsity constraint based on the standard deviation of the noise propagated to the coefficients $\boldsymbol{\alpha}$. For each wavelet coefficient α_i we set the weight w_i to the estimated standard deviation $\sigma(\alpha_i)$. As a result, the level of Soft Thresholding applied to each coefficient α_i is $\lambda'_i = \lambda w_i$ and accounts for noise variations across the field. As the threshold is proportional to the standard deviation of the noise, it can be interpreted as an hypothesis test to determine if a coefficient is due to signal or noise, assuming Gaussian statistics for the noise, which adds a powerful detection aspect to the sparsity constraint.

To estimate this noise level, and therefore set the weights $\boldsymbol{w}$, it is first necessary to understand how the noise in the data propagates to the dual variable $\boldsymbol{\alpha}$. By considering Algorithm 7.1, it can be seen that the noise at the level of the shear $\boldsymbol{\gamma}_N$ is propagated to the wavelet coefficients through the operation $\mathbf{\Phi}^t \mathfrak{R} \left(\mathbf{F} \mathbf{P}^* \mathbf{T}^* \boldsymbol{\gamma}_N \right)$. In practice, we estimate the standard deviation of wavelet coefficients by generating Monte-Carlo noise simulations, obtained by keeping the galaxies at their observed

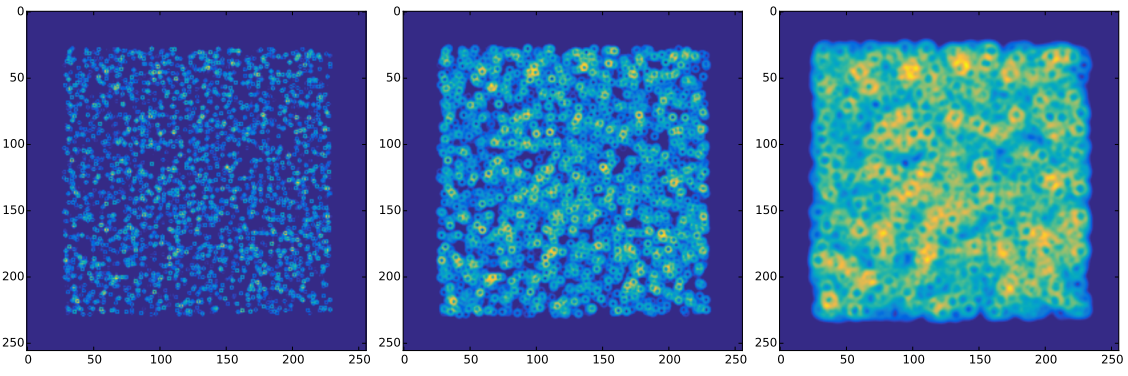

(a) Scale j=0 **(b)** Scale j=1 **(c)** Scale j=2

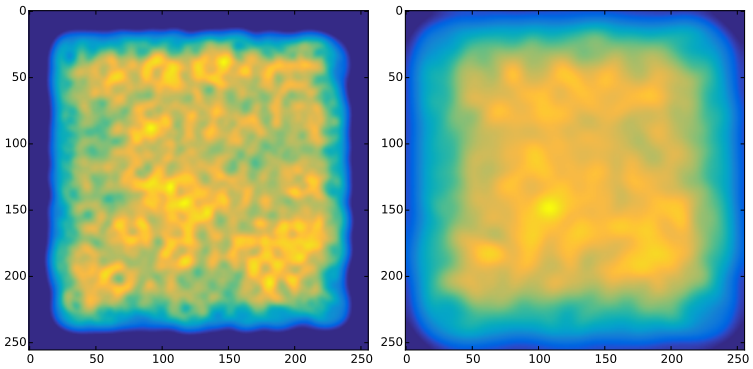

(d) Scale j=3 **(e)** Scale j=4

Figure 7.2: Standard deviation maps $\boldsymbol{w}$ used to scale the sparsity constraint, obtained by propagating the shear noise to the wavelet coefficients. We show here noise maps for five successive starlet scales.

position while randomising their orientation. Note that this step needs only to be performed once, outside of the main iteration of Algorithm 7.1. An example of the resulting standard deviation maps for different wavelet scales is shown on Figure 7.2. As can be seen, at the finest scales, the noise level has important fluctuations across the field and the contribution of individual galaxies can be seen. On larger scales, these local fluctuations are smoothed out and the noise level becomes much more homogeneous.

Using this strategy therefore allows us to tune locally the sparsity constraint to take into account the specific galaxy distribution of the survey and leaves only one free parameter λ.

7.2.5 Numerical experiment

The algorithm presented in this section is only meant to address the two first difficulties of the complete mass-mapping problem, namely the irregular sampling of the shear and the presence of noise, the complete problem being solved in the next section. In this simplified setting, we present a small numerical experiment to check the algorithm's effectiveness at solving the linear inverse problem. We compare our results to a standard Kaiser-Squires inversion and to the FLens DCT sparse inpainting technique of Pires et al. (2009).

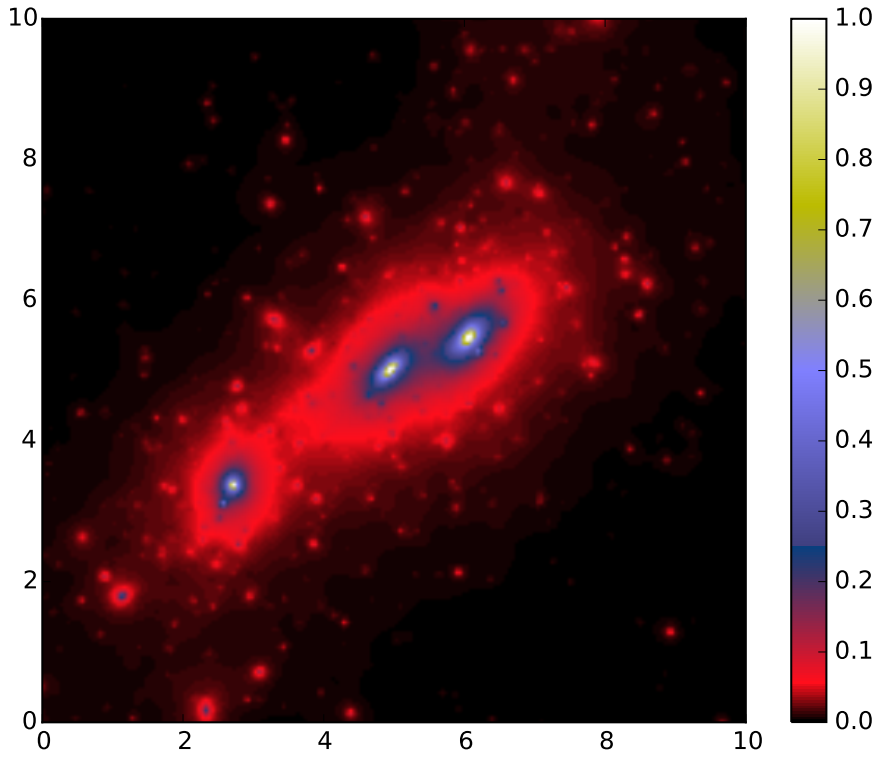

Figure 7.3: Test convergence map generated from an N-body simulation. The clusters are located at $z_l = 0.3$ while the source plane is placed at $z_s = 1.2$.

We simulate a 10×10 arcmin^2 field, containing a group of galaxy clusters extracted from the Bolshoi N-body simulations (Klypin et al., 2011) (see Section 7.4.1 for more details), as shown on Figure 7.3. These clusters are placed at redshift $z_l = 0.3$ and we simulate a lensing catalogue with randomly distributed sources on a single lens plane at redshift $z_s = 1.2$. Note that we only simulate shear measurements from the input convergence map and not the reduced shear.

We consider here the noiseless inpainting problem where the input shear is exactly known at the position of the sources. Of course this problem is unrealistic but does illustrate the impact of missing data on the mass-mapping inversion using Fourier estimators.

We generate a first shear catalogue with a density of 30 galaxies per square arcminute which corresponds to a typical Euclid density. To exacerbate the effects of missing data, we reconstruct the input convergence map on a grid with 3 arcsecond pixels. Binning the input catalogue at this resolution corresponds to an average of 0.075 galaxies per pixels, which means mostly empty pixels as illustrated by the mask on Figure 7.4a where empty pixels are shown in white. The result of a blunt Kaiser-Squires inversion with such a mask is shown on Figure 7.4b which corresponds to the solution of the inverse problem in the absence of regularisation. In order to recover the main structures we also show the result of a Gaussian smoothing with a kernel of 0.25 arcminute on Figure 7.4c. Note that on this last figure the color scale is adjusted for better visualisation but the amplitude of the reconstruction is an order of magnitude below the input signal smoothed with the same kernel. As illustrated by these plots, small scale information is lost with this approach and although using larger bins would regularise the inversion these details would still be lost.

We now apply Algorithm 7.1 to the same catalogue. As the weighting scheme presented in Section 7.2.4 is not meant to be used on noiseless data, we use a noisy version of the shear data to

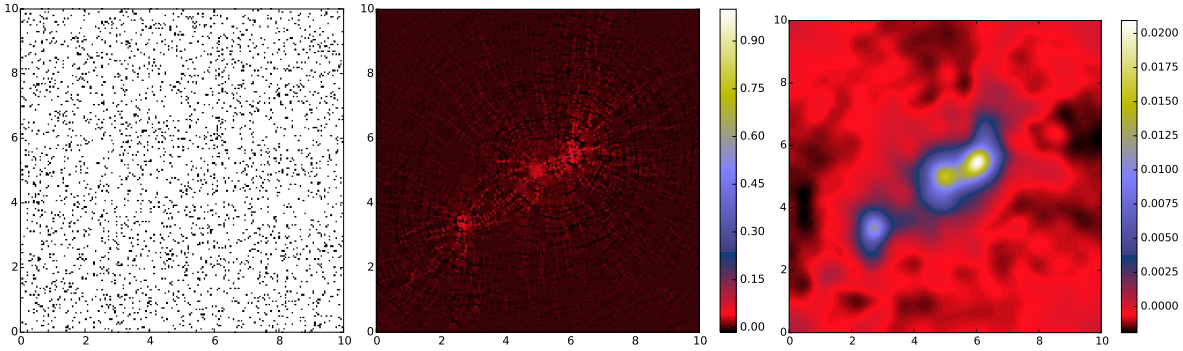

(a) Galaxy distribution **(b)** Kaiser-Squires inversion **(c)** Kaiser-Squires with a 0.25 arcmin Gaussian smoothing

Figure 7.4: Reconstruction of the input convergence map using 30 galaxies per square arcminute using a simple Kaiser-Squires estimator without additional regularisation. The left panel shows the mask applied to the binned shear map with empty pixels marked in white.

estimate the weights $\boldsymbol{w}$ and we set the regularisation parameter λ to a very low value (0.01) to apply the algorithm on noiseless data. To compare our results with a similar inpainting technique we also use the FLens method of Pires et al. (2009). Figure 7.5a and Figure 7.5b show the recovered convergence maps. In both case, the results constitute an improvement over a simple Kaiser-Squires inversion but the quality of the reconstruction is much better using the method introduced in this section, which is able to recover pixel scale details. The main reason explaining this difference of behaviour between two similar sparsity based methods is the difference of dictionaries used for the sparsity constraint. In the case of FLens, a DCT dictionary is used while a wavelet transform is used in our method. If the DCT is a good choice for capturing the lensing signal on larger scales (texture-like, mostly stationary signal) it is not adapted to individual galaxy clusters which are much more sparsely represented by wavelets (smooth signal with local singularities). This illustrates the importance of the choice of representation in sparse regularisation techniques.

We also generate a second catalogue with a higher density of 100 galaxies per square arcminute, which is more typical for HST deep lensing surveys for the study of individual clusters. As shown on Figure 7.5c and Figure 7.5d finer details are recovered by both methods.

There are two reasons explaining how pixel scale details can still be recovered using our method when about 93 % of the data is missing. First and foremost, the shear information is non local and even small structures formally impact the shear across all the field, albeit with an amplitude decaying to the square of the angular distance. The second reason is the use of *isotropic* wavelets which provide a sparse representation of the true convergence map while being morphologically distinct from the response of individual galaxies which is quadrupolar. As a result, the sparsity constraint greatly favours the true convergence map, which is smooth with small isotropic features, over anisotropic artefacts resulting from the irregular shear sampling.

The point of this numerical experiment is to show that although the shear field is randomly sampled at a relatively low rate, high frequency information can still be recovered using our method. Of course, in practice small-scale details are generally lost in the considerable amount of noise coming from intrinsic galaxy ellipticities, but not necessarily at the close vicinity of the center of galaxy clusters where the shear signal can become significant even on small scales. The ability to

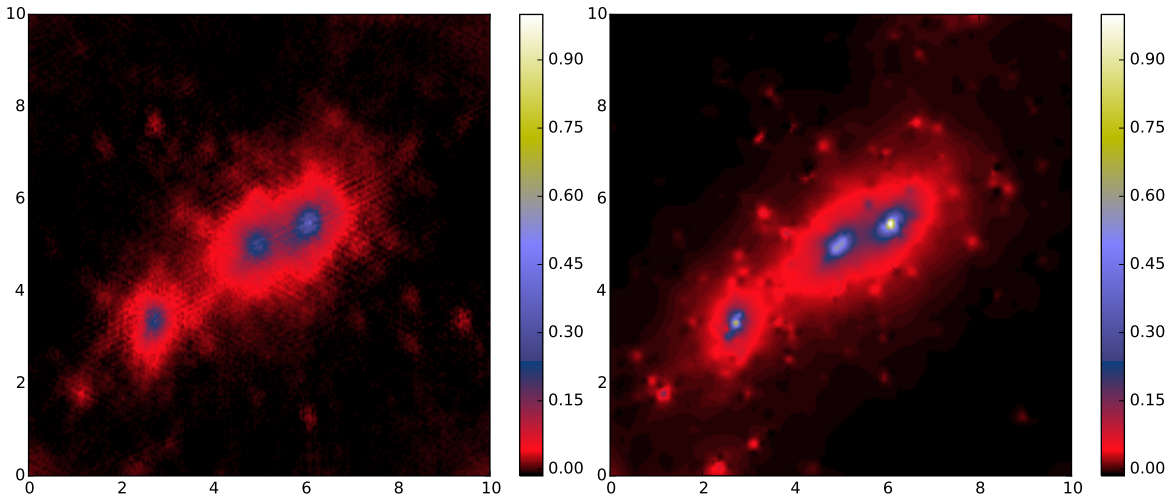

(a) FLens inpainting with 30 gal/arcmin^{-2}

(b) This method with 30 gal/arcmin^{-2}

(c) FLens inpainting with 100 gal/arcmin^{-2}

(d) This method with 100 gal/arcmin^{-2}

Figure 7.5: Recovered convergence maps using the method introduced in this chapter and the FLens inpainting method for 30 and 100 galaxies per square arcminute.

reconstruct small-scale details is even more relevant when including flexion information as the noise power spectrum of the flexion estimator drops at small scales.

7.3 CLUSTER DENSITY MAPPING

In the previous section, we have introduced an algorithm based on sparse regularisation for solving the simplified linear inversion problem. Although not realistic, this problem is still an important step towards the complete surface density reconstruction that we address now in this section.

We detail how the algorithm of the previous section can be modified to take into account the reduced, redshift information for individual galaxies, and flexion information.

7.3.1 Handling the reduced shear

The problem addressed in the previous section assumed knowledge of the shear, in which case the inverse problem remains linear. While this assumption can be made in the weak regime, it breaks down at the vicinity of the structures (galaxy clusters) we are interested in mapping. Although the method presented in the previous section no longer directly applies, we present in this section how it can be extended to take into account the reduced shear $g = \frac{\gamma}{1-\kappa}$, which makes the inversion problem non-linear.

Throughout this chapter, we will restrict ourselves to the case $|\boldsymbol{g}| \leq 1$ for simplicity, assuming that the sources which do not verify this condition can be identified and excluded from the sample. The method presented here could be extended using an iterative procedure to identify and thus treat accordingly sources lying in the region $|\boldsymbol{g}| > 1$.

Replacing the shear by the reduced shear in the inversion problem stated in Equation (7.8) yields:

$$\arg\min_{\boldsymbol{\kappa}} \frac{1}{2} \parallel \boldsymbol{g} - \frac{\mathbf{TPF}^*\boldsymbol{\kappa}}{1 - \mathbf{TF}^*\boldsymbol{\kappa}} \parallel_2^2 + \lambda \parallel \boldsymbol{w} \circ \boldsymbol{\Phi}^t \boldsymbol{\kappa} \parallel_1 + i_{\mathfrak{I}(\cdot)=0}(\boldsymbol{\kappa}) \,. \tag{7.11}$$

Note that at the denominator the NDFT operator is only used to evaluate the convergence at the position of each galaxy. In this form, the full problem cannot be directly addressed using the algorithm presented in the previous section as the operator to invert is no longer linear. Nonetheless, following a common strategy to handle this non-linearity, the term $C_\kappa^{-1} = 1/(1 - \mathbf{TF}^*\boldsymbol{\kappa})$ can be factored out and be interpreted as a diagonal covariance matrix, which depends on the signal $\boldsymbol{\kappa}$:

$$\arg\min_{\boldsymbol{\kappa}} \frac{1}{2} \parallel C_\kappa^{-1} \left[(1 - \mathbf{TF}^*\boldsymbol{\kappa})\boldsymbol{g} - \mathbf{TPF}^*\boldsymbol{\kappa}\right] \parallel_2^2 + \lambda \parallel \boldsymbol{w} \circ \boldsymbol{\Phi}^t \boldsymbol{\kappa} \parallel_1 + i_{\mathfrak{I}(\cdot)=0}(\boldsymbol{\kappa}) \,. \tag{7.12}$$

If the factor C_κ^{-1} is kept fixed, the problem is now linear and can be solved once again using the algorithms presented in Chapter 3. By iteratively solving this linearised problem, updating each time the matrix C_κ^{-1} with the current estimate of κ can recover the solution of the original problem in Equation (7.11). This is for instance the strategy adopted in Merten et al. (2009).

The gradient of the quadratic term in Equation (7.12) becomes:

$$-\frac{1}{2}\frac{\partial}{\partial \boldsymbol{\kappa}} \chi^2(\boldsymbol{\kappa}) = \mathbf{FP}^*\mathbf{T}^* C_\kappa^{-2} \left[(1 - \mathbf{TF}^*\boldsymbol{\kappa})\boldsymbol{g} - \mathbf{TPF}^*\boldsymbol{\kappa}\right] \quad + \quad \mathbf{FT}^*\boldsymbol{g}^* C_\kappa^{-2} \left[(1 - \mathbf{TF}^*\boldsymbol{\kappa})\boldsymbol{g} - \mathbf{TPF}^*\boldsymbol{\kappa}\right] \tag{7.13}$$

The first term of this expression simply corresponds to the same gradient as in the linear case but corrected for the reduced shear. In particular, the noise affecting the measured shear still propagates linearly. On the contrary, the second term is completely new compared to the linear case and is problematic in the sense that it becomes a quadratic function of the reduced shear $\boldsymbol{g}$. If one assumes the reduced shear noise to be Gaussian, the noise contribution of this second term becomes χ^2 distributed. As explained in the previous section, our regularisation scheme is defined in terms of the standard deviation of the noise propagated to the wavelet coefficients. While this scheme proves very effective for Gaussian noise, it is not appropriate for χ^2 distributed noise.

Therefore, in the approach presented here, we choose to use a suboptimal gradient for the quadratic data fidelity term by keeping only the first term of Equation (7.13), which is the price to pay in order to keep the very effective regularisation scheme introduced in the linear case. Note that this suboptimal gradient still corresponds to the reduced shear correction scheme suggested in Seitz and Schneider (1995) and presented in Section 6.3.1.

7.3.2 Including redshift information

So far we have not given any consideration to the fact that lensing sources are not located on a single plane but are distributed in redshift. As was described in Section 6.1.3, for a lens at a given redshift, the amplitude of the lensing effect experienced by each source will depend on its own redshift. This redshift dependence can be expressed as a simple cosmological weight $Z(z)$ (defined by Equation (6.43)) which scales convergence and shear according to:

$$\kappa(z) = Z(z)\,\kappa_\infty \quad ; \quad \gamma(z) = Z(z)\,\gamma_\infty \tag{7.14}$$

where κ_∞ and γ_∞ are the convergence and shear for sources at infinite redshift. In practice, when mapping known galaxy clusters, the redshift z_l of the lens is generally well known but the redshifts of the background galaxies used in weak lensing studies is often determined photometrically and is therefore not precisely known. In this case, we will define a lensing weight Z_i for a given source of i by marginalising over its photo-z distribution $p_i(z)$:

$$Z_i = \int Z(z) p_i(z)\,\mathrm{d}z = \int_{z_L}^{\infty} \frac{\Sigma^\infty_{crit}}{\Sigma_{crit}(z)} p_i(z)\,\mathrm{d}z\ , \tag{7.15}$$

where $\Sigma_{crit}(z)$ is the critical density for a lens at redshift z_l and a source at redshift $z_s = z$ while Σ^∞_{crit} is the critical density for sources at infinite redshift. With this definition, the reduced shear g_i and reduced flexion F_i for a given source can be described as a function of the shear, flexion and convergence for sources at infinite redshift γ_∞, F_∞ and κ_∞:

$$g_i = \frac{Z_i\gamma_\infty}{1 - Z_i\kappa_\infty} \quad ; \quad F_i = \frac{Z_i F_\infty}{1 - Z_i\kappa_\infty}\ . \tag{7.16}$$

Let us note $\mathbf{Z}$ the diagonal matrix of weights Z_i, the reduced shear can be computed from the convergence $\boldsymbol{\kappa}$ using the Fourier operators introduced thus far as:

$$\boldsymbol{g} = \frac{\boldsymbol{Z}\mathbf{TPF}^*\boldsymbol{\kappa}}{1 - \boldsymbol{Z}\mathbf{TF}^*\boldsymbol{\kappa}} \tag{7.17}$$

where $\boldsymbol{\kappa}$ is understood to be the convergence at infinite redshift. With this new operator, the full inversion problem becomes:

$$\arg\min_{\boldsymbol{\kappa}} \frac{1}{2} \parallel C_\kappa^{-1}\left[(1 - \boldsymbol{Z}\mathbf{TF}^*\boldsymbol{\kappa})\boldsymbol{g} - \boldsymbol{Z}\mathbf{TPF}^*\boldsymbol{\kappa}\right] \parallel_2^2 + \lambda \parallel \boldsymbol{w} \circ \boldsymbol{\Phi}^t\boldsymbol{\kappa} \parallel_1 + i_{\mathfrak{I}(\cdot)=0}(\boldsymbol{\kappa})\ . \tag{7.18}$$

where the matrix C_κ^{-1} now becomes $C_\kappa^{-1} = 1/(1 - \boldsymbol{Z}\mathbf{TF}^*\boldsymbol{\kappa})$. This is simply a generalisation of Equation (7.12), which can be recovered when no redshift information is available by setting $\boldsymbol{Z} = \mathrm{Id}$, and can be solved exactly in the same way.

The main advantage of using redshift estimates for individual galaxies is the proper scaling of the resulting convergence map, which can be translated into a physical surface mass density map Σ of the lens plane:

$$\Sigma(\boldsymbol{\theta}) = \kappa(\boldsymbol{\theta})\Sigma^\infty_{\mathrm{crit}}\ . \tag{7.19}$$

The mass of the lens can then be estimated by integrating Σ within a given radius. The second advantage of using individual redshifts is that it can help mitigate the mass-sheet degeneracy, as was described in Bradac et al. (2004) and presented in Section 6.1.4. We stress however that this

degeneracy cannot be completed lifted using the method presented in the previous section as, of the two terms in the gradient of the χ^2, we only retain the one that is insensitive to the mean. Therefore, the mean value of the field remains unconstrained by the data using our algorithm. Nevertheless, we expect the additional redshift information to locally break the degeneracy and help recover the correct amplitude for the most significant structures.

7.3.3 Improving angular resolution with flexion

We demonstrated in the first section that our sparse recovery algorithm is capable of reconstructing small scale details in the absence of noise despite the irregular galaxy sampling. In practice however, the shear is noise dominated on those scales, which makes recovering high frequency information from shear measurements unlikely.

As was explained in Section 6.3.2, while the Kaiser-Squires estimator for shear has a flat noise power spectrum, the noise power spectrum of the minimum variance estimator for flexion has a $1/k^2$ dependency. As a result, although flexion measurements are generally very noisy, the noise level of the estimated map eventually drops below that of the shear on sufficiently small scales. Flexion can therefore bring useful information but only below a given scale. While Pires and Amara (2010) claimed that flexion remained noisier than shear up to scales of a few arcseconds we show in Section 6.3.2 that flexion becomes competitive much closer to the arcminute scale, at around 45 arcsec (assuming intrinsic shear and flexion dispersion of $\sigma_\epsilon = 0.3$ and $\sigma_F = 0.04$ arcsec^{-1}). Therefore flexion is very complementary to shear and can help bridge the gap between the typical weak and strong lensing scales (i.e. between arcminutes and arcseconds) and allow us to map some cluster sub-structures below the arcminute scale.

Our aim is therefore to improve the mass map reconstruction on small scales by extending our reconstruction method to incorporate flexion information. In the interest of simplicity, we consider here only the linear problem without redshift information, to highlight the difficulties inherent to the inclusion of flexion. The full problem is solved in the next section.

Following the approach developed in the first section, we address the problem using Fourier estimators. We first introduce the diagonal operator $\mathbf{Q}$ implementing the transform from convergence κ to first flexion $\mathcal{F}$ in Fourier space, defined as:

$$\hat{\boldsymbol{\mathcal{F}}} = \mathbf{Q}\hat{\boldsymbol{\kappa}} = (k_2 - ik_1)\,\hat{\boldsymbol{\kappa}}\,, \tag{7.20}$$

where we use complex notations for the flexion with $\mathcal{F} = \mathcal{F}_1 + i\mathcal{F}_2$. Contrary to the shear operator $\mathbf{P}$, this operator $\mathbf{Q}$ is not unitary but is still invertible:

$$\hat{\boldsymbol{\kappa}} = \mathbf{Q}^{-1}\hat{\boldsymbol{\mathcal{F}}} = \frac{\mathbf{Q}^*}{k^2}\hat{\boldsymbol{\mathcal{F}}} = \frac{k_2 + ik_1}{k^2}\hat{\boldsymbol{\mathcal{F}}}\,. \tag{7.21}$$

To extend the algorithm presented in the previous section, a first straightforward approach would be to simply add a flexion term to the χ^2 of Equation (7.8) and solve the following problem:

$$\arg\min_{\boldsymbol{\kappa}} \frac{1}{2} \parallel \boldsymbol{\gamma} - \mathbf{TPF}^*\boldsymbol{\kappa} \parallel_2^2 + \frac{1}{2} \parallel \boldsymbol{\mathcal{F}} - \mathbf{TQF}^*\boldsymbol{\kappa} \parallel_2^2 + \lambda \parallel \boldsymbol{w} \circ \boldsymbol{\Phi}^t\boldsymbol{\kappa} \parallel_1 + i_{\Im(\cdot)=0}(\boldsymbol{\kappa})\,. \tag{7.22}$$

Although formally correct, solving this problem using the primal-dual algorithm presented in the previous section leads to a number of technical issues linked to the fact that the operator $\mathbf{Q}$ being

not unitary it now contributes to the difficulty of the inverse problem. In particular, this impacts our ability to robustly identify significant coefficients in the gradient of the data fidelity term (key to the regularisation strategy presented in the previous section), which is now affected by a mixture of a flat shear noise and a flexion noise with a power spectrum in k^2. Furthermore, even without considering the regularisation, the inversion of the operator $\mathbf{Q}$, which is essentially a 2D gradient, requires a large number of iterations if solved with a standard gradient descent. These considerations make the algorithm much slower and far less robust to noise than when solving the problem from shear alone.

This needs not be the case however as the inverse of this operator is explicit in Fourier space and these difficulties can be avoided if we make proper use of this explicit inverse. We propose therefore to address the combined shear and flexion reconstruction as the following sparse optimisation problem:

$$\arg\min_{\boldsymbol{\kappa},\tilde{\boldsymbol{\mathcal{F}}}} \frac{1}{2} \parallel \boldsymbol{\gamma} - \mathbf{TPF}^*\boldsymbol{\kappa} \parallel_2^2 + \frac{1}{2} \parallel \boldsymbol{\mathcal{F}} - \mathbf{TF}^*\tilde{\boldsymbol{\mathcal{F}}} \parallel_2^2 + \lambda \parallel \boldsymbol{w} \circ \boldsymbol{\Phi}^t \boldsymbol{\kappa} \parallel_1 + i_{\mathrm{Im}(\mathbf{R})}\left(\begin{bmatrix}\boldsymbol{\kappa}\\ \tilde{\boldsymbol{\mathcal{F}}}\end{bmatrix}\right) \tag{7.23}$$

where we introduce the application $\mathbf{R} : \mathbb{R}^{N\times N} \to \mathbb{C}^{2N\times N},\ \boldsymbol{\kappa} \mapsto \begin{bmatrix}\boldsymbol{\kappa}\\ \mathbf{FQF}^*\boldsymbol{\kappa}\end{bmatrix}$, with $N \times N$ the size of the reconstruction grid.

Thanks to the inclusion of the auxiliary variable $\tilde{\boldsymbol{\mathcal{F}}}$ we have now decoupled the problem of the inversion of the NDFT operator $\mathbf{T}$ from the conversion between flexion and convergence which is now addressed implicitly in the last term. Remember that the indicator function i_C is infinite outside of the set C and therefore exclude any solution which do not belong to C. In our case, we require the solution to be in the image of the operator $\mathbf{R}$:

$$\mathrm{Im}(\mathbf{R}) = \left\{ \begin{bmatrix}\boldsymbol{\kappa}\\ \tilde{\boldsymbol{\mathcal{F}}}\end{bmatrix} \in \mathbb{C}^{2N\times N} \quad | \quad \exists \boldsymbol{\kappa} \in \mathbb{R}^{N\times N}, \quad \tilde{\boldsymbol{\mathcal{F}}} = \mathbf{FQF}^*\boldsymbol{\kappa} \right\} \tag{7.24}$$

The constraint $i_{\mathrm{Im}(\mathbf{R})}$ therefore implies two conditions. First, the recovered convergence has to be real, which is equivalent to enforcing the vanishing B-modes condition already used for the shear alone inversion problem. The second is that the recovered flexion $\tilde{\boldsymbol{\mathcal{F}}}$ needs to match the flexion derived from the recovered convergence, which makes the connection between the two variables $\boldsymbol{\kappa}$ and $\tilde{\boldsymbol{\mathcal{F}}}$. In order to use the same primal-dual algorithm as in the previous section, we need the expression of the proximity operator of $i_{\mathrm{Im}(\mathbf{R})}$, which can be explicitly computed from its definition:

$$\mathrm{prox}_{i_{\mathrm{Im}(\mathbf{R})}}\left(\begin{bmatrix}\boldsymbol{\kappa}\\ \tilde{\boldsymbol{\mathcal{F}}}\end{bmatrix}\right) = \arg\min_{(\boldsymbol{x}_\kappa,\boldsymbol{x}_\mathcal{F})} \frac{1}{2} \parallel \begin{bmatrix}\boldsymbol{\kappa}\\ \tilde{\boldsymbol{\mathcal{F}}}\end{bmatrix} - \begin{bmatrix}\boldsymbol{x}_\kappa\\ \boldsymbol{x}_\mathcal{F}\end{bmatrix} \parallel_2^2 + i_{\mathrm{Im}(\mathbf{R})}\left(\begin{bmatrix}\boldsymbol{x}_\kappa\\ \boldsymbol{x}_\mathcal{F}\end{bmatrix}\right) \tag{7.25}$$

$$= \mathbf{R} \arg\min_{\boldsymbol{x}_\kappa \in \mathbb{R}^{N\times N}} \parallel \begin{bmatrix}\boldsymbol{\kappa}\\ \tilde{\boldsymbol{\mathcal{F}}}\end{bmatrix} - \begin{bmatrix}\boldsymbol{x}_\kappa\\ \mathbf{FQF}^*\boldsymbol{x}_\kappa\end{bmatrix} \parallel_2^2 \tag{7.26}$$

$$= \mathbf{R}\Re\left(\mathbf{F}\frac{1}{k^2 + \frac{\sigma_\mathcal{F}^2}{\sigma_\epsilon^2}}\left(k^2\mathbf{Q}^{-1}\mathbf{F}^*\tilde{\boldsymbol{\mathcal{F}}} + \frac{\sigma_\mathcal{F}^2}{\sigma_\epsilon^2}\mathbf{F}^*\boldsymbol{\kappa}\right)\right). \tag{7.27}$$

In this expression, σ_ϵ^2 is the variance of the intrinsic ellipticity and $\sigma_{\mathcal{F}}^2$ is the variance of the intrinsic flexion. Let us detail what is computed by this proximity operator. We first use the definition of $\mathrm{Im}(\mathbf{R})$ to rewrite the minimisation problem in terms of a single unknown, from which the output convergence and flexion can be computed by applying the operator $\mathbf{R}$. Then, remembering that in Equation (7.23) $\boldsymbol{\kappa}$ is fitted to the shear, the problem is equivalent to finding the minimum variance filter combining shear and flexion, developed in Appendix B. To summarise, this operator computes the optimal combination of the two input variables, making explicit use of the flexion Fourier operator and its inverse $\mathbf{Q}^{-1} = \mathbf{Q}^*/k^2$, thus eliminating the need to solve this additional problem as part of the main minimisation problem.

Furthermore, in the primal-dual algorithm used throughout this chapter, this proximity operator is applied to the gradient of the data fidelity term before updating the dual variable $\boldsymbol{\alpha}$. This means that only the minimum variance combination of both shear and flexion gradient terms is propagated to the wavelet coefficients, which can only increase the SNR of the signal, especially on small scales where flexion is the most useful (see the noise power spectrum of the combined estimator in Figure 6.8).

7.3.4 Complete density mapping algorithm

We now present the complete reconstruction algorithm combining reduced shear $\boldsymbol{g}$ and reduced flexion $\boldsymbol{F}$, taking into account individual redshift estimates for the sources. The complete problem we aim to solve is the following:

$$\begin{aligned}\arg\min_{\boldsymbol{\kappa},\tilde{\boldsymbol{\mathcal{F}}}} \frac{1}{2} \parallel C_\kappa^{-1}\left[(1-\boldsymbol{Z}\mathbf{T}\mathbf{F}^*\boldsymbol{\kappa})\boldsymbol{g} - \boldsymbol{Z}\mathbf{T}\mathbf{P}\mathbf{F}^*\boldsymbol{\kappa}\right] \parallel_2^2 + \frac{1}{2} \parallel C_\kappa^{-1}\left[(1-\boldsymbol{Z}\mathbf{T}\mathbf{F}^*\boldsymbol{\kappa})\boldsymbol{F} - \boldsymbol{Z}\mathbf{T}\mathbf{F}^*\tilde{\boldsymbol{\mathcal{F}}}\right] \parallel_2^2 \\ + \lambda \parallel \boldsymbol{w} \circ \boldsymbol{\Phi}^t\boldsymbol{\kappa} \parallel_1 + i_{\mathrm{Im}(\mathbf{R})}\left(\begin{bmatrix}\boldsymbol{\kappa}\\ \tilde{\boldsymbol{\mathcal{F}}}\end{bmatrix}\right).\end{aligned} \tag{7.28}$$

This problem can be solved with the same primal-dual algorithm from Vu (2013) as in the first section of this chapter and the specialisation of this algorithm to Equation (7.28) is provided in Algorithm 7.2.

Just as with the linear problem, we apply a reweighted-ℓ_1 strategy to correct for the bias caused by the ℓ_1 sparsity constraint. As the non-linear correction also requires to iteratively solve this problem we combine the update of the weights $\boldsymbol{w}$ and the update of the matrix C_κ in the following iterative procedure:

1. Set the iteration count $\ell = 0$, $C_\kappa^{(0)} = 1.0$ and initialise the weights $\boldsymbol{w}^{(0)}$ according to the procedure described in Section 7.2.4.
2. Solve the weighted ℓ_1 minimisation problem of Equation (7.28) using Algorithm 7.2, yielding a solution $\boldsymbol{\kappa}^{(\ell)}$.
3. Update the matrix $C_\kappa^{(\ell)} = (1 - \boldsymbol{Z}\mathbf{T}\mathbf{F}^*\boldsymbol{\kappa}^{(l)})$
4. Update the weights based on the wavelet transform of the solution $\boldsymbol{\alpha}^{(\ell)} = \Phi^t\boldsymbol{\kappa}^{(\ell)}$:

$$w_i^{(\ell+1)} = \begin{cases} \frac{w_i^{(0)}}{|\alpha_i^{(\ell)}|/\lambda w_i^{(0)}} & \text{if } |\alpha_i^{(\ell)}| \geq \lambda w_i^{(0)} \\ w_i^{(0)} & \text{if } |\alpha_i^{(\ell)}| < \lambda w_i^{(0)} \end{cases}, \tag{7.29}$$

Algorithm 7.2 Analysis-based density mapping algorithm from reduced shear and flexion

Require:

Reduced shear and flexion of each galaxy in the survey $\boldsymbol{g}$ and $\boldsymbol{F}$.
Redshift weights for each galaxy in the survey $\mathbf{Z}$.
Reduced shear correction matrix C_κ^{-1}.
Sparsity constraint parameter $\lambda > 0$.
Weights $w_i > 0$.
$\tau = 2/(\| \mathbf{\Phi} \|^2 + \| \mathbf{T} \|^2)$.

1: $\boldsymbol{\kappa}^{(0)} = 0$; $\tilde{\boldsymbol{\mathcal{F}}}^{(0)} = 0$
2: $\forall i, \quad \lambda_i' = \lambda w_i$
3: **for** $n = 0$ to $N_{\max} - 1$ **do**
4: $\quad \nabla^{(n)} = \mathbf{F}\mathbf{P}^*\mathbf{T}^*\mathbf{Z}C_\kappa^{-2}\left((1 - \mathbf{Z}\mathbf{T}\mathbf{F}^*\boldsymbol{\kappa}^{(n)})\boldsymbol{g} - \mathbf{Z}\mathbf{T}\mathbf{P}\mathbf{F}^*\boldsymbol{\kappa}^{(n)}\right) + \mathbf{F}\mathbf{T}^*\mathbf{Z}C_\kappa^{-2}\left((1 - \mathbf{Z}\mathbf{T}\mathbf{F}^*\boldsymbol{\kappa}^{(n)})\boldsymbol{F} - \mathbf{Z}\mathbf{T}\mathbf{F}^*\tilde{\boldsymbol{\mathcal{F}}}^{(n)}\right)$
5: $\quad \begin{pmatrix} \boldsymbol{\kappa}^{(n+1)} \\ \tilde{\boldsymbol{\mathcal{F}}}^{(n+1)} \end{pmatrix} = \mathrm{prox}_{\mathrm{Im}(\mathbf{R})}\left(\begin{pmatrix} \boldsymbol{\kappa}^{(n)} \\ \tilde{\boldsymbol{\mathcal{F}}}^{(n)} \end{pmatrix} + \tau\left(\nabla^{(n)} - \mathbf{R}\mathbf{\Phi}\boldsymbol{\alpha}^{(n)}\right)\right)$
6: $\quad \boldsymbol{\alpha}^{(n+1)} = (\mathrm{Id} - \mathrm{ST}_{\lambda'})\left(\boldsymbol{\alpha}^{(n)} + \mathbf{\Phi}^t\left(2\boldsymbol{\kappa}^{(n+1)} - \boldsymbol{\kappa}^{(n)}\right)\right)$
7: **end for**
8: **return** $\kappa^{(N_{\max})}$.

5. Terminate on convergence. Otherwise, increment ℓ and go to step 2.

As for the linear problem we find that 3 to 5 iterations of this procedure are generally sufficient to reach a satisfying solution.

7.4 VERIFICATION ON SIMULATIONS

As was mentioned in the introduction, the most advanced reconstruction technique combining shear, flexion and strong lensing yet is the method proposed in Cain et al. (2015). Therefore, in order to compare our method to state of the art results, we present in this section a set of simulations intended to mimic their test data set. However, while their simulations rely on simple analytic profiles we chose to use realistic cluster profiles extracted from N-body simulations.

7.4.1 Simulations

In order to assess the performance of the algorithm, we created a realistic test data set based on N-body simulations. Three massive clusters were extracted at z=0 from the Bolshoi simulations (Klypin et al., 2011) using the CosmoSim[2] interface. These clusters were selected because of their complex geometry with substantial substructure. In each case, at least one halo above $10^{13}h^{-1}\mathrm{M}_\odot$ can be found within the virial radius of the central halo. The masses of the three clusters we consider in this work can found in Table 4.

2 http://www.cosmosim.org

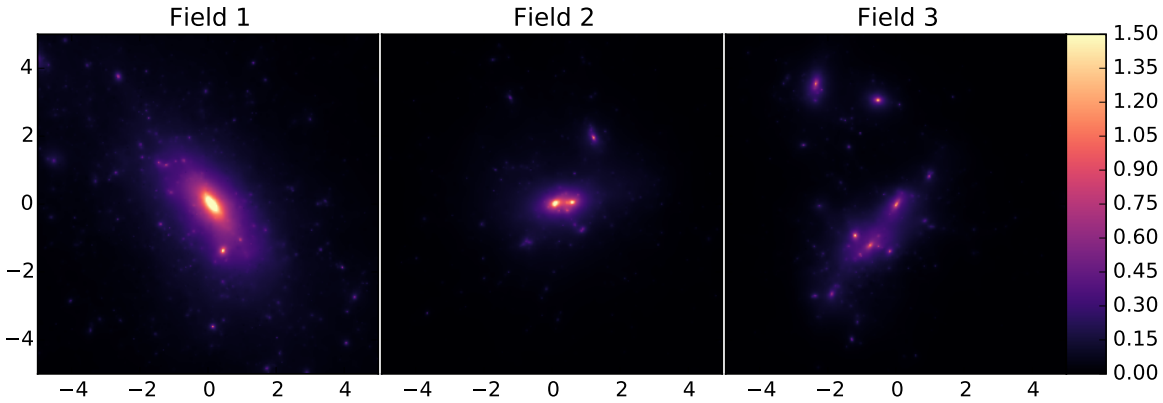

Figure 7.6: Convergence maps for sources at infinity for the three different clusters extracted from the N-body simulation.

Field number	Virial mass $[h^{-1}\mathrm{M}_\odot]$	Virial radius $[h^{-1}\mathrm{Mpc}]$
1	1.09×10^{15}	2.14
2	3.02×10^{14}	1.43
3	2.70×10^{14}	1.45

Table 4: Parameters of the three halos extracted from the Bolshoi simulation.

Density maps were obtained for these three clusters by projecting and binning the particle data with a high resolution, corresponding to pixels with an angular size of 0.5 arcsec if the clusters are located at redshift $z_L = 0.3$. A multiscale Poisson denoising was subsequently applied to the binned density maps to remove shot noise from the numerical simulation. The resulting surface density maps were then scaled to create convergence maps corresponding to clusters at a redshift $z_L = 0.3$ lensing background galaxies at infinite redshift. Although we acknowledge that these three particular clusters, selected at $z = 0$, are not necessarily representative of clusters at $z_L = 0.3$, for the purpose of testing the mapping algorithm they provide more realistic density distributions than simple models based on SIS or NFW profiles. Shear and flexion maps were then derived from these reference convergence maps and only the center 10×10 arcmin central region of each field is kept. At the redshift of the lenses, this corresponds to a physical size of 1.88×1.88 h^{-1}Mpc. Figure 7.6 illustrates the convergence maps scaled for sources at infinite redshift for the 3 fields we consider.

Finally, 100 mock galaxy catalogues were produced for each of the three fields using for each realisation a uniform spatial distribution of background galaxies with a density of 80 gal/arcmin^2 and the following redshift distribution:

$$p(z) \propto \frac{z^2}{2z_0^3} \exp(-z/z_0) \tag{7.30}$$

with $z_0 = 2/3$. The median redshift of this distribution is $z_{med} = 1.75$ and we truncate the distribution at $z_{max} = 5$ to exclude very distant faint galaxies. This particular distribution has been used in a number of different works and in particular in Cain et al. (2015) to represent the actual galaxy distribution for a typical HST/ACS field. We compute the reduced shear and flexion for each galaxy

based on their redshift and on the redshift of the lens which we choose to be $z_L = 0.3$ for the 3 different fields. In the final mock catalogues we assume photometric redshifts for each sources with Gaussian redshifts errors according to $\sigma_z = 0.05(1+z)$ and an intrinsic shape noise of $\sigma_\epsilon = 0.3$ for the reduced shear measurements and $\sigma_F = 0.029$ arcsec^{-1} for the reduced flexion measurements. This particular value for the flexion noise is in accordance with previous works (Rowe et al., 2013) and corresponds to the median dispersion of the flexion measurements obtained using the AIM method on HST data for the Abel 1689 cluster (Cain et al., 2011).

Some final cuts are applied to the galaxy catalog to exclude strongly lensed sources which our algorithm is not capable to handle and for which the flexion measurements are no longer reliable. In practice we exclude all sources which verify $|g| \geq 1$ or $|F| \geq 1.0$ arcsec^{-1}. The flexion cut is based on recommendations from Cain et al. (2015) and constitutes a simple approximation to the practical limit encountered on real data when estimating flexion for extremely lensed sources.

We have chosen this setting to match the simulations used in Cain et al. (2015) so that our results can be qualitatively compared to theirs. However, we stress that contrary to their work, our reconstruction method relies only on shear and flexion and does not include strong lensing constraint which are very powerful to map the very center of the clusters.

Throughout this section, we assume a fiducial ΛCDM model for computing distances with $\Omega_m = 0.25$, $\Omega_\Lambda = 0.75$ and $H_0 = 70$ km/s/Mpc. Cosmological computations are implemented using the NICAEA library[3].

7.4.2 Results

As demonstrated in the first section, the algorithm is extremely efficient at solving the noiseless linear problem. With the realistic simulations described in Section 7.4.1 we can now assess the performance of the complete algorithm described in Section 7.3.4 when noise and non-linear effects are included.

Parameter	Value
Pixel size	0.05 arcmin
Number of wavelet scales	7
$K-$sigma threshold	5
Number of iterations	500
Number of re-weightings	5

Table 5: Parameters of the reconstruction algorithm

For each of the mock catalogues generated for each field, we perform the inversion with and without the flexion information to assess how much can the flexion help to recover the substructure of the clusters. The same parameters are used for all the three different fields and are summarised in Table 5. The pixel size can in theory be arbitrarily small but we choose 0.05 arcmin as a good compromise between resolution and computational cost. The number of wavelet scales is not crucial to the quality of the reconstruction, we adjust it so that the maximum scale corresponds

3 http://www.cosmostat.org/software/nicaea

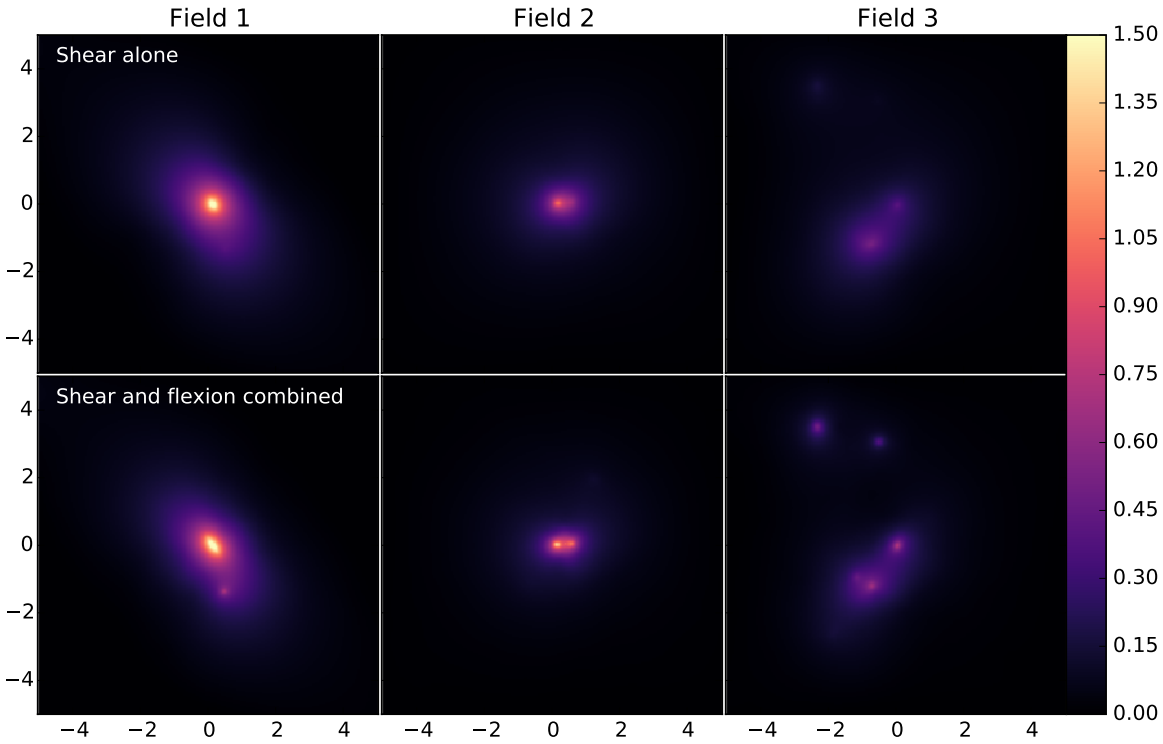

Figure 7.7: Mean of the 100 independent signal realisations for the 3 clusters. The top row corresponds to reconstructions using the shear alone. The bottom row corresponds to simulations using shear and flexion.

roughly to the order of the maximum size of the structures we want to recover. Finally, we use a sufficient number of iterations to ensure that the algorithm has converged.

Assessing the quality of the inversion from a single realisation of the galaxy catalog is difficult as the ability to detect small structures will depend on the specific positions of the galaxies in one realisation. Therefore, we compute the mean and standard deviation of the reconstructed maps over 100 realisation of both shape noise and galaxy positions.

The mean of the reconstructions for the 3 fields with and without flexion is represented in Figure 7.7. The most striking difference between the 2 sets of reconstructions is the flexion's ability to detect very small substructures at the 10 arcsec scale which are not detected from shear alone. This is clearly visible for the first halo, on the left column of Figure 7.7.

Flexion does not only enable the detection of very small structures, it also helps to constrain the small scale shape of the main halos. Indeed, in shear reconstructions, the noise dominates the signal at these scales which makes the sparsity prior apparent on small scales. As we are using isotropic wavelets, without sufficient evidence from the data, the reconstruction will be biased towards isotropic shape. This is visible for instance at the center of the first halo, on the top left image of Figure 7.7, where the elongation of the very center of the cluster, visible in Figure 7.6, is clearly lost. On the contrary, including flexion information helps to constrain the shape of the center of the cluster.

We also note that on large scales, where the shear dominates, the two sets of reconstructions agree with each other as flexion does not add information on these scales. This is illustrated by Table 6 where the aperture mass within a 2′′ radius is computed on the reconstruction with and

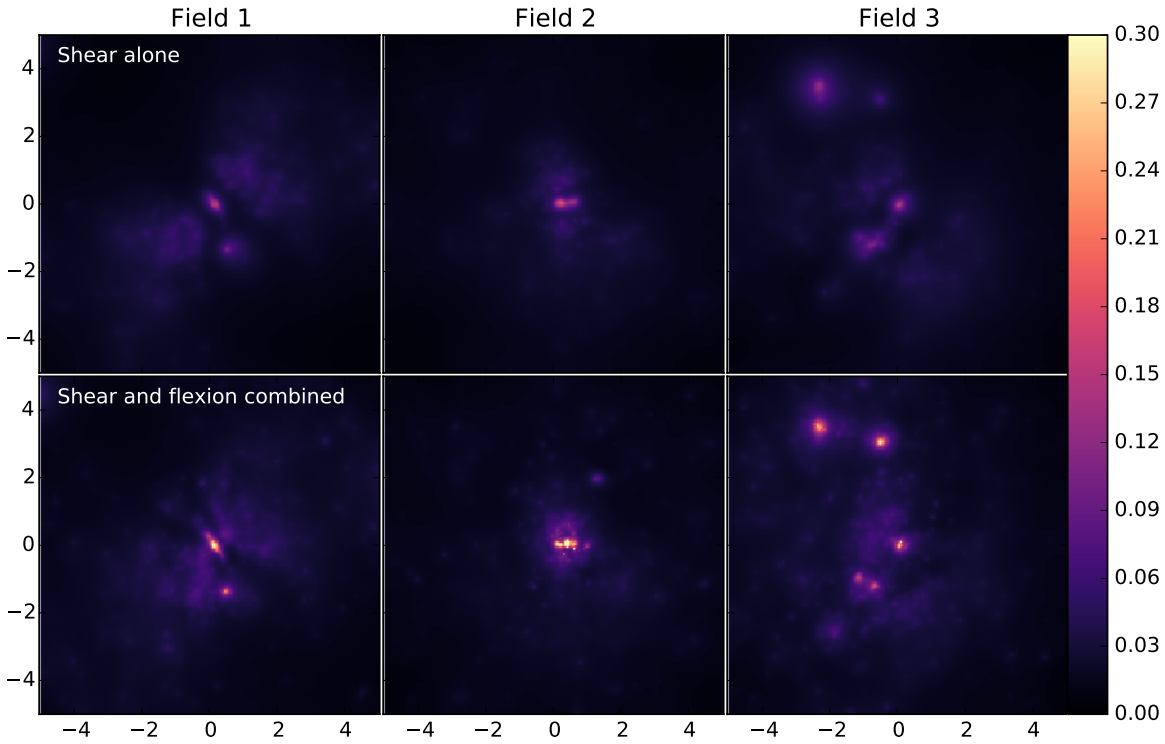

Figure 7.8: Standard deviation of the 100 independent signal realisations for the 3 clusters. The top row corresponds to reconstructions using the shear alone. The bottom row corresponds to simulations using shear and flexion.

without flexion. As can be seen, tight constraints on the halo masses can be obtained for the three fields considered from shear only and the addition of flexion does not change the mass estimates on scales larger than the arcminute.

Field	Input $M_{\rm ap}$ $10^{14}\ h^{-1}\mathrm{M}_\odot$	Shear only $M_{\rm ap}$ $10^{14}\ h^{-1}\mathrm{M}_\odot$	Combined $M_{\rm ap}$ $10^{14}\ h^{-1}\mathrm{M}_\odot$
1	4.08	3.91 ± 0.18	3.89 ± 0.17
2	1.88	1.90 ± 0.20	1.86 ± 0.18
3	1.59	1.60 ± 0.18	1.61 ± 0.18

Table 6: Aperture mass in the central 2′′ of each field.

The errors on the reconstructed mass maps are estimated by taking the standard deviation of each pixel in the 100 independent mock catalogues realisation. It is important to stress that the observed dispersion of the reconstructions are due to 2 effects: different noise realisations and different galaxy distributions in angular position and redshift. The standard deviation with and without flexion for all 3 fields is illustrated in Figure 7.8. As can be seen on this figure, we only observe significant dispersion on small scales, around detected structures. This is expected as these scales have the lowest SNR in the data which makes the reconstruction of these features strongly dependent on the specific noise and galaxy distribution realisation. Nevertheless, we find that the reconstruction is remarkably stable between realisations. Figure 7.9 shows reconstructions of the

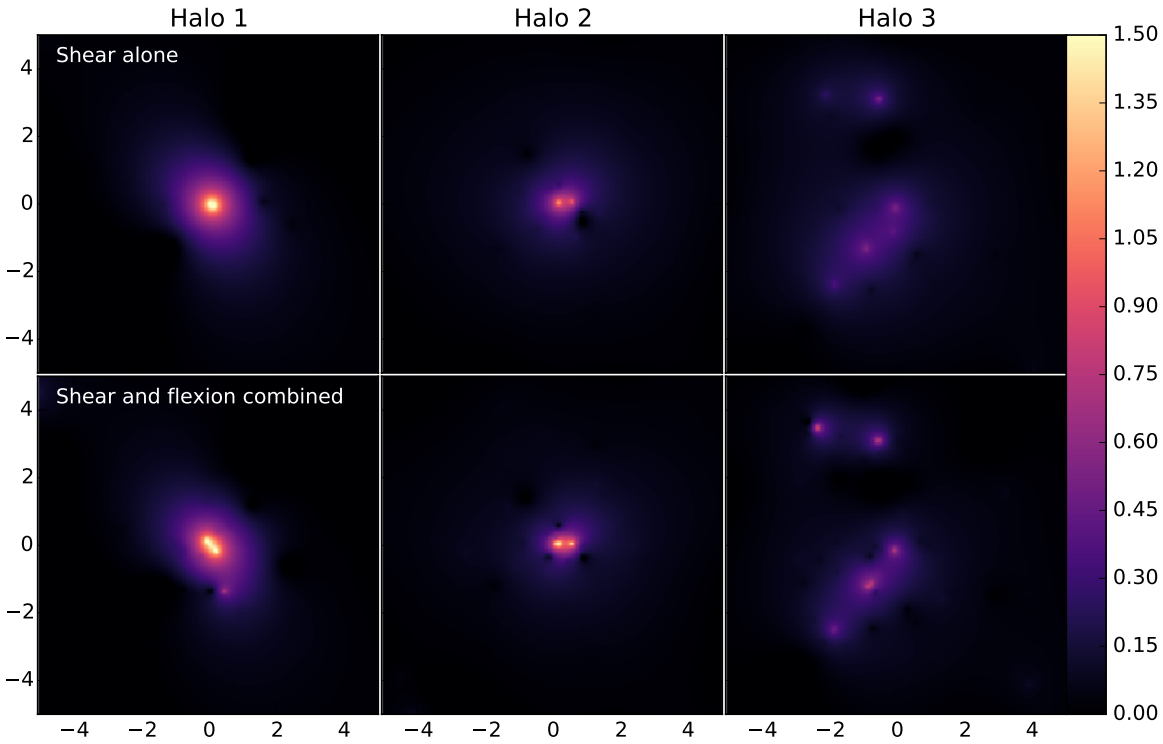

Figure 7.9: Reconstruction of the 3 clusters for one noise realisation. The top row corresponds to reconstructions using the shear alone. The bottom row corresponds to simulations using shear and flexion.

mass maps for one realisation of the galaxy catalogues. We see that for a single realisation the reconstructed maps is very close to the mean maps. One small feature is not detected in field 2 whereas it appears on the mean map while on the contrary the sub-halo at the bottom of field 3 appears stronger in this particular realisation than in the average maps. Such deviations are unavoidable but the shape of the halos can be reliably reconstructed from a single data set, as would be the case for actual data.

Finally, we compare the azimuthally averaged radial density profiles for the three different fields. Figure 7.10 shows the mean and 68% confidence level of the reconstructed profiles for the three fields with and without flexion. Again, the mean and dispersion are evaluated from the same 100 independent catalogue realisations. The reconstructed profiles show very little dispersion and the shape of the profile is accurately recovered in all cases. The impact of flexion is particularly obvious on field 3 where the small scale structure of the central region of the field cannot be recovered accurately from shear alone.

7.5 CONCLUSION

We have developed in this chapter a new 2D mass-mapping technique, based on sparse regularisation, combining shear and flexion to reconstruct the surface mass density of galaxy clusters. We find that the inclusion of flexion significantly improves the quality of the reconstruction on small scales. In particular, on simulated lensing catalogues corresponding to a typical HST/ACS

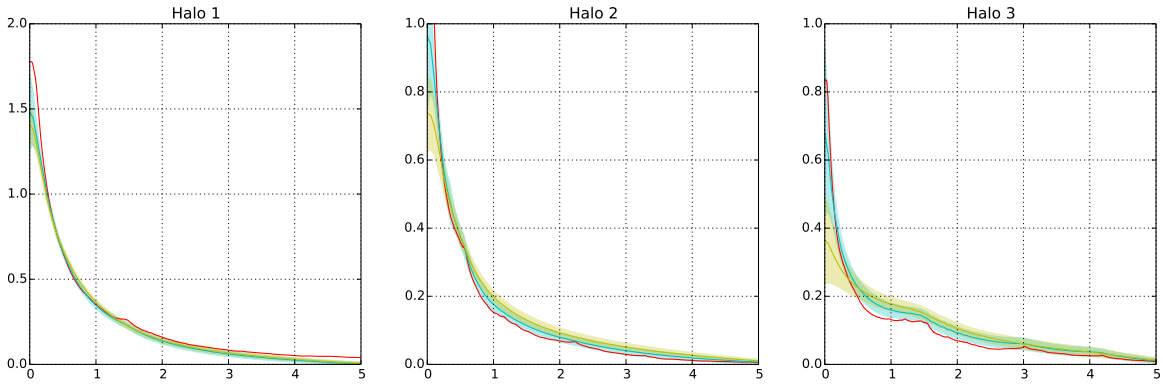

Figure 7.10: Comparison of the convergence profile recovered using shear alone (yellow line) and flexion and shear combined (blue line). Contours indicate the 1-sigma region.

cluster survey we are able to recover cluster substructures below the 10 arcsecond scale using weak lensing alone (i.e. shear and flexion), without any strong lensing information.

This technique is particularly promising for mapping the substructure of galaxy clusters outside of their Einstein radius or even when the clusters are non-critical.

8 SPARSITY BASED 3D WEAK LENSING MASS-MAPPING

CONTENTS

In the previous chapter we have focused on 2D weak lensing mass mapping in the case where information about the redshift of the lens and sources is available. Then, a density map of the lens can be recovered by scaling the reconstructed convergence map thanks to the knowledge of the relative distances between lens and sources. This density map can in turn be used to directly measure the mass of the clusters. However, prior knowledge of the lens redshift (available for instance for spectroscopically confirmed clusters) is not always available.

In this chapter, we consider a different situation where only the redshift of the sources is known. The aim is to reconstruct in three dimensions the dark matter map from gravitational lensing without external information about the position of the structures in redshift. Just as the mapping problems considered in the previous chapter, this problem still corresponds to a linear inverse problem but with an additional major degeneracy along the redshift dimension which makes constraining the redshift of the detected structures the main difficulty of 3D lensing.

We present a new algorithm for 3D mass mapping, called GLIMPSE based on the sparse regularisation framework of Chapter 3 and using wavelets introduced in Chapter 4. We compare the results of our algorithm with the state of the art methods introduced in Section 6.4 and find that we successfully overcome the main limitations of theses methods. Finally we thoroughly test the algorithm on a large set of simple simulations to demonstrate its ability to recover the redshifts and even the masses of galaxy clusters.

The results presented in this chapter were published in Leonard et al. (2014, 2015). This work has been conducted in collaboration with Adrienne Leonard and Jean-Luc Starck.

8.1 THE GLIMPSE ALGORITHM

8.1.1 Forward modelling of the 3D lensing problem

The first step in setting up the 3D reconstruction problem is to link the 3D dark matter density to the measured ellipticities. A standard tomographic binning approach, similar to Simon et al. (2009), was adopted to introduce a discrete lensing operator.

We consider that we have a catalog of measured ellipticities $\boldsymbol{\epsilon}_k$ from galaxies at angular positions $\boldsymbol{\theta}_k$ and photometric redshifts z_k. For the purpose of reconstructing the 3D density contrast, this catalog can be binned into N_z tomographic shear maps $\boldsymbol{\gamma}^{(1)}, \ldots, \boldsymbol{\gamma}^{(n)}$ of size $N_x \times N_y$ where N_x and N_y are the number of pixels on the 2D cartesian grid in the angular domain. On each pixel, the shear value is estimated by averaging the ellipticities of the galaxies belonging to this pixel:

$$\boldsymbol{\gamma}_{i,j}^{(n)} = \frac{\sum\limits_{k \in \mathbb{G}_{i,j}^{(n)}} w_k \boldsymbol{\epsilon}_k}{\sum\limits_{k \in \mathbb{G}_{i,j}^{(n)}} w_k} , \tag{8.1}$$

where w_k is a statistical weight for the measurement on galaxy k and $\mathbb{G}_{i,j}^{(n)}$ contains the indices of the galaxies belonging to the pixel (i, j) of the redshift bin n.

Previous 3D weak lensing works (VanderPlas et al., 2011; Simon et al., 2009; Bacon and Taylor, 2003) have all chosen to bin their shear catalogues in redshift bins that are of equal width Δz. We chose to adopt a different approach, and to bin our shear maps adaptively such that each redshift bin contains the same mean number of galaxies, and therefore each redshift bin has the same noise variance. However, the specific choice of binning strategy seem to have little impact on the quality of the reconstruction.

Our goal is now to establish a forward linear model to relate the 3D overdensity distribution to the measured tomographic shear. We can first compute for a given 3D distribution $\boldsymbol{\delta}$ the corresponding tomographic convergence by means of the lensing kernel. To integrate photometric redshift uncertainties into our formalism, sources are assumed to be distributed in redshift according to some probability distribution function $p(z)$. For a given density constrast $\boldsymbol{\delta}$, the effective convergence $\kappa^{(n)}$ for sources in bin n can be obtained by marginalising the lensing kernel over the distribution of sources:

$$\kappa_{i,j}^{(n)} = \frac{3H_0^2\Omega_M}{2c^2} \int_0^{r_h} \mathrm{d}r \; g^{(n)}(r) f_K(r) \frac{\delta(f_K(r)\boldsymbol{\theta}_{i,j})}{a(r)} , \tag{8.2}$$

with

$$g^{(n)}(r) = \int_r^{r_h} \mathrm{d}r' \frac{f_K(r'-r)}{f_K(r')} \left\{ p^{(n)}(z) \frac{\mathrm{d}z}{\mathrm{d}r} \right\}_{z=z(r')} . \tag{8.3}$$

To model these photometric redshifts errors we assume a common Gaussian distribution with a standard deviation which varies as a function of redshift as $\sigma_z(1+z)$ and a potential bias z_{bias}, but without catastrophic failures. Following Ma et al. (2006), the true redshift distribution $p^{(n)}(z)$ of a tomographic bin n can be obtained through:

$$p^{(n)}(z) = \frac{1}{2} p(z) \left[\mathrm{erf}(x_{n+1}) - \mathrm{erf}(x_n) \right] \tag{8.4}$$

with

$$x_n = \frac{z_{n+1} - z + z_{bias}}{\sqrt{2}\sigma_z(1+z)}. \tag{8.5}$$

We can now discretise the lensing kernel by binning the overdensity distribution into N_{lp} lens planes $\delta^{(n)}$. Equation (8.2) can be approximated by a discrete sum over lens planes along a given line of sight:

$$\kappa^{(n)}_{i,j} = \sum_{p=1}^{N_{lp}} Q_{np}\delta^{(p)}_{i,j}\ , \tag{8.6}$$

where $\mathbf{Q}$ is now a matrix operator defined by the following matrix elements:

$$Q_{np} = \frac{3H_0^2\Omega_M}{2c^2}\int_{r_p}^{r_{p+1}} \mathrm{d}r\ g^{(n)}(r)\frac{f_K(r)}{a(r)}\ . \tag{8.7}$$

Given the tomographic convergence, the corresponding tomographic shear can be computed through a 2D convolution (see the Kaiser-Squires method in Section 6.3.1). We introduce the forward 2D lensing operator $\mathbf{P}$ which preforms this inversion in the Fourier domain independently for each redshift bin so that:

$$\gamma^{(n)} = \mathbf{P}\kappa^{(n)} \tag{8.8}$$

The complete 3D lensing problem can now be stated in terms of the discretized overdensity and tomographic shear as:

$$\boldsymbol{\gamma} = \mathbf{PQ}\boldsymbol{\delta} + \boldsymbol{n} \tag{8.9}$$

where $\boldsymbol{n}$ is an additive shape noise assumed to be Gaussian with diagonal covariance matrix Σ.

8.1.2 Sparse reconstruction algorithm

The GLIMPSE algorithm aims to solve the ill-posed inverse problem in Equation (8.9) by recasting it as an instance of a BPDN problem. This class of problems, introduced in Section 3.2.1, aims at finding a solution to the problem which can be synthesised from a sparse set of atoms from an adapted dictionary. In the specific instance of 3D lensing, using the operators introduced in the previous section, this optimization problem can be written as:

$$\arg\min_{\boldsymbol{\alpha}} \frac{1}{2} \parallel \Sigma^{-1/2}\left[\boldsymbol{\gamma} - \mathbf{PQ}\Phi\boldsymbol{\alpha}\right] \parallel_2^2 + \lambda \parallel \boldsymbol{\alpha} \parallel_1 , \tag{8.10}$$

where $\boldsymbol{\alpha}$ are the coefficients of the density contrast $\boldsymbol{\delta}$ in the dictionary Φ and Σ is the diagonal covariance matrix of the shear measurements.

The choice of the dictionary Φ is a crucial step to ensure the success of the sparse recovery. For our application, it is important to use a dictionary adapted to the morphology of the 3D distribution of dark matter. Many experiments, in particular cosmological N-body simulations, have shown the dark matter to be largely distributed in halos connected by thin filaments. Our primary interest will be to map the halos as the filamentary part of the distribution is much fainter; however we note that the method presented here is entirely general, and other dictionaries can be added to our algorithm to detect other types of structure with no modification to the overall approach (e.g. ridgelets or curvelets to detect filamentary structures).

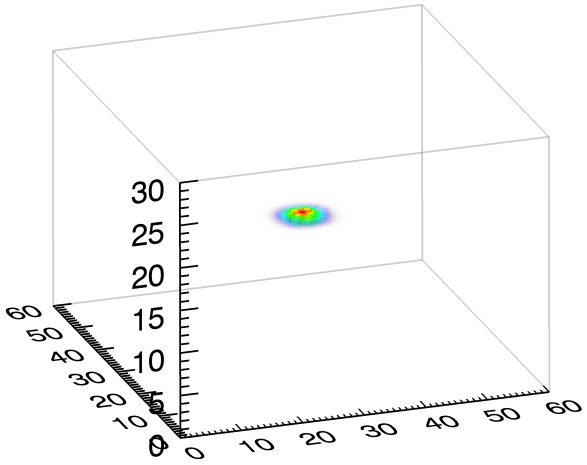

Figure 8.1: Representation of single atom of our 2D-1D dictionary composed of Isotropic Undecimated Wavelet in the angular domain and Diracs in the radial domain

At the redshift resolution we are able to attain, clusters of galaxies are small compared to the radial length of a given redshift bin. Therefore, a dark matter halo can be considered to have no radial depth and can be represented as a flat disc. This prompts us to choose a 2D-1D dictionary $\mathbf{\Phi}$ composed of Isotropic Undecimated Wavelets in the 2D angular plane and Dirac δ-functions along the radial dimension. The wavelet transform used in this work is the second generation Starlet (Starck et al., 2007) detailed in Section 4.3.3. This transform is particularly well suited to represent positive, isotropic objects. A single atom of dictionary $\mathbf{\Phi}$ is illustrated in Figure 8.1.

The optimisation problem (8.10) is relatively simple as a combination of a differentiable quadratic term and an ℓ_1 regularisation term with explicit proximity operator. In theory, it can be efficiently solved using the FISTA algorithm (Beck and Teboulle, 2009) which has been introduced in Section 3.3.2.1. The main iteration of the algorithm would therefore be:

$$X_{n+1} = \mathrm{ST}_{\lambda}\left(\alpha_n + \mu \mathbf{\Phi}^t \mathbf{Q}^t \mathbf{P}^t_{\gamma\kappa} \Sigma^{-1}\left[\gamma - \mathbf{P}_{\gamma\kappa}\mathbf{Q}\mathbf{\Phi}\alpha_n\right]\right) , \tag{8.11}$$

where ST_{λ} is the Soft Thresholding operator introduced in Chapter 3 and μ is the gradient descent step ensuring convergence of the algorithm. X_{n+1} is an updated estimation of the coefficients α of the reconstruction. The particularity of the FISTA algorithm is that the actual update of the coefficients requires an additional step amounting to a simple weighted average between current and previous estimates:

$$t_n = \frac{1+\sqrt{1+4t_n^2}}{2} , \tag{8.12}$$

$$\alpha_{n+1} = X_{n+1} + \left(\frac{t_n - 1}{t_{n+1}}(X_{n+1} - X_n)\right) , \tag{8.13}$$

with $t_0 = 1$ and $X_0 = \alpha'_0$. This algorithm converges as long as the gradient descent step μ verifies

$$0 < \mu < \frac{1}{\| \mathbf{\Phi}^t \mathbf{Q'}^t \mathbf{P}^t_{\gamma\kappa} \Sigma^{-1} \mathbf{P}_{\gamma\kappa} \mathbf{Q'} \mathbf{\Phi} \|} , \tag{8.14}$$

where $\| \cdot \|$ is the spectral norm of the operator.

In practice however, the line of sight operator $\mathbf{Q}$, the inversion of which constitutes the major difficulty in 3D lensing, is extremely badly conditioned, leading to a prohibitively slow convergence of the gradient descent in the FISTA algorithm.

To circumvent this problem, we apply a change of variable which will allow us to normalise the columns of the $\mathbf{Q}$ matrix. We define a new operator $\mathbf{Q}' = \mathbf{Q}\mathcal{N}^{-\frac{1}{2}}$ where $\mathcal{N}$ is a diagonal matrix with diagonal elements equal to the square of the ℓ_2 norm of the corresponding column of the matrix $\mathbf{Q}$:

$$\forall 1 \leq n \leq N_{lp}, \quad \mathcal{N}_{nn} = \sum_p Q_{pn}^2 \,. \tag{8.15}$$

This change of operator corresponds to the change of variable $\boldsymbol{\delta} \rightarrow \boldsymbol{\delta}'$ (and equivalently $\boldsymbol{\alpha} \rightarrow \boldsymbol{\alpha}'$) defined by $\delta'^{(n)}_{i,j} = \mathcal{N}_{nn}^{\frac{1}{2}} \delta^{(n)}_{i,j}$.

This renormalization of the lensing operator has a limited effect on the condition number of the matrix. More interestingly, applying the transpose operator $\mathbf{Q}'^t$ to a vector κ amounts to computing the correlation between the signal κ and each column of $\mathbf{Q}$ which contain the convergence profiles for density deltas at each lens planes. Therefore, now that the columns are normalised to unity, the vector $\mathbf{Q}'^t\kappa$ peaks at the redshift for which the lensing signal of a thin lens correlates the best with the vector κ. As a result, after a gradient descent step, the coefficients of highest amplitudes are expected to be at the right redshift. As discussed in the following paragraph, this will allow us to speed up the convergence of the algorithm by implementing an adaptive penalisation of the coefficients based on their amplitudes.

GLIMPSE is an implementation of FISTA which differs only in the thresholding operator used in the backward step. Instead of using standard Soft Thresholding, which is known to introduce a bias in the amplitude of the reconstruction and which does not discriminate coefficients based on their amplitude above the threshold, we use Firm Thresholding (Gao and Bruce, 1997). Firm Thresholding results in an unbiased solution by defining a significance criterion above which coefficients will no longer be altered or shrunk. Sparsity is imposed by shrinking all coefficients below the required significance level, with a shrinkage that is dependent on the amplitude of the coefficient and tends to favour the most significant coefficients.

The firm thresholding operator therefore depends on two parameters, λ_1 and λ_2, and is defined by

$$\mathrm{FT}_{\lambda_1,\lambda_2}(x) = \begin{cases} 0, & \text{if } |x| \leq \lambda_1, \\ \mathrm{sgn}(x)\frac{\lambda_2(|x|-\lambda_1)}{\lambda_2-\lambda_1}, & \text{if } \lambda_1 < |x| \leq \lambda_2, \\ x, & \text{if } |x| > \lambda_2. \end{cases} \tag{8.16}$$

This operator is used instead of the soft thresholding operator ST_λ in each iteration of the algorithm (c.f. Equation (8.11)). The parameter λ_1 sets the minimum significance level above which coefficients are retained in the solution estimate, thereby effectively denoising the solution, while λ_2 defines the significance level above which the shrinkage vanishes. The behaviour of the two thresholding operators is illustrated on Figure 8.2.

In order to provide an intuitive method for choosing λ_1, we relate it directly to the noise level $\sigma^{(p)}$ in redshift bin p, estimated at each iteration by considering the Median Absolute Deviation (MAD) of the residual

$$\sigma^{(p)} = \frac{\mu}{0.6747}\,\mathrm{MAD}\left(\left[\mathbf{Q}'^t\mathbf{P}^t_{\gamma\kappa}\Sigma^{-1}\left[\gamma - \mathbf{P}_{\gamma\kappa}\mathbf{Q}'\Phi\alpha'_n\right]\right]^{(p)}_{i,j}\right). \tag{8.17}$$

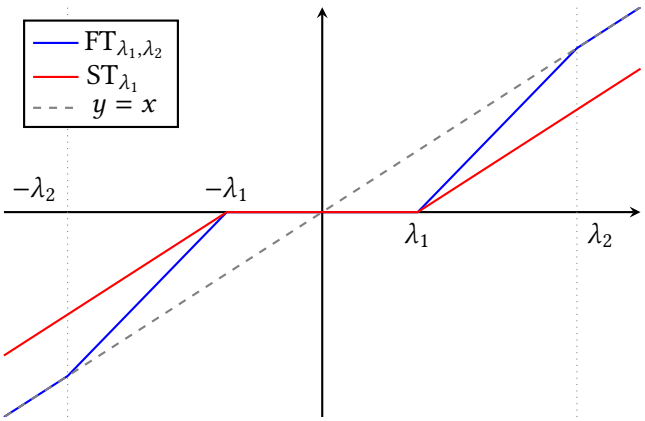

Figure 8.2: Comparison of the responses of the Firm Thresholding (blue) and Soft Thresholding (dashed red) operators.

With this estimation of the noise level, the threshold λ_1 is set for each redshift plane in the reconstruction to $\lambda_1 = k\sigma^{(p)}$, where the level of significance k is iteratively adjusted using a decreasing threshold scheme until reaching a minimum level k_{min} set by the user. For small fields, it may not be practical to estimate the MAD of the residual for each redshift bin separately. Renormalising the residual in each redshift bin by

$$\mathcal{R}^{(p)} = \left[\sum_p \mathbf{Q}'^2_{pq}\Sigma^{-1}_{pp}\right]^{-\frac{1}{2}} \tag{8.18}$$

allows the MAD to be estimated over the entire 3D field by approximately equalising the variance of the residual in each redshift bin. This method ignores correlations between redshift bins introduced by the application of the $\mathbf{Q}'^t$ operator on the data, which can be significant, but for small fields this may be preferable to computing the MAD over a statistically small number of pixels when estimating the threshold. Note that if the data are binned such that the data noise covariance Σ is proportional to the identity, $\mathcal{R}^{(p)}$ will be equal for each redshift bin p due to the normalisation of the $\mathbf{Q}'$ operator.

For the reconstructions presented here, we choose $k_{min} = 4$. We further choose $\lambda_2 = \sigma^{(p)}/\mu$, such that reconstructed coefficients with this significance level are not affected by shrinkage. This results in an unbiased estimator of the density contrast.

At the first iteration, the most significant wavelet coefficient (highest SNR) is selected and the level k is set to the average between the first and second highest SNR coefficients in the residuals. Therefore, at the beginning only this selected coefficient is allowed to enter the reconstruction. An update of the level k occurs whenever the SNR s of a different wavelet coefficient becomes greater than the SNR s_0 of this selected coefficient. This update is controlled by a sufficient decrease parameter c_{sd}, which we set to 0.1, and is only allowed when $\frac{s-s_0}{s-k_{min}} \geq c_{sd}$. When this condition is met, the most significant wavelet coefficient is selected and the threshold level is decreased to the average between the first and second highest SNR in the residuals. The procedure is repeated until the minimum thresholding level k_{min} is reached. This will lead to a slowly decreasing threshold, allowing at most only one coefficient to become active at each iteration. This scheme further promotes the sparsity of the solution and improves the speed of the reconstruction.

Finally, we do not want to impose the same sparsity constraint along the line of sight and in the tangential plane. To relax the sparsity constraint on each redshift plane, we identify the real space support of detected wavelet coefficients and we set the thresholding level k for any other wavelet coefficients within this spatial domain to the lowest level k_{min}.

The full description of the method is provided in Algorithm 8.1.

Algorithm 8.1 GLIMPSE reconstruction algorithm

Require: 3D data arrays:

- $\boldsymbol{\gamma}$: Complex 3D array of the binned shear.
- Σ: Diagonal 3D noise covariance matrix.

Require: Parameters:

- N_{scale}: Number of wavelet scales.
- k_{min}: Minimum threshold level.
- c_{sd}: Sufficient decrease parameter.

Initialisation: $\alpha'_0 = \mathbf{0}$, $t_0 = 1$, $X_0 = \alpha'_0$, $s_0 = 0$, $(x_0, y_0, z_0, j_0) = (0, 0, 0, 0)$, μ set according to Equation (8.14)
for $n = 0$ to $N_{\max} - 1$ **do**
  $\delta'_n = \Phi\alpha'_n$ *% Estimated density contrast*
  $r_n = \mu\Phi^t\mathbf{Q}'^t\mathbf{P}^t_{\gamma\kappa}\Sigma^{-1}(\boldsymbol{\gamma} - \mathbf{P}_{\gamma\kappa}\mathbf{Q}'\delta'_n)$ *% Forward gradient descent step*
  $\tilde{X}_{n+1} = \alpha_n + r_n$

  % Noise estimation on each redshift bin and each wavelet scale
  for $p = 1$ to N_{lp} and $j = 1$ to N_{scale} **do**
    $\sigma_j^{(p)} = \mathrm{MAD}(r_{n,j}^{(p)})/0.6747$
  end for

  % Update of the threshold level, $\max_2$ *gives the second highest value*
  $s_0 = |r_{n,j_0}^{(z_0)}[x_0, y_0]|/\sigma_{j_0}^{(z_0)}$
  $s = \max(|r_n|/\sigma)$
  if $s - s_0 \geq c_{sd}(s - k_{min})$ **then**
    $k = (s + \max_2(|r_n|/\sigma))/2$
    $(x_0, y_0, z_0, j_0) \leftarrow$ set to indices of maximum SNR coefficient
  end if

  % Backward thresholding step
  for all x, y, p and $j = 1$ to N_{scale} **do**
    if $\delta'^{(p)}_n[x, y] \neq 0$ **then**
      $X^{(p)}_{n+1,j}[x, y] = FT_{k_{min}\sigma_j^{(p)}, \sigma_j^{(p)}/\mu}(\tilde{X}^{(p)}_{n+1,j}[x, y])$
    else
      $X^{(p)}_{n+1,j}[x, y] = FT_{k\sigma_j^{(p)}, \sigma_j^{(p)}/\mu}(\tilde{X}^{(p)}_{n+1,j}[x, y])$
    end if
  end for

  % FISTA specific update of the coefficients
  $t_n = \frac{1+\sqrt{1+4t_n^2}}{2}$
  $\alpha'_{n+1} = X_{n+1} + \left(\frac{t_n - 1}{t_{n+1}}(X_{n+1} - X_n)\right)$
end for
return Density contrast $\delta = N\Phi\alpha'_{N_{\max}}$

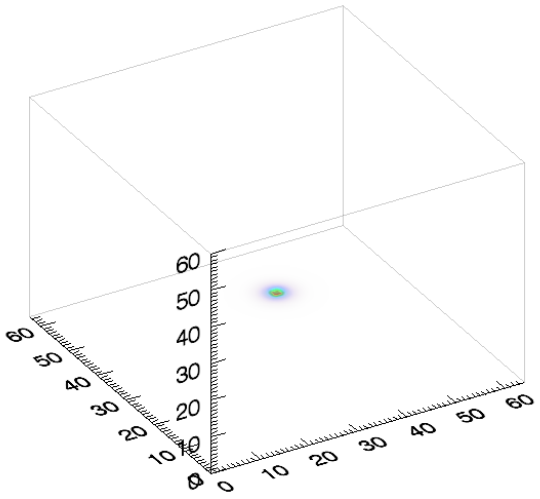

(a) Input simulated density contrast

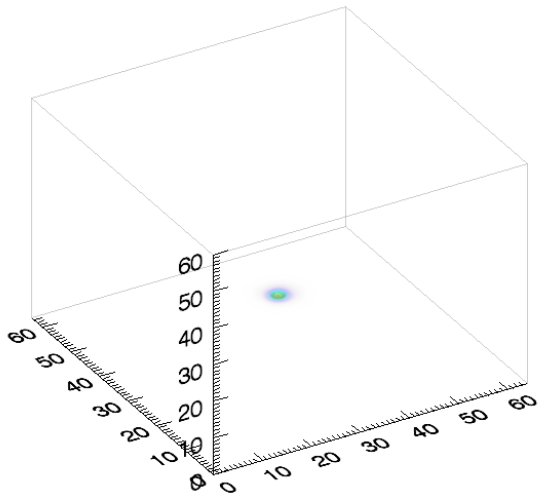

(b) Density contrast reconstruction using GLIMPSE

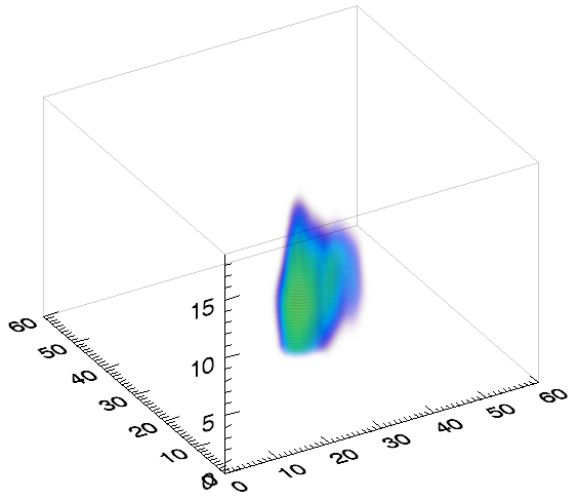

(c) SNR map thresholded at 4.5σ using Transverse Wiener Filtering

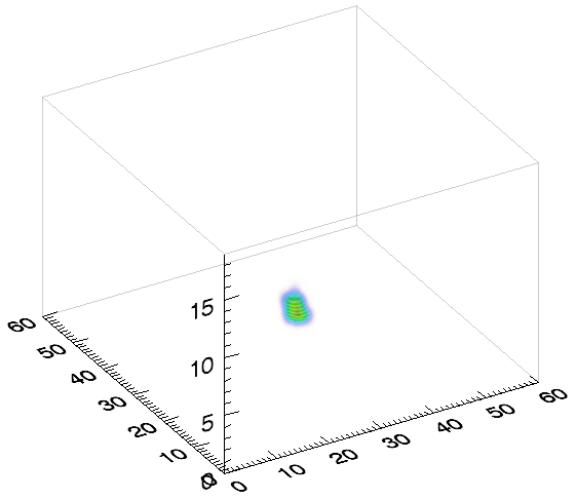

(d) SNR map thresholded at 4.5σ using Radial Wiener Filtering and Transverse Gaussian smoothing σ_t=2 arcmin

Figure 8.3: Reconstruction on an input density profile using GLIMPSE and transverse and radial Wiener filtering. The input density profile corresponds to an NFW profile for a cluster of mass $10^{15}h^{-1}M_{\odot}$ at redshift $z = 0.55$. The noise level corresponds to 30 gal/arcmin^2 and σ_ϵ. The vertical axis corresponds to redshift and ranges from $z = 0$ to $z = 2$ for all four plots but the Wiener filtering is limited to 20 redshift bins against 60 for GLIMPSE.

8.1.3 Comparison to state of the art methods

A thorough characterisation of the performance of the algorithm will be performed in the next section. As a first test, we compare in a simple case the results of the GLIMPSE algorithm to the Wiener filter method of Simon et al. (2009) presented in Section 6.4.

For this comparison, we simulate the lensing signal produced by a dark matter halo of mass $10^{15}h^{-1}M_{\odot}$ at redshift $z = 0.55$ according to an NFW profile (Navarro et al., 1997). We assume noise levels corresponding to a realistic Euclid-like space survey (number density of galaxies $n_{gal} = 30$

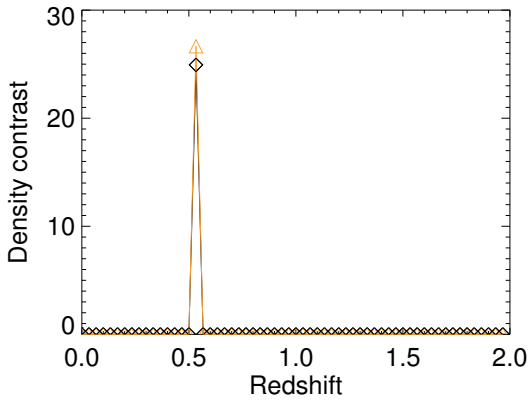

Figure 8.4: Mean density contrast along the central four lines of sight of the cluster from the input (black diamond) and the reconstruction (orange triangle) of a single noise realisation.

arcmin^{-2} with a shape noise level of $\sigma_\epsilon = 0.25$). We also assume photometric redshift errors using a centred Gaussian model for the errors, with a variance of $\sigma_z = 0.05(1+z)$. Finally, the size of the 3D field is 64×64 arcmin^2, with 1 arcmin^2 angular bins, up to a depth in redshift of $z = 2$. The input simulated density contrast is binned in redshift using 60 bins ($\Delta_z = 0.03$) while we compute the shear for only 20 redshift bins of fixed width $\Delta_z = 0.1$. A halo of this mass and at this redshift corresponds to a high SNR regime and we are consequently above the detection limit for all the reconstruction techniques considered. The input density field is shown on Figure 8.3a.

We have applied to this simulated lensing signal the two wiener filtering methods of Simon et al. (2009) and the GLIMPSE algorithm. Results are illustrated on Figure 8.3. It is important to stress that for the GLIMPSE method on Figure 8.3b it is the reconstructed density from a *single noise realisation* that is plotted whereas we plot SNR maps computed from 200 noise realisations and thresholded at 4.5σ for the Wiener filtering methods.

Even in this high SNR regime considered here, the improvement over the linear methods is obvious. The GLIMPSE algorithm has successfully reconstructed the dark matter halo, at the correct redshift and with the correct shape while the Wiener filtering can at best provide imprecise SNR maps with redshift smearing and at a much lower redshift resolution. The amplitude of the reconstructed cluster using GLIMPSE is also very close to the original amplitude as can be seen on Figure 8.4 where we plot the input and reconstructed central line of sight.

8.2 CHARACTERISATION ON CLUSTER SIMULATIONS

8.2.1 Simulations

The GLIMPSE method is nonlinear, and therefore the performance of the algorithm cannot be derived analytically, but rather needs to be assessed through numerical simulations. To do this, we generated a suite of cluster simulations similar to the one presented in Section 8.1.3, spanning a range of redshifts between $0.05 \leq z \leq 0.75$ and virial masses in the range $3 \times 10^{13} h^{-1} M_\odot \leq M_{\rm vir} \leq 10^{15} h^{-1} M_\odot$. For the density profile of the clusters we adopted an NFW (Navarro et al., 1997) profile given by

$$\rho(r) = \frac{\rho_s}{(cr/r_{\rm vir})(1+cr/r_{\rm vir})^2}\,, \tag{8.19}$$

where ρ_s is the central density parameter, $r_{\rm vir}$ is the virial radius and c is the concentration parameter. As described in Takada and Jain (2003), the number of degrees of freedom for this profile can be reduced to 2, keeping only a dependence in mass M and redshift z. First, the central density parameter ρ_s can be eliminated from the definition of the virial mass

$$M = \int_0^{r_{\rm vir}} 4\pi r^2 dr \rho(r) = \frac{4\pi\rho_s r_{\rm vir}^3}{c^3}\left[\ln(1+c) - \frac{c}{1+c}\right] . \tag{8.20}$$

Then the virial radius can be linked to the mass of the halo through the spherical collapse model

$$M = \frac{4\pi}{3}\overline{\rho}_0 \Delta_{\rm vir}(z) a^{-3} r_{\rm vir}^3 , \tag{8.21}$$

where $\overline{\rho}_0$ is the mean density of matter today and $\Delta_{\rm vir}(z)$ is the critical overdensity of collapse at redshift z. Following Coupon et al. (2012), we use the fitting formula for $\Delta_{\rm vir}(z)$ from Weinberg and Kamionkowski (2003):

$$\Delta_{\rm vir}(z) = 18\pi^2\left[1 + 0.399(\Omega_m^{-1} - 1)^{0.941}\right] . \tag{8.22}$$

Finally, we assume a mass and redshift dependance for the concentration parameter $c(M, z)$ given by

$$c(M, z) = c_0(1+z)^{-1}\left[\frac{M}{M_\star}\right]^\beta , \tag{8.23}$$

where $M_\star$ is the non-linear mass scale at present day defined by $\delta_c(z = 0) = \sigma(M_\star)$, in which δ_c is the linear critical density and $\sigma(M)$ is the RMS of density fluctuations in a sphere of radius $(3/4\pi M\overline{\rho}_0)^{1/3}$. We adopt the parameterisation from Coupon et al. (2012): $c_0 = 11$ and $\beta = 0.13$.

These relations allow us to parametrise the simulated halos only by their mass and redshift within the framework of the halo model. From the NFW density profile, we computed the corresponding shear signal, which is derived analytically in Takada and Jain (2003). The computations required to simulate the halos were performed making extensive use of the NICAEA software package[1], using a flat ΛCDM cosmology with $\Omega_{\rm M} = 0.264$, $\Omega_\Lambda = 0.736$, and $H_0 = 71$km/s/Mpc.

The shear signal was simulated for each cluster in an otherwise empty field of $1^\circ \times 1^\circ$, with an angular pixel size of $1' \times 1'$ on 30 tomographic redshift bins, with a galaxy redshift distribution given by

$$n(z) = z^\alpha \exp\left(-\left[\frac{z}{z_0}\right]^\beta\right) , \tag{8.24}$$

where we take $z_0 = 1/1.4$, $\alpha = 2$ and $\beta = 1.5$.

In this work, we consider the redshift information to be provided by photometric redshift measurements with Gaussian errors with a standard deviation that varies as a function of redshift as $\sigma_z(1+z)$, and a potential bias z_{bias}, but without catastrophic failures. Following Ma et al. (2006), the true redshift distribution $p^{(n)}(z)$ of a tomographic bin n can be obtained through:

$$p^{(n)}(z) = \frac{1}{2}p(z)\left[\mathrm{erf}(x_{n+1}) - \mathrm{erf}(x_n)\right] \tag{8.25}$$

with

$$x_n = \frac{z_{n+1} - z + z_{bias}}{\sqrt{2}\sigma_z(1+z)} . \tag{8.26}$$

1 http://www.cosmostat.org/software/nicaea/

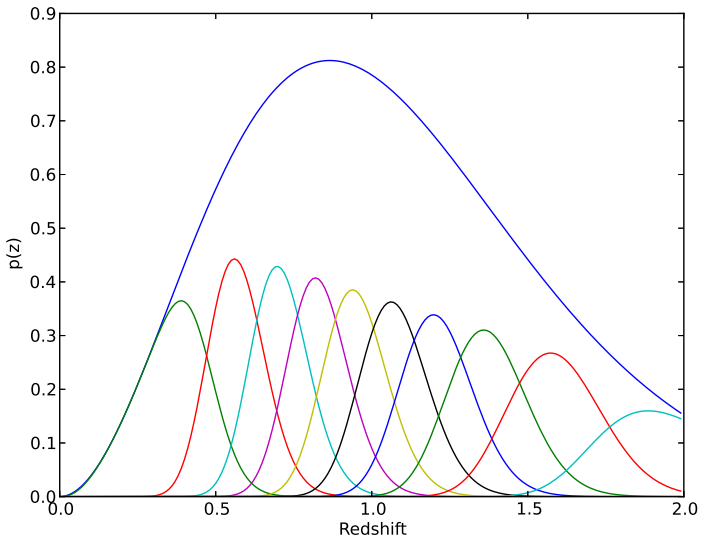

Figure 8.5: Global redshift probability distribution and 10 photometric redshift bins with an equal number of galaxies per bin assuming unbiased photometric redshift estimates with Gaussian errors whose width varies as a function of redshift as $\sigma = 0.05(1 + z)$.

For the work presented here, we took $z_{bias} = 0$ and $\sigma_z = 0.05$, the latter being in line with the minimum required accuracy for photometric redshifts in the Euclid survey. The global redshift distribution is plotted in Figure 8.5, with the galaxy distribution for 10 tomographic redshift bins overplotted, each containing the same number of galaxies per bin.

The final simulated fields were obtained by adding independent Gaussian shape noise to the shear signal with a standard deviation of $\sigma_\varepsilon = 0.25$, and we assumed a number density of galaxies of $n_g = 30/\text{arcmin}^2$. For each halo, we generated 1000 such noisy fields from independent noise realisations, in order to estimate our reconstruction errors. Figure 8.6 shows the distribution in virial mass and redshift of the 96 haloes simulated. The values were chosen to approximately trace the mass function (at the high mass end) and our estimated detection limits (at the low-mass end), but these sample points should not be taken to represent a complete sample of haloes detectable with our method.

Note that, for these simulations, we assumed pixellated data from the start, thereby implicitly assuming a uniform distribution of galaxies in x and y. This is not realistic; however, the method presented is entirely general, and can account for binned/pixellated data where the noise level per pixel varies.

Remark about the simulations

After the publication of the results presented in this section, we discovered an error in our simulations. Because of a confusion between proper distances and comoving distances, the factor a^{-3} was initially missing in Equation (8.21). The direct consequence of this error is an important underestimation of the density contrast of the simulated clusters, and thus of the lensing signal, as well as an underestimation of the virial radius. Although our simulated clusters are still NFW profiles, for a given mass, their lensing signal is smaller than for realistic clusters in the ΛCDM cosmology we consider in this chapter.

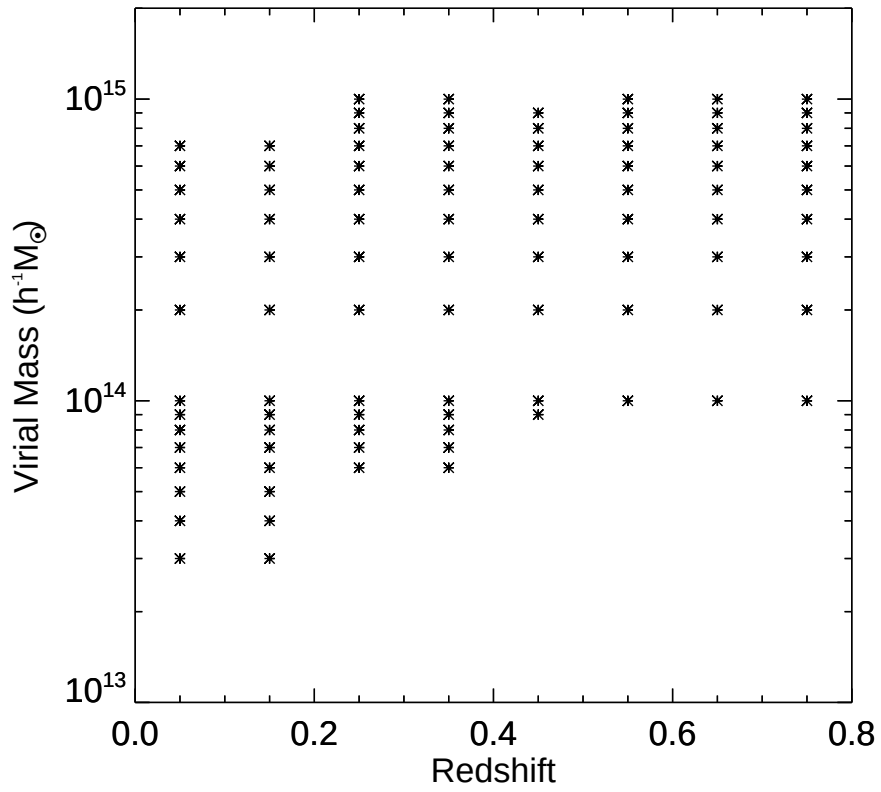

Figure 8.6: Mass and redshift distribution of simulated cluster haloes.

We have corrected this error and updated our results when possible or relevant, in particular the selection function presented in Figure 8.8.

As these simulations remain very simplistic, we have not run the complete analysis on the corrected simulations (running GLIMPSE on all 96000 fields takes between 1 and 2 months with the resources at our disposal) and have instead preferred working towards characterising our algorithm on much more realistic N-body simulations integrating the effects of the large scale structures. This work however is still ongoing. As a result, the redshift estimation in Figure 8.10 are still using the first set of simulations and are therefore very conservative, especially for clusters at high redshift, since the SNR of the simulated clusters should have been higher.

8.2.2 Processing of the reconstructions and identification of peaks

The 1000 independent noise realisations generated for each of the 96 fields were all processed by GLIMPSE using the same set of parameters described in the previous section. Having processed these simulations, it is now necessary to identify in the reconstructed fields actual detections and disentangle those from spurious structures due to noise.

Individual detections were isolated in each field by identifying peaks in the reconstructed density contrast using the Clumpfind algorithm[2] (Williams et al., 1994). This method identifies connected pixels above a given threshold within a 3D map and is able to effectively deblend overlapping structures. We use a minimum threshold of $\delta_{\min} = 1$ and we determine the 3D position of each

2 http://www.ifa.hawaii.edu/users/jpw/clumpfind.shtml

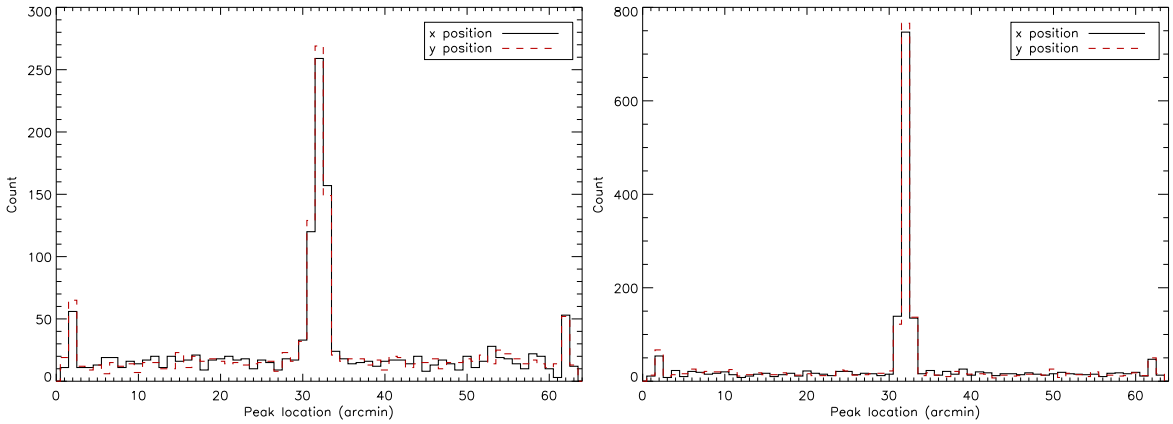

(a) Field 48 **(b)** Field 53

Figure 8.7: The distribution of the $x-y$ locations of peaks detected in 1000 reconstructions of a given cluster field. Plotted are the distributions for field 48 (left panel, $M_{vir} = 2\times10^{14}h^{-1}M_{\odot}$, $z_{48} = 0.35$) and field 53 (right panel, $M_{vir} = 7\times10^{14}h^{-1}M_{\odot}$, $z_{53} = 0.35$). In both fields, the central cluster is clearly identified, however the cluster in field 48 is less massive, and therefore the frequency of its detection is lower than that of the cluster in field 53. Away from the centre of the field, the distribution of (false) peaks is seen to be uniform, except near the edge of the field where a small overdensity is seen.

detected peak by computing the weighted centroid of the structure. This position information will allow to disentangle actual detections of the central cluster, which are expected to be located around the center of the field, from spurious noise contamination.

To illustrate the angular distribution of peaks identified by this procedure in the reconstructed fields, Figure 8.7 shows the specific case of two cluster fields. The left panel shows the distribution of x and y angular positions on the sky of all peaks detected in the 1000 realisations of the simulated data for field 48 ($M_{vir} = 2\times10^{14}h^{-1}M_{\odot}$, $z_{48} = 0.35$), while the right panel shows the distribution in x and y for the 1000 realisations of field 53 ($M_{vir} = 7\times10^{14}h^{-1}M_{\odot}$, $z_{53} = 0.35$). The distribution of the peak positions is largely seen to be uniform over the field, but with a notable excess at $\{x, y\} = \{32', 32'\}$. We can clearly recognise the detections due to the actual cluster at the center and the low level uniform background of detections due to noise. We also notice a slight excess seen near the edges of the field, which is indicative of some systematic edge effects resulting from our method, which can most probably be attributed border conditions in the computation of the wavelet transform which would require a specific treatment in our noise model. This is however not an issue as in practice the size of the reconstructed field can be increased as to avoid these border effects. For all but two fields (fields 77 and 87), the x-y peak distribution was well fitted by a Gaussian plus a constant, with the width of the Gaussian consistently found to be in the range $0.5 \lesssim [\sigma_x, \sigma_y] \lesssim 1.0$. For fields 77 and 87 (both with $M_{vir} = 10^{14}h^{-1}M_{\odot}$; $z_{77} = 0.65$, $z_{87} = 0.75$), the SNR of the central cluster was too low for consistent detection and no significant bump could be identified in the $x-y$ distribution of the detected peaks around the center of the field.

Given the measured dispersion in the position of the central peak, we consider that any peak found to lie within $30' < [x, y] < 34'$ correspond to a detection of the cluster. This cut rejects most of the spurious noise peaks but is not guaranteed to yield a perfect purity. In particular, we put no constraint on the redshift at which the peak was detected. In rare cases, a high redshift detection is seen in the reconstruction in addition to a lower-redshift peak. We opted to keep both peaks as

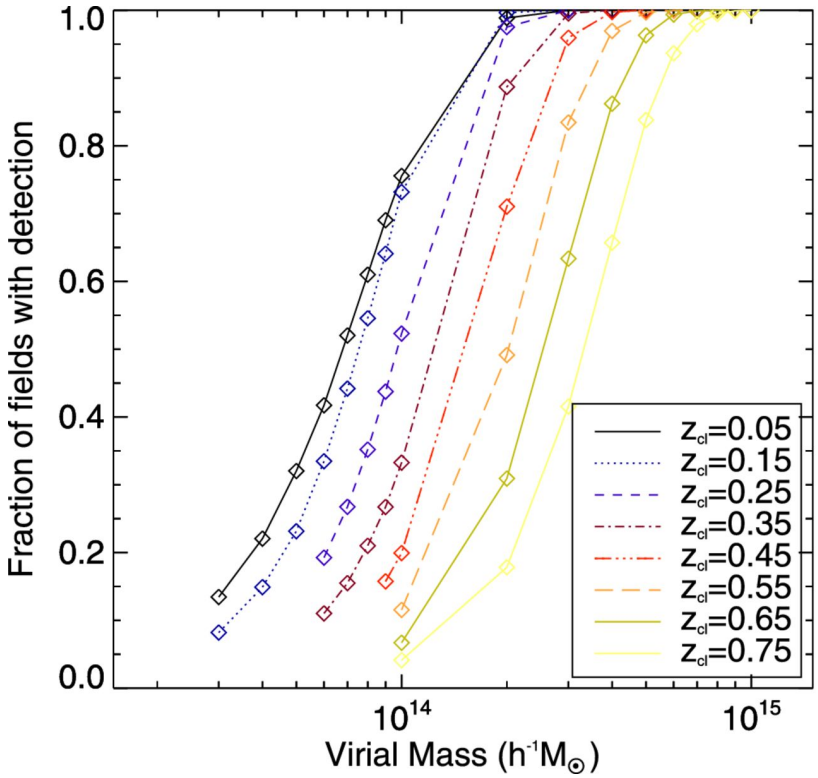

Figure 8.8: Fraction of fields with a central cluster detection, f_{det} as a function of cluster virial mass and redshift.

cluster detections in all of the analysis which follows, as there is no truly blind way to discard one or other of the peaks as being a false peak. For this work, we wish to analyse the accuracy with which we can recover the mass and redshift of the simulated clusters, and to do this in an unbiased way, we cannot preferentially discard either detection based on our prior knowledge of the true location of the cluster. This will add to the dispersion in our redshift and mass estimates, but will ensure that we are not preferentially biasing our results, and thereby unfairly shrinking our error bars.

8.2.3 Detection rates

The number of times the central cluster is detected in the ensemble of noise realisations is indicative of both the fidelity of the detection of a given structure (i.e. whether that structure is a real structure or simply due to the noise) and diagnostic of the trade-off between noise removal and sensitivity to real structures. We expect that the probability of detecting a cluster halo will depend on both the mass of that halo and its redshift, for a fixed noise and threshold level.

Plotted in Figure 8.8 is the fraction of noise realisations (out of 1000) in which one or more central density peak was detected as a function of virial mass, with different coloured curves corresponding to cluster haloes at different redshifts. We denote this fraction in the analysis that follows by f_{det}.

These curves effectively show the probability of reconstructing haloes of a given mass and redshift from a set of noisy data, given our choice of threshold parameter and shape noise level. Clearly,

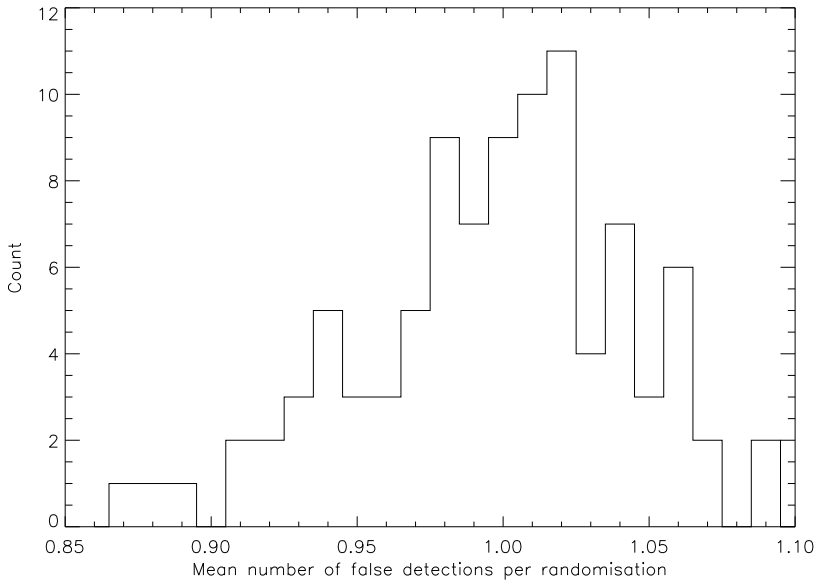

Figure 8.9: The mean number of false detections per randomisation for each of the 96 simulated cluster fields. The histogram shows the number of cluster fields (the count) with a specified mean false detection rate. This mean is taken by considering the number of peaks detected in 1000 different reconstructions whose centroids lie outside of the central 4×4 pixels.

as clusters are moved to higher redshift, they become more difficult to detect; however at the high end of the mass function ($\gtrsim 8 \times 10^{14} h^{-1} M_{\odot}$), we can still expect to detect $\gtrsim$ 60% of clusters at redshift $z_{cl} = 0.75$, and this trend is expected to continue to higher redshifts, though with an increasing mass threshold. Lowering our denoising threshold would be expected to increase the completeness at all redshifts, allowing us to detect lower mass haloes, but will come at the cost of an increased number of false detections.

Figure 8.9 shows the mean number of false detections per reconstruction for each of the 96 simulated fields. This number is consistently around 1 per reconstruction, with no dependence seen on the mass and redshift of the simulated cluster in that field. Stricter thresholding would be expected to reduce this number, at the cost of lowering the detection rates of actual clusters.

8.2.4 Redshift estimation

We now consider the redshift accuracy of our reconstructions, by assessing how closely the redshift of a reconstructed cluster matches that of the original simulation. To do this, we consider the distribution in redshift of any peaks identified as a true detection by the procedure detailed in the previous paragraph. For each cluster field, we considered the distribution in the estimated redshift over the ensemble of noise realisations.

We consider two redshift estimators for the cluster. The first is the median redshift in this distribution, for which error bars were computed by finding the redshift range, centred on the median value z_{med}, that encompassed 68.2% of the redshift estimates (thereby representing the 1σ confidence interval of the redshift estimate).

We also estimated the Probability Density Function (PDF) for the redshift estimates using an *Adaptive Kernel Density Estimate* (AKDE) with a standard Epanechnikov kernel (Silverman, 1986).

From this, we estimated cluster redshift z_{peak} as the peak in the PDF. A 1σ confidence interval was then computed from the cumulative distribution function obtained by integrating the PDF, always ensuring that 68.2% of the estimated redshift values fell within this interval.

Figure 8.10 shows the estimated vs true redshift for clusters in our sample as a function of mass. Note that the results for $M_{vir} = 3 \times 10^{13} h^{-1} M_\odot$ are not plotted in the interests of space, but are consistent with the $M_{vir} = 4 \times 10^{13} h^{-1} M_\odot$ clusters. Shown in the figure are the median estimate and the AKDE peak redshift, the latter being offset by $\delta z = 0.01$ in both directions for visualisation purposes. Both estimates clearly yield an accurate estimate of the redshift, with the peak estimate seen to have smaller error bars in many cases. No systematic bias is seen in the peak case, while the median estimate does appear to slightly overestimate the redshifts for low-mass haloes. The error bars are seen to increase in size with increasing redshift, and with decreasing halo mass, as expected.

As will be discussed more fully in the next section, an error in redshift will naturally give rise to a bias in the estimated mass of the cluster. We can parametrise this error through the relation $M(z_{est}) = (1 \pm |f_m|) M(z_{true})$. The contours labelled $|f_m| = 0.1$ and $|f_m| = 0.5$ refer to the redshift ranges within which the fractional error on the mass is less than 10% and 50%, respectively. For the 96 simulated fields, 56 [36] yield median [AKDE] redshift estimates that fall within the contours of $|f_m| = 0.1$ and 76 [77] fall within $|f_m| = 0.5$. The outliers tend to arise in low signal-to-noise clusters: for the 71 clusters with a detection fraction $f_{det} \geq 0.2$, 74.6% [46.5%] yielded median [AKDE] redshift estimates within the contours of $|f_m| = 0.1$, and 87.3% [85.9%] within $|f_m| = 0.5$.

8.2.5 Mass estimation

Given our ability to accurately estimate the redshifts of clusters using GLIMPSE, we now consider how accurately we might be able to estimate the masses of the haloes detected. Estimating the masses of clusters detected using our method is complicated by several factors, however.

Firstly, for any given noise realisation, the cluster will be detected at a given redshift z_i. The density contrast is related to the mass by

$$M_n = \sum_n \delta^{(n)}(z_i) \overline{\rho}(z_i) V_{pix}(z_i) \,, \tag{8.27}$$

where $\overline{\rho}(z_i)$ is the mean matter density at redshift z_i and $V_{pix}(z_i)$ is the comoving volume covered by one pixel in our reconstruction at that redshift, and the mass M_n is the mass enclosed by n angular pixels. Therefore, the estimated mass depends on redshift through both $\overline{\rho}$ and V_{pix}, and any redshift error in a given reconstruction will bias the mass estimate by a corresponding factor. If the cluster is detected at a different redshift, then for the same amplitude in δ, the estimated mass will change proportionately as $\overline{\rho} V_{pix}$.

Moreover, the density contrast estimate itself will be biased if the cluster is placed at the wrong redshift. Recall that $\mathbf{Q} \propto D_{ls}/D_s$. This means that if a cluster is detected, for example, at a higher redshift, it will be reconstructed with a larger amplitude in order to compensate for the decrease in the estimated D_{ls}/D_s. This scaling is encoded by the normalisation matrix $\mathcal{N}$ defined in Equation (8.15), $\mathcal{N}(z_n) \equiv \mathcal{N}_{nn} = \sum_p \mathbf{Q}_{pn}^2$, where the index n refers to a given redshift slice. As the normalised density contrast, the quantity we reconstruct with GLIMPSE, is related to δ via $\delta' = \mathcal{N}\delta$, the density contrast estimated at a given redshift (and thereby the estimated mass at that redshift) is scaled by $M_n(z) \propto \delta(z) \propto \mathcal{N}^{-1}(z)$.

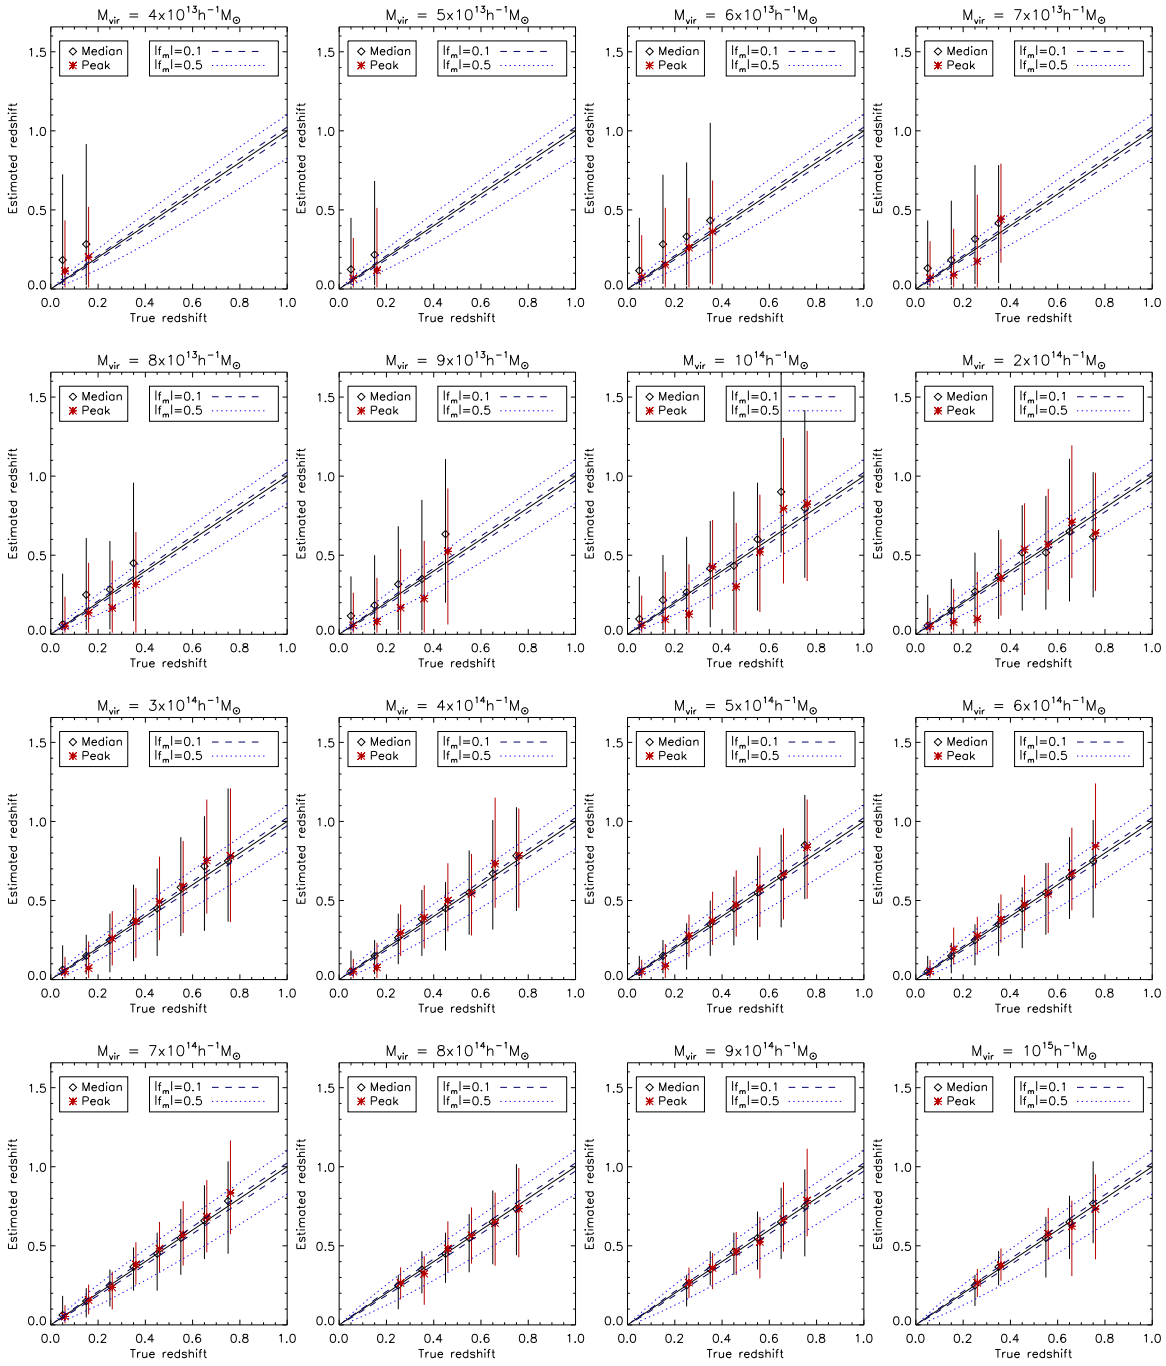

Figure 8.10: Estimated vs true redshift for the 96 clusters in our sample, separated by cluster virial mass. Red points indicate the peak of the AKDE distribution, while black points indicate the median value considering all noise realisations. The red points have been shifted by 0.01 in both directions for visualisation purposes.

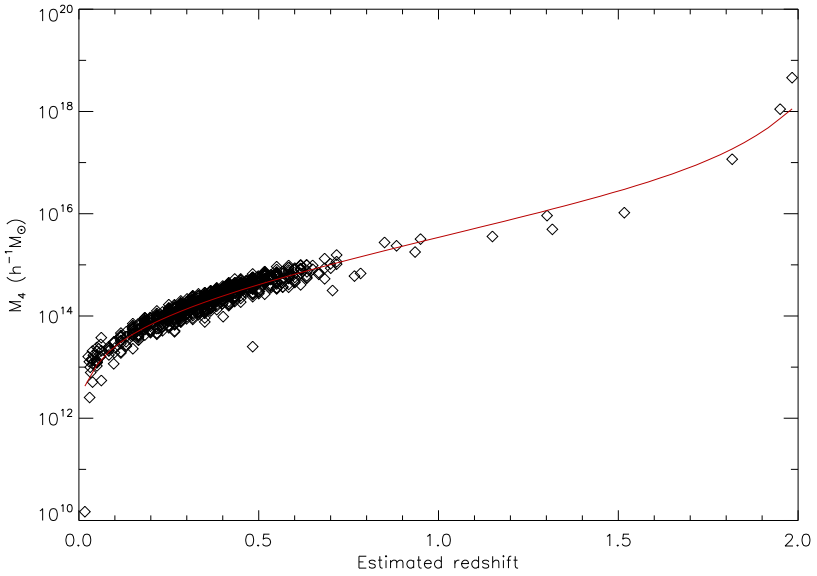

Figure 8.11: The mass estimate within the central four angular pixels M_4 as a function of estimated redshift for the reconstructions of cluster field 53. The red curve shows M_4 measured from the input simulation, and rescaled with redshift according to the relation in Equation (8.28).

The estimated mass at redshift z_{est} is therefore related to the mass estimate at the true redshift z_{true} via:

$$\frac{M_n(z_{est})}{M_n(z_{true})} = \frac{\overline{\rho}(z_{est})}{\overline{\rho}(z_{true})} \frac{V_{pix}(z_{est})}{V_{pix}(z_{true})} \frac{\mathcal{N}(z_{true})}{\mathcal{N}(z_{est})} \,. \tag{8.28}$$

In Figure 8.11, we plot the estimated mass over the central four pixels M_4 for all the reconstructions of field 53 ($M_{vir} = 7 \times 10^{14} h^{-2} M_\odot$, $z_{53} = 0.35$), computed simply by multiplying the integrated density contrast by the mean density and the pixel volume at the redshift of detection, without accounting for any biases. It is very clear that there is a strong dependence of the mass estimate on the redshift. The red curve in the figure shows M_4 estimated on the true density contrast map of the cluster, scaled as a function of redshift according to Equation (8.28).

This curve shows an excellent fit to the measured data, and we can therefore use the relationship in Equation (8.28) to rescale the mass estimates for a given cluster onto the same redshift, to allow us to compare their values, and to compute an overall mass estimate for the cluster, and error bars associated with this estimate.

In order to assess the ability of our method to constrain the angular profile of a cluster, we consider the mass within the central 2×2 pixels (M_4), 4×4 pixels (M_{16}), 6×6 pixels (M_{36}) and 8×8 pixels (M_{64}). For each detected cluster, the density contrast was rescaled by $\mathcal{N}$, and then summed within the central n angular pixels at all redshifts associated with the cluster. We then compute the mass *assuming the true redshift of the cluster*. As none of these mass measures directly corresponds to the virial mass of the simulated cluster, we also compute these values directly from the simulated maps for comparison.

Given the distribution of estimated masses across all the noise realisations, we can then compute the median and 68.2% confidence regions as for the redshift. Figure 8.12 shows our results. Included are only results for clusters in which the fraction of detections $f_{det} > 0.3$ (i.e. for which we have more than 300 detections out of the 1000 randomisations). This limit was placed to ensure that we

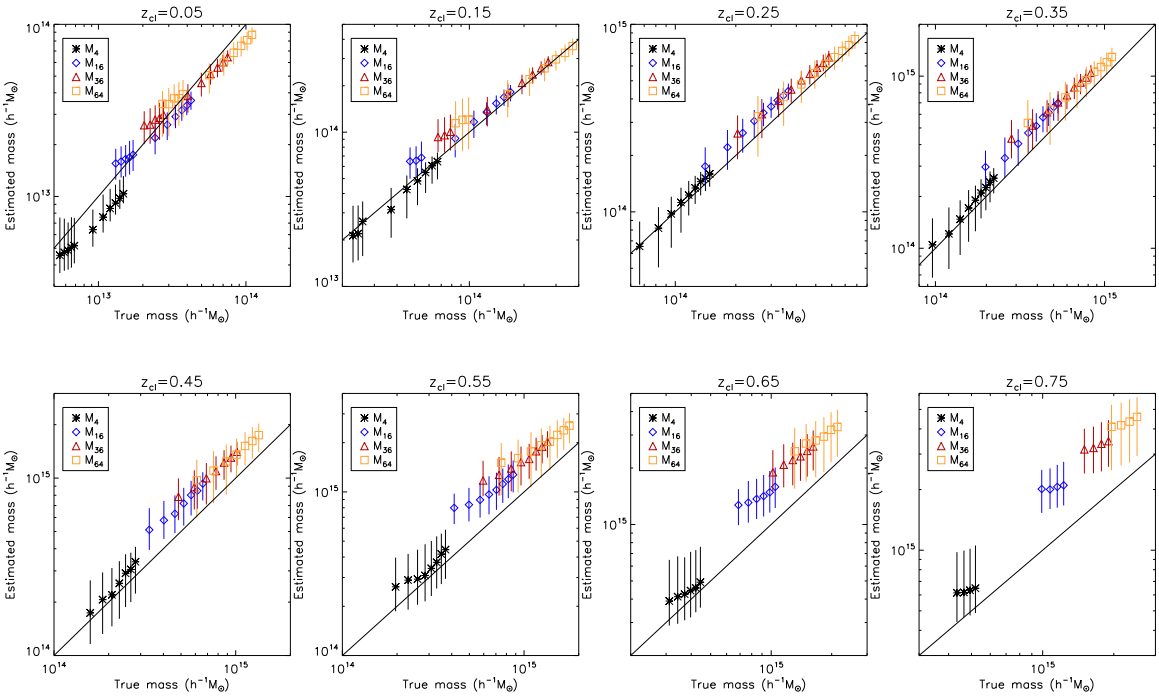

Figure 8.12: Mass estimates at four different integration scales plotted against the true mass measured from the simulation. The mass estimates use the median rescaled mass over all the realisations of the field, and the error bars are 68.2% confidence intervals computed symmetrically about the median value.

have adequate sample points for our derived statistics. In Figure 8.12, the four different estimators are shown on the same plot, and the clusters are divided up by redshift.

At the lowest simulated cluster redshift of $z_{cl} = 0.05$, we see a tendency to underestimate the masses at all scales, while the opposite trend is seen at high redshift. However, in most cases, and particularly for M_4, this is at most a 1σ effect. In the intermediate redshift regime $0.15 \leq z_{cl} \leq 0.55$, excellent agreement is seen in M_4. The trend to overestimate the mass at larger radii implies that our fit to the angular profile of the cluster is not perfect; however, the deviations seen remain at the 1σ level.

We can conclude from these results that we are able to accurately recover the masses of most clusters in our sample, even at high redshift, with very little bias seen. However, this analysis has relied on the assumption that the true redshift of the cluster is known, which it will not be in real data.

Using Equation (8.28), we can compute the expected bias on the mass that will be obtained by assuming the median or peak AKDE estimate of the redshift, instead of the true redshift. As in the previous section, we model this as

$$\frac{M(z_{est})}{M(z_{true})} = 1 + f_m \, . \tag{8.29}$$

Equation (8.28) then directly gives us a measure of f_m. In Figure 8.13 we show the histogram of f_m for the clusters included in Figure 8.12 (i.e. all the clusters with $f_{det} > 0.3$). We plot this mass bias factor for the case in which we assume as the cluster redshift the median estimate (solid black line) and the peak AKDE redshift (dashed red line). In the case of the median estimator, the vast

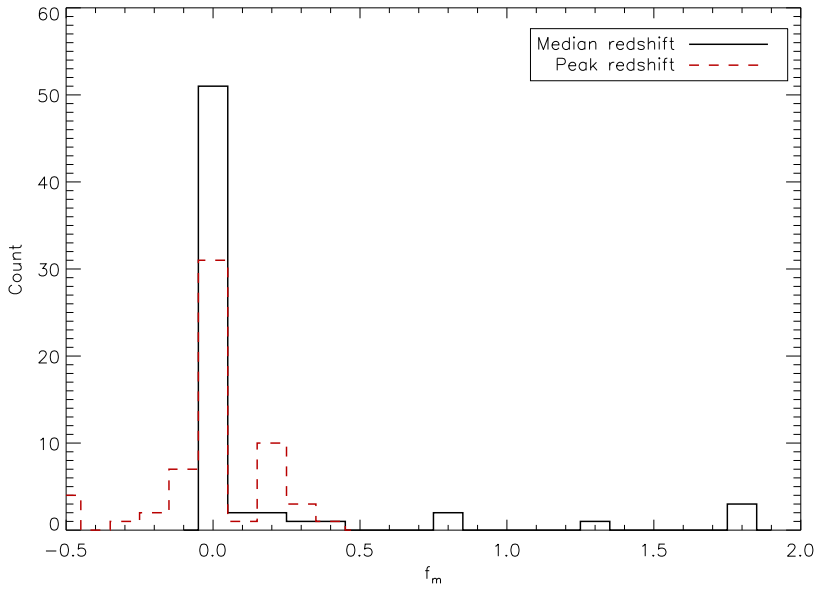

Figure 8.13: Mass bias associated with a redshift error for the clusters included in Figure 8.12, shown for both the median redshift estimate and the peak AKDE estimate. The bias is generally very low for the median estimate, with $|f_m| \leq 0.05$ in most cases. The peak redshift estimate shows a wider distribution, but no extreme outliers with $|f_m| > 0.5$.

majority of clusters fall within the ±0.05 region, implying a smaller than 5% induced bias from errors in the redshift estimate. The peak AKDE redshift shows a wider spread in f_m values, but in contrast to the median redshift does not show any extreme outliers of $|f_m| > 0.5$.

8.3 COMPARISON OF 2D AND 3D DETECTION EFFICIENCY

As was demonstrated in the previous section, GLIMPSE is capable of detecting galaxy clusters from 3D weak lensing. A natural question to ask is whether a conventional 2D mass-mapping method would fare better or worse in terms of detection efficiency. To shed some light on this question, we perform in this section a simple experiment using the simulations described in Section 8.2.1 to compare the detection efficiency of GLIMPSE with that of a conventional 2D Kaiser-Squires inversion followed by a simple wavelet filtering, performed using MRLens (Starck et al., 2006).

8.3.1 2D detection using MRLens

The MRLens method (Starck et al., 2006) relies on a Multi-Resolution Entropy filtering. Formally, the filter recovers a denoised map κ from the noisy convergence κ_n by minimising:

$$\min_{\kappa} \frac{1}{2} \parallel \kappa_n - \kappa \parallel^2 - H(\kappa) \, . \tag{8.30}$$

The second term, $H(\kappa)$ encodes the prior on the solution, expressed in wavelet space as:

$$H(\kappa) = \sum_{j=1}^{J-1} \sum_{k,l} h(w_{j,k,l}) \, , \tag{8.31}$$

where J is the number of dyadic wavelet scales and $w_{j,k,l}$ are the *starlet* coefficients of κ. The information content $H(\kappa)$ is therefore defined as the sum of the information content of each wavelet coefficient $h(w_{j,k,l})$, defined by the function h (see Starck et al., 2006, for details on the choice of this function). However, MRLens does not apply regularisation for coefficients that are clearly detected above the noise (i.e. significant coefficients). In this case, the multi-scale entropy becomes:

$$h_m(w_{j,k,l}) = \overline{M}(j,k,l)h(w_{j,k,l}) \,, \tag{8.32}$$

where $\overline{M}(j,k,l) = 1 - M(j,k,l)$ and $M(j,k,l)$ is a multiresolution support defined as:

$$M(j,k,l) = \begin{cases} 1 & \text{if } w_{j,k,l} \text{ is significant} \\ 0 & \text{if } w_{j,k,l} \text{ is not significant} \end{cases} \tag{8.33}$$

As a result, entropy based regularisation will only be applied to wavelet coefficients considered to be not significant, while no regularisation is applied to significant coefficients.

For the purpose of detecting galaxy clusters from noisy convergence maps, we are only interested in those significant wavelet coefficients. We therefore base our 2D detection method on the multiresolution mask M itself and not directly on the denoised map. We first collapse the multiresolution mask across wavelet scales to produce a binary 2D map. We then use this map to define a detection as a group of connected non-zero pixels and finally compute a convergence-weighted centroid for the detection from the denoised converge map.

In the case of Gaussian noise, the *multiresolution support* is defined by identifying as significant those wavelet coefficients $w_{j,k,l}$ that verify $|w_{j,k,l}| > K\sigma_j$, where σ_j is the noise standard deviation at a particular wavelet scale j, and the indices $[k,l]$ denote the pixel location in the 2D convergence map. Typically, this threshold K is chosen to be between 3 and 5.

Note that we use the same threshold K across all scales. This is not however the default option in MRLens which rather defines a scale dependent threshold based on a user defined False Detection Rate (FDR) but we use a fixed threshold for ease of comparison with GLIMPSE.

This 2D detection procedure using MRLens is very close to the 3D GLIMPSE detection procedure detailed in the previous section. Both methods are based on an isotropic starlet transform and use a similar $K\sigma$ detection strategy, leaving the 2D and 3D nature of the two approaches as the main difference.

8.3.2 Results

Our aim is to compare the detection efficiency of the two methods, at the same false detection rate, which is set in both methods by the $K\sigma$ level used for the detection. A full ROC curve analysis would require to run both methods for a wide set of parameter K, however given the computational cost of running the 3D reconstructions we keep the GLIMPSE detection level fixed at 4σ, thus also fixing a false detection rate of 0.78 false detections per square degree. For a meaningful comparison, we use MRLens with three different detection threshold at 3, 3.5 and 4 σ corresponding to false detections rate above and below the reference set by GLIMPSE. These rates of false detections are listed on Table 7. As can be seen, using MRLens with the 3.5σ threshold gives the number of false detections closest to that of GLIMPSE at 4σ.

Algorithm	Threshold	False detection per square degree
MRLens	3σ	2.087 ± 0.357
MRLens	3.5σ	0.373 ± 0.069
MRLens	4σ	0.056 ± 0.018
GLIMPSE	4σ	0.781 ± 0.047

Table 7: Number of false detections per field for the 96 cluster fields analysed with MRLens at 3 different denoising thresholds and GLIMPSE at a denoising threshold of 4σ.

The detection rates of the central cluster, for different masses and redshifts, using MRLens and GLIMPSE with the parameters described above are given on Figure 8.14. At low redshift (upper left panel), we can check that the GLIMPSE 4σ and MRLens 3.5σ results, which correspond to roughly the same number of false detections, are very comparable. As one would expect, the more conservative MRLens 4σ has a lower detection efficiency while at 3σ MRLens is able to recover more clusters, of course at the cost of a larger number of false detection.

The interesting result of this study is the evolution of the detection rates with the redshift of the cluster. While the hierarchy between the three MRLens detection rates remains the same, the relative efficiency of GLIMPSE starts increasing with the redshift of the cluster after $z \simeq 0.3$ to the point where it outperforms MRLens for all three parameters (see lower right panel).

Although this experiment is fairly simple, it perfectly illustrates the potential gain of exploiting the 3D information for the purpose of detecting galaxy clusters from weak lensing. This was first demonstrated in Hennawi and Spergel (2005) where the use of a tomographic matched filter enhanced the number of cluster detected with an SNR above 4.5 by as much as 76% compared to a non tomographic analysis. The reason for this important difference between 2D and 3D approaches comes the specific line-of-sight response of a thin-lens which constitutes valuable information to discriminate between noise and actual signal. In our GLIMPSE approach, the algorithm effectively applies at each iteration a matched filter for isolated density peaks in redshift. This enhances the SNR of actual clusters by adequately weighting the tomographic lensing signal compared to a simple 2D analysis.

8.4 CONCLUSION

In this chapter we have introduced a new 3D weak lensing based on sparse regularisation. While previous linear methods remained very limited, especially in their redshift resolution, we demonstrate that our non-linear reconstruction algorithm is successful at recovering not only accurate redshift but also density contrast estimates which can directly be used to constrain the masses of detected dark matter haloes.

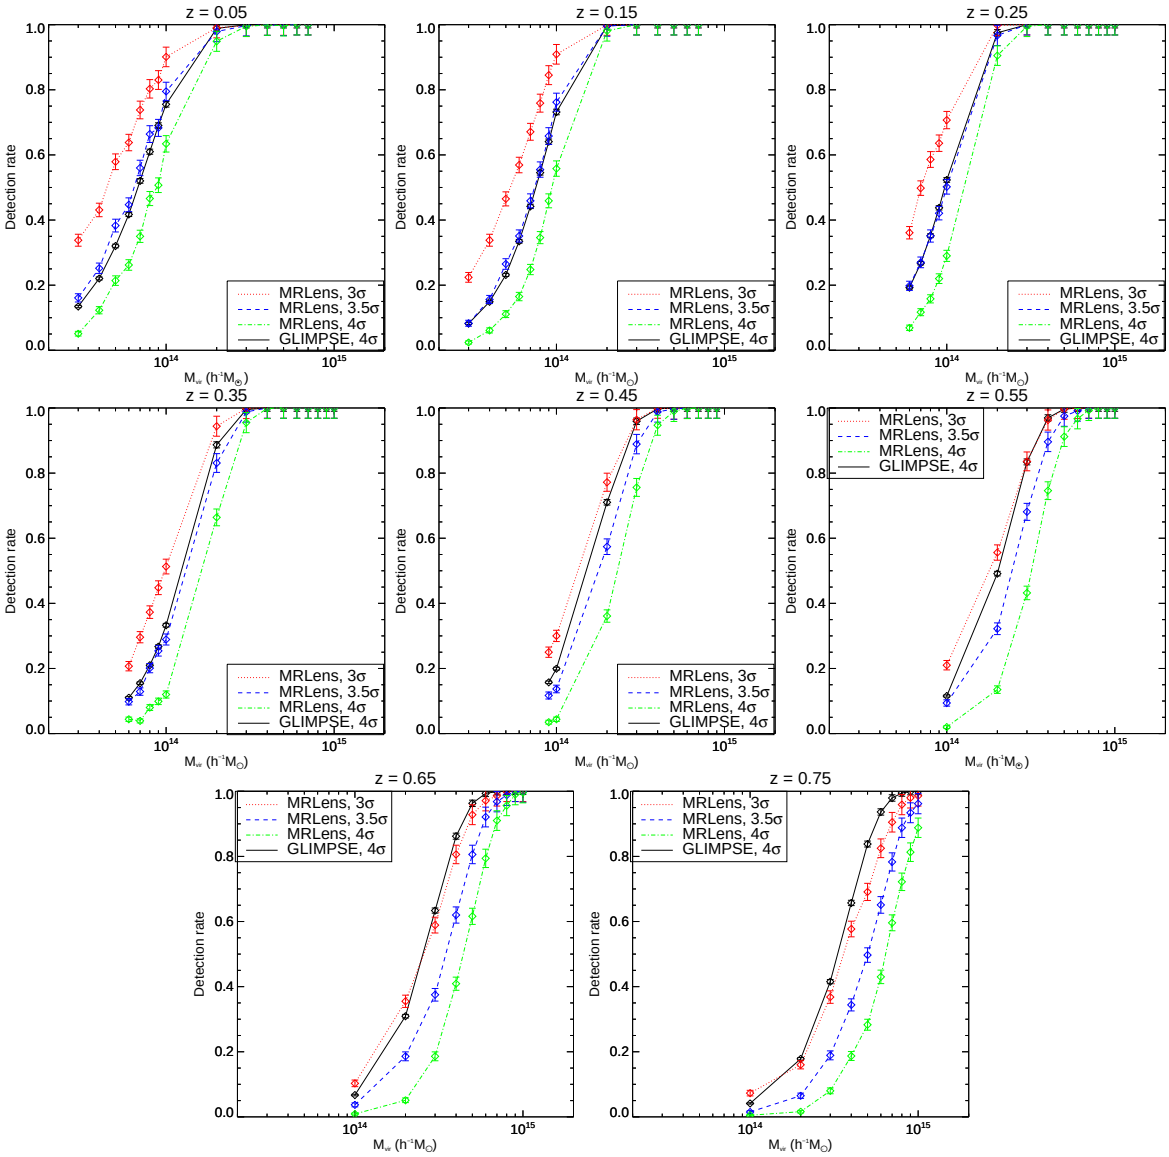

Figure 8.14: Detection rate of the central cluster using MRLens (blue) and GLIMPSE (black) for different cluster redshifts. The parameters of the two methods are calibrated to yield a similar false detection rate. At low redshift, the detection rates are similar but at high redshift, GLIMPSE clearly outperforms MRLens.

Part III

COSMOLOGY ON THE BALL

9 ISOTROPIC WAVELETS ON THE 3D BALL

CONTENTS

In this chapter, we introduce a new isotropic wavelet on the 3D ball as an extension of the 2D starlet transform introduced in Chapter 4. Such representations are of particular interest for the analysis of fields most naturally expressed in spherical coordinates. In particular, these wavelets are a promising tool for the study of current and future wide spectroscopic and photometric galaxy surveys. They open the door to a number of applications of sparse methodologies to these data sets, such as detection, denoising or deconvolution.

As these wavelets are based on the *Spherical Fourier-Bessel Transform* (SFBT), this chapter also provides an introduction to the Spherical Fourier-Bessel (SFB) framework used in the analysis presented in Chapter 10.

The results presented in this chapter were published in Lanusse et al. (2012). This work has been conducted in collaboration with Jean-Luc Starck and Anais Rassat.

9.1 SPARSE REPRESENTATIONS ON THE 3D BALL

Wavelets were introduced in Chapter 4, where we detailed the construction of 2D discrete wavelets. As was illustrated on multiple occasions throughout this thesis, wavelets are fundamental tools which find many applications, especially when combined with sparse regularisation. Based on the

framework presented in that chapter, 3D wavelets can be easily derived on a 3D Cartesian grid (see Lanusse et al. (2014) for a review of 3D sparse representations). However, these transforms are not adapted to the spherical geometry of signals living on the 3D ball. Such signals arise for instance in astrophysics in the study of the 3D distribution of galaxies (e.g. Heavens and Taylor, 1995; Rassat and Refregier, 2012). Indeed, modern cosmological surveys, either lensing or galaxy surveys, are probing the Universe in three dimensions, combining the angular position of objects on the sky with their radial distance.

While there is great interest in developing sparse representations for signals on the ball, until very recently there were no 3D wavelet on the ball which allowed analysis and synthesis in a discrete setting (although continuous wavelets on the ball had been introduced before, mainly for applications in geophysics (e.g Michel, 2005)).

In this chapter, we will present the first wavelet transform on the 3D ball which allows both synthesis and analysis in the discrete and continuous domain. This wavelet is based on the *Spherical Fourier-Bessel Transform* (SFBT), which is the natural harmonic expansion of 3D data in spherical coordinates and preserves the link between angular and radial scales. This allows us to build a shift invariant, isotropic wavelet decomposition on the 3D ball.

Very shortly after the publication of the wavelet transform presented in this chapter, a different expansion was introduced in Leistedt and McEwen (2012) which differs mainly by their choice of harmonic expansion on the ball. It is constructed from an exact sampling theorem in the angular domain based on McEwen and Wiaux (2011) and in the radial domain is based on the orthogonality of Laguerre polynomials. The resulting Fourier-Laguerre transform allows for exact decomposition and reconstruction of band limited signals on the 3D ball and is used to implement a wavelet transform (named flaglets) with exact decomposition and reconstruction formulae[1]. Due to the choice of independent basis for the radial and angular domains, flaglets probe independently angular and radial scales. However, separating angular and radial domains breaks the 3D translational invariance of the harmonic expansion. Although this separation can be desirable to disentangle purely radial and angular effects (for instance angular masks and photometric redshift errors), it is not optimal when probing for isotropic features (for instance cosmic voids).

While the SFBT presented in Section 9.2 has a number desirable properties, the main drawback of this transform however is that no exact sampling theorem exists in the radial domain (Lemoine, 1994). Contrary to the Fourier-Laguerre transform, the spherical Fourier-Bessel Transform cannot be computed exactly for a discretely sampled band limited signal on the ball. To circumvent this issue, we introduce in Section 9.3 a *Discrete Spherical Fourier-Bessel Transform* (DSBT) which allows in practice the evaluation of this transform to any desired accuracy. We then describe in Section 9.4 our isotropic wavelet on the 3D ball, which is exact in the spherical Fourier-Bessel domain and for which wavelet coefficients can be recovered in the direct domain using the DSBT.

9.2 THE SPHERICAL FOURIER-BESSEL TRANSFORM

In this section, we introduce the Spherical Fourier-Bessel transform and its two components, the Spherical Harmonics Transform and Spherical Bessel Transform. We also provide a key result

1 The source code for this transform is publicly available at `http://www.jasonmcewen.org`

regarding convolutions in SFB space which will be at the core of the wavelet transform described in Section 9.4.

9.2.1 Spherical Harmonics Transform

Laplace's spherical harmonics are a set of functions that form an orthonormal basis of the Hilbert space of square-integrable functions defined on the unit sphere. They are also eigenfunctions of the angular part of Laplace's equation which makes them very useful in Physics and especially in Quantum Mechanics.

They are indexed by two indices $\ell \in \mathbb{N}$ and $m \in [\![-\ell, \ell]\!]$ and can be formally defined as:

$$Y_{\ell m}(\theta, \varphi) = \sqrt{\frac{(2l+1)}{4\pi}\frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(\cos(\theta)) e^{im\varphi} \ , \tag{9.1}$$

where $P_{\ell m}(\cos(\theta))$ is an associated Legendre polynomial. This definition, ensures the normalisation of the spherical harmonics so that $\int |Y_{\ell m}|^2 d\Omega = 1$. These functions verify the following orthonormalisation property :

$$\int_0^{2\pi} \int_0^{\pi} Y^*_{\ell m}(\theta, \varphi) Y_{\ell' m'}(\theta, \varphi) \sin(\theta) \, \mathrm{d}\theta \, \mathrm{d}\varphi = \delta_{\ell\ell'}\delta_{mm'} \ . \tag{9.2}$$

Since they form an orthonormal basis, spherical harmonics can be used to develop any square-integrable scalar field defined on the unit sphere. Let $f(\theta, \varphi)$ be such a function, then by projecting f onto the $\{Y_{\ell m}\}_{\ell m}$ basis using the Hermitian product we get:

$$\forall \ell \in \mathbb{N}, \forall m \in [\![-\ell, \ell]\!] \quad f_{\ell m} = \int_0^{2\pi} \int_0^{\pi} Y^*_{\ell m}(\theta, \varphi) f(\theta, \varphi) \sin(\theta) \, \mathrm{d}\theta \, \mathrm{d}\varphi \ . \tag{9.3}$$

This leads to the spherical harmonics expansion of f:

$$f(\theta, \varphi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} f_{\ell m} Y_{\ell m}(\theta, \varphi) \ . \tag{9.4}$$

Relations Equation (9.3) and Equation (9.4) respectively define the *Spherical Harmonics Transform* (SHT) and its inverse.

Discrete data representation on the sphere

In practice, data acquired on the sphere needs to be pixelated in order to compute a SHT and various schemes exist in the literature. These include the Equidistant Coordinate Partition (ECP), the Icosahedron method (Tegmark, 1996), IGLOO (Crittenden and Turok, 1998), *Hierarchical Equal Area isoLatitude Pixelization of a sphere* (HEALPix) (Gorski et al., 2005), Gauss-Legendre Sky Pixelization (GLESP) (Doroshkevich et al., 2005). Some of these sampling schemes allow for exact forward and backward SHT of band limited signals on the sphere. This is for instance the case of GLESP, which directly derives its pixelisation scheme from the Gauss-Legendre quadrature leading to an exact SHT. More recently, exact sampling schemes based on an ECP grid have been proposed with a particular focus on spin spherical harmonics. In particular, McEwen (2008) introduced a sampling scheme for theoretically exact and fast spin SHT for arbitrary spin numbers but found the forward

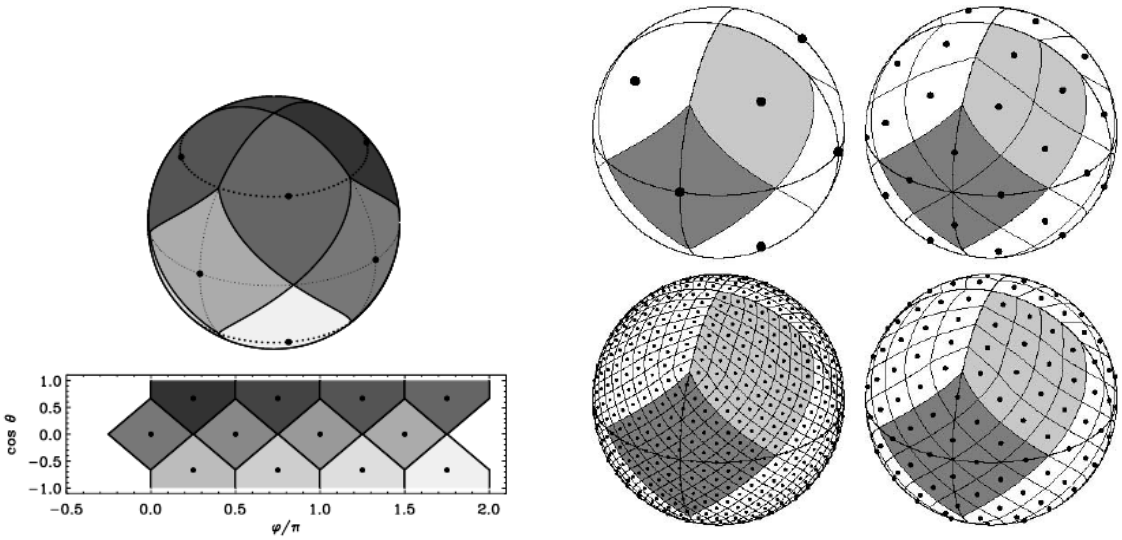

(a) HEALPix map at base resolution $N_{side} = 1$

(b) HEALPix maps for $N_{side} = 1, 2, 4, 8$

Figure 9.1: HEALPix pixelisation scheme. Credit: Gorski et al. (2005)

transform to be extremely unstable to the point that multipoles over $\ell \simeq 32$ could not be reliably computed in practice. Huffenberger and Wandelt (2010) proposed a similar exact equiangular sampling scheme while providing a stable forward algorithm thus solving the crucial limitation of the previous scheme and allowing for extremely accurate forward and backward spin SHT up to high multipoles (over $\ell \simeq 4096$). This stable forward algorithm was subsequently reused in McEwen and Wiaux (2011) but with a reduction of a factor 2 in the number of sampling points necessary to the exact quadrature scheme (of the order of $2L^2$ versus $4L^2$ for Huffenberger and Wandelt (2010), where L is the band limit of the signal).

Although any of these representations could be used for our purpose, we chose to use the popular HEALPix representation for the results presented in this chapter (for which the SHT is not exact but can still be evaluated to acceptable accuracy). This choice was essentially motivated by the excellent software package[2] publicly available but HEALPix also boasts a number of very desirable properties in practice which are not necessarily found in exact sampling schemes such as GLESP or McEwen and Wiaux (2011) such as its uniform sampling of the sphere and its hierarchical nature.

The HEALPix representation is a curvilinear hierarchical partition of the sphere into quadrilateral pixels of exactly equal area but with varying shape. The base resolution divides the sphere into 12 quadrilateral faces of equal area placed on three rings around the poles and equator. Each face is subsequently divided into N_{side}^2 pixels following a quadrilateral multiscale tree structure (see Figure 9.1). The pixel centres are located on iso-latitude rings, and pixels from the same ring are equispaced in azimuth, which is critical for efficient SHT algorithms.

9.2.2 Spherical Bessel Transform

The *Spherical Bessel Transform* (SBT), is an integral transform, part of the family of Hankel transforms which are based on a set of Bessel functions. There are several flavours of these Bessel

2 http://healpix.jpl.nasa.gov.

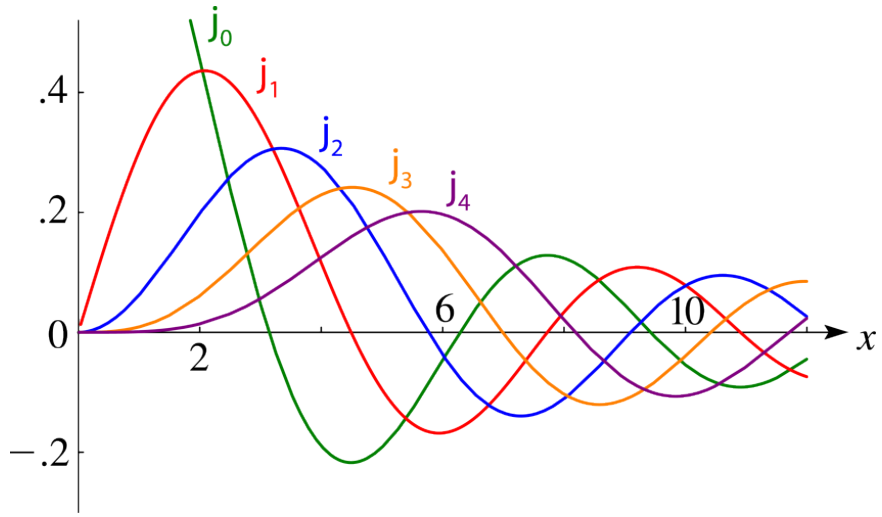

Figure 9.2: Spherical Bessel Function of the first kind $j_\ell(r)$ for $\ell = 0, 1, 2, 4$. Credit: NIST

functions which can all be expressed in terms of the Bessel function of the first kind J_ν , defined for $z \in \mathbb{C}$ and $\nu \in \mathbb{R}$ as:

$$J_\nu(z) = \sum_{k=0}^{\infty} \frac{(-1)^k z^{\nu+2k}}{2^{\nu+2k}\, k!\, \Gamma(\nu + k + 1)} \,. \tag{9.5}$$

The SBT uses specifically spherical Bessel functions of the first kind $j_\ell(r)$ which are derivatives of J_ν and can be defined for $\ell \in \mathbb{N}$ as:

$$j_\ell(r) = \sqrt{\frac{\pi}{2r}} J_{\ell+1/2}(r);. \tag{9.6}$$

Most importantly, these functions verify the following orthogonality relation:

$$\forall k, k' \in \mathbb{R}^+, \quad \int_0^\infty j_\ell(kr) j_\ell(k'r) r^2 \,\mathrm{d}r = \frac{\pi}{2k^2} \delta(k - k') \,. \tag{9.7}$$

Given a function f defined on $[0, +\infty[$ which verifies $\int_0^\infty |f(r)| r \,\mathrm{d}r < +\infty$, the SBT of order $\ell \geq 0$ and its inverse are defined as:

$$\widetilde{f}_\ell(k) = \sqrt{\frac{2}{\pi}} \int f(r) j_\ell(kr) r^2 \,\mathrm{d}r \,, \tag{9.8}$$

$$f(r) = \sqrt{\frac{2}{\pi}} \int \widetilde{f}(k) j_\ell(kr) k^2 \,\mathrm{d}k \,, \tag{9.9}$$

where we denote by $\widetilde{f}_\ell$ the SBT of f. Note that one particular property of Hankel transforms is that they are involutions. As can be seen, the inverse SBT has the same expression as the direct SBT.

An important remark about conventions can be made at this point. Based on the definition in Equation (9.8), we opt in this chapter for the *symmetrical formulation* of the SBT, which was for

instance used in Heavens and Taylor (1995). Although much more convenient for the derivations presented in this chapter, this is not the preferred convention in cosmology and most works (e.g. Castro et al., 2005; Leistedt et al., 2012; Rassat and Refregier, 2012; Pratten and Munshi, 2013) use instead the following definition for the direct and inverse transforms:

$$\widetilde{f}(k) = \sqrt{\frac{2}{\pi}} \int f(r) j_l(kr) kr^2 \, \mathrm{d}r \,, \tag{9.10}$$

$$f(r) = \sqrt{\frac{2}{\pi}} \int \widetilde{f}(k) j_l(kr) k \, \mathrm{d}k \,, \tag{9.11}$$

As can be seen these two transforms are no longer identical. This choice however is motivated by the very simple relationship between the 3D power spectrum $P(k)$ and the SFB power spectrum $C(k)$ of a SIH field: $P(k) = C(k)$ [3] . Because this convention is widespread in cosmology, it will be used in Chapter 10 which centres more on the cosmological exploitation of galaxy surveys.

9.2.3 Spherical Fourier-Bessel Transform

The *Spherical Fourier-Bessel Transform* (SFBT) arises when dealing with three dimensional problems best described in spherical coordinates. It is a projection on a set of orthogonal functions, composed of Spherical Harmonics and Spherical Bessel functions, which are solutions of Helmholtz's equation $(\Delta + k^2) f(r, \theta, \varphi) = 0$ and is therefore widely used in physics.

We define this set of orthogonal functions by the following relations:

$$\forall \ell \in \mathbb{N}, \forall m \in [\![-\ell, \ell]\!], \forall k \in \mathbb{R}^{+*}, \quad \Psi_{\ell m k}(r, \theta, \varphi) = \sqrt{\frac{2}{\pi}} j_\ell(kr) Y_{\ell m}(\theta, \varphi) \tag{9.12}$$

Using the orthogonality relations of both Spherical Harmonics and Spherical Bessel functions, we can easily show that these functions are orthogonal:

$$\int \Psi^*_{\ell m k}(\mathbf{r}) \Psi_{\ell' m' k'}(\mathbf{r}) \, \mathrm{d}\mathbf{r} = \frac{2}{\pi} \int j_\ell(k'r) j_\ell(kr) r^2 \, \mathrm{d}r \int_\Omega Y^*_{\ell m}(\theta, \varphi) Y_{\ell' m'}(\theta, \varphi) \, \mathrm{d}\Omega \tag{9.13}$$

$$= \frac{1}{k^2} \delta(k - k') \delta_{\ell \ell'} \delta_{m m'} \tag{9.14}$$

Let f be a three dimensional, square-integrable scalar field. The SFB transform of f is defined as its projection onto the set of orthogonal functions $\{\Psi_{\ell m k}\}$:

$$\begin{aligned} \widetilde{f}_{\ell m}(k) &= \int \Psi^*_{\ell m k}(r, \theta, \varphi) f(r, \theta, \varphi) r^2 \sin(\theta) \, \mathrm{d}\theta \, \mathrm{d}\varphi \, \mathrm{d}r \\ &= \int_0^{2\pi} \int_0^{\pi} \left[\sqrt{\frac{2}{\pi}} \int_0^{\infty} f(r, \theta, \varphi) j_\ell(kr) r^2 \, \mathrm{d}r \right] Y^*_{\ell m}(\theta, \varphi) \sin(\theta) \, \mathrm{d}\theta \, \mathrm{d}\varphi \qquad (9.15) \\ &= \sqrt{\frac{2}{\pi}} \int_0^{\infty} \left[\int_0^{2\pi} \int_0^{\pi} f(r, \theta, \varphi) Y^*_{\ell m}(\theta, \varphi) \sin(\theta) \, \mathrm{d}\theta \, \mathrm{d}\varphi \right] j_\ell(kr) r^2 \, \mathrm{d}r \qquad (9.16) \end{aligned}$$

We see with this definition that the SFBT is no more than the commutative composition of a SHT and a SBT. The inversion formula for the SFBT is as follows:

$$f(r, \theta, \varphi) = \sqrt{\frac{2}{\pi}} \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \int \widetilde{f}_{\ell m}(k) k^2 j_\ell(kr) \, \mathrm{d}k \; Y_{\ell m}(\theta, \varphi) \tag{9.17}$$

3 With the symmetric convention of this chapter, this relation would become $P(k) = k^2 C(k)$.

9.2.4 Convolution in the Spherical Fourier-Bessel domain

A key property of the SFBT is the existence of an SFB expression for the convolution $h = f * g$ of two functions $f, g \in \mathbf{L}^2(\mathbb{R}^3)$ which reduces to a very simple formula in the case of an isotropic function g (Baddour, 2010).

The convolution in the SFB domain can be expressed from the well known expression in the Fourier domain:

$$\widehat{h}(k,\theta_k,\varphi_k) = \mathcal{F}\{f * g\}(k,\theta_k,\varphi_k) \tag{9.18}$$

$$= \sqrt{(2\pi)^3}\widehat{f}(k,\theta_k,\varphi_k)\widehat{g}(k,\theta_k,\varphi_k)\,, \tag{9.19}$$

using the following unitary convention for the Fourier Transform :

$$\widehat{f}(\boldsymbol{k}) = \frac{1}{\sqrt{(2\pi)^3}}\int f(\boldsymbol{r})e^{-i\boldsymbol{k}.\boldsymbol{r}}\,\mathrm{d}\boldsymbol{r} \quad ; \qquad f(\boldsymbol{r}) = \frac{1}{\sqrt{(2\pi)^3}}\int \widehat{f}(\boldsymbol{k})e^{i\boldsymbol{k}.\boldsymbol{r}}\,\mathrm{d}\boldsymbol{k}\,. \tag{9.20}$$

To relate Fourier and SFB coefficients, one can use the expansion of the Fourier kernel in spherical coordinates:

$$e^{-i\boldsymbol{k}.\boldsymbol{r}} = 4\pi\sum_{\ell=0}^{\infty}\sum_{m=-\ell}^{\ell}(-i)^\ell j_\ell(kr)Y^*_{\ell m}(\theta_r,\varphi_r)Y_{\ell m}(\theta_k,\varphi_k)\,. \tag{9.21}$$

When injected in the definition of the Fourier transform (see Section C.1), this expression directly leads to the following relation between Fourier and SFB transforms:

$$\widehat{f}(k,\theta_k,\varphi_k) = \sum_{\ell=0}^{\infty}\sum_{m=-\ell}^{\ell}\left[(-i)^\ell\widetilde{f}_{\ell m}(k)\right]Y_{\ell m}(\theta_k,\varphi_k)\,. \tag{9.22}$$

It is worth noticing that the SFBT $\widetilde{f}_{\ell m}(k)$ of f is merely an SHT applied on shells of radii k in Fourier space (up to a factor $(-i)^l$): $\widehat{f}_{\ell m}(k) = (-i)^\ell\widetilde{f}_{\ell m}(k)$.

This expression for the Fourier transform combined with the Fourier convolution formula in Equation (9.19) yields the SFB convolution formula (see Section C.2 for the full derivation):

$$\widetilde{h}_{\ell m}(k) = (i)^\ell\sqrt{(2\pi)^3}\sum_{\ell'=0}^{\infty}\sum_{m'=-\ell'}^{\ell'}(-i)^{\ell'}\widetilde{f}_{\ell' m'}(k)\sum_{\ell''=|\ell-\ell'|}^{\ell+\ell'}c^{\ell''}(\ell,m,\ell',m')(-i)^{\ell''}\widetilde{g}_{\ell'' m-m'}(k)\,, \tag{9.23}$$

where $c^{\ell''}(\ell,m,\ell',m')$ are Slater integrals:

$$c^{\ell''}(\ell,m,\ell',m') = \iint Y^*_{\ell m}(\theta,\varphi)Y_{\ell' m'}(\theta,\varphi)Y^{m-m'}_{\ell''}(\theta,\varphi)d\Omega\,. \tag{9.24}$$

These integrals are only non-zero for $|\ell-\ell'| \le \ell'' \le \ell+\ell'$.

As previously mentioned, this expression reduces to a simple form when g is isotropic. In this case, g has no angular dependence in the Fourier domain therefore $\widehat{g}$ is constant on spherical shells and $\widehat{g}_{\ell m}(k) = 0 = \widetilde{g}_{\ell m}(k)$ for all $(\ell,m) \neq (0,0)$. Then, knowing that $c^0(\ell,m,\ell,m) = 1/\sqrt{4\pi}$ Equation (C.8) becomes:

$$\widetilde{h}_{\ell m}(k) = \sqrt{2\pi}\widetilde{g}_{00}(k)\widetilde{f}_{\ell m}(k)\,. \tag{9.25}$$

This expression can therefore be used to express in the spherical Fourier-Bessel domain a convolution by any isotropic filter g.

9.3 DISCRETE SPHERICAL FOURIER-BESSEL TRANSFORM

The transform introduced so far benefit from a natural discretisation in the angular domain thanks to the spherical harmonics, however in the radial domain, the SBT is purely continuous. In order to implement wavelets in the harmonic domain and to be able to compute wavelet coefficients in the direct domain, a discretisation scheme for the SBT is required. The main difficulty comes from the lack of an exact quadrature formula for this radial transform and therefore the lack of an exact sampling theorem. To circumvent this issue, we propose an approximated discrete SBT for a radially limited signal, extension of the discrete Bessel Transform introduced in Lemoine (1994). Although this discrete transform is not exact, it can be evaluated to any desired accuracy by increasing the number of sampling points. Combined with the HEALpix (Gorski et al., 2005) sampling in the angular domain we build a sampling grid in spherical coordinates which allows for back and forth computation of the SFBT.

9.3.1 The 1D Discrete Spherical Bessel Transform

The transform described here is an extension to the Spherical Bessel Transform of the discrete Bessel transform from Lemoine (1994). This discretisation of the SBT uses the well known orthogonality property of the spherical Bessel functions on the interval $[0, R]$. If f is a continuous function defined on $[0, R]$ which verifies the boundary condition $f(R) = 0$ then the spherical Bessel transform defined Equation (9.8) can be expressed using spherical Fourier-Bessel series:

$$\widetilde{f}_\ell(k_{\ell n}) = \sqrt{\frac{2}{\pi}} \int_0^R f(r) j_\ell(k_{\ell n} r) r^2 \, \mathrm{d}r \,, \tag{9.26}$$

$$f(r) = \sum_{n=1}^{\infty} \widetilde{f}_\ell(k_{\ell n}) \rho_{\ell n} j_\ell(k_{\ell n} r) \,. \tag{9.27}$$

In this expression, $k_{\ell n} = \frac{q_{\ell n}}{R}$ where $q_{\ell n}$ is the nth zero of the Bessel function of the first kind of order ℓ and the weights $\rho_{\ell n}$ are defined as:

$$\rho_{\ell n} = \frac{\sqrt{2\pi} R^{-3}}{j^2_{\ell+1}(q_{\ell n})} \,. \tag{9.28}$$

Although this formulation provides a discretisation of the inverse SBT and of the k spectrum, the direct transform is still continuous and another discretisation step is necessary. Assuming that a boundary condition of the same kind can be applied to $\widetilde{f}_\ell(k)$ so that $\widetilde{f}_\ell(K_\ell) = 0$, then by using the same result, the spherical Fourier-Bessel expansion of $\widetilde{f}_\ell(k)$ is obtained by:

$$\widetilde{\widetilde{f}}_\ell(r_{\ell n}) = \sqrt{\frac{2}{\pi}} \int_0^K \widetilde{f}_\ell(k) j_\ell(r_{\ell n} k) k^2 \, \mathrm{d}k \tag{9.29}$$

$$\widetilde{f}_\ell(k) = \sum_{n=1}^{\infty} \widetilde{\widetilde{f}}_\ell(r_{\ell n}) \kappa_{\ell n} j_\ell(r_{\ell n} k) \,, \tag{9.30}$$

where $r_{\ell n} = \frac{q_{\ell n}}{K_l}$ and where the weights $\rho_{\ell n}$ are defined as:

$$\kappa_{\ell n} = \frac{\sqrt{2\pi} K_\ell^{-3}}{j^2_{\ell+1}(q_{\ell n})} \,. \tag{9.31}$$

The SBT being an involution, $\widetilde{\widetilde{f}} = f$ so that $\widetilde{\widetilde{f}}_\ell(r_{\ell n}) = f(r_{\ell n})$. Much like the previous set of equations had introduced a discrete $k_{\ell n}$ grid, a discrete $r_{\ell n}$ grid is obtained for the radial component. Since Equation (9.27) and Equation (9.30) can be used to compute f and $\widetilde{f}_\ell$ for any value of r and k, they can in particular be used to compute $f(r_{\ell n})$ and $\widetilde{f}_\ell(r_{\ell' n})$ where ℓ' *does not have to match* ℓ.

The Spherical Bessel Transform and its inverse can then be expressed only in terms of series:

$$\widetilde{f}_\ell(k_{\ell' n}) = \sum_{p=1}^{\infty} f(r_{\ell p}) \kappa_{\ell p} j_\ell(r_{\ell p} k_{\ell' n}) \tag{9.32}$$

$$f(r_{\ell' n}) = \sum_{p=1}^{\infty} \widetilde{f}_\ell(k_{\ell p}) \rho_{\ell p} j_\ell(r_{\ell' n} k_{\ell p}) \,. \tag{9.33}$$

Thanks to this last set of equations one can compute the SBT and its inverse without the need of evaluating any integral. Furthermore only discrete values of f and $\widetilde{f}$ respectively sampled on $r_{\ell n}$ and $k_{\ell n}$ are required.

However, this expression of the direct and inverse SBT is only valid if f is both band limited ($\widetilde{f}_\ell(K_\ell) = 0$) and radially limited ($f(R) = 0$). It is well known that these two conditions can never be verified at the same time. The same problem arises for the Fourier transform, a band limited signal necessarily has an infinite time support. In practice, by increasing the band limit K_ℓ to any arbitrary value, one can recover an approximation of the exact transform to any required accuracy.

The second difficulty comes from the infinite sums over p in Equation (9.32) and Equation (9.33). In practical applications, for a given value of ℓ only a limited number N of $\widetilde{f}_\ell(k_{\ell n})$ and $f(r_{\ell n})$ coefficients can be stored so that $r_{\ell N} = R$ and $k_{\ell N} = K_\ell$. Since $r_{\ell n}$ is defined by $r_{\ell n} = \frac{q_{\ell n}}{K_\ell}$, for $n = N$, R and K_ℓ are bound by the following relation:

$$q_{\ell N} = K_\ell R \,. \tag{9.34}$$

Therefore, the value of K_ℓ is fixed for a choice of N and R.

Nevertheless, the main point remains that any desired accuracy in the evaluation of the direct and inverse transform can be reached by increasing the number of points N and artificially increasing R above the actual radial limit of the signal.

The truncation of the direct and inverse series to N coefficients yields a convenient matrix formulation to define the *Discrete Spherical Bessel Transfrom* and its inverse. Consider the following transform matrix $\mathbf{T}^{\ell\ell'}$:

$$T^{\ell\ell'}_{pq} = \left(\frac{\sqrt{2\pi}}{j^2_{\ell+1}(q_{\ell q})} j_\ell\left(\frac{q_{\ell' p} q_{\ell q}}{q_{\ell N}}\right) \right)_{pq} \,. \tag{9.35}$$

We define the direct discrete SBT as:

$$\begin{bmatrix} \widetilde{f}_\ell(k_{\ell' 1}) \\ \widetilde{f}_\ell(k_{\ell' 2}) \\ \vdots \\ \widetilde{f}_\ell(k_{\ell' N}) \end{bmatrix} = \frac{1}{K^3_\ell} \mathbf{T}^{\ell\ell'} \begin{bmatrix} f(r_{\ell 1}) \\ f(r_{\ell 2}) \\ \vdots \\ f(r_{\ell N}) \end{bmatrix} \,. \tag{9.36}$$

Reciprocally, the values of f can be recovered on any $r_{\ell' n}$ grid from $\widetilde{f}_\ell$ sampled on $k_{\ell n}$ using the inverse discrete SBT, based on the exact same matrix:

$$\begin{bmatrix} f(r_{\ell' 1}) \\ f(r_{\ell' 2}) \\ \vdots \\ f(r_{\ell' N}) \end{bmatrix} = \frac{1}{R^3} \mathbf{T}^{\ell\ell'} \begin{bmatrix} \widetilde{f}_\ell(k_{\ell 1}) \\ \widetilde{f}_\ell(k_{\ell 2}) \\ \vdots \\ \widetilde{f}_\ell(k_{\ell N}) \end{bmatrix} . \tag{9.37}$$

A simplified form of the transform could have been defined only for $\ell = \ell'$ so that $\mathbf{T}^{\ell\ell'} = \mathbf{T}^\ell$. However, keeping the distinction between the order of the transform ℓ and the order of the grid on which the results are provided ℓ' will be crucial to the implementation of the 3D transform introduced in the next section. Indeed, the order of the grid on which the function is sampled has to match the order of the transform but the resulting transform coefficients do not. Therefore, it will be possible to compute the result of the inverse SBT of any order ℓ on a grid of order ℓ_0 so that only one radial grid of order ℓ_0 will be required. Nevertheless, for the direct transform, if the field is sampled on the radial grid of order ℓ_0, only the transform of order ℓ_0 can be computed. An additional result is required to be able to relate the SBT of different orders. This is achieved by combining Equation (9.26) and Equation (9.27):

$$\begin{aligned} \widetilde{f}_\ell(k_{\ell n}) &= \sqrt{\frac{2}{\pi}} \int_0^R \left[\sum_{m=1}^{\infty} \widetilde{f}_{\ell_0}(k_{\ell_0 m}) \rho_{\ell_0 m} j_{\ell_0}(k_{\ell_0 m} r) \right] j_\ell(k_{\ell n} r) r^2 \, \mathrm{d}r \, , \\ &= \sqrt{\frac{2}{\pi}} \sum_{m=1}^{\infty} \widetilde{f}_{\ell_0}(k_{\ell_0 m}) \rho_{\ell_0 m} \int_0^R j_{\ell_0}(k_{\ell_0 m} r) j_\ell(k_{\ell n} r) r^2 \, \mathrm{d}r \, , \\ &= \sum_{m=1}^{\infty} \hat{f}_{\ell_0}(k_{\ell_0 m}) \frac{2}{j_{\ell_0+1}^2(q_{\ell_0 m})} \int_0^1 j_{\ell_0}(q_{\ell_0 m} x) j_\ell(q_{\ell n} x) x^2 \, \mathrm{d}x \, , \\ &= \sum_{m=1}^{\infty} \hat{f}_{\ell_0}(k_{\ell_0 m}) \frac{2}{j_{\ell_0+1}^2(q_{\ell_0 m})} W_{nm}^{\ell_0 \ell} \, , \end{aligned} \tag{9.38}$$

where the weights $W_{nm}^{\ell\ell'}$ are defined as:

$$W_{nm}^{\ell_0 \ell} = \int_0^1 j_{\ell_0}(q_{\ell_0 m} x) j_\ell(q_{\ell n} x) x^2 \, \mathrm{d}x \, . \tag{9.39}$$

The final expression in Equation (9.38) is an important result which shows that the SBT of a given order can be expressed as the sum of the coefficients obtained for a different order of the transform, with the appropriate weighting. This means we can convert the Spherical Bessel coefficients of order ℓ_0 into coefficients of any other order ℓ, which considerably speeds up calculations for the SBT. It is also worth noticing that the weights $W_{nm}^{\ell\ell'}$ are simply geometric terms, i.e. independent of the field and can thus be tabulated.

We note that this approach is an extension of the Discrete Bessel Transform introduced in Lemoine (1994) but using spherical Bessel functions and where the transform in Lemoine (1994) can be considered as a special case where $\ell = \ell'$.

9.3.2 The 3D Discrete Spherical Fourier-Bessel Transform

As presented in section 9.2.3, the SFBT is the composition of a SHT for the angular component and a SBT for the radial component. Since these two transforms can commute, they can be treated independently and by combining discrete algorithms for both transforms, one can build a *Discrete Spherical Fourier-Bessel Transform*. A convenient choice for the angular part of the transform is the HEALPix (Gorski et al., 2005) pixelisation scheme introduced in Section 9.2.1. The radial component can be discretised using the discrete SBT algorithm presented in the previous section. The choice of these two algorithms introduces a discretisation of the Fourier-Bessel coefficients as well as a pixelisation of the 3D space in spherical coordinates.

The SFB coefficients $\widetilde{f}_{\ell m}(k)$ are defined by Equation (9.15) for continuous values of k. Assuming a boundary condition on the density field f, the discrete SBT can be used to discretise the values of k. The discrete SFB coefficients are therefore defined as:

$$a_{\ell mn} = \widetilde{f}_{\ell m}(k_{\ell n}) \,, \tag{9.40}$$

for $0 \leq \ell \leq L_{max}$, $-\ell \leq m \leq \ell$ and $1 \leq n \leq N_{max}$. These discrete coefficients are simply obtained by sampling the continuous coefficients on the $k_{\ell n}$ grid introduced in the previous section.

To this discretised Fourier-Bessel space corresponds a dual grid of the 3D space defined by combining the HEALPix pixelisation scheme and the discrete SBT.

In the angular domain, for a given value of r, the field $f(r, \theta, \varphi)$ can be sampled on a finite number of points using HEALPix. The radial component of the transform is conveniently performed using the discrete SBT. Indeed, this algorithm introduces a radial grid compatible with the discretised $k_{\ell n}$ spectrum. Although this radial grid $r_{\ell n}$ depends on the order ℓ of the SBT, it will be justified in the next section that only one grid $r_{\ell_0 n}$ is required to sample the field along the radial dimension. The value of ℓ_0 is set to $\ell_0 = 0$ for convenience as the properties of the zeros of the Bessel function ensure that r_{0n} will be regularly spaced between 0 and R:

$$r_{0n} = \frac{n}{N_{max}} R \,. \tag{9.41}$$

For given values of θ_i and φ_j, the field $f(r, \theta_i, \varphi_j)$ can now be sampled on discrete values of $r = r_{0n}$.

Combining angular and radial grids, the 3D spherical grid is defined as a set of N_{max} HEALPix maps equally spaced between 0 and R. An illustration of this grid is provided on Figure 9.3 where only one quarter of the space is represented for clarity.

Using this 3D grid it becomes possible to compute back and forth the SFB transform between a density field and its SFB coefficients. Here, a detailed description of the algorithm for both the direct and inverse discrete SFB transform is provided below.

INVERSE TRANSFORM Let $a_{\ell mn}$ be the discrete SFB coefficients of the density field f. The reconstruction of f on the spherical 3D grid requires two steps:

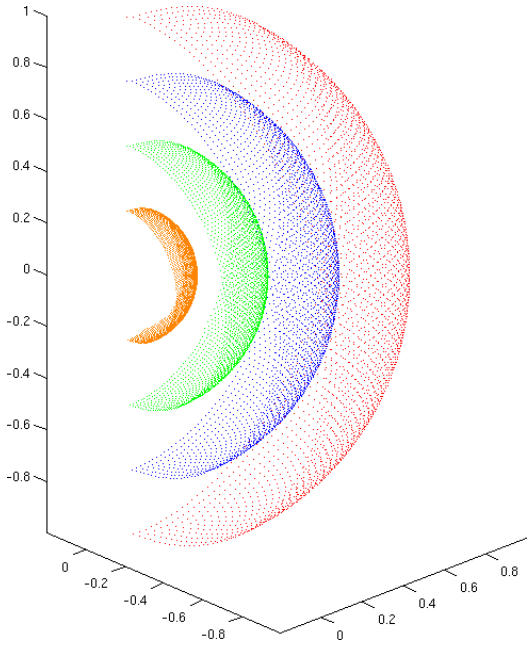

Figure 9.3: Representation of the spherical 3D grid for the Discrete Spherical Fourier-Bessel Transform ($R = 1$ and $N_{max} = 4$)

1) First, from the $a_{\ell mn}$, the inverse discrete SBT is computed for all ℓ and m. This transform can easily be evaluated thanks to a matrix product:

$$\begin{cases} \forall & 0 & \leq & \ell & \leq & L_{max} \\ \forall & -\ell & \leq & m & \leq & \ell \end{cases}, \quad \begin{bmatrix} f_{\ell m}(r_{\ell_0 1}) \\ f_{\ell m}(r_{\ell_0 2}) \\ \vdots \\ f_{\ell m}(r_{\ell_0 N_{max}}) \end{bmatrix} = \frac{\mathbf{T}^{\ell \ell_0}}{R^3} \begin{bmatrix} a_{\ell m 1} \\ a_{\ell m 2} \\ \vdots \\ a_{\ell m N_{max}} \end{bmatrix}. \tag{9.42}$$

Here, it is worth noticing that the matrix $\mathbf{T}^{\ell \ell_0}$ allows the evaluation of the SBT of order ℓ and provides the results on the grid of order ℓ_0.

2) From the spherical harmonics coefficients $f_{\ell m}(r_{\ell_0 n})$ given at specific radial distances $r_{\ell_0 n}$ it is possible to compute the inverse SHT. For each n between 1 and N_{max} the HEALPix inverse SHT is performed on the set of coefficients $\{f_{\ell m}(r_{\ell_0 n})\}_{\ell,m}$. This yields N_{max} HEALPix maps which constitute the sampling of the reconstructed density field on the 3D spherical grid.

DIRECT TRANSFORM Given a density field f sampled on the spherical 3D grid, the SFB coefficients $a_{\ell mn}$ are computed in three steps:

1) For each n between 1 and N_{max} the SHT of the HEALPix map of radius $r_{\ell_0 n}$ is computed. This yields $f_{\ell m}(r_{\ell_0 n})$ coefficients.

2) The next step is to compute the SBT of order ℓ_0 from the $f_{\ell m}(r_{\ell_0 n})$ coefficients for every (ℓ, m). Again, this operation is a simple matrix product:

$$\begin{Bmatrix} \forall & 0 & \leq & \ell & \leq & L_{max} \\ \forall & -\ell & \leq & m & \leq & \ell \end{Bmatrix}, \begin{bmatrix} \widetilde{f}^{\ell_0}_{\ell m}(k_{\ell_0 1}) \\ \widetilde{f}^{\ell_0}_{\ell m}(k_{\ell_0 2}) \\ \vdots \\ \widetilde{f}^{\ell_0}_{\ell m}(k_{\ell_0 N_{max}}) \end{bmatrix} = \frac{\mathbf{T}^{\ell_0 \ell_0}}{K^3} \begin{bmatrix} f_{\ell m}(r_{\ell_0 1}) \\ f_{\ell m}(r_{\ell_0 2}) \\ \vdots \\ f_{\ell m}(r_{\ell_0 N_{max}}) \end{bmatrix}. \tag{9.43}$$

This operation yields $\widetilde{f}^{\ell_0}_{\ell m}(k_{\ell_0 n})$ coefficients which are not yet SFB coefficients because the order of the spherical Bessel coefficients ℓ_0 does not match the order of the spherical harmonics coefficients ℓ. An additional step is necessary.

3) The last step required to gain access to the SFB coefficients $a_{\ell m n}$ is to convert the spherical Bessel coefficients for order ℓ_0 to the correct order ℓ that matches the spherical harmonics order. This is done by using relation Equation (9.38):

$$\begin{Bmatrix} \forall & 0 & \leq & \ell & \leq & L_{max} \\ \forall & -\ell & \leq & m & \leq & \ell \\ \forall & 1 & \leq & n & \leq & N_{max} \end{Bmatrix}, \quad \widetilde{f}_{\ell m}(k_{\ell n}) = \sum_{p=1}^{N_{max}} \widetilde{f}^{\ell_0}_{\ell m}(k_{\ell_0 p}) \frac{2W^{\ell_0 \ell}_{np}}{j^2_{\ell_0+1}(q_{\ell p})}, \tag{9.44}$$

where $W^{\ell_0 \ell}_{np}$ are defined by Equation (9.39). This operation finally yields the $a_{\ell m n} = \widetilde{f}_{\ell m}(k_{\ell n})$ coefficients.

9.4 ISOTROPIC WAVELETS ON THE 3D BALL

The aim of this section is to introduce a new wavelet transform, based on the SFB transform, which extends to the 3D ball the concepts of the 2D Starlet introduced in Chapter 4. Indeed, the isotropic wavelet transform behind the Starlet can be fully defined using isotropic filters which are simple to express in the SFB domain as seen in section Section 9.2.4. Based on this idea, we build an isotropic 3D wavelet on the ball implemented in SFB space, which is only made possible by the practical algorithm introduced in the previous section to evaluate the direct and inverse SFB transform.

9.4.1 Wavelet decomposition

As was detailed in Chapter 4, the key to the discrete wavelet transform is the ability to recursively compute wavelet and approximation coefficients from one resolution to the next by simple convolution with a filter bank (h, g) which eliminated the need to explicitly compute the hermitian product between the signal and the wavelets. When the strict requirements of non-redundant orthogonal wavelet transforms are relaxed, one is free to design a wavelet simply by defining these two filters h and g which must only verify the exact reconstruction formula Equation (4.50).

Taking full advantage of this freedom of design, Starck and Murtagh (2006) proposed a 2D isotropic undecimated wavelet transform entirely defined by an isotropic filter h, and such that

$g = \delta - h$ (which automatically satisfies the perfect reconstruction requirement). As the filter h is linked to the scaling function ϕ through the scaling equation in Equation (4.16), this transform is equivalently defined in terms of a scaling function ϕ. As one particular example of such an isotropic wavelet transform, Starck and Murtagh (2006) introduced the Starlet transform (see Section 4.3.2) entirely defined by the choice of a B-spline of order 3 for the scaling function. This transform was subsequently transposed to the sphere Starck and Murtagh (2006), thus defining an isotropic undecimated wavelet on the sphere, by implementing the convolution with the filters h and g in the spherical harmonics domain. While convolution by a filter on the sphere is generally difficult to perform, it becomes very simple in the spherical harmonics domain when the filter is azimuthaly symmetric.

In much the same way that the 2D undecimated isotropic wavelet was transposed to the sphere, it can now be transposed to the 3D ball since the convolution by an isotropic filter is trivial in SFB space from Equation (9.25).

Our Isotropic 3D wavelet on the ball is defined from an isotropic scaling function $\phi(r, \theta_r, \varphi_r)$. In complete analogy to the wavelet decomposition described in Chapter 4, we use the scaling equation given in Equation (4.16) to define a low-pass filter h as:

$$\frac{1}{\sqrt{2}}\phi\left(\frac{\boldsymbol{r}}{2}\right) = h * \phi(\boldsymbol{r}) \,. \tag{9.45}$$

Thanks to the isotropy of the scaling function, the convolution formula in Equation (9.25) applies and this expression can equivalently be written in the SFB domain where the filter h can be made explicit in terms of the scaling function:

$$\widetilde{h}_{00}(k) = \frac{1}{\pi}\frac{\widetilde{\phi}_{00}(2k)}{\widetilde{\phi}_{00}(k)} \,, \tag{9.46}$$

and $\widetilde{h}_{\ell m} = 0$ for all $(\ell, m) \neq 0$. At a given scale 2^j, we define the rescaled version of this filter as:

$$\widetilde{h}^{(j)}_{00}(k) = \frac{1}{\pi}\frac{\widetilde{\phi}_{00}(2^{j+1}k)}{\widetilde{\phi}_{00}(2^j k)} = \widetilde{h}_{00}(2^j k) \,. \tag{9.47}$$

Following the same approach as for the Starlet, we now define a high-pass filter g as $g = \delta - h$, which implies the following choice of wavelet function $\psi(\boldsymbol{r}) = \phi(\boldsymbol{r}) - \frac{1}{\sqrt{2}}\phi(\frac{\boldsymbol{r}}{2})$. The filter g shares the isotropy of h and its expression in SFB space is also very simple:

$$\widetilde{g}^{(j)}_{00}(k) = \frac{1}{\pi}\left(\frac{1}{\sqrt{2}} - \frac{\widetilde{\phi}_{00}(2^{j+1}k)}{\widetilde{\phi}_{00}(2^j k)}\right) \,. \tag{9.48}$$

Just as for the 2D undecimated wavelet transform, the approximation and detail coefficients can recursively be computed from one resolution to the next by applying the filter bank (h, g):

$$a^{j+1}(\boldsymbol{r}) = h^{(j)} * a^j(\boldsymbol{r}) \,, \tag{9.49}$$

$$d^{j+1}(\boldsymbol{r}) = g^{(j)} * a^j(\boldsymbol{r}) \,, \tag{9.50}$$

The strength of the Cartesian discrete undecimated wavelet transform was the ability to efficiently compute this convolution product in direct space, which is no longer the case in a spherical coordinate system. Nevertheless, we can still evaluate this product in SFB space, which yields:

$$\widetilde{a}^{j+1}_{\ell m}(k) = \sqrt{2}\pi\, \widetilde{h}^{(j)}_{00}(k)\, \widetilde{a}^j_{\ell m}(k) \,, \tag{9.51}$$

$$\widetilde{d}^{j+1}_{\ell m}(k) = \sqrt{2}\pi\, \widetilde{g}^{(j)}_{00}(k)\, \widetilde{a}^j_{\ell m}(k) \,, \tag{9.52}$$

These relations allow us to implement the wavelet decomposition recursively directly in SFB space. However, in order to apply these formulae, the SFBT of the signal $f(r)$ must first be performed and then the wavelet coefficients themselves $d^j(r)$ must be recovered from their SFB transform $d^j_{\ell m}(k)$. This is made possible by the discrete SFBT introduced in the previous section and its associated sampling of the 3D ball.

Consider a function f defined on the ball, such that $f(R) = 0$, then a discrete set of SFB coefficients can be computed using only a discrete sampling of f on the ball as described in Section 9.3.2. This leads to the complete algorithm for computing the 3D isotropic wavelet on the ball presented in Algorithm 9.1. Using this algorithm, we can in practice compute the wavelet decomposition $\mathcal{W} = \{d^1, d^2, \cdots, d^J, a^J\}$ of a function f sampled on the 3D spherical grid introduced in the previous section.

Algorithm 9.1 Discrete 3D isotropic wavelet on the ball

Require:
Input data $f(r, \theta, \varphi)$ sampled on the spherical grid defined in Section 9.3.2.
Number of scales J.

Set $a^0 = \mathrm{f}$.
Tabulate the filter $h^{(j)}_{00}(k_{0n})$ for $j \in [1, J]$.
Compute $a^0_{\ell mn}$, the discrete SFBT of a^0.

for $j = 0$ to $J - 1$ **do**
    Compute approximation coefficients: $\widetilde{a}^{j+1}_{\ell mn} = \sqrt{2}\pi\, \widetilde{h}^{(j)}_{00}(k_{0n})\widetilde{a}^j_{\ell mn}$.
    Compute wavelet coefficients: $\widetilde{d}^{j+1}_{\ell mn} = \widetilde{a}^j_{\ell mn} - \widetilde{a}^{j+1}_{\ell mn}$.
end for
Compute the inverse discrete SFBT of $\{d^1_{\ell mn}, d^2_{\ell mn}, \cdots, d^J_{\ell mn}, a^J_{\ell mn}\}$ to get $\{d^1, d^2, \cdots, d^J, a^J\}$.

return The wavelet decomposition $\{d^1, d^2, \cdots, d^J, a^J\}$, sampled on the 3D spherical grid.

9.4.2 Inverse Transform

Since the wavelet coefficients are defined as the difference between two resolutions, the reconstruction from the wavelet decomposition $\mathcal{W} = \{d^1, \ldots, d^J, a^J\}$ is straightforward and corresponds to the reconstruction formula of the *à trous* algorithm:

$$a^0 = a^J + \sum_{j=1}^{J} d^j \ . \tag{9.53}$$

However, given the redundancy of the transform, the reconstruction is not unique. It is possible to take advantage of this redundancy to reconstruct a^j from a^{j+1} and d^{j+1} by using a least squares estimate.

From the recursive wavelet decomposition defined in Equation (9.51) and Equation (9.52), by respectively multiplying these equations by $\widetilde{h}^{*j}_{\ell m}(k)$ and $\widetilde{g}^{*j}_{\ell m}(k)$, the following expression is obtained for the least squares estimate of a^j from a^{j+1} and d^{j+1}:

$$\widetilde{a}^j_{\ell m}(k) = \widetilde{a}^{j+1}_{\ell m}(k)\widetilde{H}^{(j)}_{\ell m}(k) + \widetilde{d}^{j+1}_{\ell m}\widetilde{G}^{(j)}_{\ell m}(k) \ , \tag{9.54}$$

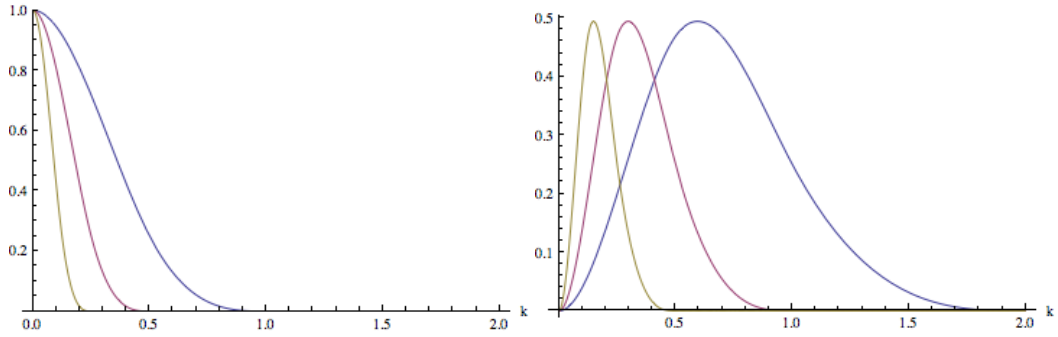

(a) Scaling function $\widetilde{\phi}_{00}^{2^{-j}k_c}(k)$ for $j = 0, 1, 2$ **(b)** Wavelet function $\widetilde{\psi}_{00}^{2^{-j}k_c}(k)$ for $j = 0, 1, 2$

Figure 9.4: Scaling function and Wavelet function for $k_c = 1$

where $\widetilde{H}^{(j)}$ and $\widetilde{G}^{(j)}$ are defined as follows:

$$\widetilde{H}_{\ell m}^{(j)}(k) = \frac{1}{\sqrt{2}\pi} \frac{\widetilde{h}_{\ell m}^{*j}(k)}{|\widetilde{h}_{\ell m}^{j}(k)|^2 + |\widetilde{g}_{\ell m}^{j}(k)|^2} , \tag{9.55}$$

$$\widetilde{G}_{\ell m}^{(j)}(k) = \frac{1}{\sqrt{2}\pi} \frac{\widetilde{g}_{\ell m}^{*j}(k)}{|\widetilde{h}_{\ell m}^{j}(k)|^2 + |\widetilde{g}_{\ell m}^{j}(k)|^2} , \tag{9.56}$$

for $(\ell, m) = 0$, otherwise $\widetilde{H}_{\ell m}^{(j)}(k) = 0$ and $\widetilde{G}_{\ell m}^{(j)}(k) = 0$.

9.4.3 Choice of a scaling function

Any function with spherical symmetry and a cut-off frequency k_c would do as a scaling function but we choose to use a B-spline function of order 3 to define our scaling function:

$$\widetilde{\phi}_{\ell m}^{k_c}(k) = \frac{3}{2} B_3\left(\frac{2k}{k_c}\right) \delta_{\ell 0}\delta_{m0} . \tag{9.57}$$

where

$$B_3(x) = \frac{1}{12}\left(|x-2|^3 - 4|x-1|^3 + 6|x|^3 - 4|x+1|^3 + |x+2|^3\right) . \tag{9.58}$$

This scaling function and its corresponding wavelet function are plotted in SFB space for different values of j in Figure 9.4

Other functions such as Meyer wavelets or the needlet function Marinucci et al. (2008) can be used as well. In particular, needlet wavelet functions have a much better frequency localisation than the wavelet function derived from the B_3-spline, but the price to pay is more oscillations in the direct space. To illustrate this, we show in Figure 9.5 two different wavelet functions. Figure 9.5 left shows the 1D profile of the spline (continuous line) and needlet (dotted line) wavelet functions at a given scale. Figure 9.5 right shows the same function, but we have plotted the absolute value in order to better visualize their respective ringing. As can be seen, for wavelet functions with the same main lob, the needlet wavelet oscillate much more than the spline wavelet. Hence, the best wavelet choice certainly depends on the final applications. For statistical analysis, detection or restoration applications, we may prefer to use a wavelet which does not oscillate too much and with a smaller support, and the spline wavelet is clearly the correct choice. For spectral or bispectral

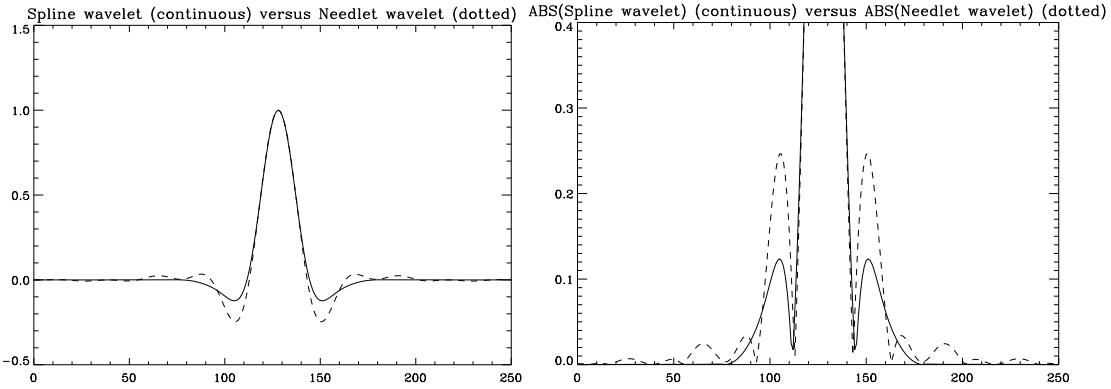

Figure 9.5: Comparaison between spline and needlet wavelet functions on the sphere.

analysis, where the frequency localization is fundamental, then needlet should be preferred to the spline wavelet.

To illustrate this wavelet transform, a set of SFB coefficients was extracted from a 3D density field using our discrete SFBT algorithm. The test density field, described more in details in the next section, was provided by a cosmological n-body simulation which was carried out by the Virgo Supercomputing Consortium using computers based at Computing Centre of the Max-Planck Society in Garching and at the Edinburgh Parallel Computing Centre[4].
The wavelet decomposition presented above can then be computed from the SFB coefficients of the test density field and yields the SFB coefficients of the various wavelet scales and smoothed away density. Using the inverse discrete SFBT, the actual wavelet coefficients can be retrieved in the form of 3D density fields. These density fields are shown on Figure 9.6.

9.4.4 Toy experiment

In this section, we present a simple wavelet denoising application on a density field in spherical coordinates using the Isotropic 3D Wavelet transform on the ball.

Denoising using sparse transforms can be performed very easily, by applying a simple thresholding on the coefficients. One can use a *soft* or *hard* thresholding (see Chapter 3) according to whether we want more accuracy or less artefacts. The threshold level is typically such that, for an additive gaussian noise, the thresholding operator kills all noise coefficients except a small percentage, keeping only the significant coefficients which contain most of the information. The threshold we use is typically a simple $\kappa\sigma$, with $\kappa \in [3, 4]$, which corresponds respectively to 0.27% and $6.3 \cdot 10^{-5}$ false detections. Other methods exist, that estimate automatically the threshold to use in each band like the False Discovery Rate (see Benjamini and Hochberg (1995); Miller et al. (2001)). The correlation between neighbour coefficients intra-band and/or inter-band may also be taken into account (see Sendur and Selesnick (2002b,a)).

This experiment is performed on the same N-body simulation from the Virgo Consortium as the one presented in the previous section on Figure 9.6. The Virgo large box simulation[5] provides us with a Cartesian density cube. The SFB coefficients of the test density field are first computed by

4 The data is publicly available at http://www.mpa-garching.mpg.de/Virgo/VLS.html

5 a ΛCDM simulation at $z = 0$, which was calculated using 512^3 particles for the following cosmology: $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$, $H_0 = 70\text{km s}^{-1}\text{Mpc}^{-1}$,$\sigma_8 = 0.9$. The data cube provided is 479 h^{-1}Mpc in length.

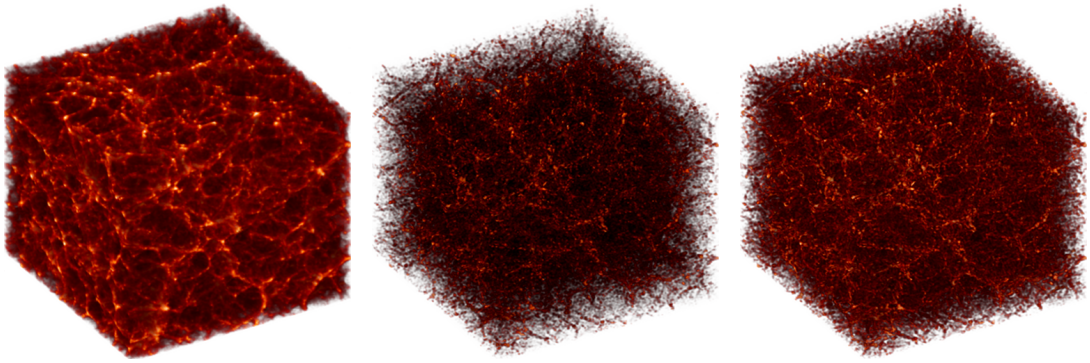

(a) Density from the initial spherical Fourier-Bessel coefficients

(b) First wavelet scale

(c) Second wavelet scale

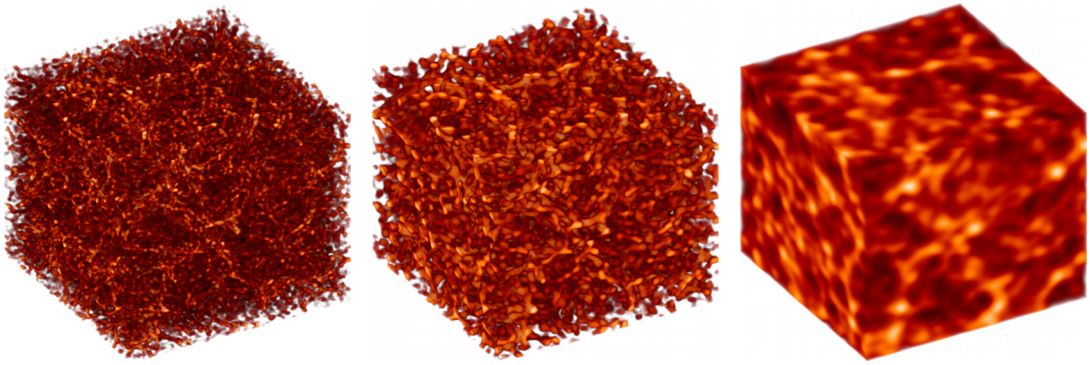

(d) Third wavelet scale

(e) Fourth wavelet scale

(f) Smooth approximation

Figure 9.6: Isotropic Spherical 3D wavelet decomposition of a density field. Only a cube at the center of the spherical field is displayed.

sampling the Virgo density field on the spherical 3D grid illustrated on Figure 9.3, for $n_{side} = 2048$, $l_{max} = 1023$ and $n_{max} = 512$. In order to perform the SFB decomposition, the observer is placed at the center of the box, and the SFB coefficients are calculated out to $R = 479/2$ h^{-1}Mpc, setting the density field to zero outside of this spherical volume.

A Gaussian noise was then added to the SFB coefficients to produce a noisy density field. Figure 9.7a and Figure 9.7b show the central portion of slices taken in the middle of respectively the original and noisy spherical density fields. The level of the noise is comparable to the amplitude of the faint filamentary structures that can be seen in the original density field on Figure 9.7a. Using Hard Thresholding of the wavelet coefficients, the noisy field is filtered to yield the restored density displayed on Figure 9.7c. The residuals after denoising are shown on Figure 9.7d. The artificially added noise is successfully removed, without much loss to the large scale structure, though some of the smaller filamentary structures are removed. This however is to be expected given the isotropic nature of the wavelet transform used here, better suited to restore more isotropic features such as clusters.

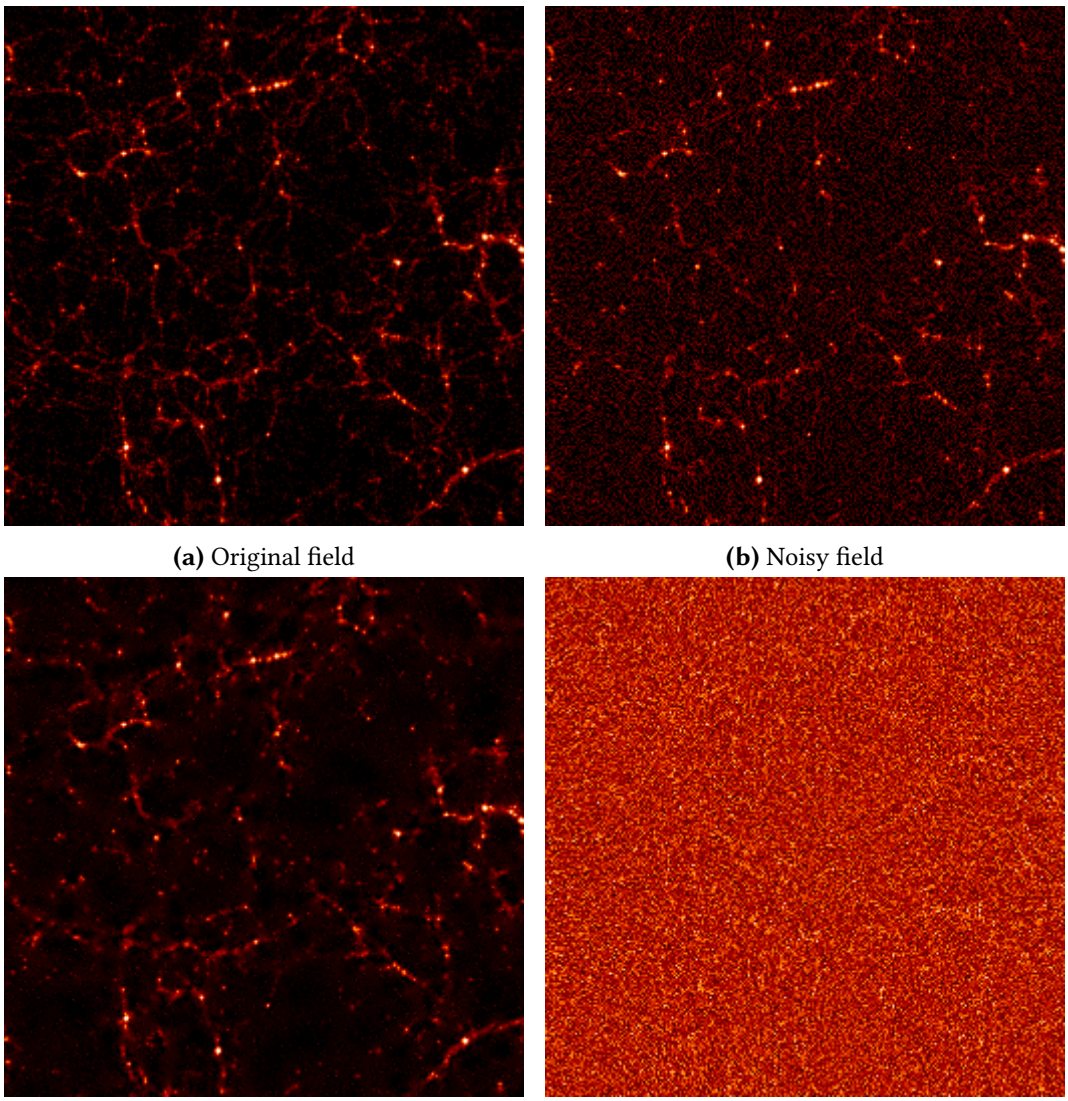

(a) Original field **(b)** Noisy field

(c) Denoised field **(d)** Residuals

Figure 9.7: Isotropic Undecimated Spherical 3D wavelet Hard thresholding applied to a test density field.

9.5 CONCLUSION

Modern cosmology requires the analysis of 3D fields on large areas of the sky, which are best described in spherical coordinates. In this configuration, the *Spherical Fourier-Bessel Transform* (SFBT) is the most natural and effective way to statistically analyse the field. As an example, we will see in Chapter 10 that an SFB analysis is better suited than a tomographic analysis in spherical shells to extract cosmological information from spectroscopic galaxy surveys.

On the other hand, as was illustrated throughout this thesis, wavelets are fundamental tools of modern signal processing. In particular, combined with the sparse regularisation framework introduced in Chapter 3, they find many applications such as denoising, deconvolution or inpainting. Therefore, there is great interest in developing wavelets adapted to the analysis of cosmological

fields. For instance, the extension of wavelets to the sphere has lead to many applications for the analysis of the CMB.

In this chapter, we have presented a new wavelet on the 3D ball, which is based on the SFBT, and is therefore ideally suited for the analysis of large scale galaxy surveys. In order to derive this wavelet, we have first presented a new sampling scheme of the 3D ball which enables the back and forth computation of the discrete SFBT for radially limited fields. The derivation of this discrete SFBT has required the introduction of a new discrete algorithm for the evaluation of the 1D *Spherical Bessel Transform* (SBT), which can find many applications, including the computation of theoretical SFB power spectra of galaxy clustering as will be illustrated in the next chapter.

To the best of our knowledge, the wavelet presented in this chapter is the only *isotropic* wavelet on the 3D ball. Furthermore, due to our particular choice of scaling function, our wavelet exhibit *minimal oscillations* in direct space. These two properties are very desirable for applications such as the detection of voids or galaxy clusters in galaxy surveys.

10 | 3D ANALYSIS OF GALAXY CLUSTERING

CONTENTS

In this chapter, we compare two different types of 3D analyses of spectroscopic galaxy surveys. Although there is theoretical motivation and novel tools to explore these surveys using a 3D *Spherical Fourier-Bessel* (SFB) expansion of the data (introduced in Chapter 9), most survey optimisations and forecasts are based on a tomographic spherical harmonics analysis. Our aim is to compare these two methods in terms of cosmological parameter constraints and in particular their robustness to unavoidable systematics due to uncertainties on the galaxy bias.

We find that in the absence of systematics on the galaxy bias, both techniques can recover comparable constraints. However, the 3D SFB seems to be more robust to this source of systematics, suggesting that this approach should be preferred for the design and analysis of future wide-field spectroscopic surveys.

The results presented in this chapter were published in Lanusse et al. (2015). This analysis has been conducted in collaboration with Anais Rassat and Jean-Luc Starck.

10.1 INTRODUCTION

Understanding the nature of the dark universe is one of the fundamental challenges of modern cosmology today. Galaxy clustering - the statistical analysis of the spatial distribution of galaxy number counts - has been identified as one of the most promising probes available to explore

this (Peebles, 1980; Albrecht et al., 2006; Peacock et al., 2006), with spectroscopic surveys being particularly useful in probing both tangential and radial modes in the Universe.

Galaxy number counts have been extensively studied with current and planned future surveys, and the analysis can be performed in various spaces, for example, Fourier space (Seo and Eisenstein, 2003, 2007), configuration space (Eisenstein et al., 2005; Slosar et al., 2009; Xu et al., 2010), and spherical harmonic space (e.g., Dolney et al., 2006; Kirk et al., 2012). For future wide-field spectroscopic surveys, the galaxy field will cover large areas on the sky so that an analysis in spherical space provides a natural decomposition for certain physical effects as well as selection effects. For wide-field spectroscopic surveys, the depth of the survey means that a 3D spherical Fourier-Bessel (SFB) analysis is the most natural to perform (Fisher et al., 1995; Heavens and Taylor, 1995; Rassat and Refregier, 2012).

Previous SFB analyses of the local Universe (e.g., Erdogdu et al., 2006a,b) used relatively small data sets, where straightforward summation methods were sufficient to measure the SFB coefficients. Today, novel numerical methods for 3D spherical analysis are available (Leistedt et al., 2012; Lanusse et al., 2012) to prepare for future wide-field surveys that will map the large-scale structure of the Universe with a large number of galaxies. The 3D SFB analysis can also be applied to other probes, for instance weak-lensing (Heavens, 2003; Castro et al., 2005; Kitching et al., 2008b, 2011; Merkel and Schafer, 2013; Grassi and Schafer, 2014; Kitching et al., 2014) and the integrated Sachs-Wolfe effect (e.g. Shapiro et al., 2012), which will be crucial for high-precision probe combinations.

Even given the existing 3D SFB tools and the theoretical motivation for this approach, most existing forecasts and survey optimisation for future wide-field surveys focus on a tomographic analysis, that is, one where the survey is split into redshift bins, and 2D spherical harmonic auto- and cross-power spectra $C^{ij}(\ell)$ are measured (e.g., Laureijs et al., 2011; Kirk et al., 2012). One of the advantages of a tomographic spherical harmonics analysis is that there are several available codes to rapidly calculate the tomographic spectra, either for galaxy correlations or for other complementary probes (e.g., Di Dio et al., 2013; Refregier et al., 2011); another advantage is that it is straightforward to convert survey observables (θ, ϕ, z) into a power spectrum measurement without assuming a distance-redshift relation (which can be based on a fiducial cosmology), while the 3D SFB analysis requires such a relation to translate the observables into the 3D SFB spectrum. However, in the tomographic analysis, some of the radial information may be lost as a result of redshift binning, while the 3D SFB analysis potentially uses the entire 3D information, especially for a spectroscopic survey. With this in mind, a natural hypothesis is that a 3D SFB spectroscopic analysis might extract more information than a tomographic one.

Several studies have already investigated this, for example, Di Dio et al. (2014) and Asorey et al. (2012) found that a tomographic analysis returned equivalent or better constraints than a 3D Fourier power spectrum analysis. They concluded that the tomographic approach should be preferred as it avoids the need to assume a particular cosmology to convert redshifts into comoving distances and simplifies the combination with other probes such as weak-lensing. Nevertheless, they both acknowledged that for a spectroscopic survey the tomographic analysis would require a large number of redshift bins to recover the full 3D information, which is limited by shot noise problems. For the first time, Nicola et al. (2014) compared the tomographic analysis to a 3D SFB analysis and found the tomographic constraints to be superior, but they still noted that the 3D SFB approach was stable with regard to the choice of fiducial cosmology for the necessary conversion from redshift to comoving distance. However, their treatment of the non-linear scale cut-off

used in the Fisher matrix comparison is not equivalent between the tomographic and the 3D SFB analysis. While the 3D power spectra is cut off at a physical scale (k) corresponding to nonlinear effects, the tomographic power spectra are truncated at fixed arbitrary angular scales. This ignores the interplay between the redshift of the tomographic bins and the wavenumber of the SFB spectrum. As a result, the non-linear cut-off in Nicola et al. (2014) does not allow a fair comparison between 2D and 3D methods, which means understanding the strength of each method is still an open question. We address this question by carefully excluding non-linear scales.

Understanding how best to extract information for a 3D galaxy survey is of utmost importance to address the fundamental questions in modern cosmology today, and also to ensure that future planned surveys are efficiently analysed as well as optimised. To address this pressing question, we propose here a new investigation of the information that can be extracted from a spectroscopic galaxy survey by tomographic vs. 3D SFB analysis. Our approach focuses on the seven common parameters that are currently used in wide-field survey optimisation and planning, that is, on $\vec{\theta} = \{\Omega_m, h, w_0, w_a, \sigma_8, \Omega_b, n_s\}$, while putting forward a coherent approach regarding the exclusion of non-linear scales for both the 2D and 3D methods for the first time. In addition, we investigate for the first time how tomographic and 3D SFB methods are affected by nuisance parameters related to the galaxy bias, which we allow to be both redshift- and scale-dependent. However, we do not include *Redshift Space Distortions* (RSD) or relativistic effects in our study. Including RSDs, which will be present in the data, provides an additional probe, which improves constraints. Although a prescription for RSDs in SFB space exists (Heavens and Taylor, 1995), as a first approach, we do not include them here in either the tomographic or the SFB analysis to ensure that we compare like with like. Their impact should nonetheless be assessed, which we plan to do in a future work.

Finally, in the spirit of reproducible research, we make available all tomographic and 3D SFB codes used for this analysis, along with the scripts to reproduce our results.

This chapter is structured as follows: in Section 10.2, we briefly review the theory behind the statistical analysis of galaxy number counts, including the prescription for the tomographic analysis and the 3D SFB. In Section 10.3, we provide an overview of the Fisher matrix forecasting-approach that we used to compare the relative constraining power of each method, and include the description of the future spectroscopic wide-field survey for which we calculate forecasts, of the question of non-linear scale treatment, and galaxy bias nuisance parameters. In Section 10.4, we present the comparison between the constraining power of the 3D SFB and tomographic methods and investigate how this comparison holds in the presence of galaxy bias nuisance parameters. We also determine how this affects a future wide-field survey optimisation. In Section 10.5, we present our conclusions in the context of high-precision cosmology with future wide-field surveys.

10.2 MODELLING OF SPECTROSCOPIC GALAXY SURVEYS

In this section we describe the formalism behind the analysis of galaxy clustering in the context of a spectroscopic survey. We present the two methodologies compared in this chapter, one based on a tomographic analysis of angular correlations, the other based on the correlations of the 3D expansion of the galaxy field on a spherical Fourier-Bessel basis.

10.2.1 Galaxy and matter fields

In a galaxy survey, the quantity observed is the galaxy number density $n(\mathbf{r} = (r, \theta, \varphi))$, which can be defined in terms of the galaxy overdensity δ_g through

$$n(\mathbf{r}) = \bar{n}(r)(1 + \delta_g(\mathbf{r}, z(r))) \,, \tag{10.1}$$

where $\bar{n}(r)$ is the mean number density of observed galaxies at comoving distance r. In this expression, the time dependence of the observed overdensity as a function of comoving distance is made explicit through the $z(r)$ relation. The mean number density $\bar{n}(r)$ can be expressed in terms of the survey selection function $\phi(r)$ as

$$\bar{n}(r) = \phi(r)\bar{n} = \frac{N}{V}\phi(r) \,, \tag{10.2}$$

with $\bar{n}$ the mean number density of observed galaxies, N the total number of observed galaxies, and V the volume of the survey that fulfils $V = \int \phi(r)\mathrm{d}\mathbf{r}$. Note that in the general case, the selection function has both an angular and a radial dependence (see Section 10.2.4), but we do not consider here the full impact of an angular mask and we will only account for partial coverage of the sky through a multiplicative f_{sky} factor.

In Equation (10.1), the time (or redshift) dependence of the galaxy overdensity is due to the growth of structure and the evolution of galaxy bias with respect to the matter density field with time. Following the approach of Rassat and Refregier (2012), in the linear regime this dependence on redshift can be separated in the form of growth and bias prefactors,

$$\delta_g(\mathbf{r}, z(r)) = b(r, k)D(r)\delta(\mathbf{r}) + \epsilon(\mathbf{r}) \,, \tag{10.3}$$

where $b(r, k)$ is a bias with a possible scale dependence, $D(r)$ is the growth factor, $\delta(\mathbf{r}) = \delta(\mathbf{r}, z = 0)$ is the matter overdensity field at present day, and $\epsilon(\mathbf{r})$ is a Poisson noise term arising from the discrete nature of the observed galaxy number density. As in Rassat et al. (2008), we considered the linear relation Equation (10.3) to hold in the standard cosmological model on large scales up to a redshift-dependent $k_{\max}(z)$ with $k_{\max}(z = 0) \simeq 0.12h\mathrm{Mpc}^{-1}$ and $k_{\max}(z = 2) \simeq 0.25h\mathrm{Mpc}^{-1}$. We then proceeded to define a modified selection function that includes the effects of bias and growth in the linear regime,

$$\phi^{\mathrm{evol}} = b(r, k)D(r)\phi(r) \,. \tag{10.4}$$

Using this modified selection function, the observed galaxy density can now be expressed directly as a function of the true matter overdensity at present time:

$$\frac{n(\mathbf{r})}{\bar{n}} = \phi(r) + \phi^{\mathrm{evol}}(r, k)\delta(\mathbf{r}) + \phi(r)\epsilon(\mathbf{r}) \,. \tag{10.5}$$

10.2.2 Tomographic analysis of galaxy clustering

In the tomographic analysis, the survey is decomposed into spectroscopic redshift bins from which are computed classical angular correlation functions. The angular number density for one spectroscopic bin (i) limited between $z^{(i)}_{\min}$ and $z^{(i)}_{\max}$ is defined as

$$n^{(i)}(\boldsymbol{\theta}) = \bar{n}^{(i)}\left(1 + \delta^{(i)}(\boldsymbol{\theta})\right) = \int_{z^{(i)}_{\min}}^{z^{(i)}_{\max}} n(z, \boldsymbol{\theta})\mathrm{d}z \,, \tag{10.6}$$

where $\bar{n}^{(i)}$ is the average galaxy number density per steradians in tomographic bin (i) and $\delta^{(i)}(\boldsymbol{\theta})$ is the angular galaxy overdensity in bin (i). Expanding the angular overdensity in spherical harmonics yields

$$n^{(i)}_{\ell m} = \int n^{(i)}(\boldsymbol{\theta}) Y^*_{\ell m}(\boldsymbol{\theta}) \mathrm{d}\boldsymbol{\theta} \,, \tag{10.7}$$

From this spherical harmonics expansion, the tomographic angular correlation functions between bins (i) and (j), noted $\overline{C}^{(ij)}_\ell$, is defined for $\ell \geq 1$ as

$$\overline{C}^{mm'(ij)}_{\ell\ell'} \equiv \frac{1}{\bar{n}^{(i)2}} < n^{(i)}_{\ell m} n^{(j)*}_{\ell' m'} > \,, \tag{10.8}$$

$$= \left(C^{(ij)}_\ell + \frac{\delta^K_{ij}}{\bar{n}^{(i)}} \right) \delta^K_{\ell\ell'} \delta^K_{mm'} \,, \tag{10.9}$$

where δ^K is the Kronecker symbol. In the last equation, the first term $C^{(ij)}_\ell$ is the contribution from galaxy clustering and the second term $\frac{1}{\bar{n}^{(i)}}$ is the contribution from shot noise, which only affects the auto-correlation power spectra. Note that different angular modes are predicted to be uncorrelated in linear theory for a Gaussian random field; these can become correlated as a result of non-linearities or lack of full-sky coverage, effects that we did not consider in this work. Formally, the correlation functions $C^{(ij)}_\ell$ are related to the matter power spectrum $P(k)$ at $z = 0$, in the linear regime, according to

$$C^{(ij)}_\ell = \frac{2}{\pi} \int \mathrm{d}k P(k) k^2 \int w^{(i)}_{\mathrm{evol}}(r,k) j_\ell(kr) \mathrm{d}r \int w^{(j)}_{\mathrm{evol}}(r',k) j_\ell(kr') \mathrm{d}r' \,, \tag{10.10}$$

where $w^{(i)}_{\mathrm{evol}}$ is a window function for bin (i), which includes the effects of spectroscopic selection, linear growth, and bias:

$$w^{(i)}_{\mathrm{evol}} = \phi^{\mathrm{evol}}(r,k) s^{(i)}(r) \,, \tag{10.11}$$

with ϕ^{evol} is the modified selection function including growth and bias introduced in Equation (10.4) and $s^{(i)}$ is the spectroscopic selection function that defines the redshift bin i , that is, $s^{(i)}(z) = 1$ if $z \in [z^{(i)}_{\mathrm{min}}, z^{(i)}_{\mathrm{max}}]$, $s^{(i)}(z) = 0$ otherwise.

This expression is the full general expression of the tomographic angular power spectrum. However, it is common to evaluate the angular power spectrum through the well-known Limber approximation. To the first order (Loverde and Afshordi, 2008), the Limber approximation applied to the previous equation yields

$$C^{(ij)}_{\ell|_{\mathrm{Limber}}} = \int \frac{\mathrm{d}r}{r^2} P\left(\frac{\ell+1/2}{r}\right) w^{(j)}_{\mathrm{evol}}\left(r, \frac{\ell+1/2}{r}\right) w^{(i)}_{\mathrm{evol}}\left(r, \frac{\ell+1/2}{r}\right) \,. \tag{10.12}$$

The Limber approximation holds to very good accuracy for the auto-correlations under the assumption that the bin window functions do not vary too rapidly or that the overlap between bins is not too small.

10.2.3 3D spherical Fourier-Bessel analysis of galaxy clustering

The Spherical Fourier-Bessel decomposition is the natural expansion of 3D fields described in spherical coordinates and was introduced in details in Section 9.2.3. We now present how this

expansion can be applied to the analysis of galaxy clustering. The *Spherical Fourier-Bessel Transform* (SFBT) of the galaxy number density $n(\mathbf{r})$ is defined as

$$n_{\ell m}(k) = \sqrt{\frac{2}{\pi}} \int n(\mathbf{r}) k j_\ell(kr) Y^*_{\ell m}(\theta, \varphi) \mathrm{d}\mathbf{r} \,, \tag{10.13}$$

where j_ℓ are spherical Bessel functions, $Y_{\ell m}$ are spherical harmonics, ℓ and m are multipole moments, and k is the wavenumber. Note that contrary to the symmetric convention adopted in the previous chapter (which matched for instance the formalism of Heavens and Taylor (1995), see Section 9.2.2), we use here the widespread orthonormal convention for the SFB, as in Rassat and Refregier (2012), Fisher et al. (1995), or Pratten and Munshi (2013). From the SFB coefficients $n_{\ell m}(k)$, the number density can be recovered through the inverse SFB transform as

$$n(r, \theta, \varphi) = \sqrt{\frac{2}{\pi}} \sum_{\ell, m} \int n_{\ell m}(k) k j_\ell(kr) \mathrm{d}k Y_{\ell m}(\theta, \varphi) \,. \tag{10.14}$$

Although the SFB expansion is performed in comoving space, in practice, the galaxy number density is only observed in redshift space. This means that a fiducial cosmology has to be assumed to relate observed redshift and comoving distance of the galaxies in the survey. To distinguish between true comoving distance r and estimated comoving distance, we introduce the notation

$$\tilde{r} \equiv r_{|\mathrm{fid}}(z) \,. \tag{10.15}$$

When the fiducial cosmology exactly corresponds to the true cosmology, $\tilde{r} = r$, but in general, this is not the case. The importance of making this distinction has been stressed in Heavens et al. (2006), especially when constraining dark energy parameters, which are very sensitive to the $r(z)$ relation.

For multipoles of order $\ell \geq 1$, the 3D SFB spectrum of the observed galaxy density can be expressed in the form

$$\overline{C}^{mm'}_{\ell\ell'}(k, k') \equiv \frac{1}{\bar{n}^2} < n_{\ell m}(k) n^*_{\ell' m'}(k') > \,, \tag{10.16}$$

$$= \left(C_\ell(k, k') + N_\ell(k, k')\right) \delta_{\ell\ell'} \delta_{mm'} \,. \tag{10.17}$$

This expression can be directly compared to the definition of the tomographic power spectra in Equation (10.9). Just like in the tomographic case, different angular multipoles are not correlated when an angular mask is neglected. In this expression, the signal power spectrum $C_\ell(k, k')$ takes the form (see Rassat and Refregier (2012) for this exact prescription or Heavens and Taylor (1995))

$$C_\ell(k, k') = \left(\frac{2}{\pi}\right)^2 \int k''^2 P(k'') W^{\mathrm{evol}}_\ell(k, k'') W^{\mathrm{evol}}_\ell(k', k'') \mathrm{d}k'' \,, \tag{10.18}$$

where the following window function includes the effects of linear growth and bias and the fiducial redshift-comoving distance relation:

$$W^{\mathrm{evol}}_\ell(k, k'') = k \int \phi^{\mathrm{evol}}(r, k'') j_\ell(k\tilde{r}) j_\ell(k''r) r^2 \mathrm{d}r \,. \tag{10.19}$$

The noise covariance matrix can be expressed as

$$N_\ell(k, k') = \frac{2kk'}{\bar{n}\pi} \int \phi(r) j_\ell(k\tilde{r}) j_\ell(k'\tilde{r}) r^2 \mathrm{d}r \,. \tag{10.20}$$

This expression is equivalent to that used in Yoo and Desjacques (2013) and a derivation can be found in Section D.3.

When considering a realistic galaxy survey with finite depth, the observed galaxy number density vanishes above a given $r_{\max}$ and fulfils

$$\forall(\theta,\phi), \quad n(r_{\max},\theta,\phi) = 0 \,. \tag{10.21}$$

Under this boundary condition, the spherical Fourier-Bessel transform can be inverted from discretely sampled coefficients $n_{\ell m}(k_{\ell n})$ and Equation (10.14) becomes

$$n(r,\theta,\phi) = \sum_{\ell,m,n} \kappa_{\ell n} n_{\ell m}(k_{\ell n}) k_{\ell n} j_\ell(k_{\ell n} r) Y_{\ell m}(\theta,\varphi) \,, \tag{10.22}$$

where the discrete wavenumbers $k_{\ell n}$ are defined in terms of the zeros of the spherical Bessel function $q_{\ell n}$ as

$$k_{\ell n} = \frac{q_{\ell n}}{r_{\max}} \,, \tag{10.23}$$

and the normalisation factors $\kappa_{\ell n}$ are defined as $\kappa_{\ell n} = \frac{\sqrt{2\pi} r_{\max}^{-3}}{j_{l+1}^2(q_{ln})}$ (Fisher et al., 1995).

In the context of Fisher matrix forecasting, the main consequence of this discretisation is that it imposes a discrete sampling of the SFB spectrum that can be represented in matrix form $C_\ell(n,n') = C_\ell(k_{\ell n}, k_{\ell n'})$ without loss of information. Furthermore, our discrete *Spherical Bessel Transform* (SBT) algorithm defined in Section 9.3.1 applies in this case and allows us to efficiently compute $C_\ell(n,n')$ in practice.

10.2.4 Effect of partial sky coverage

So far, we have assumed complete sky coverage. However, obscuration and confusion due to our own galaxy means that only a portion of the sky is observable in practice. This effect can be modelled in a similar way for both tomographic and SFB derivations by applying an angular weighting function $M(\theta,\varphi)$ to the galaxy density field, for instance with $M(\theta,\varphi) = 0$ in masked areas and $M(\theta,\varphi) = 1$ otherwise. The effect of such a mask on angular power spectra is well known and results in a coupling of angular modes that would otherwise remain uncorrelated. Formally, the signal part of both tomographic and SFB power spectra becomes

$$C_{\ell\ell'}^{mm'(ij)} = \sum_{\ell''m''} M_{\ell m \ell'' m''} M_{\ell' m' \ell'' m''} C_{\ell''}^{(ij)} \,, \tag{10.24}$$

$$\begin{aligned} C_{\ell\ell'}^{mm'}(k,k') = \sum_{\ell''m''} M_{\ell m \ell'' m''} M_{\ell' m' \ell'' m''} \iint K_{\ell\ell''}(k,k_1) K_{\ell'\ell''}(k',k_2) \\ \times \quad C_{\ell''}(k_1,k_2) \quad \mathrm{d}k_1 \mathrm{d}k_2 \,, \end{aligned} \tag{10.25}$$

where $M_{\ell m \ell'' m''}$ is an angular coupling kernel defined in terms of the angular mask as

$$M_{\ell m \ell' m'} = \int_\Omega Y_{\ell' m'}(\Omega) M(\Omega) Y_{\ell m}^*(\Omega) \,\mathrm{d}\Omega \,, \tag{10.26}$$

and $K_{\ell\ell'}(k,k')$ is a wavenumber coupling kernel defined as

$$K_{\ell\ell'}(k,k') = \frac{2}{\pi} k'^2 \int_r j_\ell(kr) j_{\ell'}(k'r) r^2 \,\mathrm{d}r \,. \tag{10.27}$$

The noise part of the power spectra due to shot noise can be approximated as only affected by a simple area scaling (Kitching et al., 2014). Note that the coupling matrix $K_{\ell\ell'}(k,k')$ reduces to a Dirac delta function when $\ell = \ell'$ thanks to the orthogonality of the spherical Bessel functions. In both cases, the mask will induce a coupling of angular modes. In the SFB case, this also means that the coupling kernels $K_{\ell\ell'}(k,k')$ can no longer be considered as Dirac delta functions and induce an additional coupling between different wavenumbers.

In practice, the impact of the mask can conveniently be taken into account using the pseudo-C_ℓ methodology, which is well known for studies of the CMB (Hivon et al., 2002). In the tomographic as well as the SFB analysis, the pseudo-C_ℓ estimator can be linked to the theoretical C_ℓ power spectrum using either a 2D or 3D mixing matrix:

$$< \widetilde{C}_\ell(k_{\ell n}, k_{\ell n'}) > = \sum_{\ell' n_1 n_2} M^{3D}_{\ell\ell' n n_1 n' n_2} C_{\ell'}(k_{\ell' n_1}, k_{\ell' n_2}) \,, \tag{10.28}$$

$$< \widetilde{C}^{(ij)}_\ell > = \sum_{\ell\ell'} M^{2D}_{\ell\ell'} C^{(ij)}_{\ell'} \,. \tag{10.29}$$

A derivation of the 2D mixing matrix can be found in Hivon et al. (2002), while the 3D matrix for the galaxy clustering SFB power spectrum is derived in Pratten and Munshi (2013).

It is important to point out that the effect of the mask can be taken into account in a similar way using pseudo-C_ℓs for the two methodologies explored here. Consequently, in a likelihood analyses using these expressions for the measured power spectra, the effect of the mask should be equivalent for the tomographic and SFB approaches. Therefore we only took partial sky coverage through the common f_{sky} scaling factor into account here for simplicity. This is standard practice for Fisher matrix analyses.

10.3 FORECASTING COSMOLOGICAL CONSTRAINTS

10.3.1 Fisher matrix forecasting

Expected cosmological constraints using the two different analysis techniques introduced in the previous section can be estimated with the Fisher matrix formalism (Tegmark et al., 1997). The Fisher information matrix provides a lower bound on the expected errors on cosmological parameters under the assumption that the likelihood can be approximated by a Gaussian at its peaks. It is formally defined as the expectation value of the second derivative of the logarithmic likelihood with respect to the parameters $\Theta_\alpha, \Theta_\beta$:

$$F_{\alpha\beta} = -\left\langle \frac{\partial^2 \ln L}{\partial\Theta_\alpha \partial\Theta_\beta} \right\rangle \,. \tag{10.30}$$

From this matrix, the marginal error on parameter Θ_α in particular can be extracted as $\sqrt{(F^{-1})_{\alpha\alpha}}$, and the error on Θ_α, all other parameters being fixed, is bounded by $(F_{\alpha\alpha})^{-1/2}$.

The Fisher matrix may be computed from the covariance matrix of the observable and its derivatives as

$$F_{\alpha\beta} = \frac{1}{2} Tr[C^{-1} C_{,\alpha} C^{-1} C_{,\beta}] \,. \tag{10.31}$$

10.3.1.1 *Implementing the tomographic Fisher matrix*

For the tomographic spectra $C_\ell^{(ij)}$, we computed the covariances between spectra under the Gaussian approximation following the approach of Hu and Jain (2004) and Joachimi and Bridle (2010). Denoting by $\Delta C_\ell^{(ij)}$ the difference between the ensemble average of the spectrum and its estimator, the tomographic power spectra covariance is defined as

$$\mathrm{Cov}_\ell^{(ijkl)} \equiv \left\langle \Delta C_\ell^{(ij)} \Delta C_\ell^{(kl)} \right\rangle \,, \tag{10.32}$$

$$= \frac{\delta_{\ell\ell'}}{f_{sky}(2\ell+1)} \left[\bar{C}_\ell^{(ik)} \bar{C}_\ell^{(jl)} + \bar{C}_\ell^{(il)} \bar{C}_\ell^{(jk)} \right] \,, \tag{10.33}$$

where f_{sky} accounts for partial coverage of the sky and $\bar{C}_\ell^{(ij)}$ is the tomographic power spectrum including shot noise defined in Equation (10.9). The expression of the tomographic Fisher matrix becomes

$$F_{\alpha\beta}^{tomo} = \sum_{(ij),(kl)} \sum_{\ell}^{\ell_{\max}(ijkl)} \frac{\partial C_\ell^{(ij)}}{\partial \Theta_\alpha} \mathrm{Cov}^{-1}{}_\ell^{(ijkl)} \frac{\partial C_\ell^{(kl)}}{\partial \Theta_\beta} \,, \tag{10.34}$$

where the sum over $(ij), (kl)$ indices loops over all $N_{zbins}(N_{zbins}+1)/2$ combinations of bins, and ℓ_{max} is a cut in multipole. The aim of this cut is to restrict the Fisher matrix to linear scales. Several strategies are possible to define $\ell_{\max}$; we describe the one adopted in this work in Section 10.3.1.4.

For the binning strategy, we chose to use equal galaxy density bins with no overlap. This choice led to bins with irregular widths, but constant shot noise.

10.3.1.2 *Implementing the SFB Fisher matrix*

The Fisher matrix for the 3D SFB spectra was computed using the non-diagonal covariance matrix obtained by discretising wavenumbers k under the boundary condition $n(r_{\max}) = 0$ as explained in Section 10.2.3. Details of computing the non-diagonal covariance matrix are given in Section D.2. In the absence of angular mask, Equation (10.17) shows that the SFB coefficients are uncorrelated between different angular multipoles ℓ. Therefore, the Fisher matrix for the SFB spectra takes the following form:

$$F_{\alpha\beta}^{SFB} = f_{sky} \sum_\ell \frac{(2\ell+1)}{2} \mathrm{Tr} \left[\widehat{C}_\ell^{-1} \frac{\partial \widehat{C}_\ell}{\partial \Theta_\alpha} \widehat{C}_\ell^{-1} \frac{\partial \widehat{C}_\ell}{\partial \Theta_\beta} \right] \,, \tag{10.35}$$

where the matrices $\widehat{C}_\ell$ are defined as

$$\widehat{C}_\ell = \begin{bmatrix} \overline{C}_\ell(0,0) & \overline{C}_\ell(0,1) & \dots & \overline{C}_\ell(0, n_{\max}^\ell) \\ \overline{C}_\ell(1,0) & \overline{C}_\ell(1,1) & \dots & \overline{C}_\ell(1, n_{\max}^\ell) \\ \vdots & \vdots & \ddots & \vdots \\ \overline{C}_\ell(n_{\max}^\ell, 0) & \overline{C}_\ell(n_{\max}^\ell, 1) & \dots & \overline{C}_\ell(n_{\max}^\ell, n_{\max}^\ell) \end{bmatrix} \,, \tag{10.36}$$

with $\overline{C}_\ell(n,p) = C_\ell(k_{\ell n}, k_{\ell p}) + N_\ell(k_{\ell n}, k_{\ell p})$. The size of each of this matrix $\widehat{C}_\ell$ is $n_{\max}^\ell \times n_{\max}^\ell$, where $n_{\max}^\ell$ defines the maximum wavenumber included in the Fisher analysis for each multipole ℓ. This allows us to restrict the analysis to linear scales. Again, different strategies can be adopted to define this cut in wavenumber; they are described in Section 10.3.1.4.

10.3.1.3 Fisher analysis baseline

To conduct this study, we adopted as a fiducial model a 'Vanilla' concordance flat cosmology with $h = 0.7$, $\Omega_b = 0.045$, $\Omega_m = 0.25$, $\Omega_\Lambda = 0.75$, $\Omega_b = 0.045$, $w_0 = -0.95$, $w_a = 0$, $n_s = 1$, $\tau = 0.09$, and $\sigma_8 = 0.8$. We adopted the standard parametrisation for the dark energy equation of state (Chevallier and Polarski, 2001),

$$w(a) = w_0 + w_a(1 - a) \, . \tag{10.37}$$

This fiducial cosmology was also used to compute the $\tilde{r} = r_{\rm fid}(z)$. In this model, we computed the linear matter power spectra, including baryonic oscillations, using the fitting formula of Eisenstein and Hu (1998).
We performed our Fisher analysis, under the constraint of a flat cosmology, on the following parameters $\Theta = (h, \Omega_m, w_0, w_a, \Omega_b, n_s, \sigma_8)$

For our baseline analysis we considered a spectroscopic survey with a very small redshift uncertainty $\sigma_z = 0.003(1 + z)$ and a Smail-type galaxy distribution $p(z)$ (Smail et al., 1994),

$$p(z) \propto z^2 e^{-\left(\frac{z}{0.708}\right)^{1.5}} \, , \tag{10.38}$$

which corresponds to a median redshift of $z_{\rm med} = 1$, and we used a mean number density of galaxies of $\bar{n} = 0.9$ gal. arcmin^{-2}. To account for partial coverage of the sky, we scaled the Fisher information by $f_{sky} = 0.3636$, which corresponds to a survey size of 15,000 square degrees. This setting was chosen to correspond to the specification of the Euclid spectroscopic survey (Laureijs et al., 2011).

Finally, we adopted a redshift dependent fiducial galaxy bias of the form

$$b(z, k) = \sqrt{1 + z} \, , \tag{10.39}$$

as in Rassat et al. (2008). In Section 10.3.1.5 we describe how we accounted for our lack of knowledge on the actual galaxy bias by parametrising this relation through nuisance parameters.

10.3.1.4 Restriction to linear scales

The constraints we aim to extract from a galaxy survey result from the information contained in the matter power spectrum. However, since the galaxies are only biased tracers of the actual underlying matter density, our knowledge of the matter power spectrum is limited by our understanding of the bias. This bias becomes more uncertain on small non-linear scales. Assuming an optimistic knowledge of the bias could result in overestimated or cosmologically biased constraints. Hence, following previous galaxy clustering studies (e.g., Rassat et al., 2007, 2008; Joachimi and Bridle, 2010), we completely discarded the mildly to non-linear scales and express our uncertainty of the bias on large scales by using nuisance parameters in the next section.

As the aim of this work is to compare the constraining power of two different approaches to galaxy clustering analysis, it is important to apply the exclusion of non-linear scales to the two methods in a coherent way to avoid biasing our results towards the method with the less conservative cut. Following the approach taken in Joachimi and Bridle (2010), which was based on results from Rassat et al. (2008), we aim to only retain linear scales through the following redshift-dependent cut in wavenumber $k_{\rm lin}^{\rm max}$:

$$k_{\rm lin}^{\rm max}(z) \approx \min[0.132z, \quad 0.25] \quad h{\rm Mpc}^{-1} \, . \tag{10.40}$$

This formula is a linear fit to the non-linearity scale in Fig. 2 of Rassat et al. (2008), which was computed as a function of redshift by selecting scales that fulfil $\sigma(R) < 0.20$ and $k_{\rm max} < 0.25 h{\rm Mpc}^{-1}$, where $\sigma(R)$ corresponds to the amplitude of fluctuations at $R\ h{\rm Mpc}^{-1}$. However, it provides a conservative cut at lower redshift (below z=0.5). Since the purpose of this work is to compare two methodologies given the same framework and set of assumptions, we used this model for the sake of simplicity. An accurate computation of the non-linear scale could be used just as well, but this is not expected to change the conclusions of the comparative analysis.

Because we computed the tomographic power spectra within the Limber approximation, we related wavenumbers k to angular modes ℓ through $k = \frac{\ell+1/2}{r}$. As a result, the non-linear scale cut translates into multipoles ℓ for redshift bin (i) as

$$\ell^{(i)}_{max} = k^{\rm max}_{\rm lin}(z^{(i)}_{\rm med}) r(z^{(i)}_{\rm min}) \ . \tag{10.41}$$

This cut allows us to reject all the multipoles for a given bin (i) that is affected by scales above $k^{\rm max}_{\rm lin}(z^{(i)}_{\rm med})$. When computing the correlation function between two different bins (i), (j) we applied the most conservative cut: $\ell^{(ij)}_{max} = \min(\ell^{(i)}_{max}, \ell^{(j)}_{max})$.

In the SFB framework, applying a corresponding wavenumber cut leads to an ℓ dependent maximum number of discrete wavenumbers $k_{\ell n}$, noted $n^{\ell}_{\rm max}$, which can be obtained as the solution of the equation

$$k_{\ell n^{\ell}_{\rm max}} r\left(\frac{0.132}{k_{\ell n^{\ell}_{\rm max}}}\right) = \ell \ , \tag{10.42}$$

under the constraint $k_{\ell n^{\ell}_{\rm max}} \leq 0.25 \quad h{\rm Mpc}^{-1}$. Both cuts are illustrated in Figure 10.1.

Thanks to this prescription, the same scales are excluded from the tomographic and SFB analysis. This point is the main difference between our work and the analysis performed in Nicola et al. (2014), where the exclusion of non-linear scales is not coherent between the two methodologies. In their work, the angular power spectra are truncated at $\ell_{max} = 50$ for all redshifts, whereas as shown in Figure 10.1, ℓ_{max} should be a function of the median redshifts of the tomographic bins to take into account the time evolution of the non-linear scale as well as the physical size of angular modes as a function of redshift. Similarly, in the SFB analysis performed in their work, a fixed cut at $k_{\rm max} = 0.20\ h{\rm Mpc}^{-1}$ was applied, which not only ignores the interplay between angular modes and z illustrated by Figure 10.1, but is also incoherent with the cut applied in the tomographic analysis.

10.3.1.5 Nuisance parameters

As mentioned in the previous section, restricting the study to linear scales avoids the high uncertainty on the bias that arises in the non-linear regime. Nevertheless, we also wish to express our uncertainty on the bias even on linear scales. Following the approach of Bridle and King (2007), Joachimi and Bridle (2010), and Kirk et al. (2012), we parametrised the bias in redshift and scale using a grid of nuisance parameters such that the galaxy bias becomes

$$b(k,z) = AQ(k,z) b_0(k,z) \ , \tag{10.43}$$

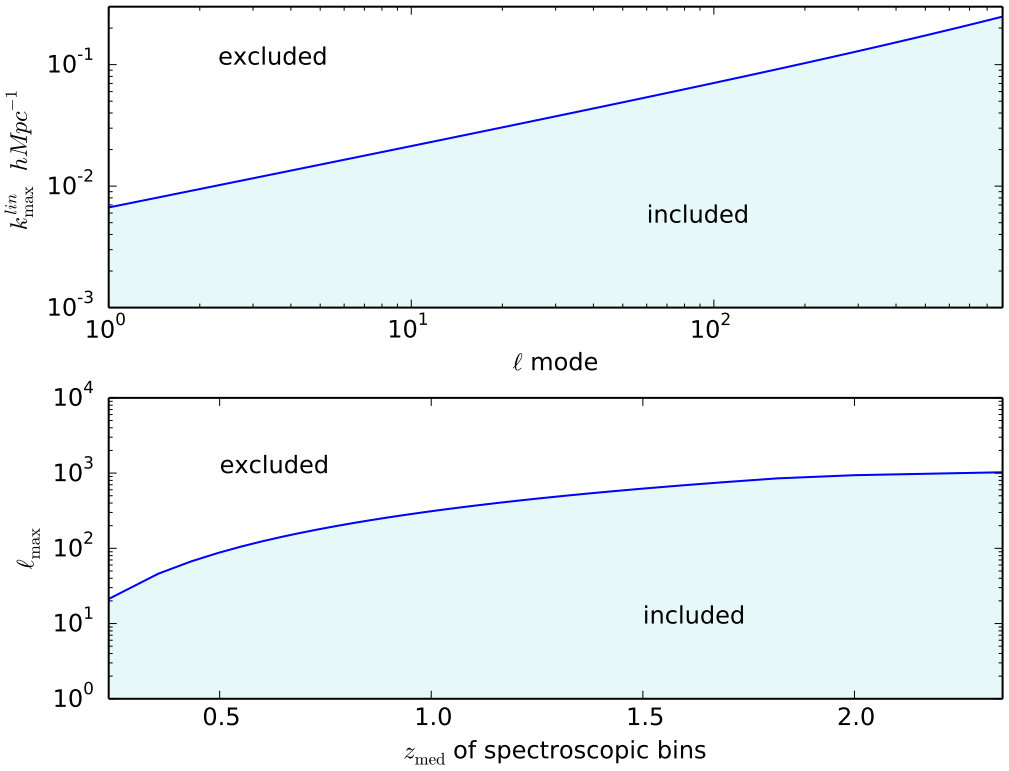

Figure 10.1: Top: Linear-scale limit in k for the SFB power spectra as a function of angular modes. Bottom: Linear-scale limit $\ell_{\max}$ for the tomographic power spectra $C_\ell^{(ij)}$ as a function of the lowest median redshift of the spectroscopic bins i and j. The regions above the lines are excluded from the Fisher analysis.

where b_0 is our fiducial bias relation introduced in Equation (10.39), A is an overall amplitude and $Q(k,z)$ encodes perturbations around the fiducial bias and is defined in terms of an $N_z \times N_k$ grid of parameters B_{ij}:

$$\begin{aligned}\ln Q(k,z) = K_i(k)Z_j(z)B_{ij} + [1-K_i(k)]Z_j B_{(i+1)j} \\ + K_i[1-Z_j(z)]B_{i(j+1)} + [1-K_i(k)][1-Z_j(z)]B_{(i+1)(j+1)}\,,\end{aligned} \tag{10.44}$$

for $k_i \le k \le k_{i+1}$ and $z_j < z \le z_{j+1}$, where the coefficients Z_j and K_i are expressed as

$$K_i(k) = \frac{\ln(k) - \ln(k_i)}{\ln(k_{i+1}) - \ln(k_i)}\,, \tag{10.45}$$

$$Z_j(k) = \frac{\ln(1+z) - \ln(1+z_j)}{\ln(1+z_{j+1}) - \ln(1+z_j)}\,. \tag{10.46}$$

The k_i and z_j fix the nodes of the grid and are spaced logarithmically in the intervals $k \in [10^{-4}, 1.0]$ and $z \in [0,5]$ such that $k_0 = k_{\min}$, $k_{N_k+1} = k_{\max}$ and $z_0 = z_{\min}$, $z_{N_z+1} = z_{\max}$.
The Fisher matrices are then obtained by marginalising over these $N_k \times N_z + 1$ nuisance parameters $\left(A, B_{00}, B_{01}, B_{10}, \ldots, B_{N_k N_z}\right)$.

10.3.2 Figures of merit

Throughout the rest of this work we compare the constraining power of the tomographic and SFB methods by evaluating their respective *Figure of Merit* (FoM). We consider two FoMs, first the total figure of merit $\mathrm{FoM}_{\mathrm{TOT}}$ defined according to Joachimi and Bridle (2010) as

$$\mathrm{FoM}_{\mathrm{TOT}} = \ln\left(\frac{1}{\det(F^{-1})}\right) , \tag{10.47}$$

and second, the dark energy figure of merit recommended by the report of the *Dark Energy Task Force* (DETF) (Albrecht et al., 2006),

$$\mathrm{FoM}_{\mathrm{DETF}} = \frac{1}{\sqrt{\det(F^{-1})_{w_0 w_a}}} . \tag{10.48}$$

The DETF FoM was designed to measure the strength of a given future survey or probe in constraining cosmological parameters related to the nature of dark energy, such that a large $\mathrm{FoM}_{\mathrm{DETF}}$ value meant a high constraining power on w_0 and w_a. The total FoM (i.e. $\mathrm{FoM}_{\mathrm{TOT}}$) was designed to encompass the strength of a future survey or probe in constraining several parameters across different sectors of cosmology, such as the nature of dark matter and dark energy and initial conditions. A high value of $\mathrm{FoM}_{\mathrm{TOT}}$ therefore means a good constraining power across all cosmological sectors. The parameter is taken as an ln value, since we consider this number for seven cosmological parameters, and the $\mathrm{FoM}_{\mathrm{TOT}}$ value would grow very quickly otherwise.

10.4 RESULTS: SFB VS. TOMOGRAPHIC ANALYSIS

10.4.1 Comparison of SFB and tomographic analysis in the absence of systematics

We now compare the relative constraining power of the tomographic and SFB analysis of galaxy clustering presented in Section 10.2 using the Fisher matrix formalism and the fiducial cosmology and survey baseline described in Section 10.3. We investigate first the impact of the number of redshift bins and whether the same constraints can be recovered from the two different analysis. Figure 10.2 shows the FoMs obtained using both methods as a function of number of tomographic spectroscopic bins when assuming perfect knowledge of the bias (in dark blue).

As expected, the two FoMs for the tomographic analysis increase with the number of redshift bins and eventually reach the performance of the SFB analysis for 30 redshift bins. Not only do the two methodologies yield equivalent figures of merit for this number of bins, but the 1σ contours for all cosmological parameters are extremely similar, both in size of the ellipse and for the direction of the degeneracies. Figure 10.3 shows the 1σ contours on all pairs of cosmological parameters considered for the two analysis techniques using 30 tomographic bins with and without nuisance parameters for the bias. For the fixed bias, the contours obtained by the tomographic analysis are plotted in red and are almost indistinguishable from the contours for the SFB analysis, which are depicted in orange.

We conclude that exactly the same information is extracted from the two methodologies for an appropriate number of redshift bins, 30 in our case.

This result disagrees with the conclusions of Nicola et al. (2014), who found that the SFB analysis is weaker than a tomographic analysis and not capable of extracting the same radial information.

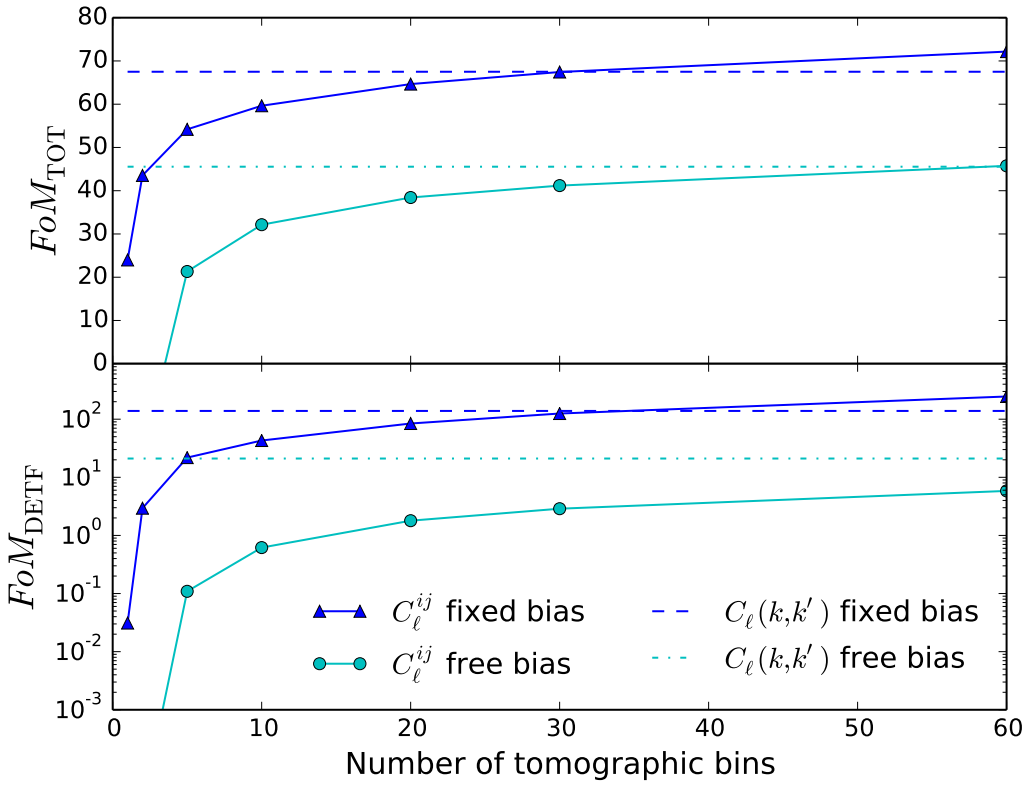

Figure 10.2: Comparison of the total $\mathrm{FoM}_{\mathrm{TOT}}$ (top) and dark energy $\mathrm{FoM}_{\mathrm{DETF}}$ (bottom) figures of merit for the 3D Fourier-Bessel analysis (horizontal dashed lines) vs. tomographic analysis (solid lines) as a function of number of redshift bins. The upper lines (dashed and triangle) result from assuming a fixed bias, the lower lines (dotted and circle) are obtained when assuming a grid of 5×5 nuisance parameters in scale and redshift described in Section 10.3.1.5.

The difference in these conclusions is probably related to the choice of non-linear prescription. In Nicola et al. (2014), the tomographic SHT analysis was limited to a fixed ℓ_{max} for all bins, while the SFB analysis was restricted to a fixed k_{max} for all multipoles. However, we show in Figure 10.1 that to apply equivalent cuts for the tomographic and SFB analysis, redshift-dependent $\ell_{\mathrm{max}}(z_{med})$ and ℓ-dependent $k_{\mathrm{max}}(\ell)$ cuts need to be used.

Additionally, Figure 10.2 shows that when the number of bins is increased, the tomographic analysis eventually surpasses the SFB analysis. This behaviour is expected, because when the width of the redshift bins reaches the non-linearity scale, the tomographic analysis probes more modes than a 3D analysis (Asorey et al., 2012; Di Dio et al., 2014). Indeed, only non-linear angular scales are excluded from the tomographic analysis, but for very thin redshift bins, small radial scales are being probed that are potentially beyond the non-linear cut-off. Asorey et al. (2012) found that a tomographic analysis with a bin width of $\Delta r \simeq 0.8 \frac{2\pi}{k_{\mathrm{max}}}$ was equivalent to a 3D power spectrum analysis including scales up to k_{max}. We found that the tomographic analysis recovers the information from the 3D analysis for about 30 redshift bins. If one expects the two methodologies to give similar results for $\Delta r \simeq \frac{2\pi}{k_{max}}$, then one would expect a larger number of tomographic bins to

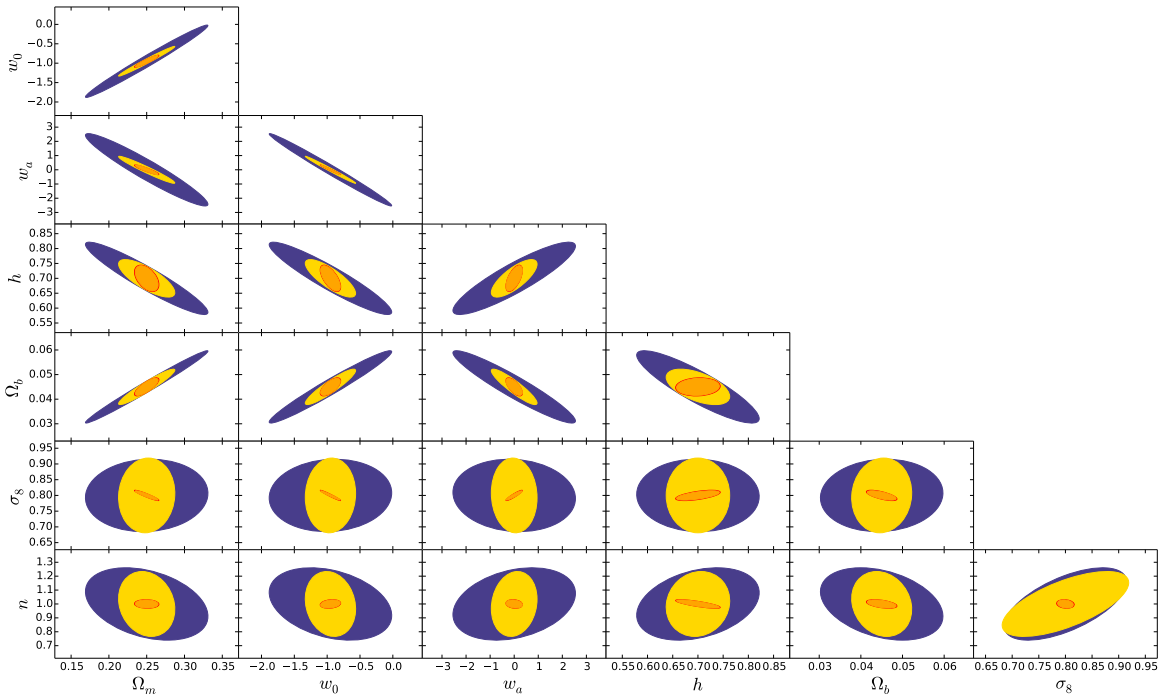

Figure 10.3: 1σ contours for all pairs of cosmological parameters for the SFB analysis and the tomographic analysis for 30 bins with and without nuisance parameters. Inner orange and red contours (almost indistinguishable) result from the SFB and tomographic analysis when assuming a fixed bias. Outer yellow and purple contours are obtained from the SFB and tomographic analysis when using a 5×5 nuisance parameter grid in scale and redshift for the bias.

be necessary. Here, our 30 bins correspond to a minimum bin width $\Delta r \simeq 0.55 \frac{2\pi}{k_{max}}$, which is not as close to the non-linearity scale as the results from Asorey et al. (2012), but remains of the same order of magnitude.

However, we stress that such a direct comparison is subject to several factors that complicate the interpretation. Firstly, the tomographic spectra are computed within the Limber approximation, which may not be accurate for a large number of thin bins. A recent study of the effect of the Limber approximation for a spectroscopic survey can be found in Eriksen and Gaztanaga (2015). Because we restricted our analysis to large linear scales, we limited the number of tomographic bins to 30 in the rest of the analysis, which corresponds to redshift widths between $\Delta_z = 0.1$ and $\Delta_z = 0.05$. In this case, according to Eriksen and Gaztanaga (2015), the error of approximation remains limited (below 15% for most bins). Therefore, we do not expect the full computation to significantly alter the results of the comparative study lead in this work. Nevertheless, this point should be kept in mind and deserves a thorough analysis, which we will include in future work. We also stress that although care has been taken to apply similar non-linear cuts, they are not strictly equivalent, and different strategies to restrict angular modes in the tomographic analysis would affect the results.

Therefore, we consider that for a fixed bias, both analysis methodologies recover the same information for 30 tomographic bins, which corresponds to a minimum bin width of the order of the non-linearity scale. We also acknowledge that the exact number of bins is likely to change for different binning strategy, computation techniques of angular power spectra, restrictions of

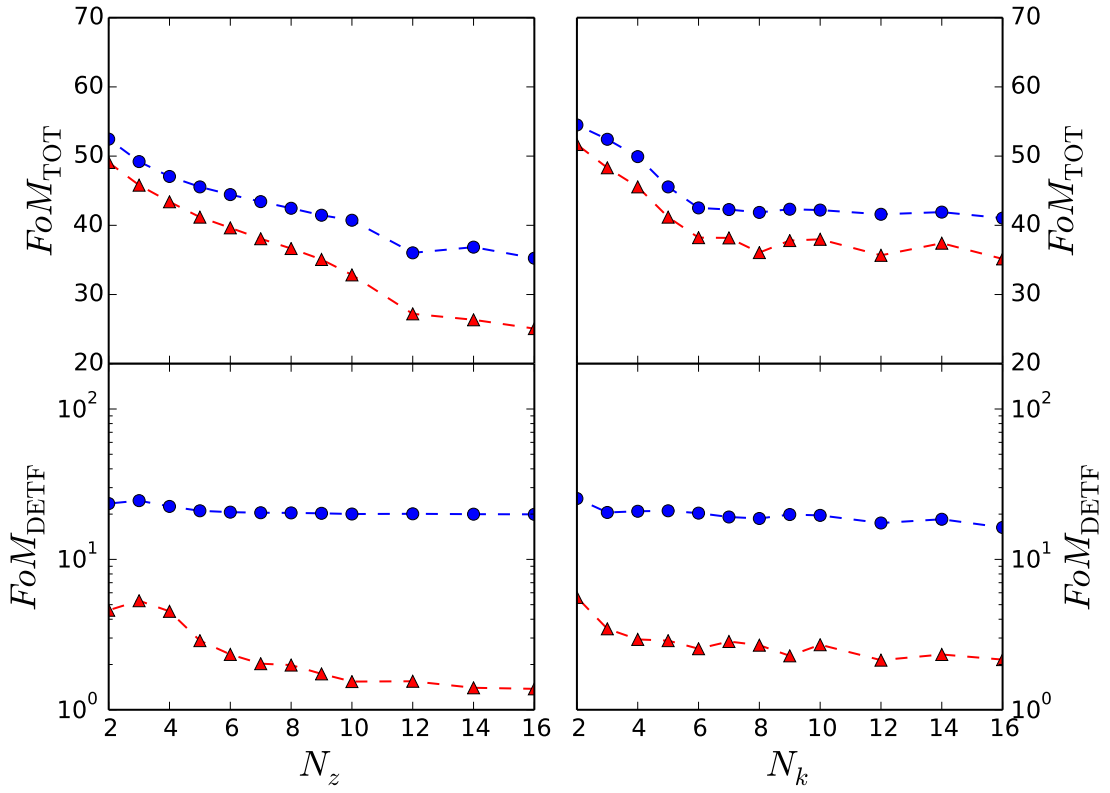

Figure 10.4: Total $\mathrm{FoM}_{\mathrm{TOT}}$ (top) and dark energy $\mathrm{FoM}_{\mathrm{DETF}}$ (bottom) figures of merit as a function of the number of nuisance parameters in redshift (left) and scale (right), for a tomographic with 30 bins (red triangle) and an SFB (blue dot) analysis. When varying the number of nuisance parameters in scale or redshift, the other number of parameters is kept fixed at 5.

non-linear scales and with the inclusion of additional effects such as redshift space distortions or relativistic effects. A thorough study of all these effects will be addressed in a future work.

10.4.2 Impact of systematics due to galaxy bias

After establishing that the same information can be recovered from both methodologies in the absence of systematics on the bias, we now investigate the impact of an unknown bias. As described in Section 10.3.1.5, we include in the analysis an uncertainty on the galaxy bias using a grid of nuisance parameters in scale and redshift. Figure 10.2 demonstrates how the FoMs for both analysis are degraded when using a free bias parametrised in scale and redshift by a 5×5 nuisance parameter grid (in cyan). Whereas the FoMs were equivalent with 30 tomographic bins in the fixed bias case, the tomographic analysis can no longer recover the same information as the SFB analysis in the free bias case, even with 60 redshift bins. The tomographic analysis is much more sensitive to systematics resulting from the unknown bias than the SFB analysis.

We investigated the effect of the number of nuisance parameters in scale and redshift $N_k \times N_z$ on the FoMs for the tomographic (red triangle) and SFB (blue dot) analysis in Figure 10.4. We

varied N_k and N_z independently while keeping the other parameter fixed to 5. When the number of nuisance parameters increases, the constraints from both analyses decreases, although the FoMs from the tomographic analysis degrade faster than for the SFB analysis.

Although the FoMs reach a plateau at about $N_z = 12$ and $N_k = 6$, these numbers would correspond to a very conservative model of the galaxy bias and therefore are probably unrealistic. Indeed, the evolution of galaxy bias should be smooth on large scales, which prompts us to limit the fiducial parameter grid used in this section to $N_z = 5$ and $N_k = 5$. Since the trends in FoMs do not change with the number of nuisance parameters, a more complex grid (increasing either N_z or N_k) would not change the conclusions on the relative strength of the two methodologies investigated here (the SFB FoM remains higher for any choice of nuisance parameters).

The effect of the free bias on the 1σ contours on cosmological parameters is shown in Figure 10.3, where the purple and yellow contours are computed from the 30-bin tomographic analysis and the SFB analysis. Interestingly, the constraints on σ_8 and n_s are affected in the same way for the two methodologies by the inclusion of nuisance parameters; the contours are almost equivalent for (n_s, σ_8) with or without nuisance parameters. In contrast, all other parameters are much more degraded by the including nuisance parameters in the case of the tomographic analysis compared to the SFB analysis. This is particularly true for the dark energy parameters w_0 and w_a.

These results agree with Asorey et al. (2012), who noted that the tomographic constraints degrade faster than a 3D power spectrum analysis when a single nuisance parameter on the amplitude of the bias was included. We find a similar behaviour with a more flexible parameterisation of the bias and for the 3D SFB analysis.

Furthermore, these results highlight the well-known sensitivity of galaxy clustering studies to the galaxy bias, which is one of its most important systematics. Although other approaches such as the measurement of the BAO scale are less sensitive to the galaxy bias, this results in the usual trade-off between systematics and statistical constraining power, so that BAO studies alone (i.e. using only BAO scale measurement) only provide conservative constraints without relying on external priors (Rassat et al., 2008).

10.4.3 Optimisation of a stage-IV survey

Since we have shown in Section 10.4.2 that the 3D SFB and tomographic methods depend differently on nuisance parameters, we are interested in investigating whether there are other differences in using one method or the other to plan for future wide-field surveys.

In this section we investigate the influence of the median redshift on the constraining power of a stage-IV spectroscopic survey using the two techniques. To perform this comparison, we used the same 5×5 nuisance parameter grid for the bias as in the previous section. We also adapted the number of tomographic bins to the median redshift of the survey to preserve the equivalence between tomographic and SFB constraints in the absence of systematics found in Section 10.4.1. The smallest radial scales probed by a tomographic analysis depend on the depth of the survey and on the number of bins. Therefore, to remain coherent for different median redshifts with the SFB analysis, the number of bins needs to be adjusted to the median redshift. We find that for a median redshift of $z_{med} \simeq 0.4$, the number of bins of the tomographic analysis should be $N = 26$ and for $z_{med} \simeq 1.7$ this number increases to $N = 42$. To illustrate this point, we plot in Figure 10.5 the FoMs as red triangles as a function of the median redshift using this adapted number of bins. The cyan line shows the evolution of the FoMs when keeping the number of bins fixed at $N = 30$. Since this

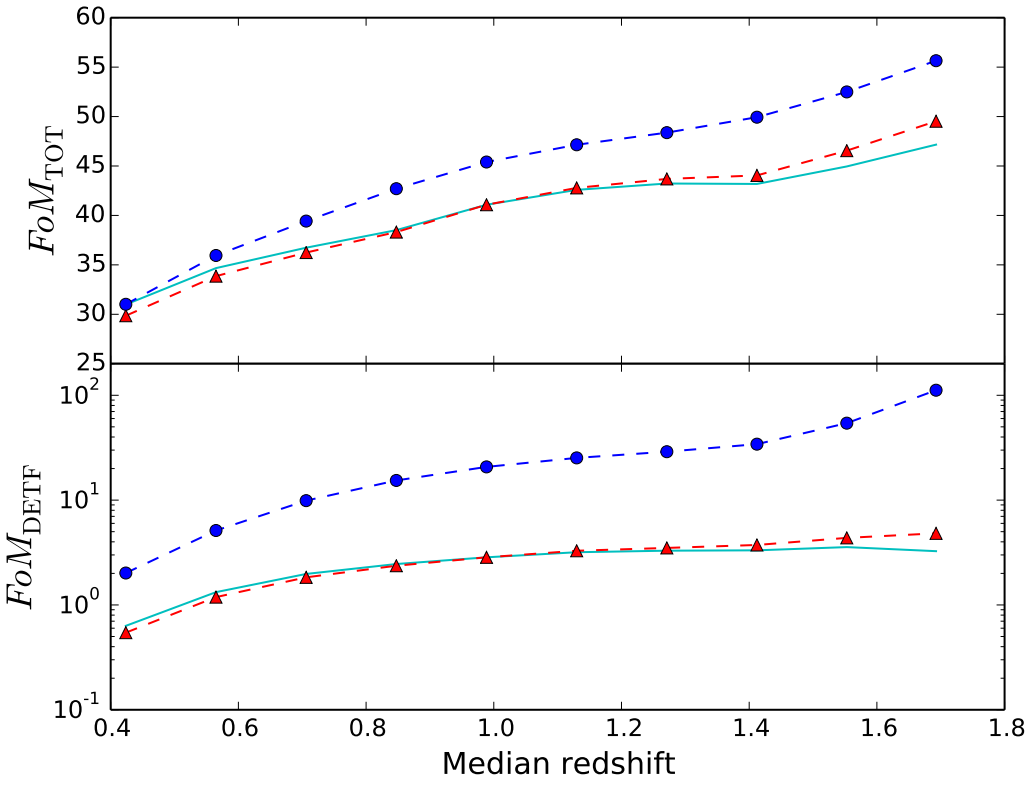

Figure 10.5: Total $\mathrm{FoM}_{\mathrm{TOT}}$ (top) and dark energy $\mathrm{FoM}_{\mathrm{DETF}}$ (bottom) figures of merit as a function of median redshift of a stage-IV spectroscopic survey using a tomographic analysis (red triangles) and an SFB analysis (blue dots). The cyan line shows the figures of merit for the tomographic analysis when the number of bins is not adapted to the depth of the survey and kept at 30 bins. For the red-dashed lines, the number of bins has been adapted to each median redshift. In all cases, a grid of 5×5 nuisance parameters in scale and redshift is used to parametrise the galaxy bias.

number of 30 tomographic bins was chosen in the previous section based on our fiducial survey with a median redshift of 1, we see that the red and cyan curves cross at $z_{med} = 1$. However, using 30 bins below $z_{med} = 1$ means probing smaller radial scales, which are beyond the scales probed by the SFB analysis, and this increases the FoMs. In contrast, above $z_{med} = 1$, this means using wider tomographic bins and thus probing larger radial scales, which lowers the FoMs compared to when the number of bins is adapted.

We also plot in Figure 10.5 the 3D SFB FoMs as a function of the median redshift of the survey, using blue circles. This curve should be compared to the red triangles showing the FoMs for the tomographic analysis where we have adapted the number of tomographic bins based on the median redshift, as described above. The two techniques exhibit a similar scaling with the median redshift of the survey, although the SFB constraints are consistently better than the tomographic constraints. Interestingly, for median redshifts above $z_{med} = 1.4$, the SFB dark energy FoM exhibits a better scaling than the tomographic one.

In conclusion, in the presence of galaxy bias systematics, any desired FoM level can be reached for shallower surveys if a 3D SFB analysis is performed. Furthermore, increasing the depth of the

survey is more profitable in terms of $\mathrm{FOM}_{\mathrm{DETF}}$ for the 3D SFB analysis because the tomographic method reaches a plateau somewhat after $z = 1$, whereas the 3D method continues to increase significantly up to $z = 1.8$, potentially pushing the optimisations towards higher median redshifts.

10.5 CONCLUSION

We have compared two different approaches to the three-dimensional analysis of galaxy clustering in the context of wide and deep future spectroscopic galaxy surveys. Based on the Fisher matrix analysis, we have compared the tomographic spherical harmonics and Spherical Fourier-Bessel methodologies in terms of figures of merit and cosmological parameter constraints.

We focused on the seven common parameters that are currently used in wide-field survey optimisation and planning: $\vec{\theta} = \{\Omega_m, h, w_0, w_a, \sigma_8, \Omega_b, n_s\}$, while putting forward a coherent and realistic approach regarding the exclusion of non-linear scales for both the 2D and 3D methods. In addition, we investigated for the first time how tomographic and 3D SFB methods are affected by nuisance parameters related to the galaxy bias, which we allowed to be both redshift- and scale-dependent.

In the absence of systematics, for an appropriate number of tomographic bins the two methodologies are equivalent and are able to recover the exact same constraints - both in value and in direction of degeneracy between different parameters. Increasing the number of redshift bins further leads to stronger constraints for the tomographic analysis, as seen by Asorey et al. (2012), Di Dio et al. (2014), and Nicola et al. (2014). Nevertheless, this effect could result from including radial scales in the tomographic analysis that are beyond the non-linear cut-off applied to the SFB analysis, and should be investigated further.

On the other hand, when we included unavoidable systematics due to the galaxy bias through a grid of nuisance parameters in scale and redshift, we found that the SFB analysis is more robust than the tomographic analysis, whose constraints suffer more from including nuisance parameters. As a result, we found that when we optimised the median redshift of a stage-IV type spectroscopic galaxy survey, a given level of accuracy can be achieved for shallower surveys if a 3D SFB analysis is performed. Moreover, the scaling of the dark energy figure of merit with median redshift is better for the 3D SFB analysis in the presence of systematics, which means that a given increase of the survey depth yields more information using an SFB analysis than a tomographic analysis.

Our results suggest that an SFB analysis is preferable to a tomographic analysis for realistic future spectroscopic wide-field surveys where the galaxy bias can be both redshift- and scale-dependent, and is unknown. These conclusions should be investigated in more detail, for example regarding the potential effect of the exact computation of angular power spectra, binning strategy, and including RSD.

In the spirit of reproducible research, the Python package **CosmicPy** developed to produce all the results presented in this work is freely available at

`http://cosmicpy.github.io`.

This package allows for simple and interactive computation of tomographic and 3D SFB power spectra as well as Fisher matrices while relying on a fast C++ implementation of Fourier-Bessel related computations.

11 CONCLUSION

The main goal of this thesis has been the development of new algorithms, based on recent advances in the field of sparse regularisation, to address several inverse problems in cosmology, with a particular focus on weak lensing mass-mapping.

The first part of this thesis introduced the concepts and tools of sparse regularisation and presents an application of these methods to the reconstruction of the power spectrum of primordial perturbations. The method I developed is non parametric and capable of reconstructing small deviations from a power law spectrum. I applied this new technique on WMAP-9 year and Planck data but do not find any significant deviations from the currently preferred near scale invariant power spectrum. These results are in accordance with the most recent Planck analysis.

The second part of this thesis focused on the particular problem of the reconstruction of the dark matter mass-map from weak lensing measurements. The partial and irregular sampling of the shear field make this inverse problem ill-posed, especially when reconstructing the dark matter density in three dimensions. I first developed a new 2D mass mapping technique, based on sparse regularisation, specifically designed for the high resolution density mapping of galaxy clusters from weak lensing. To complement the shear signal, which is noise dominated on small scales, this method can incorporate flexion information (higher order lensing deformations) to improve angular resolution and recover some of the substructures, as was demonstrated on a set of cluster simulations.

I then considered the more complex problem of the 3D reconstruction of the dark matter distribution. Compared to the 2D case, this problem involves an additional layer of complexity linked to the inversion of the radial lensing operator, which is particularly difficult. Previous attempts at 3D mass-mapping, based on linear methods, have had mitigated success and exhibited in particular extremely poor redshift resolution. For a long time, the limitations of linear methods have hindered potential applications of 3D mass-maps. To overcome these limitations, I developed a new non-linear reconstruction technique which dramatically improves the quality of 3D maps. As demonstrated on a large set of simulations, this new method is not only capable of accurately recovering the redshift of dark matter halos but can also be used to directly estimate their masses, which had never been considered with previous techniques.

In the final part of this thesis, I presented developments for the 3D analysis of cosmological surveys in spherical coordinates. I first constructed a new isotropic wavelet on the 3D ball based on the SFB expansion, ideally suited for the analysis of galaxy surveys. This wavelet boasts a number of very desirable properties, in particular translation invariance, isotropy and compact support in direct space with minimal oscillations. As wavelets are fundamental tools in data analysis but also building blocs for developing sparsity based methods, this new wavelets opens a wide range of applications, from the detection of voids to the estimation of the 3D dark matter distribution. Finally, I conducted a comparison of methods for the 3D analysis of spectroscopic galaxy surveys,

investigating the relative constraining power of a tomographic spherical harmonics analysis and a full 3D SFB analysis. If both methods can recover similar constraints, we find the SFB analysis to be more robust to unavoidable systematics related to the galaxy bias and allows a much more rigorous rejection of non-linear scales.

OUTLOOK

PRIMORDIAL POWER SPECTRUM RECONSTRUCTION The method presented in this thesis can be improved and extended in a number of ways. First and foremost, while we have only considered the reconstruction problem from CMB temperature anisotropies, additional data sets can be incorporated to improve the precision, sensitivity as well as the wavenumber range of the reconstruction. A first straightforward extension would be the inclusion of CMB polarisation data from Planck but LSS data, for instance from the future Euclid mission, can increase the range of scales included in the reconstruction. Another important point which we have not addressed in this work is the joint estimation of the primordial power spectrum with the other cosmological parameters, which could be implemented by integrating our reconstruction algorithm as part of an MCMC exploration, similar to the approach followed in Nicholson et al. (2010).

2D CLUSTER DENSITY MAPPING The 2D mapping technique developed in this thesis is one of the very first methods combining shear and flexion. With the recent improvement of flexion measurement methods, it will become possible to map in much more details a wider range of galaxy clusters, especially when constraints from strong lensing are poor. This is of particular interest for investigations into the nature of dark matter (Massey et al., 2015; Harvey et al., 2015). This method can be further developed in a number of ways, starting with the inclusion of additional constraints from strong lensing, when available, but also by using specifically designed dictionaries instead of the generic wavelets used in this work.

3D MASS MAPPING This thesis has demonstrated that the fundamental limitations which seemed to plague 3D weak lensing mass-mapping are not insurmountable and rather inherent to the linear methods used in previous reconstruction techniques. With the sparse recovery algorithm I developed in this thesis it becomes possible to detect, locate in redshift, and weight dark matter halos using a 3D reconstruction. This opens a range of very interesting new applications and in particular these maps could be used in cluster number counts analyses in the future Euclid mission, as a complement of optical cluster finders. Although already very effective, the algorithm developed in this thesis can be further improved by incorporating the various developments made for the 2D mapping algorithm. In particular, the tomographic binning of the input shear catalogue can be entirely avoided in order to fully take into account individual redshift estimates for each lensing source.

3D WAVELETS ON THE BALL Similarly to the numerous applications made possible by the development of wavelets on the sphere, wavelets on the ball, such as the one that I have developed in this thesis, have a wide range of applications for future cosmological surveys, which include not only galaxy surveys but also 21 cm radio surveys. The application of wavelets on the ball to the detection of cosmic voids was already suggested in Leistedt et al. (2013), a task for which

our isotropic, minimally oscillating wavelets would be extremely well adapted. Another possible application would be the estimation of the matter density field from galaxy surveys, combining the Multi-Scale Variance Stabilisation Transform for Poisson processes developed in Zhang et al. (2007) and sparse inpainting algorithms to accommodate missing data due to angular masks.

3D SFB ANALYSIS There is now a significant amount of literature on the analysis of galaxy clustering and weak lensing using a full SFB expansion. An interesting prospect is the joint analysis of both probes using this framework for the self-calibration of intrinsic alignment and galaxy bias, as a direct extension of the method proposed in Joachimi and Bridle (2010). As the results of the analysis performed in this thesis seem to suggest, an SFB expansion should outperform a standard tomographic analysis for this task.

Part IV

APPENDIX

A SNIA DETECTION IN THE SNLS PHOTOMETRIC ANALYSIS USING MORPHOLOGICAL COMPONENT ANALYSIS

The following article presents an application of sparse signal processing to the removal of artefacts corrupting the image stacks used for the detection of type Ia supernovae in the deferred photometric pipeline of the SuperNova Legacy Survey (SNLS). The main problematic in this application is to robustly identify supernovae candidates from so called subtracted images i.e. images from which a reference frame of the background sky has been subtracted. These subtracted images should only contain transient objects such as supernovae. However, in practice due to imperfect image subtractions, these subtracted contain a number of artefacts which lead to an overwhelming number of false positives in the detection pipeline. Although these false positives can be removed latter on in the analysis by considering the light curves, they considerably slow down the analysis of the data.

Our goal in this application is to take advantage of the morphological difference between artefacts and supernovae candidates in order to disentangle the signal of interest from subtraction artefacts. We apply a *Morphological Component Analysis* (MCA) decomposition algorithm which disentangles the different morphological components of an image based on their relative sparsity in different dictionaries. This algorithm allows us to filter out most of the artefacts present in the original data and therefore reduce the number of false positives.

An example of a subtracted image before and after MCA cleaning is shown on Figure A.1. The left panels are the stack images while the right panels are detection maps. As can be clearly seen in this example the large stellar residuals present in the original image are successfully removed by the separation algorithm.

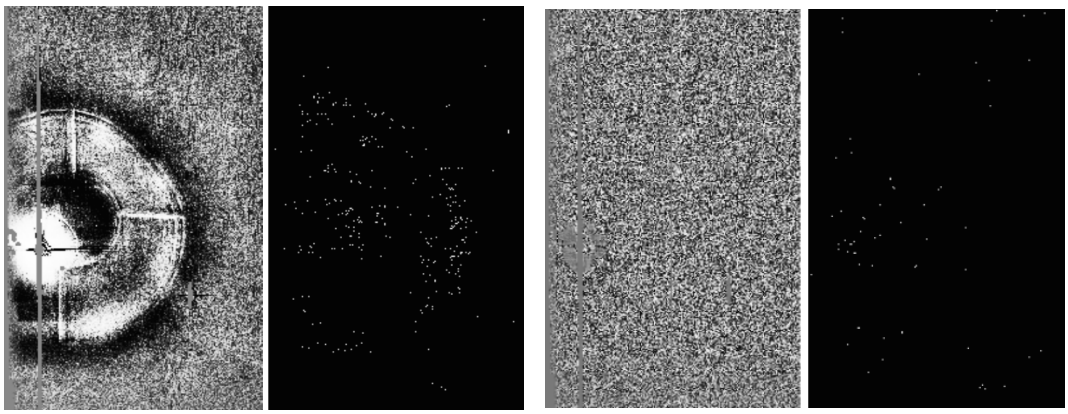

(a) Before MCA cleaning **(b)** After MCA cleaning

Figure A.1: Image stack for field D4 (left) and its corresponding detection map (right) before (a) and after (b) MCA cleaning. Before cleaning, detections clearly follow the shadow of the star while after cleaning most of these spurious detections are eliminated.

Prepared for submission to JCAP

SNIa detection in the SNLS photometric analysis using Morphological Component Analysis

A. Möller,[a,b,1] **V. Ruhlmann-Kleider,**[a] **F. Lanusse,**[c] **J. Neveu**[a] **N. Palanque-Delabrouille**[a] **and J.-L. Starck**[c]

[a]Irfu, SPP, CEA Saclay,
F-91191 Gif sur Yvette cedex, France.
[b]Université Paris Diderot - Paris 7,
75013 Paris, France.
[c]Laboratoire AIM, UMR CEA-CNRS-Paris 7, Irfu, SAp, CEA Saclay,
F-91191 Gif sur Yvette cedex, France.
E-mail: anais.moller@cea.fr, vanina.ruhlmann-kleider@cea.fr, francois.lanusse@cea.fr, jeremy.neveu@cea.fr, nathalie.palanque-delabrouille@cea.fr, jstarck@cea.fr

Abstract. Detection of supernovae (SNe) and, more generally, of transient events in large surveys can provide numerous false detections. In the case of a deferred processing of survey images, this implies reconstructing complete light curves for all detections, requiring sizable processing time and resources. Optimizing the detection of transient events is thus an important issue for both present and future surveys. We present here the optimization done in the SuperNova Legacy Survey (SNLS) for the 5-year data deferred photometric analysis. In this analysis, detections are derived from stacks of subtracted images with one stack per lunation. The 3-year analysis provided 300,000 detections dominated by signals of bright objects that were not perfectly subtracted. Allowing these artifacts to be detected leads not only to a waste of resources but also to possible signal coordinate contamination. We developed a subtracted image stack treatment to reduce the number of non SN-like events using morphological component analysis. This technique exploits the morphological diversity of objects to be detected to extract the signal of interest. At the level of our subtraction stacks, SN-like events are rather circular objects while most spurious detections exhibit different shapes. A two-step procedure was necessary to have a proper evaluation of the noise in the subtracted image stacks and thus a reliable signal extraction. We also set up a new detection strategy to obtain coordinates with good resolution for the extracted signal. SNIa Monte-Carlo (MC) generated images were used to study detection efficiency and coordinate resolution. When tested on SNLS 3-year data this procedure decreases the number of detections by a factor of two, while losing only 10% of SN-like events, almost all faint ones. MC results show that SNIa detection efficiency is equivalent to that of the original method for bright events, while the coordinate resolution is improved.

[1]Corresponding author.

Contents

1 Introduction

Surveys of distant type Ia supernovae (SNe Ia) revealed at the end of the twentieth century the acceleration of the expansion of the Universe [1, 2]. Since then, other surveys such as SNLS and SDSS-II [3] have been set in place to obtain measurements of SNe Ia with higher precision. The first step for detecting SNe events is to make a sample of transient events to be later classified. Detection using only photometry with difference images in one filter, where a reference image is subtracted, provides a good approach. However, difference images are filled with various artifacts from instrumental defects and incomplete subtraction of permanent objects. Disentangling real transient events from artifacts becomes an important requirement especially for photometric only pipelines such as the one developed in the deferred analysis of SNLS [4]. This is also of interest for future surveys which will process large amounts of data, such as LSST which expects to detect one million SNe per year [5].

SNLS is part of the Deep Synoptic Survey conducted on the Canada-France-Hawaii Telescope (CFHT). It was designed for detecting hundreds of SNe Ia in a redshift range between 0.2 and 1. Using the MegaCam imager [6], an array of 36 CCD with 340 million of pixels, four one square degree fields were targeted throughout 5 to 7 consecutive lunations per year using four different broadband filters g_M, r_M, i_M and z_M in the wavelength range from 400 to 1000 nm. Images were preprocessed at CFHT to perform flat-fielding and to remove defects. These pre-processed images were submitted to different pipelines. The real-time one provides detections of SNIa candidates and includes the result of spectroscopic follow-up for further classification and redshift determination [7]. This pipeline will not be addressed in this work.

The deferred photometric pipeline is independent of this real time analysis. Transient events are detected in one filter and multi-band light-curves are processed for all detections. Then, these light-curves are used to select SN-like events which are assigned host galaxy

photometric redshifts from an external catalogue. Light-curves and redshifts are then used to classify objects in SN Ia and core collapse types. The feasibility of detecting SNIa with this deferred analysis was proven for the 3-year SNLS data in [4]. In the era of large future surveys, spectroscopic resources will be limited for candidate follow-up and classification, which makes photometric pipelines interesting to study, e.g. [8].

The SNLS photometric pipeline is described in more detail in [4]. In the following we will summarize the main features of the detection step which are relevant to our study. Detection of transient events is done only in the i_M filter because distant SNe in SNLS have their maximum flux in this band. Reference images are constructed for each field from a set of best quality images which are coadded. Each image of the survey has the reference image subtracted. The subtraction is done using determination of the sky background and a convolution kernel which allows the subtraction to be adapted to different observing conditions. In order to increase the signal-to-noise ratio, subtracted images for each lunation are stacked. Lunation detection catalogues are constructed from these subtracted image stacks keeping only events which have a S/N ratio of 2.5 σ w.r.t. the sky background. The final detection catalogue is obtained by merging all lunation catalogues. All detection catalogues are constructed using SExtractor [9].

For the SNLS 3-year (SNLS3) analysis [4], the detection resulted in 300,000 transient candidates for which four-filter light curves were reconstructed. However, detections were dominated by spurious objects due to bad subtraction. Spurious detections came mostly from imperfectly subtracted objects such as bright stars, resampling defects and masks (see e.g. Figures 1, 2). Processing light curves of such a large number of detections knowing that 80% of those will be rejected by the early steps of the scientific analysis and do not contribute to science results [4] represents a waste of time. Therefore, in order to reduce the number of detections, it is necessary to disentangle true signal from artifacts in subtracted image stacks.

In this paper we present a new approach for improving transient event detection based on morphological component analysis [10] for difference image stacks in the SNLS deferred processing. Our goal is to obtain a reduction of the number of detections while limiting the loss of SNe Ia in the detection sample. We exploit the different morphologies of objects in the stacks to separate transient objects from artifacts. Other methods to achieve such a goal exist, such as the one recently introduced by [11] based on machine learning and principal component analysis using SDSS images. We also describe a new strategy for extracting signal coordinates from our cleaned image stacks with a good resolution.

The outline of the paper is as follows. Morphological component analysis is introduced in Section 2. The method proposed to clean subtracted image stacks in order to severe spurious detections is presented in Section 3. Section 4 describes the new detection strategy based on the cleaned stacks. Results on SNLS3 data and MC efficiency and coordinate resolution studies are presented in Section 5.

2 Morphological Component Analysis

The Morphological Component Analysis (MCA) framework assumes that an observed image can be described as the sum of several components, each exhibiting a distinct morphology. The aim of MCA is to leverage these characteristic morphologies to disentangle the different components of an image. More formally, given an image X, assumed to be the sum of K

morphologically different components x_k,

$$X = \sum_{k=1}^{K} x_k \ , \tag{2.1}$$

morphological component analysis can be used to recover each individual contribution x_k. In the context of SN detection in image stacks, this approach can be used to disentangle transient objects from artifacts as they exhibit different morphologies.

To actually perform this separation, MCA relies on the theory of sparse representation of signals. Any signal can be represented in a number of different domains (e.g. time domain, Fourier domain, wavelet domain) but the coefficients of this signal will exhibit different properties depending on the domain. One property of particular interest is the so called sparsity of the coefficients, i.e. the property that only a small number of coefficients are non zero. As a general rule, the coefficients of the signal will be sparse when the basis functions of the domain (so called *atoms*) are very similar to the signal itself.

In fact, this sparsity property is extremely desirable as it can be used as a very powerful prior in the regularization of a wide range of inverse problems. Some applications to astronomy and astrophysics include denoising [12], deconvolution [13], blind source separation for CMB analysis [14], weak gravitational lensing [15].

More formally, let us denote α the coefficients of a signal x in a *dictionary* $\mathbf{\Phi}$ (a dictionary being the set of atoms ϕ_i of a given domain):

$$x = \mathbf{\Phi}\boldsymbol{\alpha} = \sum_i \phi_i \alpha_i \ , \tag{2.2}$$

If x is sparse in dictionary $\mathbf{\Phi}$ then only a small number of coefficients in $\boldsymbol{\alpha}$ are non zero. Given image X defined in (2.1) as the sum of K different morphological components, let us introduce K different dictionaries $\mathbf{\Phi}_k$, each adapted only to the particular morphology of component x_k i.e. such that the $\boldsymbol{\alpha}_k$ coefficients, $\{\alpha_{ki}\}$, of x_k in $\mathbf{\Phi}_k$ are sparse but not the coefficients of x_l for $l \neq k$. Then performing the separation of the different morphological components can be achieved by finding an optimal set of coefficients $\boldsymbol{\alpha}_k$ maximizing the sparsity of the decomposition of each component in the corresponding dictionary.

The Morphological Component Analysis (MCA) algorithm has been proposed by [16] as a practical way to perform this decomposition as the solution of an ℓ_1 minimisation problem, where ℓ_1 denotes the ℓ_1-norm. Observed images, Y, are assumed to be a combination of signals, X, plus some noise, N. The decomposition algorithm solves iteratively the following optimization problem:

$$\min_{\boldsymbol{x}_1,\ldots,\boldsymbol{x}_K} \sum_{k=1}^{K} \parallel \mathbf{\Phi}^* \boldsymbol{x}_k \parallel_1 \quad \textit{such that} \quad \parallel Y - \sum_{k=1}^{K} \boldsymbol{x}_k \parallel_2 \leq \sigma \ , \tag{2.3}$$

where $\mathbf{\Phi}^* \boldsymbol{x}_k = \boldsymbol{\alpha_k}$ and σ is the standard deviation of the noise contaminating the data, assumed to be stationary and Gaussian distributed. Note that dictionaries and scales are chosen prior to minimization. The ℓ_1-norm promotes the sparsity of the decomposition of each component [17]. At convergence, each morphological component is obtained as $x_k = \mathbf{\Phi}_k \boldsymbol{\alpha}_k$. This component reconstruction can be restricted to a sub-sample of $\{\alpha_{ki}\}$, for example to some size scales in a given dictionary.

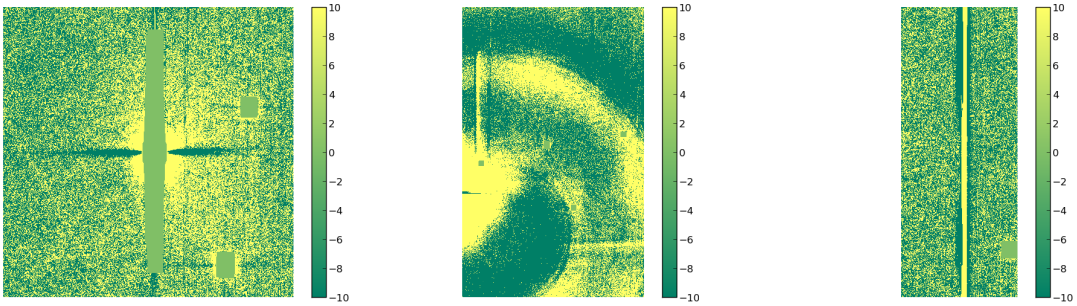

(a) (b) (c)

Figure 1: Different defects on the subtracted image stacks that yield spurious detections on large scale: (a) shows a saturated star with some areas masked by subtraction, (b) a saturated star plus the camera mounting shadow and (c) defects from sampling and dead pixel lines.

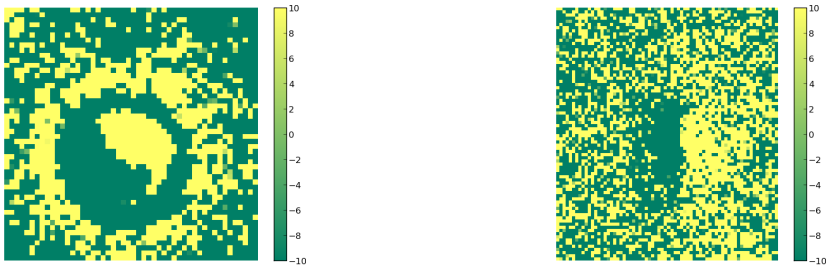

(a) (b)

Figure 2: Defects on the subtracted image stacks that yield spurious detections on small scale: (a) and (b) dipoles from imperfect galaxy subtraction. These are adjacent positive and negative areas on the stacks.

3 Reducing artifacts in SNLS subtraction stacks

Morphological component analysis allows to disentangle artifacts from other signals and can be adapted to treat subtracted image stacks in SNLS. First, we choose dictionaries which characterize signal and artifacts distinctively at different size scales. Then, we present a two-step treatment designed to extract interesting SN-like signals and rejects spurious detections.

A sub-sample of SNLS3 data was used to characterize artifacts and to assign the algorithm parameter values. Field D4 was chosen since it is a summer field with very good observing conditions and a large number of both detections and events classified as SN-like candidates as in [4].

3.1 Choice of dictionaries

The aim of the filtering approach presented in this section is to separate the signal of interest (SN-like events) from a complex background. The latter is constituted by noise, defects that cannot be subtracted (e.g. Figure 1) and features from imperfect subtractions (e.g. Figure 2).

Because these artifacts are structured, a naive strategy based on detection through a simple threshold on signal-over-noise ratio yields a large number of spurious detections. Our aim is to leverage additional morphological information to separate the signal of interest from artifacts and noise, by exploiting their stark contrast in both shape and scale.

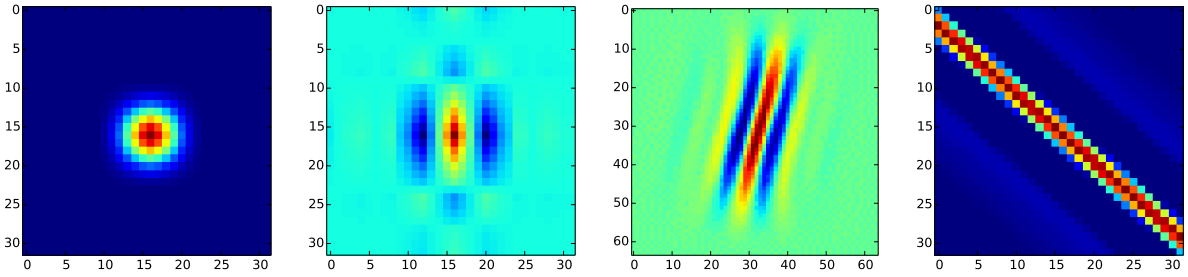

(a) Starlet (b) Bi-orthogonal (c) Curvelet (d) Ridgelet

Figure 3: Typical atoms from the dictionaries used in the MCA algorithm. (a) starlet atom representing circular-like signals, (b) bi-orthogonal wavelets for dipole features, (c) curvelets for elliptical signals and (d) ridgelets representing line features.

As explained in the previous section, the MCA algorithm separates images into a number of morphological components, using the sparsity level of each component in appropriate dictionaries as a discriminant. Therefore, in the case of the SNLS data, it is important to select, on one hand, a dictionary adapted to the morphology of the signal of interest and, on the other hand, additional dictionaries adapted to the artifacts we want to reject. More information on available dictionaries can be found in [18].

SN-like signals are small scale circular type shaped objects. A wavelet based dictionary is suited to this kind of morphology. We choose in particular the starlet dictionary since it is composed of isotropic atoms, especially efficient for representing positive structures such as our SN candidates. An example of a starlet atom is presented on Figure 3a.

For the small scale artifacts presented in Figure 2 we adopt a bi-orthogonal wavelet dictionary (Figure 3b). These artifacts result from improper subtraction of galaxies which lead to characteristic dipole features. The bi-orthogonal dictionary has the advantage of representing such features more efficiently than the starlet, enabling us to discriminate these artifacts from the signal.

For large scale curved or line artifacts such as the one in Figure 1, we adopt curvelet and ridgelet dictionaries. The curvelet dictionary is composed of localized, elongated atoms, at different scales, which are known to provide a sparse representation for curved features, see Figure 3c. The ridgelet atoms are line of different widths and orientations (see Figure 3d) which are perfect to represent the second type of artifacts.

An important aspect of all the dictionaries introduced here is that they are based on multiresolution transforms, meaning that they can be used to probe features at different discrete scales $j \in [\![0, N]\!]$. Typically, atoms of these transforms at scale j have a characteristic size of 2^j, starting with the finest resolution with details at the pixel scale for $j = 0$. Note that this j scale index is embedded together with pixel coordinates in the i index in equation 2.2. The advantage of choosing different scales for each dictionary is that we are able to separate small scale signals from large scale defects. We make use of this scale information within the MCA algorithm as explained in the next section.

3.2 First treatment: removal of main artifacts

The MCA algorithm in [16] was adapted to treat our subtracted image stacks. To disentangle signal artifacts we chose the previously described dictionaries: starlets, bi-orthogonal wavelets, curvelets and ridgelets. SNLS3 D4 test sample was used to decompose known artifacts and

select the scales in each dictionary that allowed the best characterization. The best choices were 5 scales in dictionaries representing mostly artifacts (curvelets, ridgelets, bi-orthogonal), and 3 scales in the starlet dictionary which efficiently decompose circular-like signals as can be seen in Figure 4. Both positive and negative signals were decomposed since some SNe may have part of their flux included in the references, in which case subtraction yields a negative signal.

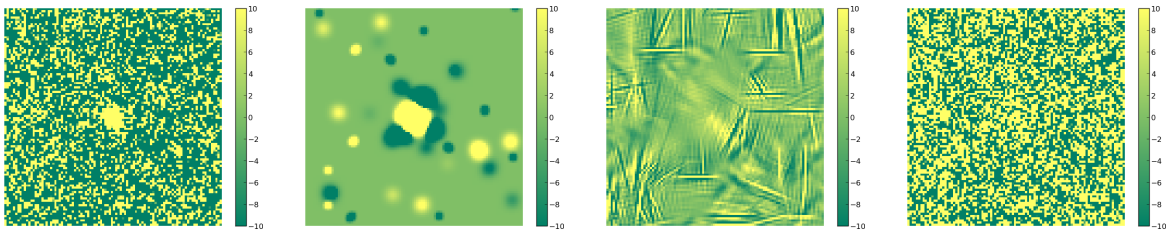

(a) (b) (c) (d)

Figure 4: MCA Decomposition of a SNIa event. (a) shows the original subtracted image stack centered on the SN event (yellow spot), (b) the starlet component, where the SN signal (yellow spot) is surrounded by remaining galaxy residuals (green spots), (c) the curvelet component and (d) the residuals left after the decomposition.

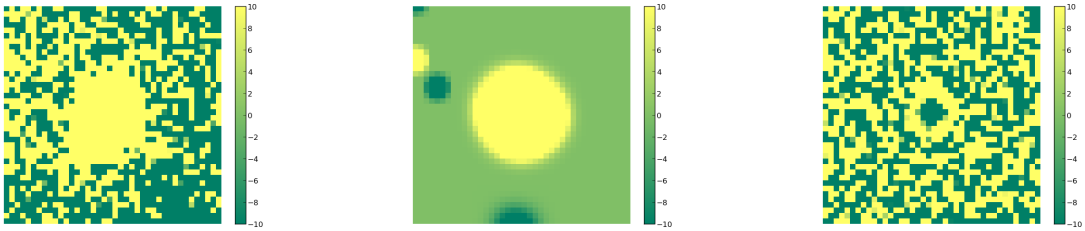

(a) (b) (c)

Figure 5: MCA Decomposition of a SNIa event where some part of the signal leaks into residuals. (a) shows the original subtracted image stack, (b) the starlet component and (c) the residuals after decomposition containing part of the signal.

Once dictionaries and scales are set, other algorithm parameters must be chosen such as the number of iterations in the optimization. The choice of parameters resulted from a trade-off between reducing the total number of detections and keeping most of the SN-like objects in the D4 test sample. A compromise between number of iterations and computation time was achieved with 30 iterations for the decomposition.

The transforms used in the algorithm, especially that of the curvelet dictionary, do not scale well with the image size and too much CPU time and memory would be required for a SNLS survey image of 2176 by 4912 pixels. Therefore, tiling of the images was done both to reduce time and memory resources and to allow parallel processing. For reference, one SNLS image divided in 8 tiles requires on average 6 days of HS06 CPU time and 500 Mb of virtual memory to be treated.

The algorithm assumes a stationary and Gaussian noise in the input images which is not the case for our subtracted image stacks, that are built from the coaddition of subtracted images spanning several weeks of observations. Thus, the signal we aimed at recovering was

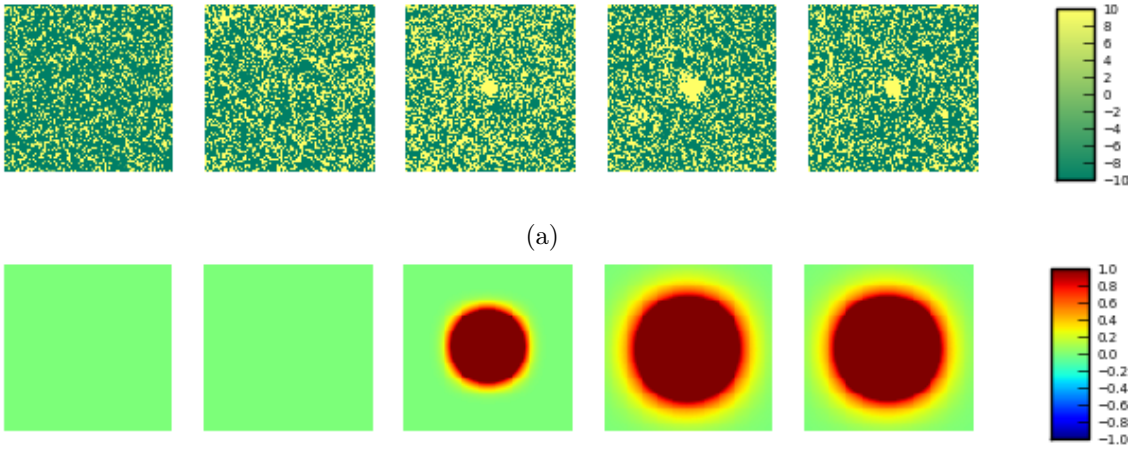

(a)

(b)

Figure 6: A SNIa event shown in different lunations around maximum light in the original subtracted image stack (a) and after both treatments of the cleaning procedure (b).

not properly decomposed and was partially in the residuals, e.g. Figure 5. To tackle this, a second treatment was developed which uses as input the starlet component and the residuals of our first treatment decomposition (e.g. components (b) and (d) of Figure 4).

3.3 Second treatment: signal extraction with varying noise

A utility based on the algorithm in [18] was developed. It handles non-stationary noise and exploits further morphological decomposition. Non-stationary noise requires varying the threshold in the decomposition depending on the position of the analyzed pixel. Such a feature can be easily implemented in the wavelet dictionary since it can handle actual noise maps. The latter were computed from the first treatment output images using a median absolute deviation estimator. This computation used a sliding window with a fixed size larger than what is expected for a SN-like signal. Since some of the SNe may have part of their flux included in the references, both positive and negative signals were treated.

We used again the starlet dictionary but this time handling varying noise, in order to select significant coefficients. All signals present in the output can be considered as morphologically compatible with circular-like objects. An example can be seen in Figure 6.

This utility does not require tiling images since only the wavelet dictionary is used. One SNLS image takes on average 3 hours of HS06 CPU time and 100 Mb of virtual memory to be treated.

4 New Detection strategy

A detection strategy includes both extracting events from an image and reconstructing their coordinates. Event extraction depends on the image and its characteristics, e.g. its local noise information. The TERAPIX tool SExtractor [9] was used for the whole detection strategy both in the original procedure and the new procedure adjusting its parameters accordingly.

In [4] the detection strategy consisted in constructing lunation catalogues with SExtractor with deblending, requiring for each detection at least 4 pixels with signal of more than

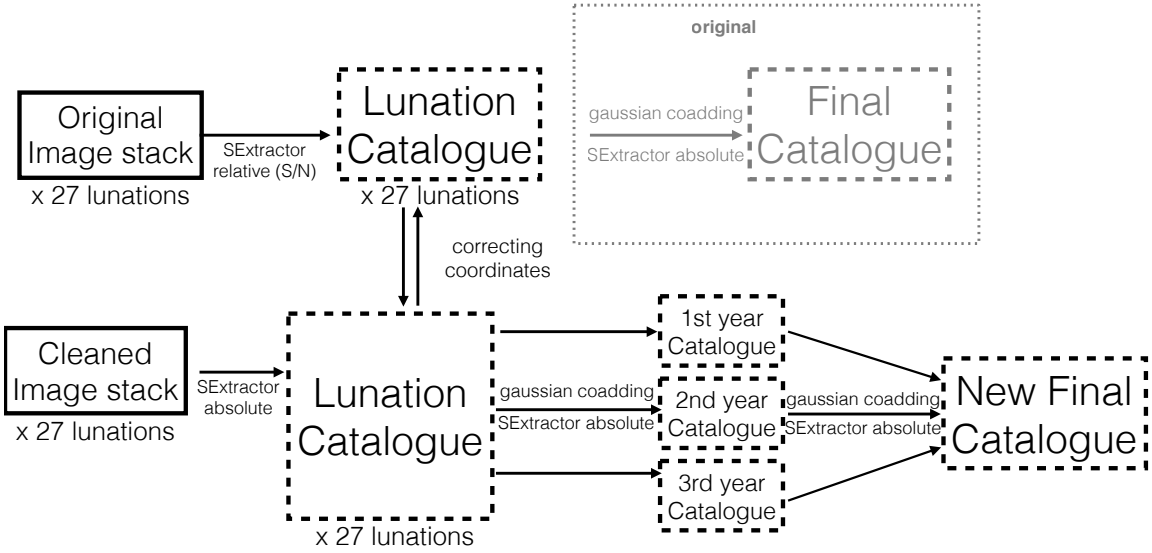

Figure 7: New detection strategy schema. Doted lines represent catalogue ASCII files while continuous lines images.

2.5σ w.r.t. sky background. A final detection catalogue was obtained by merging all lunation catalogues obtained in three years and converting the result into an image where each detection was replaced by a Gaussian of height and width of 1. This image was processed with SExtractor selecting only pixels with a content above a value of 0.01 and deblending objects. In this way, any object detected on several lunations at the same position (within a pixel) gave only one detection, with a position averaged over all lunation stacks. This is described on the top part of Figure 7. Even though the lunation stacking reduces the number of detections (typically by a factor 3), when adding many years it can degrade signal coordinate resolution due to close-by spurious detections as can be seen in Figure 8a.

Our two-step treatment outputs image stacks which do not have the same properties as the original subtracted image stacks. The noise has been removed and sources are reconstructed using inverse transformations. As transformed images they have less objects but coordinates cannot be extracted accurately from them. We thus propose a new detection strategy (see Figure 7 bottom) which also addresses the degradation of coordinate resolution when using several years of data. This strategy was set up on SNLS3 D4 data.

Lunation catalogs are constructed from our cleaned subtraction stacks using SExtractor, requiring at least 200 pixels with a signal value above one to confirm an object. Deblending is imposed in order to separate adjacent objects. The values of the SExtractor parameters were tuned using the SNLS3 test sample. They resulted from a trade-off between the reduction of the total number of detections and the number of SN-like objects detected on the cleaned image stacks. To each object detected in a lunation we assign the coordinates of the closest detection in the same lunation catalogue of the original procedure. In this way we maintain the reduced number of candidates while having precise coordinates.

In the original procedure coordinates were averaged over all lunations which degrades signal coordinate resolution due to close-by spurious detections. The latter are not always

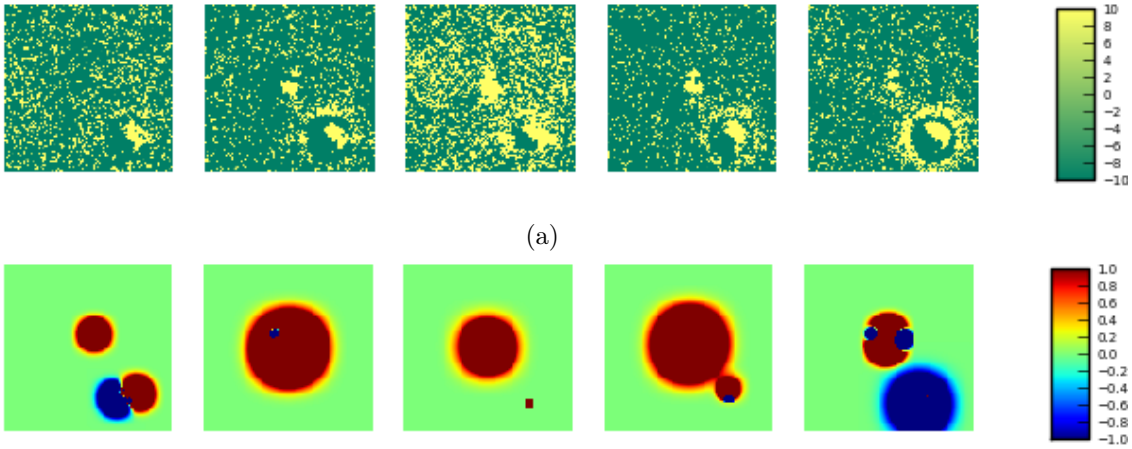

(a)

(b)

Figure 8: A SNIa event (center of the image) with galaxy residuals shown in different lunations around maximum light in the original subtracted image stacks (a) and after cleaning (b). In the original stacks, galaxy residuals are present in all lunations. The cleaning removes them in some cases.

completely removed by cleaning as can be seen in Figure 8b. When adding data from other seasons, the coordinate resolution degradation becomes even more important. Real SN-like events can be present in at most three adjacent lunation catalogs but not over several seasons. Hence, to address this in the new procedure we first build a catalogue for each season as we did for the final catalogue in the original procedure. Then, we build the new final catalogue from the season catalogues in the same way.

In this way, coordinate averaging is done first for a season were a transient object can be present and then detections are added from other years. It is equivalent to assigning a weight for a given detection taking into account that a SN will be detected only during one season.

5 Results

5.1 SNLS data

The SNLS3 D4 test sample contained 90,971 detections from which 362 events were extracted as SN-like objects as described in [4]. After our processing, the number of detections is reduced to 40,575. This represents more than a factor 2 reduction on the number of candidates to be further processed. Loss of SN-like candidates is less than 5% and all lost events are faint (observed magnitude at peak in $i_M > 24.2$) and so not suitable for further cosmological analysis. The complete procedure with parameters as determined from D4 data was then applied to the three other fields. Results are summarized in Table 1. The reduction of the number of detections is similar in all fields. The loss of SN-like events is less than 5% in D3 and 15% in D1 and D2. It must be noted that D1 and D2 are the fall and winter fields which have less suitable weather conditions than D3 and D4. All lost events are faint with the exception of one medium brightness event in D1 which is lost during our new detection procedure. This

	Old procedure		New procedure	
Field	# detected	# SN-like	# detected	# SN-like
D1	76,806	444	34,314	382
D2	64,763	300	28,627	258
D3	70,447	377	29,292	359
D4	90,971	362	40,575	346
All	302,987	1,483	127,808	1,345

Table 1: Number of detected and SN-like events for the original and new proceedures applied on SNLS3 data.

	Old procedure		**New procedure**			
			No season stacks		With season stacks	
Stack	coordinate resolution ±0.002	magnitude bias ±0.0001	coordinate resolution ±0.002	magnitude bias ±0.0001	coordinate resolution ±0.002	magnitude bias ±0.0001
1-year	0.374	0.0093	0.368	0.0090	0.368	0.0090
3-year	0.392	0.0104	0.388	0.0100	0.381	0.0096
5-year	0.408	0.0111	0.406	0.0109	0.398	0.0105

Table 2: Coordinate resolutions (pixels) and corresponding magnitude bias of SNIa detection original, new procedure with no season stacks or complete new procedure with season stacks: for year 3 MC data (1-year stack), adding two additional years of data (3-year stack) and adding 4 additional years of data (5-year stack). Uncertainties shown here are from the statistics of generated SNIa.

event is found on the output images of the two-step treatment but the number of pixels above threshold is smaller than our criteria to validate a detection.

5.2 Monte-Carlo efficiency and coordinate resolution

The performance of our treatment was studied using Monte Carlo (MC) artificial images in the i_M filter generated in the D1 field [19]. The MC images were constructed by adding simulated supernovae to real raw images, on host galaxies identified from deep image stacks of the CFHT-LS Deep Fields [20]. Using a two-dimensional gaussian for modeling the galaxies, SN positions within their hosts were randomly generated up to a distance of 2σ from the host galaxy centers. This method provided compatible simulated and observed SN-host galaxy angular distance distributions. The redshift assigned to each SN was that of its host galaxy and was restricted to the range between 0.2 and 1.2. For each SN the i_M light curve was simulated according to the SN properties (redshift, color, stretch) and the generated SN flux as deduced from the light curve at each observation date was added to the corresponding raw image. The MC images were then processed by the deferred pipeline as real images. Subtracted MC image stacks were then processed by our optimized pipeline. The corresponding results were compared to the ones of the original procedure in [4].

Detection efficiency was defined as the fraction of recovered simulated supernovae at the end of each processing. For both the original and the new procedure we computed the detection efficiency for one year of simulated SNe Ia as a function of the generated SN peak magnitude as can be seen in Figure 9. The efficiency is nearly magnitude independent up

to $m_{0i} = 23.5$ and then steeply declines at faint magnitudes. When compared to the old procedure, the new procedure corresponds to a loss of 0.5% in the plateau efficiency and a 0.2 downward shift of the magnitude corresponding to 0.50 efficiency. Note that after the new procedure the efficiency behavior as a function of magnitude is close to the original one and is consistent with that expected from a magnitude limited survey. This MC result will allow us to correct the Malmquist bias of the photometric sample to be derived from the new procedure in order to perform a cosmological analysis.

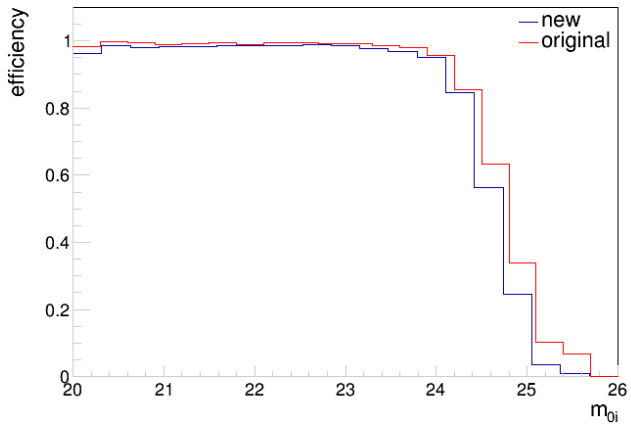

Figure 9: Efficiency of detection as a function of the generated peak magnitude in i_M. The new procedure (blue line) is compared to the original one (red line).

The SNIa coordinate resolution was studied for new and old procedures as can be seen in Table 2. The resolution was given by the RMS of the distance between the coordinate at generation and at detection. Coordinate resolution is improved first by reducing spurious detection with our cleaning procedure (Table 2, column 2). Further gain is obtained by the modified detection strategy (Table 2, column 3).

The coordinate resolution for one year simulated SNe Ia in our complete optimized pipeline was found to be 0.368 ± 0.002 pixels to be compared to 0.374 ± 0.002 pixels in the original pipeline. We also studied the effect of adding other years of survey data (without simulated SN signal), constructing catalogues with two or 4 additional years. A degradation of coordinate resolution is seen but the new procedure handles better many years of data than the original one.

Position measurement inaccuracy leads to underestimated fluxes. Using appendix B of [21] we computed an indicative magnitude bias corresponding to our coordinate resolution. Thanks to the improved coordinate resolution, the new procedure applied on 5-year stacks has similar performance to the old procedure applied on 3-year stacks.The latter was found accurate enough for photometric typing as shown in [4]. The application of the new procedure to the whole set of SNLS data will be the subject of a future work.

6 Discussion

Morphological component analysis has proven to be a useful approach for cleaning subtracted image stacks such as the ones in the SNLS deferred processing. Our experience shows that

the precise nature of the input images was a key point when choosing and adapting this type of algorithms. The choice of algorithm was based on the availability of a robust tool that could decompose our subtracted image stacks efficiently and within our CPU and time resources. For adapting the algorithms to the defects present in our input images we had to use a two-step procedure. In the first step we needed several dictionaries and scales to eliminate the various artifacts. Note that many of these defects came from the fact that we used subtracted images that usually have many residuals. For the second algorithm the goal was to handle non-stationary noise (typical from stacks) in addition to SN-like signals which provided a natural choice of the starlet dictionary for the decomposition. Finally, the choice of algorithm parameters (e.g. number of iterations) was heavily dependent on the efficiency and purity we wanted to achieve and on computing resources.

We note that besides improving the subtraction algorithm itself, eliminating artifacts at the level of subtracted images instead of stacks can provide a higher reduction of the number of detections. However, this should be applied at the beginning of the survey. For implementing such methods, a thorough analysis must be done of the trade-off between gain on signal extraction and removal of artifacts with respect to the high computational and time costs of processing using dictionary decomposition.

Future surveys like LSST may detect around ten thousands SNe Ia a year [22], which is two orders of magnitude higher than in SNLS. Extrapolating what we experienced in the deferred processing of SNLS, the anticipated number of detections in LSST may be as high as 10^7 per year which is too large to process. To reduce the number of candidate transient events to process further, cleaning images with a fast multi-resolution method can be of interest. But due to the huge number of detections, additional multi-band and temporal information will be necessary. The above arguments are valid for both real-time and deferred processings, which will both face too large numbers of detections to process. Differences between the two approaches would affect the choice of cleaning algorithms and selection criteria based on multi-band and temporal information.

7 Conclusions

In this paper we presented a new procedure for detecting supernovae in the SNLS photometric analysis. We developed a two-step procedure for cleaning subtracted image stacks, reducing artifacts and extracting SN-like signals using morphological component analysis. A new detection strategy, adapted to the cleaned image stacks was also presented.

The performance of the new procedure was evaluated using MC artificial images. Detection efficiency of SNeIa in the old and the new procedure is almost unchanged for bright events. However, there is a small reduction for events at higher magnitudes, which is expected since signal separation is not perfect and some SN-like signal may not be properly transformed. When applied to real SNLS3 data, 10% percent of SN-like events were lost while the number of detections was reduced by more than a factor two. Almost all lost events were faint with the exception of one medium brightness event which was lost in the detection step. This result agrees with MC findings.

Coordinate resolution of SNIa events was equivalent for one year of MC for both procedures. Furthermore, since SNLS is a five-year survey, coordinate resolution was also studied adding other years of data. The new procedure yields slightly better SNIa coordinate resolution with respect to the original procedure when adding 4 additional years of data, simulating a five-year stack. Therefore, for a five-year photometric analysis this new procedure yields

a slightly smaller magnitude bias for SNe Ia when compared to the original procedure. The new procedure presented in this work will be applied to the final SNLS 5-year photometric analysis which will be the subject of a forthcoming paper.

This work is a first step on morphological component analysis applied for SN-like signal detection and may be used as a starting point for future surveys. For those surveys that will detect a large number of events such as the LSST a fast multi resolution algorithm can be of interest, provided additional information (e.g. other filters, partial light curves) is also used.

Acknowledgments

This work was done based on observations obtained with MegaPrime/MegaCam, a joint project of CFHT and CEA/IRFU, at the Canada-France-Hawaii Telescope (CFHT) which is operated by the National Research Council (NRC) of Canada, the Institut National des Science de l'Univers of the Centre National de la Recherche Scientifique (CNRS) of France, and the University of Hawaii. This work is based in part on data products produced at Terapix available at the Canadian Astronomy Data Centre as part of the Canada-France-Hawaii Telescope Legacy Survey, a collaborative project of NRC and CNRS.

References

[1] **Supernova Search Team** Collaboration, A. G. Riess et al., *Observational evidence from supernovae for an accelerating universe and a cosmological constant*, *Astron.J.* **116** (1998) 1009–1038, [astro-ph/9805201].

[2] **Supernova Cosmology Project** Collaboration, S. Perlmutter et al., *Measurements of Omega and Lambda from 42 high redshift supernovae*, *Astrophys.J.* **517** (1999) 565–586, [astro-ph/9812133].

[3] **SDSS** Collaboration, M. Betoule et al., *Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples*, *Astron.Astrophys.* **568** (2014) A22, [arXiv:1401.4064].

[4] G. Bazin, V. Ruhlmann-Kleider, N. Palanque-Delabrouille, J. Rich, E. Aubourg, et al., *Photometric selection of Type Ia supernovae in the Supernova Legacy Survey*, *Astron.Astrophys.* **534** (2011) A43, [arXiv:1109.0948].

[5] **LSST Science Collaborations, LSST Project** Collaboration, P. A. Abell et al., *LSST Science Book, Version 2.0*, arXiv:0912.0201.

[6] O. Boulade, X. Charlot, P. Abbon, S. Aune, P. Borgeaud, P.-H. Carton, M. Carty, J. Da Costa, H. Deschamps, D. Desforge, D. Eppelle, P. Gallais, L. Gosset, R. Granelli, M. Gros, J. de Kat, D. Loiseau, J. . Ritou, J. Y. Rousse, P. Starzynski, N. Vignal, and L. G. Vigroux, *Megacam: the new canada-france-hawaii telescope wide-field imaging camera*, vol. 4841, pp. 72–81, 2003.

[7] **SNLS** Collaboration, P. Astier et al., *The Supernova legacy survey: Measurement of omega(m), omega(lambda) and W from the first year data set*, *Astron.Astrophys.* **447** (2006) 31–48, [astro-ph/0510447].

[8] H. Campbell, C. B. D'Andrea, R. C. Nichol, M. Sako, M. Smith, et al., *Cosmology with Photometrically-Classified Type Ia Supernovae from the SDSS-II Supernova Survey*, *Astrophys.J.* **763** (2013) 88, [arXiv:1211.4480].

[9] E. Bertin and S. Arnouts, *SExtractor: Software for source extraction*, *Astron.Astrophys.Suppl.Ser.* **117** (1996) 393–404.

[10] J.-L. Starck, M. Elad, and D. Donoho, *Redundant multiscale transforms and their application for morphological component separation*, *Adv.Imaging Electron Phys.* (2004).

[11] L. d. Buisson, N. Sivanandam, B. Bassett, and M. Smith, *Machine Learning Classification of SDSS Transient Survey Images*, arXiv:1407.4118.

[12] S. Beckouche, J.-L. Starck, and J. Fadili, *Astronomical Image Denoising Using Dictionary Learning*, *Astron.Astrophys.* **556** (2013) A132, [arXiv:1304.3573].

[13] J. Schmitt, J.-L. Starck, J.-M. Casandjian, J. Fadili, and I. Grenier, *Multichannel Poisson denoising and deconvolution on the sphere : Application to the Fermi Gamma Ray Space Telescope*, *Astron.Astrophys.* **546** (2012) A114, [arXiv:1206.2787].

[14] J. Bobin, J.-L. Starck, F. Sureau, and S. Basak, *Sparse component separation for accurate CMB map estimation*, *Astron.Astrophys.* **550** (2013) 73, [arXiv:1206.1773].

[15] A. Leonard, F. Lanusse, and J.-L. Starck, *GLIMPSE: Accurate 3D weak lensing reconstructions using sparsity*, *Mon.Not.Roy.Astron.Soc.* **440** (2014) 1281, [arXiv:1308.1353].

[16] J.-L. Starck, M. Elad, and D. Donoho, *Image decomposition via the combination of sparse representations and a variational approach*, *IEEE Trans.Image Process.* **14** (2005) 1570–1582.

[17] J.-L. Starck, F. Murtagh, and J. Fadili, *Sparse Image and Signal Processing: Wavelets, Curvelets, Morphological Diversity*. Cambridge University Press, New York, NY, USA, 2010.

[18] J. Starck, F. Murtagh, and A. Bijaoui, *Image Processing and Data Analysis: The Multiscale Approach*. Cambridge Univ. Press, 1998.

[19] P. Ripoche, *PhD thesis: Mesure de l'évolution du taux d'explosion des supernovae de type Ia en fonction du décalage spectral dans le SuperNova Legacy Survey*. Université de Marseille, 2007.

[20] O. Ilbert, S. Arnouts, H. McCracken, M. Bolzonella, E. Bertin, et al., *Accurate photometric redshifts for the cfht legacy survey calibrated using the vimos vlt deep survey*, *Astron.Astrophys.* **457** (2006) 841–856, [astro-ph/0603217].

[21] **SNLS** Collaboration, J. Guy et al., *The Supernova Legacy Survey 3-year sample: Type Ia Supernovae photometric distances and cosmological constraints*, *Astron.Astrophys.* **523** (2010) A7, [arXiv:1010.4743].

[22] C. M. Carroll, E. Gawiser, P. L. Kurczynski, R. A. Bailey, R. Biswas, et al., *Improving the LSST dithering pattern and cadence for dark energy studies*, *Proc.SPIE Int.Soc.Opt.Eng.* **9149** (2014) 91490C, [arXiv:1501.0473].

B MINIMUM VARIANCE FILTER FOR SHEAR AND FLEXION

The convergence can be estimated from flexion and shear by combining the following operators:

$$\tilde{\kappa} = \frac{i}{k_1}\tilde{\mathcal{F}_1} \quad ; \quad \tilde{\kappa} = \frac{i}{k_2}\tilde{\mathcal{F}_2} \quad ; \quad \tilde{\kappa} = \frac{k^2}{k_1^2 - k_2^2}\tilde{\gamma}_1 \quad ; \quad \tilde{\kappa} = \frac{k^2}{2k_1k_2}\tilde{\gamma}_2 \tag{B.1}$$

where we note $k^2 = k_1^2 + k_2^2$. We are looking for a minimum variance filter of the form:

$$\hat{\tilde{\kappa}} = \frac{ia}{k_1}\tilde{\mathcal{F}_1} + \frac{ib}{k_2}\tilde{\mathcal{F}_2} + \frac{ck^2}{k_1^2 - k_2^2}\tilde{\gamma}_1 + \frac{(1-a-b-c)k^2}{2k_1k_2}\tilde{\gamma}_2 \tag{B.2}$$

The optimum values for a, b and c should minimize the variance of the estimator, which is:

$$< \hat{\tilde{\kappa}}^2 > = \frac{a^2}{k_1^2} < \tilde{\mathcal{F}_1}^2 > + \frac{b^2}{k_2^2} < \tilde{\mathcal{F}_2}^2 > + \frac{c^2k^4}{(k_1^2 - k_2^2)^2} < \tilde{\gamma}_1^2 > + \frac{(1-a-b-c)^2k^4}{4k_1^2k_2^2} < \tilde{\gamma}_2^2 > \tag{B.3}$$

$$= \frac{a^2}{k_1^2}\sigma_F^2 + \frac{b^2}{k_2^2}\sigma_F^2 + \frac{c^2k^4}{(k_1^2 - k_2^2)^2}\sigma_\gamma^2 + \frac{(1-a-b-c)^2k^4}{4k_1^2k_2^2}\sigma_\gamma^2 \tag{B.4}$$

where we assume no intrinsic correlation between shear and flexion components. We now find the minimum variance by differentiating each variable:

$$\partial_a < \hat{\tilde{\kappa}}^2 >= 0 \Leftrightarrow \frac{2a}{k_1^2}\sigma_F^2 - 2\frac{(1-a-b-c)k^4}{4k_1^2k_2^2}\sigma_\gamma^2 = 0 \tag{B.5}$$

$$\Leftrightarrow a\left(\frac{\sigma_F^2}{k_1^2} + \frac{k^4\sigma_\gamma^2}{4k_1^2k_2^2}\right) = \frac{(1-b-c)k^4\sigma_\gamma^2}{4k_1^2k_2^2} \tag{B.6}$$

$$\Leftrightarrow a(4k_2^2\sigma_F^2 + k^4\sigma_\gamma^2) = (1-b-c)k^4\sigma_\gamma^2 \tag{B.7}$$

$$\Leftrightarrow \boxed{a = \frac{1-b-c}{1 + k_2^2\frac{4\sigma_F^2}{k^4\sigma_\gamma^2}}} \tag{B.8}$$

$$\partial_b < \hat{\tilde{\kappa}}^2 >= 0 \Leftrightarrow \frac{2b}{k_2^2}\sigma_F^2 - 2\frac{(1-a-b-c)k^4}{4k_1^2k_2^2}\sigma_\gamma^2 = 0 \tag{B.9}$$

$$\Leftrightarrow b\left(\frac{\sigma_F^2}{k_2^2} + \frac{k^4\sigma_\gamma^2}{4k_1^2k_2^2}\right) = \frac{(1-a-c)k^4\sigma_\gamma^2}{4k_1^2k_2^2} \tag{B.10}$$

$$\Leftrightarrow b(4k_1^2\sigma_F^2 + k^4\sigma_\gamma^2) = (1-a-c)k^4\sigma_\gamma^2 \tag{B.11}$$

$$\Leftrightarrow \boxed{b = \frac{1-a-c}{1 + k_1^2\frac{4\sigma_F^2}{k^4\sigma_\gamma^2}}} \tag{B.12}$$

$$\partial_c < \hat{\tilde{\kappa}}^2 >= 0 \Leftrightarrow \frac{2ck^4}{(k_1^2-k_2^2)^2}\sigma_\gamma^2 - 2\frac{(1-a-b-c)k^4}{4k_1^2k_2^2}\sigma_\gamma^2 \tag{B.13}$$

$$\Leftrightarrow c\left(\frac{k^4\sigma_\gamma^2}{(k_1^2-k_2^2)^2} + \frac{k^4\sigma_\gamma^2}{4k_1^2k_2^2}\right) = \frac{(1-a-b)k^4\sigma_\gamma^2}{4k_1^2k_2^2} \tag{B.14}$$

$$\Leftrightarrow \boxed{c = \frac{1-a-b}{1+\frac{4k_1^2k_2^2}{(k_1^2-k_2^2)^2}}} \tag{B.15}$$

Let us note $\lambda = \frac{4\sigma_F^2}{k^4\sigma_\gamma^2}$ and $\mu = \frac{4k_1^2k_2^2}{(k_1^2-k_2^2)^2}$, we now have to solve the following system of equations:

$$\begin{cases} (1+k_2^2\lambda) & a & = & 1-b-c & (1) \\ (1+k_1^2\lambda) & b & = & 1-a-c & (2) \\ (1+\mu) & c & = & 1-a-b & (3) \end{cases} \tag{B.16}$$

From this system, we can first obtain the following:

$$(1)-(2) \Rightarrow b = \frac{k_2^2}{k_1^2}a \qquad ; \qquad (3)-(2) \Rightarrow c = \frac{k_2^2\lambda}{\mu}a \tag{B.17}$$

which results in the following expression for a:

$$(1+k_2^2\lambda)a = 1-b-c \tag{B.18}$$

$$(1+k_2^2\lambda)a = 1-\frac{k_2^2}{k_1^2}a - \frac{k_2^2\lambda}{\mu}a \tag{B.19}$$

$$a = \frac{k_1^2}{k^2 + k_1^2k_2^2\lambda(1+\mu^{-1})} \tag{B.20}$$

From which we directly derive the expressions for all variables a, b and c:

$$a = \frac{1}{k^2+\sigma_F^2/\sigma_\gamma^2}k_1^2 \tag{B.21}$$

$$b = \frac{1}{k^2+\sigma_F^2/\sigma_\gamma^2}k_2^2 \tag{B.22}$$

$$c = \frac{\sigma_F^2/\sigma_\gamma^2}{k^2+\sigma_F^2/\sigma_\gamma^2}\frac{(k_1^2-k_2^2)^2}{k^4} \tag{B.23}$$

$$d = 1-a-b-c = \frac{\sigma_F^2/\sigma_\gamma^2}{k^2+\sigma_F^2/\sigma_\gamma^2}\frac{4k_1^2k_2^2}{k^4} \tag{B.24}$$

Which yields the following expression for the minimum variance filter:

$$\hat{\tilde{\kappa}} = \frac{1}{k^2+\sigma_F^2/\sigma_\gamma^2}\left(ik_1\tilde{\mathcal{F}}_1 + ik_2\tilde{\mathcal{F}}_2 + \frac{\sigma_F^2}{\sigma_\gamma^2}\left(\frac{k_1^2-k_2^2}{k^2}\tilde{\gamma}_1 + \frac{2k_1k_2}{k^2}\tilde{\gamma}_2\right)\right) \tag{B.25}$$

C SPHERICAL FOURIER-BESSEL TRANSFORM AND 3D CONVOLUTION

C.1 RELATION TO THE 3D FOURIER TRANSFORM

In order to have a better understanding of the SFB coefficients and of how to use them to perform filtering, the SFB transform can be related to the 3D Fourier transform. We follow a similar definition as the one presented in Baddour (2010), but using our conventions for the different transforms.The following convention will be used for the Fourier transform:

$$F(\vec{k}) = \frac{1}{\sqrt{(2\pi)^3}} \int f(\vec{r}) e^{-i\vec{k}.\vec{r}} d\vec{r}, \tag{C.1}$$

where F denotes the Fourier transform of f. This formulation does not assume any coordinate system. However, to relate this transform to the SFB transform, it is possible to express this equation in spherical coordinates using the following expansion for the Fourier kernel:

$$e^{-i\vec{k}.\vec{r}} = 4\pi \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} (-i)^{\ell} j_{\ell}(kr) \overline{Y_{\ell m}(\theta_r, \varphi_r)} Y_{\ell m}(\theta_k, \varphi_k), \tag{C.2}$$

where (k, θ_k, φ_k) and (r, θ_r, φ_r) are respectively the spherical coordinates of vectors $\vec{k}$ and $\vec{r}$.

Substituting this expression for the kernel in the definition of the 3D Fourier transform yields:

$$\begin{aligned} F(k, \theta_k, \varphi_k) &= \frac{4\pi}{\sqrt{(2\pi)^3}} \int_0^{\infty} \int_{\Omega} \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} (-i)^{\ell} f(r, \theta_r, \varphi_r) j_{\ell}(kr) \overline{Y_{\ell m}(\theta_r, \varphi_r)} Y_{\ell m}(\theta_k, \varphi_k) d\Omega r^2 dr, \\ &= \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} (-i)^{\ell} \left[\sqrt{\frac{2}{\pi}} \int_0^{\infty} \int_{\Omega} f(r, \theta_r, \varphi_r) j_{\ell}(kr) \overline{Y_{\ell m}(\theta_r, \varphi_r)} d\Omega r^2 dr \right] Y_{\ell m}(\theta_k, \varphi_k), \\ &= \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \left[(-i)^l \hat{f}_{\ell m}(k) \right] Y_{\ell m}(\theta_k, \varphi_k). \end{aligned} \tag{C.3}$$

In the last equation, the expression of the Spherical Harmonics Expansion of $F(k, \theta_k, \varphi_k)$ for a given value of k can be recognised. In the Fourier space, the $(-i)^{\ell} \hat{f}_{\ell m}(k)$ are the Spherical Harmonics coefficients of F on a sphere of given radius k. In other words, the Spherical Harmonics coefficients $F_{\ell m}(k)$ of the 3D Fourier transform $F(k, \theta_k, \varphi_k)$ on a sphere of given radius k in Fourier space are the SFB coefficients $\hat{f}_{\ell m}(k)$ for the same value k but multiplied by factor $(-i)^{\ell}$.

The relationship between 3D Fourier transform and SFB transform is therefore very simple. The SFB transform can be thought of as a mere Fourier transform in spherical coordinates. In the next sections, this relationship will be used to derive convolution and filtering relations for the SFB transform using the well known relations verified by the Fourier transform.

C.2 3D CONVOLUTION PRODUCT USING THE SFB TRANSFORM

A prerequisite to the establishment of filtering relations is the expression of a 3D convolution product in terms of SFB coefficients. Let $v(r,\theta_r,\varphi_r)$ be the 3D convolution of $f(r,\theta_r,\varphi_r)$ and $u(r,\theta_r,\varphi_r)$. Then the 3D Fourier transform of v verifies:

$$\begin{aligned} V(k,\theta_k,\varphi_k) &= \mathcal{F}\{f*u\}(k,\theta_k,\varphi_k)\,, \\ &= \sqrt{(2\pi)^3}F(k,\theta_k,\varphi_k)U(k,\theta_k,\varphi_k)\,, \end{aligned} \tag{C.4}$$

where $\mathcal{F}$ denotes the 3D Fourier transform. From Equation (C.3) the expression of the 3D Fourier transform in spherical coordinates is known in terms of SFB coefficients. Applying this relationship to $V(k,\theta_k,\varphi_k)$ in the last equation yields:

$$(-i)^l\hat{v}_{lm}(k) = \int_0^{2\pi}\int_0^{\pi}\sqrt{(2\pi)^3}F(k,\theta_k,\varphi_k)U(k,\theta_k,\varphi_k)\overline{Y_l^m(\theta_k,\varphi_k)}\sin(\varphi_k)d\varphi_k d\theta_k\,. \tag{C.5}$$

Then, by applying Equation (C.3) to F and U one gets:

$$\begin{aligned} \hat{v}_{lm}(k) &= (i)^l\sqrt{(2\pi)^3}\iint\sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}(-i)^{l'}\hat{f}_{l'm'}(k)Y_{l'}^{m'}(\theta_k,\varphi_k) \\ &\quad\times\sum_{l''=0}^{\infty}\sum_{m''=-l''}^{l''}(-i)^{l''}\hat{u}_{l''m''}(k)Y_{l''}^{m''}(\theta_k,\varphi_k) \\ &\quad\times\overline{Y_l^m(\theta_k,\varphi_k)}\sin(\varphi_k)d\varphi_k d\theta_k, \\ &= (i)^l\sqrt{(2\pi)^3}\sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}(-i)^{l'}\hat{f}_{l'm'}(k) \\ &\quad\times\sum_{l''=0}^{\infty}\sum_{m''=-l''}^{l''}(-i)^{l''}\hat{u}_{l''m''}(k) \\ &\quad\times\iint Y_{l'}^{m'}(\theta_k,\varphi_k)Y_{l''}^{m''}(\theta_k,\varphi_k)\overline{Y_l^m(\theta_k,\varphi_k)}d\Omega_k\,. \end{aligned} \tag{C.6}$$

The last integral over the two angular variables can be expressed as a Slater integral (which is a special case of the Gaunt integral) defined as:

$$c^{l''}(l,m,l',m') = \iint\overline{Y_l^m(\theta,\varphi)}Y_{l'}^{m'}(\theta,\varphi)Y_{l''}^{m-m'}(\theta,\varphi)d\Omega\,. \tag{C.7}$$

The Slater integrals are only nonzero for $|l-l'|\leq l''\leq l+l'$ which simplifies the expression of $\hat{v}_{lm}(k)$.

The SFB transform of the 3D convolution product is therefore:

$$\widehat{(f*u)}_{lm}(k) = (i)^l\sqrt{(2\pi)^3}\sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}(-i)^{l'}\hat{f}_{l'm'}(k)\sum_{l''=|l-l'|}^{l+l'}c^{l''}(l,m,l',m')(-i)^{l''}\hat{u}_{l''m-m'}(k). \tag{C.8}$$

D COSMICPY : A PYTHON PACKAGE FOR INTERACTIVE COSMOLOGY

CONTENTS

D.1 THE COSMICPY PACKAGE

CosmicPy is an interactive Python package that allows for simple cosmological computations. Designed to be modular, well-documented and easily extensible, this package aims to be a convenient tool for forecasting cosmological parameter constraints for different probes and different statistics. Currently, the package includes basic functionalities such as cosmological distances and matter power spectra (based on Eisenstein and Hu (1998) and Smith et al. (2003)), and facilities for computing tomographic (using the Limber approximation) and 3D SFB power spectra for galaxy clustering and the associated Fisher matrices.

Listing 1 illustrates how CosmicPy can be used to easily compute the 3D SFB Fisher matrix, extract the figure of merit, and generate the associated corner plot similar to Figure 10.3.

The full documentation of the package and a number of tutorials demonstrating how to use the different functionalities and reproduce the results of Chapter 10 is provided at the CosmicPy webpage: `http://cosmicpy.github.io` .

Although CosmicPy is primarily written in Python for code readability, it also includes a simple interface to C/C++, allowing critical parts of the codes to have a fast C++ implementation as well as enabling existing codes to be easily interfaced with CosmicPy.

Contributions to the package are very welcome and can be in the form of feedback, requests for additional features, documentation, or even code contributions. This is made simple through the GitHub hosting of the project at

`https://github.com/cosmicpy/cosmicpy` .

D.2 COMPUTING THE SFB COVARIANCE MATRIX

Performing a Fisher analysis requires computing the SFB covariance matrix, and more importantly, computing the inverse of this matrix. This last step can be quite challenging as the covariance of the spherical Fourier-Bessel coefficients is a continuous quantity $C_\ell(k, k')$. Two approaches can be considered to define a covariance matrix in this situation: (i) only using the diagonal covariance $C_\ell(k_i, k_i)$ at discrete points k_i (advocated by Nicola et al. (2014)), or (ii) binning $C_\ell(k, k')$ into bins

```
>>> from cosmicpy import *

# Create a standard cosmology
>>> cosmo = cosmology()

# Setup a spectroscopic survey specifying the redshift distribution,
# fsky, redshift errors and galaxy density per square arcmin.
>>> surv = survey(nzparams={'type':'smail',
                            'a':2.0,
                            'b':1.5,
                            'z0':0.71},
                  fsky=0.3636, zphot_sig=1e-3, ngal=0.9, nzbins=30)

# Cosmological parameters to include in the Fisher matrix analysis
>>> params = ('Omega_m', 'w0', 'wa', 'h', 'Omega_b', 'sigma8', 'n')

# Create a 3D SFB Fisher matrix given a cosmology and survey.
>>> f3d = fisher3d(cosmo, surv, params)

# Create a Tomographic Fisher matrix given a cosmology and survey.
>>> ftomo = fisherTomo(cosmo, surv, params, 'g',
                       cutNonLinearScales='realistic'))

# Output the total Figure of Merit
>>> f3d.FoM
67.516476905327863

>>> ftomo.FoM
67.495159354357327

# Display a corner plot for both fisher matrices
>>> ftomo.corner_plot(nstd=1)

>>> f3d.corner_plot(nstd=1,color='r')
```

Listing 1: Example of 3D SFB and tomographic Fisher matrix computations using CosmicPy.

of size Δ_k. However, by neglecting the correlation between neighbouring wavenumbers, the first approach overestimates the information content if the interval between wavenumbers is too small, while the second approach would lose information for bins of increasing size and become numerically challenging to invert for bins too small. Another problem is to select the largest scale $k_{\min}$ to include in the covariance matrix. Indeed, $C_\ell(k,k)$ becomes extremely small and numerically challenging to compute for very small k, but small wavenumbers can still potentially contribute to the Fisher information. A careful study is necessary to select a $k_{\min}$ that does not lose information.

Instead, using the k_{ln} sampling defined by Equation (10.23) naturally introduces a minimum wavenumber and a discrete sampling of scales that preserves all the information. As an added benefit, this approach yields numerically invertible covariance matrices in practice for sensible choices of the boundary condition $r_{\max}$. Indeed, as $k_{ln} = \frac{q_{ln}}{r_{\max}}$, the choice of cut-off radius sets

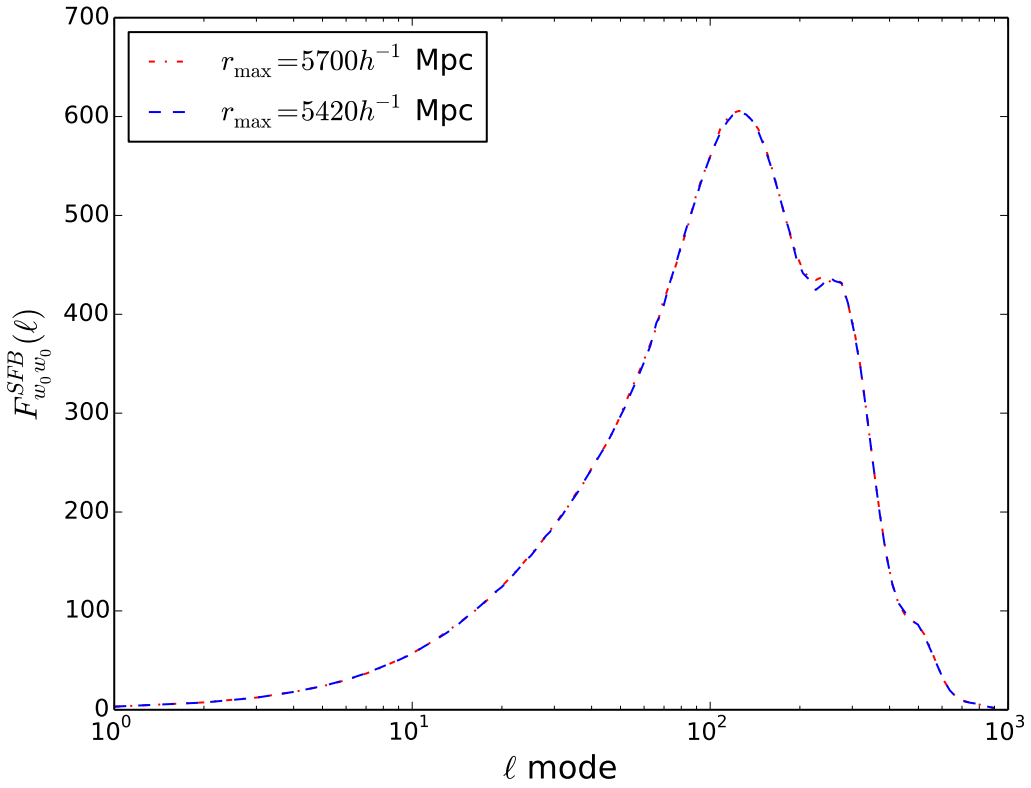

Figure D.1: Contribution to the SFB Fisher matrix element $F^{SFB}_{w_0 w_0}$ as a function of angular mode, computed with different values of $r_{\max}$. The excellent agreement between the two curves shows that our computation of the Fisher matrix is robust to our arbitrary choice of $r_{\max}$.

the fineness of the $C_\ell(n, n')$ matrix and affects its condition number. However, we find that the Fisher information remains largely unaffected by varying $r_{\max}$ above a certain distance because cutting the very end of the galaxy distribution has little effect. In practice, we have arbitrarily set $r_{\max}$ to the comoving distance at which $\phi(r)$ reaches 10^{-5} of its maximum value. This choice has proven stable in all situations considered in this work. The robustness of our computation of the Fisher matrix with respect to the choice of $r_{\max}$ is illustrated in Figure D.1, where we show the contributions of each angular mode to the Fisher matrix element $F^{SFB}_{w_0 w_0}$. Our empirical choice for $r_{\max}$ in this case is 5420 h^{-1}Mpc, but the results are not affected by increasing $r_{\max}$ to 5700 h^{-1} Mpc even more.

D.3 DERIVING THE SPHERICAL FOURIER-BESSEL SHOT NOISE POWER SPECTRUM

Here, we derive the expression of the shot noise by discretising the survey in cells that either contain one or zero galaxies (Peebles, 1980). This method was used in Heavens et al. (2006) to yield the expression of the shot noise in the case of 3D cosmic shear. We considered a point process

defined on small cells c, each of which contains $n_c = 0$ or 1 depending on whether the cell contains a galaxy or not:

$$n(\mathbf{r}) = \frac{1}{V}\sum_c n_c \delta_c(\mathbf{r}) \,, \tag{D.1}$$

where $\delta_c(\mathbf{r}) = 1$ if $\mathbf{r}$ is within the cell c, 0 otherwise, and where n_c fulfils (Peebles, 1980)

$$< n_c >=< n_c^2 >= \bar{\rho}_g^{obs}\phi(r_c)\Delta_c \,, \tag{D.2}$$

where Δ_c is the volume of cell c and $\bar{\rho}_g^{obs}\phi(r_c)$ is the average number density of galaxies of the survey at distance r_c. Furthermore, the cross-term for $c \neq d$ is

$$< n_c n_d >= \bar{\rho}_g^{obs\,2}\,\phi(r_c)\phi(r_d)\Delta_c\Delta_d[1 + \xi(|\mathbf{r_c} - \mathbf{r_d}|)] \,. \tag{D.3}$$

The SFB expansion of the density field can now be expressed as a sum over small cells c:

$$n_{\ell m}(\mathbf{k}) = \sqrt{\frac{2}{\pi}}\sum_c n_c k j_\ell(kr_c) Y_{\ell m}(\Omega_\mathbf{c}) \,. \tag{D.4}$$

From this expression, we can derive the two-point correlation function of this field:

$$\begin{aligned} < n_{\ell m}(\mathbf{k}) n_{\ell' m'}(\mathbf{k}') > &= \frac{2}{\pi}\sum_{c,d} < n_c n_d > kk' j_\ell(kr_c) j_{\ell'}(k'r_d) Y^*_{\ell m}(\Omega_\mathbf{c}) Y_{\ell' m'}(\Omega_\mathbf{d}) \\ &= \frac{2}{\pi}\sum_{c=d} \bar{n}\phi(r_c)\Delta_c kk' j_\ell(kr_c) j_{\ell'}(k'r_c) Y^*_{\ell m}(\Omega_\mathbf{c}) Y_{\ell' m'}(\Omega_\mathbf{c}) \\ &\quad + \frac{2}{\pi}\sum_{c\neq d} \bar{n}^{\,2}\,\phi(r_c)\phi(r_d)\Delta_c\Delta_d[1 + \xi_g(|\mathbf{r_c} - \mathbf{r_d}|)] \\ &\quad \times kk' j_\ell(kr_c) j_{\ell'}(k'r_d) Y^*_{\ell m}(\Omega_\mathbf{c}) Y^*_{\ell' m'}(\Omega_\mathbf{d}) \,. \end{aligned} \tag{D.5}$$

In the last equation, the first term for $c = d$ contains the shot noise contribution and the second term contains the monopole contribution and the correlation function of the density fluctuations. Returning to continuous integration by decreasing the volume of cells Δ_c, we have

$$\begin{aligned} \frac{< n_{\ell m}(\mathbf{k}) n_{\ell' m'}(\mathbf{k}') >}{\bar{n}^{\,2}} &= \underbrace{\frac{2kk'}{\pi}\int \frac{\phi(r)}{\bar{n}} j_\ell(kr) j_\ell(k'r) r^2 dr\; \delta_{\ell\ell'}\delta_{mm'}}_{\text{Shot noise}} \\ &\quad + \underbrace{\frac{2}{\pi}\int \phi(r) k j_\ell(kr) r^2 dr \int \phi(r) k' j_{\ell'}(k'r) r^2 dr\; \delta_{\ell 0}\delta_{m0}\delta_{\ell' 0}\delta_{m' 0}}_{\text{Monopole contribution}} \\ &\quad + \frac{2}{\pi}\iint \xi_g(|\mathbf{r} - \mathbf{r}'|)\phi(r)\phi(r') kk' j_\ell(kr) j_{\ell'}(k'r) Y^*_{\ell m}(\Omega) Y_{\ell' m'}(\Omega') d\mathbf{r} d\mathbf{r}' \,. \end{aligned} \tag{D.6}$$

In this expression, we recognise in order three terms:

- the shot noise contribution, only for $l = l'$ and $m = m'$:

$$\frac{2kk'}{\pi\bar{n}}\int \phi(r, k') j_\ell(kr) j_\ell(k'r) r^2 \mathrm{d}r \,, \tag{D.7}$$

- the monopole contribution, only for $l = 0$ and $m = 0$:

$$M_{\ell m}(k) M_{\ell' m'}(k') \,, \tag{D.8}$$

  with $\sqrt{\frac{2}{\pi}} \int \phi(r) k j_\ell(kr) r^2 dr \delta_{\ell 0} \delta_{m 0}$.

- the contribution from the power spectrum, only for $l = l'$ and $m = m'$:

$$C_\ell(k, k') \,. \tag{D.9}$$

BIBLIOGRAPHY

Adams, J., B. Cresswell, and R. Easther
2001. Inflationary perturbations from a potential with a step. *Physical Review D*, 64(12):123514. (Cited on page 56.)

Aharon, M., M. Elad, and A. Bruckstein
2006. K-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322. (Cited on pages 24 and 27.)

Albrecht, A., G. Bernstein, R. Cahn, W. L. Freedman, J. Hewitt, W. Hu, J. Huth, M. Kamionkowski, E. W. Kolb, L. Knox, J. C. Mather, S. Staggs, and N. B. Suntzeff
2006. Report of the Dark Energy Task Force. *ArXiv e-prints*. (Cited on pages 176 and 187.)

Ashoorioon, A. and A. Krause
2006. Power Spectrum and Signatures for Cascade Inflation. *ArXiv e-prints*, Pp. hep–th/0607001. (Cited on page 56.)

Ashoorioon, A., A. Krause, and K. Turzynski
2009. Energy transfer in multi field inflation and cosmological perturbations. *Journal of Cosmology and Astro-Particle Physics*, 2009(2):014. (Cited on page 56.)

Asorey, J., M. Crocce, E. GaztaA+-aga, and A. Lewis
2012. Recovering 3D clustering information with angular correlations. *Monthly Notices of the Royal Astronomical Society*, 427(3):1891. (Cited on pages 176, 188, 189, 191, and 193.)

Bacon, D. J., D. M. Goldberg, B. T. P. Rowe, and A. N. Taylor
2006. Weak gravitational flexion. *Monthly Notices of the Royal Astronomical Society*, 365(2):414. (Cited on pages 94, 96, and 106.)

Bacon, D. J. and A. N. Taylor
2003. Mapping the 3D dark matter potential with weak shear. *Monthly Notices of the Royal Astronomical Society*, 344(4):1307. (Cited on pages 98, 99, 100, and 130.)

Baddour, N.
2010. Operational and convolution properties of three-dimensional Fourier transforms in spherical polar coordinates. *Journal of the Optical Society of America A*, 27(10):2144. (Cited on pages 161 and 219.)

Bartelmann, M.
1995. Cluster mass estimates from weak lensing. *Astronomy and Astrophysics*, 303:643. (Cited on page 88.)

Bartelmann, M.
2010. Gravitational lensing. *Classical and Quantum Gravity*, 27(23):233001. (Cited on pages 79 and 85.)

Bartelmann, M., R. Narayan, S. Seitz, and P. Schneider
1996. Maximum-likelihood Cluster Reconstruction. *The Astrophysical Journal*, 464:L115. (Cited on page 96.)

Bartelmann, M. and P. Schneider
2001. Weak gravitational lensing. *Physics Reports*, 340(4-5):291. (Cited on page 86.)

Battle, G.
1987. A block spin construction of ondelettes. Part I: Lemarie functions. *Communications in Mathematical Physics*, 110(4):601–615. (Cited on page 45.)

Bauschke, H. H. and P. L. Combettes
2011. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, CMS Books in Mathematics. New York, NY: Springer New York. (Cited on page 29.)

Beck, A. and M. Teboulle
2009. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202. (Cited on pages 32, 64, and 132.)

Beckouche, S., J. L. Starck, and J. Fadili
2013. Astronomical image denoising using dictionary learning. *Astronomy and Astrophysics*, 556:A132. (Cited on page 24.)

Benjamini, Y. and Y. Hochberg
1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300. (Cited on page 171.)

Bobin, J., F. Sureau, P. Paykari, A. Rassat, S. Basak, and J. L. Starck
2013. WMAP nine-year CMB estimation using sparsity. *Astronomy and Astrophysics*, 553:L4. (Cited on pages 55, 65, 66, and 70.)

Bobin, J., F. Sureau, J. L. Starck, A. Rassat, and P. Paykari
2014. Joint Planck and WMAP CMB map reconstruction. *Astronomy and Astrophysics*, 563:A105. (Cited on page 71.)

Bradac, M., M. Lombardi, and P. Schneider
2004. Mass-sheet degeneracy: Fundamental limit on the cluster mass reconstruction from statistical (weak) lensing. *Astronomy and Astrophysics*, 424(1):13–22. (Cited on pages 88 and 118.)

Bradac, M., P. Schneider, M. Lombardi, and T. Erben
2005. Strong and weak lensing united. *Astronomy and Astrophysics*, 437(1):39–48. (Cited on page 98.)

Bridges, M., F. Feroz, M. P. Hobson, and A. N. Lasenby
2009. Bayesian optimal reconstruction of the primordial power spectrum. *Monthly Notices of the Royal Astronomical Society*, 400(2):1075. (Cited on page 57.)

Bridges, M., A. N. Lasenby, and M. P. Hobson
2006. A Bayesian analysis of the primordial power spectrum. *Monthly Notices of the Royal Astronomical Society*, 369(3):1123. (Cited on page 57.)

Bridges, M., A. N. Lasenby, and M. P. Hobson
2007. WMAP 3-yr primordial power spectrum. *Monthly Notices of the Royal Astronomical Society*, 381(1):68. (Cited on page 57.)

Bridle, S. and L. King
2007. Dark energy constraints from cosmic shear power spectra: impact of intrinsic alignments on photometric redshift requirements. *New Journal of Physics*, 9(12):444. (Cited on page 185.)

Bridle, S. L., A. M. Lewis, J. Weller, and G. Efstathiou
2003. Reconstructing the primordial power spectrum. *Monthly Notices of the Royal Astronomical Society*, 342(4):L72–L78. (Cited on page 57.)

Cain, B., M. Bradac, and R. Levinson
2015. Reconstruction of small-scale galaxy cluster substructure with lensing flexion. *ArXiv e-prints*, P. arXiv:1503.08218. (Cited on pages 106, 122, 123, and 124.)

Cain, B., P. L. Schechter, and M. W. Bautz
2011. Measuring Gravitational Lensing Flexion in A1689 Using an Analytic Image Model. *The Astrophysical Journal*, 736(1):43. (Cited on pages 106 and 124.)

Candès, E. J. and D. L. Donoho
1999. Curvelets – A Surprisingly Effective Nonadaptive Representation For Objects with Edges. In *Curve and Surface Fitting: Saint-Malo 1999*, A. Cohen, C. Rabut, and L. L. Schumaker, eds., Nashville, TN. Vanderbilt University Press. (Cited on page 24.)

Candes, E. J. and D. L. Donoho
1999. Ridgelets: a key to higher-dimensional intermittency? *Royal Society of London Philosophical Transactions Series A*, 357(1760):2495. (Cited on page 23.)

Candès, E. J., Y. C. Eldar, D. Needell, and P. Randall
2011. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1):59–73. (Cited on pages 28 and 29.)

Candès, E. J., M. B. Wakin, and S. P. Boyd
2008. Enhancing Sparsity by Reweighted L1 Minimization. *Journal of Fourier Analysis and Applications*, 14(5-6):877–905. (Cited on pages 28, 29, 63, and 110.)

Castro, P. G., A. F. Heavens, and T. D. Kitching
2005. Weak lensing analysis in three dimensions. *Physical Review D*, 72(2):023516. (Cited on pages 160 and 176.)

Chambolle, A. and T. Pock
2011. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145. (Cited on page 33.)

Chevallier, M. and D. Polarski
2001. Accelerating Universes with Scaling Dark Matter. *International Journal of Modern Physics D*, 10(2):213. (Cited on page 184.)

Choudhury, S. and A. Mazumdar
2014. An accurate bound on tensor-to-scalar ratio and the scale of inflation. *Nuclear Physics B*, 882:386. (Cited on page 56.)

Choudhury, S., A. Mazumdar, and S. Pal
2013. Low and high scale MSSM inflation, gravitational waves and constraints from Planck. *Journal of Cosmology and Astroparticle Physics*, 2013(07):041–041. (Cited on page 56.)

Clowe, D., M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky
2006. A Direct Empirical Proof of the Existence of Dark Matter. *The Astrophysical Journal*, 648(2):L109–L113. (Cited on page 5.)

Clowe, D., A. Gonzalez, and M. Markevitch
2004. Weak-Lensing Mass Reconstruction of the Interacting Cluster 1E 0657-558: Direct Evidence for the Existence of Dark Matter. *The Astrophysical Journal*, 604(2):596. (Cited on page 5.)

Cohen, A., I. Daubechies, and J.-C. Feauveau
1992. Biorthogonal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, 45(5):485–560. (Cited on page 46.)

Combettes, P. L. and J.-C. Pesquet
2011. Proximal Splitting Methods in Signal Processing. In *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, volume 49, Pp. 185–212. Springer. (Cited on pages 30 and 31.)

Combettes, P. L. and V. R. Wajs
2005. Signal Recovery by Proximal Forward-Backward Splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200. (Cited on page 32.)

Coupon, J., M. Kilbinger, H. J. McCracken, O. Ilbert, S. Arnouts, Y. Mellier, U. Abbas, S. de la Torre, Y. Goranova, P. Hudelot, J. P. Kneib, and O. Le Fevre
2012. Galaxy clustering in the CFHTLS-Wide: the changing relationship between galaxies and haloes since z ˜ 1.2[?]. *Astronomy and Astrophysics*, 542:A5. (Cited on page 139.)

Covi, L., J. Hamann, A. Melchiorri, A. Slosar, and I. Sorbera
2006. Inflation and WMAP three year data: Features are still present. *Physical Review D*, 74(8):083509. (Cited on page 57.)

Crittenden, R. G. and N. G. Turok
1998. Exactly Azimuthal Pixelizations of the Sky. *ArXiv e-prints*, Pp. astro–ph/9806374. (Cited on page 157.)

Daubechies, I.
1988. Orthonormal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, 41(7):909–996. (Cited on pages 44 and 45.)

Debono, I., A. Rassat, A. Réfrégier, A. Amara, and T. D. Kitching
2010. Weak lensing forecasts for dark energy, neutrinos and initial conditions. *Monthly Notices of the Royal Astronomical Society*, 404:110–119. (Cited on page 56.)

Di Dio, E., F. Montanari, R. Durrer, and J. Lesgourgues
2014. Cosmological parameter estimation with large scale structure observations. *Journal of Cosmology and Astro-Particle Physics*, 2014(1):042. (Cited on pages 176, 188, and 193.)

Di Dio, E., F. Montanari, J. Lesgourgues, and R. Durrer
2013. The CLASSgal code for relativistic cosmological large scale structure. *Journal of Cosmology and Astro-Particle Physics*, 2013(11):044. (Cited on page 176.)

Dodelson, S.
2003. *Modern cosmology*. Academic Press. (Cited on page 13.)

Dolney, D., B. Jain, and M. Takada
2006. Baryon oscillations and dark-energy constraints from imaging surveys. *Monthly Notices of the Royal Astronomical Society*, 366(3):884. (Cited on page 176.)

Donoho, D. L. and X. Huo
2001. Uncertainty principles and ideal atomic decomposition. *IEEE Transactions on Information Theory*, 47(7):2845–2862. (Cited on page 26.)

Doroshkevich, A. G., P. D. Naselsky, O. V. Verkhodanov, D. I. Novikov, V. I. Turchaninov, I. D. Novikov, P. R. Christensen, and L. Y. Chiang
2005. Gauss-Legendre Sky Pixelization (glesp) for CMB Maps. *International Journal of Modern Physics D*, 14(2):275. (Cited on page 157.)

Dyson, F. W., A. S. Eddington, and C. Davidson
1920. A Determination of the Deflection of Light by the Sun's Gravitational Field, from Observations Made at the Total Eclipse of May 29, 1919. (Cited on page 77.)

Eisenstein, D. J. and W. Hu
1998. Baryonic Features in the Matter Transfer Function. *The Astrophysical Journal*, 496(2):605. (Cited on pages 14, 15, 184, and 221.)

Eisenstein, D. J., I. Zehavi, D. W. Hogg, R. Scoccimarro, M. R. Blanton, R. C. Nichol, R. Scranton, H.-J. Seo, M. Tegmark, Z. Zheng, S. F. Anderson, J. Annis, N. Bahcall, J. Brinkmann, S. Burles, F. J. Castander, A. Connolly, I. Csabai, M. Doi, M. Fukugita, J. A. Frieman, K. Glazebrook, J. E. Gunn, J. S. Hendry, G. Hennessy, Z. IveziR, S. Kent, G. R. Knapp, H. Lin, Y.-S. Loh, R. H. Lupton, B. Margon, T. A. McKay, A. Meiksin, J. A. Munn, A. Pope, M. W. Richmond, D. Schlegel, D. P. Schneider, K. Shimasaku, C. Stoughton, M. A. Strauss, M. SubbaRao, A. S. Szalay, I. Szapudi, D. L. Tucker, B. Yanny, and D. G. York
2005. Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. *The Astrophysical Journal*, 633(2):560. (Cited on page 176.)

Elad, M., P. Milanfar, and R. Rubinstein
2007. Analysis versus synthesis in signal priors. *Inverse Problems*, 23(3):947. (Cited on page 27.)

Erdogdu, P., J. P. Huchra, O. Lahav, M. Colless, R. M. Cutri, E. Falco, T. George, T. Jarrett, D. H. Jones, C. S. Kochanek, L. Macri, J. Mader, N. Martimbeau, M. Pahre, Q. Parker, A. Rassat, and W. Saunders
2006a. The dipole anisotropy of the 2 Micron All-Sky Redshift Survey. *Monthly Notices of the Royal Astronomical Society*, 368(4):1515–1526. (Cited on page 176.)

Erdogdu, P., O. Lahav, J. P. Huchra, M. Colless, R. M. Cutri, E. Falco, T. George, T. Jarrett, D. H. Jones, L. M. Macri, J. Mader, N. Martimbeau, M. A. Pahre, Q. A. Parker, A. Rassat, and W. Saunders
2006b. Reconstructed density and velocity fields from the 2MASS Redshift Survey. *Monthly Notices of the Royal Astronomical Society*, 373(1):45–64. (Cited on page 176.)

Eriksen, M. and E. Gaztanaga
2015. Combining spectroscopic and photometric surveys: Same or different sky? *Monthly Notices of the Royal Astronomical Society*, 451(2):1553–1560. (Cited on page 189.)

Feng, B. and X. Zhang
2003. Double inflation and the low CMB quadrupole. *Physics Letters B*, 570(3-4):145. (Cited on page 56.)

Fisher, K. B., O. Lahav, Y. Hoffman, D. Lynden-Bell, and S. Zaroubi
1995. Wiener reconstruction of density, velocity and potential fields from all-sky galaxy redshift surveys. *Monthly Notices of the Royal Astronomical Society*, 272(4):885. (Cited on pages 176, 180, and 181.)

Gao, H.-Y. H. and A. G. Bruce
1997. WaveShrink with firm shrinkage. *Statistica Sinica*, 7:855–874. (Cited on page 133.)

Gauthier, C. and M. Bucher
2012. Reconstructing the primordial power spectrum from the CMB. *Journal of Cosmology and Astro-Particle Physics*, 2012(10):050. (Cited on page 57.)

Gentile, M., F. Courbin, and G. Meylan
2012. A simple fitting method (gfit) for galaxy shape measurement in weak lensing surveys. *ArXiv e-prints*, P. arXiv:1211.4847. (Cited on page 92.)

Goldberg, D. M. and D. J. Bacon
2005. Galaxy-Galaxy Flexion: Weak Lensing to Second Order. *The Astrophysical Journal*, 619(2):741. (Cited on pages 84 and 106.)

Gorski, K. M., E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann
2005. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal*, 622(2):759–771. (Cited on pages 157, 158, 162, and 165.)

Goswami, G. and J. Prasad
2013. Maximum entropy deconvolution of primordial power spectrum. *Physical Review D*, 88(2):023522. (Cited on page 57.)

Grassi, A. and B. M. Schafer
2014. Detecting baryon acoustic oscillations by 3d weak lensing. *Monthly Notices of the Royal Astronomical Society*, 437(3):2632. (Cited on page 176.)

Gray, M. E., C. Wolf, M. Barden, C. Y. Peng, B. Haussler, E. F. Bell, D. H. McIntosh, Y. Guo, J. A. R. Caldwell, D. Bacon, M. Balogh, F. D. Barazza, A. Bohm, C. Heymans, K. Jahnke, S. Jogee, E. van Kampen, K. Lane, K. Meisenheimer, S. F. Sanchez, A. Taylor, L. Wisotzki, X. Zheng, D. A. Green,

R. J. Beswick, D. J. Saikia, R. Gilmour, B. D. Johnson, and C. Papovich
2009. STAGES: the Space Telescope A901/2 Galaxy Evolution Survey. *Monthly Notices of the Royal Astronomical Society*, 393(4):1275. (Cited on page 101.)

Grossmann, A. and J. Morlet
1984. Decomposition of Hardy Functions into Square Integrable Wavelets of Constant Shape. *SIAM Journal on Mathematical Analysis*, 15(4):723–736. (Cited on page 38.)

Guo, Z.-K., D. J. Schwarz, and Y.-Z. Zhang
2011. Reconstruction of the primordial power spectrum from CMB data. *Journal of Cosmology and Astro-Particle Physics*, 2011(8):031. (Cited on page 57.)

Guth, A. H.
1981. Inflationary universe: A possible solution to the horizon and flatness problems. *Physical Review D*, 23(2):347–356. (Cited on pages 11 and 56.)

Haar, A.
1910. Zur Theorie der orthogonalen Funktionensysteme. *Mathematische Annalen*, 69(3):38–53. (Cited on page 39.)

Hamann, J., A. Shafieloo, and T. Souradeep
2010. Features in the primordial power spectrum? A frequentist analysis. *Journal of Cosmology and Astroparticle Physics*, 2010(04):010–010. (Cited on page 57.)

Hannestad, S.
2001. Reconstructing the inflationary power spectrum from cosmic microwave background radiation data. *Physical Review D*, 63(4):043009. (Cited on page 57.)

Hannestad, S.
2003. Can cosmology detect hierarchical neutrino masses? *Physical Review D*, 67(8):085017. (Cited on page 57.)

Harrison, E. R.
1970. Fluctuations at the Threshold of Classical Cosmology. *Physical Review D*, 1(10):2726–2730. (Cited on pages 12 and 56.)

Harvey, D., R. Massey, T. Kitching, A. Taylor, and E. Tittley
2015. The nongravitational interactions of dark matter in colliding galaxy clusters. *Science*, 347(6229):1462. (Cited on page 196.)

Hazra, D. K., M. Aich, R. K. Jain, L. Sriramkumar, and T. Souradeep
2010. Primordial features due to a step in the inflaton potential. *Journal of Cosmology and Astro-Particle Physics*, 2010(10):008. (Cited on page 57.)

Hazra, D. K., A. Shafieloo, and T. Souradeep
2013. Primordial power spectrum: a complete analysis with the WMAP nine-year data. *Journal of Cosmology and Astro-Particle Physics*, 2013(7):031. (Cited on page 57.)

Heavens, A.
2003. 3D weak lensing. *Monthly Notices of the Royal Astronomical Society*, 343(4):1327. (Cited on page 176.)

Heavens, A. F., T. D. Kitching, and A. N. Taylor
2006. Measuring dark energy properties with 3D cosmic shear. *Monthly Notices of the Royal Astronomical Society*, 373(1):105. (Cited on pages 180 and 223.)

Heavens, A. F. and A. N. Taylor
1995. A spherical harmonic analysis of redshift space. *Monthly Notices of the Royal Astronomical Society*, 275(2):483. (Cited on pages 156, 160, 176, 177, and 180.)

Hennawi, J. F. and D. N. Spergel
2005. Shear-selected Cluster Cosmology: Tomography and Optimal Filtering. *The Astrophysical Journal*, 624(1):59. (Cited on page 151.)

Heymans, C., L. Van Waerbeke, D. Bacon, J. Berge, G. Bernstein, E. Bertin, S. Bridle, M. L. Brown, D. Clowe, H. Dahle, T. Erben, M. Gray, M. Hetterscheidt, H. Hoekstra, P. Hudelot, M. Jarvis, K. Kuijken, V. Margoniner, R. Massey, Y. Mellier, R. Nakajima, A. Refregier, J. Rhodes, T. Schrabback, and D. Wittman
2006. The Shear Testing Programme - I. Weak lensing analysis of simulated ground-based observations. *Monthly Notices of the Royal Astronomical Society*, 368(3):1323. (Cited on page 92.)

Hinshaw, G., D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, L. Page, K. M. Smith, J. L. Weiland, B. Gold, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, G. S. Tucker, E. Wollack, and E. L. Wright
2013. Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *The Astrophysical Journal Supplement Series*, 208(2):19. (Cited on pages 64, 65, 69, and 71.)

Hivon, E., K. M. Grnrski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen
2002. MASTER of the Cosmic Microwave Background Anisotropy Power Spectrum: A Fast Method for Statistical Analysis of Large and Complex Cosmic Microwave Background Data Sets. *The Astrophysical Journal*, 567(1):2. (Cited on pages 59 and 182.)

Hoekstra, H., M. Franx, K. Kuijken, and G. Squires
1998. Weak lensing analysis of Cl 1358+62. *New Astronomy Reviews*, 42(2):137–140. (Cited on page 92.)

Holschneider, M., R. Kronland-Martinet, J. Morlet, and P. Tchamitchian
1989. A Real-Time Algorithm for Signal Analysis with the Help of the Wavelet Transform. In *Wavelets. Time-Frequency Methods and Phase Space*, J. M. Combes, A. Grossmann, and P. Tchamitchian, eds., P. 286. Springer-Verlag. (Cited on page 50.)

Hou, Z., C. L. Reichardt, K. T. Story, B. Follin, R. Keisler, K. a. Aird, B. a. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H.-M. Cho, T. M. Crawford, a. T. Crites, T. de Haan, R. de Putter, M. a. Dobbs, S. Dodelson, J. Dudley, E. M. George, N. W. Halverson, G. P. Holder, W. L. Holzapfel, S. Hoover, J. D. Hrubes, M. Joy, L. Knox, a. T. Lee, E. M. Leitch, M. Lueker, D. Luong-Van, J. J. McMahon, J. Mehl, S. S. Meyer, M. Millea, J. J. Mohr, T. E. Montroy, S. Padin, T. Plagge, C. Pryke, J. E. Ruhl, J. T. Sayre, K. K. Schaffer, L. Shaw, E. Shirokoff, H. G. Spieler, Z. Staniszewski, a. a. Stark, A. van Engelen, K. Vanderlinde, J. D. Vieira, R. Williamson, and O. Zahn
2014. Constraints on Cosmology from the Cosmic Microwave Background Power Spectrum of the 2500 deg2 SPT-SZ Survey. *The Astrophysical Journal*, 782(2):74. (Cited on pages 57 and 65.)

Hu, W. and B. Jain
2004. Joint galaxy-lensing observables and the dark energy. *Physical Review D*, 70(4):043009. (Cited on page 183.)

Hu, W. and C. R. Keeton
2002. Three-dimensional mapping of dark matter. *Physical Review D*, 66(6):063506. (Cited on pages 98, 99, and 100.)

Hu, W. and T. Okamoto
2004. Principal power of the CMB. *Physical Review D*, 69(4):043004. (Cited on page 57.)

Hubble, E.
1929. A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, 15(3):168. (Cited on page 9.)

Huffenberger, K. M. and B. D. Wandelt
2010. Fast and Exact Spin-s Spherical Harmonic Transforms. *The Astrophysical Journal Supplement Series*, 189(2):255. (Cited on page 158.)

Hunt, P. and S. Sarkar
2004. Multiple inflation and the WMAP "glitches". *Physical Review D*, 70(10):103518. (Cited on page 56.)

Hunt, P. and S. Sarkar
2007. Multiple inflation and the WMAP "glitches". II. Data analysis and cosmological parameter extraction. *Physical Review D*, 76(12):123504. (Cited on page 56.)

Ichiki, K. and R. Nagata
2009. Brute force reconstruction of the primordial fluctuation spectrum from five-year Wilkinson Microwave Anisotropy Probe observations. *Physical Review D*, 80(8):083002. (Cited on page 57.)

Ichiki, K., R. Nagata, and J. Yokoyama
2010. Cosmic discordance: Detection of a modulation in the primordial fluctuation spectrum. *Physical Review D*, 81(8):083010. (Cited on page 57.)

Jain, R. K., P. Chingangbam, J.-O. Gong, L. Sriramkumar, and T. Souradeep
2009. Punctuated inflation and the low CMB multipoles. *Journal of Cosmology and Astroparticle Physics*, 2009(01):009–009. (Cited on page 56.)

Joachimi, B. and S. L. Bridle
2010. Simultaneous measurement of cosmology and intrinsic alignments using joint cosmic shear and galaxy number density correlations. *Astronomy and Astrophysics*, 523:A1. (Cited on pages 183, 184, 185, 187, and 197.)

Joy, M., V. Sahni, and A. A. Starobinsky
2008. New universal local feature in the inflationary perturbation spectrum. *Physical Review D*, 77(2):023514. (Cited on page 56.)

Joy, M., A. Shafieloo, V. Sahni, and A. A. Starobinsky
2009. Is a step in the primordial spectral index favoured by CMB data? *Journal of Cosmology and Astroparticle Physics*, 2009(06):028–028. (Cited on page 57.)

Kaiser, N. and G. Squires
1993. Mapping the Dark Matter with Weak Gravitational Lensing. *The Astrophysical Journal*, 404:441. (Cited on page 93.)

Kaiser, N., G. Squires, and T. Broadhurst
1995. A Method for Weak Lensing Observations. *The Astrophysical Journal*, 449:460. (Cited on page 92.)

Keiner, J., S. Kunis, and D. Potts
2009. Using NFFT 3—A Software Library for Various Nonequispaced Fast Fourier Transforms. *ACM Transactions on Mathematical Software*, 36(4):1–30. (Cited on page 109.)

Kirk, D., A. Rassat, O. Host, and S. Bridle
2012. The cosmological impact of intrinsic alignment model choice for cosmic shear. *Monthly Notices of the Royal Astronomical Society*, 424(3):1647. (Cited on pages 176 and 185.)

Kitching, T. D., A. F. Heavens, J. Alsing, T. Erben, C. Heymans, H. Hildebrandt, H. Hoekstra, A. Jaffe, A. Kiessling, Y. Mellier, L. Miller, L. van Waerbeke, J. Benjamin, J. Coupon, L. Fu, M. J. Hudson, M. Kilbinger, K. Kuijken, B. T. P. Rowe, T. Schrabback, E. Semboloni, and M. Velander
2014. 3D cosmic shear: cosmology from CFHTLenS. *Monthly Notices of the Royal Astronomical Society*, 442(2):1326. (Cited on pages 176 and 182.)

Kitching, T. D., A. F. Heavens, and L. Miller
2011. 3D photometric cosmic shear. *Monthly Notices of the Royal Astronomical Society*, 413(4):2923. (Cited on pages 90 and 176.)

Kitching, T. D., L. Miller, C. E. Heymans, L. van Waerbeke, and A. F. Heavens
2008a. Bayesian galaxy shape measurement for weak lensing surveys - II. Application to simulations. *Monthly Notices of the Royal Astronomical Society*, 390(1):149. (Cited on page 92.)

Kitching, T. D., A. N. Taylor, and A. F. Heavens
2008b. Systematic effects on dark energy from 3D weak shear. *Monthly Notices of the Royal Astronomical Society*, 389(1):173. (Cited on page 176.)

Klypin, A. A., S. Trujillo-Gomez, and J. Primack
2011. Dark Matter Halos in the Standard Cosmological Model: Results from the Bolshoi Simulation. *The Astrophysical Journal*, 740(2):102. (Cited on pages 114 and 122.)

Kogo, N., M. Matsumiya, M. Sasaki, and J. Yokoyama
2004a. Reconstructing the Primordial Spectrum from WMAP Data by the Cosmic Inversion Method. *The Astrophysical Journal*, 607(1):32–39. (Cited on page 57.)

Kogo, N., M. Sasaki, and J. Yokoyama
2004b. Reconstructing the primordial spectrum with CMB temperature and polarization. *Physical Review D*, 70(10):103001. (Cited on page 57.)

Kuijken, K.
1999. Weak weak lensing: correcting weak shear measurements accurately for PSF anisotropy. *Astronomy and Astrophysics*, 352:355. (Cited on page 92.)

Kumazaki, K., S. Yokoyama, and N. Sugiyama
2011. Fine features in the primordial power spectrum. *Journal of Cosmology and Astro-Particle Physics*, 2011(12):008. (Cited on page 56.)

Lanusse, F., P. Paykari, J. L. Starck, F. Sureau, J. Bobin, and A. Rassat
2014. PRISM: Recovery of the primordial spectrum from Planck data. *Astronomy and Astrophysics*, 571:L1. (Cited on pages 55, 64, and 156.)

Lanusse, F., A. Rassat, and J. L. Starck
2012. Spherical 3D isotropic wavelets. *Astronomy and Astrophysics*, 540:A92. (Cited on pages 155 and 176.)

Lanusse, F., A. Rassat, and J. L. Starck
2015. 3D galaxy clustering with future wide-field surveys: Advantages of a spherical Fourier-Bessel analysis. *Astronomy and Astrophysics*, 578:A10. (Cited on page 175.)

Laureijs, R., J. Amiaux, S. Arduini, J. L. Augueres, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet, B. Garilli, P. Gondoin, L. Guzzo, J. Hoar, H. Hoekstra, R. Holmes, T. Kitching, T. Maciaszek, Y. Mellier, F. Pasian, W. Percival, J. Rhodes, G. Saavedra Criado, M. Sauvage, R. Scaramella, L. Valenziano, S. Warren, R. Bender, F. Castander, A. Cimatti, O. Le Fevre, H. Kurki-Suonio, M. Levi, P. Lilje, G. Meylan, R. Nichol, K. Pedersen, V. Popa, R. Rebolo Lopez, H. W. Rix, H. Rottgering, W. Zeilinger, F. Grupp, P. Hudelot, R. Massey, M. Meneghetti, L. Miller, S. Paltani, S. Paulin-Henriksson, S. Pires, C. Saxton, T. Schrabback, G. Seidel, J. Walsh, N. Aghanim, L. Amendola, J. Bartlett, C. Baccigalupi, J. P. Beaulieu, K. Benabed, J. G. Cuby, D. Elbaz, P. Fosalba, G. Gavazzi, A. Helmi, I. Hook, M. Irwin, J. P. Kneib, M. Kunz, F. Mannucci, L. Moscardini, C. Tao, R. Teyssier, J. Weller, G. Zamorani, M. R. Zapatero Osorio, O. Boulade, J. J. Foumond, A. Di Giorgio, P. Guttridge, A. James, M. Kemp, J. Martignac, A. Spencer, D. Walton, T. Blumchen, C. Bonoli, F. Bortoletto, C. Cerna, L. Corcione, C. Fabron, K. Jahnke, S. Ligori, F. Madrid, L. Martin, G. Morgante, T. Pamplona, E. Prieto, M. Riva, R. Toledo, M. Trifoglio, F. Zerbi, F. Abdalla, M. Douspis, C. Grenet, S. Borgani, R. Bouwens, F. Courbin, J. M. Delouis, P. Dubath, A. Fontana, M. Frailis, A. Grazian, J. Koppenhofer, O. Mansutti, M. Melchior, M. Mignoli, J. Mohr, C. Neissner, K. Noddle, M. Poncet, M. Scodeggio, S. Serrano, N. Shane, J. L. Starck, C. Surace, A. Taylor, G. Verdoes-Kleijn, C. Vuerli, O. R. Williams, A. Zacchei, B. Altieri, I. Escudero Sanz, R. Kohley, T. Oosterbroek, P. Astier, D. Bacon, S. Bardelli, C. Baugh, F. Bellagamba, C. Benoist, D. Bianchi, A. Biviano, E. Branchini, C. Carbone, V. Cardone, D. Clements, S. Colombi, C. Conselice, G. Cresci, N. Deacon, J. Dunlop, C. Fedeli, F. Fontanot, P. Franzetti, C. Giocoli, J. Garcia-Bellido, J. Gow, A. Heavens, P. Hewett, C. Heymans, A. Holland, Z. Huang, O. Ilbert, B. Joachimi, E. Jennins, E. Kerins, A. Kiessling, D. Kirk, R. Kotak, O. Krause, O. Lahav, F. van Leeuwen, J. Lesgourgues, M. Lombardi, M. Magliocchetti, K. Maguire, E. Majerotto, R. Maoli, F. Marulli, S. Maurogordato, H. McCracken, R. McLure, A. Melchiorri, A. Merson, M. Moresco, M. Nonino, P. Norberg, J. Peacock, R. Pello, M. Penny, V. Pettorino, C. Di Porto, L. Pozzetti, C. Quercellini, M. Radovich, A. Rassat, N. Roche, S. Ronayette, E. Rossetti, B. Sartoris, P. Schneider, E. Semboloni, S. Serjeant, F. Simpson, C. Skordis, G. Smadja, S. Smartt, P. Spano, S. Spiro, M. Sullivan, A. Tilquin, R. Trotta, L. Verde, Y. Wang,

G. Williger, G. Zhao, J. Zoubian, and E. Zucca
2011. Euclid Definition Study Report. *ArXiv e-prints*, P. arXiv:1110.3193. (Cited on pages 176 and 184.)

Leach, S.
2006. Measuring the primordial power spectrum: principal component analysis of the cosmic microwave background. *Monthly Notices of the Royal Astronomical Society*, 372(2):646. (Cited on page 57.)

Leistedt, B. and J. D. McEwen
2012. Exact Wavelets on the Ball. *IEEE Transactions on Signal Processing*, 60(12):6257–6269. (Cited on page 156.)

Leistedt, B., H. V. Peiris, and J. D. McEwen
2013. Flaglets for studying the large-scale structure of the Universe. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, D. Van De Ville, V. K. Goyal, and M. Papadakis, eds., volume 8858, P. 88580J. (Cited on page 196.)

Leistedt, B., A. Rassat, A. Refregier, and J. L. Starck
2012. 3DEX: a code for fast spherical Fourier-Bessel decomposition of 3D surveys. *Astronomy and Astrophysics*, 540:A60. (Cited on pages 160 and 176.)

Lemarié, P. G.
1988. Ondelettes à localisation exponentielle. *Journal de Mathématiques Pures et Appliquées*, 67:227–236. (Cited on page 45.)

Lemoine, D.
1994. The discrete Bessel transform algorithm. *Journal of Chemical Physics*, 101(5):3936. (Cited on pages 156, 162, and 164.)

Leonard, A., F.-X. Dupé, and J.-L. Starck
2012. A compressed sensing approach to 3D weak lensing. *Astronomy & Astrophysics*, 539:A85. (Cited on page 98.)

Leonard, A., D. M. Goldberg, J. L. Haaga, and R. Massey
2007. Gravitational Shear, Flexion and Strong Lensing in Abell 1689. *The Astrophysical Journal*, 666(1):17. (Cited on page 106.)

Leonard, A. and L. J. King
2010. A new tool to determine masses and mass profiles using gravitational flexion. *Monthly Notices of the Royal Astronomical Society*, 405(3):1854. (Cited on page 106.)

Leonard, A., L. J. King, and S. M. Wilkins
2009. Detecting mass substructure in galaxy clusters: an aperture mass statistic for gravitational flexion. *Monthly Notices of the Royal Astronomical Society*, 395(3):1438. (Cited on page 106.)

Leonard, A., F. Lanusse, and J.-L. Starck
2014. GLIMPSE: accurate 3D weak lensing reconstructions using sparsity. *Monthly Notices of the Royal Astronomical Society*, 440(2):1281–1294. (Cited on page 129.)

Leonard, A., F. Lanusse, and J.-L. Starck
2015. Weak lensing reconstructions in 2D and 3D: implications for cluster studies. *Monthly Notices of the Royal Astronomical Society*, 449(1):1146. (Cited on page 129.)

Lerner, R. N. and J. McDonald
2009. Space-dependent step features: Transient breakdown of slow-roll, homogeneity, and isotropy during inflation. *Physical Review D*, 79(2):023511. (Cited on page 56.)

Lesgourgues, J.
2000. Features in the primordial power spectrum of double D-term inflation. *Nuclear Physics B*, 582(1-3):593. (Cited on page 56.)

Lesgourgues, J.
2011. The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview. *ArXiv e-prints*, P. arXiv:1104.2932. (Cited on pages 14 and 66.)

Lewis, A., A. Challinor, and A. Lasenby
2000. Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models. *The Astrophysical Journal*, 538(2):473–476. (Cited on page 14.)

Linde, A.
1982. A new inflationary Universe scenario: a possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108:389–393. (Cited on pages 11 and 56.)

Linder, E. V.
2003. Exploring the Expansion History of the Universe. *Physical Review Letters*, 90(9):091301. (Cited on page 8.)

Loverde, M. and N. Afshordi
2008. Extended Limber approximation. *Physical Review D*, 78(12):123506. (Cited on page 179.)

Luppino, G. A. and N. Kaiser
1997. Detection of Weak Lensing by a Cluster of Galaxies at z = 0.83. *The Astrophysical Journal*, 475(1):20. (Cited on page 92.)

Ma, Z., W. Hu, and D. Huterer
2006. Effects of Photometric Redshift Uncertainties on Weak-Lensing Tomography. *The Astrophysical Journal*, 636(1):21. (Cited on pages 130 and 139.)

Mallat, S.
1999. *A Wavelet Tour of Signal Processing*. ACADEMIC PRESS. (Cited on page 47.)

Mallat, S. G.
1989. Multiresolution Approximations and Wavelet Orthonormal Bases of L 2 (R). *Transactions of the American Mathematical Society*, 315(1):69. (Cited on pages 40, 41, 42, and 43.)

Mandelbaum, R., B. Rowe, R. Armstrong, D. Bard, E. Bertin, J. Bosch, D. Boutigny, F. Courbin, W. A. Dawson, A. Donnarumma, I. Fenech Conti, R. Gavazzi, M. Gentile, M. S. S. Gill, D. W. Hogg, E. M.

Huff, M. J. Jee, T. Kacprzak, M. Kilbinger, T. Kuntzer, D. Lang, W. Luo, M. C. March, P. J. Marshall, J. E. Meyers, L. Miller, H. Miyatake, R. Nakajima, F. M. Ngole Mboula, G. Nurbaeva, Y. Okura, S. Paulin-Henriksson, J. Rhodes, M. D. Schneider, H. Shan, E. S. Sheldon, M. Simet, J.-L. Starck, F. Sureau, M. Tewes, K. Zarb Adami, J. Zhang, and J. Zuntz
2015. GREAT3 results - I. Systematic errors in shear estimation and the impact of real galaxy morphology. *Monthly Notices of the Royal Astronomical Society*, 450(3):2963. (Cited on page 92.)

Mandelbaum, R., B. Rowe, J. Bosch, C. Chang, F. Courbin, M. Gill, M. Jarvis, A. Kannawadi, T. Kacprzak, C. Lackner, A. Leauthaud, H. Miyatake, R. Nakajima, J. Rhodes, M. Simet, J. Zuntz, B. Armstrong, S. Bridle, J. Coupon, J. P. Dietrich, M. Gentile, C. Heymans, A. S. Jurling, S. M. Kent, D. Kirkby, D. Margala, R. Massey, P. Melchior, J. Peterson, A. Roodman, and T. Schrabback
2014. The Third Gravitational Lensing Accuracy Testing (GREAT3) Challenge Handbook. *The Astrophysical Journal Supplement Series*, 212(1):5. (Cited on pages 90 and 92.)

Marinucci, D., D. Pietrobon, A. Balbi, P. Baldi, P. Cabella, G. Kerkyacharian, P. Natoli, D. Picard, and N. Vittorio
2008. Spherical needlets for cosmic microwave background data analysis. *Monthly Notices of the Royal Astronomical Society*, 383(2):539. (Cited on page 170.)

Massey, R., J. Rhodes, R. Ellis, N. Scoville, A. Leauthaud, A. Finoguenov, P. Capak, D. Bacon, H. Aussel, J.-P. Kneib, A. Koekemoer, H. McCracken, B. Mobasher, S. Pires, A. Refregier, S. Sasaki, J.-L. Starck, Y. Taniguchi, A. Taylor, and J. Taylor
2007. Dark matter maps reveal cosmic scaffolding. *Nature*, 445(7125):286. (Cited on page 99.)

Massey, R., L. Williams, R. Smit, M. Swinbank, T. D. Kitching, D. Harvey, M. Jauzac, H. Israel, D. Clowe, A. Edge, M. Hilton, E. Jullo, A. Leonard, J. Liesenborgs, J. Merten, I. Mohammed, D. Nagai, J. Richard, A. Robertson, P. Saha, R. Santana, J. Stott, and E. Tittley
2015. The behaviour of dark matter associated with four bright cluster galaxies in the 10 kpc core of Abell 3827. *Monthly Notices of the Royal Astronomical Society*, 449(4):3393–3406. (Cited on page 196.)

Mathews, G. J., D. J. Chung, K. Ichiki, T. Kajino, and M. Orito
2004. Constraints on resonant particle production during inflation from the matter and CMB power spectra. *Physical Review D*, 70(8):083505. (Cited on page 56.)

Matsumiya, M., M. Sasaki, and J. Yokoyama
2002. Cosmic inversion: Reconstructing the primordial spectrum from CMB anisotropy. *Physical Review D*, 65(8):083007. (Cited on page 57.)

McEwen, J. D.
2008. Fast, exact (but unstable) spin spherical harmonic transforms. *ArXiv e-prints*, P. arXiv:0807.4494. (Cited on page 157.)

McEwen, J. D. and Y. Wiaux
2011. A Novel Sampling Theorem on the Sphere. *IEEE Transactions on Signal Processing*, 59(12):5876. (Cited on pages 156 and 158.)

Meerburg, P. D., R. A. M. J. Wijers, and J. P. van der Schaar
2012. WMAP7 constraints on oscillations in the primordial power spectrum. *Monthly Notices of the Royal Astronomical Society*, 421(1):369. (Cited on page 56.)

Merkel, P. M. and B. M. Schafer
2013. Intrinsic alignments and 3d weak gravitational lensing. *Monthly Notices of the Royal Astronomical Society*, 434(2):1808. (Cited on page 176.)

Merten, J., M. Cacciato, M. Meneghetti, C. Mignone, and M. Bartelmann
2009. Combining weak and strong cluster lensing: applications to simulations and MS 2137. *Astronomy and Astrophysics*, 500(2):681. (Cited on page 117.)

Meyer, Y.
1992. *Wavelets and operators*. Cambridge University Press. (Cited on pages 40, 41, 42, and 43.)

Michel, V.
2005. Wavelets on the 3-dimensional Ball. *PAMM*, 5(1):775–776. (Cited on page 156.)

Miller, C. J., C. Genovese, R. C. Nichol, L. Wasserman, A. Connolly, D. Reichart, A. Hopkins, J. Schneider, and A. Moore
2001. Controlling the False-Discovery Rate in Astrophysical Data Analysis. *The Astronomical Journal*, 122(6):3492. (Cited on page 171.)

Miller, L., T. D. Kitching, C. Heymans, A. F. Heavens, and L. van Waerbeke
2007. Bayesian galaxy shape measurement for weak lensing surveys - I. Methodology and a fast-fitting algorithm. *Monthly Notices of the Royal Astronomical Society*, 382(1):315. (Cited on page 92.)

Misner, C., K. Thorne, and J. Wheeler
1973. *Gravitation*. Macmillan. (Cited on page 84.)

Mollerach, S. and E. Roulet
2002. Gravitational lensing theory. In *Gravitational Lensing and Microlensing*, volume 22, Pp. 29–68. WORLD SCIENTIFIC. (Cited on page 84.)

Moreau, J.-J.
1962. Fonctions convexes duales et points proximaux dans un espace hilbertien. *CRAS Sér. A Math.*, 255:2897–2899. (Cited on page 30.)

Mortonson, M. J., C. Dvorkin, H. V. Peiris, and W. Hu
2009. CMB polarization features from inflation versus reionization. *Physical Review D*, 79(10):103519. (Cited on page 57.)

Mukherjee, P. and Y. Wang
2003. Model-independent Reconstruction of the Primordial Power Spectrum from Wilkinson Microwave Anistropy Probe Data. *The Astrophysical Journal*, 599(1):1–6. (Cited on page 57.)

Mukherjee, P. and Y. Wang
2005. Primordial power spectrum reconstruction. *Journal of Cosmology and Astro-Particle Physics*, 2005(12):007. (Cited on page 57.)

Nagata, R. and J. Yokoyama
2008. Reconstruction of the primordial fluctuation spectrum from the five-year WMAP data by the cosmic inversion method with band-power decorrelation analysis. *Physical Review D*, 78(12):123002. (Cited on page 57.)

Nagata, R. and J. Yokoyama
2009. Band-power reconstruction of the primordial fluctuation spectrum by the maximum likelihood reconstruction method. *Physical Review D*, 79(4):043010. (Cited on page 57.)

Nam, S., M. Davies, M. Elad, and R. Gribonval
2013. The cosparse analysis model and algorithms. *Applied and Computational Harmonic Analysis*, 34(1):30–56. (Cited on page 28.)

Nam, S., M. E. Davies, M. Elad, and R. Gribonval
2011. Cosparse analysis modeling - uniqueness and algorithms. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 2, Pp. 5804–5807. IEEE. (Cited on page 28.)

Natarajan, B. K.
1995. Sparse Approximate Solutions to Linear Systems. *SIAM Journal on Computing*, 24(2):227–234. (Cited on page 26.)

Navarro, J. F., C. S. Frenk, and S. D. M. White
1997. A Universal Density Profile from Hierarchical Clustering. *The Astrophysical Journal*, 490(2):493. (Cited on pages 137 and 138.)

Nicholson, G. and C. R. Contaldi
2008. The large scale cosmic microwave background cut-off and the tensor-to-scalar ratio. *Journal of Cosmology and Astroparticle Physics*, 2008(01):002. (Cited on page 56.)

Nicholson, G. and C. R. Contaldi
2009. Reconstruction of the primordial power spectrum using temperature and polarisation data from multiple experiments. *Journal of Cosmology and Astro-Particle Physics*, 2009(7):011. (Cited on pages 57 and 65.)

Nicholson, G., C. R. Contaldi, and P. Paykari
2010. Reconstruction of the primordial power spectrum by direct inversion. *Journal of Cosmology and Astro-Particle Physics*, 2010(1):016. (Cited on pages 57 and 196.)

Nicola, A., A. Refregier, A. Amara, and A. Paranjape
2014. Three-dimensional spherical analyses of cosmological spectroscopic surveys. *Physical Review D*, 90(6):063515. (Cited on pages 176, 177, 185, 187, 188, 193, and 221.)

Okura, Y., K. Umetsu, and T. Futamase
2007. A New Measure for Weak-Lensing Flexion. *The Astrophysical Journal*, 660(2):995. (Cited on page 106.)

Pahud, C., M. Kamionkowski, and A. R. Liddle
2009. Oscillations in the inflaton potential? *Physical Review D*, 79(8):083503. (Cited on page 56.)

Parkinson, D., S. Tsujikawa, B. A. Bassett, and L. Amendola
2005. Testing for double inflation with WMAP. *Physical Review D*, 71(6):063524. (Cited on page 57.)

Paykari, P. and A. H. Jaffe
2010. Optimal Binning of the Primordial Power Spectrum. *The Astrophysical Journal*, 711(1):1. (Cited on page 57.)

Paykari, P., F. Lanusse, J. L. Starck, F. Sureau, and J. Bobin
2014. PRISM: Sparse recovery of the primordial power spectrum. *Astronomy and Astrophysics*, 566:A77. (Cited on pages 55 and 64.)

Paykari, P., J.-L. Starck, and M. J. Fadili
2012. True cosmic microwave background power spectrum estimation. *Astronomy & Astrophysics*, 541:A74. (Cited on pages 60 and 70.)

Peacock, J., P. Schneider, G. Efstathiou, J. Ellis, B. Leibundgut, S. Lilly, and Y. Mellier
2006. ESA-ESO Working Group on "Fundamental Cosmology". Technical report, Ed. Universirty. (Cited on page 176.)

Peebles, P.
1980. *The large-scale structure of the universe.* (Cited on pages 13, 176, 223, and 224.)

Peiris, H. V. and L. Verde
2010. The shape of the primordial power spectrum: A last stand before Planck data. *Physical Review D*, 81(2):021302. (Cited on page 57.)

Penzias, A. A. and R. W. Wilson
1965. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *The Astrophysical Journal*, 142:419. (Cited on page 4.)

Perlmutter, S., G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer, B. J. Boyle, A. V. Filippenko, T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, W. J. Couch, and T. S. C. Project
1999. Measurements of Omega and Lambda from 42 High-Redshift Supernovae. *The Astrophysical Journal*, 517(2):565–586. (Cited on page 4.)

Piao, Y.-S., B. Feng, and X. Zhang
2004. Suppressing the CMB quadrupole with a bounce from the contracting phase to inflation. *Physical Review D*, 69(10):103520. (Cited on page 56.)

Pires, S. and A. Amara
2010. Weak Lensing Mass Reconstruction: Flexion Versus Shear. *The Astrophysical Journal*, 723(2):1507. (Cited on pages 95 and 119.)

Pires, S., J. L. Starck, A. Amara, R. Teyssier, A. Refregier, and J. Fadili
2009. FAst STatistics for weak Lensing (FASTLens): fast method for weak lensing statistics and map making. *Monthly Notices of the Royal Astronomical Society*, 395(3):1265. (Cited on pages 94, 111, 113, and 115.)

Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo, E. Battaner, K. Benabed, A. Benoit, A. Benoit-Levy, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. Bobin, J. J. Bock, A. Bonaldi, J. R. Bond, J. Borrill, F. R. Bouchet, M. Bridges, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, A. Catalano, A. Challinor, A. Chamballu, H. C. Chiang, L. Y. Chiang, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, F. Couchot, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, J. M. Delouis, F. X. Desert, C. Dickinson, J. M. Diego, H. Dole, S. Donzelli, O. Dore, M. Douspis, J. Dunkley, X. Dupac, G. Efstathiou, T. A. Ensslin, H. K. Eriksen, F. Finelli, O. Forni, M. Frailis, E. Franceschi, S. Galeotta, K. Ganga, C. Gauthier, M. Giard, G. Giardino, Y. Giraud-Heraud, J. Gonzalez-Nuevo, K. M. Gorski, S. Gratton, A. Gregorio, A. Gruppuso, J. Hamann, F. K. Hansen, D. Hanson, D. Harrison, S. Henrot-Versille, C. Hernandez-Monteagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, W. Hovest, K. M. Huffenberger, A. H. Jaffe, T. R. Jaffe, W. C. Jones, M. Juvela, E. Keihanen, R. Keskitalo, T. S. Kisner, R. Kneissl, J. Knoche, L. Knox, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lahteenmaki, J. M. Lamarre, A. Lasenby, R. J. Laureijs, C. R. Lawrence, S. Leach, J. P. Leahy, R. Leonardi, J. Lesgourgues, A. Lewis, M. Liguori, P. B. Lilje, M. Linden-Vornle, M. Lopez-Caniego, P. M. Lubin, J. F. Macias-Perez, B. Maffei, D. Maino, N. Mandolesi, M. Maris, D. J. Marshall, P. G. Martin, E. Martinez-Gonzalez, S. Masi, M. Massardi, S. Matarrese, F. Matthai, P. Mazzotta, P. R. Meinhold, A. Melchiorri, L. Mendes, A. Mennella, M. Migliaccio, S. Mitra, M. A. Miville-Deschenes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Norgaard-Nielsen, F. Noviello, D. Novikov, I. Novikov, I. J. O'Dwyer, S. Osborne, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, R. Paladini, S. Pandolfi, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, H. V. Peiris, O. Perdereau, L. Perotto, F. Perrotta, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, N. Ponthieu, L. Popa, T. Poutanen, G. W. Pratt, G. Prezeau, S. Prunet, J. L. Puget, J. P. Rachen, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, S. Ricciardi, T. Riller, I. Ristorcelli, G. Rocha, C. Rosset, G. Roudier, M. Rowan-Robinson, J. A. Rubino-Martin, B. Rusholme, M. Sandri, D. Santos, M. Savelainen, G. Savini, D. Scott, M. D. Seiffert, E. P. S. Shellard, L. D. Spencer, J. L. Starck, V. Stolyarov, R. Stompor, R. Sudiwala, R. Sunyaev, F. Sureau, D. Sutton, A. S. Suur-Uski, J. F. Sygnet, J. A. Tauber, D. Tavagnacco, L. Terenzi, L. Toffolatti, M. Tomasi, J. Treguer-Goudineau, M. Tristram, M. Tucci, J. Tuovinen, L. Valenziano, J. Valiviita, B. Van Tent, J. Varis, P. Vielva, F. Villa, N. Vittorio, L. A. Wade, B. D. Wandelt, M. White, A. Wilkinson, D. Yvon, A. Zacchei, J. P. Zibin, and A. Zonca
2014a. Planck 2013 results. XXII. Constraints on inflation. *Astronomy and Astrophysics*, 571:A22. (Cited on page 57.)

Planck Collaboration, P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, E. Battaner, K. Benabed, A. Benoit, A. Benoit-Levy, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. Bobin, J. J. Bock, A. Bonaldi, J. R. Bond, J. Borrill, F. R. Bouchet, M. Bridges, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, B. Cappellini, J. F. Cardoso, A. Catalano, A. Challinor, A. Chamballu, R. R. Chary, X. Chen, H. C. Chiang, L. Y. Chiang, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, F. Couchot, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, J. M. Delouis, F. X. Desert, C. Dickinson, J. M. Diego, K. Dolag, H. Dole, S. Donzelli, O. Dore, M. Douspis,

J. Dunkley, X. Dupac, G. Efstathiou, F. Elsner, T. A. Ensslin, H. K. Eriksen, F. Finelli, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, T. C. Gaier, S. Galeotta, S. Galli, K. Ganga, M. Giard, G. Giardino, Y. Giraud-Heraud, E. Gjerlow, J. Gonzalez-Nuevo, K. M. Gorski, S. Gratton, A. Gregorio, A. Gruppuso, J. E. Gudmundsson, J. Haissinski, J. Hamann, F. K. Hansen, D. Hanson, D. Harrison, S. Henrot-Versille, C. Hernandez-Monteagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, Z. Hou, W. Hovest, K. M. Huffenberger, A. H. Jaffe, T. R. Jaffe, J. Jewell, W. C. Jones, M. Juvela, E. Keihanen, R. Keskitalo, T. S. Kisner, R. Kneissl, J. Knoche, L. Knox, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lahteenmaki, J. M. Lamarre, A. Lasenby, M. Lattanzi, R. J. Laureijs, C. R. Lawrence, S. Leach, J. P. Leahy, R. Leonardi, J. Leon-Tavares, J. Lesgourgues, A. Lewis, M. Liguori, P. B. Lilje, M. Linden-Vornle, M. Lopez-Caniego, P. M. Lubin, J. F. Macias-Perez, B. Maffei, D. Maino, N. Mandolesi, M. Maris, D. J. Marshall, P. G. Martin, E. Martinez-Gonzalez, S. Masi, M. Massardi, S. Matarrese, F. Matthai, P. Mazzotta, P. R. Meinhold, A. Melchiorri, J. B. Melin, L. Mendes, E. Menegoni, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M. A. Miville-Deschenes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Norgaard-Nielsen, F. Noviello, D. Novikov, I. Novikov, I. J. O'Dwyer, S. Osborne, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, R. Paladini, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, D. Pearson, T. J. Pearson, H. V. Peiris, O. Perdereau, L. Perotto, F. Perrotta, V. Pettorino, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, P. Platania, E. Pointecouteau, G. Polenta, N. Ponthieu, L. Popa, T. Poutanen, G. W. Pratt, G. Prezeau, S. Prunet, J. L. Puget, J. P. Rachen, W. T. Reach, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, S. Ricciardi, T. Riller, I. Ristorcelli, G. Rocha, C. Rosset, G. Roudier, M. Rowan-Robinson, J. A. Rubino-Martin, B. Rusholme, M. Sandri, D. Santos, M. Savelainen, G. Savini, D. Scott, M. D. Seiffert, E. P. S. Shellard, L. D. Spencer, J. L. Starck, V. Stolyarov, R. Stompor, R. Sudiwala, R. Sunyaev, F. Sureau, D. Sutton, A. S. Suur-Uski, J. F. Sygnet, J. A. Tauber, D. Tavagnacco, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, M. Tucci, J. Tuovinen, M. Turler, G. Umana, L. Valenziano, J. Valiviita, B. Van Tent, P. Vielva, F. Villa, N. Vittorio, L. A. Wade, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Wilkinson, D. Yvon, A. Zacchei, and A. Zonca

2014b. Planck 2013 results. XVI. Cosmological parameters. *Astronomy and Astrophysics*, 571:A16. (Cited on page 73.)

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, F. Arroja, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, E. Battaner, K. Benabed, A. Benoit, A. Benoit-Levy, J. P. Bernard, M. Bersanelli, P. Bielewicz, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, A. Catalano, A. Challinor, A. Chamballu, R. R. Chary, H. C. Chiang, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, C. Combet, D. Contreras, F. Couchot, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, F. X. Desert, J. M. Diego, H. Dole, S. Donzelli, O. Dore, M. Douspis, A. Ducout, X. Dupac, G. Efstathiou, F. Elsner, T. A. Ensslin, H. K. Eriksen, J. Fergusson, F. Finelli, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frejsel, A. Frolov, S. Galeotta, S. Galli, K. Ganga, C. Gauthier, M. Giard, Y. Giraud-Heraud, E. Gjerlow, J. Gonzalez-Nuevo, K. M. Gorski, S. Gratton, A. Gregorio, A. Gruppuso, J. E. Gudmundsson, J. Hamann, W. Handley, F. K. Hansen, D. Hanson, D. L. Harrison, S. Henrot-Versille, C. Hernandez-Monteagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, W. Hovest, Z. Huang, K. M. Huffenberger, G. Hurier, A. H. Jaffe, T. R. Jaffe, W. C. Jones, M. Juvela, E. Keiha-

nen, R. Keskitalo, J. Kim, T. S. Kisner, R. Kneissl, J. Knoche, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lahteenmaki, J. M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, R. Leonardi, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Linden-Vornle, M. Lopez-Caniego, P. M. Lubin, Y. Z. Ma, J. F. Macias-Perez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, P. G. Martin, E. Martinez-Gonzalez, S. Masi, S. Matarrese, P. Mazzotta, P. McGehee, P. R. Meinhold, A. Melchiorri, L. Mendes, A. Mennella, M. Migliaccio, S. Mitra, M. A. Miville-Deschenes, D. Molinari, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, M. Munchmeyer, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Norgaard-Nielsen, F. Noviello, D. Novikov, I. Novikov, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, R. Paladini, S. Pandolfi, D. Paoletti, F. Pasian, G. Patanchon, T. J. Pearson, H. V. Peiris, O. Perdereau, L. Perotto, F. Perrotta, V. Pettorino, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, L. Popa, G. W. Pratt, G. Prezeau, S. Prunet, J. L. Puget, J. P. Rachen, W. T. Reach, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, A. Renzi, I. Ristorcelli, G. Rocha, C. Rosset, M. Rossetti, G. Roudier, M. Rowan-Robinson, J. A. Rubino-Martin, B. Rusholme, M. Sandri, D. Santos, M. Savelainen, G. Savini, D. Scott, M. D. Seiffert, E. P. S. Shellard, M. Shiraishi, L. D. Spencer, V. Stolyarov, R. Stompor, R. Sudiwala, R. Sunyaev, D. Sutton, A. S. Suur-Uski, J. F. Sygnet, J. A. Tauber, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, T. Trombetti, M. Tucci, J. Tuovinen, L. Valenziano, J. Valiviita, B. Van Tent, P. Vielva, F. Villa, L. A. Wade, B. D. Wandelt, I. K. Wehus, M. White, D. Yvon, A. Zacchei, J. P. Zibin, and A. Zonca
2015a. Planck 2015 results. XX. Constraints on inflation. *ArXiv e-prints*, P. arXiv:1502.02114. (Cited on pages 55 and 58.)

Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, N. Bartolo, E. Battaner, R. Battye, K. Benabed, A. Benoit, A. Benoit-Levy, J. P. Bernard, M. Bersanelli, P. Bielewicz, A. Bonaldi, L. Bonavera, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, A. Catalano, A. Challinor, A. Chamballu, R. R. Chary, H. C. Chiang, J. Chluba, P. R. Christensen, S. Church, D. L. Clements, S. Colombi, L. P. L. Colombo, C. Combet, A. Coulais, B. P. Crill, A. Curto, F. Cuttaia, L. Danese, R. D. Davies, R. J. Davis, P. de Bernardis, A. de Rosa, G. de Zotti, J. Delabrouille, F. X. Desert, E. Di Valentino, C. Dickinson, J. M. Diego, K. Dolag, H. Dole, S. Donzelli, O. Dore, M. Douspis, A. Ducout, J. Dunkley, X. Dupac, G. Efstathiou, F. Elsner, T. A. Ensslin, H. K. Eriksen, M. Farhang, J. Fergusson, F. Finelli, O. Forni, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frejsel, S. Galeotta, S. Galli, K. Ganga, C. Gauthier, M. Gerbino, T. Ghosh, M. Giard, Y. Giraud-Heraud, E. Giusarma, E. Gjerlow, J. Gonzalez-Nuevo, K. M. Gorski, S. Gratton, A. Gregorio, A. Gruppuso, J. E. Gudmundsson, J. Hamann, F. K. Hansen, D. Hanson, D. L. Harrison, G. Helou, S. Henrot-Versille, C. Hernandez-Monteagudo, D. Herranz, S. R. Hildebrandt, E. Hivon, M. Hobson, W. A. Holmes, A. Hornstrup, W. Hovest, Z. Huang, K. M. Huffenberger, G. Hurier, A. H. Jaffe, T. R. Jaffe, W. C. Jones, M. Juvela, E. Keihanen, R. Keskitalo, T. S. Kisner, R. Kneissl, J. Knoche, L. Knox, M. Kunz, H. Kurki-Suonio, G. Lagache, A. Lahteenmaki, J. M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, J. P. Leahy, R. Leonardi, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Linden-Vornle, M. Lopez-Caniego, P. M. Lubin, J. F. Macias-Perez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marchini, P. G. Martin, M. Martinelli, E. Martinez-Gonzalez, S. Masi, S. Matarrese, P. Mazzotta, P. McGehee, P. R. Meinhold, A. Melchiorri, J. B. Melin, L. Mendes, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M. A. Miville-Deschenes, A. Moneti, L. Montier, G. Morgante, D. Mortlock, A. Moss, D. Munshi, J. A. Murphy, P. Naselsky, F. Nati, P. Natoli, C. B. Netterfield, H. U. Norgaard-Nielsen, F. Noviello,

D. Novikov, I. Novikov, C. A. Oxborrow, F. Paci, L. Pagano, F. Pajot, R. Paladini, D. Paoletti, B. Partridge, F. Pasian, G. Patanchon, T. J. Pearson, O. Perdereau, L. Perotto, F. Perrotta, V. Pettorino, F. Piacentini, M. Piat, E. Pierpaoli, D. Pietrobon, S. Plaszczynski, E. Pointecouteau, G. Polenta, L. Popa, G. W. Pratt, G. Prezeau, S. Prunet, J. L. Puget, J. P. Rachen, W. T. Reach, R. Rebolo, M. Reinecke, M. Remazeilles, C. Renault, A. Renzi, I. Ristorcelli, G. Rocha, C. Rosset, M. Rossetti, G. Roudier, B. Rouille d'Orfeuil, M. Rowan-Robinson, J. A. Rubino-Martin, B. Rusholme, N. Said, V. Salvatelli, L. Salvati, M. Sandri, D. Santos, M. Savelainen, G. Savini, D. Scott, M. D. Seiffert, P. Serra, E. P. S. Shellard, L. D. Spencer, M. Spinelli, V. Stolyarov, R. Stompor, R. Sudiwala, R. Sunyaev, D. Sutton, A. S. Suur-Uski, J. F. Sygnet, J. A. Tauber, L. Terenzi, L. Toffolatti, M. Tomasi, M. Tristram, T. Trombetti, M. Tucci, J. Tuovinen, M. Turler, G. Umana, L. Valenziano, J. Valiviita, B. Van Tent, P. Vielva, F. Villa, L. A. Wade, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Wilkinson, D. Yvon, A. Zacchei, and A. Zonca
2015b. Planck 2015 results. XIII. Cosmological parameters. *ArXiv e-prints*, P. arXiv:1502.01589. (Cited on pages 5, 12, 56, and 57.)

Powell, B. A. and W. H. Kinney
2007. Pre-inflationary vacuum in the cosmic microwave background. *Physical Review D*, 76(6):063512. (Cited on page 56.)

Pratten, G. and D. Munshi
2013. Effects of linear redshift space distortions and perturbation theory on BAOs: a 3D spherical analysis. *Monthly Notices of the Royal Astronomical Society*, 436(4):3792. (Cited on pages 160, 180, and 182.)

Raguet, H., J. Fadili, and G. Peyré
2013. A Generalized Forward-Backward Splitting. *SIAM Journal on Imaging Sciences*, 6(3):1199–1226. (Cited on page 33.)

Rapin, J., J. Bobin, A. Larue, and J.-L. Starck
2014. NMF with Sparse Regularizations in Transformed Domains. *ArXiv e-prints*, P. arXiv:1407.7691. (Cited on page 31.)

Rassat, A., A. Amara, L. Amendola, F. J. Castander, T. Kitching, M. Kunz, A. Refregier, Y. Wang, and J. Weller
2008. Deconstructing Baryon Acoustic Oscillations: A Comparison of Methods. *ArXiv e-prints*, P. arXiv:0810.0003. (Cited on pages 178, 184, 185, and 191.)

Rassat, A., K. Land, O. Lahav, and F. B. Abdalla
2007. Cross-correlation of 2MASS and WMAP 3: implications for the integrated Sachs-Wolfe effect. *Monthly Notices of the Royal Astronomical Society*, 377(3):1085. (Cited on page 184.)

Rassat, A. and A. Refregier
2012. 3D spherical analysis of baryon acoustic oscillations. *Astronomy and Astrophysics*, 540:A115. (Cited on pages 156, 160, 176, 178, and 180.)

Refregier, A., A. Amara, T. D. Kitching, and A. Rassat
2011. iCosmo: an interactive cosmology package. *Astronomy & Astrophysics*, 528:A33. (Cited on page 176.)

Rhodes, J. D., R. J. Massey, J. Albert, N. Collins, R. S. Ellis, C. Heymans, J. P. Gardner, J.-P. Kneib, A. Koekemoer, A. Leauthaud, Y. Mellier, A. Refregier, J. E. Taylor, and L. Van Waerbeke
2007. The Stability of the Point-Spread Function of the Advanced Camera for Surveys on the Hubble Space Telescope and Implications for Weak Gravitational Lensing. *The Astrophysical Journal Supplement Series*, 172(1):203. (Cited on page 90.)

Riess, A. G., A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry
1998. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *The Astronomical Journal*, 116(3):1009. (Cited on page 4.)

Romano, A. E. and M. Sasaki
2008. Effects of particle production during inflation. *Physical Review D*, 78(10):103522. (Cited on page 56.)

Rowe, B., D. Bacon, R. Massey, C. Heymans, B. Haussler, A. Taylor, J. Rhodes, and Y. Mellier
2013. Flexion measurement in simulations of Hubble Space Telescope data. *Monthly Notices of the Royal Astronomical Society*, 435(1):822. (Cited on page 124.)

Schneider, P.
2006. Weak Gravitational Lensing. In *Gravitational Lensing: Strong, Weak and Micro*, Pp. 269–451. Berlin, Heidelberg: Springer Berlin Heidelberg. (Cited on page 94.)

Schneider, P.
2014. Generalized multi-plane gravitational lensing: time delays, recursive lens equation, and the mass-sheet transformation. *ArXiv e-prints*, P. arXiv:1409.0015. (Cited on page 86.)

Scoville, N., R. G. Abraham, H. Aussel, J. E. Barnes, A. Benson, A. W. Blain, D. Calzetti, A. Comastri, P. Capak, C. Carilli, J. E. Carlstrom, C. M. Carollo, J. Colbert, E. Daddi, R. S. Ellis, M. Elvis, S. P. Ewald, M. Fall, A. Franceschini, M. Giavalisco, W. Green, R. E. Griffiths, L. Guzzo, G. Hasinger, C. Impey, J. P. Kneib, J. Koda, A. Koekemoer, O. Lefevre, S. Lilly, C. T. Liu, H. J. McCracken, R. Massey, Y. Mellier, S. Miyazaki, B. Mobasher, J. Mould, C. Norman, A. Refregier, A. Renzini, J. Rhodes, M. Rich, D. B. Sanders, D. Schiminovich, E. Schinnerer, M. Scodeggio, K. Sheth, P. L. Shopbell, Y. Taniguchi, N. D. Tyson, C. M. Urry, L. Van Waerbeke, P. Vettolani, S. D. M. White, and L. Yan
2007. COSMOS: Hubble Space Telescope Observations. *The Astrophysical Journal Supplement Series*, 172(1):38. (Cited on page 99.)

Sealfon, C., L. Verde, and R. Jimenez
2005. Smoothing spline primordial power spectrum reconstruction. *Physical Review D*, 72(10):103520. (Cited on page 57.)

Seitz, C. and P. Schneider
1995. Steps towards nonlinear cluster inversion through gravitational distortions II. Generalization of the Kaiser and Squires method. *Astronomy and Astrophysics*, 297:287. (Cited on pages 94 and 117.)

Seitz, C. and P. Schneider
1997. Steps towards nonlinear cluster inversion through gravitational distortions. III. Including a redshift distribution of the sources. *Astronomy and Astrophysics*, 318:687. (Cited on page 91.)

Seitz, S. and P. Schneider
1996. Cluster lens reconstruction using only observed local data: an improved finite-field inversion technique. *Astronomy and Astrophysics*, 305:383. (Cited on page 94.)

Seitz, S., P. Schneider, and M. Bartelmann
1998. Entropy-regularized maximum-likelihood cluster mass reconstruction. *Astronomy and Astrophysics*, 337:325. (Cited on page 98.)

Seitz, S., P. Schneider, and J. Ehlers
1994. Light propagation in arbitrary spacetimes and the gravitational lens approximation. *Classical and Quantum Gravity*, 11(9):2345. (Cited on pages 79 and 80.)

Selesnick, I. W. and M. A. T. Figueiredo
2009. Signal restoration with overcomplete wavelet transforms: comparison of analysis and synthesis priors. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, V. K. Goyal, M. Papadakis, and D. Van De Ville, eds., volume 7446, P. 74460D. (Cited on page 27.)

Sendur, L. and I. W. Selesnick
2002a. Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency. *IEEE Transactions on Signal Processing*, 50(11):2744. (Cited on page 171.)

Sendur, L. and I. W. Selesnick
2002b. Bivariate shrinkage with local variance estimation. *IEEE Signal Processing Letters*, 9(12):438. (Cited on page 171.)

Seo, H.-J. and D. J. Eisenstein
2003. Probing Dark Energy with Baryonic Acoustic Oscillations from Future Large Galaxy Redshift Surveys. *The Astrophysical Journal*, 598(2):720. (Cited on page 176.)

Seo, H.-J. and D. J. Eisenstein
2007. Improved Forecasts for the Baryon Acoustic Oscillations and Cosmological Distance Scale. *The Astrophysical Journal*, 665(1):14. (Cited on page 176.)

Shafieloo, A. and T. Souradeep
2004. Primordial power spectrum from WMAP. *Physical Review D*, 70(4):043523. (Cited on page 57.)

Shafieloo, A. and T. Souradeep
2008. Estimation of primordial spectrum with post-WMAP 3-year data. *Physical Review D*, 78(2):023511. (Cited on page 57.)

Shafieloo, A., T. Souradeep, P. Manimaran, P. K. Panigrahi, and R. Rangarajan
2007. Features in the primordial spectrum from WMAP: A wavelet analysis. *Physical Review D*, 75(12):123502. (Cited on page 57.)

Shapiro, C., R. G. Crittenden, and W. J. Percival
2012. The complementarity of redshift-space distortions and the integrated Sachs-Wolfe effect: a 3D spherical analysis. *Monthly Notices of the Royal Astronomical Society*, 422(3):2341. (Cited on page 176.)

Shensa, M. J.
1992. The discrete wavelet transform: wedding the a trous and Mallat algorithms. *IEEE Transactions on Signal Processing*, 40(10):2464. (Cited on page 50.)

Silverman, B.
1986. *Density estimation for statistics and data analysis.* (Cited on page 144.)

Simon, P., C. Heymans, T. Schrabback, A. N. Taylor, M. E. Gray, L. van Waerbeke, C. Wolf, D. Bacon, M. Barden, A. Bohm, B. Haussler, K. Jahnke, S. Jogee, E. van Kampen, K. Meisenheimer, and C. Y. Peng
2012. Spatial matter density mapping of the STAGES Abell A901/2 supercluster field with 3D lensing. *Monthly Notices of the Royal Astronomical Society*, 419(2):998. (Cited on page 101.)

Simon, P., A. N. Taylor, and J. Hartlap
2009. Unfolding the matter distribution using three-dimensional weak gravitational lensing. *Monthly Notices of the Royal Astronomical Society*, 399(1):48. (Cited on pages 98, 100, 101, 102, 130, 137, and 138.)

Sinha, R. and T. Souradeep
2006. Post-WMAP assessment of infrared cutoff in the primordial spectrum from inflation. *Physical Review D*, 74(4):043518. (Cited on page 57.)

Slosar, A., S. Ho, M. White, and T. Louis
2009. The acoustic peak in the Lyman alpha forest. *Journal of Cosmology and Astro-Particle Physics*, 2009(10):019. (Cited on page 176.)

Smail, I., R. S. Ellis, and M. J. Fitchett
1994. Gravitational lensing of distant field galaxies by rich clusters - I. Faint galaxy redshift distributions. *Monthly Notices of the Royal Astronomical Society*, 270(2):245–270. (Cited on page 184.)

Smith, R. E., J. A. Peacock, A. Jenkins, S. D. M. White, C. S. Frenk, F. R. Pearce, P. A. Thomas, G. Efstathiou, and H. M. P. Couchman
2003. Stable clustering, the halo model and non-linear cosmological power spectra. *Monthly Notices of the Royal Astronomical Society*, 341(4):1311. (Cited on pages 14, 15, and 221.)

Starck, J.-L., E. J. Candes, and D. L. Donoho
2002. The curvelet transform for image denoising. *IEEE Transactions on Image Processing*, 11(6):670. (Cited on page 23.)

Starck, J. L., M. Elad, and D. Donoho
2004. Redundant multiscale transforms and their application for morphological component separation. *Advances in Imaging and Electron Physics*, 132(1):287–348. (Cited on page 27.)

Starck, J. L., M. Elad, and D. L. Donoho
2005. Image decomposition via the combination of sparse representations and a variational approach. *IEEE Transactions on Image Processing*, 14(10):1570. (Cited on page 27.)

Starck, J.-L., J. Fadili, and F. Murtagh
2007. The Undecimated Wavelet Decomposition and its Reconstruction. *IEEE Transactions on Image Processing*, 16(2):297. (Cited on pages 51, 53, and 132.)

Starck, J.-L. and F. Murtagh
2006. *Astronomical Image and Data Analysis*. Astronomical image and data analysis, by J.-L.˜Starck and F.˜Murtagh.˜Astronomy and astrophysics library.˜ Berlin: Springer, 2006. (Cited on pages 50, 51, 167, and 168.)

Starck, J. L., S. Pires, and A. Refregier
2006. Weak lensing mass reconstruction using wavelets. *Astronomy and Astrophysics*, 451(3):1139. (Cited on pages 111, 149, and 150.)

Starobinsky, A. A.
1992. Spectrum of adiabatic perturbations in the universe when there are singularities in the inflation potential. *JETP Lett.*, 55:489–494. (Cited on page 56.)

Takada, M. and B. Jain
2003. Three-point correlations in weak lensing surveys: model predictions and applications. *Monthly Notices of the Royal Astronomical Society*, 344(3):857. (Cited on page 139.)

Taylor, A. N.
2001. Imaging the 3-D cosmological mass distribution with weak gravitational lensing. *ArXiv e-prints*, Pp. astro–ph/0111605. (Cited on pages 98 and 99.)

Taylor, A. N., D. J. Bacon, M. E. Gray, C. Wolf, K. Meisenheimer, S. Dye, A. Borch, M. Kleinheinrich, Z. Kovacs, and L. Wisotzki
2004. Mapping the 3D dark matter with weak lensing in COMBO-17. *Monthly Notices of the Royal Astronomical Society*, 353(4):1176. (Cited on page 99.)

Tegmark, M.
1996. An Icosahedron-Based Method for Pixelizing the Celestial Sphere. *The Astrophysical Journal*, 470:L81. (Cited on page 157.)

Tegmark, M., A. N. Taylor, and A. F. Heavens
1997. Karhunen-Loeve Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets? *The Astrophysical Journal*, 480(1):22–35. (Cited on page 182.)

Tocchini-Valentini, D., M. Douspis, and J. Silk
2005. A new search for features in the primordial power spectrum. *Monthly Notices of the Royal Astronomical Society*, 359(1):31–35. (Cited on page 57.)

Tocchini-Valentini, D., Y. Hoffman, and J. Silk
2006. Non-parametric reconstruction of the primordial power spectrum at horizon scales from WMAP data. *Monthly Notices of the Royal Astronomical Society*, 367(3):1095. (Cited on page 57.)

Vaiter, S., G. Peyre, C. Dossal, and J. Fadili
2013. Robust Sparse Analysis Regularization. *IEEE Transactions on Information Theory*, 59(4):2001–2016. (Cited on page 28.)

VanderPlas, J. T., A. J. Connolly, B. Jain, and M. Jarvis
2011. Three-dimensional Reconstruction of the Density Field: An SVD Approach to Weak-lensing Tomography. *The Astrophysical Journal*, 727(2):118. (Cited on pages 98, 101, 102, and 130.)

Vazquez, J. A., M. Bridges, M. P. Hobson, and A. N. Lasenby
2012. Model selection applied to reconstruction of the Primordial Power Spectrum. *Journal of Cosmology and Astro-Particle Physics*, 2012(6):006. (Cited on page 57.)

Verde, L. and H. Peiris
2008. On minimally parametric primordial power spectrum reconstruction and the evidence for a red tilt. *Journal of Cosmology and Astro-Particle Physics*, 2008(7):009. (Cited on page 57.)

Vetterli, M.
1986. Filter banks allowing perfect reconstruction. *Signal Processing*, 10(3):219–244. (Cited on page 46.)

Vu, B. C.
2013. A splitting algorithm for dual monotone inclusions involving cocoercive operators. *Advances in Computational Mathematics*, 38(3):667–681. (Cited on pages 34, 110, and 121.)

Wang, X., B. Feng, M. Li, X.-L. Chen, and X. Zhang
2005. Natural Inflation, Planck Scale Physics and Oscillating Primordial Spectrum. *International Journal of Modern Physics D*, 14(8):1347. (Cited on page 56.)

Wang, Y. and G. J. Mathews
2002. Model-independent Primordial Power Spectrum from MAXIMA, BOOMERANG, and DASI Data. *The Astrophysical Journal*, 573(1):1–6. (Cited on page 57.)

Weinberg, N. N. and M. Kamionkowski
2003. Constraining dark energy from the abundance of weak gravitational lenses. *Monthly Notices of the Royal Astronomical Society*, 341(1):251–262. (Cited on page 139.)

Williams, J. P., E. J. de Geus, and L. Blitz
1994. Determining structure in molecular clouds. *The Astrophysical Journal*, 428:693. (Cited on page 141.)

Wolf, C., K. Meisenheimer, H. W. Rix, A. Borch, S. Dye, and M. Kleinheinrich
2003. The COMBO-17 survey: Evolution of the galaxy luminosity function from 25 000 galaxies with 0.2< z <1.2. *Astronomy and Astrophysics*, 401:73. (Cited on page 99.)

Xu, X., M. White, N. Padmanabhan, D. J. Eisenstein, J. Eckel, K. Mehta, M. Metchnik, P. Pinto, and H. J. Seo
2010. A New Statistic for Analyzing Baryon Acoustic Oscillations. *The Astrophysical Journal*, 718(2):1224. (Cited on page 176.)

Yang, X., J. M. Kratochvil, S. Wang, E. A. Lim, Z. Haiman, and M. May
2011. Cosmological information in weak lensing peaks. *Physical Review D*, 84(4):043529. (Cited on page 86.)

Yoo, J. and V. Desjacques
2013. All-sky analysis of the general relativistic galaxy power spectrum. *Physical Review D*, 88(2):023502. (Cited on page 181.)

Zalinescu, C.
2002. *Convex Analysis in General Vector Spaces*. WORLD SCIENTIFIC. (Cited on page 29.)

Zeldovich, Y. B.
1972. A hypothesis, unifying the structure and the entropy of the Universe. *Monthly Notices of the Royal Astronomical Society*, 160:1P. (Cited on pages 12 and 56.)

Zhang, B., M. J. Fadili, J.-L. Starck, and J.-C. Olivo-Marin
2007. Multiscale Variance-Stabilizing Transform for Mixed-Poisson-Gaussian Processes and its Applications in Bioimaging. In *2007 IEEE International Conference on Image Processing*, Pp. VI – 233–VI – 236. IEEE. (Cited on page 197.)

université PARIS-SACLAY

ÉCOLE DOCTORALE
Astronomie et Astrophysique d'Île-de-France

Titre : Reconstruction de la carte de masse de matière noire par effet de lentille gravitationnelle.
Mots clés : Effet de lentille gravitationnelle faible, cosmologie, parcimonie

Resumé : L'effet de lentille gravitationnelle, qui se traduit par une déformation des images nous parvenant de galaxies lointaines, constitue l'une des techniques les plus prometteuse pour répondre aux nombreuses questions portant sur la nature de l'énergie sombre et de la matière noire. Cet effet de lentille étant sensible à la masse totale, il permet de sonder directement la distribution de matière noire, qui resterait autrement invisible. En mesurant la forme d'un grand nombre de galaxies lointaines, il est possible d'estimer statistiquement les déformations causées par l'effet de lentille gravitationnelle puis d'en inférer la distribution de masse de la lentille. La reconstruction de ces cartes de masse constitue un problème inverse qui se trouve être mal posé dans un certain nombre de situations d'intérêt, en particulier lors de la reconstruction de la carte de masse aux petites échelles ou en trois dimensions. Dans ces situations, il devient impossible de reconstruire une carte sans l'ajout d'information a priori.
Une classe particulière de méthodes, basées sur un a priori de parcimonie, s'est révélée remarquablement efficace pour résoudre des problèmes inverses similaires pour un large champ d'applications tels que la géophysique et l'imagerie médicale. Le but principal de cette thèse est donc d'adapter ces techniques de régularisation parcimonieuses au problème de la cartographie de la matière noire afin de developper une nouvelle generation de méthodes. Nous développons en particulier de nouveaux algorithmes permettant la reconstruction de cartes de masse bi-dimensionnelles de haute résolution ainsi que de cartes de masse tri-dimensionnelles.
Nous appliquons de plus les mêmes méthodes de régularisation parcimonieuse au problème de la reconstruction du spectre de puissance des fluctuations primordiales de densités à partir de mesures du fond diffus cosmologique, ce qui constitue un problème inverse particulièrement difficile à résoudre. Nous développons un nouvel algorithme pour résoudre ce problème, que nous appliquons aux données du satellite Planck.
Enfin, nous investiguons de nouvelles méthodes pour l'analyse de relevés cosmologiques exprimés en coordonnées sphériques. Nous développons une nouvelle transformée en ondelettes pour champs scalaires exprimés sur la boule 3D et nous comparons différentes méthodes pour l'analyse cosmologique de relevés de galaxies spectroscopiques.

Title : Sparse reconstruction of the dark matter mass map rom weak gravitational lensing.
Keywords : Weak gravitational lensing, cosmology, sparsity

Resumé : Gravitational lensing, that is the distortion of the images of distant galaxies by intervening massive objects, has been identified as one of the most promising probes to help answer questions relative to the nature of dark matter and dark energy. As the lensing effect is caused by the total matter content, it can directly probe the distribution of the otherwise invisible dark matter. By measuring the shapes of distant galaxies and statistically estimating the deformations caused by gravitational lensing, it is possible to reconstruct the distribution of the intervening mass. This mass-mapping process can be seen as an instance of a linear inverse problem, which can be ill-posed in many situations of interest, especially when mapping the dark matter on small angular scales or in three dimensions. As a result, recovering a meaningful mass-map in these situations is not possible without prior information.
In recent years, a class of methods based on a so-called sparse prior has proven remarkably successful at solving similar linear inverse problems in a wide range of fields such as medical imaging or geophysics. The primary goal of this thesis is to apply these sparse regularisation techniques to the gravitational lensing problem in order to build next-generation dark matter mass-mapping tools. We propose in particular new algorithms for the reconstruction of high-resolution 2D mass-maps and 3D mass-maps and demonstrate in both cases the effectiveness of the sparse prior. We also apply the same sparse methodologies to the reconstruction the primordial density fluctuation power spectrum from measurements of the Cosmic Microwave Background which constitutes another notoriously difficult inverse problem. We apply the resulting algorithm to reconstruct the primordial power spectrum using data from the Planck satellite.
Finally, we investigate new methodologies for the analysis of cosmological surveys in spherical coordinates. We develop a new wavelet transform for the analysis of scalar fields on the 3D ball. We also conduct a comparison of methods for the 3D analysis of spectroscopic galaxy survey.